



VERISIGN®

Open Resolvers in COM/NET Resolution

Duane Wessels, Aziz Mohaisen

DNS-OARC 2014 Spring Workshop

Warsaw, Poland

Outline

- Why do we care about Open Resolvers?
- Surveys at Verisign
- Characterizing Open Resolvers
- Intersection with COM/NET query sources
- Geographic distribution
- Discussion


Why do we care?

- Exploited in DDoS attacks
 - Makes cache poisoning attacks much easier
 - Cache snooping
 - Analogous to open mail relays
-
- Note: we're talking about unintentionally open resolvers here...

Two Surveys of IPv4 Open Resolvers

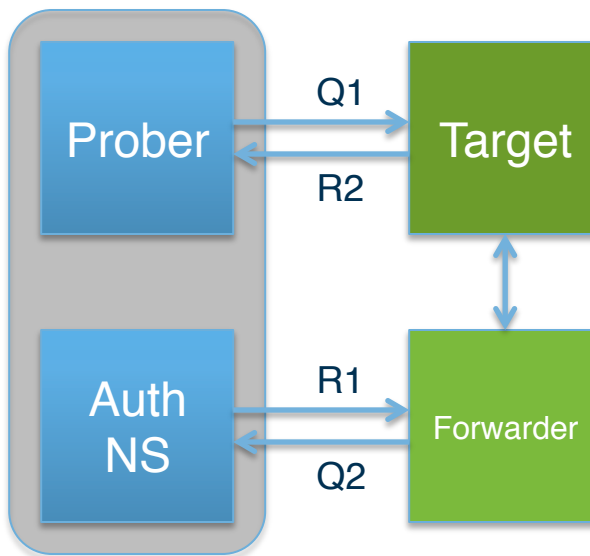
Models

- Target forwards query directly to Authority



Models


- Target forwards to a “forwarder”




October 2013 Survey

- From Amazon Web Services
- Took 173 Hours
 - 2013-10-28 14:00 – 2013-11-04 18:00
- Sent 3,676,739,504 Q1 probes
 - All IPv4 space, except class D/E, RFC1918 and do-not-probe list
- Received 43,538,209 Q2's
 - For 28,897,054 distinct probes
 - From 277,049 distinct IP addresses
- Received 34,604,998 R2's
 - For 32,040,586 distinct probes
 - From 31,424,854 distinct IP addresses

Query and Reply Rates during Open Resolver Scan




Query and Reply Rates during Open Resolver Scan



May 2014 Survey

- From Verisign
- Took 17 hours
 - 2014-05-01 18:20 – 2014-05-02 11:30
- Sent 3,676,724,690 Q1 probes
 - All IPv4 space, except class D/E, RFC1918, and do-not-probe list
- Received 38,079,578 Q2's
 - For 24,553,785 distinct probes
 - From 230,417 distinct IP addresses
- Received 28,426,251 R2's
 - For 27,905,762 distinct probes
 - From 27,281,623 distinct IP addresses

Query and Reply Rates during Open Resolver Scan




Data Analysis

- Data is collected with pcap while scan runs
- Pcap files are then parsed into whitespace delimited text
 - Separate files for Q1, Q2, R1, R2
- The text files are loaded onto Hadoop
- Analyzed with Hive (SQL statements)
 - Lots of large, multi-table joins

Closed Targets


- When the probe results in neither a Q1 nor an R2.



	Oct 2013	May 2014
Closed %	99.1	99.2

Open Targets

- When the probe results in either a Q1 or an R2.




	Oct 2013	May 2014
Open Count	33,660,906	29,292,597

	Oct 2013	May 2014
openresolverproject	32,673,337	27,454,609

Simple Open Resolver


- Q2 source address equals Target address
- i.e., Target does not forward elsewhere



	Oct 2013	May 2014
Simple	0.6 %	0.6 %

Forwarder

- Q2 source address differs from Target address



	Oct 2013	May 2014
Simple	0.6 %	0.6 %
Forwarder	79.8 %	78.0 %


- How many to Google?

