Mitigating the Risk of DNS Namespace Collisions

A Study on Namespace Collisions in the Global Internet DNS Namespace and a Framework for Risk Mitigation

Final Report

28 OCTOBER 2015

TABLE OF CONTENTS

1 Preface to Final Report

JAS would like to thank ICANN and the ICANN Community for their patience in the months since the publication of our Phase One report. JAS, together with ICANN and Microsoft, elected to hold the publication of the complete Final Report until Microsoft released a fix to the critical MS15-011 [1] ("JASBUG") vulnerability. As a result, the overall impact of this critical vulnerability was materially reduced.

Microsoft offers its appreciation to the [Coordinated Vulnerability Disclosure] CVD community and a special thanks to the reporters of the issue which has resulted in UNC Hardening: Jeff Schmidt of JAS Global Advisors, Dr. Arnoldo Muller-Molina of simMachines, The Internet Corporation for Assigned Names and Numbers (ICANN) and Luke Jennings from MWR Labs. [2]

The following pages contain updates throughout as issues related to the Microsoft vulnerability may now be discussed. Material that did not appear in the Phase One report appears in Sections 4 and 5 and Appendices A and B of this report. None of JAS' recommendations have changed.

[1] https://technet.microsoft.com/en-us/library/security/ms15-011.aspx
[2] http://blogs.technet.com/b/srd/archive/2015/02/10/ms15-011-amp-ms15-014-hardening-group-policy.aspx

2 Summary

Collisions in the global Domain Name System (DNS) namespace have the potential to expose serious security-related issues for users of the DNS. This report dives right into the technical discussion and is targeted at readers who have been following the issue. Those new to the issue should first read the introductory documents located at: http://www.icann.org/en/help/name-collision.

We do not find that the addition of new Top Level Domains (TLDs) fundamentally or significantly increases or changes the risks associated with DNS namespace collisions. The modalities, risks, and etiologies of the inevitable DNS namespace collisions in new TLD namespaces will resemble the collisions that already occur routinely in the other parts of the DNS. The addition of multiple new TLDs over the past decade (generic and country code) has not suggested that new failure modalities might exist; rather, the indication is that the failure modalities are similar in all parts of the DNS namespace. Our research has shown that a very few root causes are responsible for nearly all collisions, and these root causes appear in nearly every classification of TLD, albeit in varying proportions.

That said, DNS namespace collisions are a complex and pervasive occurrence that manifests throughout the global Internet DNS namespace. Collisions in all TLDs and at all levels within the global Internet DNS namespace have the ability to expose potentially serious security and availability problems and deserve serious attention. While current efforts to expand the global DNS namespace have collision-related implications, the collision problem is bigger than new TLDs and must be viewed in this context.

In summary, our recommendations describe a comprehensive approach to reducing current and future DNS namespace collisions, alerting operators of potential DNS namespace related issues, and providing emergency response capabilities in the event that critical (e.g., life safety) systems are adversely impacted.

DNS namespace collisions exist outside of, and independently from, the current efforts to expand the DNS namespace. These collisions have almost certainly existed since the emergence of a global public DNS. As early as 2003, multiple researchers have pointed to the existence of queries into undelegated space received at the root.[3,4,5,6] Our research shows that every TLD that has been added to the root since

---

[3] *Understanding DNS Evolution*, Castro, Zhang, John, Wessels, claffy, 2010, http://www.caida.org/publications/papers/2010/understanding_dns_evolution/understanding_dns_evolution.pdf

[4] *Is Your Caching Resolver Polluting the Internet?*, Wessels, 2004, http://dns.measurement-factory.com/writings/wessels-netts2004-paper.pdf

[5] *RFC 4697: Observed DNS Resolution Misbehavior*, Larson, Barber, 2006, http://tools.ietf.org/html/rfc4697

consistent data collection has occurred (2007) has exhibited some symptoms of collision activity prior to delegation.

