Revised Study 2 Proposal for the Name Collision Analysis Project (NCAP)

22 January 2021

Table of Contents

Executive Summary 2
1. Introduction 2
2. Proposed Changes 3
3. Revised Proposal 5
4. Project Management Matters 6
   4.1 Method of Engagement 6
   4.2 Conflicts of Interest 6
   4.3 Estimated Project Timeline 7
   4.4 Estimated Project Resources 7
Appendix 1 – Table of Board Question vs Study Two 9
Appendix 2 – NCAP Gap Analysis Brief 11
Appendix 3 - Additional Details on Study Two Proposal 13
Appendix 4 - Original NCAP Project Study Two Proposal 18

Status of the Report

This proposal is prepared by the Name Collision Project Analysis Discussion Group (NCAP DG) and approved by the ICANN Security and Stability Advisory Committee (SSAC) on XX 2020.

Executive Summary

This document is a revision to the NCAP proposal originally produced by SSAC in September 2018 for the ICANN Board Technical Committee (BTC). The original proposal details SSAC's proposed approach for studying name collision in response to the ICANN Board's request in resolutions 2017.11.02.29 – 2017.11.02.31.

On the 17th of June 2020, the draft final version of the Study One report was published for public comment. The report on this public comment recommended that Studies Two and Three should "not be performed as currently designed." The NCAP Discussion Group agrees with this assessment and proposes four alterations to Study Two that would address these concerns:

- Removal of the original Study Two Goal of "Building a data repository.".
- Removal of the Study Two Tasks to "Build a test system which can be used for impact analysis and to test possible mitigation strategies.".
- Expansion of the Study Two Task "Conduct an impact analysis." to detail the activities this Task involves.
- Having the NCAP Discussion Group undertake most of the work which was slated for paid contractors in the original version of the Study Two proposal.

The results of these modifications will dramatically reduce the scope, level of effort, total costs, and resources to execute Study Two.

Study Two will undertake to:

- Perform a study of ICANN Collision Reports.
- Perform a Impact and Data Sensitivity Analyses with respect to name collisions.
- Respond to Board Questions Relating to Study Two.
- Produce a final report on Study Two.

It is proposed that these tasks will be conducted over a period of 18 months beginning in January 2021 and ending in June 2022.

1. Introduction

Since the new gTLD program in 2012, the internet's DNS ecosystem has evolved as well as the DNS community's understanding of the more nuanced concerns, threats, vulnerabilities, and underlying causes of name collisions. This evolution has resulted in a gap of substantive data resources and knowledge that should be considered when assessing the risk profile and mitigating controls to deploy with respect to future TLD delegations by ICANN. These changes include at least the following:

1. several new technologies and recommended best practices within the DNS ecosystem now have a significant impact on the volume and quality of DNS queries observed at nameservers in the DNS hierarchy, and

2. various new data sets have become available that may provide additional information to allow for better understanding and assessing of name collision risks.

As previously put forth to the ICANN community, NCAP Study Two was designed to understand the root cause of most of the name collisions and to also understand the impact of any choice made regarding .CORP, .HOME, and .MAIL. However, given the technical and knowledge gap identified by the NCAP Discussion Group as well as the commentary provided by the Study One Final Report, the scope and design of Study Two, and subsequently Study Three, should be re-examined to ensure they are aligned to address and provide advice to ICANN Board's questions in their resolution 2017.11.02-29 – 2017.11.02.30.[1]

This brief puts forth a proposed revised Study Two plan. It is organized as follows. In Section 2, we explain, at a high level, the proposed changes of study 2 and the rationale for these changes. The details of the updated proposal are covered in Section 3 and Section 4.

The document contains several important appendixes. Appendix 1 maps how Study Two, and to some extent Study Three, would answer the ICANN Board's questions in their resolutions 2017.11.02-29 – 2017.11.02.30. Appendix 2 identifies the set of substantive data resources and knowledge gaps between the 2012 round of new gTLDs that motivated the work party to redefine the work tasks and objectives of Study Two to better answer the Board's questions. Appendix 3 contains suggested questions and corresponding data measurements and research tasks to help guide the Study Two impact and data sensitivity analysis of name collision data. Finally, to provide clarity and context, Appendix 4 lists the original scope and tasks of Study Two.

