Additional Comments on "Mitigating the Risk of DNS Namespace Collisions" Phase One Report

VeriSign, Inc.
April 21, 2014

1 Introduction

The public comment period on JAS Global Advisors' Phase One Report [1] has produced a range of comments [2], providing helpful feedback on the recommendations in the report in preparation for a full review of the name collision management framework also including the Phase Two Report, which is expected in June. (As noted in the preliminary comments that Verisign submitted on March 31 [3], the Phase One Report does not fulfill the stated objectives and intent of the name collision occurrence management framework that ICANN resolved to develop [4], hence the need for the Phase Two Report to be available before it would be appropriate to consider the framework as having been reviewed by the public.)

The comments are for the most part either supportive or complementary to the preliminary comments in [3], in some cases offering further perspective on the name collision issue and potential mitigations.

In the interest of focus, a few highlights among the comments are covered here where a clarification of Verisign's preliminary comments may be helpful to the discussion.

2 Internal Addresses vs. External Honeypots

Verisign maintains its position that directing requesters to an internal address during the controlled interruption period is preferable to an external honeypot, because as previously stated, it avoids "controlled exfiltration" where sensitive traffic from an installed system – without the advance consent of the user or system administrator – may be drawn outside the local network. This risk is acknowledged in Google Registry's comments advocating for the external honeypot [5]:

- *"Unfortunately, some protocols will send sensitive information unsolicited (e.g., login.example/login.php?user=fred and HTTP cookies). The honeypot will specifically not log this sort of information, but this doesn't change the fact that the information has been communicated over the Internet."*

It's important to keep in mind that it's not the security of the honeypot that's so much a problem (although it could be), but the potential disclosure of sensitive information over the network connection to the honeypot. The installed system won't be expecting to connect to a honeypot, it will just be

conducting what it thinks is business as usual, but to a new IP address, so there's no way to assure that the traffic will be encrypted. Although an installed system may well send traffic over unsecured networks all the time, it shouldn't be "controlled" into doing so without its consent, especially without demonstrable evidence that no lower-risk mitigation measure is available. If a system administrator actually wants to direct traffic to a honeypot, it's easy enough to do so by applying appropriate rewriting rules to the DNS responses received during the controlled interruption period, as Jeff Schmidt observed in his technical presentation on the Phase One Report ([6], slides 37-38).

The main objectives motivating an external honeypot are to make it easier to notify users and system administrators of an impending change to the DNS, and to make it easier to monitor the effectiveness of controlled interruption. Like controlled interruption itself, however, a honeypot is a major operational activity that requires careful design and implementation to reach these objectives. The outreach to users and system administrators currently in ICANN's plan is a more manageable way to achieve the first objective. The analysis of controlled interruption traffic proposed in Verisign's preliminary comments – where samples of controlled interruption traffic are provided to a forum like DNS-OARC[1] for independent evaluation by the research community (see [3], Section 5, as well as the Internet Service Providers & Connectivity Providers Constituency (ISPCP)'s second comment [7]) – is a more collaborative way to achieve the second.

Other stakeholders [9][10] have expressed reservations about the 127.0.53.53 address itself. These concerns should be evaluated further, as well as the point about the lack of an IPv6 address [11]. Verisign's position is not so much in favor of one particular address, but a preference for an internal address over an external honeypot. And in any case, a qualitative assessment of name collision risk per new gTLD and SLD, as ICANN set out to accomplish, followed by a targeted mitigation of the risk, would be much preferable to either of the choices contemplated in this discussion.

3 An Uneven Playing Field for Whom?

In a previous public comment period, Eric Osterweil summarized key differences between established and new gTLDs as they affect name collision risks [12]. Namespaces associated with established TLDs, he observed, represent "well known and measurable real estate" that system administrators can plan for. In contrast, namespaces associated with applied-for strings including new gTLDs, Osterweil continued, "inherently have no well-known policies and structure" – other than the assumption that they weren't expected to be delegated in the future foreseeable to system administrators.

