Preliminary Comments on “Mitigating the Risk of DNS Namespace Collisions” Phase One Report

VeriSign, Inc.
March 31, 2014

1 Introduction

On October 8, 2013, the New gTLD Program Committee (NGPC) of the ICANN Board resolved [1] to develop a name collision occurrence management framework to address the potential risks of name collisions related to the introduction of new generic top-level domains (gTLDs), a subject that had been highlighted in a number of reports earlier in the year [2][3][4][5][6][7][8][9].

The NGPC-approved *New gTLD Collision Occurrence Management Plan* [10] set forth the following objectives:

- *“ICANN will commission a study to develop a name collision occurrence management framework. The framework will include appropriate parameters and processes to assess both probability and severity of impact resulting from name collision occurrences. Examples of the parameters include number of DNS requests, type of DNS requests, type of queries, diversity of query source and appearances in internal name certificates.*
- *“The framework will specify a set of name collision occurrence assessments and corresponding mitigation measures if any, that ICANN or TLD applicants may need to implement per second level domain name (SLD) seen in the DITL and other relevant dataset (e.g., information from Certificate Authorities regarding the issuance of internal name certificates) . The proposed name collision management framework will be made available for public comment”*

The framework was intended to be applied to each applied-for new gTLD to produce a name collision occurrence assessment and suggested mitigation measures for specific second-level domains (SLD) within the new gTLD. Examples mitigation measures given included blocking the SLD indefinitely; blocking it temporarily; conducting a trial delegation; and if only one entity is impacted by collisions with the SLD, making the SLD available to that entity.

Five weeks later, on November 11, ICANN announced that it had awarded a contract for the study to JAS Global Advisors [11][12].

JAS Global Advisors published the first of two parts of the report, the “Phase One Report,” on February 26, 2014 [13][14]. ICANN has asked for public comment on the first part with comments due March 31 and replies to comments due April 21 [15]. The second part, the “Phase Two Report,” has not yet been published because, according to the first part, it discloses a security vulnerability. Per ICANN policy, parties impacted by a security vulnerability are notified and given a certain amount of time to remediate the vulnerability prior to public disclosure [16]. Based on the report author's comments at ICANN 49 in Singapore on March 24, the Phase Two Report may not be available until June [17].

The purpose of this note is to provide preliminary comments on the Phase One Report. The comments are subject to change pending the completion of the unpublished Phase Two Report. In the interest of being constructive, various specific comments are offered on the suggested mitigation measures in the Phase One Report. The comments should not be taken as an endorsement of the measures, which would be premature. However, they are nevertheless intended to be helpful should ICANN choose to move forward with the measures.

2 Name Collision Framework Not Yet Provided

Comparing the Phase One Report to the stated objectives and intent of the framework, it is clear that there is as yet no framework to comment on. No "parameters and processes to assess both probability and severity of impact" have been presented, nor any "name collision occurrence assessments." The statement of work reinforces the expectations in calling for these deliverables [12]:

- *"1.1 Develop a Risk Assessment Model*

    1.1.1 Impact of malware/adware/clickfraud tools

    1.1.2 Analysis of Collisions in previous TLD delegations

    1.1.3 Analysis of Collisions in existing TLDs

    1.1.4 Monte Carlo Analysis

    1.1.5 Survey Instruments

    1.1.6 Develop a Taxonomy of Queries

- *"1.2 Options to manage risks"*

The Phase One Report does not detail any of the analysis mentioned under item 1.1.

According to the workflow in the October 2013 plan, the study was to develop a framework; then ICANN was to apply the framework to a list of SLDs for each applied-for new gTLD. Instead, the Phase One Report has gone directly to application, arriving at the following suggested mitigation measures:

- If the new gTLD is .CORP, .HOME, or .MAIL, then the entire new gTLD must be blocked indefinitely (indeed, "permanently")
- If the new gTLD hasn't already been delegated, then the entire new gTLD must undergo a technique called **controlled interruption**

- If the new gTLD has already been delegated, i.e., via the "alternate path" with an SLD block list [2][18], then the SLDs on the block list must undergo controlled interruption

None of this is "per second level domain name (SLD)," varying based on any qualitative assessment of the risk of any specific SLD. The only distinction made is whether the SLD is on the block list. Thus, the reviewer must assume the following assessment has been reached in the unpublished Phase Two Report: that regardless of the actual risk associated with a given new gTLD and SLD, the best way to mitigate the risk[1], in terms of actions under ICANN's control, consists of the three mitigation measures proposed in the Phase One Report. If the assessment were otherwise, then the Phase One Report would not make the suggestions that it does.

Until the Phase Two Report is published, it is not possible to verify if this is indeed the assessment made in the report, and even if it is, that the analysis leading to the assessment is correct. However, it is possible to provide some initial evidence that the assessment may not be correct. This can be done by offering two counterexamples where the controlled interruption technique may not in fact mitigate risk.

3 Controlled Interruption Is Untested, May Not Be Effective

The Phase One Report offers an intriguing approach to mitigating name collisions, framed in terms of managing the transition from one set of system conventions to another. In his first public presentation on the report at WPNC '14 in London on March 10 [19], report author Jeff Schmidt drew an analogy between name collision risk mitigation and historic transitions in the phone numbering and postal code systems. (A brief mention of the analogy was included in his March 24 presentation at ICANN 49 in Singapore [20], where he referred again to the "Anti-Digit-Dialing League" that opposed a 1960s phone numbering transition.)