2. Proposed Changes

Learning from the studies that have been published since the last new gTLD round and considering what we know about the technical changes to the DNS and Internet infrastructures, we are proposing the removal of one of the three original goals, the removal of a Study task, and we add substantial detail to one other Study task. The details of the original NCAP Study 2 goals and tasks are described in Appendix 4 of this document.

Study Two Goals:

1. ~~Build a data repository~~
2. Understand the root cause of most name collisions
3. Understand the impact of name collisions

Study Two Tasks:

1. Conduct root cause analysis
2. ~~Build a test system which can be used for impact analysis and to test possible mitigation strategies~~
3. Conduct impact analysis

[1] https://www.icann.org/resources/board-material/resolutions-2017-11-02-en#2.a.rationale

4. Produce a report on the results of Study Two
5. Undertake a formal public consultation on the results of Study Two

Rationale to drop "Build a data repository"
The principal observation to be made is that confidentiality has become a significant overarching concern throughout the Internet, and in Internet standard protocols. This raises significant questions both as to whether or not the data would be accessible at all, as well as the question of whether a centralized location for all such relevant data is practical. Instead we propose a sub-task under the impact analysis to conduct a data sensitivity analysis to consider the question of data availability in a future with increasingly less data available for analysis.

Rationale to drop "Build a test system ..."
Given the research and findings provided in the Study One Final Report dated May 5, 2020, the scope, resource requirements, and overall feasibility of successfully constructing a sustainable, repeatable, all-encompassing test system does not seem achievable. Instead, we suggest this task be incorporated into the Study Three body of work as a tabletop exercise.

Additional details for root cause analysis (Study Two Task 1)
Study Two Task 1 is to conduct a root cause analysis. We propose for this to be a review, study, and detailed analysis of all name collision reports that ICANN has received via its portal.[2] To complete this task it will be necessary to contact the reporter to obtain as much information as possible about what happened, how it was discovered, and how it was mitigated. Additional investigation should be conducted as needed to ensure as complete an analysis as possible. This will be used as part of Study Three to review possible mitigation strategies.

Additional details for impact and data sensitivity analysis (Study Two Task 3)
The bulk of the work in Study Two is to conduct an impact and data sensitivity analysis. We propose the following sub-tasks as described below. Additional details and specifics are provided in Appendix 3.

- Using the similar data sources and methodologies by JAS Global Advisors[3] and Interisle Consulting Group,[4] perform updated case studies of the CORP, MAIL, HOME, and other strings[5]. The study should highlight changes over time of the properties of DNS queries, and traffic alterations as a result of DNS evolution.

- Perform a data sensitivity analysis to (1) identify the minimum data requirements for analysis to allow the Board to make decisions about Collision Strings per its questions to the SSAC; and (2) make recommendations regarding best practices for data handling and processing.

---

[2] https://forms.icann.org/en/help/name-collision/report-problems

[3] https://www.icann.org/news/announcement-2-2015-11-30-en

[4] www.icann.org%2Fen%2Fsystem%2Ffiles%2Ffiles%2Fname-collision-02aug13-en.pdf&usg=AOvVaw12vKxcxpqOeU33Vy_Hgm2e

[5] Using this threshold and DNS query data from A and J root servers, this will result in six strings (after the explicit inclusion of .mail): .local, .home, .internal, .lan, .corp, and .mail.

Additional details for "Produce a report on the results of Study Two" (Study Two Task 4)
The Study Two report will contain (1) the results of the root cause analysis, (2) the results of the impact and data sensitivity analysis, and (3) answers from the NCAP discussion group regarding the following Board questions, based on the data of these studies:

- the role that negative answers currently returned from queries to the root for these strings play in the experience of the end user, including in the operation of existing end systems;
- the harm to existing users that may occur if Collision Strings were to be delegated, including harm due to end systems no longer receiving a negative response and additional potential harm if the delegated registry accidentally or purposely exploited subsequent queries from these end systems, and any other types of harm;
- possible courses of action that might mitigate harm;
- factors that affect potential success of the courses of actions to mitigate harm; and
- potential residual risks of delegating Collision Strings even after taking actions to mitigate harm;

The NCAP Discussion Group has held numerous discussions about the Board questions and how to best answer them. These details are documented in Appendix 3.