Osterweil's points are important to keep in mind, because they apply just as much to one of Donuts' comments [13] in this public review period as they did to comments by the same stakeholder in the previous period [14].

A better understanding of the situation starts with clear definitions. A name collision occurs when one system assumes that a name is in one name space, another system assumes that the name is in another

[1] DNS-OARC has expressed its willingness to facilitate the analysis of mitigation measures for name collisions under "the principle that any significant changes to the DNS should be monitored via appropriate data-gathering" [8].

name space, and the two systems interact unaware of their difference in assumptions. One of the reasons they may not be aware is that the assumptions of both systems were historically the same, and then the assumptions of one of the systems changed. ICANN's Security and Stability Advisory Committee (SSAC) expresses the definition as follows in SAC062 [15]:

> *"The term "name collision" refers to the situation in which a name that is properly defined in one operational domain or naming scope may appear in another domain (in which it is also syntactically valid), where users, software, or other functions in that domain may misinterpret it as if it correctly belonged there."*

With this definition in mind, it's useful to highlight two situations that are not the same as name collisions.

First, **a change in the registrant for a domain name that already exists in the global DNS isn't a name collision**. Installed systems may get a different resource record in response to query after the change, but they've already assumed that domain names that exist in the global DNS are in an external name space; they don't need to change their assumptions when the registrant changes. To counter one specific concern raised along these lines ([16], page 10), the "retirement" of a domain name for which query traffic continues doesn't mean that a collision occurs when the domain name is registered again. The domain name was external when its registration was not renewed, and it will continue to be external when it's registered again; there's no change in assumptions about the placement of the domain name in a name space.[2]

Second, **a query that results in an NXDOMAIN response isn't a name collision.** The installed system that generated the query may have assumed that the domain name is in an internal name space, or it may have assumed that it's in an external one. It's not possible to tell directly from a query what assumptions were made, and therefore whether those assumptions would be at risk if the response were to change from NXDOMAIN to a resource record. This is why there's a need for qualitative analysis (and, in particular, the name collision management framework that ICANN resolved to develop).

So it should be clear that periodic changes in registrants as well as evidence of NXDOMAIN activity should not be confused with actual name collision risks, which depend on analysis of the queries themselves. Rather, one must look on the other side of the ecosystem, at the installed systems that interact with the global DNS and the evolution of their assumptions about internal vs. external name spaces. The difference between established gTLDs and new gTLDs becomes obvious in light of the differences in "expectations of their usage and policies" [12], as illustrated by the three examples:

- In February 2013, the CA/Browser Forum resolved that certificate authorities should no longer issue certificates bearing internal domain names [17]. This was a direct consequence of the potential that those names might collide with the new gTLDs, as SSAC had noted in SAC057 [18].

[2] Furthermore, if there's any notification needed that the status – not the namespace – of the domain name has changed, then the change in response from a resource record to NXDOMAIN during the period between retirement and re-registration surely serves as its own form of "controlled interruption."

There had been no previous need for SSAC or the CA/B Forum to sway certificate authorities away from issuing internal-name certificates on the basis that they might collide with established TLDs. Internal name spaces and the global DNS were already understood to be separate.

- In December 2013, ICANN issued a 20-page document advising IT professionals on how to mitigate name collision risks [19]. This was also a response to the name collision issue for new gTLDs. There had been no previous need for ICANN to offer such advice about established TLDs, again because of the understanding that internal name spaces were safely separate. Indeed, had there been such a need, a statement along these lines surely could have been included by SSAC in SAC045, which provided initial guidance on the name collision issue in 2010 [20]. Such a statement would likely have elicited a more immediate response than it did, given that SSAC, in this hypothetical scenario, would have been drawing attention to both potential future and current dangers. Instead, SSAC focused on the impending risks of new gTLDs, and the later ICANN guidance was issued accordingly.
- In January 2014, the DBOUND mailing list [21] was established within the IETF to discuss the challenges in managing policy information and other metadata about domain names, including, among other things, the public suffix list [22] that specifies the administrative boundaries between domain name registries and managed domains. To quote from the mailing list charter: "Most of the existing mechanisms are managed semi-manually, and there are good reasons to suppose that the limits of such management are either about to be exceeded, or already have been" – again a direct consequence of the introduction of new gTLDs, or to quote further, the fact that "[t]he DNS root is expanding rapidly."

