VERISIGN™

12061 Bluemont Way
Reston, VA 20190
t: 703-948-3200

VerisignInc.com

September 15, 2013

Board of Directors
Internet Corporation for Assigned Names and Numbers
12025 Waterfront Drive, Suite 300
Los Angeles, CA 90094-2536

Re: ICANN's Proposal to Mitigate Name Collision Risks – .CBA Case Study

Dear Members of the ICANN Board:

On August 27, 2013, Verisign submitted several comments regarding ICANN's New gTLD Collision Risk Mitigation proposal. In our comment, "New gTLD Security, Stability, Resiliency Update: Exploratory Consumer Impact Analysis," we demonstrated that query volume alone is an inadequate measure of risk arising from name collisions; provided a candidate risk assessment matrix; and advocated that ICANN finally implement the prior recommendations from ICANN's own Security and Stability Advisory Committee ("SSAC"), including the Scaling the Root studies. These recommendations were in line with Interisle Consulting Group's observation in its August 2013 report prepared at the direction of the ICANN Board that stated "[a]n additional qualitative analysis of the harms that might ensue from [name] collisions would be necessary to definitively establish the risk of delegating any particular string as a new TLD label…"[1]

Others choosing to prioritize speed over a secure and stable DNS operation have also submitted comments critical of ICANN's risk categorization based solely upon query volume, but have drawn the incorrect conclusion. These applicants contend that ICANN's proposal is too conservative and that delegations should be expedited, because, in their view, the risk of internal network name collisions is acceptable, even absent the qualitative analysis of each string as discussed in the Interisle Report. Other than Verisign's Exploratory Consumer Impact Analysis,

[1] See *Name Collision in the DNS (Version 1.5, August 2, 2013)*, Interisle Consulting Group, LLC at 2-3 ("Interisle Report").

no effort has been made to perform a qualitative analysis as recommended repeatedly, and absent this analysis, the potentially harmful consequences associated with the delegation of a new TLD label, and the associated risks, simply cannot be assessed.

Verisign's analysis has focused on identifying some of the systematic risks that will be exposed by the new gTLD program and who the impacted parties are likely to be. In this context, we were intrigued when the Commonwealth Bank of Australia ("CBA") filed a public comment taking responsibility for the name collisions identified in the Interisle Report and associated with its .CBA application.

Commonwealth Bank of Australia - .CBA

CBA has applied to operate the TLD string .CBA. The .CBA string was placed by ICANN in the "uncalculated risk" group (specifically at position 153) based upon the analysis set forth in the Interisle Report. In an effort to follow the guidelines of ICANN's New gTLD Collision Risk Mitigation proposal, and thus assess the risk of delegation of the .CBA string, CBA concluded in an August 23, 2013 comment to ICANN that after "some internal investigation," the cause of the name collisions is "primarily from CBA internal systems" and "that it is within the CBA realm of control to detect and remediate said systems and internal certificate use." Thus, CBA concluded that it could self-mitigate the name collision risks resulting from delegation of the .CBA TLD and that ICANN should therefore move the .CBA string to the "low risk" group of applications. CBA also represented that it would undertake a further investigation and verify the origins of the .CBA requests, and that it would conduct appropriate remediation upon the completion of its internal investigation. While we await that investigation and remediation plan, we conducted our own analysis of the .CBA queries. We believe our data and analysis shows without a doubt that CBA's initial conclusions are incorrect.

More broadly, we believe that the .CBA assessment can serve as a useful case to test the components of ICANN's risk mitigation proposal. Using .CBA as an example, we can assess whether ICANN's interpretation of the Interisle Report and ICANN's risk categorization is appropriate. Further, we can assess how effective ICANN's plan to require applicants (and not ICANN) to conduct risk mitigation is likely to be.

Analysis of .CBA Queries

Verisign conducted a focused study on .CBA and then further narrowed its focus on two of the most active geographies (namely Japan and Brazil) in terms of errant queries to the root server system for .CBA for seven weeks based on data from 1.5 roots (100% of "A," and ~50% of "J"). While the coverage from the vantage points of these servers suggests that our data has a correspondingly limited scope, we believe our analysis demonstrates that CBA's analysis is measurably unreliable and inaccurate, as root server system query behaviors simply do not exist within CBA's observation space, and are themselves source-anonymized or otherwise incomplete within the annual "day in the life (DITL)" repositories. Furthermore, we believe that a reasonable conclusion to draw is that ICANN's risk mitigation proposal is not a practical or reliable option.

