

Agenda

NCAP Discussion Group | 6 January 19:00 UTC

1. Update on Study 2
2. Work items for the discussion group:
 - a. Case studies on CORP, MAIL, HOME and other strings
 - b. Data sensitivity analysis
3. Review measurements and data that reflect the current state of the DNS
4. Review JAS and Interisle measurements and discuss to what extent do we repeat and modify those measurements on current data

Table of Contents

ADMINISTRATION.....	2
CURRENT STATE OF THE DNS.....	2
State of DNS Today in Name Collision Context	2
Slide 2: Reserved Strings.....	2
Slide 3: ITHI Measurement	3
Slide 4: Current Top Strings	3
Slide 5: ITHI Measurement table.....	4
Slide 6: RSSAC002 Data	4
Slide 7: Qtype Changes.....	5
Slide 8: DNS Consolidation	6
Slide 8: DNS Consolidation 2.....	6
Slide 10: DNS Consolidation 3.....	7
Slide 11: Qname Minimization.....	7
Slide 12: Qname Minimization 2.....	8
Slide 13: Qname Minimization and resolver characteristics	8
Slide 14 & 15: .HOME & .INTERNAL.....	9
Slide 16: CORP AND HOME	9
Slide 17: Hardcoded Labels and DNS Service Discovery.....	9
Slide 18: Hardcoded Labels and DNS Service Discovery 2	10
ACTION ITEMS:	10

ADMINISTRATION

lead by Matthew Thomas

SOI Changes

- James Galvin (Afilias): Affiliates has been acquired by donuts.

Matthew Thomas is now a member of SSAC

CO-CHAIRS

Charter calls for 3 co-chairs and now need someone from the community for one of the co-chairs

CURRENT STATE OF THE DNS

SSAC has until 11 Jan to respond to Study proposal and will present the BTC with it on 28 Jan meeting. Asking BTC to fund Study 2, and OCTO will take ownership of project. Proposal includes tech writer with goal of 1 March to start.

Data analysis will be based on case studies ono.corp, .mail, .home
For consideration

- perform what type of data sensitivity analysis measurements?
- What benchmarks. compare DNS telemetry data from the previous rounds to current state?
- how did Consolidation or q&a minimization or other protocol changes impacted the data?

State of DNS Today in Name Collision Context

https://docs.google.com/presentation/d/1y9ydZYLSM_wa0RjI7j5i8a10ZhlzZia7UeB0XWer2Ag/edit#slide=id.gaa6cc16472_0_17

Highlights of slides (focus is route server traffic, no recursive data)

Slide 2: Reserved Strings

Reserved Strings

RFC 2606 – Reserved Top Level DNS Names

To safely satisfy these needs, four domain names are reserved as listed and described below.

```
.test  
.example  
.invalid  
.localhost
```

.test is recommended for use in testing of current or new DNS related code.

.example is recommended for use in documentation or as examples.

.invalid is intended for use in online construction of domain names that are sure to be invalid and which it is obvious at a glance are invalid.

The *.localhost* TLD has traditionally been statically defined in host DNS implementations as having an A record pointing to the loop back IP address and is reserved for such use. Any other use would conflict with widely deployed code which assumes this use.

Name	Reference
0tsch.arpa.	[RFC-ref-0tsch-minimal-security-15]
10.in-addr.arpa.	[RFC6761]
16.172.in-addr.arpa.	[RFC6761]
17.172.in-addr.arpa.	[RFC6761]
18.172.in-addr.arpa.	[RFC6761]
19.172.in-addr.arpa.	[RFC6761]
20.172.in-addr.arpa.	[RFC6761]
21.172.in-addr.arpa.	[RFC6761]
22.172.in-addr.arpa.	[RFC6761]
23.172.in-addr.arpa.	[RFC6761]
24.172.in-addr.arpa.	[RFC6761]
25.172.in-addr.arpa.	[RFC6761]
26.172.in-addr.arpa.	[RFC6761]
27.172.in-addr.arpa.	[RFC6761]
28.172.in-addr.arpa.	[RFC6761]
29.172.in-addr.arpa.	[RFC6761]
30.172.in-addr.arpa.	[RFC6761]
31.172.in-addr.arpa.	[RFC6761]
168.162.in-addr.arpa.	[RFC6761]
170.0.0.192.in-addr.arpa.	[RFC6761]
171.0.0.192.in-addr.arpa.	[RFC6761]
254.169.in-addr.arpa.	[RFC6762]
8.a.f66.arpa.	[RFC6762]
9.a.f66.arpa.	[RFC6762]
a.a.f66.arpa.	[RFC6762]
b.a.f66.arpa.	[RFC6762]
home.arpa.	[RFC6762]
example.	[RFC6761]
example.com.	[RFC6761]
example.net.	[RFC6761]
example.org.	[RFC6761]
invalid.	[RFC6761]
ip-v4only.arpa.	[RFC6762]
local.	[RFC6762]
localhost.	[RFC6761]
onion.	[RFC6766]
test.	[RFC6761]