While there are several problems with this analogy (for example, ICANN doesn't have a provider/subscriber relationship with users, and the transition is not from one ICANN-managed namespace to another but from a non-ICANN-managed space to an ICANN-managed one), it does offer a helpful way to describe the proposed controlled interruption technique, as well as its limitations.

In his WPNC '14 presentation, Schmidt described prior transitions as generally consisting of three phases:

- **advance notification**, where users are informed that the set of system conventions will be changed in the near future, and that users continuing under the old set may be at risk of different system behavior after the change

[1] A reasonable argument can be made that a mitigation measure cannot be evaluated in isolation on a per-gTLD and per-SLD basis, because measures are not applied in isolation. Rather, a collection of measures must be evaluated together – perhaps at the per-gTLD level or even across all new gTLDs. Thus the question of interest to the reviewer is not whether the suggested measures are best for each and every gTLD/SLD combination individually, but whether they are best collectively for the present situation. Still, even that assessment requires technical justification.

- **grace period**, where users continuing under the old conventions receive an error message or "negative acknowledgement" indicating that the old practice is no longer supported and reminding them of the correct new conventions, but where the new conventions have not yet been activated
- **activation**, where the new conventions are in effect and where users continuing under the old convention are no longer warned, but instead are subject to the predicted risks

A recurring example in Schmidt's presentation concerns transitions of the "regional" or "area code" portion of phone numbers. From time to time, phone systems may change the regional portion of some of their subscribers' phone numbers, in order to make room for new subscribers or phone lines at the old numbers. The grace period would make a defined separation between the time when phone numbers are associated with their original subscribers, and when they could potentially be assigned to new subscribers. In the interim, the numbers would be assigned to neither. Instead, subscribers who dialed the old numbers would get an error message indicating that they must dial the new numbers instead.

The introduction of new gTLDs has some similarities. Here, at a certain time, ICANN is adding a new gTLD to the global DNS root. In the past, domain names in the new gTLD didn't resolve in the global DNS; instead, they generated an NXDOMAIN response. In the future, such domain names may resolve to an IP address. Installed systems in many cases have been relying on the NXDOMAIN response under system conventions where the new gTLD is assumed not to be in the global DNS, and such systems could be at risk due to the change in conventions. The grace period moderates the risk by, it is hoped, providing a "negative acknowledgment" that the response is no longer NXDOMAIN, before the response directs the requester to an IP address controlled by someone else. The controlled interruption technique thus signals the change.

It should be noted that a technique for notifying users that a change to the DNS is about to occur by returning a novel internal IP address is unprecedented, a point that is elaborated further in Section 4. Putting this concern aside for the moment, however, there are at least two scenarios where it appears plausible that users and system administrators might not actually get notified that a change is forthcoming. In both scenarios, further justification is called for in order to draw the conclusion that the Phase One Report apparently makes, which is either that the frequency of occurrence of these scenarios is not significant, or that the controlled interruption technique would nevertheless mitigate the risk, the limitations notwithstanding.

3.1 Scenario 1: SLD block lists

The first example has to do with SLD block lists. Recall that ICANN established SLD blocking based on DNS-OARC's Day-in-the-Life (DITL) data as an alternate path for mitigating name collision risks [2]. As previously observed, the DITL data is not statistically valid for determining which queries from installed systems may be at risk of a name collision when a new gTLD is delegated [21]. Matthew Thomas, Yannis Labrou and Andrew Simpson summarize the research findings supporting this observation in their WPNC '14 paper [22]. As they document, the set of SLDs queried by installed systems continues to evolve over time; a 48-hour snapshot every year cannot capture all of them.

According to Recommendation 7 of the Phase One Report, if a new gTLD has already been delegated – meaning that is in production and is already delegating SLDs, except for those on the block list – then, during a 120-day controlled interruption period, the new gTLD must return a designated controlled interruption IP address for queries for SLDs on the block list. The intent is that installed systems that are querying for these SLDs will receive the controlled interruption IP address instead of NXDOMAIN, attempt to connect to the IP address rather than looking up a different domain name, and then "break," hopefully gracefully. A user or system administrator would then try to diagnose the break, notice the controlled interruption IP address, and have an initial clue to what's going on.

This approach may or may not work for installed systems that query for domain names whose SLDs are on the block list. However, it certainly won't work for installed systems that don't query for such SLDs at all. As Thomas *et al.* report, some of those queries could be at risk of name collisions as well. Thus, in order for the controlled interruption technique to be effective, any system that generates at-risk queries whose SLDs aren't on the block list, must also generate at-risk queries whose SLDs are on the list. Otherwise the system wouldn't receive a controlled interruption response, and the user or system administrator wouldn't know that anything was changing (at least as a result of this mitigation measure). Without the qualitative analysis promised in the framework, it's hard to know whether systems meet this requirement, and thus it's not possible to conclude that the controlled interruption technique would be effective for them.

Applying the analogy, this would be like the phone system returning the error message only for a fraction of the phone numbers that are about to be changed – and letting all the others go through without warning right up to the last minute. Hopefully no caller always misses the warning that the old numbers might be reassigned to new subscribers – or sometimes does get the warning, but doesn't make the connection that it applies to other old numbers as well.

3.2 Scenario 2: WPAD protocol

The second example has to do with whether the controlled interruption response actually "breaks" the installed system in a way that is visible to a user or application, so that the user or system administrator will in fact be motivated to take action.

