Independent Submission
Request for Comments: 8023
Category: Informational
ISSN: 2070-1721

M. Thomas

A. Mankin
Salesforce

L. Zhang
UCLA

November 2016

Report from the Workshop and Prize on Root Causes and Mitigation of Name Collisions

Abstract

This document provides context and a report on the workshop on "Root Causes and Mitigation of Name Collisions", which took place in London, United Kingdom, from March 8 to 10, 2014. The main goal of the workshop was to foster a discussion on the causes and potential mitigations of domain name collisions. This report provides a small amount of background and context; then, it provides a summary of the workshop's discussions.

Status of This Memo

This document is not an Internet Standards Track specification; it is published for informational purposes.

This is a contribution to the RFC Series, independently of any other RFC stream. The RFC Editor has chosen to publish this document at its discretion and makes no statement about its value for implementation or deployment. Documents approved for publication by the RFC Editor are not a candidate for any level of Internet Standard; see Section 2 of RFC 7841.

Information about the current status of this document, any errata, and how to provide feedback on it may be obtained at http://www.rfc-editor.org/info/rfc8023.

Copyright Notice

Copyright (c) 2016 IETF Trust and the persons identified as the document authors. All rights reserved.

This document is subject to BCP 78 and the IETF Trust's Legal Provisions Relating to IETF Documents (http://trustee.ietf.org/license-info) in effect on the date of publication of this document. Please review these documents carefully, as they describe your rights and restrictions with respect to this document.

Table of Contents

1. Introduction

It has been well known within the Internet research and engineering community that many installed systems in the Internet query the domain name system (DNS) root for names under a wide range of top-level domains (TLDs). Many of these TLDs are not delegated, which results in a response indicating that the name queried does not exist (commonly called an NXDOMAIN response [RFC7719]). In the Internet Corporation for Assigned Names and Numbers (ICANN) community, it was observed as early as November 2010 by the Security and Stability Advisory Committee (SSAC) report [SAC045] that the addition of new TLDs in the DNS root could result in so-called name collisions for names used in environments other than the global Internet. Some installed systems, following established (albeit not vetted) operational practices, generate queries to the global DNS with name suffixes that, under seemingly reasonable assumptions at the time the systems were designed or configured, were not expected to be delegated as TLDs. Many of these installed systems depend explicitly

or implicitly on the indication from the global DNS that the domain name suffix does not exist. After a new TLD is delegated, the global DNS may give a different response to the query involving the TLD than it did prior to the TLD's delegation.

A name collision occurs when an attempt to resolve a name used in a private namespace results in a query to the public DNS, and the response indicates that the name is in the global DNS [NCRI]. In other words, the overlap of public and private namespaces may result in potential unintended (and, therefore, potentially harmful) resolution results. The impact of the global change on installed systems will be varied; risks to installed systems introduced by name collisions may arise due to varied causes.

In a globally distributed system, such as the Internet, it is difficult, yet critical, to agree on policies for demarking boundaries of ownership and autonomy. Name space governance is critical to ensure predictable use of names in the global DNS.

In order to help ensure this uniqueness and interoperability, ICANN, through its coordination of the IANA functions, is responsible for administration of certain responsibilities associated with Internet DNS root zone management, such as generic and country code Top-Level Domains (gTLDs and ccTLDs). Prior to ICANN's creation in 1998, seven generic TLDs were defined in the early development of the Internet [RFC1591]. Since the formation of ICANN, the delegations of generic, internationalized and country code TLDs have been administered and delegated by ICANN. During these delegations, it quickly became apparent within the IETF community that there was a need to reserve name spaces that can be used for creating limited sets of internal names without fear of conflicts with current or future TLD name spaces in the global DNS [RFC2606].

While the reserved TLDs [RFC2606] aimed to enable operators to use them only as a small set of reserved names internally, with limited uses, educational awareness and operational best practices did not achieve the goal of reserving special-use domain names [RFC6761]; other suffixes, not reserved though at the time not in conflict, were often employed instead. Faulty assumptions, or encouragement in some cases by vendor documentation, of "we only use this name internally and there is no possibility of leakage to the global DNS" were made by numerous operators or administrators. Numerous reports and findings have clearly disproved these faulty assumptions by showing substantial "DNS leakage" into the global DNS through mechanisms such as search lists.