3. Revised Proposal

The revised proposal reduces the number of Study Two goals down to two and reduces the number of Study tasks from five to four. It also provides additional detail about the study tasks to facilitate the development of a statement of work suitable for seeking the engagement of a contractor.

The proposed Study Two goals include the following:

1. Understand the root cause of most name collisions
2. Understand the impact of name collisions

The proposed Study Two tasks include the following:

- Study of ICANN Collision Reports
    - Perform an analysis of ICANN Collision Reports to determine the underlying cause of these collisions.
    - Produce a report on the results of the analysis.
    - Performed by: Technical Investigator

- Impact and Data Sensitivity Analyses
    - Research the impact of collisions with regards to Root servers and Resolvers for CORP, HOME and MAIL.
    - Research the impact of collisions with regards to Root servers and Resolvers for other selected strings.

    - Based on the above research, evaluate the effectiveness of using multiple sources of collision data with regards to the assessing the impact of collisions.
    - Undertake a public consultation on the findings relative to .CORP, .HOME and .MAIL.
    - Produce a report on the results of this work.
    - Performed by: Discussion Group (DG) and Technical Investigator (in the capacity guided by the DG / Admin team).

- Response to Board Questions Relating to Study 2
    - Respond to Board questions based on the results of the Study of ICANN Collision Reports and Impact and Data Sensitivity Analyses.
    - Produce report on the responses to Board questions.
    - Performed by: DG

- Final Report
    - Produce the final report for Study Two
    - Undertake a public consultation on the draft version of this report.

4. Project Management Matters

4.1 Method of Engagement

Both the original SSAC project proposal and the OCTO revised project proposal had the same Section 3.1 that identified several teams that would be used to fulfill the Board's requirement that the studies be conducted in a thorough and inclusive manner that includes technical experts (such as members of the IETF working groups, technical members of the GNSO, and other technologists). No changes are proposed for these teams, their membership, or their working methods. Specifically:

- All teams have mailing lists in support of engagement. The Discussion Group mailing list archive is publicly available.
- All teams will meet weekly via virtual means to progress the work of the NCAP project. To observe or participate in the meeting, one must join the discussion group.
- The Discussion Group will host 2 one day in-person meetings on the day before the official start of an general ICANN in-person meeting. The meeting will be open to observers.

4.2 Conflicts of Interest

The original NCAP proposal included Section 3.2 "Conflicts of Interest" to deal with three issues:

- Controlling the access to confidential information supplied by operators in the proposed data repository.
- Ensuring impartiality of contracted resources performing data analysis to present to the NCAP WP and NCAP DG.
- Restricting the access to financial and resource requirements for the project prior to RFPs being published.

With respect to these COI issues, it should be noted that the revised Study 2 proposal:

- No longer includes a data repository and thus eliminates this confidentiality concern.
- Moves the data analysis from being performed by contractors to being performed by the NCAP Discussion Group and having a contractor only document the Discussion Group analysis, thus eliminating the COI concerns associated with the hiring of contractors by ICANN to perform the analysis.
- Maintains the restricted access to financial and resources requirements for the NCAP project to the NCAP Admin Committee.

It is important to note that the SOI requirements presented in section 3.1 of the original proposal, including the NCAP specific questions, are maintained.

4.3 Estimated Project Timeline

The project is scheduled to start in January 2021 and end in June 2022.

