HyperLogLog inspired three-minute scan of DNSSEC delegations worldwide

DNS-OARC 26, Madrid
@PowerDNS_Bert
bert.hubert@powerdns.com

POWERDNS
AN OX COMPANY

Agenda

- DNSSEC: existing measurements for secured delegations
  - Actively published/gathered/scraped
- Introduction to HyperLogLog
- NSEC3 “by eye”
- HyperLogLog for DNSSEC
  - NSEC3 HyperLogLog
  - NSEC applicability
- Live DEMO!
- Presentation of results
- What does it mean?

Whoami

- PowerDNS - open source nameservers since 1999
  - Commercially supported, sustainable software
  - Per-user, per-subscriber malware filtering & associated tooling available
- PowerDNS Authoritative Server
  - Leading DNSSEC signer, 30%-50% domain share
- PowerDNS Recursor
  - Highly flexible. RPZ, DNSSEC.
- Dnsdist: “highly DNS-, DoS- and abuse-aware loadbalancer. Its goal in life is to route traffic to the best server, delivering top performance to legitimate users while shunting or blocking abusive traffic."

DNSSEC: existing measurements

- Various registries expose or actively publish metrics
  - Varies from “can be scraped” to “official URL with numbers”
  - Other zones are public and can be counted
- Popular sources:
  - Kees Monshouwer’s DNSSEC graph for .NL: https://www.monshouwer.eu/dnssec-nl-graph/
  - Frederic Cambus’s StatDNS: https://www.statdns.com/
  - Rick Lamb’s exhaustive list: http://rick.eng.br/dnssecstat/
- These are numbers of TLDs willing (or forced) to publish statistics
  - Takes a lot of work to gather
- **There is some bias: TLDs that have reason to be proud will be more likely to publish**

DNSSEC: probing measurements

- Through statistical measures, it is possible to gather statistics even for zones that do not publish numbers
- Such measurements work reasonably well, in descending order of accuracy:
  - NSEC3 (opt-in, opt-out)
  - NSEC
  - Non-DNSSEC
- The techniques used relate to the awesome but generally not very well explained HyperLogLog algorithm
- Utility of measurements is that they work for all zones & are unbiased
- Downside is reduced accuracy
  - NSEC3 measurements can be ‘arbitrarily precise’ for large zones & delegation counts

Introduction to HyperLogLog

- “HyperLogLog is an algorithm for the count-distinct problem, approximating the number of distinct elements in a multiset”
- In short: smarter way of doing:
  - $ sort -u < big-data | wc -l
- “How many **distinct** things are in this list”
  - Without using a terabyte of memory
- HyperLogLog is simultaneously:
  - Utterly amazing and magical
  - Nothing more than a way to approximate number of distinct things
- **Tomorrow, Alexander Mayrhofer goes into more depth on HyperLogLog**
  - This presentation has the ‘hand-waving’ intro

Core concept for numbers

- Imagine bowl full of marbles, with random numbers between 1 and 1 billion
  - A typical marble I pick may have number 242,123,456
- If I happen to pick a marble with number 12 (out of 1 billion), I'd be pretty surprised
  - Even though I should not be! Could happen!
- The core idea of HyperLogLog: if I pick sufficient marbles, the **lowest** number I see tells me the number of distinct numbers in the entire bowl
  - Or I can pick another number than 0. The closest number to 500,000,000 also works
- Reality check: if there are 1 billion distinct numbers in the bowl, I'll eventually pick number 1
  - And if there is only 1 number in there, chance of this happening is 10^{-9}

More accuracy, and now for non-numbers

- As noted before, simply counting the lowest number ever seen can tell us size of distinct set
  - Pretty strong element of chance however
- Also noted: can also count numbers closest to other points than 0
- To improve accuracy: count closest number to 10 million, 20 million, 30 million etc
  - And take the mean of the predictions coming from that
- So how does this apply to non-numbers?
- Take a hash of course, and use that as a number
- Works really well

That does not sound too magical. But wait.

- To determine distinct elements in a terabyte of data naively will require more than 1 terabyte of storage
- How much memory does HyperLogLog use for reasonable precision?
  - **1.5 kilobyte**
- This is possible because statistics tell us we don't actually need the whole lowest number on the marbles we count
- If we denote the number 12 as "0000000012", turns out all we need to do is average the number of leading zeroes for all the "closeness" measurements we are doing
  - **1.5 kilobyte**!!

