

AS112 Operations Update

The Name Servers

at the end of the Universe

William Sotomayor, Keith Mitchell

**Yeti DNS Workshop
Yokohama**

31st October 2015

www.dns-oarc.net



DNS-OARC

Domain Name System Operations Analysis and Research Center

OARC's Mission

The Domain Name System Operations Analysis and Research Center (DNS-OARC) is a non-profit, membership organization that seeks to improve the security, stability, and understanding of the Internet's DNS infrastructure.

DNS-OARC's mission is to:

- promote and conduct research with operational relevance through data collection and analysis*
- offer useful services and tools*
- build relationships among its community of members*
- facilitate an environment where information can be shared responsibly*
- enable knowledge transfer by organizing open workshops*
- increase public awareness of the DNS's significance*



DNS-OARC

Domain Name System Operations Analysis and Research Center

OARC's Value to the Community



OARC is a prestigious library specializing in the DNS. We do our best to index a large collection of data in perpetuity, and because we're a neutral organization we strive to make the library accessible and welcoming to all researchers and operators of the DNS. For our members we offer private meeting rooms (mailing lists) where they can discuss and collaborate on current DNS Operations. At our upper levels you'll find our most precious archives and artifacts, set aside behind glass for safe keeping. We also have a Friends of the Library Program, where an exclusive group of members and guests gather on a semi-annual basis at our workshops to share their ideas, incentivize and encourage research, and catch up on the latest industry news.



OARC's Functions

- Facilitate co-ordination of DNS operations community
- Ongoing data gathering
- Operate community info-sharing resources
- Maintain/host DNS software tools
- Outreach via external and shared meetings
 - <https://indico.dns-oarc.net/event/22/>
 - Buenos Aires, 1-2 April 2016 (*IETF95*)



DNS-OARC

Domain Name System Operations Analysis and Research Center

What is...

- A global, loosely organised, volunteer-driven effort to divert 'junk' DNS reverse lookups from the root servers to decrease their workload
- Junk being the in-addr.arpa queries of RFC1918 network addresses coming from many resolvers
- Also absorbs dynamic DNS update attempts from certain mis-configured operating systems
- The name AS112 is taken from the autonomous system number used for the project, using a well-known set of network prefixes, that attracts this type of DNS traffic
 - The service is also anycasted for better performance.

AS112 Update

- Recent progress made on AS112 with 2 new Informational RFCs
- RFC 7534 updates since RFC 6304 include:
 - IPv6 transport
 - some new advice for reserving the AS112 project AS112 project ASN and prefixes in the IANA special use registry
 - advice on logging, and on potentially leaking information on internal infrastructure behind a NAT, plus other fixes
- RFC 7535 addresses the problem of delegating new zones to AS112 operations without relying on AS112 operators making required changes to their nodes for every change
- <https://www.as112.net/>