- Where reserved strings are falling within the the top distribution of names at the root servers

Slide 3: ITHI Measurement

ITHI Measurement

Queries to RFC 6761 reserved names (?)

In the following table, the *current value* is the fraction of queries to the root directed at RFC 6761 names in the current month. The table also provides the average value over the 3 previous months, and the "historical" minimum and maximum observed since the beginning of the measurements.

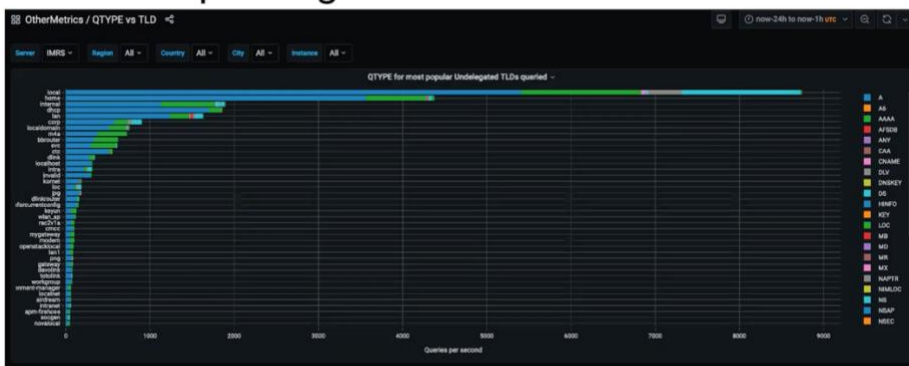
RFC 6761 name	As of Dec 2020	Past 3 months	Historic Low	Historic High
LOCAL	6.106%	4.607%	2.360%	5.079%
LOCALHOST	0.370%	0.265%	0.206%	0.561%
INVALID	0.260%	0.203%	0.191%	0.341%
TEST	0.075%	0.043%	0.008%	0.049%
ONION	0.011%	0.007%	0.002%	0.008%
EXAMPLE	0.003%	0.002%	0.001%	0.020%

<https://ithi.research.icann.org/graph-m3.html>

- Data from ICANN’s Identifier Technology Health Indicators (ITHI)
What is happening with strings that have had much change vs those that didn’t

Slide 4: Current Top Strings

Current Top Strings



<https://stats.dns.icann.org/stats/d/wom-ext-5minagg-qtype-vs-tld/qtype-vs-tld?orgId=1>

- Data from ICANN’s IMRS system, looking at non-delegated, non-existence domains.
 - shows top n based queries per second that L Root is receiving
 - .home is interest to NCAP is up at #2 position. Corp is #5, .mail is in top 200. A few strings from our case studies are in this data analysis

-not all strings have similar Qtype distributions; various Qtypes. What does the type distribution for that string actually entail and tell you about the risk factor of that particular strain. Are those key types associated with service discovery protocols?

Slide 5: ITHI Measurement table

ITHI Measurement

Queries to frequently leaked strings (2)

In the following table, the current value is the fraction of queries to the root directed at frequently used strings in the current month. The table also provides the average value over the 3 previous months, and the "historical" minimum and maximum observed since the beginning of the measurements.