As noted in Verisign's preliminary comments, the combinational complexity of changes in certificate practices, installed system configuration, the public suffix list and other system components to support new gTLDs, dramatically increases the risk to users and system administrators for new gTLDs compared to established TLDs, where each of these controls have been worked out over time.

The potential risks are real, as Osterweil further observed, citing a well-known recipe for mounting a man-in-the-middle attack by exploiting name collisions and internal-name certificates that was offered by an audience member at a TLD Security Forum meeting in August 2013 ([23] at 1:27:20).

If there's any unevenness in the domain name landscape, then, it's a result of the tectonic interruptions that are requiring users, system administrators, network operators, infrastructure providers and platform and application developers across the globe[3] to update their installed systems to accommodate 1400 or more new gTLDs. The parties who rely on the global DNS are the ones whose playing field is out of balance due to the largest operational change to the global DNS in its 30-year history.

[3] Including, as insightfully noted in United TLD's comments, applications that analyze data from DNS and WHOIS, and which therefore could be at risk of unexpected inconsistencies if the DNS and WHOIS data become wildly out of sync during the unusual circumstance of the controlled interruption period.

4 Wildcards and Dotless Domains

Controlled interruption, both as specified in the Phase One Report for new gTLDs that have not already been delegated, and potentially, depending on the variant, for those that have been delegated as well, is essentially an experiment in the use of wildcard records to implement policy for DNS resolution.

For new gTLDs that have not already been delegated, the Phase One Report recommends that a wildcard record be placed in the zone file during the controlled interruption period, so that queries for any SLD – with the exception of "NIC.", which is required due to other ICANN policy – return the controlled interruption IP address mentioned above. Because there are too many possible SLDs that could be queried by an installed system to enumerate them all in advance[4], the only practical way to interrupt all of them is with a wildcard record.

The situation with respect to the Phase One Report is different for new gTLDs that have already been delegated, because there, the only SLDs that are interrupted are those on the SLD block list [25][26], which is already enumerated. An explicit record could therefore be placed in the zone file for each of them.

One of the recommendations in Verisign's preliminary comments (see [3],Section 4) is that a new gTLD operator should be given the option *not* to interrupt an SLD if it has no intention to delegate the SLD, in order to minimize unnecessary harm to systems that otherwise would not be at risk of name collisions. The option took the form of two alternate approaches, both under the heading of "selective interruption": an **SLD white list** where only a defined set of SLDs that are intended to be delegated are interrupted; and an **SLD black list** where all but a defined set of SLDs that are *not* intended to be delegated are interrupted[5].

Viewed more generally, these approaches take a step of moderation toward the analysis expected in the full name collision management framework, where mitigation measures are designed based on actual risk associated with a given combination of new gTLD and SLD rather than one or perhaps two sizes fitting all.

Another alternate approach is recommended in the comments by United TLD [27], the New TLD Applicant Group (NTAG) [28] and CNNIC [29], which propose that a new gTLD operator whose new gTLD has already been delegated should be given the option to interrupt *all* SLDs except those already delegated – as opposed to just the ones on the SLD block list[6]. This expansion in the set of SLDs subject

[4] The April 2014 Domain Name Industry Brief [24] offers a somewhat related analysis of how many SLDs could exist in principle within a given TLD.