DNS-OARC

Domain Name System Operations Analysis and Research Center

AS112 Operations

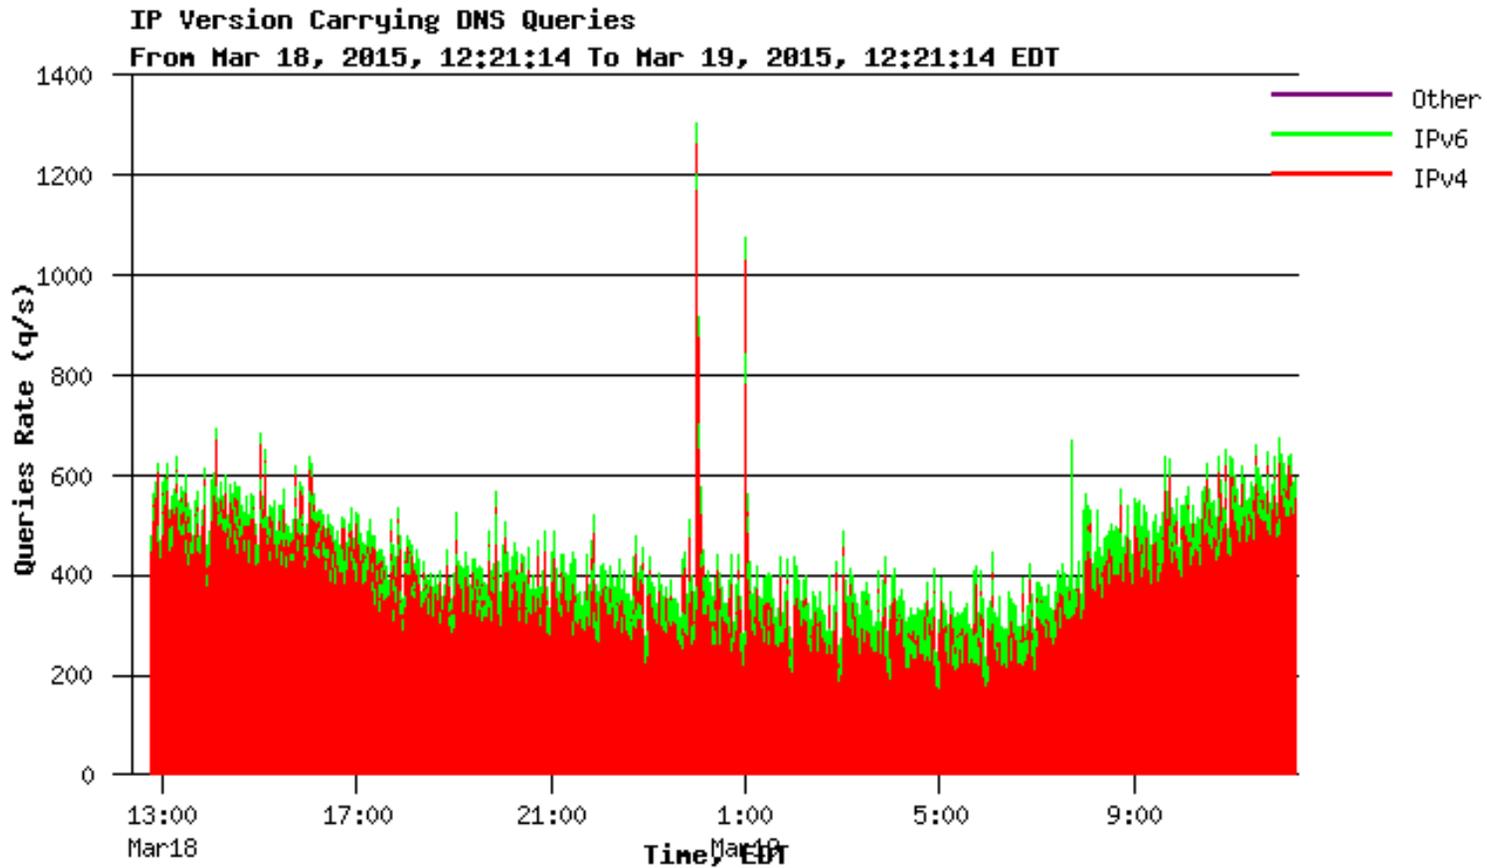
- As a result of being in the RFC Editor's hands, on March 17 2015, certain IANA actions were (suddenly) triggered:
 - AS112 nodes must now expect queries on IPv6 sockets, and AS112 operators should follow new directions in RFC 7534
 - There are new IPv4 and IPv6 prefixes for use as described in RFC 7535 to attach and listen for delegated traffic
- Yes, there are IPv4 NATs using IPv6 DNS to ask about IPv4 reverse junk...note the irony



DNS-OARC

Domain Name System Operations Analysis and Research Center

AS112 Dual-stack



AS112 Surveys, 2015

- For the past 4 years, AS112 surveys have been conducted by my self
 - The official as112.net operators list is never 100% accurate.
- Based on the following question: How does one find an unpublished public or 'private' AS112 server?
 - The same way as detecting a blackhole – by inference.
- Open DNS resolvers are so very, very useful
 - Ask the open resolver if it can identify which AS112 node is responding to it and maybe get some extra info too
- Route Views is good for this too
 - Automated login to various route views servers and query for the AS112 IPv4 route to derive first upstream ASN
 - But not as useful as open resolvers to find localised AS112 nodes.
- Results have been posted to the AS112 website for everyone's enjoyment
- Trivia: There is an AS112 server in New Caledonia! Yes, really!