In 2012, ICANN created a new gTLD program to add a potentially unlimited number of new gTLDs to the root zone as a mechanism to

enhance competition, innovation, and consumer choice. With the potential of many new gTLDs becoming delegated in the global DNS, operators or administrators who elected to use a non-delegated name space internally may face potential "name collision" problems.

This document is primarily a report on the March 2014 workshop that set out to examine the causes and mitigation of such name collisions and their associated risks. It is a companion to the Workshop and Prize on Root Causes and Mitigation of Name Collisions proceedings [WPNC], and it also provides some additional background and context.

2. Background and Context

When the workshop was convened, the context and status of the work around name collisions could be described as follows.

Since early 2008, there had been numerous lengthy discussions within the ICANN community about the ability of the DNS root to scale to accommodate new gTLDs and the impact of making those changes on the DNS ecosystem. In March 2008, the Internet Architecture Board (IAB) observed that the introduction of suffixes in use in a number of environments could lead to instability [IAB2008]. In December 2010, the Security and Stability Advisory Committee (SSAC) issued their report on root scaling in which the committee formalized several recommendations based on "actual measurement, monitoring, and data-sharing capabilities of root zone performance" to help determine the feasibility of root scaling [SAC046]. Separately, the Root Server System Advisory Committee [RSSAC] agreed in late 2010 on the need to establish standard metrics to be collected and reported by all operators. This effort would provide the community with a baseline measure of the entire root server system’s performance. With such an established baseline, any possible negative effect from additional TLDs within the root could potentially be identified. In late 2012, the ICANN Board affirmed the need to work with the root server operators via RSSAC to complete the documentation of the interactions between ICANN and the root server operators with respect to root zone scaling [IR2012].

In March 2013, SSAC published an advisory titled "SSAC Advisory on Internal Name Certificates," which identified a Certificate Authority (CA) practice that, if widely exploited, "could pose a significant risk to the privacy and integrity of secure Internet communications" [SAC057]. The ICANN Board acknowledged the issues identified in the advisory report on internal name certificates [SAC057] as part of a more general category of issues. These issues included installed systems utilizing a namespace in a private network that includes a non-delegated TLD that is later delegated into the root. In May 2013, the ICANN Board commissioned a study on the use within private

name spaces of TLDs that are not currently delegated at the root level of the global DNS [ISTUDY]. This study was focused on potential name collision events between applied-for new gTLDs and non-delegated TLDs potentially used in private namespaces. The study also examined the potential possibility of name collisions arising from the use of digital certificates referenced in the SSAC report on internal name certificates [SAC057].

Between the RSSAC's and SSAC's advisory statements ([RSSAC] [SAC046]) and the ICANN commissioning of a study in May 2013, there was significant progress on establishing formalized, coordinated monitoring and measurement of the root. RSSAC approached its finalization of the specific metrics that each root operator will collect and initiated discussions about where the operators will send their data for analysis once collected. To properly gauge the risks of new gTLD delegations to the root, an established baseline of normal performance of the system would be required to start sufficiently ahead of the new delegations. The execution of these RSSAC and SSAC recommendations was timed poorly with the commissioned study, resulting in a limited pool of data repositories from which any baseline and risk measurements could be established.

It is common practice for each root operator to monitor its own root server, and some operators report the status and performance of their services publicly. As of ICANN's study commissioned in May 2013 [ISTUDY], there was no mechanism in place to allow a detailed view of the entire root system, short of the annual "Day in the Life" ([DITL]) data repository, which contains root DNS data over a short coordinated time period from a varying subset of root operators and was intended to be used for research purposes, not to provide overall monitoring and an operational view of system health. Due to the lack of a more comprehensive and desirable data repository for baseline and collision analysis DITL has become the de facto referential dataset for root traffic analysis.