The rationale given for returning an IP address that won't respond to further protocol interactions correctly is that something will "break," forcing user or system administrator attention. Because DNS errors are a known source of breakage, a user or system administrator, when encountering a break, would typically search DNS logs or try further queries, and, it is hoped, eventually discover evidence such as the unusual 127.0.53.53 address.

This all assumes, however, that something "breaks." But it's not necessarily the case that a controlled interruption response will result in an immediate, user- or application-visible error. If the user or system administrator isn't expecting anything to break (and this is one reason that outreach is so important), the system administrator can't be expected to be looking proactively for the controlled interruption response. So unless the controlled interruption response affects user- or application-visible behavior, the change to the DNS may go unnoticed, perhaps even for the full controlled interruption period.

In another paper at WPNC '14 [23], Andy Simpson describes precisely such a situation where the controlled interruption technique may go undetected: the Web Proxy Auto-Discovery protocol (WPAD) [24]. In WPAD, a non-standard but nearly ubiquitous protocol, an application looks for a web proxy by attempting to download a file from a web server at a domain names of the form WPAD.<suffix>, where <suffix> is selected from a domain name search list. If the domain name doesn't resolve – or if it does resolve to an IP address, but there's no web server at that IP address or no file of the appropriate form -- the protocol just goes on and tries another suffix, and eventually stops looking for a web proxy entirely. In any of these cases, no error message is reported to the user.

Based on analyzing DITL data, Simpson shows significant evidence of queries for domain names of this form where <suffix> ends with an applied-for new gTLD. This suggests that applications in installed systems are looking for web proxy configuration files at domain names that could collide with those that may be delegated in new gTLDs. The recommended version of the WPAD protocol constructs the search list directly from the fully qualified domain name of the host computer that is running the protocol, so it would only risk a potential name collision if the host name itself ends with a new gTLD. However, implementations of WPAD vary widely and it is possible that other search list techniques are employed in practice. In particular, the prevalence of WPAD queries in the DITL data for domain names that end with applied-for new gTLDs suggests that many installed systems currently query for web proxies using search lists that contain new gTLDs.

Because of the permissive nature of the WPAD protocol, a controlled interruption response will have the same effect in terms of the flow of the protocol as an NXDOMAIN response: there won't be a web server at the controlled interruption IP address, so the protocol will just go on to the next suffix in the search list, without reporting an error. The installed system will never actually be "interrupted." The user or system administrator may thus have a false sense of security after the controlled interruption period is complete, despite the fact that the installed system remains at risk if one of the queried domain names is subsequently delegated and a web proxy is set up at that domain name.

In order for the controlled interruption technique to be effective in this second scenario, any system that generates WPAD queries that end with a new gTLD must also generate other queries where the controlled interruption response will directly "break" something. Again, without the promised qualitative analysis, it's hard to know where systems stand in terms of the "breakage" that could occur from a controlled interruption. (Also note also that WPAD is just one example of a "hidden" service discovery protocol that relies on the DNS. Others of these, such as ISATAP [25], may be subject to the risk described here as well.)

Returning to the phone system analogy, readers may recall the 1990s-era automated dialers that tried several phone numbers in succession in order to connect to an online service. If one of these couldn't connect – either because of a busy signal, or, for the present analogy, because of a negative acknowledgement indicating that the phone number was changing – then the dialer would just skip the number and move on to the next one. Other than a delay, the user wouldn't directly know that the number was being taken out of service. If the phone number were assigned to a new subscriber at some later point, the user might then inadvertently be connected to the new subscriber without actually having been explicitly notified of this risk beforehand.[2]

4 Controlled Interruption May Break Systems that Are Not at Risk

During the controlled interruption period for a new gTLD, the response of the global DNS to queries involving the new gTLD will change as follows:

- If the new gTLD hasn't already been delegated, then the then the response to queries with any domain name under the new gTLD will change from NXDOMAIN to the controlled interruption IP address
- If the new gTLD has already been delegated, then responses to queries whose domain name includes an SLDs on the block list for that new gTLD will change to the controlled interruption IP address

This is a broad change, affecting *every* SLD in the first case and every SLD on the block list in the second.

The intent of controlled interruption is to notify users and system administrators that a change to the DNS is about to occur. However, the actual change that is about to occur is not that every possible SLD will be delegated, nor even that every SLD on the block list will necessarily be, but rather that *some* SLDs are going to be delegated. This could be a small number or a large number, but in general it won't involve every possibility.

It is important to note that controlled interruption is, by definition, a kind of name collision. Viewed in terms of the effect on a single SLD, a controlled interruption is arguably the least impactful kind of name collision, because it directs a client just to connect to itself (via a loopback address), not to another server, either internal or external to the client's network. Viewed in terms of the effect of the entire new gTLD, however, controlled interruption as proposed is arguably the *most* impactful kind of name collision, because it affects every SLD under the new gTLD (or if the new gTLD has already been delegated, every SLD on the block list).

The Phase One Report argues that this broad and shallow profile of controlled interruption is the right balance, because it ensures that a large fraction of installed systems will be notified, with a small impact on each one. But still, consider the "small" impact: protocols may break, causing some harm, including a loss of availability, one of the three pillars of information security.

