

# Building and operating a global DNS anycast network

Sara Hassan

[sara@pch.net](mailto:sara@pch.net)

Packet Clearing House (PCH)

# #whoami

- Co-Founder of Sudan NOG ([sdnog.sd](http://sdnog.sd))
- Interconnection Manager at PCH
- outreach team for Quad9 ([9.9.9.9](http://9.9.9.9))

# About PCH

Packet Clearing House (PCH) is the global intergovernmental treaty organization responsible for providing operational support and security to critical Internet infrastructure, including Internet exchange points and the core of the DNS, since 1994.

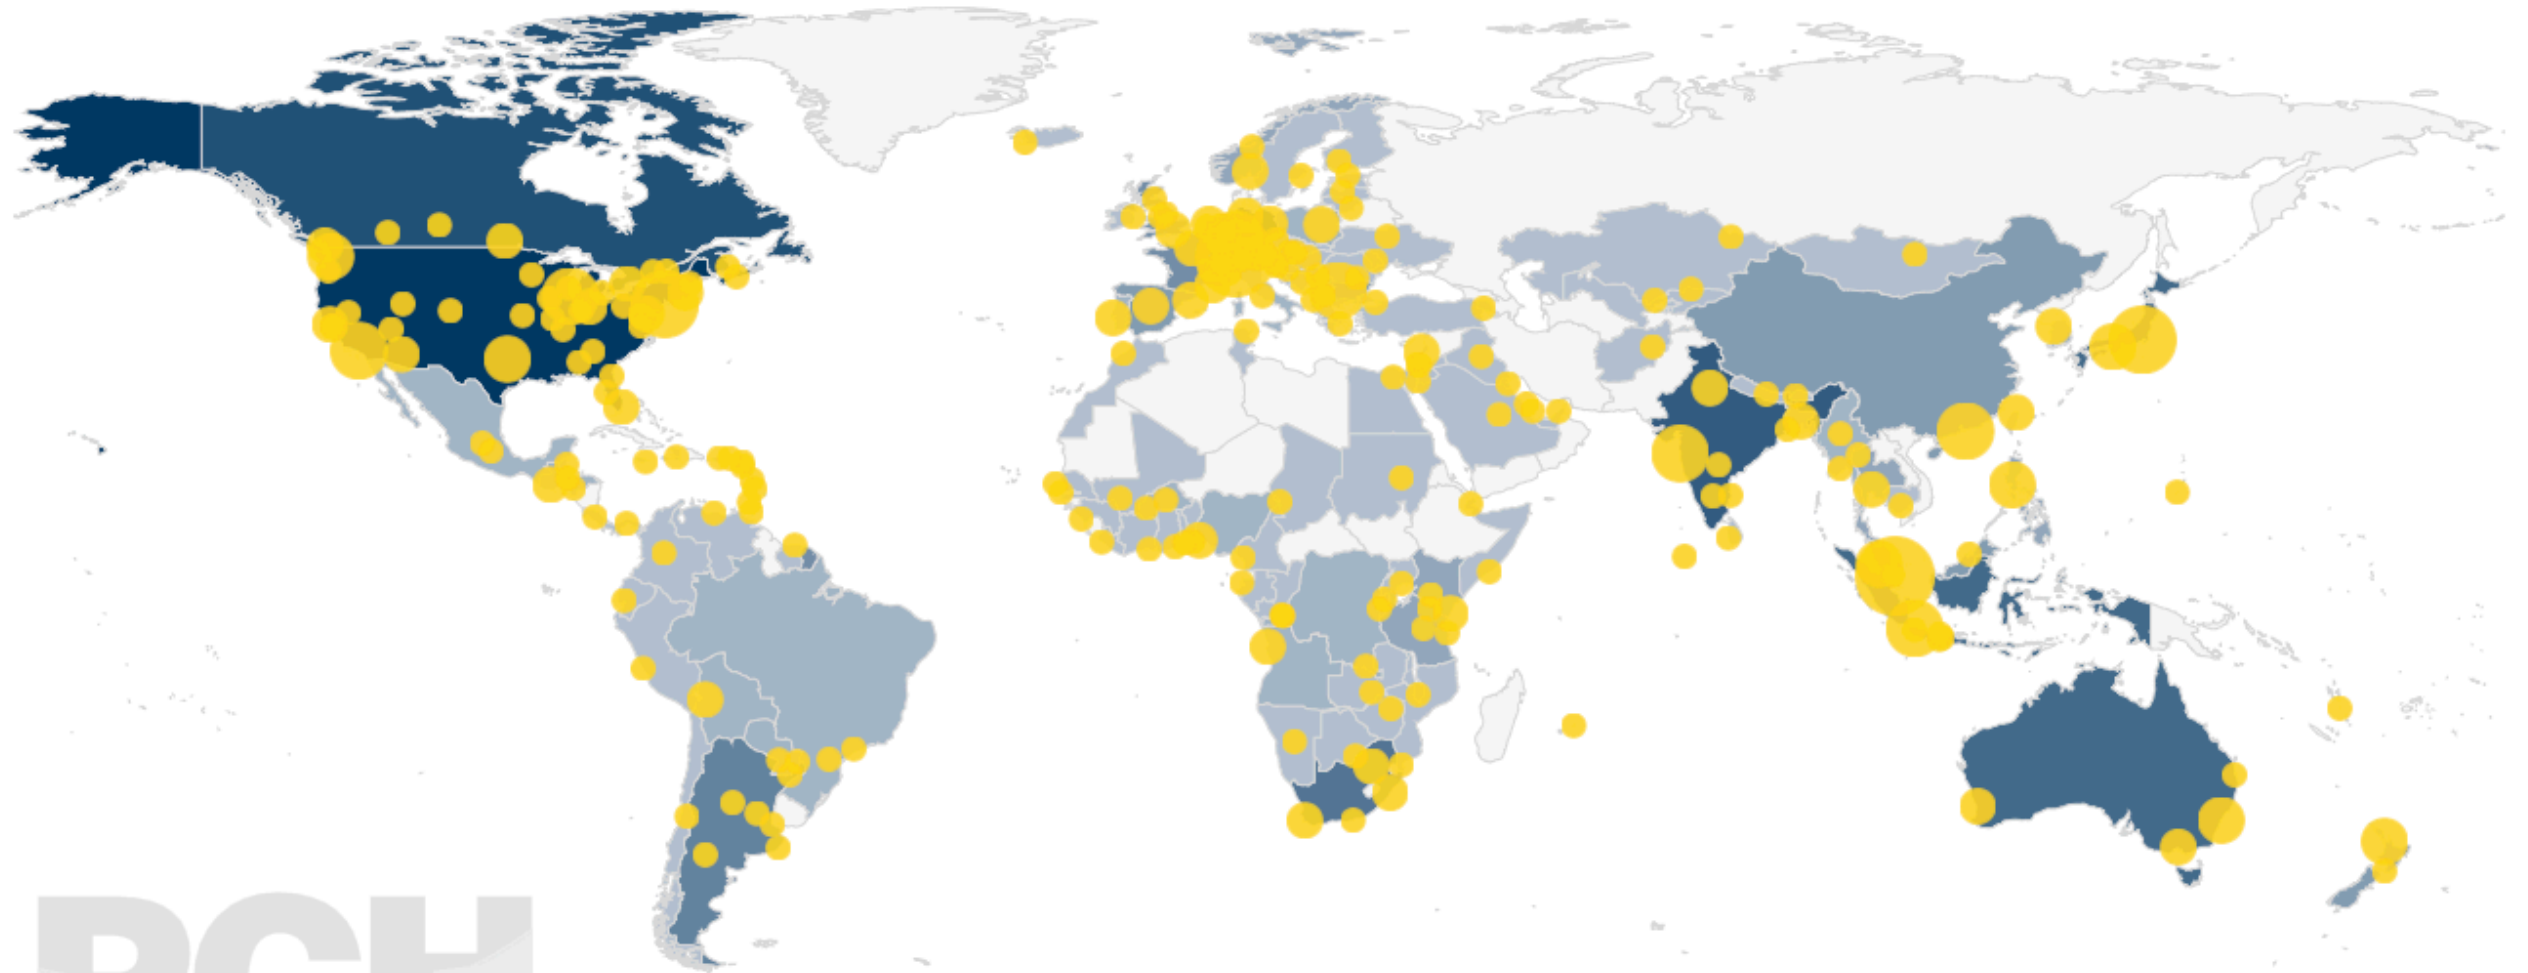
# About PCH

- Providing DNS infrastructure for 2 root server administrators
  - instance D - University of Maryland
  - instance E - NASA
- about ~300 TLDs including 128 ccTLDs
- Quad9 recursive resolver:
  - IPv4: 9.9.9.9
  - IPv6: 2620:fe::fe

# Anycast technology


- Anycast is a \*network\* trick that is used to have the same IP address online at various locations, at the same time.
- Via BGP, the Internet's global routing system routes users to the anycast node that is closest (topographically) to them.
- Anycast is a popular trick, especially with DNS.