The following graphic illustrates the timeline per project component:

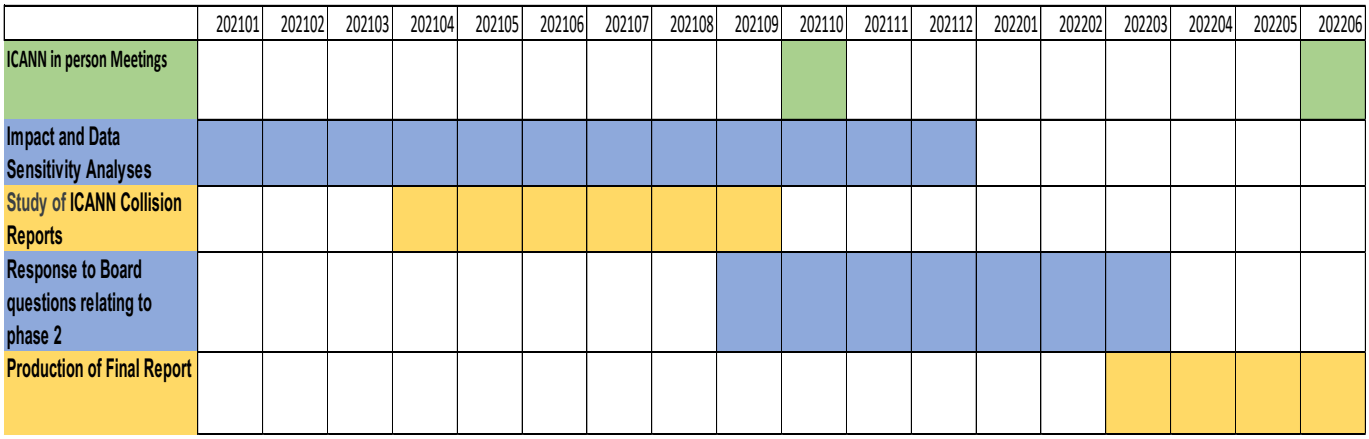

4.4 Estimated Project Resources

To reduce the cost required to execute the revised Study Two tasks, NCAP Discussion Group members are responsible for obtaining access to DNS data or writing code to measure various aspects outlined in Appendix 3. They are also responsible to perform those measurements and to present their findings to the group.

The following personnel resources are required for the study:

- Technical Investigator - This person will take all the reports (approximately 40) of name collisions that ICANN has received at its portal and produce a work product that is a root cause analysis that answers at least the following three questions for each incident as completely as possible.
    - What happened? How was it detected? This should include reviewing the prior collision studies and looking for any signal that might have been missed.
    - How was the issue resolved? What else was considered and rejected?
    - What lessons were learned from this experience? Are there any open questions or unresolved consequential issues? Based on this analysis, would collision prevention mechanisms, such as controlled interruption, be effective or could other collision prevention mechanisms be deployed/implemented to better prevent the collision.

  Assist as required with the work for Impact and Data Sensitivity Analyses. This will likely require data analysis skills of DNS data stored in data repositories, such as the DNS-OARC DITL
- A Technical Writer – Prepare various reports, including the final Study Two report, as well as the two public consultation documents and webinar and other Power Point type documents.
- A Project Manager to work with the discussion group to develop and execute the detailed project timeline and deliverables.
- A Project Secretary that manages all aspects of the logistics of the discussion group operations, including teleconferences and in-person meetings.

In addition, resources are needed to support 2 one day in person meetings on the day before the official start of an ICANN in-person meeting. The support and resources include:

- Logistics to include the meeting room, media services in the room, remote participation availability, and hosted breaks and lunch
- An additional hotel night for arrival the day before for all attendees who are otherwise supported
- Travel support for up to 5 attendees selected by the NCAP Admin Committee to attend the in-person meetings
    - Support only available for those not otherwise supported, i.e., this is separate from the additional night that should be included for those already supported per item 2.b above and it may not be used in lieu of support available elsewhere
    - Includes airfare, hotel, and other expenses according to ICANN Travel Guidelines
    - Includes support for the entire week of the ICANN meeting

Appendix 1 – Table of Board Question vs Study Two

The matrix below enumerates the set of ten questions put forth by the Board and maps each question to corresponding proposed Study Two and Three tasks. The NCAP Discussion Group believes that the revised Study Two and Three tasks will provide enough information to create appropriate guidance to the Board's questions.

The Board's first question was already answered as part of the final Study One report deliverable. The remaining unanswered questions are mapped to specific work items in the proposed Study Two and Three plans. Board Questions two through six heavily rely on the proposed Study Two task of "conducting an impact analysis." Details of the impact analysis are included in Appendix 3, which outlines various research questions, measurements, and other investigative / tabletop tasks.

