Putting
IPv6
to work

North American IPv6 Summit

Grand Hyatt, Denver, Colorado
September 23-25, 2014

Rocky Mountain IPv6 Task Force
RMv6TF

Implications of IPv6-centric networks for DNS

Andrew Sullivan
asullivan@dyn.com

Isn’t this done?

- AAAA (since 1995)
- A6 (since 2000, already obsolete)
- Transport Just Works
- What else could there possibly be?

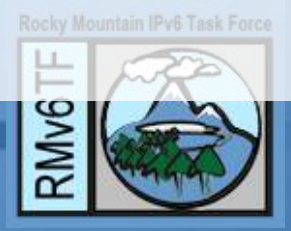

Ok, a little operations pain

- Stupid servers sending NXDOMAIN
- Dual stack + transition technologies
  - disable-aaaa-on-v4-transport
- NAT64/DNS64 is far from perfect
- Still mostly done

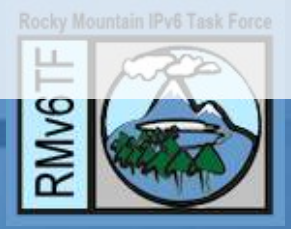

Not quite that simple

- Small networks no longer that small & no longer optional
- The more relied-upon the network is, the higher the expectation for reliability
- How do you build reputation in an IPv6 network?

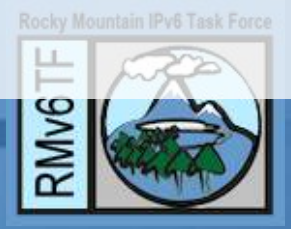

The v6-centric network changes things

- Networked device is decreasingly a general-purpose computer
- Mobile devices have poor network-troubleshooting interfaces
- Light switches have basically no such interface
- Most end-users of bathrooms don’t know about plumbing, either

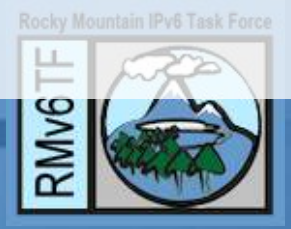

More changes

- “Small” networks become medium-sized, but still unmanaged
- Unreasonable to have to remember dozens of devices – service discovery critical
- Customer support: “Enter in the box 2001:cdba:0000:0000:0000:0000:3257:9652”

IETF: not totally wrong

- MIF produced multi-interface DNS guide (RFC 6731)
- HOMENET's architecture assumes multiple networks at a given "site"
- DNSSD trying to make Bonjour useful beyond local link

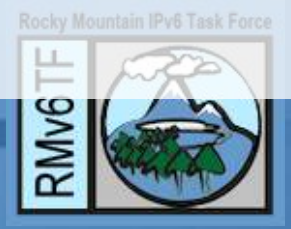

“Small” networks

- “Homenet” with multiple local links
- DNS service needed even within homenet
- Dynamic update is apparently too hard
  - dyndns, no-ip and similar services
  - CPE with DNS responder
- Service discovery uses more than just A & AAAA

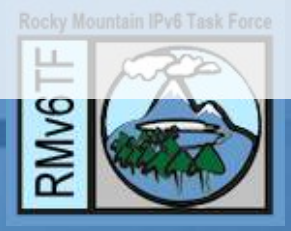

More on “small” networks

- DNS and mDNS proxies
  - Will expose things that weren’t intended to be
- Think of the opportunities for password mischief!
- Will reverse tree actually be important?
- Port 53 access & “typo interception”?
- DNSSEC?

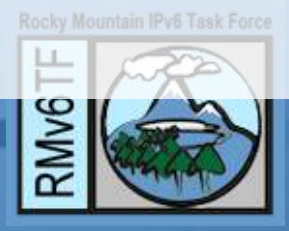

Always on means *always*

- Homenet becomes more important, it must be more reliable
- “Internet is down” cannot be the experience
- If your system is down, it will reflect badly on you, not the network

Who do ISP customers call?

- If a service is broken, call the service operator (maybe)
- If a device is broken, call the store (maybe)
- If the network is broken, call the ISP (maybe)
- Get your teenage child to look at it
- If none of those things work, you're stuck
  - Stop using whatever it is

DNS is not about lookup

- Nobody uses names directly any more
- Search, apps
- Perversely, this makes DNS more important
- Basic piece of infrastructure management

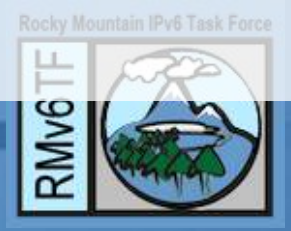

A service-management tool

- CNAME chains to final destination
- Useful in directing traffic globally
- Anycast works well with DNS
- A single eventually-consistent global DNS is important and desirable
  - Let's assume it's dead

At least there's no NAT

- We've heard the rhetoric
  - No more middleware mess!
  - No more NAT!
- The network is just not that clean
  - Outages in dusty corners
  - Outages in bright, well-lit rooms
  - Outages from network disputes

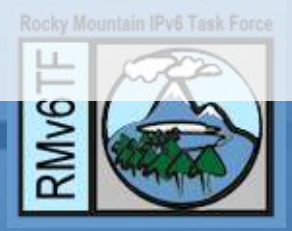

Local or remote?