Survey Methodology

- Automation is pretty much key, as well as trying to get the results quickly
 - Fetch the open resolvers file
 - Find number of lines, calculate the number of factors, then use those to split the file up into 100 or more pieces of even size by number of lines
 - Initiate parallel 'dig' queries and save results in unique files
 - Resulting output combined
 - Import results into database or spreadsheet and start count analysis
 - Remove duplicates
- Then publish the final result

AS112 Survey Results

- Of 7,601,333 allegedly open resolvers tested, 926,555 responded to any of 2 questions (January 2015)
- A number of observations stated on the website, but in summary:
 - AS112 nodes need to better identify themselves, as per the RFC
 - Eg, “Osaka, JAPAN” or “Widgets, Anytown, AnyCountry”, are not useful to operators, etc.
 - The mechanism of the test is infrequent (once a year)
 - Many AS112 nodes were geographically close to AS112 clients, but because of poor local peering (deliberate or otherwise), many clients cross half the world to reach their ‘nearest’ AS112 server
 - There seems to be a small churn of AS112 nodes, some drop off, some new ones come online but the number remains within a range of 71-73 nodes.
- Number 1 AS112 node appears to be Hivane’s
- ‘Nodes’ is a term used very loosely as there could be more than one instance of a poorly described AS112 server duplicated in these counts.

