KNOWLEDGEIT.CO.UK

Deploying a RIPE Atlas Probe (The Hard Way)

Chris Russell
Pulsant, Newcastle

UKNOF 41
Edinburgh, September 2018

Me, Myself, I and the company I worked for

@kit_chrisr

- Senior Networks Engineer, Pulsant (Onyx, Knowledge I.T)
- Programme Committee Chair, UKNOF
- Box Ticker on the Autism Spectrum

KNOWLEDGEIT.CO.UK

- Managed / Professional Services company at heart (Cisco, Cloud, Support)
- 4 pops (North East), 120 Rack DC, ISP, Cabling Division…
- Circa ~ 12m Turnover

.

The shortest presentation ever…

Thank you! Any Questions ?

What this is really about …

A Long Time Ago In a Galaxy, Far Far Away (well, Washington, New York and ….. Sunderland) in 2014 …

A Customer, a Business Incubator with ~ 100 Small Companies over a 9 Building Campus – requested a network refresh & split from our core network – mutual benefits

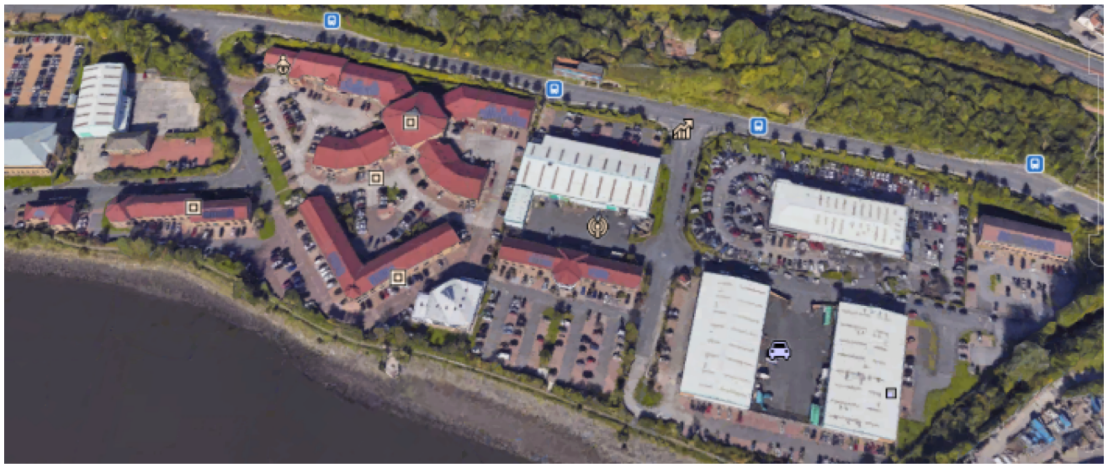

Primarily funding via services (Tenants) and Grants (EU mainly) – Value Required in any investment & resilience essential (good level of occupancy based on good reputation for business support && connectivity)

Close working relationship – challenging tenants – (“Your ISP is broken, they are assigning us Microsoft address overriding our DHCP so are clearly clueless”)

Out with the old, in with the new …

The Old (Justified & Ancient)

- Mix of 2950 and 3500XLs
  - Some horrific bridging / spanning tree fun (including our Core Network at the time)
- PIX 515E's
- OM2 1Gbps (barely) fiber
  - Riverside Campus – Rat's aren't friendly
- Some Interesting Switch locations
  - External cabinets (with heaters)
  - Facilities cupboards (technical term – a bit manky)

Out with the old, in with the new …

The New

- 3750X (collapsed core), 3560X access
- New (semi Diverse) SMF (Hub/Spoke)
- 2x DHCP Servers (DHCPD)
- We had plenty of DL360s
- Previously everything was statically assigned
- ASA5515X
- 3925E

Then things got a little convoluted

- I had been attending UKNOF for a little while and taking in a lot of things we'd never seen before in … (UKNOF19, AQL, Leeds, Apr 2011)
  First Timers - I knew what Andy Davidson looked like, that's about it!
  The Adelphi drinks…

- The Technical Director Moved On

- Onyx came in for us (~ 2 year process)

- Customer started construction of a new building off-site

- I started thinking - about things I saw at UKNOF, about Onyx …

- I started redesigning things…..