Considering the implied principle that a small harm can be justified if it avoids a greater harm (such as leakage of information to or exploitation by an external system), one must assess whether all the small

[2] Of course, readers who recall 1990s-era automated dialers may also remember listening to the phone connecting with its distinctive modem handshake, so may well have heard these error messages. Moreover, the analogy breaks down in that the online service itself maintained the list of phone numbers and could update them before they were at risk of being reassigned. If only ICANN could update the "phone numbers" – the internal domain names – in installed systems to mitigate name collision risks!

harm is necessary.[3] In particular, consider an installed system whose queries within a new gTLD involving a known set of SLDs. If it is also known that none of those SLDs will be delegated, then the installed system is not at risk of a name collision (as long as the commitment of non-delegation is maintained). Accordingly, the cure is worse than the (non-)disease for that installed system. Controlled interruption would unnecessarily harm the installed system in the present, even though there's no risk of leakage or exploitation due to name collisions for that system in the future, at least with respect to this particular new gTLD.

There is therefore a reasonable case to be made, at least for some new gTLDs and SLDs, that the controlled interruption should be done more selectively.

For new gTLDs that have already been delegated, where the controlled interruption as currently specified is limited to the SLD block list, it may be reasonable for the new gTLD operator to apply controlled interruption only to a defined subset of the block list, not the entire list. Only the SLDs in the subset would then be eligible to be delegated (in addition to the ones that are not on the block list). The new gTLD operator would commit to continue to block, i.e., not to delegate, any SLDs on the block list on which controlled interruption is not performed**.**

For new gTLDs that have not already been delegated, where the controlled interruption would as currently specified apply to all SLDs, there are two options. The first option is, as in the previous case, that the new gTLD operator applies controlled interruption *only* to a defined subset of all SLDs – in effect, an **SLD white list**. As in the previous case, only those SLDs would be eligible to be delegated after the controlled interruption period, and the rest would have to be blocked. The second option is that the new gTLD operator applies controlled interruption to all SLDs *except* for a defined subset[4] -- in effect, and **SLD black list**. The new gTLD operator would commit to continue to block any SLDs on this "exclusion" list after the controlled interruption period, but could delegate everything else.

These options may give a new gTLD operator (or ICANN, as the party specifying the controls), more flexibility and a better balance in defining the profile of the controlled interruption. The intent remains to ensure that a large fraction of installed systems will be notified, with a small impact on each one, but with the further improvement that if an installed system will not be at risk of a name collision, then it is not at risk of harm from a controlled interruption either.

Determining when "selective interruption" is appropriate requires careful qualitative analysis. The typical case where selective interruption may be appropriate is one where the operator of the new gTLD also operates certain installed systems that generate queries for the new gTLDs, where the queries all

[3] The Phase One Report quotes a similar adage in reminding that one must be careful about scenarios where the "cure is worse than the disease."

[4] This would require a modified name server, because 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, i.e., no "wildcard-with-exclusions" option.

involve SLDs in a known set.[5] The operator of the new gTLD would commit not to delegate those SLDs, and therefore has a reasonable argument for not interrupting any of the installed systems.

It is also important to note some secondary risks with the selective interruption proposal. First, by not interrupting certain SLDs, it is possible that other installed systems that also query these SLDs will not be notified during the controlled interruption period. This is another expression of the concern noted in Section 3.1. Second, the ongoing blocking of an SLD on the basis that it is in use as an internal name validates a mixing of worlds where some names within a new gTLD are explicitly understood to be internal, and others to be external. Unless this mix is carefully managed, it could, over time, encourage bad DNS practices in other installed systems. Therefore, although there are reasonable arguments for the practice, it should be handled with care.

5 Risk Management Requires Feedback

An essential element of any risk management process is a feedback mechanism that provides evidence of whether, in fact, the risks of concern have actually been mitigated. The Phase One Report does propose a feedback mechanism, but it's only to confirm that the new gTLD operator has implemented the controlled interruption technique correctly. It does not confirm that the mitigation measure has its intended effect.

The ISO 31000 family of standards [26] provides a general model for risk management in enterprises and other organizations (see also [27] for an industry perspective). The model is based on a continuous feedback loop. As described in ICANN's DNS Risk Management Map [28], which is adapted from ISO 31000, the feedback loop is initiated with a **mandate and commitment** by an organization to manage risk. Based on this mandate and commitment, the organization then follows a repeated cycle of these four steps:

1. **Design a framework**, consisting of one or more **risk management processes**
2. **Implement the framework** according to the design
3. **Monitor and review the framework**
4. **Improve the framework**

A risk management process according to ISO 31000 consists of three phases:

1. **Establishing the context** for assessing and mitigating a particular set of risks
2. **Risk assessment —** including **risk identification**, **risk analysis**, and **risk evaluation**
3. **Risk treatment**

[5] Note that this case is the *opposite* of the one covered in the CBA study [28] where a different party operated some of the installed systems that generated queries for the new gTLD. Not delegating an SLD on the basis that it's already in use internally by one's own systems is a "failsafe" option, because the blocking will avoid unintended consequences on others' systems. However, delegating an SLD on the basis on the basis that it's *only* in use by one's own system requires additional assurance that it's not also in use in others' (or at least, a way to mitigate the risk that it may be).

The phases are supported by two feedback mechanisms: the **communication and consultation** with stakeholders; and **monitoring and review** of the risk management process to ensure that it meets its objectives.