During the seven-week study period, Verisign observed approximately 10,000 root queries for the .CBA TLD per day. Many of the queries we observed (~80%) related to .CBA are from systems or devices utilizing DNS-based Service Discovery protocols such as BONJOUR, which, when signaled, typically respond with available services. These services often include printers, smart home and industrial automation systems, or other specialized devices/services. (BONJOUR predates ICANN's new gTLD program by several years, and DNS-SD was recently published as a Standards Track RFC by the IETF.) Other queries appear to be the result of standards-based DNS resolution search list processing issues, and originate from McAfee's popular anti-virus software, which transmits queries to confirm in real-time that certain resources are virus-free prior to execution and processing. Finally, the originator of many of the queries simply could not be identified absent direct interaction, which was deemed out of scope of this study but considered necessary to fully understand the scope of the problem in any qualitative manner.

It should be noted that the Internet Service Providers ("ISPs") that operate the network connectivity services and recursive name server infrastructure (including NTT-ME, Stelmat, AT&T, Comcast, Telefonica, Telstra, Level(3) Communications, Embratel, Cox, Bell Canada, etc.) for the observed .CBA query sources may be impacted by the end system operators. This is because these end system operators will most likely contact these ISPs for support services, and problem resolution assistance resulting from impacts of name collisions. The disruption to the ISP's business and residential networks and their subscribers could be substantial. Under ICANN's proposed risk mitigation plan, there is no mechanism to assist these infrastructure operators, or any acknowledgment of them as impacted parties. Commonwealth Bank of Australia, as registry operator, will bear sole financial liability to any third party, whether it be an injured ISP, network systems operator or otherwise, arising from the results of name collisions that occur as the proximate result of delegation and operation of the .CBA TLD. Indeed, should an injured third party seek to hold ICANN accountable for its losses in this regard, Commonwealth Bank of Australia as registry operator will be obligated to indemnify and defend ICANN and its agents for such claims.[2] Furthermore, it should be noted that Commonwealth Bank of Australia's indemnification obligation to ICANN, and as a result its liability, is uncapped.[3] These uncapped indemnification obligations, as well as uncapped liability, exist for all applicants that execute the New gTLD Registry Agreement.

While there were two major functional namespaces observed and identified, .CBA is used by more than 100 others in Japan alone. A sampling of the DNS-SD and McAfee queries can also be seen in the accompanying .CBA Focused Analysis presentation. These represent two new functional namespaces, beyond those identified in Verisign's Exploratory Consumer Impact Analysis. It should be noted that whatever the nature of many of the applications relying on non-delegation of .CBA, some at least appear sufficiently important to justify the use of real-time malicious software (malcode) protection techniques. The reliance on the non-delegation of .CBA for the correct operation of the malcode protection software itself is likely a result of standard search list processing issues that permeate nearly all of the applied-for strings. In

[2] See New gTLD Registry Agreement §7.1(a) ("Registry Operator shall indemnify and defend ICANN and its directors, officers, employees, and agents from and against any and all third-party claims, damages, liabilities, costs, and expenses including reasonable legal fees and expenses, arising out of or relating to…the delegation of the TLD to Registry Operator, Registry Operator's operation of the registry for the TLD or Registry Operator's provision of Registry Services…").

[3] See New gTLD Registry Agreement §5.3.

addition to inherent privacy issues with potentially sensitive information leakage, this finding may more specifically signal the expanded potential and ease of launching Man-In-The-Middle (MiTM) attacks as well as more easily enable the injection of malicious software when .CBA begins to resolve. Simple explanations of how spoofing could target these networks can be found in the public comment on the Proposal to Mitigate Name Collision Risks submitted by Dr. Eric Osterweil on September 11, 2013.[4]

Our measurement study reveals evidence of a substantial Internet-connected infrastructure in Japan that lies beneath the surface of the public-facing Internet, which appears to rely on the non-resolution of the string .CBA. This infrastructure appears hierarchical and seems to include municipal and private administrative and service networks associated with electronic resource management for office and residential building facilities, as well as consumer devices. However, further study is required to determine the nature and full extent of this infrastructure. A sampling of these queries can be seen in the accompanying .CBA Focused Analysis presentation.

While predominantly originating in Japan and Brazil, the data shows .CBA queries originating from more than 170 countries over less than two months of observation within a fraction of the root server system. Most registry operators much like Commonwealth Bank of Australia and their new gTLD partners, would face significant challenges in studying and identifying queries that originate from all over the globe. Preparing effective notifications in the official languages of the more than 170 countries from which .CBA queries originated in order to warn potentially impacted parties is a challenging task for anyone and even more so under the short, 30-day time frame proposed by ICANN. Further, assuming Commonwealth Bank of Australia successfully notifies the impacted parties, ICANN's plan provides limited time for the impacted parties to actually remediate the name collisions and prevent disruptions or the introduction of vulnerabilities, and does not even provide for any means of educating potentially impacted parties of how to mitigate the issue once they are notified. And, finally, under ICANN's proposal, ICANN is and will remain unaware of the naming collisions that Commonwealth Bank of Australia identifies and seeks to remediate.