Frequently used string	As of Dec 2020	Past 3 months	Historic Low	Historic High
HOME	3.218%	3.034%	2.478%	3.674%
INTERNAL	1.744%	0.714%	0.301%	0.971%
LAN	1.096%	1.134%	0.499%	1.306%
DHCP	0.890%	0.787%	0.206%	1.285%
CORP	0.808%	0.439%	0.191%	0.572%
LOCALDOMAIN	0.492%	0.516%	0.298%	0.804%
DLINK	0.289%	0.244%	0.141%	0.341%
BBROUTER	0.283%	0.279%	0.000%	0.379%
CTC	0.279%	0.138%	0.000%	0.232%
INTRA	0.240%	0.174%	0.000%	0.270%
MMA	0.208%	0.223%	0.000%	0.278%
SVC	0.184%	0.000%	0.000%	0.000%
LOC	0.179%	0.131%	0.000%	0.178%
JPG	0.152%	0.039%	0.000%	0.116%
OPENSTACKLOCAL	0.122%	0.000%	0.000%	0.396%
GETCACHEDDHCPRESULTSFORCURRENTCONFIG	0.121%	0.076%	0.000%	0.129%
DLINKROUTER	0.118%	0.151%	0.123%	0.220%
IP	0.000%	0.000%	0.000%	1.010%
DHCP HOST	0.000%	0.000%	0.000%	0.877%
GATEWAY	0.000%	0.152%	0.000%	0.234%
_SMB	0.000%	0.000%	0.000%	0.304%
BELKIN	0.000%	0.000%	0.000%	0.188%

WLAN_AP	0.000%	0.077%	0.000%	0.181%
LAN1	0.000%	0.044%	0.000%	0.169%
HOMESTATION	0.000%	0.000%	0.000%	0.190%
MODEM	0.000%	0.166%	0.000%	0.222%
NULL	0.000%	0.000%	0.000%	0.220%
MYGATEWAY	0.000%	0.101%	0.000%	0.188%
DAVOLINK	0.000%	0.000%	0.000%	0.191%
DNSBL	0.000%	0.000%	0.000%	0.227%
WORKGROUP	0.000%	0.000%	0.000%	0.153%
COMM	0.000%	0.000%	0.000%	0.367%
CONSUL	0.000%	0.000%	0.000%	0.137%
SJY	0.000%	0.000%	0.000%	0.346%
KEYUN	0.000%	0.000%	0.000%	0.167%
RACV1A	0.000%	0.042%	0.000%	0.126%
MYNET	0.000%	0.000%	0.000%	0.121%
KSCN	0.000%	0.000%	0.000%	0.114%
DOESNOTEXIST	0.000%	0.000%	0.000%	0.185%
LOCAL_OK	0.000%	0.000%	0.000%	0.153%
MRDKA	0.000%	0.000%	0.000%	0.130%
NO-USE-NOW	0.000%	0.000%	0.000%	0.110%
TOTOLINK	0.000%	0.033%	0.000%	0.100%


<https://ithi.research.icann.org/graph-m3.html>

- table from report published monthly, -this gives another set of data points to reaffirm probably break priority listing of what the most popular prevalent Leaking strengths is

-Where are ITHI measurements are coming from? Mostly L Route data, but is a small amount of data being used, and OCTO is looking for more sources to add to ITHI

Slide 6: RSSAC002 Data

RSSAC002 Data




<https://www.potaroo.net/ispcol/2020-09/root-fig1.png>

- Server Operators are a Data source (in name collision context) to illustrate how route system has scaled to since the data analysis back in 2012-2013 so far. Sample is this chart.

-shows query increase

Slide 7: Qtype Changes

Qtype Changes




- Graph by Mathew Thomas from recent data he pulled of of A&J.
- a couple months ago, Apple iOS 14 launched, and made change in DNS query behavior by doing an additional DNS query for what is known as an HTTP s q type (value 65) now in certain conditions where the application is specifically trying to make a DLS connection or the user has entered HTTPS-like-URL so that DNS query is sent an addition to the A and quad A record that's coming out of those Apple devices.

- previously relatively stable set of Q types at the root; now is so disrupted by HTTPS record it is 5th most popular qType you see at the route
 - o what is causing leakage in Apple devices

Slide 8: DNS Consolidation

DNS Consolidation




<https://ant.isi.edu/~imana/presentations/lmana18b.pdf>

- Graph shows data from recent measurement made by West H. at B Root. Distribution of resolvers that B Root saw.
 - how does the consolidation of the DNS infrastructure has changed over time.
 - graph shows a few resolvers making a lot of queries; less resolver that make only a few queries
 - interesting for comparison year by year and what does that impact on data sensitivity analysis.
- Cash efficiencies to prevent leaking queries to root system.