Board Questions seven, eight, and nine are dependent on the final reports of Study Two and Study Three. Providing guidance to these three questions will require the culmination of qualitative and quantitative inputs undertaken in Studies Two and Three as well as their final reports.

Finally, the revised Study Two proposal now recommends an explicit case study be conducted during the impact analysis for the .CORP, .MAIL, and .HOME strings so guidance and advice can be given to the Board's final question.

Board Questions	**Study Two Tasks**
(1) a proper definition for name collision and the underlying reasons why strings that manifest name collisions are so heavily used;	Completed during Study One but subject to revision according to analysis in Study Two
(2) the role that negative answers currently returned from queries to the root for these strings play in the experience of the end user, including in the operation of existing end systems;	Conduct impact analysis
(3) the harm to existing users that may occur if Collision Strings were to be delegated, including harm due to end systems no longer receiving a negative response and additional potential harm if the delegated registry accidentally or purposely exploited subsequent queries from these end systems, and any other types of harm;	Conduct root cause analysis Conduct impact analysis
(4) possible courses of action that might mitigate harm;	Conduct root cause analysis Conduct impact analysis

	● Study Three Tasks to follow
(5) factors that affect potential success of the courses of actions to mitigate harm;	● Study Three Tasks to follow
(6) potential residual risks of delegating Collision Strings even after taking actions to mitigate harm;	Conduct impact analysis ● Study Three Tasks to follow
(7) suggested criteria for determining whether an undelegated string should be considered a string that manifest name collisions, (i.e.) placed in the category of a Collision String;	Produce a report on the results of Study Two
(8) suggested criteria for determining whether a Collision String should not be delegated, and suggested criteria for determining how remove an undelegated string from the list of Collision Strings; and	Produce a report on the results of Study Two
(9) measures to protect against intentional or unintentional creation of situations, such as queries for undelegated strings, which might cause such strings to be placed in a Collision String category, and research into risk of possible negative effects, if any, of creation of such a collision string list. (10) to present data, analysis and points of view, and provide advice to the Board regarding the risks posed to users and end systems if .CORP, .HOME, .MAIL strings were to be delegated in the root, as well as possible courses of action that might mitigate the identified risks.	Produce a report on the results of Study Two ● Study Three Tasks to follow Produce a report on the results of Study Two

Appendix 2 – NCAP Gap Analysis Brief

The Study One report of the Name Collision Analysis Project (NCAP) provides a concrete definition of the term "name collision" and serves as a summary report on the topic in which it brings forth important knowledge from prior work in the area. While some name collision research was conducted in the years prior to the new gTLD program in 2012, the field was and still remains an esoteric field of cybersecurity research. However, over the course of the last decade numerous peer reviewed academic proceedings and industry reports have been published that highlight the more nuanced concerns, threats, vulnerabilities, and underlying causes of name collisions within the DNS. Furthermore, the internet's DNS ecosystem has evolved since the previous round of TLD delegations to a state in which there is more nameserver consolidation as well as protocol bifurcation and alterations that will directly impair the observational capacity to conduct name collision risk assessments. To that end, there is a gap of substantive data resources and knowledge between the 2012 round and now that should be considered when assessing the risk profile and mitigating controls to deploy for future TLD delegations by ICANN.

This brief serves as an initial foundation to highlight major areas that should be considered to address the knowledge and data gaps through subsequent studies in NCAP Study Phases 2 and 3. Those studies will incorporate this knowledge and data sources that were not utilized to quantitatively or qualitatively assess name collision risks in the 2012 program and help provide guidance to ICANN Board's questions in their resolution 2017.11.02-29 – 2017.11.02.30. A non-exhaustive list of gaps has been identified by the participants of the NCAP Discussion Group and roughly organized into the following major categories:

1. Data Sets: Since the new gTLD program, various new data sets have become available that may provide additional telemetry to better understand and assess name collision risks. The new gTLD name collision risk assessment was conducted against a few years of Day In the Life of the Internet (DITL) DNS traffic data. Unfortunately, the DITL data set has several limitations, as it only provides a few days per year of authoritative root server DNS traffic, is contributed by root server operators on a voluntary basis, and may be anonymized due to privacy concerns. Since the last TLD round, the collection of DITL data has continued and may provide better longitudinal measurements pre/post the new TLD delegations. Other entities have also started to retain high fidelity root DNS traffic that may provide better insights. The emergence of popular open recursive resolvers has also transpired and dramatically shaped the DNS ecosystem since the new gTLD delegations. These recursive services may provide a richer and more complete understanding of name collisions if they can be utilized for analysis. Other potential data repositories of interest would also include the ORDINAL DNS data as well as Certificate Transparency records, neither of which existed during the previous assessment.

