




# Port and Message ID Analysis of Resolvers Querying *.com/.net* Name Servers

David Blacka

Matt Larson

September 24, 2008

DNS OARC Meeting, Ottawa, Canada



# Goals

## + Primary:

- Determine how many servers patched in wake of summer 2008 cache poisoning vulnerability
  - Examine UDP source port and DNS message ID distribution

## + Secondary:

- Examine EDNS0 usage
  - Interesting and easy to add to analysis
- Examine other *.com/.net* query metrics
  - Unique source IPs, recursive queries

# Methodology


- + Analyze *.com/.net* queries for 24 hours
  - Beginning midnight UTC on September 5, 2008
  - 11 of 13 names in *.com/.net* NS RRSet
- + Count each <source IP, source port, message ID, RD> tuple
  - Custom libpcap application
  - Roll up counts from each name server to create grand totals across all name servers
- + Also, for each source IP:
  - Count queries with OPT RR (EDNS0 capable)
  - Count queries with DO (“DNSSEC OK”) bit set
  - Track advertised maximum UDP buffer size

# Totals

- + 34.2 billion queries analyzed
- + 28.3 billion unique <source IP, source port, message ID, RD> tuples
- + 4,950,579 unique IP addresses
  - 2632 bogons (0.053%) (Team Cymru definition)
  - Bogons mostly RFC 1918
- + 3,004,936 addresses (60%) sent at least 10 queries
  - Our minimum threshold for port/message ID analysis
  - 1455 bogons (0.048%) in this set
- + A lot of data:

```
$ ls -lh distilled.2008-09-05  
-r--r--r-- 1 matt matt 817G Sep 10 17:46 distilled.2008-09-05
```

# Query Distribution



# Query Distribution (2)

Queries received	Source IP addresses
<= 10	1,992,508
<= 100	1,036,271
<= 1,000	1,082,181
<= 10,000	583,276
<= 100,000	218,293
<= 1,000,000	33,026
<= 10,000,000	4,799
<= 100,000,000	225

# Top Queriers

Source IP	Domain Name	Queries	Q/sec
65.32.5.74	dns-cac-lb-01.tampabay.rr.com.	94,709,254	1,096
67.18.167.2	svr01.wwm.net.	66,414,877	769
66.75.164.90	dns-cac-lb-01.orange.rr.com.	64,976,327	752
65.24.7.3	dns-cac-lb-01.ohiordc.rr.com.	62,258,376	721
70.84.138.226	e2.8a.5446.static.theplanet.com.	60,557,732	701
74.52.217.34	22.d9.344a.static.theplanet.com.	57,121,972	661
208.111.154.15	crawl1.nat.svl.searchme.com.	52,803,926	611
208.69.36.14	bld4.chi.opendns.com.	51,063,423	591
24.92.226.9	dns-cac-lb-0.nyroc.rr.com.	50,629,095	586
24.93.41.125	dns-cac-lb-01.texas.rr.com.	47,290,908	547
212.19.48.14	ns.plusline.de.	44,736,109	518
24.25.5.150	dns-cac-lb-01.southeast.rr.com.	41,840,629	484
193.110.28.100	(No PTR record)	40,640,366	470
208.80.194.27	(Timed out)	39,098,191	453
209.235.152.127	mail937c35.nsolutionszone.com.	36,925,210	427
206.248.154.22	dns.pppoe.ca.	36,023,743	417
83.170.94.31	ns4.uk2.net.	34,477,392	399
208.138.27.134	echo2.cwjamaica.com.	33,548,963	388
202.126.40.9	(No PTR record)	33,250,654	385
209.235.152.126	mail936c35.nsolutionszone.com.	33,182,629	384
80.12.195.55	(No PTR record)	33,144,749	384
203.146.237.88	(No PTR record)	32,334,639	374
208.69.36.13	bld3.chi.opendns.com.	32,029,180	371
208.69.36.12	bld2.chi.opendns.com.	30,436,211	352
212.217.0.14	adslrabort3.iam.net.ma.	29,366,154	340
24.29.103.10	dns-cac-lb-01.rdc-nyc.rr.com.	29,363,106	340
209.235.146.139	mail369c25.carrierzone.com.	29,083,164	337
209.235.146.130	mail360c25.carrierzone.com.	28,652,569	332
204.179.96.100	(No PTR record)	27,836,893	322