Because it addresses a particular set of risks, name collisions, the name collision occurrence management "framework" that ICANN resolved to develop in October 2013 would, under ISO 31000, be part of a risk management "process." The realization of an ISO 31000 risk management framework in terms of ICANN's more general activities would be its DNS risk management framework [27]. Approved for implementation by the ICANN Board on November 21, 2013, the framework recommends that on a quarterly basis, ICANN should:

- *Provide evidence that risk treatments are successful in reducing / managing DNS risk levels, and that emerging risks are identified as soon as possible ([28], p. 28)*

ICANN's resolution in October 2013 may be considered a decision to design and implement a risk management process for name collisions. The design of the process was then assigned to JAS Global Advisors. With the ISO 31000 model as a guide, it is clear that some, but not all, of the elements of the risk management process have been defined. ICANN's October 2013 resolution and the body of work that preceded it set the context, and the Phase One Report has further elaborated it. The risk assessment presumably will be reported, at least in part, in the unpublished Phase Two Report. Risk treatments are suggested in The Phase One Report. The first feedback mechanism, communication and consultation are achieved, in principle, through the public review of the reports as well as the outreach to affected parties.

The second feedback mechanism, monitoring and review, remains incomplete. Although Recommendation 9 specifies that ICANN should "monitor the implementation of controlled interruption by each registry," the intent of the monitoring is not to ensure that the risk management *process* meets its objectives, but rather to ensure that the risk management *treatment* – controlled interruption – is implemented correctly. That aspect of monitoring is necessary, but it is not sufficient, because of the uncertainty of whether the treatment, even if implemented correctly, in fact will mitigate risk. This is the reason for the feedback loop. The monitoring and review feedback is needed not only to confirm whether the application of the risk treatment in a given case has been effective, but also to determine, based on evidence from a body of applications, whether the risk treatment itself – as well the context establishment and risk assessment that preceded it – needs to be revised.

The purpose of the still-unpublished framework was, to quote again, "to assess both probability and severity of impact resulting from name collision occurrences." If the controlled interruption technique is indeed effective, then the combination of probability and severity of impact should demonstrably decrease over the course of the interruption period as users and system administrators are notified and remediate their systems. (An example of the "risk reduction curve" one would be looking for and what a good outcome looks like, consider the reduction of potentially at-risk queries for the .CBA string, following notification of the network operator generating the queries [29].) It should be possible for a new gTLD operator and researchers, using similar techniques as developed for the name collision occurrence management framework, to assess risk both before and after controlled interruption is applied, and therefore to understand how the risk has changed. This not only provides assurance that the intervention has been worthwhile, but also gives an indication of the residual risk that may still need to be mitigated (which, one hopes would ideally be close to zero). In addition, the feedback would provide valuable guidance for improving the mitigation measure for future new gTLDs, including guidance on how long the interruption period needs to be.

To reflect the best practice of a risk management feedback loop, Recommendation 9 should be improved as follows (additions in bold):

> RECOMMENDATION 9: ICANN monitor the implementation of controlled interruption by each registry to ensure proper implementation and compliance, **and to assess effectiveness in mitigating risk**.

One practical way to assess the effectiveness of controlled interruption for the new gTLD operator to provide periodic samples of DNS queries and responses for analysis. Similar to the DITL project that seeks to understand DNS activity at the root servers, an organization like DNS-OARC could run an ongoing project to study "Day-in-the-Controlled-Interruption" data sets provided by registry operators, root server operators and other participants in the DNS ecosystem.

Researchers would need to bear in mind the possible presence of synthesized traffic (along with all the other variety of DNS queries) in assessing what changes may need to be made to the overall risk management process. But gaming tactics would not directly affect the approval of specific registrations.

Finally, it should be kept in mind that controlled interruption has never been deployed at the scale proposed, where the responses for queries to potentially hundreds of new gTLDs and hundreds of thousands of SLDs (or potentially *all* SLDs, with the wildcard option), are all changed at the same time to a novel IP address. There is no operational experience to indicate how users and system administrators will detect that a controlled interruption has occurred, nor how long it may take them, after detection, to remediate their systems. The recent research focus on name collisions has provided valuable insight into the extent of the problem, but further research is needed to understand the effectiveness of solutions as they are deployed. The feedback loop will help accomplish this objective.

6 Additional Comments and Questions

The previous sections covered the main observations and recommendations offered as preliminary comments on the Phase One Report. Here, some additional, brief comments are given on the specific recommendations in the report. As previously, they should not be taken as endorsement of the mitigation measures, but are intended as constructive feedback.

6.1 Blocking .CORP, .HOME and .MAIL

- *RECOMMENDATION 1: The TLDs .corp, .home, and .mail be permanently reserved for internal use and receive RFC 1918-like protection/treatment, potentially via RFC 6761.*[6]

Although it may be clear (pending publication of Phase Two Report) that these three applied-for new gTLDs are categorically at higher risk than all the rest, is it also the case that there are no SLDs in all the other applied-for new gTLDs that are of high enough risk to consider blocking indefinitely? The risk doesn't need to be high on average, just for enough installed systems. But without the benefit of the risk criteria Phase Two Report, there's not enough information on which to draw a conclusion.

6.2 Outreach efforts

- *RECOMMENDATION 2: ICANN continue efforts to make technical information available in fora frequented by system operators (e.g., network operations groups, system administration-related conferences, etc.) regarding the introduction of new gTLDs and the issues surrounding DNS namespace collisions.*

Outreach is essential for any transition of the present magnitude, and this particular aspect is one area very much under ICANN's control. The only practical way to engage with the vast community of potentially affected parties is with consistent communications from one party. It's not practical for hundreds of new gTLD operators all to be contacting system administrators on their own.

The security vulnerability notification that's causing the delay in the publication of the Phase Two Report ironically illustrates the importance of timely engagement with affected parties. A similarly motivated, though broader and longer-term obligation, should be taken into account when considering the need for outreach to DNS users.