[4] See http://forum.icann.org/lists/comments-name-collision-05aug13/msg00038.html.

Conclusions

While we believe even this analysis could be far more comprehensive, these results support the following conclusions:

- Root server system instrumentation is critically important, as ICANN's own SSAC has recommended repeatedly since 2010. Such a capability would allow issues such as those described here to be surveyed, studied, and identified, and responded to in a consistent, cohesive and deliberate manner for all applied-for strings by experienced analysts prior to delegation and immediately upon the enablement of this capability.[5]

- Our study vindicates the Interisle Report's observation that additional qualitative analyses of the harms that might ensue from naming collisions is necessary to definitively establish the risk of delegating any particular string as a new TLD label.

- ICANN's risk mitigation plan improperly categorizes strings by arbitrary risk groups based on no apparent analysis beyond query volume, and with no survey whatsoever of the possible impacts.

- Applicants lack the experience and thus are a poor choice to perform such risk assessments and to operationalize the ICANN-prescribed "customer service" (which ICANN has not defined). Lacking root instrumentation, and thus unable to see much beyond their own internal usage of .CBA, the bank wrote to the ICANN public forum claiming that .CBA collisions could be self-mitigated. However, our analysis demonstrates that the bank is the source of at most 6% of the observed query volume.

[5] Verisign continues to convey our eagerness to stand with ICANN in their leadership role here, as well as in their role as L-root operator, and align with all other willing root operators to as quickly as possible develop a sustainable root server system measurement apparatus to provide early warning and instrumentation across the system at the root. We would also note with dismay that this work is NOT currently underway within the existent work plan, nor has it been.

We are unaware of how many of the 952 or more Internal Names Certificates they possess, a topic that would need separate study and analysis.

- Our data vindicates the observation that applicants face increased risk of liability from end users and network operators upon delegation. Under ICANN's current constructs, applicants will bear this risk alone, and will indemnify ICANN should the delegation give rise to claims against ICANN.

Verisign's risk analysis of just one string of more than 1,400 in just two of 170 geographies demonstrates that the as-yet unknown potential risks from new strings is real, not just theoretical. ICANN should consider this new data and analysis, and carefully review Verisign's Exploratory Consumer Impact Analysis, integrating this with a comprehensive risk matrix and community input of what weights should be applied to each element of the risk matrix.

ICANN's mission and primary priority is to coordinate and ensure the secure and stable operation of the DNS. ICANN, after five years and hundreds of millions of dollars invested by applicants and others in the new gTLD program, continues to ignore sound warnings from industry experts and their very own advisory committees, and as a result have failed to identify the readily discoverable and at-risk DNS usage described here. The community is owed an explanation as to why ICANN made the decision to limit the scope of the Interisle Report, and then subsequently decided not to expand the scope of the Report when its shortcomings, identified by Interisle itself as well as others, were plainly evident.

This analysis of .CBA demonstrates clearly how little can be known confidently until qualitative analysis of each individual string is conducted, and that what the community does not know can have unforeseen consequences, which could be severe. This is why SSAC, Interisle, and many others strongly advise that individual string risk analysis - the only way to categorize a string as anything other than unknown risk - should be performed and assessed prior to any delegation. The analysis also validates the concerns of parties such as General Electric, Verizon, the American Insurance Association, and the ISP and Connectivity Provider (ISPCP) constituency, all of whom asked ICANN through submitted comments for additional time and study prior to proceeding with ICANN's plan as proposed.

ICANN has reacted late to name collisions in a manner that does not properly prioritize security and stability. The facts make this clear: ICANN failed to act on SSAC advice; the Interisle study was late, limited, with no follow-up to its findings; no SSAC analysis has been published that would inform community comments; no effort has been made to develop qualitative risk analysis; new responsibilities and risks are transferred to applicants, without the tools to address them; and the broader community of affected parties is only now beginning to research the potential for impact for themselves and their clients.

With its risk mitigation plan, ICANN proposes to transfer certain security and stability responsibilities to applicants - a policy that subverts ICANN's core mission. This should be soundly rejected by the ICANN board, not ratified and continued.

This study of only one string shows that no one, ICANN or others, should assert or assume that collision risk is understood and acceptable without conducting proper risk analysis and incorporating informed community input.

Very truly yours,

Patrick S. Kane
Senior Vice President, Naming and Directory
Services
VeriSign, Inc.