[5] Footnote 4 in [3] stated that this approach would require a modified name server, but may not have been sufficiently clear on the modification required. There is no provision in the standard zone file format for specifying that a response should be returned for all SLDs except those on a defined list, such that SLDs on the defined list get the NXDOMAIN error message – the zone file format specifies when to return a positive response, but not a negative one. The modification would specify on a per-TLD basis when to return a negative response.

[6] The CONAC comments [30] also speak to controlled interruption, but rather than proposing an alternate approach to the mitigation measure in general, the comments request an alternate classification of SLDs within

to controlled interruption is motivated by the practicalities of managing zone file records for so many SLDs outside the normal delegation process. But in a sense it also brings a moderation toward the more precise analysis expected in the full framework, given that it recognizes (though not necessarily by the commenters' intent) that domain names outside the SLD block list may also require mitigation measures.

The following table compares the four approaches – Phase One Report, white list, black list, and the United TLD / NTAG / CNNIC "alternate path to delegation (APD) wildcard" approach. It's offered to facilitate the development and review of these and similar measures, without necessarily endorsing controlled interruption itself. As noted above, specific mitigation measures based on qualitative assessment of name collision risk are preferable to any of these.

Status of new gTLD and SLD				Action of controlled interruption approaches			
New gTLD Already Delegated under APD?	**SLD on Block List?**	**SLD Already Delegated?**	**SLD Intended to Be Available for Delegation?**	**JAS Phase One Report**	**White-List**	**Black-List**	**APD Wildcard**
Yes	Yes	No	Yes	Per-SLD	Per-SLD	Per-SLD	Wildcard
"	"	"	No	Per-SLD	None	None	Wildcard
"	No	Yes	Yes	Delegation	Delegation	Delegation	Delegation
"	"	No	Yes	None	None	None	Wildcard
"	"	"	No	None	None	None	Wildcard
No	n/a	No	Yes	Wildcard	Per-SLD	Wildcard	Wildcard
"	"	"	No	Wildcard	None	NXDOMAIN	Wildcard

In the table, the first four columns indicate the status of the new gTLD / SLD combination. Only seven status arrangements are possible, because SLDs on the block list ("Yes" in column 2) would not yet be delegated ("No" in column 3), SLDs already delegated ("Yes" in column 3) are intended to be available for delegation ("Yes" in column 4), and for new gTLDs that have not already been delegated ("No" in column 1), the SLD block list is not relevant in any of the proposed approaches ("n/a" in column 2) and SLDs would not yet be delegated ("No" in column 3).

---

specific applied-for new gTLDs, 政务 and .公益, such that interruption would not be required. This again points to the need for the full name collision management framework.

Akram Atallah, president of ICANN's Global Domains Division, affirmed the need in a response to the comments on the same issue by CONAC at the ICANN Public Forum in Singapore ([31], page 67): "The Board has not made -- the NGPC has not made a decision yet on the collision's *(sic)* framework. Once that decision is made, if you want to have an exception to that, you can submit to us a request and we will discuss that with you based on your specific case. But right now there is -- you know, we cannot -- we don't have the framework to move forward yet." Atallah did not elaborate on the proposed "exception" process, but given that the framework as commissioned by ICANN was intended to enable ICANN to make risk assessments on a per-new-gTLD/per-SLD basis according to a defined methodology, it's not clear why an exception would be required.

The last four columns indicate whether controlled interruption is applied for each status combination, and if so, how:

- **Per-SLD** means that a separate domain name record mapping to the controlled interruption IP address is included for the SLD in the zone file for the new gTLD.
- **Wildcard** means that a wildcard record mapping to the controlled interruption IP address is included in the zone file for the new gTLD, covering all SLDs with the given status[7].
- **Delegation** means that a normal domain name record is included for the SLD.
- **None** means that no domain name records are included, so the NXDOMAIN error message is returned by default
- **NXDOMAIN** refers to the modification discussed above where the name server returns the NXDOMAIN name error message explicitly (it couldn't be returned by default, because of the wildcard covering other SLDs in the zone file).