# Definition: Standard Deviation

+ A measure of the *variability* or *dispersion* of a set of data

+ For a discrete data set (like ports, query IDs),

calculated as: 
$$\sqrt{\sum_{i=1}^N \frac{1}{N} (x - \mu)^2}$$


+ Zero = no variation in data (e.g., just one port)

+ Uniform discrete distribution calculated as: 
$$\sqrt{\frac{N^2 - 1}{12}}$$

+  $\sigma$  (standard dev.) of 0→65535 = 18918.61361

+ Low  $\sigma$  = data clustered near mean

+ High  $\sigma$  = data clustered away from mean



sample data: mean is 50,  
standard deviation is 20



## Definition: Q

+ A normalized form of standard deviation:

$$Q = 1 - \frac{|s - \sigma|}{\sigma}$$

+  $\sigma$  = standard deviation of the uniform distribution

+  $s$  = the calculated standard deviation from data

+ Basically, folds high std. dev. over  $\sigma$ , then normalizes to 0→1

+ Low Q = not close to uniform distribution

+ High Q = close to uniform distribution

+ Not a measurement of randomness

- E.g., a non-uniform distribution could also have high Q


# Definition: “bits”

- + Attempts to be a measure of how many bits of field are being used
- + Similar to formula used by *entropy.dns-oarc.net*
- + Based on modified range value:
  - $\text{Range} = \text{max} - \text{min}$
  - $M = \# \text{ of unique ports} / \min(\text{total ports}, 65536)$
  - $\text{MR} = \text{Range} * M$
  - “bits” =  $\log_2(\text{MR})$
- + Substitute “query id” for “ports”, etc.
- + High “bits” = wide range of ports, mostly different ports
- + Low “bits” = narrow range of port and/or not many different ports
- + Not a measure of randomness
  - E.g., a sequential series would have high “bits” value


# Analysis Methodology

- + Examined all source IPs sending at least ten queries
  - 3,004,936 IP addresses remained
- + Calculated standard deviation, Q and “bits” across each IP address’s:
  - UDP source ports (16 bits)
  - DNS message IDs (16 bits)
  - <source port, message ID> tuples (32 bits)
- + Attempted to classify patched vs. unpatched resolvers
  - Primarily using source port
  - Hard!
- + Examined message ID variability
  - True randomness calculation impossible, since query order lost in data collection method