Thomas C. Indelicarto
Vice President and Associate General Counsel
VeriSign, Inc.

Danny McPherson
Vice President, Chief Security Officer
VeriSign, Inc.

VERISIGN®

Focused Analysis on New Applied-For gTLDs (Focus: .cba)

September 12, 2013

The new gTLD program: risk vs. reward

- The new gTLD program offers a lot of positive opportunities
  - This is why Verisign has applied for over a dozen strings and is contracted to act as a backend registry for nearly 200 others
- But, the security and stability of the DNS is serious, and risk-taking at the DNS root has global implications
- ICANN’s community has fractured: do applied-for gTLD strings pose risks? “Yes,” “no,” “yes, but…”
- *Evidence* and *measurements* are critical in clarifying what the appropriate level of caution should be

Assessing risk can be done broadly and also in a focused manner

- In our recent Technical Report, we examined the broad "*spread*" of risk across all applied-for strings with a candidate *Risk Matrix*

- But, now we add a focused (per-gTLD) methodology
  - We add "who is impacted" to our "what's the problem" analysis
- We use network-level information (such as ASNs) and semantic information (namespaces) to identify impacted parties employing applied-for strings
  - How many namespaces are actually going to be impacted, and what might actually break for them
  - For example, why do we see DNS Service Discovery (SD) queries and virus scans that *seem* to be from residential apartment complexes?

To illustrate the efficacy of this approach…

- Hoping to alleviate concerns, some have sought to exonerate individual applied-for strings
- Commonwealth Bank of Australia claimed to be the primary source of .cba queries

  "As the cause of the name collision is primarily from CBA… it is within the CBA realm of control…"

  http://forum.icann.org/lists/comments-name-collision-05aug13/msg00004.html
- Interisle Consulting Group ranked .cba as 153rd out of all strings, w/ 952 Internal Names Certificates, and classified as "uncalculated risk"
- However, without broad root server system instrumentation and qualitative analysis, Commonwealth Bank simply cannot *know*
- This presentation illustrates a deeper analysis of .cba

Outline

- Reprise our analysis of network spread for .cba
- Namespace Definition and Characterization
- Findings and Future Work

Spread: Sources of queries for .cba

- We observed 504K queries for .cba related domains
  - NXD root traffic from 12 sites carrying both A+J between 7/16 and 9/5/2013
- Global interest but intense activity out of Japan

Country	Query Count	Percent
Japan	399044	79%
United States	47198	9%
Brazil	11591	2%
Australia	8255	2%
Canada	4728	1%
China	3525	1%
Germany	3332	1%
OTHER	26549	5%

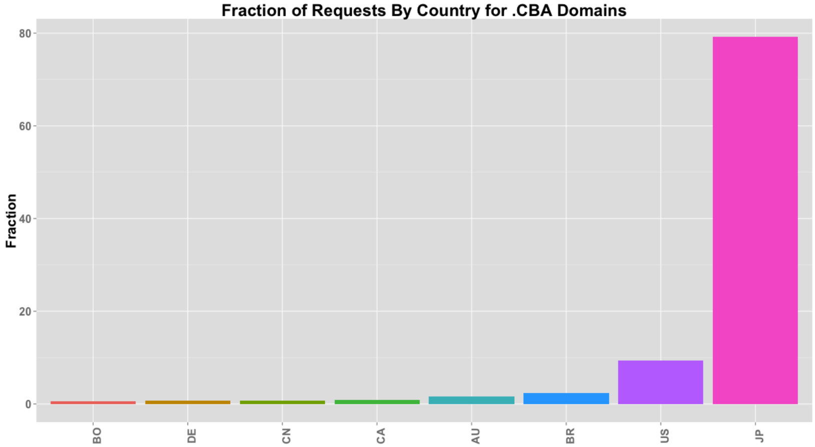

- Operational constraint resulted in limited visibility across our sites, impacting accuracy and depth of our analysis
  - Queries from NTT-ME in Chiba, Japan primarily hit Tokyo site; when it was added it changed the scope of the findings
  - Current study covered 100% of A root and an estimated 50% of J root

