IPv6 Only Hosting

Pete Stevens
Mythic Beasts Ltd

What's wrong with IPv4?

- 2005: One IP per server.
- 2010: One IP per VM. Single server now requires ~ 50 IPs.
- 2015: Ideally one IP per container. Single VM now requires 30+ IPs. Single server can consume 1000+ IPs.
- This is unaffordable – Overlay networks on overlay networks. RFC1918 inside RFC1918. NAT inside NAT.

The seven(ish) layer OSI model

- Layer 1 : physical
- Layer 2 : ethernet
- Layer 3 : UDP
- Layer 2 : overlay ethernet / VXLANS
- Layer 3 : UDP
- Layer 2 : overlay flannel / dockernet etc.
- Layer 3 : TCP
- Layer 4+ : HTTP et al

Economics

- A new hosting company can get a /22 of address space.
- VM prices are ~ £10/month.
- A new VM hosting company is limited to £100kpa income per year before it runs out of IPv4 addresses
- We could offer £1/month virtual servers / containers if IPv4 addresses were free.
- IPv6 addresses effectively are free!

Computers get cheaper

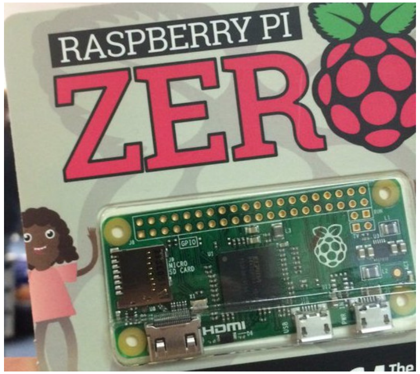

- 93.93.128.1

This computer costs $5

This IP address costs $10

IPv6

- Our VMs can talk IPv6 or IPv4, there's both on the network.
- IPv4 allocated statically and via a static dhcp server.
- Allocate customers a block of IPv6 addresses
- SLAAC doesn't give predictable server addresses – hopeless for inbound services
- SLAAC makes every machine auto-configure IPv6 even if they don't want it – customers go mad.

IPv6 only hosting

- Initially static addressing
- Need IPv6 resolvers so you can download updates
- Advertise gateways with IPv6 route advertisements
- Problems with mirror services – not all package mirrors have IPv6, the mirror directors aren't protocol aware
- Many other services don't have IPv6 (twitter, akismet, newrelic, anything in AWS/Azure etc.)
- Not very useful unless everything you talk to is also IPv6

NAT64

- Normal resolver
- dig AAAA www.cam.ac.uk
  - no answer
- NAT64
- dig AAAA www.cam.ac.uk +short
  - 2a00:1098:0:80:1000:3a:836f:9619
- Our resolver proxies 131.111.150.25
- Outbound to IPv4 hosts works!

Inbound Proxy

- proxy.mythic-beasts.com
- Haproxy, auto configured from our control panel
- IPv4 / IPv6 connections terminate on our load balancers, we forward them to the IPv6 only back end.
- Forwards any SSL service that uses SNI
- Forwards HTTP
- Doesn't yet forward ssh

DHCP6

- Typing v6 addresses is very annoying
- Configure a DHCP6 server to auto allocate addresses to machines
- Virtually all client implementations expect to get IPv6 + DNS servers + gateway from the DHCP6 server
- If you don't supply the gateway and expect it to pick it up from the router advertisements mostly it doesn't
- This is very annoying, back to VRRP for a fixed gateway

Useful

- We have an IPv6 only VM
- It has full outbound via NAT64
- It has inbound for SSL & HTTP via our proxy service
- You can host real websites with it
- Like this one, https://www.raspberrypi.org/
- 40+ VMs, we don't have to route a layer 2 private network between data centres – they can talk to each other over IPv6 +SSL.