# UDP Source Ports: Standard Deviation (1)




# UDP Source Ports: Standard Deviation (2)




# UDP Source Ports: Q



# UDP Source Ports: "bits"




# DNS Message IDs: Standard Deviation






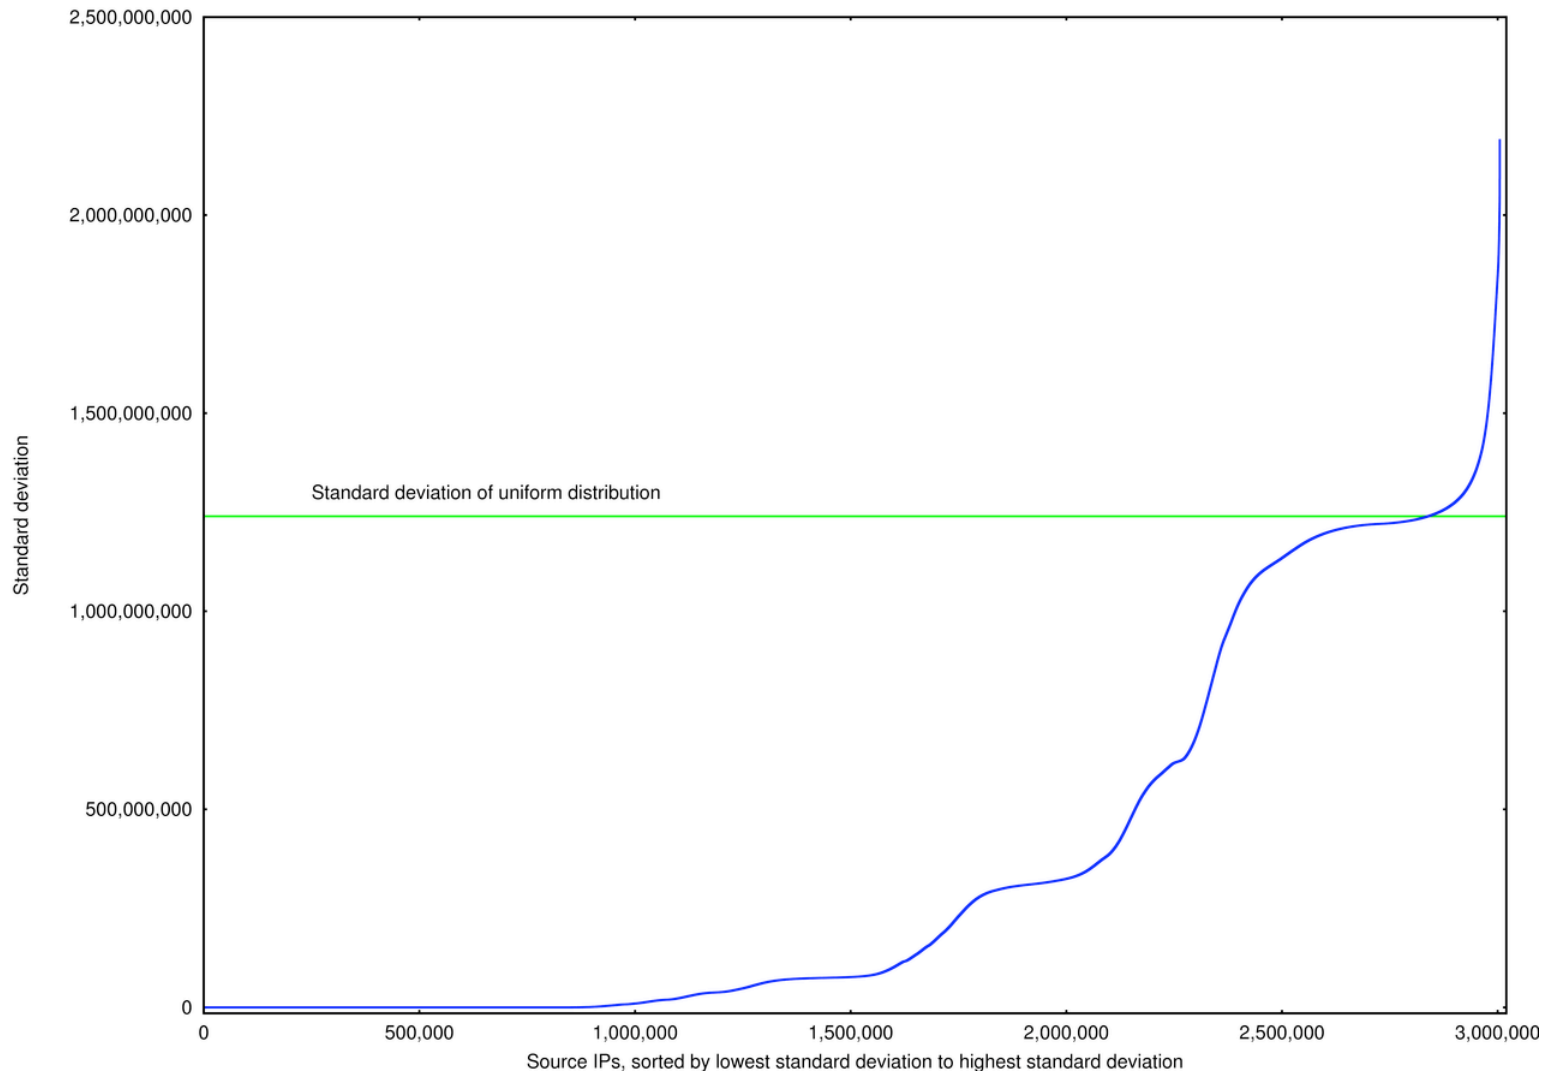
# DNS Message IDs: Q




# DNS Message IDs: "bits"



# Port + Message ID: Standard Deviation



# Port + Message ID: "bits"



# Conclusions

- + Message IDs: most queriers look good
  - A very few have both constant Message ID and source port, 1341
- + Likely Patched vs. Obviously Unpatched vs. Maybe Patched
  - Likely Patched = wide range of ports, not much repetition
  - Obviously Unpatched = very narrow range, much repetition
  - Maybe Patched = narrow range, not much repetition
- + “bits” Metric (ports)
  - Likely Patched,  $> 15.4$  bits = **18.9%**
  - Obviously Unpatched, 0 bits = **29.7%**
  - Maybe Patched = **51.4%**
- + Q Metric (ports)
  - Likely Patched,  $> 0.8$  = **18.3%**
  - Obviously Unpatched, is  $< 0.1$  = **28.0%**
  - Maybe Patched = **53.7%**


# EDNS0

- + Calculated EDNS0 capability of:
  - Total queries received
  - Total source IP addresses seen
- + Surprisingly low EDNS0 deployment

<b>Total queries</b>	34,251,224,131	
<b>EDNS0 queries</b>	19,717,298,077	<b>57.57%</b>
<b>DO bit queries</b>	17,030,829,145	<b>49.72%</b> (of total queries)
		<b>86.38%</b> (of EDNS0 queries)

<b>Total queriers (unique IPs)</b>	4,950,579	
<b>EDNS0 queriers</b>	1,409,778	<b>28.48%</b>
<b>DO bit queriers</b>	948,820	<b>19.17%</b> (of total queriers)