Heat map of query sources for .cba

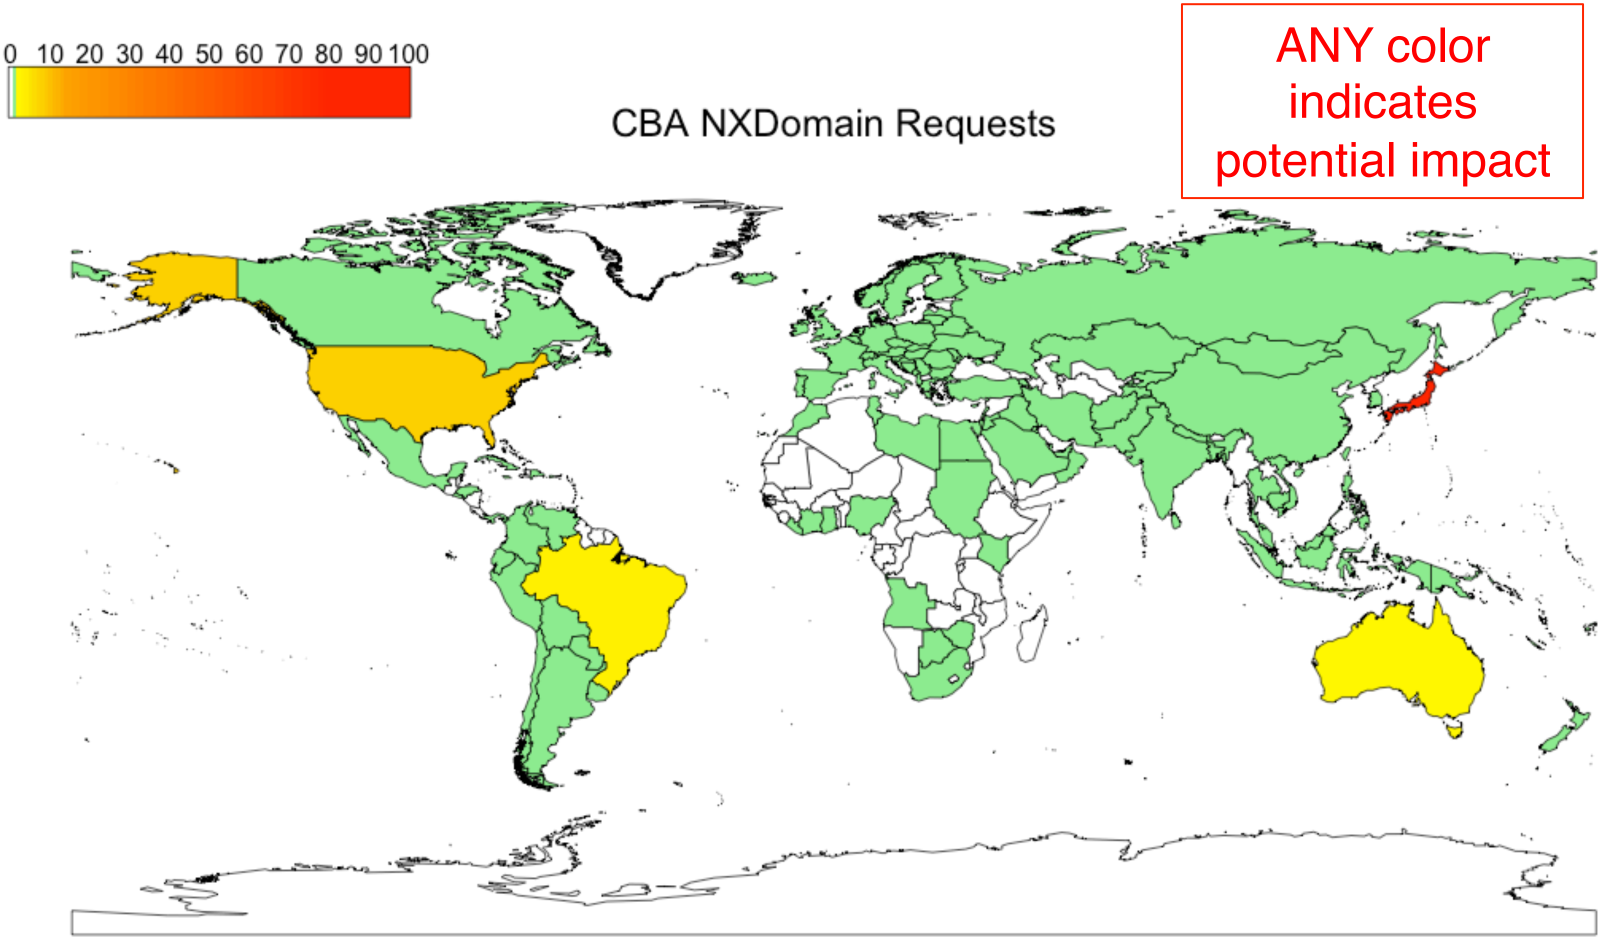

ANY color indicates potential impact

What network sources are making these queries

- 2,639 unique ASN’s across 171 countries responsible for queries
  - 1,785 (~68%) ASNs made more than 1 query over the collection period
  - Top 20 ASN’s account for ~90% of all queries
  - NTT-ME Corporation in Japan generates ~79% of all queries