# PCH's Anycast Cloud (AS42)




# A day in PCH's anycast network

DNS queries by protocol




DNS queries by IP version




# A day in PCH's anycast network (ii)

DNS queries processed by global and rest of nodes

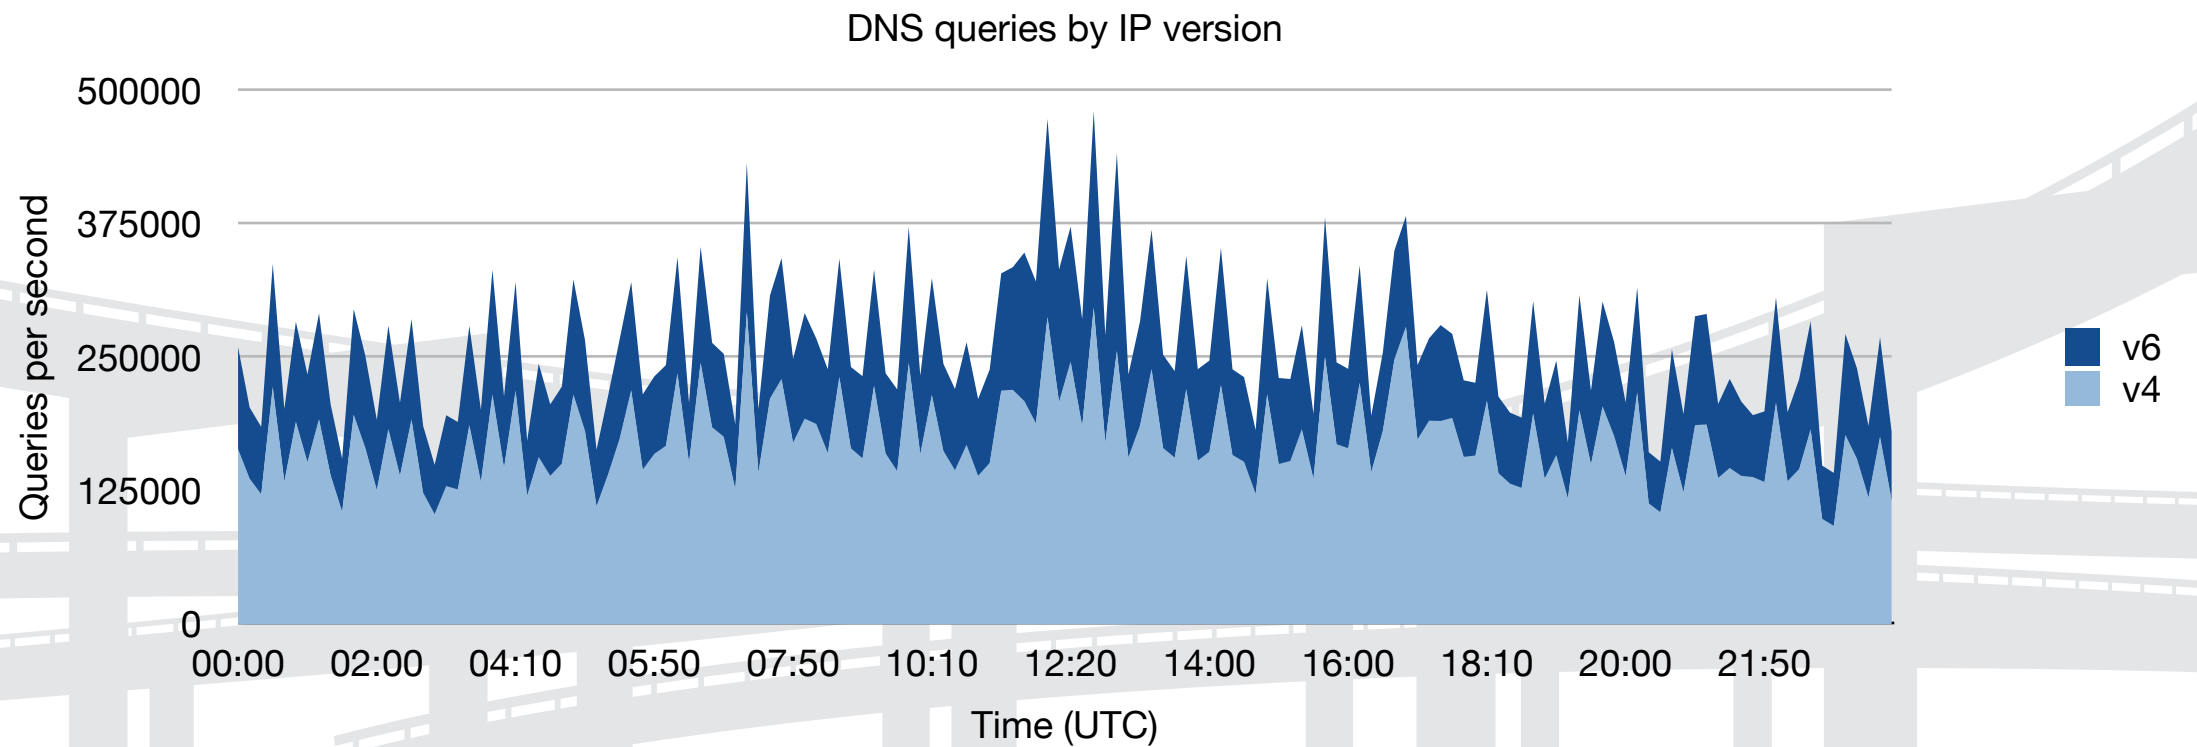



DNS queries by region





# A day in PCH's anycast network (.za)



# A day in PCH's anycast network (.za)

Key	UDPv4 Queries ▼	UDPv4 Responses	TCPv4 Queries	TCPv4 Responses	UDPv6 Queries	UDPv6 Responses	TCPv6 Queries	TCPv6 Responses
Hamburg	7,273,684	4,795,056	1,847,481	1,810,512	4,461,189	2,423,121	1,780,970	1,747,239
Johannesburg	2,316,272	2,270,769	51,654	41,474	505,190	498,813	2,762	1,062
Ashburn	1,525,971	1,518,584	10,322	6,075	280,004	273,368	5,464	4,152
Vienna	944,353	924,652	3,437	1,286	155,310	154,788	2,579	821
Frankfurt	888,072	851,382	21,492	10,265	338,461	335,488	5,920	5,387
Cape Town	782,859	744,897	14,397	11,412	216,533	189,359	1,161	778
Jakarta	729,299	727,070	5,595	1,357	18,776	17,544	854	494
Oslo	613,281	544,313	74,099	60,755	647,345	549,714	88,892	85,598
Tokyo	573,972	568,313	12,744	7,593	101,179	100,898	3,006	738
Santiago	492,150	491,843	1,553	410	127,859	127,792	1,079	144

# Planning Anycast Nodes





- Considerations when planning for new sites
  - PCH prefers to host infrastructure at IXPs
  - Perform testing to measure performance
  - Work closely with IX operator to obtain maximum outreach
- Limited resources make new deployments challenging
- Use of external (off-country) resolvers negates benefits of local nodes.

# How to measure performance?

Simple experiment :

