

IPv6 Only Hosting – part 2

Pete Stevens
Mythic Beasts Ltd

Virtual servers

- .Our VMs can talk IPv6 or IPv4, there's both on the network.
- .IPv4 allocated statically and via a static dhcp server
- .IPv6 allocated statically with a /64 per customer
- .HAProxy auto-configured for inbound HTTP/HTTPS and other SSL based protocols
- .NAT64/DNS64 for outbound connectivity

Dual Stack is rubbish

- .You can run IPv4 and IPv6 on the same host
- .IPv6 only is cheaper as we charge for IPv4 addresses
- .Our management tools assume IPv6 on every host we manage
- .You can either do dual stack, or save yourself money and time by single stack IPv6

What doesn't work at all?

- Mail. Lots of providers only deliver/receive mail from IPv4. Forwarding through another MX removes your ability to reject at SMTP time
- Ftp. This really should be turned off everywhere but 'web developers'
- Hadoop

What doesn't work well

- Node.js
- Docker
- Snap
- Many things will always prefer v4 if available, so you have to break v4 completely to get them to work v6 only

Things we got wrong

- /64 per site LAN
- Gave customers a /96 from it
- This worked nicely until blocklists settled on /64
- Now /48 per customer
- Renumbering was not much fun

Filtering

- IPv4 all appears to originate from the same IPv6 address
- Less of a problem now ~40% of traffic comes in over v6
- IP blacklisting occasionally blocks the proxy service and the whole of IPv4
- Can't selectively block abusive IPv4 addresses with the firewall
- Have to do in apache after proxy protocol decodes the source IP

Filtering #2

- Proof of concept proxy that doesn't lose IP information
 - Maps the source IPv4/port into the v6 address
 - 32 bits = IPv4 address
 - 16 bits = destination port number
 - 16 bits = unused
- Could do with this in HAproxy

More things that went badly

- Large on-link prefixes

- /64 per customer means they can use a different IP for every outbound connection

- And use this to evade rate limits while aggressively scanning other providers

- 2000+ new IPv6 addresses appearing per second causes neighbour discovery issues

- Inbound address space scanning also hurts

Neighbour discovery woes

- Your router can't fit a /64 of v6 address → mac address mappings
- Ram exhaustion, garbage collection, network stalls
- Linux can stall the entire network stack while the v6 neighbour table garbage collects for tens of milliseconds
- Anyone in the same layer 2 domain can force you to neighbour discover billions of addresses

Why do we do neighbour discovery/ARP?

- To match IPv[4 | 6] addresses to mac addresses
 - Our billing database already knows about every server/VM, what port it's attached to and it's mac address
 - Any answer from arp/nd that doesn't match billing is a security violation and should be ignored
 - So, why do we do neighbour discovery / arp at all?

Static config in the switch

•Statically configure this in your switch with every port a different vlan

-ipv6 neighbor <link-local> Vlan11 <mac>

-Arp <ipv4> <mac> arpa

•BGP daemon knows the v6 range → link local mapping

-route <b6 range> via <link-local>%vlan16 {bgp_med = 0;};


-When the link comes up, advertise out

All layer3

- All /64s and IPv4 addresses are now portable anywhere in our new fully routed network
- Currently 2/5 data centres enabled, rest coming Real Soon Now

single/dual/single stack?

- Infrastructure and customers
- where Mythic Beasts manage
- the server



Not quite as good as it seems

- V6 only tends to be
 - Simpler applications / newer builds
 - Distributed rather than monolithic applications

Retiring legacy things is hard

- We're slowly picking them off

- Another 3% due to go by end of this month – v4 only with CentOS 5 (end of life 2017). '6 month' migration plan started in 2017, completed Sept 2022

- In preparing this talk I found a handful of Mythic Beasts public facing services that aren't dual stack

- Now due to be fixed

V4 only will always be here

- Some things are too expensive and too hard to migrate to v6
 - We recently installed a new IPv4 only setup for a bank to do SWIFT transactions
 - They have to pay extra for IPv4
 - They're a **bank**, they're never going to run out of money for IPv4
 - \$100/month per IPv4 would be **fine**

Quick interlude from finance land

- This is not financial advice. Do not make investment decisions based on technical conference presentations.
- When we started v6 only options in 2014 we set the price of an IPv4 add on at £2/month
- Appears we set an industry standard

IPv4 rental prices

- AWS elastic IP: \$0.005/hour = \$3.60/month
- Azure static IP: \$0.0036/hour = \$2.60/month
- Google Cloud static IP: \$0.004/hour = \$2.88/month
- OVH: \$0.0027/hour = \$1.94/month
- Hetzner: €1.70/month + setup
- Digital Ocean, free if in use: \$5/month when not
- Zen Internet: £0.83 /month

IPv4 address as income stream

• Suppose you ran a pension firm and you needed a steady monthly income stream to pay your pensioners

– 10 year UK Government bond = 3.5% return

– \$50 IP address rented at \$2/month = 48% return

– \$500 IP address rented at \$2/month = 4.8% return

– Obviously IPv4 addresses need a higher return than government bonds as they're higher risk, but >10x?