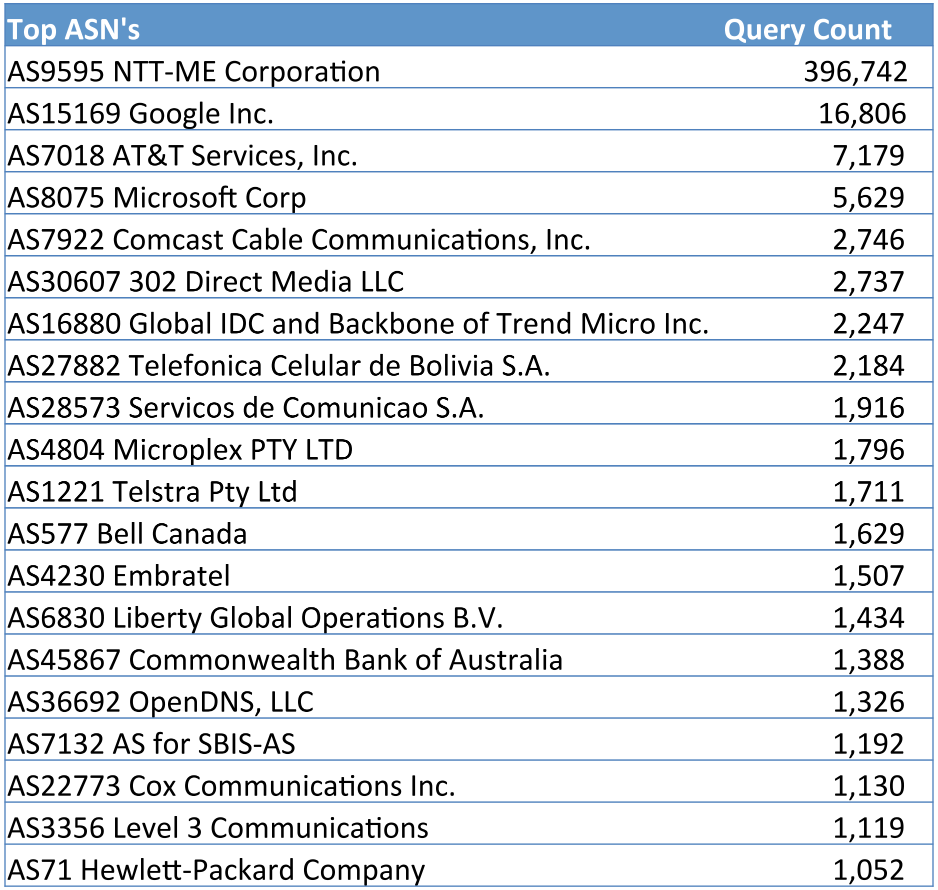

Top ASN's	Query Count
AS9595 NTT-ME Corporation	396,742
AS15169 Google Inc.	16,806
AS7018 AT&T Services, Inc.	7,179
AS8075 Microsoft Corp	5,629
AS7922 Comcast Cable Communications, Inc.	2,746
AS30607 302 Direct Media LLC	2,737
AS16880 Global IDC and Backbone of Trend Micro Inc.	2,247
AS27882 Telefonica Celular de Bolivia S.A.	2,184
AS28573 Servicos de Comunicao S.A.	1,916
AS4804 Microplex PTY LTD	1,796
AS1221 Telstra Pty Ltd	1,711
AS577 Bell Canada	1,629
AS4230 Embratel	1,507
AS6830 Liberty Global Operations B.V.	1,434
AS45867 Commonwealth Bank of Australia	1,388
AS36692 OpenDNS, LLC	1,326
AS7132 AS for SBIS-AS	1,192
AS22773 Cox Communications Inc.	1,130
AS3356 Level 3 Communications	1,119
AS71 Hewlett-Packard Company	1,052

Breaking down the full queries

- The queries vary in length
- Some query names can leak information about end user namespaces, services, protocols, and more
  - www.cba
  - cbadomain.cba
  - wpad.CBA
  - Example 14 label query:

0.0.0.157c. … .avqs.mcafee.com.gama.cba

Labels in Query	Number of Queries	Percent
1	22,192	4%
2	32,863	7%
3	217,145	43%
4	20,672	4%
5	116,814	23%
6	42,332	8%
7	9,552	2%
8	8,482	2%
9	760	0%
10	368	0%
11	23,543	5%
12	6	0%
13	2,556	1%
14	5,956	1%
15	640	0%
16	171	0%
17	132	0%
18	28	0%
19	2	0%
36	8	0%
	504,222	

What do we mean by “namespace”

- One primary goal is using semantics to identify impacted parties
  - In this work, we do this by analyzing “namespace” collisions
- Here we define namespace as a Second Level Domain (SLD) that is queried for by any of our risk vectors
  - WPAD, ISATAP, Bonjour/DNS-SD, McAfee GTI
  - Each represent automated configuration attempts or information leakage
  - Each define their namespace in the most significant labels
- We leverage this to identify the logical configuration boundaries of impacted parties
  - The semantic names of those who may suffer
- Here we focus on DNS-SD and McAfee GTI services

What is Bonjour?