DNSSEC and non-existence

- DNSSEC signs record sets. DNS encode ‘domain does not exist’ by responding with an empty RRSET
  - Which can not usefully be signed!
- Clever thought was thought up: deny existence with an NSEC record that says ‘between powerdnr.com and powerdzz.com, nothing exists’
  - And sign that
- This leads to a directory of all names that DO exist however
  - “Follow the NSEC trail”
- A hashed variant was provided called NSEC3
  - Allegedly after the number of people that really understood how it worked
- Soon we will have NSEC5 which should improve on that number

HyperLogLog: Relation to NSEC3

- NSEC3 is a hash already!
- And, much like the bowl full of marbles, we can pick random NSEC3 records from a zone
  - “Just ask random questions”
  - Each answer is a number
- Even more interesting: each answer is TWO numbers!
  - Which we’ll make good use of
- **Based on the NSEC3 hashes, we can make use of HyperLogLog-like tricks to effectively count the number of authoritative names in a zone**

NSEC3: By eye (party trick!)

```
q2hnik5kkka91nki71r0elhqabmrudoi.nl.
   600 IN NSEC3 1 1 5 68..A6
Q2HOLFFJVRSBSH0RFQR8TI89NU3N7778
   NS DS RRSIG


1at7vb94mg2eajh8rof9nndjiafo68rc.lu.
   3600 IN NSEC3 1 1 3 1B..80
1EMRUNOARIQ9D2C7328V5UFJU2QSI91F
   NS DS RRSIG
```

“Dnssecmeasure” technology

- Algorithm used:
  - Gather all nameserver names for a zone
    - In parallel gather all IPv{4,6} addresses
    - Connect over TCP/IP to all these addresses
    - Send random queries
    - Until (say) 4096 have been answered
- Upside of TCP/IP is that we get built-in rate limiting
- Spreading out over all IPv{4,6} addresss means each server sees a few hundred queries or less
- Written in modern C++ 2014 (which has a lot going for it)
- Measures a typical zone in 0.3 seconds

NSEC3 specifics

- Look at a gathered NSEC3 range
- Use name of the range as the HyperLogLog starting point
  - Remember, does not need to be zero
- Store range in unique set (to prevent duplicates for small zones)
- Once sufficient numbers are in:
  - Determine distance between end-point of range and beginning
    - We cheat and only use first 64 bits of hash!
  - Express this distance as **fraction** of the 64-bits space
  - Average all those fractions
- The average ‘fraction’ covered by an NSEC3 range is the inverse of the number of NSEC3 ranges present

```
$ time ./dnssecmeasure nl
Will send 4096 queries to: sns-pb.isc.org ns2.dns.nl ns4.dns.nl
ns5.dns.nl ns3.dns.nl ns1.dns.nl nl1.dnsnode.net ns-nl.nic.fr
2001:500:2e::1 192.5.4.1 194.171.17.10 2001:610:0:800d::10
192.93.0.4 2001:660:3005:1::1:2 213.154.241.85 95.142.99.212
2001:7b8:606::85 193.176.144.5 2a00:1188:5::212 194.0.28.53
2001:678:2c:0:194:0:28:53 194.146.106.42
2a00:d78:0:102:193:176:144:5 2001:67c:1010:10::53
nl poisson size 2.50137e+06
Based on 4096 queries, 3931 distinct ranges
Saw 3930 ranges that started secure
{"delegation-estimate": 2501373, "dnssec": true, "nsec-type":
"NSEC3", "queries": 4096, "secure-delegation-estimate":
2500737, "secure-factor": 0.9997456118036123 2, "zone": "nl"}

real   0m0.262s user 0m0.064s sys  0m0.064s
```

POWERDNS
AN OX COMPANY


```
2048:  nl poisson size 2.42579e+06
4096:  nl poisson size 2.50193e+06
8192:  nl poisson size 2.49363e+06
16384: nl poisson size 2.57947e+06
32768: nl poisson size 2.60383e+06
65536: nl poisson size 2.64482e+06
```

POWERDNS
AN OX COMPANY

NSEC3: opt-in & “in-zone stuff”

- Quite rare
  - .TOP, .DE
- For opt-in: every NS record gets an NSEC3
  - So does not tell us a lot about secure delegation or not
- For Germany, zone still contains all kinds of MX records, IP addresses etc
  - Sorta interesting service by the way
- **Dnssecmeasure** solution:
  - For each NSEC3 range, we tally if it included a DS in the type set
  - Determine ratio between NSEC3s with DS and without
  - Pro-rate the result
- Outcomes interesting

```
$ time ./dnssecmeasure de 16384
Will send 16384 queries to: n.de.net l.de.net s.de.net z.nic.de
f.nic.de a.nic.de
2001:67c:1011:1::53 2001:668:1f:11::105 195.243.137.26
194.146.107.6 77.67.63.105 194.246.96.1 2a02:568:0:2::53
81.91.164.5 194.0.0.53 2001:678:2::53
de poisson size 434992
Based on 16384 queries, 15678 distinct ranges
Saw 2569 ranges that started secure
{"delegation-estimate": 434991, "dnssec": true, "nsec-type":
"NSEC3", "queries": 16384, "secure-delegation-estimate": 71277,
"secure-factor": 0.16386018624824594, "zone": "de"}

real   0m1.961s
```