Making your own life difficult, aka, the hard way…

- Lets look to deploy at Atlas probe in the new network - ON IPV6
- Hell lets flood the network with ipv6 – including their Windows Cloud
- Lets use OSPF within the customer network rather than EIGRP (we used OSPF only on our core, even then limited)
- Can we use these magical things called VRF's (VRF-Lite in this case)… ????

The reactions when I said ‘ipv6’

Support Services

Professional Services

Management

But the Technical Director had a different way of thinking…

The Business Case for ipv6
(when you have lots of ipv4 and NAT)

This page intentionally left blank

Then the fun really began …. The Addressing Plan!

RIPE's ipv6 courses are very good – but when we did them, we were some way away from implementing ipv6 – ie: I'd forgotton nie on everything. (1st UKNOF = RIPE Course)

HE.net's ipv6 certification was also useful (helps when you run an ISP however)

Below is a way better summary of what I learned the hard way

Tom Coffeen/Veronika McKillop
UKNOF35 – Top 5 things when preparing your v6 addressing plan -
https://indico.uknof.org.uk/event/37/contribution/9/material/slides/0.pdf

The takeaways:

Think Subnets & Supernets, NOT addresses

Nibble boundaries are your friends. (/52, /56, /60)

The Addressing Plan – Mapping the SuperNets

Network	V4	V6 equiv
Firewall	5x/24s	/60 (16*/64)
Tenant	/16 supernet	/56 (256*/64)
Staff	/16 supernet	/56
IS	/29s (Outside/DMZ)	/60s (Just In case)

Supernet	Subnet	Prefix Size	Prefix Range	v4	
2001:db8:4:0::		/60	4:0 - 4:f		Covering Subnet - Firewall / Core Router Internal
	2001:db8:4:0::	/64		1.1.1.1/23	Outside Subnet
	2001:db8:4:1::	/64		172.31.1.0/24	Inside (Management) Subnet
2001:db8:4:100::		/56	4:100 - 4:1ff	4:100 - 4:1ff	Covering Supernet - Tenant (RFC 1918) Network
	2001:db8:4:100::	/64		10.16.1.0/24	Switch Management Vlan
	2001:db8:4:101::	/64		10.16.2.0/24	Services Vlan (DHCP Servers)
	2001:db8:4:103::	/64		172.16.3.0/24	Tenant Vlan 3
2001:db8:4:200::		/56	4:200 - 4:2ff		Covering Supernet - Staff (RFC 1918) Network
	2001:db8:4:200::	/64		10.17.6.0/24	Coud
	2001:db8:4:201::	/64		10.17.5.0/24	Staff Wifi Network
2001:db8:4:400	ONWARDS USED FOR IS CUSTOMER	(/60 = 16 64s)			
	2001:db8:4:400	/60		2.2.2.0/29	IS 1
	2001:db8:4:411	/60		3.3.3.0/29	IS2

https://www.ripe.net/manage-ips-and-asns/ipv6/ipv6-subnetting-card

We should probably test, *something*..

- Not lots of lab network equipment to play with, had to be creative…

DL360 G7 – Dual Hex Core, 56GB

- ESXI, Ubuntu VM – GNS3, some 7200 images
- Taught me the basics of OSPFv3, eigrpv3 -> ospfv3 && ipv6 config

- Later on used with the Ubuntu VM + IOS XRv, CSR1000v to lab the Onyx 7200 –> ASR 9/1K migration (Many Many virtual routers talking to each other)