Slide 8: DNS Consolidation 2

DNS Consolidation

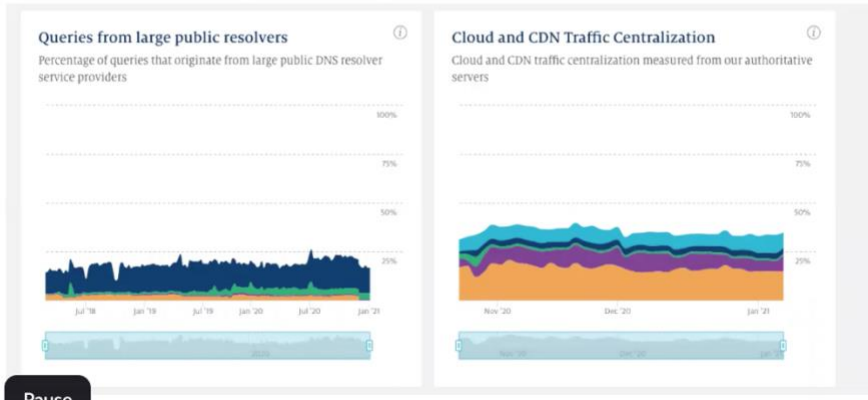


- Data from same company as last chart, shows that popular open recursive or cloud providers are accounting for more of a larger % of traffic seen at upper levels of DNS hierarchy

Slide 10: DNS Consolidation 3


- Chart shows that queries from large public recursive resolvers are becoming more than norm of their traffic

DNS Consolidation



Slide 11: Qname Minimization

Qname Minimization



<https://blog.verisign.com/security/maximizing-qname-minimization-a-new-chapter-in-dns-protocol-evolution/>

- Impact of qname minimization on DNS which has become turned on as a Default setting.

As resolvers upgrade or update their software we hope to see the impacts of qname min. for privacy reasons. For privacy reasons, we will be blocked from pulling info in labels. 50% queries only come in with 1 label.

Slide 12: Qname Minimization 2

Qname Minimization



Slide 13: Qname Minimization and resolver characteristics

Qname Minimization and resolver characteristics

Characteristics of resolvers seen at the root

Additional metrics characterize the options found in queries sent to the root, such as whether resolvers use extended DNS (M3.4.1), what EDNS options they use (M3.4.2), whether they set the DNSSEC OK bit in queries (M3.5), and whether they appear to enforce QName Minimization (M3.6)


Metric		As of Dec 2020	Past 3 months	Historic Low	Historic High
M3.4.1 (2)	%resolvers using Extended DNS (EDNS)	89.707%	90.360%	88.683%	91.821%
	%resolvers using COOKIE(10)	8.394%	7.965%	3.146%	8.243%
	%resolvers using edns-client-subnet(8)	0.000%	0.072%	0.000%	0.316%
	%resolvers using 65001	0.000%	0.000%	0.000%	0.012%
M3.4.2 (2)	%resolvers using EDNS EXPIRE(9)	0.000%	0.000%	0.000%	0.002%
	%resolvers using NSID(3)	0.000%	0.000%	0.000%	0.001%
	%resolvers using 65433	0.000%	0.000%	0.000%	0.000%
	%resolvers using 65435	0.000%	0.000%	0.000%	0.000%
M3.5(2)	%resolvers setting DNSSEC OK (DO) flag	84.521%	84.533%	81.920%	85.881%
M3.6(2)	%resolvers using QName minimization	34.540%	29.555%	11.700%	32.582%

<https://iithi.research.icann.org/graph-m3.html>

- Additional measurement from Site In (?) looking at .nl
- From CHI (?), they have more stringent requirements on the # of queries coming out of the resolver.


Slide 14 & 15: .HOME & .INTERNAL

.HOME



<https://www.icann.org/en/system/files/files/sac-113-en.pdf>

.INTERNAL




<https://www.icann.org/en/system/files/files/sac-113-en.pdf>

- Traffic for .home: upwards of 4% of traffic at A&J was for .home. 1% of all AJ's traffic is for .internal
 - Discussion on 1 cause of spikes is expiring from large providers' caches. Based on how query comes in, and how it is cached upon the length of labels that caused the original lockup
 - .internal: traffic comes from diverse set of sources (150 -200 distinct ASNs were sending traffic, but 95% of traffic from one ASN, a large Cloud provider who says it is coming from their infrastructure and they will look into fixing it.
 - o Shows potential of reaching out to others to get them to clean up their naming collisions.

Slide 16: CORP AND HOME

CORP and HOME



- Both have seen upward growth.

Slide 17: Hardcoded Labels and DNS Service Discovery