	Oct 2013	May 2014
Google Fwds	8.3 %	8.9 %

No Q2, R2 Error

- Didn't get a Q2 query and got an Error response
- Usually REFUSED, which is good!


	Oct 2013	May 2014
Simple	0.6 %	0.6 %
Forwarder	79.8 %	78.0 %
Err No Forward	10.8 %	12.6 %



RCODE	Oct 2013	May 2014
1 FORMERR	0.0 %	0.0 %
2 SERVFAIL	10.0 %	9.1 %
3 NXDOMAIN	3.0 %	3.6 %
4 NOTIMPL	0.0 %	0.0 %
5 REFUSED	86.9 %	87.3 %
7	0.0 %	0.0 %
9	0.0 %	0.0 %
10	0.0 %	

Got Q2, but R2 error code

- Received the Q2 query, but then got an error response.
- Usually SERVFAIL




	Oct 2013	May 2014
Simple	0.6 %	0.6 %
Forwarder	79.8 %	78.0 %
Err No Forward	10.8 %	12.6 %
Err w/ Forward	0.7 %	0.5 %

RCODE	Oct 2013	May 2014
1 FORMERR	0.1 %	0.4 %
2 SERVFAIL	77.5 %	75.9 %
3 NXDOMAIN	0.4 %	0.1 %
4 NOTIMPL	0.0 %	
5 REFUSED	22.0 %	23.6 %
13	0.0 %	

R2 Blocked

- Received Q2
- But no R2




	Oct 2013	May 2014
Simple	0.6 %	0.6 %
Forwarder	79.8 %	78.0 %
Err No Forward	10.8 %	12.6 %
Err w/ Forward	0.7 %	0.5 %
R2 Blocked	4.8 %	4.7 %

Synthesized Answers

- No Q2
- R2 had an Answer section with an A record, but wrong value.
- Many answer with their own IP


	Oct 2013	May 2014
Simple	0.6 %	0.6 %
Forwarder	79.8 %	78.0 %
Err No Forward	10.8 %	12.6 %
Err w/ Forward	0.7 %	0.5 %
R2 Blocked	4.8 %	4.7 %
Synthesized	3.4 %	3.6 %



Q2 Missing

- No Q2, but R2 had an Answer section with correct A record!
- How?
 - Data collection problem
 - Lucky guess


	Oct 2013	May 2014
Simple	0.6 %	0.6 %
Forwarder	79.8 %	78.0 %
Err No Forward	10.8 %	12.6 %
Err w/ Forward	0.7 %	0.5 %
R2 Blocked	4.8 %	4.7 %
Synthesized	3.4 %	3.6 %
Q2 Missing	0.0 %	0.0 %
Totals	100 %	100 %



- 120 times in Oct 2013 survey
- 1109 times in May 2014 survey

Weirdness: R2 not from Target


- Sent query to x.x.x.x
- Got response from y.y.y.y



	Oct 2013	May 2014
IP Changed	2.1 %	2.4 %

Weirdness: Local Port Changed


- Query sent from port X
- Response sent to port Y
- 1560 cases



	Oct 2013	May 2014
Local Port	1560	4936

Weirdness: Remote Port Changed

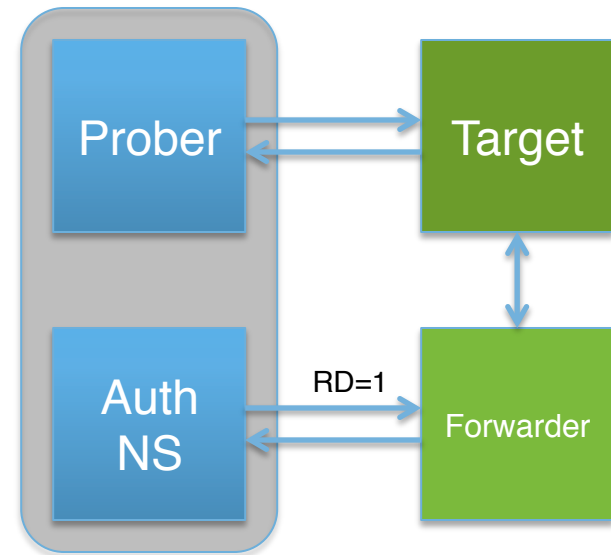
- Query to port 53
- Response from port != 53