The commissioned study, conducted by the Interisle Consulting Group, was published in August of 2013. Their report "Name Collisions in the DNS" [INTERISLE], based on [DITL] measurements, addressed name collisions in the DNS and also recommended options to mitigate the various name collision risks. The study identified categories of strings according to the risk they represent: low risk (80 percent of applied-for strings), uncalculated risk (20 percent of applied-for strings), and high risk (2 applied-for strings).

At the same time as the [INTERISLE] study, ICANN published a proposal, titled "New gTLD Collision Occurrence Management Plan" [NGCOMP], to manage the risk of name collisions within the applied-for gTLDs. Based on measurements, ICANN deemed two strings, .home

and .corp, to be high risk because of their widespread use within internal networks and would indefinitely delay their delegation [INTERISLE]. Those strings within the uncalculated-risk classification would be delayed 2 to 3 months in their application process while ICANN conducted more research into whether the string is of high- or low-risk classification. Those in the low-risk classification would face a delay in activating domains until 120 days after contracting with ICANN to allow for the change in certificate authority practices recommended in the SSAC report on internal name certificates [SAC057].

Within the ICANN proposal [NGCOMP], an approach termed the "alternative path to delegation" was outlined, in which a registry operator could elect to proceed with delegation, provided it initially blocked all second-level domains (SLDs) that appeared in the certain DITL datasets pending the completion of the assessment. The majority of new gTLD applicants that were eligible elected this alternative path once otherwise approved for delegation. The plan also outlined an outreach campaign to educate system administrators, software developers, and other engineers about the name collision issue and possible mitigation measures.

As a further provision, the "New gTLD Collision Occurrence Management Plan" called for a follow-up study that would develop a "Name Collision Occurrence Management Framework" [NCOMF]. In February 2014, the document, "Mitigating the Risk of DNS Namespace Collisions: Phase One Report," was published by the ICANN-contracted group JAS Global Advisors [MRDNC]. The report provides a number of recommendations for addressing the name collision issue focusing on a technique termed "controlled interruption," in which a registry would temporarily resolve all SLDs (or all SLDs present in the block list) to a specific IP: 127.0.53.53. The report also makes provisions to implement an emergency plan and strategy in case name collisions had a "clear danger to human life."

2.1. Brief Update

In the time frame after the workshop, a final version of the Phase One Report was released in June 2014 [MRDNC].

In July 2014, after a community review phase, a final recommendation was issued by ICANN [NCOMFINAL]; this has been followed by the publication of management documents for the implementation of a controlled interrupt for new gTLD delegations [NOCA] [NCSLDCIV] [ADDNOCA].

Much of the framework called for in the Name Collision Occurrence Management Framework [NCOMF] was not released by the time of writing this document, and the Phase One Report [MRDNC] indicated that its publication was delayed due to a security vulnerability [JASBUG] identified during the course of the work.

Broad community efforts to measure the impact of name collisions were not included in the final recommendation issued by ICANN [NCOMFINAL]. At the time of this writing, RSSAC has just published its specification of common measurements to be collected by root operators, meeting one part of the needs for measurements of the root server system [RSSAC002].

3. Workshop Structure

The Workshop and Prize on Root Causes and Mitigation of Name Collisions [WPNC], sponsored by Verisign, took place March 8-10, 2014 in London, United Kingdom. The WPNC was open to the public, and it gathered subject-area specialists, researchers, and practitioners to discuss and present their views, concerns, and ideas surrounding the name collision issue. Proceedings are published at the workshop's website [WPNC].

The workshop focused on studies of name collision risks and mitigations with the expectation to advance the global community's insight into operational uses of name suffixes that can result in name collisions and to gain a stronger understanding of the potential risks for the users of the installed systems. Additional emphasis and attention was given to discussions that might advance the state of knowledge about the architecture and impacts of DNS namespaces with multiple scopes or resolution contexts and the utilization of new methods of monitoring and understanding the needs and methods for mitigating emerging Internet risks around name collisions. A technical program committee, whose members spanned a variety of organizations and universities, was assembled. The committee issued a call for papers and evaluated all submissions to ensure the highest level of quality.