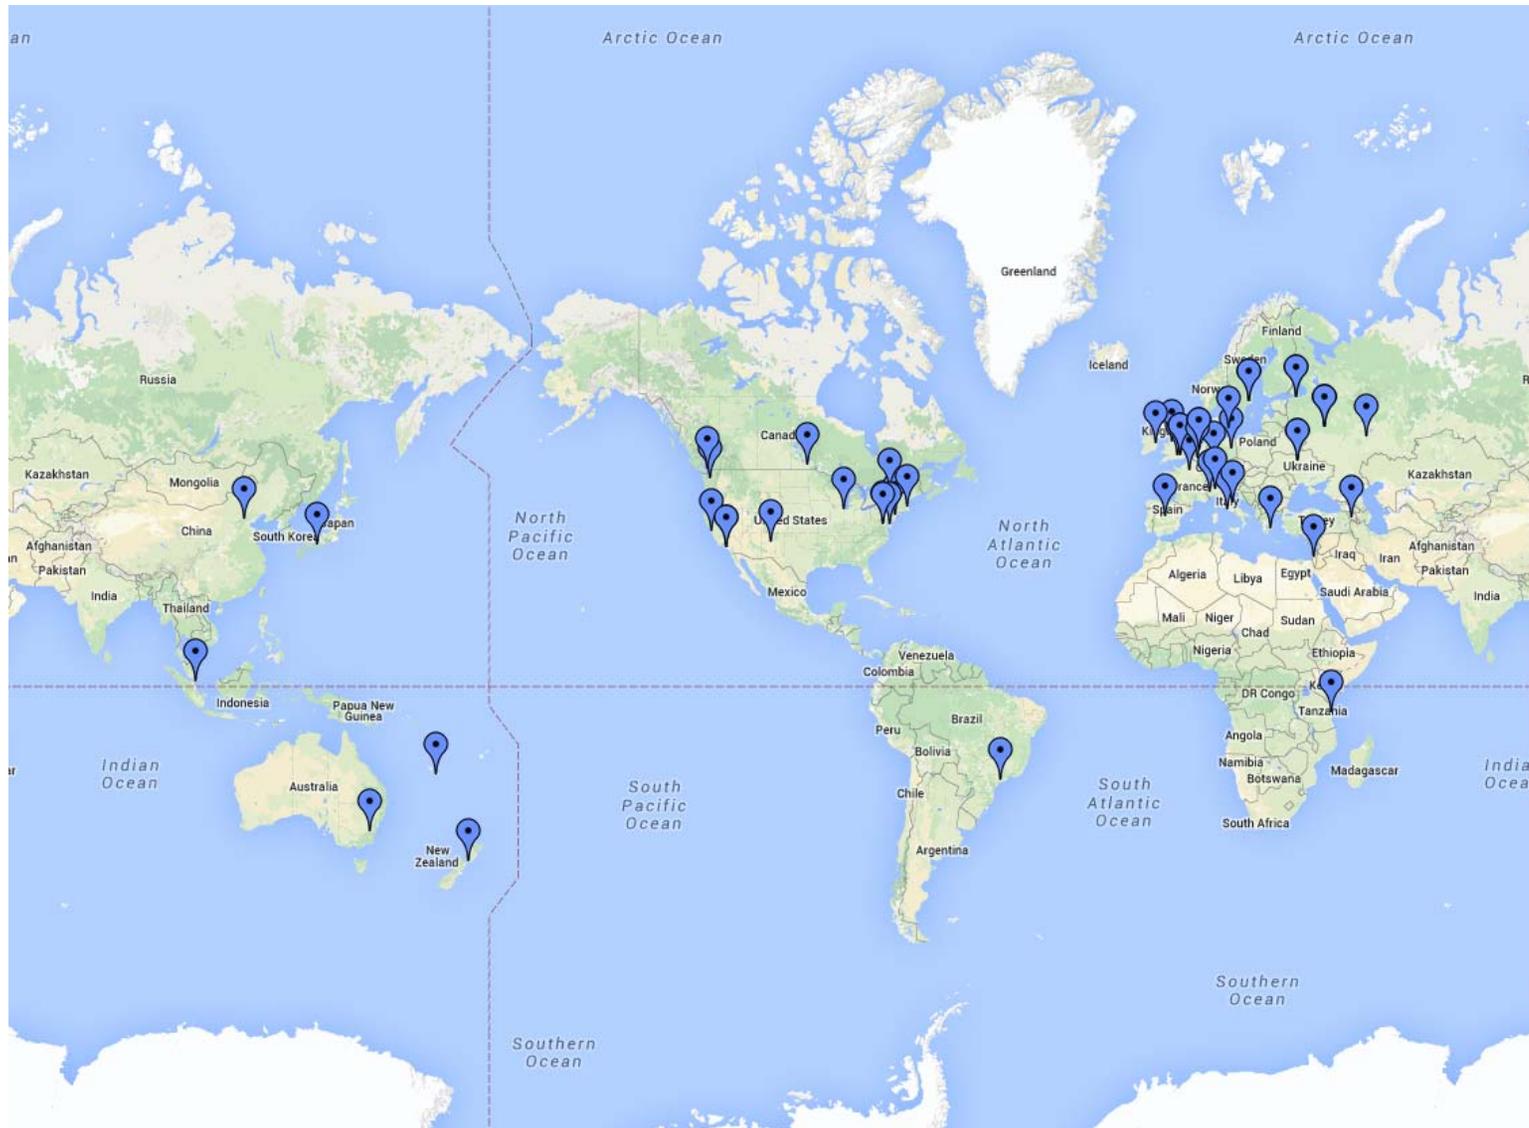
AS112 Ops list vs Survey

- What about the operator's listing?
- Using Google Maps, Red markers for ops listing, blue markers for survey listing just to compare
- Note that despite the AS112 server density in Europe, not all European clients use them.
- Also note there of course more nodes per organisation, for example Verisign has several.

2014 AS112 Ops Listing



2014 AS112 Survey Results



AS112 Maps

- They don't look all that different, but the devil is of course in the details, as there are fewer
 - In fact for 2015 they're the same, respectively
- Despite this, it is useful to visualise not just the proximity of these AS112 nodes but also how these nodes are reached
- However this would require a greater degree of co-operation between AS112 operators to contribute data to pinpoint proximities between clients and nodes to optimise network placement
- An example follows

AS112 and NRENs

- For over 15 years, a some-what separate global Internet has been constructed for research and education only
 - BGP counts hold steady at around 16000 routes.
- Examples are AARNet, CANARIE, CUDI, Internet2, GÉANT, etc.
 - All of these networks have one policy in common: You must provide your own commodity Internet access
- It has long been known that there are few root or TLD servers located on that network (Project: Find them!)
- To what extent is a lone (is it alone?) AS112 node heard on such a network if located there – with no default route of its own?
- Do DNS queries leak beyond those R&E networks into the commodity Internet?
- What is the reach of such leaky traffic?
- Why not measure using AS112?
- It stands to reason that for every open resolver on the Internet, there *must* be corresponding ones on the R&E networks, right?