ICANN's *Guide to Name Collision Identification and Mitigation for IT Professionals* [30] provides helpful orientation for system administrators who may not yet be aware of the name collision issue. The materials were developed prior to the release of the Phase One Report, however, so would need to be updated to indicate how to detect and manage controlled interruption responses should that mitigation method be adopted.

ICANN will also need an aggressive media strategy to ensure that communications related to controlled interruption lead people to ICANN resources, and not to malicious content.

6.3 "Clear and present danger" standard

- *RECOMMENDATION 3: Emergency response options are limited to situations where there is a reasonable belief that the DNS namespace collision presents a clear and present danger to human life.*

The rationale makes sense, and it also makes sense that this will be one of the more contentious recommendations. Even if emergency response options are limited, however, affected parties may still

[6] This and the nine other recommendations are quoted directly from the Phase One Report.

seek recourse in other situations. To that end, it would be helpful to have a standard form of "impact report" that affected parties can prepare or obtain to document their concerns.

6.4 No root-level de-delegation

- *RECOMMENDATION 4: Root-level de-delegation of a production TLD is not considered as an emergency response mechanism under any circumstances.*

The rationale again makes sense.

6.5 EBERO functionality

- *RECOMMENDATION 5: ICANN leverage the EBERO mechanisms and functionality to respond to DNS namespace-related issues. ICANN must have the following capabilities on a 24x7x365, emergency basis: 1). Analyze a specific report/incident to confirm a reasonable clear and present danger to human life; 2). Direct the registry on an emergency basis to alter, revert, or suspend the problematic registrations as required by the specific situation; 3). Ensure that the registry complies in a timely manner; and 4). Evaluate and monitor the specific situation for additional required actions. Furthermore, we recommend that ICANN develop policies and procedures for emergency transition to an EBERO provider and/or emergency root-level de-delegation in the event the registry is unable and/or unwilling to comply. We recommend ICANN maintain this capability indefinitely.*

This is reasonable given Recommendation 3, up to second half of the "Furthermore". However, given Recommendation 4, if root-level de-delegation is not an option, then why is it included here? Is this an "ability" vs. "advisability" distinction?

6.6 Controlled interruption wildcards for newly delegated gTLDs

- *RECOMMENDATION 6: ICANN require new TLD registries to publish the controlled interruption zone immediately upon delegation in the root zone. After the 120-day period, there shall be no further collision-related restrictions on the registry.*

The previous sections have presented several reservations and recommendations for improvement on this technique, which, nevertheless, is noteworthy for its creativity.

As far as the length of the controlled interruption period, the rationale for 120 days based on the amount of time it may take a user or system administrator to detect the break and then fix it – potentially across a large corporate network – seems quite reasonable as starting point for an untested technique. With more operational experience – which could be based in part on the analysis proposed in Section 5 – it may be possible to justify a shorter period.

From a risk/benefit perspective, starting the controlled interruption period at either contracting or delegation seems reasonable, given that the new gTLD operator's role and intentions will then be clear. Controlled interruption prior to delegation – for instance at the time a new gTLD is first applied for – imposes potential harm on installed systems to mitigate a name collision risk that may never occur.

A loopback address such as 127.0.53.53 is preferable to an internal network address because it's easier for a general user to manage. An external honeypot address should not be used. If controlled interruption is, in principle, a name collision, then controlled interruption with an external honeypot address is a **controlled exfiltration** – potentially drawing sensitive personal and corporate data to the collection site over an unencrypted path over the Internet.[7] As Schmidt noted in his WPNC '14 presentation, Response Policy Zones can be used to rewrite loopback address responses with a an internal network address if a system administrator wants to set up an internal honeypot, or attach intentionally to an external one.

6.7 Controlled interruption for SLDs on block list for new gTLDs in production

- *RECOMMENDATION 7: ICANN require registries that have elected the "alternative path to delegation," rather than a wildcard, instead publish appropriate A and SRV resource records for the labels in the ICANN 2LD Block List to the TLD's zone with the 127.0.53.53 address for a period of 120 days. After the 120-day period, there shall be no further collision-related restrictions on the registry.*

The general comments are the same as those on Recommendation 6, but there are additional concerns specific to this special provision for the alternate path.

First, as mentioned in Section 3.1, it is not clear from the Phase One Report alone that controlled interruption only of names on the SLD block lists will provide broad enough notification to users and administrators of installed systems that generate at-risk queries to the new gTLD involving SLDs that are not on the block list. For this reason alone, it seems clear that the risk treatment will be less effective for new gTLDs that have already been delegated and are thus subject to Recommendation 7, than for those that have not been and are subject to Recommendation 6.

Second, because Recommendation 7 is significantly less restrictive on a new gTLD operator than Recommendation 6, it is likely that new gTLD applicants will rush to the alternate path in order to be able to delegate some SLDs right away rather than having their business interrupted for 120 days. Accordingly, game theory suggests that the availability of the option in Recommendation 7 will increase overall risk. On the other hand, the rationale of not interrupting a business that's already underway (i.e., by putting the equivalent of a wildcard-with-exceptions in a production zone) is hard to argue with.

6.8 Wildcards temporarily allowed

- *RECOMMENDATION 8: ICANN relieve the prohibition on wildcard records during the controlled interruption period.*

[7] The honeypot could naturally be designed not to draw out sensitive information explicitly. But without further qualitative analysis, it remains unclear whether installed systems would offer sensitive information voluntarily, believing the honeypot to be an internal resource. This is a function both of protocol and implementation details specific to the particular installed system.

