Google Registry

March 31, 2014

Via Electronic Mail
comments-name-collision-26feb14@icann.org

Re: Public Comment on Mitigating the Risk of DNS Namespace Collisions

We appreciate the opportunity to comment on the recommendations presented in the study on namespace collisions in the global Internet DNS and a framework for risk mitigation produced by JAS Advisors.

In general, JAS Advisors has done a thorough and insightful job of analyzing potential collisions due to the delegation of new top-level domains (TLD), and their proposed approach of adopting a controlled interruption period seems to be an appropriate precaution as part of the process of introducing new gTLDs. Notably, the controlled interruption framework should be superior at detecting problems and provides a better balance between the usability of new gTLDs and the protection of existing computer systems and technical process, compared to the blocklist approach employed in the Alternative Path to Delegation. Although we support the overall approach and encourage ICANN to proceed with a controlled-interruption-based approach to detect and mitigate name collisions as quickly as possible, below we suggest a number of improvements to the framework as suggested by the draft report.

1. Consider the use of a "honeypot" to perform controlled interruption rather than resolving names to loopback interfaces

In computer security, a "honeypot" is a system that is deployed to attract and capture attack traffic. The name comes from the English idiom, "You catch more flies with honey than you do with vinegar." Security researchers often deploy a honeypot to lure attackers, and then use logs from the honeypot to analyze and track the attacks.

Although the proposal to resolve names to the 127.0.53.53 IP address would likely allow for the detection of problems, we believe the use of a hosted honeypot as described in SAC62 (and described in more detail below), provides a better opportunity to inform users of impending problems, while at the same time the honeypot gathers information regarding the usage of the TLD. Although Section 2.1.5 of the JAS report provides some reasons for preferring the use of 127/8 addresses over a honeypot, we believe concerns about the logging of data do not supersede the superior notification characteristics of the honeypot. We also believe the concerns can be mitigated by a selective, audited approach to logging. A more detailed explanation of our approach is provided in Appendix A.

2. Separate notice and remediation periods

The JAS report proposes that the duration of the controlled interruption period should be 120 days. We believe that this period is far too long, if considered as a notice period, but may be insufficient in some cases if it is also intended to allow remediation of any problems that may be discovered. Systems that would be both adversely affected by name collision and would seriously impair the use of the DNS would almost certainly be discovered quickly—in days, if not hours. Even periodic monthly jobs that fail and require extensive diagnosis might require at most 45 days (31 days to cover all possible monthly frequencies plus 14 days for diagnosis) of controlled interruption before problems related to critical systems would reasonably be detected.

In particular, we consider it unlikely that: 1) periodic jobs that operate on a frequency of less than once a month would be essential to business operations; 2) would otherwise relate to the Internet's security and stability; or 3) would be otherwise isolated from other systems that would offer opportunities for much more rapid detection. At the same time, 120 days may not be a sufficient amount of time to remediate serious problems caused by name collision (for example, a large Active Directory namespace rooted in a proposed TLD or a widely-used subdomain that may result in name collision due to search path processing).

Rather than relying on a single, 120-day controlled interruption period, we propose a multi-phase approach to collision detection and remediation. In the controlled interruption phase, we propose no more than 24 hours for the first interruption and then removal of the wildcard record. This would allow enough time to detect the majority of problems and direct affected users to sites with more information on how to resolve this, while not causing lengthy outages for the same users. After some amount of time, longer controlled interruptions would be performed. Having a one-day outage for a company and then having time to implement mitigation options is preferable to having an outage for however long it takes to rename/reconfigure all the systems, etc.

There are no clear remediation steps and thus leaving end users in an interrupted state is suboptimal. The controlled interruption period would then be followed by a mitigation period to allow this round of problems to be resolved. Since the solutions may be non-obvious, time consuming, non-trivial, and require outside assistance/consultants, the resolution period will need to be long enough to accommodate these variables. The subsequent controlled interruption period would be 72 hours to catch additional problems as well as to confirm that the first round of solutions were effective. We would recommend longer and longer periods of controlled interruption, until the TLD is delegated. This will allow more issues to be caught and fixed in a controlled manner, rather than on an urgent basis after TLD delegation. In the event that ICANN elected to use a sustained, post-delegation controlled interruption period rather than this approach, we see no reason for a controlled interruption period longer than 45 days, which should be adequate to detect any serious problems caused as a result of the TLD’s delegation.