A synthesis of the accepted papers and conference proceedings is captured in the subsections below. Another informal synopsis of the workshop combined with individual statements and observations is available online [COMMENTARY].

3.1. Research Findings

Many of the research papers focused on the analysis of DITL data to better understand various aspects of the root NXDOMAIN traffic ([TECHNIQUES], [RARDBITS], [BLOCKLISTS], [MODELING], and [SEARCHLISTS]). Note: all workshop contributions are listed in Appendix B; full papers and slides are available at the website [WPNC].

While the DITL data has become the de facto referential dataset for root traffic analysis, some presenters echoed concerns that the dataset may have become biased or polluted with "artificial" queries after the ICANN "Reveal Day," in which the list of applied-for gTLD strings was publicly disclosed. No conclusive or empirical evidence of tampering was presented; however, concerns about the integrity and reliability of future DITL collections and analysis for purposes related to new gTLDs were echoed by some panelists [IESCPANEL]. Furthermore, the statistical accuracy and completeness of DITL data -- used to draw inferential conclusions or more specifically create SLD block lists -- was examined. The efficacy of blocking domains based on sampled DNS data, e.g., DITL, was investigated by comparing measurements of SLDs within DITL and that of a multi-month root NXDOMAIN collection at the A and J roots [BLOCKLISTS]. The findings provided insights into SLD-root affinities, SLD temporal query patterns and occurrence frequencies that demonstrated the ineffectiveness of block listing domains based on sampled DNS data such as [DITL].

Measurements of queries specifying the recursion desired (RD) bit to the roots in DITL were quantified to identify the level and nature of naive DNS clients and to determine and assess potential impacts that could arise from the proposed SLD blocking technique to these naive clients [RARDBITS]. A substantial proportion of the root server request traffic contained queries with the RD bit specified. Both in absolute and relative terms, requests specifying the RD bit for applied-for gTLDs were found to be significantly lower when compared to existing TLDs. The root cause determination of what system or mechanism is responsible for generating the queries was inconclusive and only speculative explanations of faulty implementations of a DNS resolving server were hypothesized. However, the analysis was also not able to identify instances of actual or potential harm resulting from these naive clients, suggesting if SLD blocking techniques were to be utilized, it is unlikely there would be any negative impact to these naive clients.

3.2. System Analysis

Comparison of elements can often help us to understand a system as a whole. A passive study of the DNS traffic in a provisioned domain such as "corp.com" may elucidate certain name collision parallels [CORPCOM]. Such measurements were presented as a proxy for the ".corp" potential new gTLD. According to the study, significant DNS traffic volume was directed at a variety of third-level domains under "corp.com". This prompted a series of questions surrounding how name collisions can be identified, as most end-users won't recognize that problems may be due to a name collision. How will users know that the problem they are experiencing is a result of a new, colliding gTLD? Will support groups be able to diagnose a name collision event from reported symptom(s)? Will a collision-based security hole be detectable?

These questions, upon which underpinnings rely on communication and educational awareness, may find recommendations or parallels from other system references during the workshop [JASFRAMEWORK] -- such as the postal and telephone system. Most telephone and postal systems have evolved over time, requiring individuals to alter the way they address their parcels or place their calls. Both systems implemented their changes in such a way that prior to the change, educational material is distributed and communicated and for a period of time and after the change, compliance of the previous standard is temporarily accepted. While the telephone and postal system operate in a very different way than the DNS, these parallels of "advanced notification, education and communication, and a grace period" were insightful for how other similar systems transitioned.

3.3. Frameworks: Modeling, Analysis, and Mitigation

Statements from several TLD operators during the conference reverberated a theme for the need of improved tooling, education, and communication surrounding name collisions. The delegation of new gTLDs is an ongoing event, and there is a clear and immediate need for these operators to have visibility to monitor and measure the effects of these new gTLD delegations. A lack of tools, shared data, communication, and education surrounding name collisions has handicapped operators in their ability to quantitatively measure and proactively provide any steps for mitigation of risks. To this end, numerous techniques, frameworks, and models that focused on the concepts of analyzing, detecting, and measuring various name collision risk factors were presented and reviewed with the hope of understanding these underlying concerns and issues ([TECHNIQUES] [MODELING] [SEARCHLISTS] [DNSENDUSER] [ENTNETWORK]).