		<b>67.30%</b> (of EDNS0 queriers)

# EDNS0 Buffer Sizes



# Recursive Queries

## + Lots of recursive queries

- **17.6%** of total queries had RD set
- **25.8%** of total queriers (source IP addresses) sent exclusively recursive queries
- **50,714** queriers (**1.02%**) sent a combination of recursive and non-recursive
  - But that could be running *dig* or *nslookup* on a host also running a recursive name server

## + What sends exclusively recursive queries?

- Lots of different implementations, according to *fpdns*
- No smoking gun
- Hypothesis: misguided administrators configure a *.com/.net* name server IP address as a forwarder



# Recursive Querying Source IPs Fingerprinted

Percent	Count of IPs	Fingerprint result
94.37%	1,196,842	TIMEOUT
1.92%	24,287	No match found
1.69%	21,454	ISC BIND 9.2.3rc1 -- 9.4.0a0 [recursion enabled]
0.91%	11,523	Nominum CNS
0.67%	8,456	ISC BIND 9.2.3rc1 -- 9.4.0a0
0.15%	1,841	Mikrotik dsl/cable
0.07%	913	VeriSign ATLAS
0.04%	487	Paul Rombouts pdnsd
0.03%	345	ISC BIND 9.2.0rc7 -- 9.2.2-P3 [recursion enabled]
0.02%	250	ISC BIND 8.3.0-RC1 -- 8.4.4 [recursion enabled]
0.02%	226	vermicelli totd
0.02%	208	DJ Bernstein TinyDNS 1.05
0.02%	204	ISC BIND 9.1.0 -- 9.1.3 [recursion enabled]
0.01%	146	ATOS Stargate ADSL
0.01%	142	robtex Viking DNS module
0.01%	126	Microsoft Windows DNS 2000
0.01%	124	ISC BIND 4.9.3 -- 4.9.11
0.01%	102	Microsoft Windows DNS 2003
0.01%	97	ISC BIND 8.1-REL -- 8.2.1-T4B [recursion enabled]
0.01%	96	Runtop dsl/cable

# Questions + Answers



QUESTIONS + ANSWERS