Cautions on the Picture

- Sparse, isn't it?
- But there were 2837 open, unique resolvers.
 - GeolIP mining yielded many overlapping latitudes and longitudes (if that is to be believed)
 - Yet many IPs are within ranges of networks, so we can deduce that there are a number of open resolvers on campuses – Let's hear it for consistent IT policies!
 - On a high-speed research network, quite the abuse opportunity particularly when bandwidth is typically 10Gb/s, 40Gb/s or 100Gb/s
- Remember: these are open resolvers telling us about the potential for reach, it doesn't tell us anything about clients heard at a particular AS112 node on a research network

Clients on Research Networks

- The picture is vastly different if we have a DITL capture instead
- This would give a view of the clients on the research networks using a particular AS112 server
- But research networks tend to be quite leaky
 - Routing asymmetry
 - Secret academic handshakes with backdoors
 - Very, *very* special projects with multiple types of connectivity
 - Etc.
- How much AS112 traffic is not symmetrical may be difficult to answer and is really a separate question
 - But there are quite a number of clients for whom this AS112 node has SYN_SENT in its network state table
- Thus with some Open Street Map and JavaScript magic...

AS112 NREN Clients

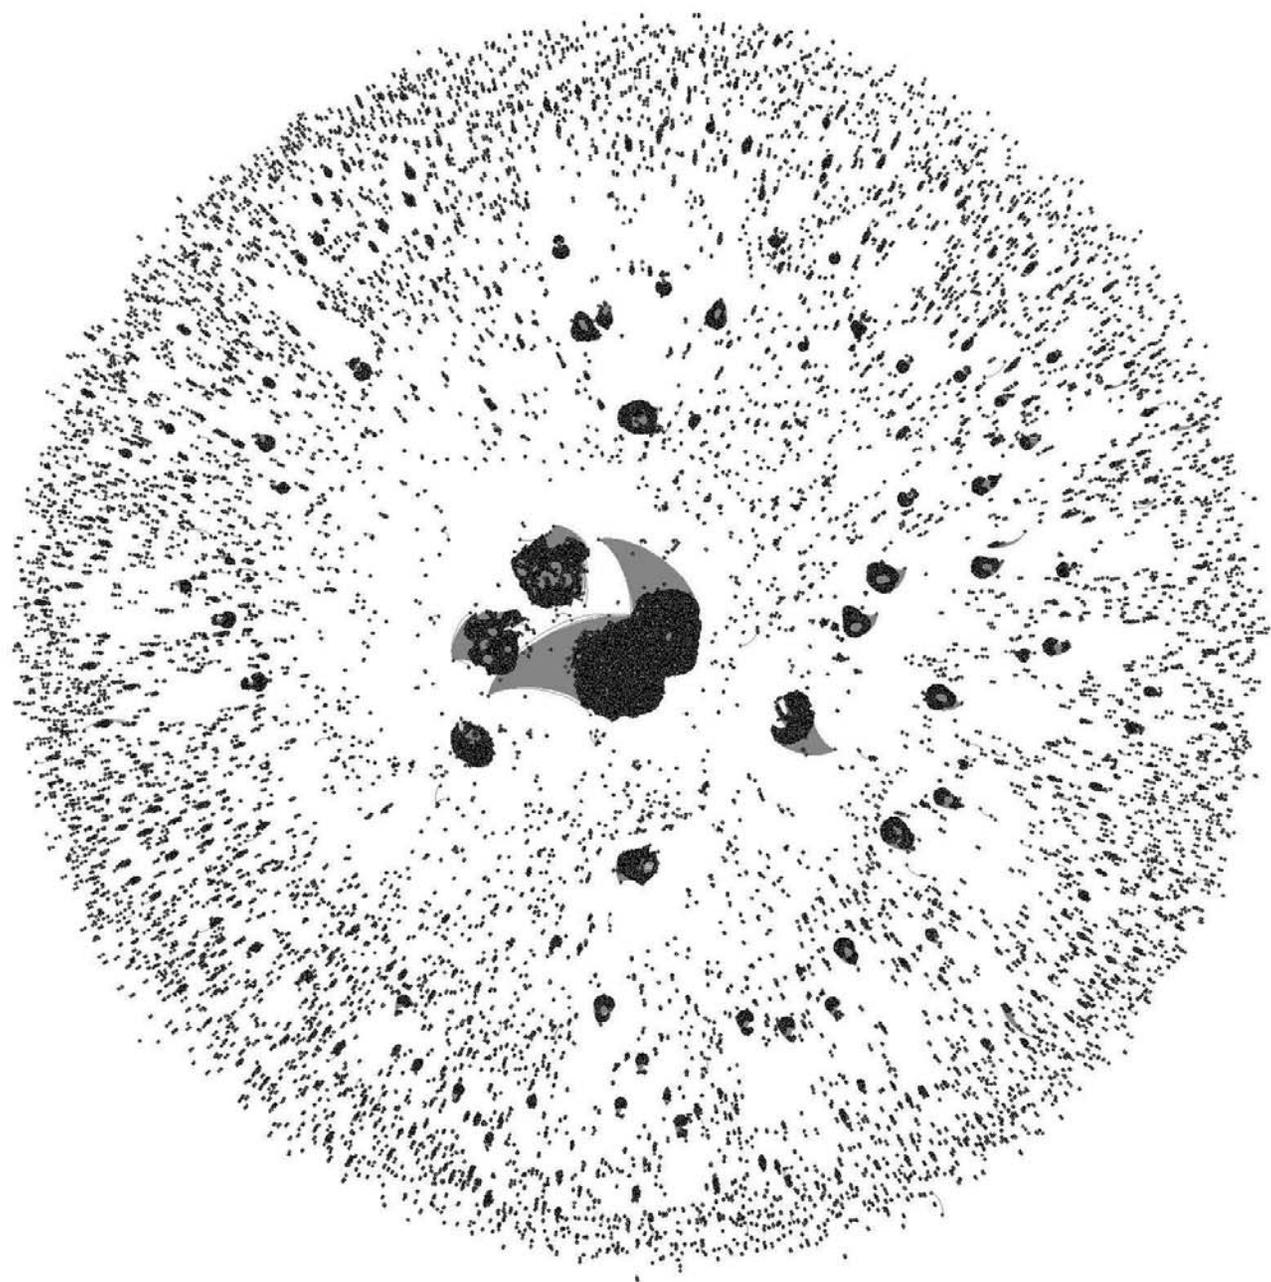


AS112 NREN Conclusions

- Answer: It's usually the end-user institutions' routing policy to prefer the faster NRENs over commodity Internet
- Moral of this story: We need more AS112 nodes on the NRENs, for example
 - EuroIX could do more to have IXPs with AS112 nodes peer with NRENs and GEANT
 - AARNet could deploy an AS112 node for their members
 - Regional NRENs in Latin America can do their part as well
- This might infer that we also need more root and TLD servers in the right places
- Having a list of open IPv6 resolvers would be useful
- Or, just fix your DNS and help make AS112 obsolete

AS112 Traffic on One Node