Data-driven analysis and mitigation require operators to be versed and skilled with data analysis techniques to better understand the contextual intent and ownership of DNS queries. An overview of various DNS analysis techniques in which ways of decomposing names, measuring temporal distributions between queries, and detecting organizational/geographical affinities was presented [TECHNIQUES]. More-specific techniques were also showcased, such as a systematic way of observing and characterizing the impact of search lists within root DNS traffic allowing operators to quantify the number of unique entities that may be reliant on a particular name space [SEARCHLISTS]. While not exhaustive, the techniques presented have been proven to elucidate patterns within root DNS traffic data and could serve as the potential building blocks of a DNS analysis framework.

Most of the previously published work focused on name collisions has produced various quantitative analyses based on observations of Internet traffic and data, including DNS queries and web content, in which behavior and associated risks have been inferred. An understanding of the inverse of the process by starting with a fundamental model of name resolution at the client was proposed as an alternative means to define risk [MODELING]. This model deconstructed the process of name resolution at the resolver library of a client system and formalized a model from which derived metrics could be used to define and quantify associated risks. While the model presented is only a piece of the greater name collision puzzle, it provides potentially new insights into what may otherwise be considered a missing piece.

Just as important as understanding the root causes of name collisions, providing effective mitigation strategies is a critical piece of the name collision puzzle. Mitigation can be achieved from both higher levels, such as ICANN, as well as the enterprise level. Proposed strategies for mitigating name collisions at both of these levels were presented. While the technical details for each proposed strategy varies, underlying dependencies in both strategies require operators to monitor and educate/train their users.

3.4. Conclusions and Next Steps

In their concluding statement [NEXTSTEPS], the workshop committee stated:

> It occurs to the program committee that the analysis of the interactions between the different uses of domain names within local or global context is almost a nonexistent topic of research. This may have to do with the lack of accessible data, lack of theory of root causes, a lack of interest, or a bias in the participation of the workshop. We think that this is evidence that this study of the global centrally important technical system needs to be ramped up.

Follow-on commentary [NEXTSTEPS] from the attendees reaffirmed this opinion with recurring messages of a need to understand the root causes of name collision and the need to overcome shortcomings within our shared data collection, monitoring, and analysis of the DNS.

Many name collision unknowns still exist. What are the root causes of these queries? What is going on within a recursive name server? What vulnerabilities or subtle attack vectors do these new gTLD delegations enable? The limited datasets available to researchers and operators are not sufficient to draw baseline measurements for these questions, forcing the community to make inferences and rank guesses as to what is going on within the DNS. Using these suboptimal data repositories to create solutions such as block lists is only dealing with the symptoms of the problem and not addressing the root cause. To properly answer these questions, the community needs to address the issue of a shortage of funding and data collection/analysis. Communication and educational outreach programs need to be improved in order raise the awareness of impacted parties and broaden participation and sharing.

4. Security Considerations

Workshop participants discussed security aspects related to root cause analysis and mitigation techniques of potential name collision events. As noted in several papers and presentations, security concerns may both arise and be addressed with name collision mitigation techniques. Follow-on measurement-based research is important to security considerations for name collisions.

5. Informative References

[ADDNOCA] ICANN, "Addendum To Name Collision Occurrence Assessment", November 2014, <http://newgtlds.icann.org/sites/default/files/agreements/name-collision-assessment-addendum-14nov14-en.htm>.

[BLOCKLISTS] Thomas, M., Labrou, Y., and A. Simpson, "The Effectiveness of Block Lists in Preventing Collisions", March 2014, <http://namecollisions.net/program/index.html>.

[COMMENTARY] Kaliski, B., "Proceedings of Name Collisions Workshop Available", March 2014, <http://www.circleid.com/posts/20140326_proceedings_of_name_collisions_workshop_available/>.

[CORPCOM] Strutt, C., "Looking at corp.com as a proxy for .corp", March 2014, <http://namecollisions.net/program/index.html>.