POWERDNS
AN OX COMPANY

```
$ time ./dnssecmeasure top 16384
Will send 16384 queries to: g.zdnscloud.com f.zdnscloud.com
j.zdnscloud.com b.zdnscloud.com i.zdnscloud.com c.zdnscloud.com
e.zdnscloud.com a.zdnscloud.com d.zdnscloud.com
42.62.2.16 182.131.23.22 2401:8d00:2::1 2401:8d00:1::1
119.167.248.154 1.8.240.1 1.8.242.1 1.8.241.1 1.8.243.1

top poisson size 3.64338e+06
Based on 16384 queries, 16210 distinct ranges
Saw 10 ranges that started secure
{"delegation-estimate": 3643381, "dnssec": true, "nsec-type":
"NSEC3", "queries": 16384, "secure-delegation-estimate": 2247,
"secure-factor": 0.00061690314620604567, "zone": "top"}

real    0m10.484s
```

POWERDNS
AN OX COMPANY

NSEC? Can we do the same?

- Still NSEC zones around
  - Which we could theoretically just “walk” to get results
  - Although they try to make this somewhat harder
- **An NSEC name.. somewhat looks like a number if you squint!**
- Have to take a little bit of care to convert so no huge “gaps” are left
  - So, map ‘0’-’9’ onto 0-9
  - And ‘a’-’z’ to 10-36
  - Map - to 37
- Only look at first 12 characters
- Next apply same algorithm, determine how much of the NSEC “range” gets covered (37^12)

Bert, this can't **POSSIBLY** work

- For NSEC3 we had the luxury of working with nice hashes
  - Well distributed, lots of iterations
  - Fully random
  - Leave no trace of human foibles
- NSEC is **nothing** like that
  - The ranges are heavily influenced by human naming conventions
  - Not random at all
- Behold, the central limit theorem:
- **"In probability theory, the central limit theorem (CLT) establishes that, for the most commonly studied scenarios, when independent random variables are added, their sum tends toward a normal distribution (commonly known as a *bell curve*) even if the original variables themselves are not normally distributed."**

```
$ time ./dnssecmeasure se 16384
Will send 16384 queries to: j.ns.se e.ns.se g.ns.se c.ns.se
b.ns.se a.ns.se d.ns.se x.ns.se i.ns.se f.ns.se
81.228.10.57 130.239.5.114 199.254.63.1 192.36.135.107
2001:6b0:e:3::1 2001:500:2c::1 192.36.144.107
2001:67c:254c:301::53 2001:67c:2554:301::53 2a01:3f0:0:301::53
192.36.133.107 194.146.106.22 81.228.8.16 2001:67c:1010:5::53
213.108.25.4 2001:67c:124c:e000::4 2a01:3f0:0:305::53
192.71.53.53
se poisson size 1.43591e+06
Based on 16384 queries, 7884 distinct ranges
Saw 3312 ranges that started secure
{"delegation-estimate": 1435908, "dnssec": true, "nsec-type":
"NSEC", "queries": 16384, "secure-delegation-estimate": 603212,
"secure-factor": 0.4200913242009322, "zone": "se"}
```