The complexity of the approaches discussed here should serve as a reminder why the Internet technical community has recommended against the use of wildcards [32]. They're powerful constructions, but they're hard to use correctly, can easily go wrong. Indeed, wildcards head in the same, potentially insecure and unstable direction as dotless domains [33][34], which are disallowed by ICANN [35] (see also the discussion under the fourth comment of [4]). To quote from the Internet Architecture Board's statement *Dotless Domains Considered Harmful* [34]:

> *1.The IAB strongly recommends against considering, implementing, or deploying dotless domains.*
>
> *2.The IAB believes that dotless domains are inherently harmful to Internet security.*
>
> *3.Applications and platforms that apply a suffix search list to a single-label name are in conformance with IETF standards track RFCs. Furthermore, applications and platforms that do not query DNS for a TLD are in conformance with IETF standards track recommendations intended to minimize security vulnerabilities and reduce load on the root servers.*

ICANN's decision not to allow dotless domains [35] becomes all the more important to maintain when the risk of dotless domains to Internet security and stability is considered in combination with the other risks enumerated in Section 3, which is essential in the context of name collision management. Indeed, "controlled interruption" sets a potentially risky precedent of relying on wildcards (and by extension dotless domains) to implement Internet policy (which is why even the Phase One Report acknowledges in its Recommendation 8 that the precedent would require temporary relief from current ICANN requirements). Therefore, if it's necessary to use wildcards for name collision mitigation, that practice

[7] In the case of new gTLDs that haven't already been delegated, this is with the exception of "NIC.", as noted above. For the APD wildcard option, it's not clear in the comments whether delegation would be allowed to continue for SLDs not on the block list while controlled interruption is performed, but if so, the wildcard would only have its effect on a given SLD until delegation occurs. CNNIC proposes not to delegate any SLDs until after the controlled interruption period, making the APD wildcard option for new gTLDs that have already been delegated essentially the same as the JAS Phase One Report approach for new gTLDs that haven't.

should be a well-managed and widely monitored interruption to the norm of staying with explicit per domain-name records.

5 How Long to Interrupt?

The length of the controlled interruption period appears to be among the more controversial of the topics. Some say it's too long, at least one says it's too short, and it's not clear that anyone thinks it's just right. Given that there is no prior operational experience with controlled interruption at this scale (a point also made in Verisign's preliminary comments ([3], Section 5)), it's not surprising that it would be unclear how long the period should be.

There does appear to be general consensus that the controlled interruption period, if the mitigation measure is adopted, should begin as soon as the registry agreement is executed for a new gTLD, which would allow, quoting NTAG's comments [28], "for the maximum opportunity for third parties to assess unlikely leakages while minimizing the disruption of the Registry's business model."

If controlled interruption is adopted, the only way to get a better understanding of the appropriate period is by qualitative analysis of the effectiveness of the mitigation measure in practice. This is consistent with the idea that third parties – including the research community – should be engaged in risk assessment. Such a feedback loop will also provide a platform for identifying and testing further optimizations such as the promising one suggested by Google Registry (see [5], Section 4) where controlled interruption is done in intervals, rather than continuously, with time allowed for remediation in between the intervals and afterward. Indeed, a feedback loop can enable innovations both in controlled interruption method and in other mitigation measures, supporting NTAG's fifth comment recommending that alternate measures be allowed, with the motivation, expressed in NTAG's seventh comment, "that [registries] can select what solutions fits their technical systems and business needs" [28]. The name collision management framework, when fully made available as ICANN intended, should give the Internet community this kind of confident flexibility.