- Lots of Internet end-users follows a similar pattern
- If all your YouTube, Facebook, Netflix works, then do you notice if remote things are broken?
- Who do you blame if your expectations are not met?

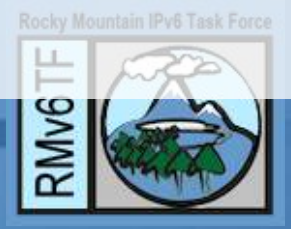

The “global” Internet

- Everyone worries about latency
- Caches and media servers are all over
- If we’re doing this right, the *content* is global and the *traffic* is still mostly local

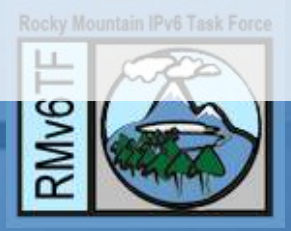

How much anycast is enough?

How many ASes?

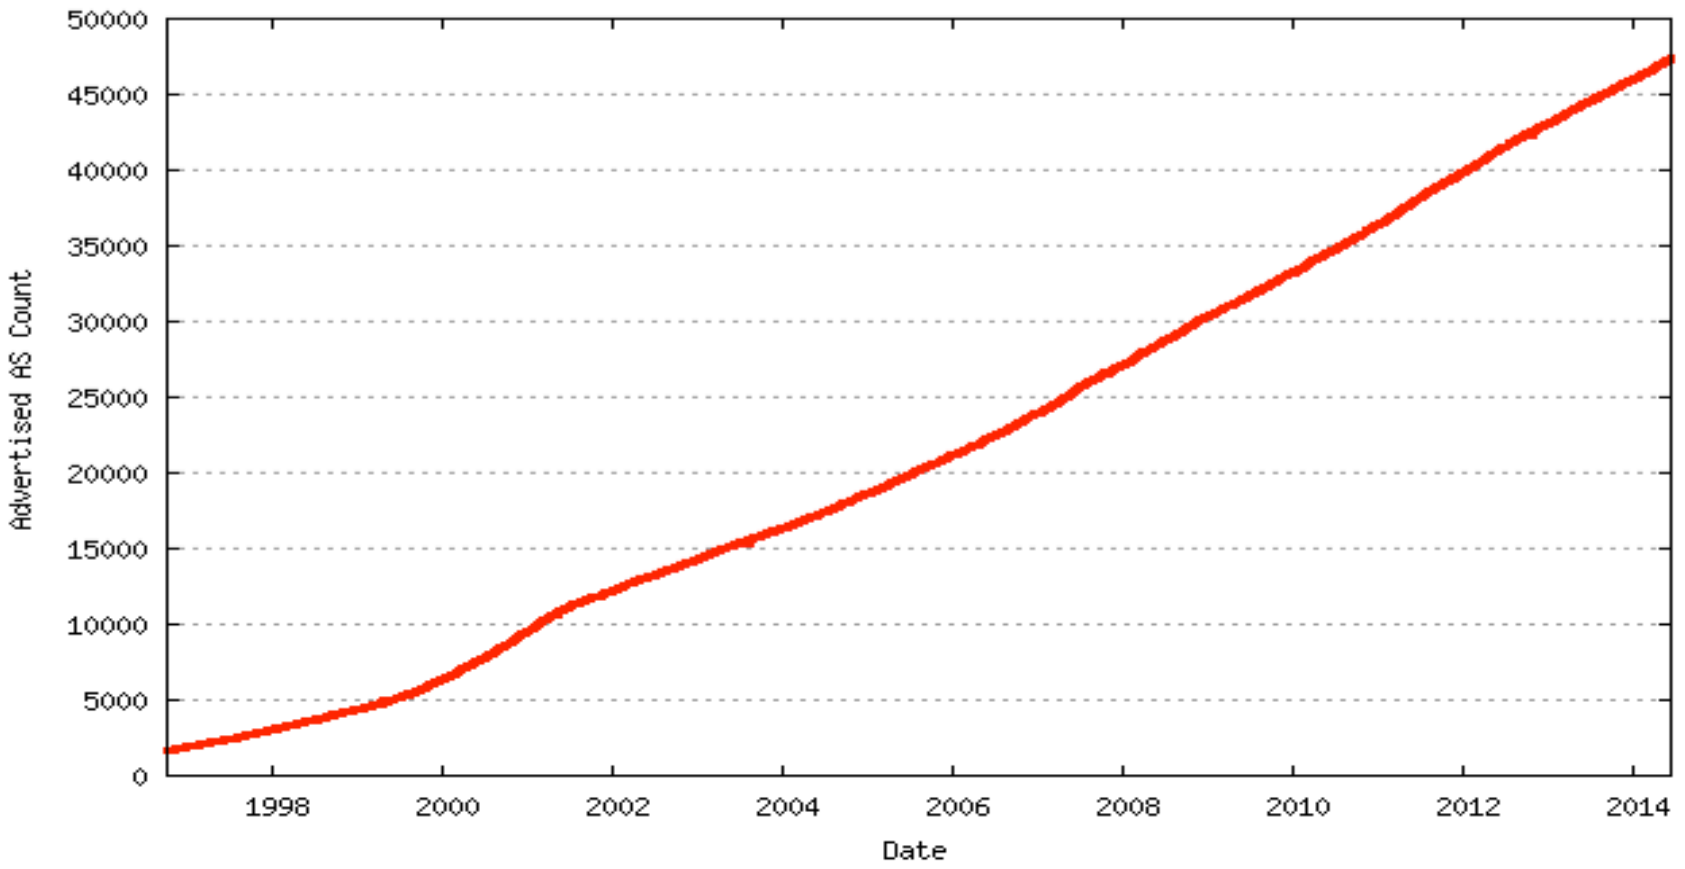

Globally-local scale

- To keep things local, we need to anycast at roughly 1 node per AS
  - Ok, any AS that wants one
- Keep attack flows on-net
- Keep latency as low as possible
- Need cheap hardware for this

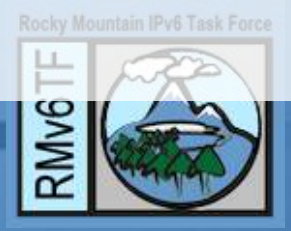

How is Dyn building these?

- Self-service
- Ship directly to sites
- Low power, 2 post, no shelves or rails
- Regular, frequent, bare-metal installs

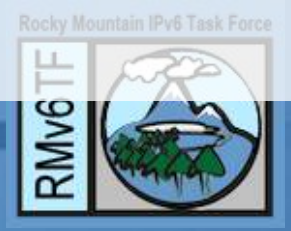

What can we do with it?

- Very widely-distributed DNS service network
  - Use this for service management
- Very widely-distributed service measurement system