IIS .SE number of DNSSEC domains

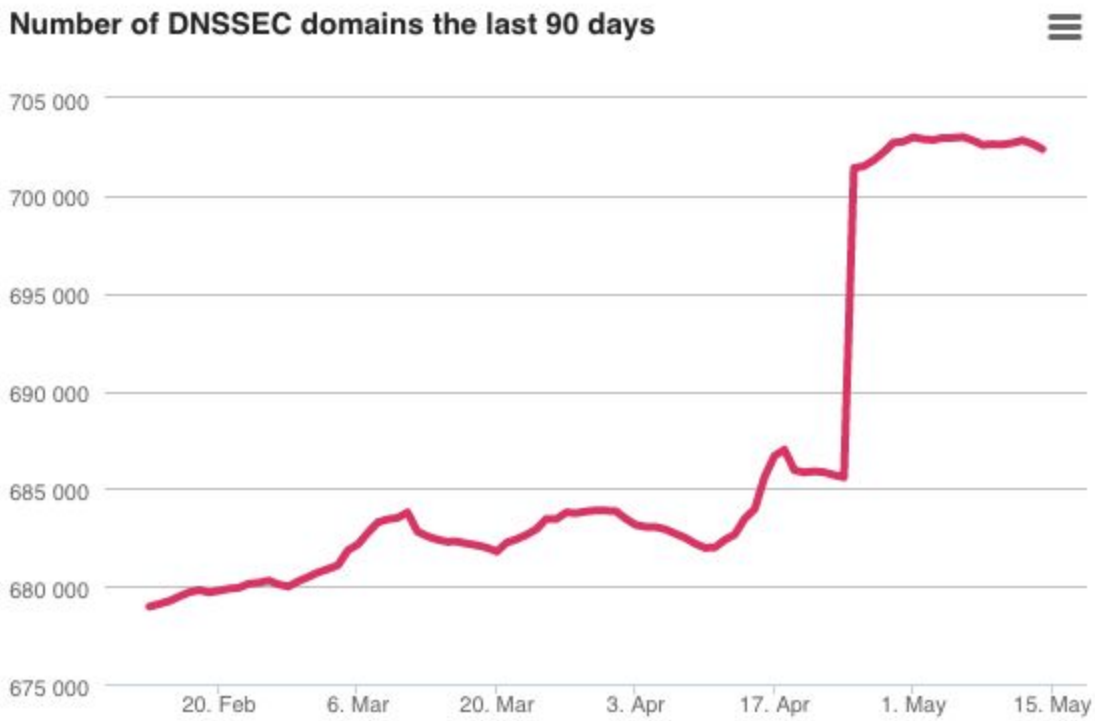

POWERDNS
AN OX COMPANY

A few words on precision

- NSEC3 appears to be highly precise
  - Thanks to the cryptographic properties of the hashes used
  - Can be made <1% precise with ~65k queries
- NSEC measurement is vulnerable towards outlier gaps
  - If you have no domain that starts with a ‘c’, the gap between ‘b’ and ‘d’ will make your zone extremely small (your language might not have a c)
  - Conversely, if you have powerdns12345.se and powerdns12346.se in your zone, this creates the impression of 37^{12} delegations
  - NSEC measurements must therefore be post-processed and weighed for likelihood
- Important to note that these measurements determine number of signed DS record names
  - This may not exactly be what gets counted by registry as “secured delegations”

DEMO

If it all works

Total number of secure delegations: 7375455

Zone	DNSSEC	NSEC(3)	Signed
nl.	true	NSEC3	2618262
com.br	true	NSEC3	771183
cz.	true	NSEC3	638970
se	true	NSEC	607290
com.	true	NSEC3	588136
no.	true	NSEC3	431968
eu.	true	NSEC3	365356
fr.	true	NSEC3	337751
be.	true	NSEC3	132395
net.	true	NSEC3	113548
hu.	true	NSEC3	111436
nu.	true	NSEC3	84151
org.	true	NSEC3	74039
de.	true	NSEC3	62775
pl.	true	NSEC3	32193
info.	true	NSEC3	28651
dk.	true	NSEC3	22422
ovh.	true	NSEC3	21505
gov.	true	NSEC3	19157
co.uk	true	NSEC3	17909
ch.	true	NSEC3	16764
pt.	true	NSEC3	15334
mx.	true	NSEC3	9861

A few words on what this means

- The reported numbers are **pretty dire** outside of the known success zones
  - CZ, SE, NL, COM.BR, NO, EU, FR
  - Registrars from those zones “radiate out” to other TLDs
    - Most BE DNSSEC registrations come from Dutch registrars (perhaps NET too)
- Success in these regions is due to incentive programs
  - Dutch registrars signed their **low value** domains first
    - Since “customers” unlikely to complain
- First thing that happens on sign of trouble: remove DS
  - And kill your .ORG delegation
- **DNSSEC is still on life support**
- Monitoring numbers important at this stage
  - This tool may be helpful

Further reading & where to get this software

- Longer writeup: https://ds9a.nl/hypernsec3/
- https://powerdns.org/dnssec-stats/ with numbers
- Actual software used:
  - https://github.com/ahupowerdns/pdns/tree/measurensec2
  - ./bootstrap ; ./configure --with-modules=””; make ; cd pdns ; make dnssecmeasure

HyperLogLog inspired three-minute scan of DNSSEC delegations worldwide

DNS-OARC 26, Madrid
@PowerDNS_Bert
bert.hubert@powerdns.com

POWERDNS
AN OX COMPANY