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

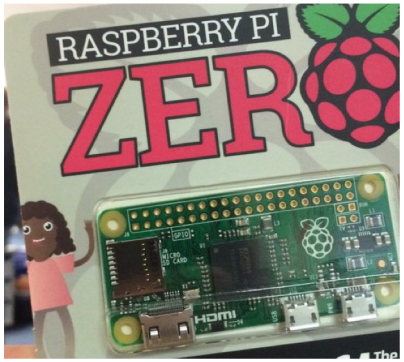

•93.93.128.1

This computer costs $5

This IP address costs
~~$10 $20~~ $24

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

- 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 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

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.

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.

Link local

- An server image with no networking configured has link-local addressing.
- So you can pull your configuration from a server using the link local address and write the static configuration into the VM on first boot.
- No need to try and make DHCP6 or SLAAC work, because to a first approximation they don’t.
- You can support IPv6 and Ipv4 with a single call on bootup.

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

Dual stack is rubbish

- Dual stack has significantly increased configuration complexity over single stack.
- NAT also has significantly increased configuration complexity over single stack IPv6.
- Nobody sensible will deploy dual stack by choice if single stack will do.
- Motivation for IPv6 only is it's less horrible than NAT.

Does it work in practice

- Only if your application and operating system support single stack IPv6 operation.
- We're a Linux shop, all our managed services are Debian/Ubuntu or CentOS.
- Each customer gets a /64 in each data centre.
- The customer must have IPv6 enabled for our managed services.

LAMP

- Linux / Apache / Mysql / PHP
- All work well with IPv6.
- So do Python, Perl (with Socket6).
- So does Nginx.
- So does PostgreSQL.

Wordpress

- Started here with www.raspberrypi.org in 2012
- Core works fine.
- Plugins are of highly variable quality and might not work
  - With IPv6...
  - At an acceptable speed...
  - At all...

Login Lockdown

README.md

login-lockdown-mb

Fork of Wordpress Login Lockdown plugin with support for IPv6.

The original plugin has a serious bug, whereby if there are any IP addresses at all that have been locked down, all attempts to login from any IPv6 address will be blocked. This fork fixes that issue, and performs blacklisting of IPv6 addresses at the /64 level for too many login failures. IPv4 addresses continue to be blacklisted in /24 blocks.

Plugins

- WooCommerce – sell stuff through Wordpress.
- Gravity Forms – easy form filling.
- SuperCache – performance.
- Yoast SEO – game Google for extra hits.
- Akismet – spam filtering.
- Nearly every well used plugin works fine.

Wordpress

- We offer Wordpress as a managed service.
- We apply RIPE rules, you need a technical justification for an additional IPv4 address.
- But we default to 0 IPv4 addresses rather than 1.
- Valid technical justifications are very rare.

Mediawiki

- So far single server installs.
- Apache + MySQL + PHP
- This all works fine.
- Visual editor, which means you don't have to learn wiki syntax adds node.js and parasoid.
- Internally talks to [::]:8000
- Only public facing install we have is https://wiki.uknof.org.uk

Let’s Encrypt

- Free SSL, better than paid for ones.
  - No customer invoicing and payment hassle.
  - No validation via email – done through the webserver or DNS.
  - Automatically renew.
  - Works out of the box in our setup.
  - Also now works pure IPv6.

LINX

- The London Internet Exchange has a new Wordpress based site on our IPv6 only infrastructure.
- They have full IPv6 support internally.
- We forgot to enable our Ipv4 proxy for their dev site.
- 28 days before anybody noticed!

rpf.io

- URL Shortener for Raspberry Pi
- Asked bit.ly if we could use our own domain on their service
- Bit.ly has a slightly nicer interface that a .htaccess file

Bit.ly

Ben Nuttall <ben@raspberrypi.org> 18/12/2015

to joe

Hi Joseph

Can you confirm the price is $695/month? Did you miss a decimal point? $6.95 would make more sense.

YOURLS

- PHP app that implements URL redirection.
  - User accounts we can hand to employees.
  - Trivial to install, works fine on a v6 only VM.
- If you hate PHP and want to write your own in Python the uknof.uk shortner works fine too.
  - https://github.com/uknof/shortener

Etherpad

- Realtime shared document editing.
- Written in Node.js
- Listens fine on [::]:9001
- Apache forwards from [::]:80 to it
- Outbound it’s a bit special

Node.js in Debian Jessie

- HappyEyeballs++
  - AAAA registry.npmjs.org (NAT64).
  - A registry.npmjs.org.
  - Fails to connect over IPv4 and stops.
- Workaround, put AAAA registry.npmjs.org into /etc/hosts
  - Now it connects outwards through our NAT64 proxy.

Indico

- Conference booking system, Node.js again.
- Also adds redis (works fine) and memcached (works fine).
- Cache early, cache often…
- https://indico.uknof.org.uk/
- Other organisations are now using this too.