  - Use this to inform the DNS answers/tricks

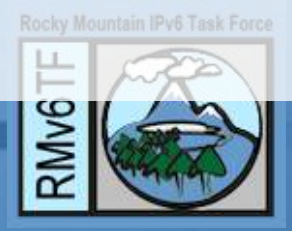

Effects

- In the approaching network reality, reliability expectations are high
- Keeping traffic local reduces latency
- Keeping traffic local contains attacks
- Very highly-automated service (re)configuration keeps availability high
- Combine with multiple CDNs for optimal availability and cost

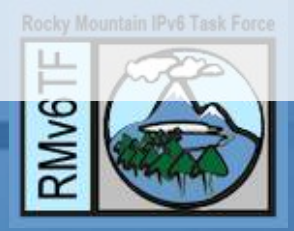

DNSBL: it's over

- DNSBLs are ingenious
- They're doomed: you'll never populate them in time
- Anyway, if you could, the user's address would change

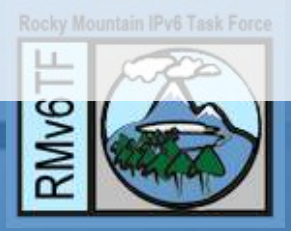

Reputation in a v6 world

- The DNS is still the database that's available
- New reputation systems that publish in the DNS are going to be needed
- Can't be based on address
- Reverse mapping (again)?
- Keys and other tokens in forward zones?

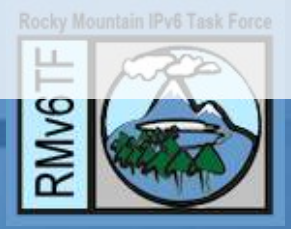

Would policy assertions help?

- Today little binding between name and expected behaviour
- DMARC is a step away from this
- Would general-purpose policy mechanisms help?
- If they were there, could we trust them?

Thank you

Andrew Sullivan

asullivan@dyn.com

Thanks to Joe Abley