	Oct 2013	May 2014
Remote Port	46.2 %	46.7 %

Weirdness: Q2 with RD=1

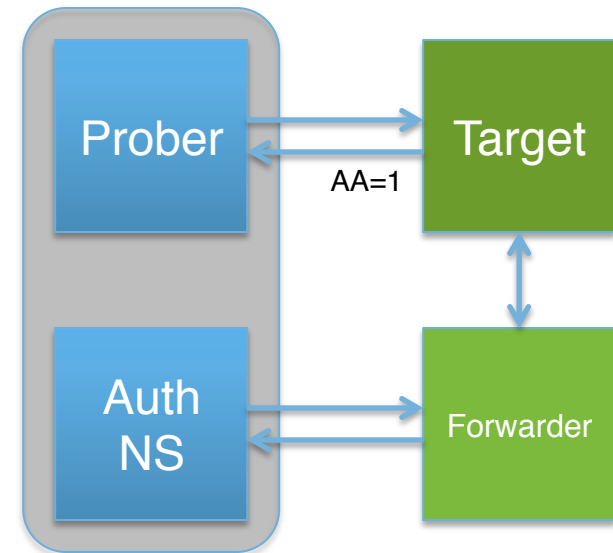
- Usually queries to Authoritative name servers have RD=0



	Oct 2013	May 2014
Q2 RD=1	6079	5186

Weirdness: R2 with AA=1

- Usually responses from recursive name servers have AA=0



	Oct 2013	May 2014
R2 AA=1	0.7 %	0.8 %

Intersection with COM/NET Queriers

COM/NET Query Data

- Four Verisign “big” sites

Site	Server
Amsterdam	h.gtld-servers.net
Wash DC	i.gtld-servers.net
New York	c.gtld-servers.net
San Francisco	g.gtld-servers.net

- Only 4 of 13 gtld-servers.net letters


Intersection of open resolvers and COM/NET (Oct 2013)

- What percent of open resolver exit Ips appear in the COM/NET query data?


Site	%OR IPs	%COM/NET IPs	%COM/NET Queries
Amsterdam	64.3	5.2	51.7
Wash DC	66.4	5.5	48.4
New York	66.2	5.2	46.2
San Francisco	64.6	5.4	45.2

- Example: At Amsterdam, we see 64.3% of the open resolvers IPs in one day. This is 5.2% of all COM/NET IPs seen there. Those IPs are responsible for 51.7% of COM/NET queries at the site.


Percent of Open Resolver Exit IPs found in COM/NET Queries



Percent of COM/NET query IPs found in Open Resolvers



Percent of COM/NET queries coming from Open Resolver IPs



Intersection of open resolvers and COM/NET (May 2014)