- In all seriousness, the foregoing is a visualisation of one AS112 node's traffic (it happens to be the one on a NREN, the subject of this presentation)
- Each galactic 'bloom' in this universe represents one server making in.addr-arpa queries
- The bigger the bloom, the greater the number of queries from a single source to that AS112 node's universe in a given time period
- Do the larger ones represent an overly aggressive set of resolvers, or infer the *size* of a particular NATed network??
- 'Sir, the second blob to the left is us leaking internal data'



Questions ?

AS112 and IXPs

- DNS-OARC and some major IXP operators exploring working together on AS112 operations
- IXP would provide ports and server for AS112 IXP open peering
- OARC will operate these and gather research data generated from queries
- Helps IXP and its members localize and sink junk traffic
- Mutual associate membership status
- Please talk to us if you operate an large IXP (or eyeball ISP) and are interested in doing this



DNS-OARC

Domain Name System Operations Analysis and Research Center

Data Interpretation

- How can we best interpret this part of the data?
- Why are so many reaching this particular AS112 node?
 - Notice Germany, Spain, UK, Ireland, Greece and Italy have no clients
 - Good DNS admin? Good local peering to IXPs with AS112?
- How much of these are actually unreachable by this AS112 node?
 - Hint: There is no default route in this node's routing table, so clients whose routes are not seen in its BGP tables will never get an answer back. If there are routing leaks and asymmetries...?
- Given the proximity of other AS112 nodes to these clients, why aren't they using local nodes via commodity Internet instead?
- And would a longer capture infer a consistent asymmetry between NRENs and commodity Internet resolvers or signal instabilities between the two?
 - Can they be located and local routing policies changed?
 - So many other questions, so little resources and time
- With all this data, what can one learn about so many networks?