References

[1] *Mitigating the Risk of DNS Namespace Collisions: Phase One Report.* JAS Global Advisors, February 24, 2014. http://www.icann.org/en/about/staff/security/ssr/name-collision-mitigation-26feb14-en.pdf

[2] *Mitigating the Risk of DNS Namespace Collisions (public comments)*. ICANN, http://www.icann.org/en/news/public-comment/name-collision-26feb14-en.htm

[3] Burt Kaliski. *Verisign preliminary comments on "Mitigating the Risk of DNS Namespace Collisions" Phase One Report.* comments-name-collision-26feb14 discussion thread, March 31, 2014. http://forum.icann.org/lists/comments-name-collision-26feb14/msg00010.html

[4] *NGPC Resolution for Addressing the Consequences of Name Collisions.* ICANN, October 8, 2013. http://www.icann.org/en/news/announcements/announcement-08oct13-en.htm

[5] Sarah Falvey. *Google Registry's Public Comment on Mitigating the Risk of DNS Namespace Collisions.* comments-name-collision-26feb14 discussion thread, March 31, 2014. http://forum.icann.org/lists/comments-name-collision-26feb14/msg00012.html

[6] Jeff Schmidt. *Mitigating the Risk of DNS Name Space Collisions*. Presented at Workshop and Prize on Root Causes and Mitigation of Name Collisions (WPNC '14), London, United Kingdom, March 8-10, 2014. http://namecollisions.net/downloads/wpnc14_slides_jas_framework_session.pdf

[7] Christian Dawson. *ISPCP Public Comments to JAS Report.* comments-name-collision-26feb14 discussion thread, April 1, 2014. http://forum.icann.org/lists/comments-name-collision-26feb14/msg00014.html

[8] Keith Mitchell. *DNS-OARC Comments on JAS Name Collisions Report.* Comments-name-collison-26feb14 discussion thread, April 20, 2014. http://forum.icann.org/lists/comments-name-collision-26feb14/msg00021.html

[9] Aaron Beck. *Returning 127.0.53.53 is potentially dangerous, and my suggestions for an alternative.* comments-name-collision-26feb14 discussion thread, February 28, 2014. http://forum.icann.org/lists/comments-name-collision-26feb14/msg00000.html

[10] Jason A Fesler. *Re: Returning 127.0.53.53 is potentially dangerous, and my suggestions for an alternative.* comments-name-collision-26feb14 discussion thread, March 28, 2014. http://forum.icann.org/lists/comments-name-collision-26feb14/

[11] Martin J. Levy. *If the IPv4 answer is 127.0.53.53 then what is the IPv6 answer? - While trivial; the lack of IPv6 support is something to consider.* comments-name-collision-26feb14 discussion thread, March 28, 2014. http://forum.icann.org/lists/comments-name-collision-26feb14/msg00003.html

[12] Eric Osterweil. *NXDomain responses under existent TLDs are _not_ the same as NXDomain responses under applied-for strings.* comments-name-collision-05aug13 discussion thread, September 11, 2013. http://forum.icann.org/lists/comments-name-collision-05aug13/msg00038.html

[13] Mason Cole. *Donuts Comment on JAS Name Collision Report.* comments-name-collision-26feb14 discussion thread, March 31, 2014. http://forum.icann.org/lists/comments-name-collision-26feb14/msg00009.html

[14] Paul Stahura. *Donuts' Comments.* comments-name-collision-05aug13 discussion thread, August 27, 2013. http://forum.icann.org/lists/comments-name-collision-05aug13/msg00030.html

[15] *SAC062: SSAC Advisory Concerning the Mitigation of Name Collision Risk.* ICANN Security and Stability Advisory Committee, November 7, 2013. http://www.icann.org/en/groups/ssac/documents/sac-062-en.pdf

[16] *Name Collision Mitigation | Transcript.* ICANN, March 24, 2014. http://singapore49.icann.org/en/schedule/mon-name-collision/transcript-name-collision-24mar14-en.pdf