[DITL] Center for Applied Internet Data Analysis, "A Day in the Life of the Internet (DITL)", July 2011, <http://www.caida.org/projects/ditl/>.

[DNS-OARC] Mitchell, K., "DNS-OARC", March 2014, <http://namecollisions.net/program/index.html>.

[DNSENDUSER] Huston, G., "Measuring DNS Behaviors from the End User Perspective", March 2014, <http://namecollisions.net/program/index.html>.

[ENTNETWORK] Hoffman, P., "Name Collision Mitigation for Enterprise Networks", March 2014, <http://namecollisions.net/program/index.html>.

[IAB2008] IAB, "The IAB’s response to ICANN’s solicitation on DNS stability", March 2008, <https://www.iab.org/documents/correspondence-reports-documents/docs2008/2008-03-07-icann-new-gtlds/>.

[IESCPANEL] Woolf, S., Koch, P., Kolkman, O., Kumari, W., and J. Levine, "Internet Engineering and Standards Considerations", March 2014, <http://namecollisions.net/program/index.html>.

[INTERISLE] ICANN, "Name Collision in the DNS", Version 1.5, August 2013, <https://www.icann.org/en/about/staff/security/ssr/name-collision-02aug13-en.pdf>.

[IR2012] ICANN, "Preliminary Report | Regular Meeting of the ICANN Board", September 2012, <http://www.icann.org/en/groups/board/documents/prelim-report-13sep12-en.htm>.

[ISTUDY] ICANN, "Security Studies on the Use of Non-Delegated TLDs, and Dotless Names", May 2013, <https://www.icann.org/en/news/announcements/announcement-28may13-en.htm>.

[JASBUG] Common Vulnerabilities and Exposures, "Group Policy Remote Code Execution Vulnerability", CVE-2015-0008, February 2015, <http://www.cve.mitre.org/cgi-bin/cvename.cgi?name=CVE-2015-0008>.

[JASFRAMEWORK] Schmidt, J., "Name Collisions Management Framework", March 2014, <http://namecollisions.net/program/index.html>.

[KEEPEYE] Schneier, B., "Keeping an Eye on Name Collisions", March 2014, <http://namecollisions.net/program/index.html>.

[MODELING] Deccio, C. and D. Wessels, "What’s in a Name (Collision): Modeling and Quantifying Collision Potential", March 2014, <http://namecollisions.net/program/index.html>.

[MRDNC] ICANN, "Mitigating the Risk of DNS Namespace Collisions: A Study on Namespace Collisions in the Global Internet DNS Namespace and a Framework for Risk Mitigation", February 2014, <https://www.icann.org/en/about/staff/security/ssr/name-collision-mitigation-26feb14-en.pdf>.

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

[NCOMFINAL] ICANN, "Name Collision Occurrence Management Framework", July 2014, <https://www.icann.org/en/system/files/files/name-collision-framework-30jul14-en.pdf>.

[NCRI] ICANN, "Name Collision Resources & Information", <http://www.icann.org/en/help/name-collision>.

[NCSLDCIV] ICANN, "Name Collision SLD Controlled Interruption Variations", September 2014, <http://newgtlds.icann.org/sites/default/files/agreements/name-collision-sld-controlled-interruption-12sep14-en.htm>.

[NEXTSTEPS] Kaliski, B., "Workshop Wrap-Up and Next Steps", March 2014, <http://namecollisions.net/program/index.html>.

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

[NOCA] ICANN, "Name Collision Occurrence Assessment", August 2014, <http://newgtlds.icann.org/sites/default/files/agreements/name-collision-assessment-04aug14-en.htm>.

[RARDBITS] Reid, J., "Analysing the Use of the RA and RD bits in Queries to Root Servers", March 2014, <http://namecollisions.net/program/index.html>.

[RFC1591] Postel, J., "Domain Name System Structure and Delegation", RFC 1591, DOI 10.17487/RFC1591, March 1994, <http://www.rfc-editor.org/info/rfc1591>.