The Implementation

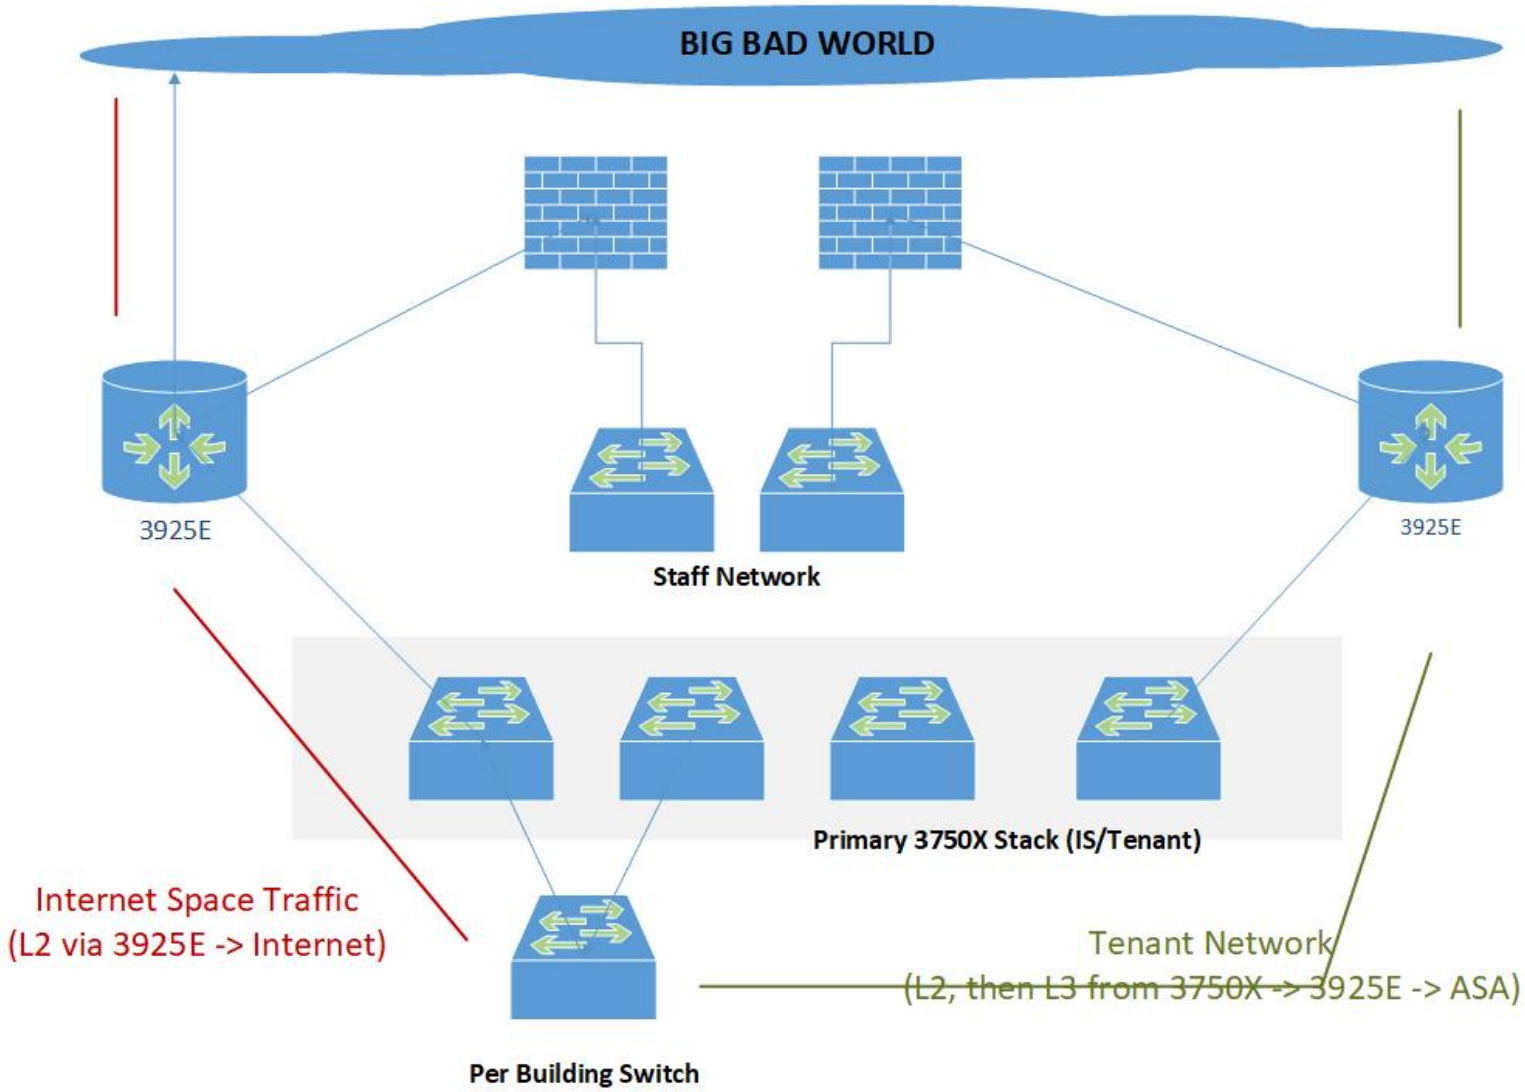

The Rollout

The v4

- We did v4 first - we wanted to know we had a stable platform before we started adding in v6
- Firewalls went in as a direct replacement, staff was upgraded, new staff Cloud servers were built, new switches for tenants put in alongside routers
- Tenant network joined to New Tenant network via l3 routed link and routing changed to route from firewalls via 3925E's then to new and over to old
- Tenants migrated over a number of early mornings (6am – 8am) building by building
- No real downtime to clients due to windows – everyone happy

- Admittedly we did play with 6-in-4 tunneling with anyconnect and Nat64 at various points (on non used networks)

The Rollout

The v6

- Firewalls enabled, then staff – with full Windows entirement (despite much rumbling from the Windows guys)

- RDS infrastructure tested on the RIPE v6 only SSID – it worked

- Customer told we'd enable v6 in the near future

2 weeks later a conversation:

Customer: 'Are you still planning to enable ipv6'

Me: Can you ping google for me ?

Customer: What's this thing which colons in it ?

The Pseudo Automation

```
foreach ($tdo as $key=>$value) {
# ACLs
 printf("ipv6 access-list v6_t_vlan%03d_in\n",$key);
 print " deny icmp any any router-advertisement sequence 10\n";
 print "permit icmp any any sequence 20\n";
 print "permit ipv6 2A02:2B38:4:" . $value . "::/64 2A02:2B38:4:101::/64 sequence 30\n";
 print "permit udp any any eq 546 sequence 40\npermit udp any any eq 547  sequence 50\n permit ipv6 fe80::/10 any seq 60\n";