We strongly recommend that this approach is rolled out sooner rather than later, so that mitigation/resolution periods can be longer. As there are currently no clear cut solutions, we cannot direct affected users to a specific remediation; therefore, the mitigation/resolution periods will need to be lengthy. In the event that the delegation did, in fact, cause harmful collisions, affected parties could request that ICANN delay the launch of the registry's normal operations up to 180 days in order to allow for appropriate remediation. This will be better determined after the first phases of collision detection. Because TLDs affected by harmful collisions should be quite rare, rather than trying to develop elaborate threshold criteria to determine whether or not to allow for this remediation period, ICANN could simply require affected parties to identify themselves, provide a simple explanation of the problem, and request a specific time period up to 180 days in order to resolve the problems. Upon verifying that the source of any problems is, in fact, name collision, ICANN would automatically grant any requested extensions up to the 180 day maximum. This approach balances the need for appropriate notification through controlled collision, registry operators' desire to move quickly to launch after significant delays throughout the new gTLD program, and the likely rare necessity of providing extended time for collision mitigation.

3. Begin controlled interruption as early as possible

There is no reason to wait to begin controlled interruption until after the registry operator has completed pre-delegation testing and other aspects of onboarding designed to ensure effective operation under ICANN's registry/registrar model. As described above and in Appendix A, we believe controlled interruption could begin immediately for all proposed TLDs using nameservers and a process supervised by ICANN. Even if current IANA procedures would not allow such an approach, ICANN should, at minimum, begin controlled interruption immediately upon the execution of a TLD's Registry Agreement. This can be done while continuing to restrict a registry's ability to accept registrations and insert other names into the TLD's zone file until the registry operator has completed the relevant procedures. In addition to providing potentially affected users notice of problems as early as possible, this approach allows the controlled interruption period to run simultaneously with ICANN's existing requirement, specifically that registries do not add names to the TLD zone files for the first 120 days after the execution of the Registry Agreement in order to allow certificate authorities the opportunity to revoke internal certificates in the TLD's namespace.

4. Move consideration of reserved namespaces into the technical standards-setting process

Various IETF RFCs reserve a number of TLDs for specific uses. As the draft report notes, there is currently no namespace reserved for internal use like there is for IP address space. We agree that a TLD reserved for internal usage is desirable and encourage ICANN to work with the IETF standards-setting process to establish such a namespace. At this time, we believe that a new namespace such as .INTERNAL would be a superior approach to reserving applied-for TLDs including CORP, HOME, and MAIL. We encourage ICANN to continue to place these particular TLDs on hold pending additional standards-setting discussions. In the event that they were reserved through the IETF standards-setting process, it would obviously be necessary to terminate application processes for these TLDs, but establishing an alternative namespace may make TLDs such as CORP and HOME safe to delegate at some point in the future.

In particular, we note that MAIL differs markedly from CORP and HOME in that almost all of the queries are for the hostname “mail” rather than for subdomains within the TLD (99.8% of queries based on our previous investigation into this topic). Based on this difference in usage, as well as the fact that the use of “dotless” TLDs is disallowed by ICANN, we believe that it may be safe to delegate MAIL and it highlights the need for additional discussion through the technical standards-setting process in order to establish reserved namespaces rather than relying on a single report, however well done.

Conclusion

While we appreciate all of the work by JAS Advisors in producing its report on a topic that is important to the ICANN community, we still feel that improvements can be made to many of the recommendations.

We look forward to working with ICANN as we move forward into the public auction phase of the new gTLD program.

Sincerely,

Sarah Falvey

Sarah Falvey
Policy Manager and Primary Contact

APPENDIX A

Our approach is designed to minimize the impact on end users, and to provide as controlled a disruption as possible. While we cannot offer or provide specific solutions to the collisions that may occur, we hope that the honeypot approach results in a smoother transition for those end users who are affected by namespace collisions.

Any of the controlled interruption systems, by definition, interrupt the operation of colliding systems. This means that an organization using the colliding name internally will be broken until they understand, design, and mitigate the problem. We believe that the vast majority of collisions will be detected in a very short time (hours), but the mitigation could also take time.