This is reasonable given Recommendation 7. Note also the discussion in Section 4 on wildcard-with-exception capabilities.

6.9 Monitoring

- ➢ *RECOMMENDATION 9: ICANN monitor the implementation of controlled interruption by each registry to ensure proper implementation and compliance.*

Monitoring is essential to any risk management process. As discussed further in Section 5, the monitoring must be part of a feedback process that determines whether the risk treatment is effective and guides future improvements.

Given that new gTLD operators have already demonstrated non-compliance with ICANN policy by delegating SLDs on the block list [31], it is important to ensure not only that registry operators are in compliance, but also that ICANN is effectively checking.

6.10 Medium-latency root summary feed

- ➢ *RECOMMENDATION 10: ICANN, DNS-OARC, and the root operators explore a medium-latency, aggregated summary feed describing queries reaching the DNS root.*

For similar reasons as to Recommendation 9, this improvement is also essential, both for the present problem and for ongoing security, stability and resiliency of the root server system. The implementation of RSSAC001 by all root server operators is an important part of accomplishing this goal.

6.11 Authoritative archive of historical root data

- ➢ RECOMMENDATION 11: ICANN, DNS-OARC, and the root operators explore establishment of a single, authoritative, and publicly available archive for historical data related to the root.

This is another important step forward, both for the present and for other purposes.

7 Conclusions

The name collision issue is, perhaps more than anything else, a symptom of the maturity of the Internet: a system that has grown up and prospered, while also developing its share of infirmities. JAS Global Advisors has done a credible job diagnosing the symptoms (even if the full diagnosis remains doctor-patient confidential at the moment), and has recommended a novel treatment.

As the Phase One Report is reviewed by the public, it is important to remember, as explained in Section 2, that the report alone is not the name collision management framework ICANN resolved in October 2013 that it would develop. Rather, the report suggests a generic mitigation measure, controlled interruption, to be applied to *all* new gTLDs (except for the three that are to be blocked entirely). Presumably the framework will be included in the Phase Two Report, now expected in June. But it would be premature for ICANN to act on the Phase One Report and implement its recommendations, before the actual framework that ICANN resolved to develop is available for public review. At most, any comments – including the ones in this document– should be taken as feedback on a proposed example mitigation measure – a "possible way forward" to quote the title of the announcement of the Phase One Report.

If ICANN does move forward without name collision mitigations without having made a framework available for public comment, as it resolved to do in October 2013, the action may appear to be expedient, but it would call into question ICANN's accountability to its own resolutions.

It should also be remembered that name collisions, despite the recent attention, are not the only security and stability issue for new gTLDs. Internal-name certificates [32] also remain a concern because of their potential for misuse against new gTLDs that overlap with the internal names. As a result of outreach by ICANN's Security and Stability Advisory Committee (SSAC), the CA/Browser Forum has resolved that such certificates will be revoked according to a set schedule, and new ones will no longer be issued. Like any mitigation measure, however, the resolution will need to be monitored and reviewed to ensure compliance. Also significant are the two limitations of the measure, both noted in a May 2013 update on the issue [33]. First, although certificate authorities generally follow best practices, not all are bound by CA/Browser Forum resolutions. Second, as is well known for public-key infrastructures in general, revocation is not always an effective control because many systems "fail open" when revocation information is unavailable. An adversary may therefore be able to cause such systems to continue to accept a certificate, even if it has been revoked, simply by blocking access to the revocation information.

The public suffix list [34] that defines the administrative boundary between domain name registries and independently managed domains is another concern, as it may take time for updates that reflect the zone cuts for the new gTLDs to be propagated to all relying parties. Without appropriate information about these "zone cuts" in the Public Suffix List, a user may be at risk of such threats as "browser super-cookies" that allow a rogue web site to compromise the privacy of other web sites within the same new gTLD.

The combinatorial complexity of risks in these interdependent Internet navigation and ecosystem elements, within and beyond the DNS, poses a substantial challenge to the security and stability of applications relying on a new gTLD. Effective risk management requires attention not only to each of the risks separately, but to their collective impact as well. The risks have been worked out over the years for established TLDs as the DNS has matured, but for new gTLDs, achieving a similarly balanced set of controls will take time.

Finally, on a more philosophical note: The security, stability and resiliency of the DNS is one of ICANN's priorities, and rightly so. The Phase One Report confirms, as others have previously concluded, that these properties are not at risk due to name collisions related to new gTLDs. However, the report continues,

- The remainder of our research is focused on issues from the perspective of end-systems as consumers of the global DNS.

Security, stability and resiliency of the DNS are essential, but they're a means to an end: reliability and confidence for users of the Internet. Whatever the historic causes for the present infirmities related to name collisions in the Internet, changes to the DNS must help to cure them, not make them worse. ICANN's resolution to improve the situation has set the right goal; the Phase One Report has presented a credible way forward for public review. Adding to this the engagement of the broader Internet community, it is reasonable to expect that the situation will improve.