 print "deny ipv6 2001:db8:4:" . $value . "::/64 2001:db8:4:100::/56 sequence 70\n";
 print "permit ipv6 2001:db8:4:" . $value . "::/64 any sequence 80\n";
 print "deny ipv6 any any sequence 90\n\n";
}

foreach ($tdo as $key=>$value) {
 # DHCP range
 print "# vlan " . $key . "\n";
 print "subnet6 2001:db8:4:" . $value . "::/64 {\n";
 print "range6 2001:db8:4:" . $value . ":b1c::100 2001:db8:4:" . $value . ":b1c::200;\n";
 print "}\n\n";
}


foreach ($tdo as $key=>$value) {

# interface config
print "int vlan " . $key . "\n";
print "ipv6 enable\n";
print "ipv6 address 2001:db8:4:" . $value . "::1/64\n";
print " ipv6 nd prefix 2001:db8:4:" . $value . "::/64 no-advertise\n";
print "ipv6 nd managed-config-flag\n";
print " ipv6 nd other-config-flag\n";
print "ipv6 dhcp relay destination 2001:db8:4:101::2 Vlan2\n";
print "ipv6 dhcp relay source-interface Vlan" . $key . "\n";
printf( "ipv6 traffic-filter v6_t_vlan%03d_in in\n",$key);
```

The Fun along the way – The Tenant Network

- Slowly started rolling out the tenant network – switched to OSPFv3 from the ASA's down
- Datasheets don't always tell the absolute truth
- V6 feature set not as mature as v4 – some missing features – eg: Lack of HSRP global v6 for VIP, required code upgrade
- Found a nasty memory leak with the 3750X's – somewhere between resources, vrfs and OSPFv3 within them – had to design around
- Security is interesting, some caveats but stills secure
- TCAM split on the 3750X – being careful about MAC/Route limits (required a covering ACL rather than individual per SVI ACL Set)
- Ultimately a L3 switch is not a router – expect caveats along the way