We propose that ICANN perform short interruptions for each applied-for TLD and then remove the interruption to give the affected parties time to mitigate the problem. We also think that this should occur as soon as possible to give the affected party the largest amount of time possible. Therefore, we recommend that for all applied-for, but not yet delegated, names ICANN perform 24 hour controlled interruptions and then remove the delegation (these should be rolling, continuous tests, until the TLD is delegated).

The general plan of controlled interruptions has been described in the body of this document. The honeypot outlined above would, for example, return a web page containing helpful pointers to information on namespace collisions and potential mitigations for an HTTP request.

The primary advantages to the honeypot approach over just returning an address of 127.0.53.53 is the ability to provide much clearer feedback to the user, and the ability to log (and therefore evaluate) the extent of the collisions. Logging the number of hits to the honeypot and source of the traffic will allow ICANN and the new TLD applicants to understand the source (and potential severity) of the collisions. By providing the user with more descriptive/helpful information, the user should be able to understand and mitigate the issues sooner.

Description of the honeypot

The honeypot system would comprise of a set of nameservers and small number of honeypot servers.

The nameservers would be identical to those proposed in the JAS report, but instead of returning an IP address of 127.0.53.53 they would return the public IP address of the honeypot servers. This would result in the user seeing content from the honeypot server as described below.

The honeypot itself would be a machine that runs a large number of services (as many as we can easily stand up) and provides feedback (webpages, login banners, etc) to users.

Protocols (non-exhaustive list)

There are a number of protocols for which we can easily provide a honeypot service.
For many of the protocols/servers, we can simply use "standard" daemons; for example, for HTTP/web requests we can simply use Apache or some other standard webserver. There are also existing systems that are specifically designed as honeypots, for example, the “honeyd” project developed by Niels Provos.

HTTP / Web

It would probably be the simplest/most scalable to use a normal webserver (for example, Apache) configured to answer for all names. It would respond with a web page containing a banner warning the user that they have probably reached this site in error. It would also contain links to additional information on the name collisions (for example the ICANN "Name Collision Resources & Information" page). The served page should attempt to be localized for the connecting client. This could be accomplished with IP geolocation and/or with the "Accept-Language:" browser header.

Telnet

This would probably also be easiest with a standard daemon. It would present a login banner containing similar information to the webserver. The telnet server should disconnect immediately after sending the banner, and not accept usernames and passwords.

SSH

This is basically identical to telnet. It should send a login banner, and then disconnect/not offer authentication methods. Login banners are only supported on SSH Version 2.

SMTP / Email

This could also be done with a standard server. It could be configured to reject connections at various levels, either immediately (when a client connects), after the HELO / EHLO verb, after the MAIL FROM, or after the RCPT TO. In order to limit the amount of data collected, we suggest against actually accepting the mail.

Minimizing data collection

One of the concerns with running a honeypot for namespace collisions is that the honeypot causes traffic to leave the local network, and this might include sensitive information.
To mitigate these concerns we provide a number of suggestions:

Strip URL query parameter

URLs sometimes contain query parameters that could be sensitive (for example http://www.example.com/login.html?username=fred&password=hunter3, or http://fred:password@example.com). The server should be specifically configured to strip off and not to log any of the URL parameters (it should only log the requested hostname, and probably the remote IP address).

External audit

An external, trusted auditor should be engaged to audit the information that is being collected and that the system is not logging additional information.

Causing data to leave the network

One of the justifications for using 127.0.53.53 is that this should cause a connection failure, but should not cause traffic to leave the host/network. The honeypot solution is different - it *does* cause a connection to the honeypot.

We believe that this is an acceptable risk. The honeypot will be designed to minimize the amount of actual data transmitted (by closing connections early/providing a clear login banner, etc). The honeypot also provides some data to the end user, in terms of what the problem is and where solutions might be found; the 127.0.53.53 approach simply causes a connection failure, leaving the end user with no information about what failed, why, and/or how to go about fixing the problem.

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. While not ideal, we think that it is preferable to the alternative, which is that the affected user might not know what is happening, and so won't know the risk. This means that they will not mitigate this, and once the TLD goes live, a malicious user would register names specifically to collect this information. From a security standpoint, it is preferable to send data to a trusted third party than accidentally sending it to a malicious phishing site.