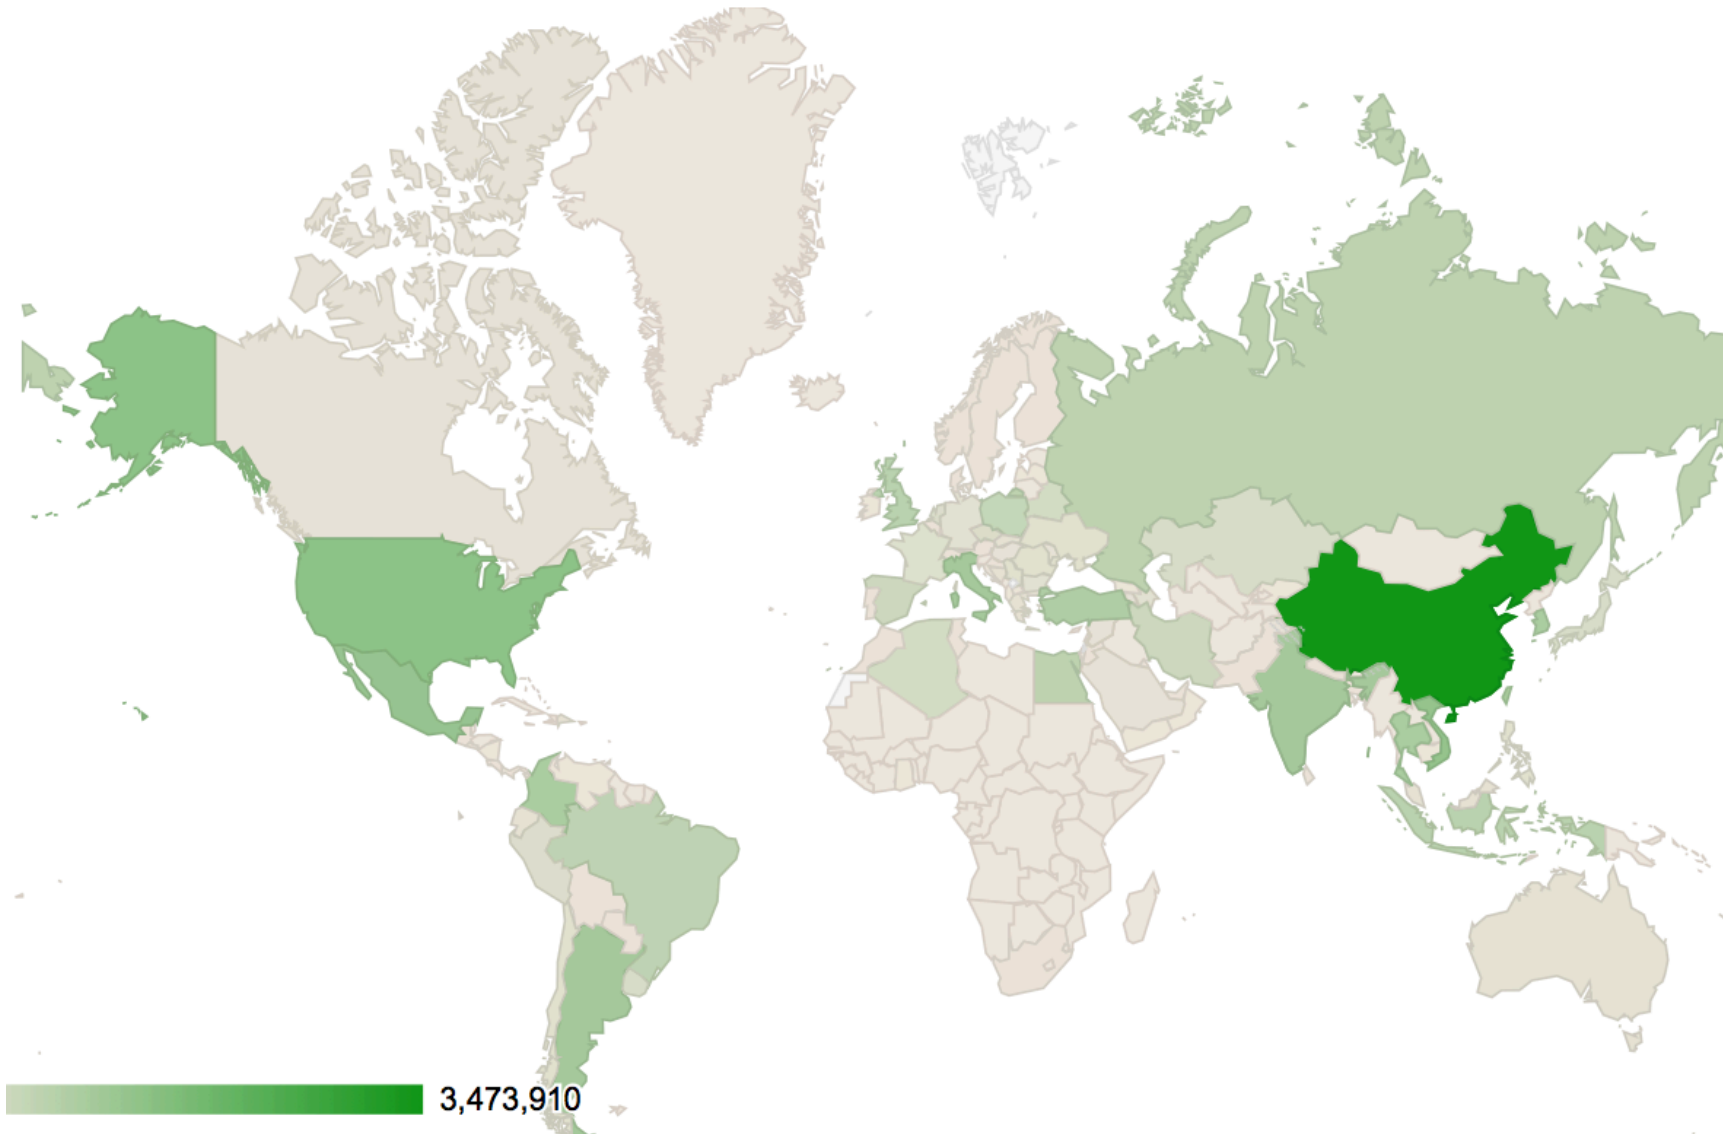
Site	%OR IPs	%COM/NET IPs	%COM/NET Queries
Amsterdam (H)	59.4	4.8	57.2
Wash DC (L)	61.3	4.8	50.6
New York (C)	down for maintenance		
San Francisco (G)	59.2	4.9	47.4