Do you need to worry about asset strippers?

- Fictional tier 1 ISP, enterprise value (shares + debt) ~\$3bn
 - Originates ~30m IP addresses
 - \$100 per IP address
 - 50% of the value is IP space!

Greed is good?



Computers keep getting cheaper




.93.93.128.1

This computer still costs \$5

This IP address now costs
~~\$10 \$20 \$24 \$30 \$40 \$50~~

Pi Cloud



Pi Cloud



mythic beasts

Pi Cloud

- .2x 3U Chassis mounted back to back
- .96 x 8GB 4 core Raspberry Pi 4 servers and two 48 port 1G PoE switches
- .Netboot with a /30 of RFC1918 space to get kernel + NFS filesystem
- .Cost prohibitive to give every \$75 computer a \$50 IPv4
- .Doesn't support v4 at all
- .All external comms is IPv6

IPv4 not allowed

.We do lose some customers

-If there was a ipv4 addon with this server, I would have gladly paid for it.

-I just purchased your Raspberry Pi hosting service, however only after paying I noticed the service is IPv6 only, which is very critical for my use case. After going back to the product page it is indeed mentioned there so I do realize this is really my fault. However, is it possible I can cancel and get a refund?

-Would you be able to proxy a UDP port on an IPv4 address to my Raspberry Pi host, like you do for web servers?

IPv4 security education

- PiCloud customers often security harden their Raspberry Pi servers by setting up stringent firewalls
- Things your IPv4 firewall can protect you from
 - Your filesystem
- This does not end well

/etc/hosts workarounds

- Node / npm application trying to send mail and install modules
 - 64:ff9b::37a:1db7 smtp.eu.mailgun.org
 - 2606:4700::6810:1723 registry.npmjs.org
- It sees the RFC1918 address for NFS and only tries v4 because 'v4 always works'
- Npm has dual stack sources and still won't download over v6

V4 workarounds

- To investigate

- DNS64 servers drop all A record requests (a customer implemented this in the resolver)

- tnat64, LD_PRELOAD library that intercepts all attempted v4 connections

- 127.0.0.1 seems to cause problems

- Clatd solves it but isn't packaged

- Need to install by default

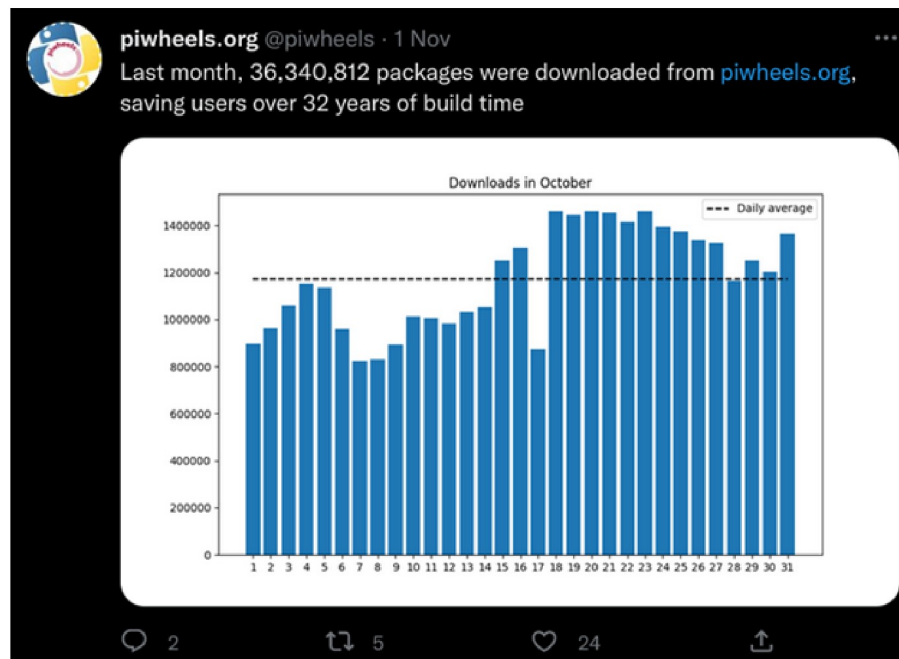
Is this just a toy?