The issue of collisions is not specific to TLDs; rather, risk exists wherever a collision crosses an administrative control boundary in the DNS. Said differently, the most dangerous DNS namespace collisions occur when *the resulting DNS query is resolved by a different administrative party than expected by the querier*. This makes intuitive sense. Because of the hierarchical nature of the DNS, the vast majority of administrative control separations occur at the TLD and Second Level Domain (2LD) levels.

Over the course of the study, JAS found no evidence to suggest that the security and stability of the global Internet DNS itself is at risk. This finding confirms the results of the *DNS Stability String Review* performed on each string during Initial Evaluation pursuant to Section 2.2.1.3.1 of the Applicant Guidebook (AGB).[7,8] The remainder of our research is focused on issues from the perspective of end-systems as consumers of the global DNS.

When faced with a range of unknowns and hypotheticals, it is important not to overlook emergent facts and experience. At the time we wrote the Phase One report, 275 New gTLDs had been delegated and over 835,000 second level registrations had been added. TLDs representative of the complete range of the taxonomy JAS developed (see Section 5) are represented. .berlin – a geographic term that our research suggests is heavily present in DNS search paths – has the third largest number of registrations of all new TLDs. .email and .link – short, technology-oriented generic terms that our research suggests are present in a number of hardcoded configurations – rank 6th and 7th respectively, each with over 30,000 2LD registrations. .company, .solutions, and .agency – terms that our research suggests are commonly hardcoded into small business-oriented configurations – are also delegated and have thousands of registrations each. Neither JAS nor ICANN is aware of even a single instance of a seriously problematic collision. Of course this fact certainly doesn't "prove the negative" but it also can't be ignored at this point.

Certainly the nature of the string impacts the drivers behind colliding behavior, and history provides lessons and data regarding the introduction of a variety of strings

---

[6] *Wow, that's a lot of packets,* Wessels, Fomenkov, 2003, http://www.caida.org/publications/papers/2003/dnspackets/wessels-pam2003.pdf

[7] *gTLD Applicant Guidebook*, ICANN, 2012, http://newgtlds.icann.org/en/applicants/agb

[8] The process followed by ICANN's vendor for this review, Interisle Consulting Group, process is documented at http://newgtlds.icann.org/en/program-status/evaluation-panels/dns-stability-process-07jun13-en.pdf

at the TLD. As we presented at Verisign's *Workshop and Prize on Root Causes and Mitigation of Name Collisions* (WPNC)[9] in London, strings with the potential to introduce new failure etiologies have been introduced into the TLD in the past. .post, (delegated in 2012) saw the most collision activity prior to delegation of any of the nine TLDs added since 2007. .post is interesting because "post" is also an HTTP method and a not insignificant proportion of the collisions appeared to be related to erroneous DNS lookups of text intended to be transmitted to an HTTP server. History provides lessons and data regarding the introduction of a variety of strings to the TLD.

We believe the introduction of new TLDs offers an opportunity to educate operators regarding DNS namespace collisions and help find and remedy potential collision-related issues that may be present in their systems. As such, we recommend implementation of a 90-day "controlled interruption" period for all approved new TLDs with the exception of .corp, .home, and .mail. Registries that have not yet been delegated to the root zone shall implement controlled interruption via wildcard records; registries that have elected the "alternative path to delegation" shall implement controlled interruption by adding appropriate resource records for the labels appearing in their respective block lists. Following the 90-day controlled interruption period, registries will not be subject to further collision-related restrictions. Like the Certificate Authority (CA) revocation approach, which may be partially implemented in parallel, we believe the 90-day controlled interruption period offers a conservative buffer between potential legacy usage of a TLD and the new usage.

Lacking clear RFC 1918-like guidance directing operators to DNS namespaces safe for internal use, several such namespaces have been "appropriated" for this purpose over the years. While the etiology is subtly different, the .corp and .home TLDs are clear outliers in this respect; the use of .corp and .home for internal namespaces/networks is so overwhelming that the inertia created by such a large "installed base" and prevalent use is not likely reversible. We also note that RFC 6762 suggests that .corp and .home are safe for use on internal networks.[10]

Given that the Internet has demonstrated a need for RFC 1918-like DNS namespaces, we recommend that .corp and .home be referred to the Internet Engineering Task Force (IETF) for potential RFC 1918-like protection/treatment.[11]

---

[9] http://namecollisions.net

[10] *RFC 6762: Multicast DNS* (appendix G), Cheshire, Krochmal, 2013, http://tools.ietf.org/html/rfc6762

[11] RFC 6761 may be the appropriate vehicle for implementing a permanent reservation.