- Using RIPE atlas probes
- Trace-route to “valid destination”
  - up and running
  - the service is operating
  - we *\*always\** peer with the route server

# Routing in AE

ASN	Connected to the IX	PCH
60924		
3491		

# AS60924

Hop	ASN	IP Address	Reverse DNS	RTT 1	RTT 2	RTT 3
1	<u>60924</u>	<u>185.92.128.4</u>		0.419 ms	0.323 ms	0.403 ms
2	<u>60924</u>	<u>185.92.128.3</u>		0.487 ms	0.454 ms	0.282 ms
3		<u>185.1.8.14</u>	router.dxb.woodynet.net	0.608 ms	0.459 ms	0.297 ms
4	<u>42</u>	<u>204.61.216.4</u>	anyns.pch.net	0.572 ms	0.354 ms	0.842 ms



# AS3491

Hop	ASN	IP Address	Reverse DNS	RTT 1	RTT 2	RTT 3
1	<u>3491</u>	<u>63.222.248.14</u>		1.15 ms	1.085 ms	1.175 ms
2	<u>3491</u>	<u>63.223.34.74</u>	HundredGEO-3-0-1.br04.sin02.as3491.net	77.945 ms	77.944 ms	77.977 ms
3	<u>3491</u>	<u>63.223.34.74</u>	HundredGEO-3-0-1.br04.sin02.as3491.net	77.717 ms	77.832 ms	77.788 ms
4	<u>38278</u>	<u>27.111.228.1</u>	42.sgw.equinix.com	77.987 ms	77.97 ms	77.974 ms
5	<u>42</u>	<u>204.61.216.4</u>	anyns.pch.net	77.752 ms	77.625 ms	77.681 ms



# Why!

- Networks don't peer at the IXP
- Mis-configurations
- Wrong AS-SET Filters

# Remember!

- `peering_local_pref > transit_local_pref`
- advertising more specifics to your transit hurts \*you\*
- Use modern tools to automate building filters from AS-SETs
- Connect more Atlas Probes and Anchors



“every time you send a packet to international destination that could be serve locally, you are subsidizing the internet eco-system in that country ”



Thanks for your attention

Sara Hassan  
Packet Clearing House  
sara@pch.net