[RFC2606] Eastlake 3rd, D. and A. Panitz, "Reserved Top Level DNS Names", BCP 32, RFC 2606, DOI 10.17487/RFC2606, June 1999, <http://www.rfc-editor.org/info/rfc2606>.

[RFC6761] Cheshire, S. and M. Krochmal, "Special-Use Domain Names", RFC 6761, DOI 10.17487/RFC6761, February 2013, <http://www.rfc-editor.org/info/rfc6761>.

[RFC7719] Hoffman, P., Sullivan, A., and K. Fujiwara, "DNS Terminology", RFC 7719, DOI 10.17487/RFC7719, December 2015, <http://www.rfc-editor.org/info/rfc7719>.

[RSSAC] Murai, J., "RSSAC response to the root scaling report", November 2010, <http://www.icann.org/en/news/correspondence/murai-to-board-25nov10-en.pdf>.

[RSSAC002] ICANN Root Server System Advisory Committee, "Advisory on Measurements of the Root Server System", November 2014, <https://www.icann.org/en/system/files/files/rssac-002-measurements-root-20nov14-en.pdf>.

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

[SAC046] ICANN Security and Stability Advisory Committee, "Report of the Security and Stability Advisory Committee on Root Scaling", SAC 046, December 2010, <https://www.icann.org/en/groups/ssac/documents/sac-046-en.pdf>.

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

[SEARCHLISTS] Simpson, A., "Detecting Search Lists in Authoritative DNS", March 2014, <http://namecollisions.net/program/index.html>.

[TECHNIQUES] Thomas, M. and A. Simpson, "Analysis Techniques for Determining Cause and Ownership of DNS Queries", March 2014, <http://namecollisions.net/program/index.html>.

[WPNC] Verisign, "Workshop and Prize on Root Causes and Mitigation of Name Collisions (WPNC)", June 2014, <http://namecollisions.net/>.

Appendix A. Program Committee

This workshop program committee consisted of Geoff Huston, Burt Kaliski, Olaf Kolkman, John Levine, Allison Mankin, Lixia Zhang, Anne-Marie Eklund Loewinder, and Andrew Sullivan.

Appendix B. Workshop Material

Main Workshop Page: <http://namecollisions.net/>

Name Collision Invited and Submitted Papers, Panels, and Videos: <http://namecollisions.net/program/index.html>

The peer-reviewed papers were:

- "Analysis Techniques for Determining Cause and Ownership of DNS Queries" [TECHNIQUES],
- "Analysing the Use of the RA and RD bits in Queries to Root Servers" [RARDBITS],
- "The Effectiveness of Block Lists in Preventing Collisions" [BLOCKLISTS],
- "What’s in a Name (Collision): Modeling and Quantifying Collision Potential" [MODELING], and
- "Detecting Search Lists in Authoritative DNS" [SEARCHLISTS].

The invited talks were:

- "Keeping an Eye on Name Collisions" [KEEPEYE],
- "Looking at corp.com as a proxy for .corp" [CORPCOM],
- "Measuring DNS Behaviors from the End User Perspective" [DNSENDUSER],
- "DNS-OARC" [DNS-OARC], and
- "Name Collision Mitigation for Enterprise Networks" [ENTNETWORK].

The panels and discussions were:

- "Internet Engineering and Standards Considerations" [IESCPANEL],
- "Name Collisions Management Framework" [JASFRAMEWORK], and

- "Workshop Wrap-Up and Next Steps" [NEXTSTEPS].

Appendix C. Workshop Participants

A list of workshop participants is provided at [WPNC].

Acknowledgments

We would like to thank both the program committee (Appendix A) and the workshop participants (Appendix C), with equal appreciation to those who spoke formally and those who joined in the lively discussions.

Additionally, we would like to thank the following people for their review comments: Burt Kaliski, Olaf Kolkman, Ed Lewis, Nevil Brownlee, Tim Wicinski, and Danny McPherson.

Authors' Addresses

Matthew Thomas
Email: mthomas@verisign.com

Allison Mankin
Salesforce
Email: allison.mankin@gmail.com

Lixia Zhang
UCLA
Email: lixia@cs.ucla.edu