The Fun along the way – Internet Space

- Enabled the Atlas probe – 1st in Sunderland, as far as we know the site was also the first v6 enabled campus in the NE –still more than likely is one of the few
- Lots of things we don't control – still a moving target 4 years later – pragmatism required
- Customer now uses Igaware (Linux SBS type system, no v6 – I keep trying)

And then we were done – Oh, wait

- New Site – finally completed, how to work out best way to integrate – VRF-lite ? – staff primary, internet vrf …
- No v6 within VRF-Lite requiring switcharound on VRF's to allow where I wanted the most v6 to be the main v6 routing table
- No budget for 3925E line cards, had to use the 3750X's for the new site – required tweaking MST instances to have both links active and BFD in OSPF ← never, ever do this unless you have too! (do not route over layer 2 spt links)
- CPE didn't support v6 – despite saying they did (Disti had hardware v3 sales blurb but supplied v1 hardware) – gradual swap out as timing/budgets allow

There's always some level of truth in Statistics ….

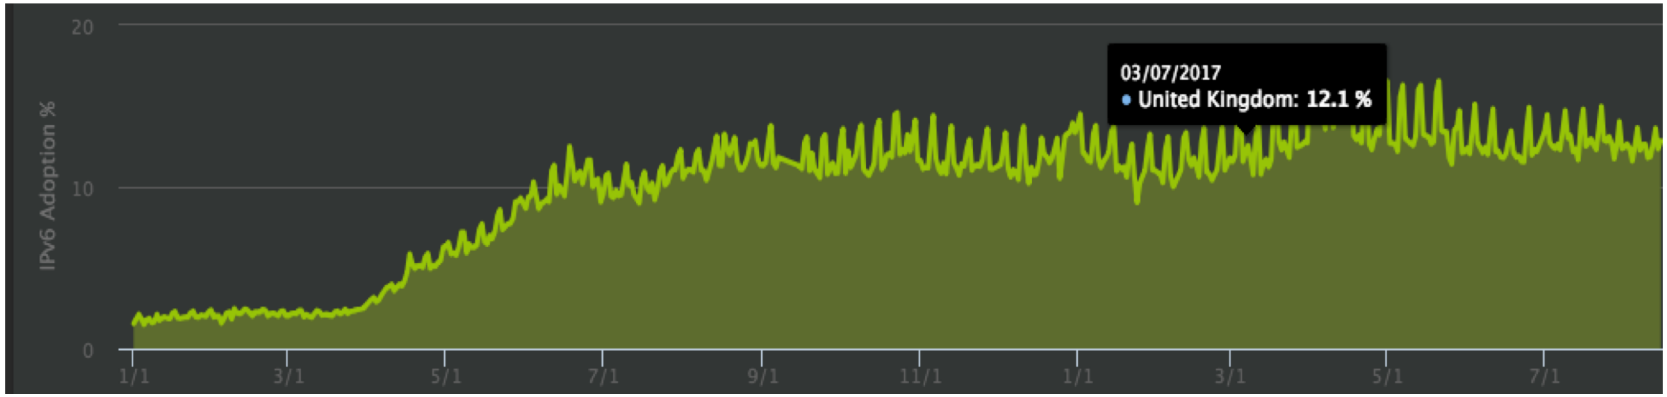

- Source: Akamai

- UKNOF/v6 Council/Industry content showing rise in v6 traffic, I wasn't seeing it – netflow logs backed this up – why ?

- How to 'force' more traffic ? – v6 enabled a pop3/imap/smtp server used by a number of tenants to see if I could see more traffic

- The 'no one can send email' phone call…. (smtp auth acl - oopsie)

- Still saw only a minor subset of the v6 enabled clients in logs … started looking at routing & DHCP…. *lightbulb*

The DHCPd Oopsie…

- DHCPd was forthcoming - some log entries from dhcp6 – unable to assign prefix / no prefixes available

- Guessed the many little netgears / dlinks between the end clients and our infrastructure were acting in routed rather than AP mode :/ - ugh, Prefix delegation required…

- Tried to manually enable – couldn’t get DHCPd to work properly (old, CentOS 6 version)

- ….. Remembered another UKNOF presentation

Revisiting how we implemented DHCP

https://indico.uknof.org.uk/event/30/contribution/14/material/slides/0.pdf

Kea Introduction (UKNOF30)