[17] *Ballot 96 – Wildcard Certificates and New gTLDs (Passed).* CA/Browser Forum, February 20, 2013. https://cabforum.org/2013/02/20/ballot-96-wildcard-certificates-and-new-gtlds/

[18] *SAC057: SSAC Advisory on Internal Name Certificates.* ICANN Security and Stability Advisory Committee, March 15, 2013. http://www.icann.org/en/groups/ssac/documents/sac-057-en.pdf

[19]*Guide to Name Collision Identification and Mitigation for IT Professionals.* ICANN, December 5, 2013. https://www.icann.org/en/about/staff/security/ssr/name-collision-mitigation-05dec13-en.pdf

[20]*SAC045: Invalid Top Level Domain Queries at the Root Level of the Domain Name System.* ICANN Security and Stability Advisory Committee, November 15, 2010. http://www.icann.org/en/groups/ssac/documents/sac-045-en.pdf

[21]*Dbound -- DNS tree bounds.* Mailing list, https://www.ietf.org/mailman/listinfo/dbound. Accessed April 16, 2014.

[22]*Public Suffix List.* http://publicsuffix.org/. Accessed March 28, 2014.

[23]*TLD Security Forum part 1.* Video transcript, August 22, 2013. http://www.youtube.com/watch?v=XRvk6ySPwTc&feature=youtu.be

[24]*The Domain Name Industry Brief.* Volume 11, Issue 1. VeriSign, Inc., April 2014. http://www.verisigninc.com/assets/domain-name-report-april2014.pdf

[25]*NGPC Resolution for Addressing the Consequences of Name Collisions.* ICANN, October 8, 2013. http://www.icann.org/en/news/announcements/announcement-08oct13-en.htm

[26]*Reports for Alternate Path to Delegation Published.* ICANN, November 17, 2013. http://newgtlds.icann.org/en/announcements-an6d-media/announcement-2-17nov13-en

[27]Statton Hammock. *United TLD Comment on Mitigating the Risk of DNS Namespace Collisions*. comments-name-collision-26feb14 discussion thread, March 31, 2014. http://forum.icann.org/lists/comments-name-collision-26feb14/msg00007.html

[28]Rubens Kuhl. *NTAG comments on JAS Namespace Collision Report.* comments-name-collision-26feb14 discussion thread, March 31, 2014. http://forum.icann.org/lists/comments-name-collision-26feb14/msg00008.html

[29]Yi (Daisy) Ding. *CNNIC comments on Name Collision.* comments-name-collision-26feb14 discussion thread, April 1, 2014. http://forum.icann.org/lists/comments-name-collision-26feb14/msg00015.html

[30]Limei Liu. *CONAC's Comments on JAS's Name Collision Mitigating Plan.* comments-name-collision-26feb14 discussion thread, March 31, 2014. http://forum.icann.org/lists/comments-name-collision-26feb14/msg00004.html

[31]*ICANN Public Forum | Transcript.* ICANN, March 27, 2014. http://singapore49.icann.org/en/schedule/thu-public-forum/transcript-public-forum-27mar14-en

[32]*IAB Commentary: Architectural Concerns on the Use of DNS Wildcards.* Internet Architecture Board, September 19, 2003. http://www.iab.org/documents/correspondence-reports-documents/docs2003/2003-09-20-dns-wildcards/

[33]*SAC053: SSAC Report on Dotless Domains.* ICANN Security and Stability Advisory Committee, February 23, 2012. http://www.icann.org/en/groups/ssac/documents/sac-053-en.pdf

[34]*IAB Statement: Dotless Domains Considered Harmful.* Internet Architecture Board, 2013. http://www.iab.org/documents/correspondence-reports-documents/2013-2/iab-statement-dotless-domains-considered-harmful/

[35]*New gTLD Dotless Domain Names Prohibited.* ICANN, August 30, 2013. http://www.icann.org/en/news/announcements/announcement-30aug13-en.htm