8 References

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

[2] *New gTLD Security and Stability Considerations.* Verisign Labs Technical Report #1130007. Version 2.2, March 28, 2013. http://www.verisigninc.com/assets/gtld-ssr-v2.1-final.pdf

[3] Danny McPherson. Part 1 of 5; Introduction: New gTLD Security and Stability Considerations. Between the Dots, May 9, 2013. http://blogs.verisigninc.com/blog/entry/part_1_of_5_introduction

[4] *Name Collision in the DNS.* Interisle Consulting Group. Version 1.5, August 2, 2013. https://www.icann.org/en/about/staff/security/ssr/name-collision-02aug13-en.pdf

[5] *New gTLD Collision Risk Mitigation.* ICANN, August 5, 2013. https://www.icann.org/en/about/staff/security/ssr/new-gtld-collision-mitigation-05aug13-en.pdf

[6] *New gTLD Security, Stability, Resiliency Update: Exploratory Consumer Impact Analysis.* Verisign Labs Technical Report #1130008. Version 1.1, August 27, 2013. http://techreports.verisignlabs.com/tr-lookup.cgi?trid=1130008&rev=1

[7] Patrick S. Kane, Thomas C. Indelicarto, and Danny McPherson. *Letter to ICANN Board of Directors re: ICANN's Proposal to Mitigate Name Collision Risks – .CBA Case Study.* September 15, 2013. http://www.verisigninc.com/assets/report-cba-analysis.pdf

[8] Danny McPherson and Warren Kumari. *On DNS Search List Processing: Perhaps the Most Misunderstood Staple of DNS Resolution.* comments-name-collision-05aug13 discussion thread, September 17, 2013. http://forum.icann.org/lists/comments-name-collision-05aug13/msg00061.html

[9] *A Methodology for Assessing Collision Risk and New gTLDs.* Neustar, undated (first published September 2013). http://www.neustar.biz/enterprise/docs/whitepapers/domain-name-registry/new-tlds-dns-collision.pdf

[10]*New gTLD Collision Occurrence Management.* ICANN, October 4, 2013. http://www.icann.org/en/groups/board/documents/resolutions-new-gtld-annex-1-07oct13-en.pdf

[11]*ICANN Selects Lead for Development of Name Collision Occurrence Management Framework.* ICANN, November 11, 2013. http://www.icann.org/en/news/announcements/announcement-2-11nov13-en.htm

[12]*Statement of Work for the Development of the Name Collision Occurrence Management Framework*. ICANN, November 11, 2013. https://www.icann.org/en/about/staff/security/ssr/name-collision-sow-11nov13-en.pdf

[13]*Independent Report Maps Possible Way Forward in Mitigating Domain Name Collisions*, February 26, 2014. http://www.icann.org/en/news/announcements/announcement-26feb14-en.htm

[14]*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

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

[16]*Coordinated Vulnerability Disclosure Reporting at ICANN.* ICANN, Version 1.0, March 11, 2013. https://www.icann.org/en/about/staff/security/vulnerability-disclosure-11mar13-en.pdf

[17]*Name Collision Mitigation.* ICANN 49, Singapore, March 24, 2014. Audio transcript, http://audio.icann.org/meetings/singapore2014/name-collision-24mar14-en.mp3

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

[19]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

[20]Jeff Schmidt. *Name Collision Mitigation Update.* Presented at ICANN 49, Singapore, March 24, 2014. http://singapore49.icann.org/en/schedule/mon-name-collision/presentation-name-collision-24mar14-en.pdf

[21]Burt Kaliski. Part 2 of 4 – DITL Data Isn't Statistically Valid for This Purpose. Between the Dots, November 8, 2013. http://blogs.verisigninc.com/blog/entry/part_2_of_4_ditl

[22]Matthew Thomas, Yannis Labrou, and Andrew Simpson. *The Effectiveness of Block Lists to Prevent 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/wpnc2014_paper_effectiveness_block_lists.pdf

[23]Andrew Simpson. *Detecting Search Lists in Authoritative DNS.* 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/wpnc2014_paper_simpson.pdf

[24]Paul Gauthier, Josh Cohen, Martin Dunsmuir, and Charles Perkins. *Web Proxy Auto-Discovery Protocol.* Internet-Draft draft-ietf-wrec-wpad-01. July 28, 1999. http://tools.ietf.org/html/draft-ietf-wrec-wpad-01

[25]Fred L. Templin, Tim Gleeson, and Dave Thaler. *Intra-Site Automatic Tunnel Addressing Protocol (ISATAP).* IETF RFC 5214, March 2008. http://www.rfc-editor.org/rfc/rfc5214.txt

[26]*ISO 31000:2009, Risk management – Principles and guidelines.* ISO. First edition, November 15, 2009. http://www.iso.org/iso/home/standards/iso31000.htm

[27]*A Structured Approach to Enterprise Risk Management (ERM) and the Requirements of ISO 31000.* The Institute of Risk Management, 2010. http://www.theirm.org/documents/SARM_FINAL.pdf

[28]*ICANN DNS Risk Management Framework.* Westlake Governance. Draft 2.0 for public comment, August 19, 2013. http://www.icann.org/en/groups/other/dns-risk-mgmt/draft-final-19aug13-en.pdf

[29]Burt Kaliski. *Part 3 of 4 – Name Collision Mitigation Requires Qualitative Analysis.* Between the Dots, November 13, 2013. http://blogs.verisigninc.com/blog/entry/part_3_of_4_name

[30]*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

[31]Burt Kaliski. *Uncontrolled Interruption? Dozens of "Blocked" Domains in New gTLDs Actually Delegated.* Between the Dots, February 26, 2014. http://blogs.verisigninc.com/blog/entry/uncontrolled_interruption_dozens_of_blocked

[32]*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

[33]Patrik Fälström and Warren Kumari. *SAC057 / non-FQDN Certs.* Presented at RIPE 66, Dublin, Ireland, May 13-15, 2013. https://ripe66.ripe.net/presentations/143-RIPE-Lightning-SSAC057-paf.pdf

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