Geographic Distribution

Open Resolvers Geographical Distribution

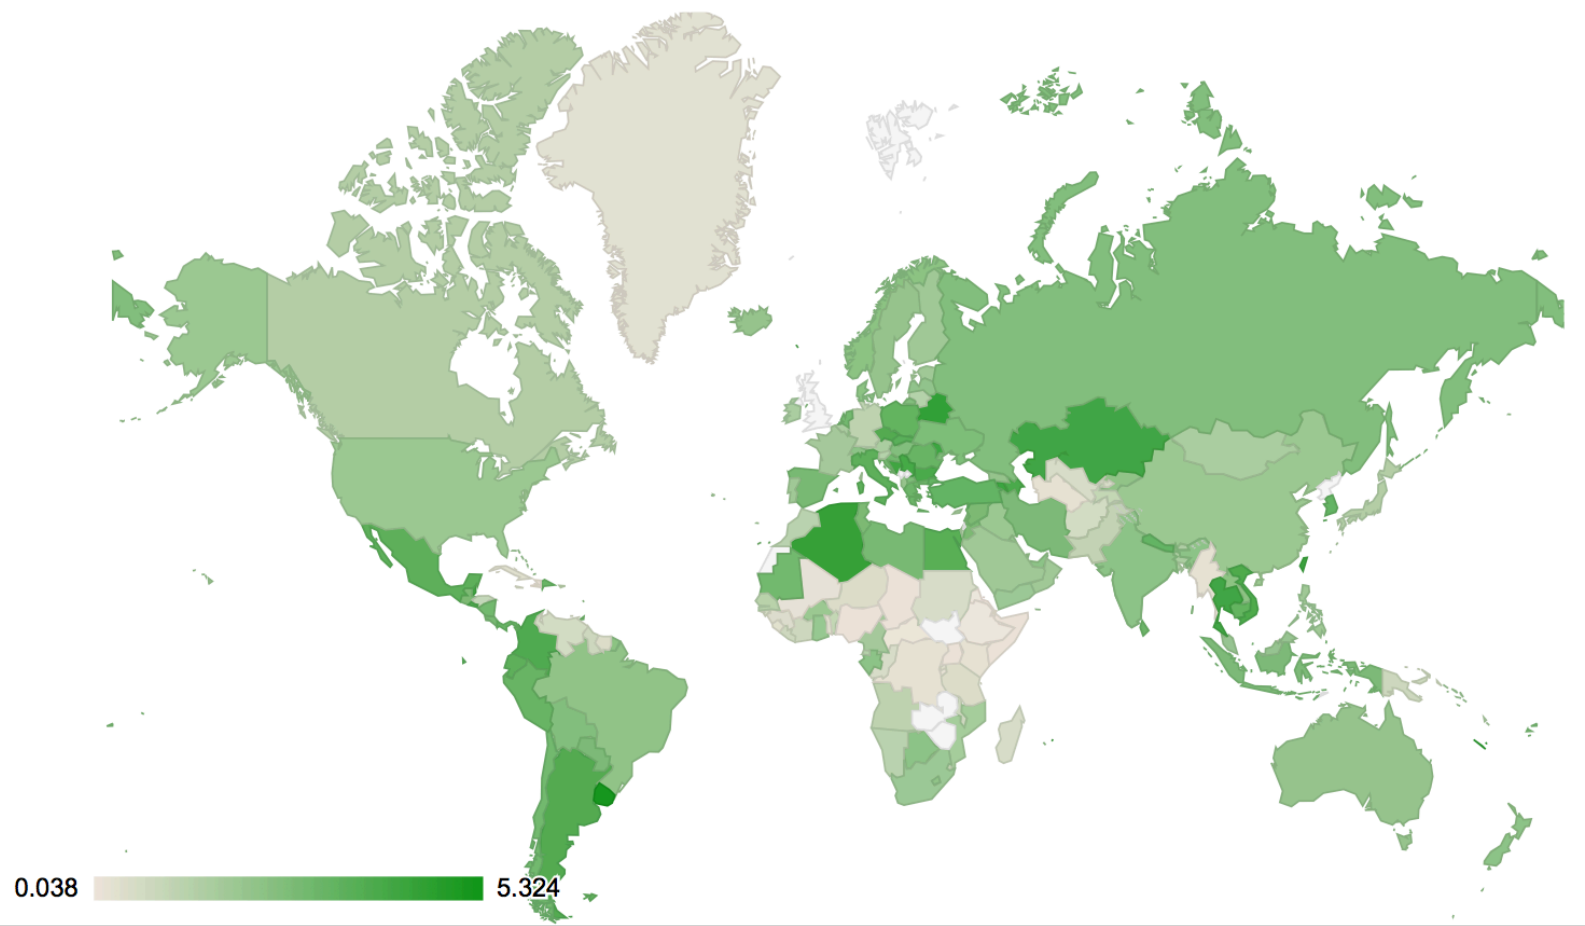
- Open resolvers are massively distributed
 - 232 countries (including special territories)
 - 10,240 different cities
 - 13,887 different organizations (including ISPs)
 - 83,407 different networks (domains)
- All distributions are heavy tailed (city, org, net, country)
- Open resolvers/forwarder associations are distributed
 - Includes across country associations
 - Not only limited to well-understood applications, but includes service providers association without territory resolvers

Open Resolvers Geographical Distribution




Open Resolvers vs. Internet Usage


- Per-user distribution is consistent with overall per-country, except in a few cases (small, hop countries)




Organization Level Distribution – Resolvers



Organization Level Distribution – Open Resolvers Per Forwarder



Organization Level Distribution - $\log_{10}(\text{Forwarders})$



Final Thoughts

Key Points

- Still many millions of Open Resolvers on the Internet
 - The trend is decreasing
- Most Open Resolvers forward to another recursive
- About half respond from the wrong port!
- Open Resolver forwarder IPs are strongly linked to COM/NET queries.
 - Responsible for 50% of the query traffic

Questions?