SugarCRM

- Was open source CRM setup, now closed source.
- Apache2 + elastic search (java/tomcat) + memcached + mysql.
- Split site – high availability, each instance has a hot spare in another data centre
- Massive single tenant infrastructure, each client gets their own pair of installs

Containers

- Implemented with LXC – Linux containers.
- Spin up a pair of containers for each end customer
- IPv6 address each, firewall so they can only talk to each other. No IPv4.
- MySQL replication with SSL, internal traffic with SSL/SSH.
- No VPN needed, no layer 2.

SugarCRM / Containers

- Much simpler than the protoype with Docker + CoreOS + overlay networking + custom persistence layer.
- Everything worked out of the box.

Hadoop

- https://wiki.apache.org/hadoop/HadoopIPv6
- Doesn't work IPv6 only.
- Most instructions still tell you to disable IPv6 completely (no link local or ::1).
- Things have improved, it now works on a machine with IPv6 enabled.
- Wondering about v4 in v6 tunnel and RFC1918 space just for this special snowflake.

Shared Web Hosting

- Joshua Bayfield, 16 years old, www.gwiddle.co.uk
- Web hosting platform, free of charge for people in full time education.
- We gave them free VMs.
- … but no IPv4 addresses.
- Has around 1000 school age children running websites in an IPv6 only environment.

Pi Rack

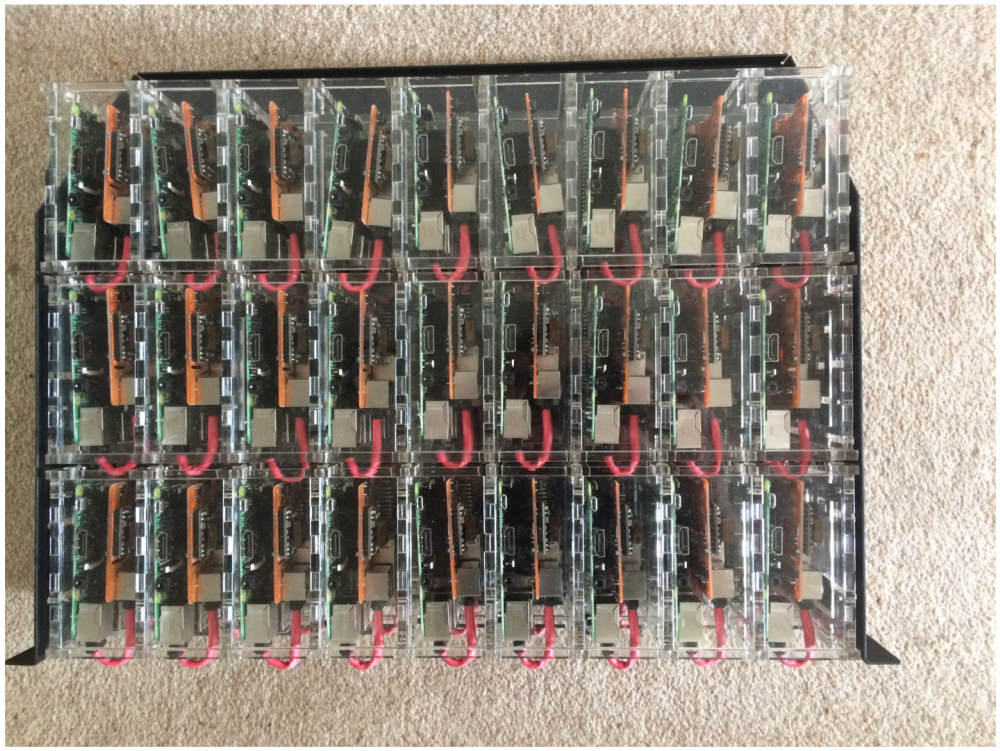

Pi3 Hosting

- 4U of rackspace including the switch
- 108 Pi3s – 108GB RAM, 432 cores
- All netboot and PoE
- Just one wire to each Pi
- 2-3W each

Pi3 Network

- /30 of RFC1918 space for network filesystem.
- v4 address + v6 address on a tagged vlan.
- **Nice idea but it doesn't work.**
- Bug in the rom – a tagged vlan crashes it.
- Computers that don't boot are very hard to sell.
- Add a /30 for each Pi3 - $80 of IP space to turn on a $35 computer. A /64 is free.

Pi3 Networking

- The v4 costs are too high for a proper setup.
- /31 or proxy arp make it possible, but horrid.
- Educational – so do it right, hacks come later.
- V6 only. We don't support direct V4.
- Ssh.petespi.hostedpi.com:XXXX → ssh (4&6).
- Www.petespi.hostedpi.com → SSL & HTTP (4&6).
- Petespi.hostedpi.com → v6 only.

Questions?

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