Management services

- We back up managed customer machines
- Enable IPv6 on the backup service, add an AAAA record – easy.
- We monitor customer machines.
- Add control panel functionality to put IPv6 addresses in
- Update libwww-mechanize for perl to a version that supports IPv6

Management Services

- Munin graphing
- Update munin to the latest version
- Add v6 to the munin server
- Watch all of your graphs break
- Add the ACL to munin-node.conf on every customer machine to allow our v6 address to communicate with the agent and update ip6tables.
- Add “allow ^::ffff:a\.b\.c\.d$” for syntax hilarity
- This was a boring few days

Management Services

- Update our code that auto-magically generates all of our munin config to correctly escape IPv6 addresses
- Address a.b.c.d
- Address [e:f:g:h:i:j:k:l]
- Find the other bits in the control panel where people had asserted that each machine had at least one IPv4 address and fix them
- Turn on IPv6 by default on all new customer installs

Management Services

- We log reporting data daily
- The source address identifies the machine it came from
- We already had a horrific blob of code to deal with machines behind NAT firewalls
- Now another nasty blog of code to match up reports coming from v4 and v6 addresses that belong to the same host

Management Services

- Jump box
- Automatically picks the correct customer key for logging into a host
- Hosts now have multiple addresses – need to mine the database further
- Since all requests go via our jump box, once our jump box has IPv6 we can access every IPv6 only server even if we don't currently have IPv6 natively.

Let's Encrypt

- Free SSL certificates
- IPv6 support due tomorrow
- Out of the box works perfectly with NAT64 + v4 Proxy
- Probably works with our DNS API pure v6 only but haven't yet tried it

Deploying new services

- Setting up scripted customer installations
- Logic for single stack v4, dual stack v4/v6, single stack v6 was getting twisted.
- Simple solution, only support single stack v6.
- Add a v4 address at the end only if required.
- New management services can be v6 only.

Customer incentives

- We itemise IPv4 connectivity at £20 per server per year
- We're starting to get accounts departments asking 'if they really need IPv4'
- Increasingly the answer is no
- The easiest way to persuade a techie to deploy IPv6 is make the alternative explaining to the accounts department why the additional expense is necessary

Customers

- Technical professionals learning IPv6 – proper supported testbed.
- DNS anycast services, BGP etc.
- Non technical managed customers who want the discount
- Roughly 5% of our servers are now IPv6 only

April Fools Pi

Gwiddle

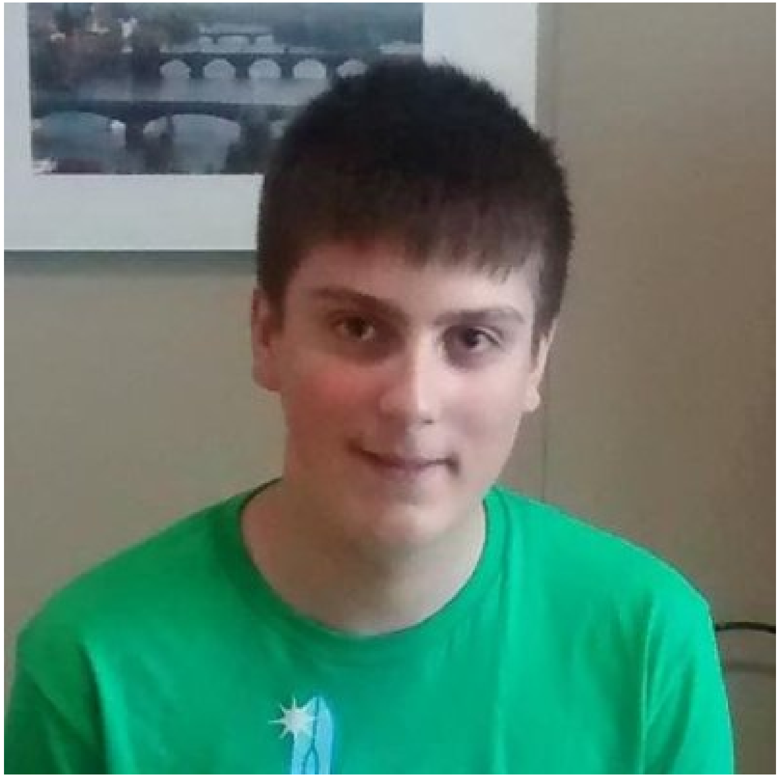

- Joshua Bayfield
- CEO Gwiddle Web Hosting
- Free accounts for educational use
- Age 15

Asking for help & reply

- “I am approaching you today to enquire about the possibility of Mythic Beasts supporting Gwiddle's mission...”

- “All your services run on IPv6 only VMs fronted by our IPv6 proxies, we'll give you a single IPv4 address for logging in for management reasons...”

Questions?

- http://blog.mythic-beasts.com/
  - We blog all of our updates
- https://twitter.com/Mythic_Beasts
- Ask me directly pete@ex-parrot.com