- Industrial users build/test Raspberry Pi facing software
 - ~25m Raspberry Pis in industrial applications and commercial products so far
 - We've sold private Raspberry Pi clouds to be CI/CD services
 - On demand Raspberrys are used by universities for distributed system research/testing

Pi Wheels

- Precompiling numpy (numerical python) reduces install time on the Pi1 from 2.5 hours to <10s
- PiWheels precompiles all python modules for the Raspberry Pi
- Everything natively built on the Raspberry Pi
- Binary distribution direct from the Raspberry Pi cloud
- Uses our proxy service to dual stack the front end

Pi Wheels



Raspberry Pi Desktop

Raspberry Pi
Educational computing in the cloud

Pay for your server annually and get 12 months for the price of 10.

☒ Pay monthly
Billed monthly in advance

☐ Pay quarterly
Billed quarterly in advance

☐ Pay yearly
Billed yearly in advance

☐ On demand
Billed monthly based on per-second usage

	Price	RAM	CPU	NIC	Monthly Bandwidth
Raspberry Pi 3	£5.75	1GB	4 x 1.2GHz	100Mbps	1TB
Raspberry Pi 4	£9.25	4GB	4 x 1.5GHz	10Gbps	2TB
Raspberry Pi 4	£10.00	8GB	4 x 1.5GHz	10Gbps	3TB
Raspberry Pi 4	£12.50	8GB	4 x 2.0GHz	10Gbps	4TB

Disk space: 10 GB [20p (2.0p per GB)]

Raspberry Pi
Our Raspberry Pi hosting service gives you access to a dedicated, 64-bit quad core ARM processor with 1GB, 4GB or 8GB of RAM, backed by network storage. The server is supplied with the Raspbian operating system pre-installed.

Features

- Raspberry Pi 3 or 4 — quad-core 64-bit ARMv8 processor

£5.95 per month and no setup [Continue](#) [Show price inclusive of UK VAT](#)



mythic beasts

Mastodon



- Decentralised social network
- Suddenly become very popular
- Every mastodon server needs to talk directly to every other server

Mastodon just stopped being a techy toy



Mastodon

- Now over ~~4000~~, 8000 servers
- They all need to talk to each other, so all need IPv4 until they're all dual stack
- Or do they?
 - V6 only + inbound proxy for IPv4 + DNS64 + NAT64 works since v4.0.2
 - In the last two weeks all the v4 bugs that required clatd have been fixed


Mastodon



Alasdair Allan @aallan@mastodon.social

41m

So @Raspberry_Pi just spun up a @Mastodon instance. I talk about why we're doing it, and (the fun bit) how we did it. Because it's running on a Pi in the Sky ☁️ on a #RaspberryPi 4 with #IPv6 only networking hosted at @beasts. #RaspberryPi is now part of the #TwitterMigration. raspberrypi.com/news/an-escape...



1+ 2 3 4 5 6 7 8 9 10 11 12 13 14 15 16 17 18 19 20 21 22 23 24 25 26 27 28 29 30 31 32 33 34 35 36 37 38 39 40 41 42 43 44 45 46 47 48 49 50 51 52 53 54 55 56 57 58 59 60 61 62 63 64 65 66 67 68 69 70 71 72 73 74 75 76 77 78 79 80 81 82 83 84 85 86 87 88 89 90 91 92 93 94 95 96 97 98 99 100



Raspberry Pi @Raspberry_Pi · 38m

Like everyone else, we've been watching the goings on at Twitter.

After a lot of debate here at Pi Towers, we've now spun up our own Mastodon instance. The best thing about it? It's running on a Raspberry Pi 4 hosted @Mythic_Beasts.



raspberrypi.com

An escape pod was jettisoned during the fighting - Raspberry Pi
We're spinning up our own Mastodon instance. Come and join us!

6 22 86

mythic beasts

Mastodon

- Mythic Beasts have a v6 only private mastodon instance
- Full bidirectional interop with every other Mastodon instance thanks to nat64/dns64/v4 proxy
- Self host v6 only is <£8/month.
- No limit on the number of instances we can deploy on dedicated hardware
- We don't need to care that the rest of the world hasn't deployed v6 yet, we have another v6 only product

Questions?

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