"Bonjour, also known as zero-configuration networking, enables automatic discovery of devices and services on a local network…" https://developer.apple.com/bonjour/

- A DNS Service Discovery protocol for network services like:
  - printers, Apple TV, etc…
- Also enables ***smart home*** automation technologies like:
  - thermostats, remote and physical access systems, energy management, alarms, etc.., e.g.,:

  http://www.marvell.com/smart-energy/assets/Marvell-Smart-Energy-Platform-Brief.pdf

- DNS-SD queries leak from the local namespace when a machine thinks it exists in a "zone" that does not exist in global DNS
  - 106,881 (21.2%) of all CBA queries represent explicit DNS-SD related queries
- Anyone could potentially answer DNS-SD queries
  - So, when a Bonjour or other DNS-SD client asks where its printers are, *who* is really answering?

Namespace identification

- Distinctive query structure in the different service related activities can help us identify “namespaces”
  - makuharibaytown-mirama3.cba comes from Bonjour control queries:
    - lb._dns-sd._udp.makuharibaytown-mirama3.cba
    - b._dns-sd._udp.makuharibaytown-mirama3.cba
    - r._dns-sd._udp.makuharibaytown-mirama3.cba
  - stelmat.local.cba from other services, perhaps as a result of standards-based search list processing or explicit configuration
    - wpad.stelmat.local.cba
    - isatap.stelmat.local.cba
  - These namespaces can each be localized
    - makuharibaytown-mirama3.cba only receives traffic from Japan
    - stelmat.local.cba primarily receives traffic from Brazil

Generalized impact statement

- The namespaces conducting Bonjour and other DNS-SD queries account for 80% of all queries seen in CBA
  - Of the 65 namespaces:
    - 49 are based in Japan
    - 5 are based in Brazil
    - 2 are based in Canada
    - Only 1 sees activity from Australia and that namespace has the most diverse set of queries.
- We automate this portion across any TLD
  - Under TLD X there are Y number of namespaces relevant in Z region or country

DNS-SD/Bonjour-based namespaces

- 63 different namespaces are making DNS-SD queries

Example Namespaces
abando-nehara1.cba
a-m-13.cba
a-t-10.cba
a-takane4.cba
a-takane5.cba
b.national.cba
b.stelmat.local.cba
esc-oyumino.cba
etc-kaihinmakuhari.cba
g-n-2.cba
gp-sonnou-4.cba
gt-hikarigaoka.cba
ichikawaminami-3.cba
ichikawaminami-4.cba
maehara-b2.cba
makuhari-bavtown.cba
makuharibaymirama-ru.cba
makuharibaytown-mirama3.cba

- Stelmat is a networking company based in Cuiaba, Brazil (0.5% of CBA queries)

- Makuhari Baytown High-rise in Chiba, Japan (4.9% of CBA queries)

McAfee Global Threat Intelligence (GTI) leaks

“GTI provides the most up-to-date malware detection … receives the request from … endpoint, it determines if this program is malicious and responds appropriately.” https://kc.mcafee.com/corporate/index?page=content&id=KB53735

- GTI clients emit DNS queries whenever files (exe’s, pdfs, apks, etc.) are being checked for malware, essentially piggybacking on the DNS
  - 9.y-0.<label>.<label>.157c.1beb.3ea1.210.0.<label>.avts.mcafee.com.winsinage2.cba
  - 9-0.<label>.<label>.157c.1beb.3ea1.410.0.<label>.avts.mcafee.com.parkside-kamagaya.cba

- Monitoring queries leaks info about files being scanned
- Active responses tell clients if malware is detected
- Blocking responses impacts malware service
- Possible result of standards-based search list processing

McAfee GTI based namespaces

- 30 different namespaces are making McAfee queries

Example Namespaces
parkside-kamagaya
makuharibaytown-mirama3
maehara-b2
winstown-inage45
makuharibaymirama-ru
ilink-ichikawa-03
wt-inagekaigan3
w-g-c-2
yachiyo-pc12
gp-sonnou-4
w-g-c-1
Winsinage2
ilink-ichikawa-04
ver-ichi6
Toyoshiki2
ilink-ichikawa-02
takanedai-9
g-n-2
ichikawaminami-3
a-takane5

Parkside Kamagaya – Rental property in Chiba, JP (parkside-kamagaya)

Makuhari Baytown High-rise in Chiba

(makuharibaytown-mirama3 and makuharibaymirama-ru)

I-Link Ichikawa The Towers East (ilink-ichikawa-03)

So, does CBA own .cba?