RECOMMENDATION 1: The TLDs .corp, .home, and .mail be referred to the Internet Engineering Task Force (IETF) for potential RFC 1918-like protection/treatment.

Like .corp and .home, the TLD .mail also exhibits prevalent, widespread use at a level materially greater than all other applied-for TLDs. Our research found that .mail has been hardcoded into a number of installations, provided in a number of example configuration scripts/defaults, and has a large global "installed base" that is likely to have significant inertia comparable to .corp and .home. As such, we believe .mail's prevalent internal use is also likely irreversible and recommend reservation similar to .corp and .home and similarly recommend ICANN not delegate that TLD at this time.

JAS uncovered a vulnerability not directly related to ICANN's New gTLD Program nor to new TLDs in general that has the potential to impact end-systems. Pursuant to ICANN's Coordinated Vulnerability Disclosure Process,[12] ICANN shall: "...privately disclose information relating to a discovered vulnerability to a product vendor or service provider ("affected party") and allow the affected party time to investigate the claim, and identify and test a remedy or recourse before coordinating the release of a public disclosure of the vulnerability with the reporter." Furthermore, ICANN's process states: "All parties to the disclosure generally agree to refrain from disclosing the vulnerability to the public until a remedy is identified and tested or until the threat is considered contained."

After extensive discussions with impacted vendors and ICANN executives, JAS is concerned that publication of the experimental methods and data contained in the complete JAS report may accelerate discovery of the vulnerability and/or serve to facilitate exploitation of the vulnerability after it is discovered. As such, pursuant to ICANN's process and out of an abundance of caution, JAS published the report in two phases: a Phase One report published in June, 2014 and this Final Report published after the impacted vendor addressed the vulnerability.

2.1 Summary of Recommendations

RECOMMENDATION 1: The TLDs .corp, .home, and .mail be referred to the Internet Engineering Task Force (IETF) for potential RFC 1918-like protection/treatment.

RECOMMENDATION 2: ICANN continue efforts to make technical information available in fora frequented by system operators (e.g., network operations groups,

[12] *Coordinated Vulnerability Disclosure Reporting at ICANN*, ICANN, 2013, https://www.icann.org/en/about/staff/security/vulnerability-disclosure-05aug13-en.pdf

system administration-related conferences, etc.) regarding the introduction of new gTLDs and the issues surrounding DNS namespace collisions.

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.

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

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 in the event the registry is unable and/or unwilling to comply. We recommend ICANN maintain this capability indefinitely.

RECOMMENDATION 6: ICANN require new TLD registries to publish the controlled interruption zone immediately upon delegation in the root zone. After the 90-day period, there shall be no further collision-related restrictions on the registry.

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 90 days. After the 90-day period, there shall be no further collision-related restrictions on the registry.

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

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

RECOMMENDATION 10: ICANN work with the IETF to identify a mechanism for IPv6 that provides similar functionality to that available in IPv4's "localhost" reserved prefix.

RECOMMENDATION 11: ICANN, DNS-OARC, and the root operators explore a medium-latency, aggregated summary feed describing queries reaching the DNS root.

RECOMMENDATION 12: ICANN, DNS-OARC, and the root operators explore establishment of a single, authoritative, and publicly available archive for historical data related to the root.

RECOMMENDATION 13: ICANN explore collecting NXDOMAIN entries in DNS query logs from registry operators and contribute them to an independent data repository such as DNS-OARC for further analysis. To limit the potential for commercial gaming or use by malicious parties, we recommend that logs be provided six months in arrears.

RECOMMENDATION 14: ICANN request that the appropriate bodies further explore issues relating to collisions in existing DNS namespace, the practice of "domain drop catching," and the associated data feeds that may be leveraged by attackers when attempting to exploit collisions.