- Built a Debian 8 VM, wrote a basic kea config – routed another /48 – used another PHP script to generate the rest of the kea scopes

```php
<?php

$text = "";
// iterate through vlans ($x = vlan)
for ($x = 4; $x < 218; $x++)
{
    $text .= sprintf("#Vlan %d - 2001:db8:4:1%02X\n",$x,$x);
    print "{\n";
    printf("\"subnet\": \"2001:db8:4:1%02x::/64\",\n",$x);
    printf(" \"pools\": [ { \"pool\": \"2001:db8:4:1%02x:b1c::/80\" } ],\n",$x);
    printf(" \"pd-pools\": [ { \"prefix\": \"2001:db8:8:%02x0::\", \"prefix-len\": 60, \"delegated-len\": 64 } ]\n",$x);
    printf("},\n");
}

print $text;

?>
```

- Another script to change SVI DHCP relay server, then lots of delegations in logs within 20 mins, PD relay agent on the 3750X worked flawlessly – thankfully one feature which did work as it should

We have charts and graphs…

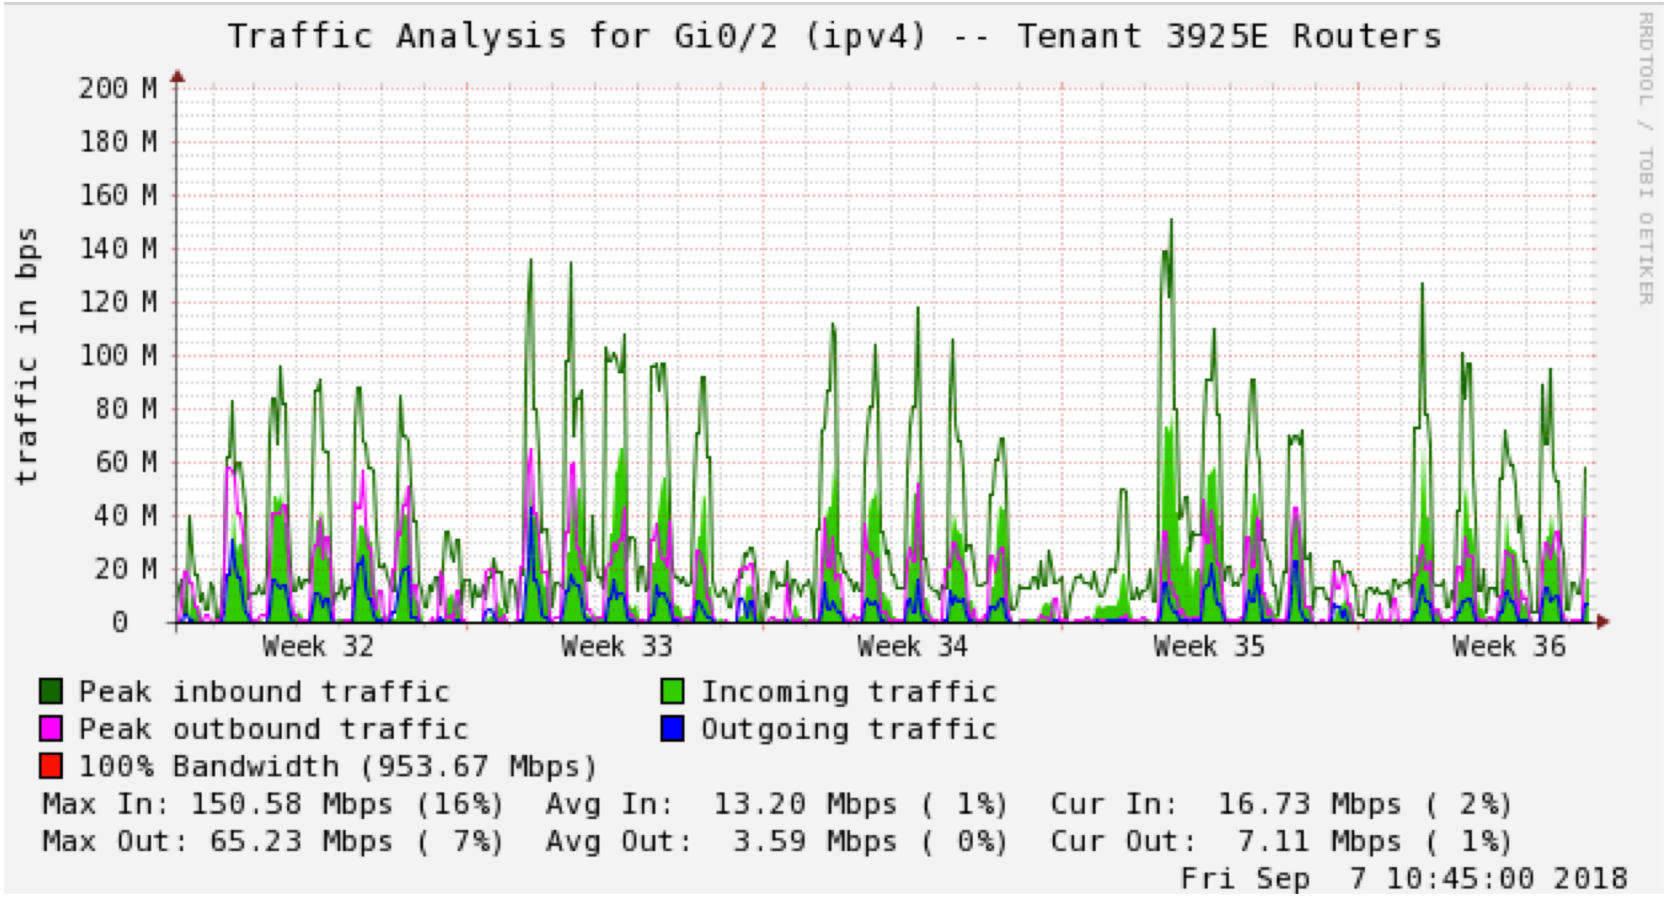

We have charts and graphs…

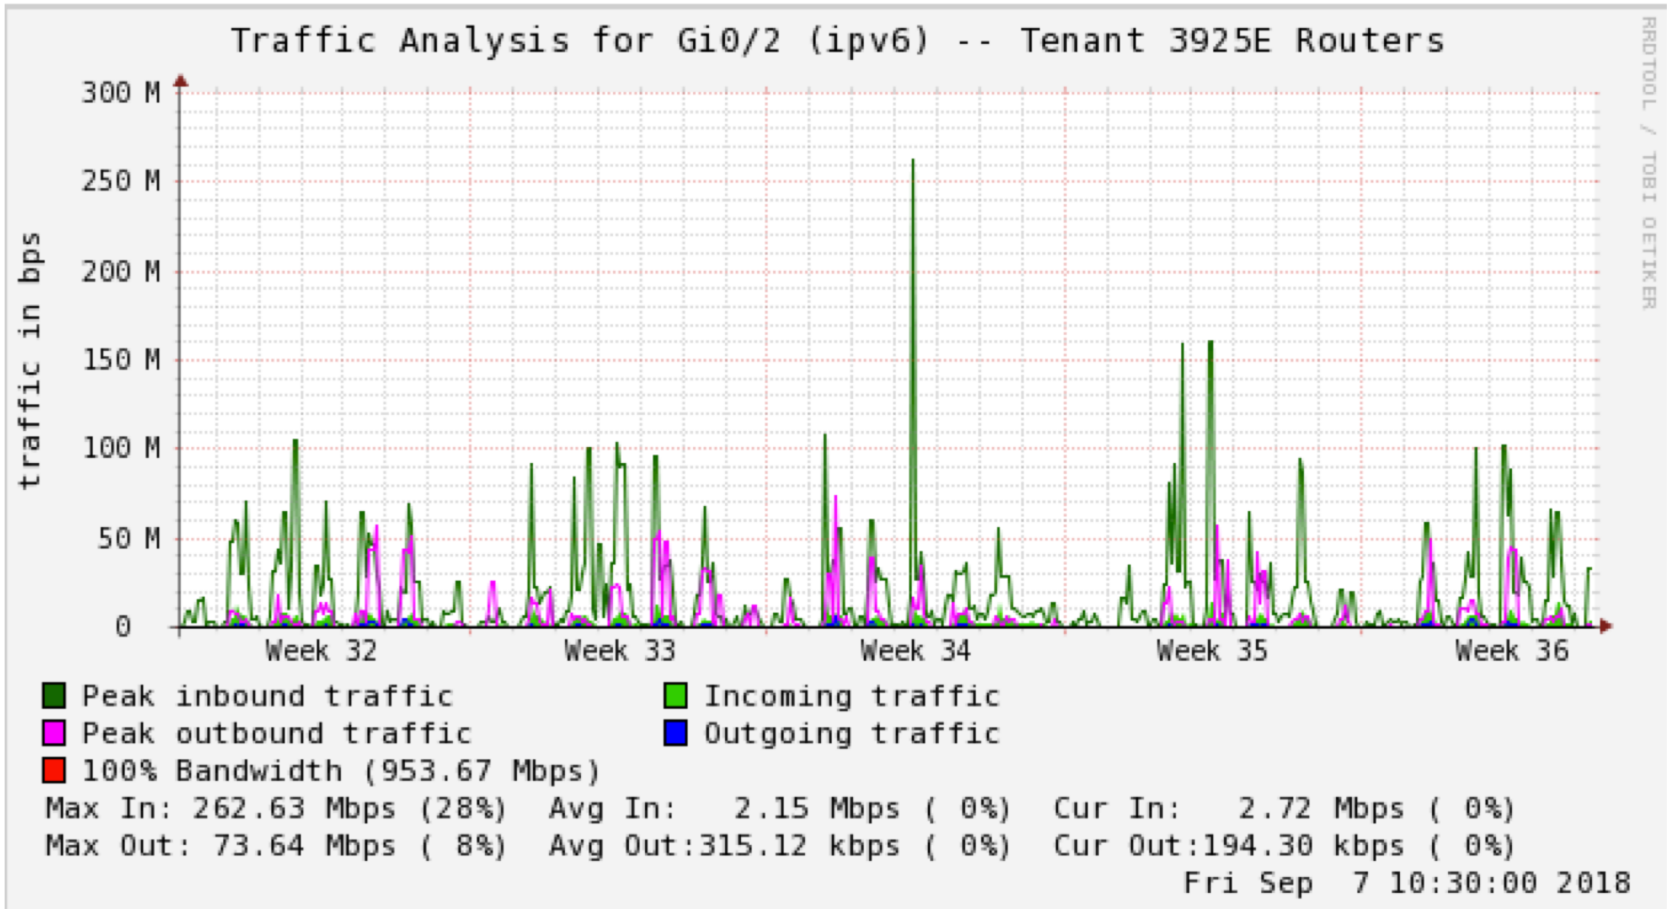

Best Days: 45% of traffic is v6, worst is 5% - average at 16% I can live with
(non v6 routers still and non v6 client endpoints too)

The Business Case for ipv6
(when you have lots of ipv4 and NAT)

This page intentionally left blank

The ‘I Told You so’ moment….

Another Customer:
’We’ve just bought a new door entry system, its Chinese and it only supports ipv6
we need to roll it out ASAP, can you help ? …’

Me:

‘Of course we can …’

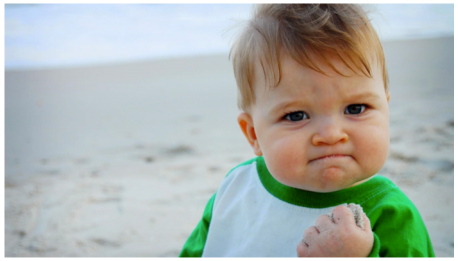

Project completed a week later..

- Some tweaks to statically address servers
- 1 legacy application using ipv4 broadcast for SQL Servers (Vendor issue)

The Summary …

- Do infrastructure companies even ask the v6 question ? Probably not, should they, probably.

- Education still a factor – the initial hurdle of the v6 lightbulb moment across teams
- There are ways to test without having lots of routers to play with
- It mostly does just work
- Remember your v6 security, esp FHS
- Your customers *WILL* start asking you for this at some point – we're *finally* starting to see requests

Any Questions ?