CommonwealthBank

- There are no labels being queried in .cba that explicitly indicate “Commonwealth Bank of Australia” or obvious derivatives
  - Inspecting the hostnames it appears they may be operating:
    - commnet.cba
      - http://www.ours
    - commnet2.cba
    - commsec.cba
    - commsec-it.cba
    - commsee.cba

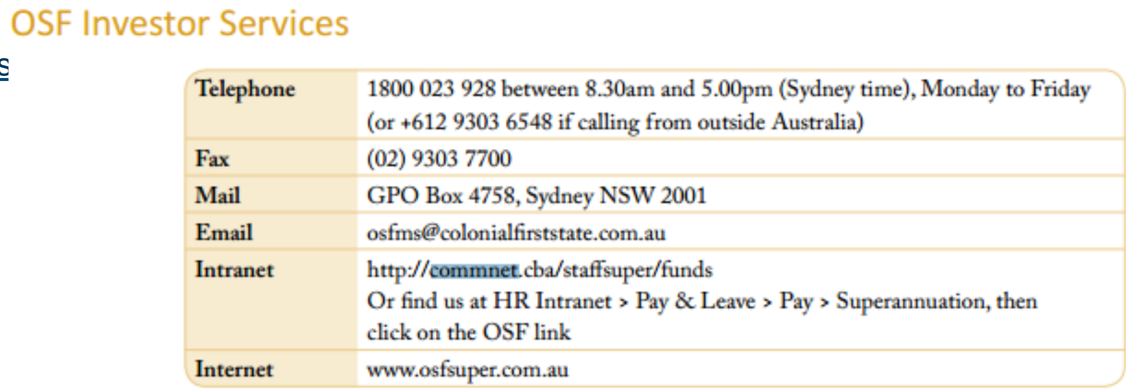

OSF Investor Services

Telephone	1800 023 928 between 8.30am and 5.00pm (Sydney time), Monday to Friday (or +612 9303 6548 if calling from outside Australia)
Fax	(02) 9303 7700
Mail	GPO Box 4758, Sydney NSW 2001
Email	osfms@colonialfirststate.com.au
Intranet	http://commnet.cba/staffsuper/funds Or find us at HR Intranet > Pay & Leave > Pay > Superannuation, then click on the OSF link
Internet	www.osfsuper.com.au

  - Some Australian queries for “.CBA”
- Liberal estimates for their traffic range from ~2-6% of .cba queries based on above namespace volumes

Initial findings

- Focused analysis of gTLD strings enables more nuanced and qualitative analysis
  - Has allowed us to discover additional service behaviors and namespace communities
- A possible relation between learned namespaces in .cba and residential and commercial entities
  - Stelmat, Makuhari Baytown High-rise, Parkside Kamagaya, I-Link Ichikawa The Towers East, etc.
- Our new systematic techniques allow us to discover namespaces in a repeatable/automated way
  - Uses protocol/service-specific techniques to learn namespaces
- Under ICANN's risk mitigation plan, the bank would appear to assume all liability for delegation of.cba

Conclusions

- Most applicants do not seem to be qualified to assess the risks of delegating their strings without visibility to root server system data and qualitative analysis; indicators simply not within their current observation space
  - Necessitates implementation of SAC045 & SAC046 recommendations regarding early warning and instrumentation across root server system
- X.509 certificates serve as an indicator of usage for a given string and vulnerabilities exist until ALL certificates expire (revocation alone is insufficient)
- DNS Service Discovery and apparent standards-based search list interactions account for a large number of the queries at the root for these and most other applied-for strings; this may pose considerable risks
  - Given types of devices that employ DNS-SD, notification and upgrade/ corrections could be costly + resource intensive and should begin immediately
- Our analysis proves the wisdom of Interisle's warning in their report's Executive Summary in that:
  - *".. additional qualitative analysis of the harms that might ensue from those collisions would be necessary to definitively establish the risk of delegating any particular string as a new TLD label, and in some cases the consequential harm might be apparent only after a new TLD label had been delegated."*
  - Furthermore, we believe much can be accomplished by ICANN directly before ANY delegation occur, thereby both minimizing risks to consumers and liability to applicants considerably

VERISIGN®

Thank You