2. General DNS Evolution and Observational Impairments: DNS usage monitoring provides insight into time-resolved traffic evolution patterns useful in the quantification of system stability and performance as well as detecting aberrant events. Longitudinal measurements and usage trends, however, are increasingly difficult to leverage as the underlying system evolves or as bifurcation within the system occurs. These system changes may result in non-symmetric system usage, partial or even total impairments in DNS measurements, and ultimately confound the interpretability of the system's usage

metrics. Since the last round of TLD delegations, several new technologies and recommended best practices within the DNS ecosystem now have a significant impact on the volume and fidelity of DNS queries observed at nameservers in the DNS hierarchy. These technologies include running Root on Loopback (RFC 7706), Aggressive Use of DNSSEC-Validated Cache (RFC 8198), DNS Query Name Minimization (RFC 7816), and DNS Queries over HTTPS (RFC 8484). It is in the DNS community's best interest to develop a better understanding of how these standards and technology changes will influence data collection capabilities as well as their impacts to data analysis of DNS traffic in an ever evolving, technologically fragmented, and highly distributed system.

3. Controlled Interruption Efficacy and Data Analysis: While the NCAP Study One Report highlights some anecdotal reports around the efficacy of Controlled Interruption, a thorough assessment of the framework has yet to be started. The collected reports should at a minimum be analyzed to better understand any trends, commonalities, faulty assumptions, and success attributes. Understanding the nature of these reports with a re-examination of previous DITL data may help identify key signals in the DNS that could better inform name collision risk assessments moving forward. Some applications, including popular browsers, have implemented specific DNS controls to signal when Controlled Interruption events occur. To that end, efforts should be made to identify and contact such vendors to see if instrumentation data is available. Finally, a study should be made to provide evidence that Controlled Interruption was a successful mitigation model, which may include creating and running simulation test beds.

4. Vulnerability Understanding and Mitigation Strategies: Since the last delegation of TLDs, various peer reviewed academic and industry papers have been published that elucidate some of the more detailed nuances of name collisions, specifically as they relate to various risks and vulnerabilities. Specifically, many of these publications directly identify known DNS query patterns, typically associated with zero-configuration protocols such as DNS-SD, that can be weaponized and exploited in a name collision environment. This new knowledge should be applied to future TLD delegation risk assessments as it builds upon a foundational understanding of the intent of the DNS queries as opposed to the volume of queries that was originally used in the new gTLD risk assessment.

Appendix 3 - Additional Details on Study Two Proposal

The sections below contain additional details on Study 2. It includes suggested questions and corresponding data measurements and research tasks to help guide the Study Two impact analysis of name collision data.

Additional Details and Rationale For DNS Traffic Evolution – Re-examining 2012 and measuring present day

The community established a set of criteria in the Interisle and JAS reports for determining the risk of various strings based on data observed (primarily) using root operator data analyzed on the DNS-OARC Day-in-the-Life (DITL) data repository. These criteria focused on measuring occurrence rates and estimating the severity of the consequences based on contextual information within the query data.

Given the evolution of the DNS infrastructure and protocol, how would those risk assessments look today? Alterations in the DNS ecosystem, such as recursive resolver consolidation, qname minimization deployment, NSEC caching, will directly influence the quality and availability of the data. To what extent has that occurred and how does it impact our ability to answer the Board's questions?

We propose that we do an initial assessment of data sets available in 2012 vs those currently available to identify changes in metadata. The identification of these changes will help establish a baseline from which evaluations of risk on a per string basis can be performed.

The re-examination of 2012 and measurement of present-day DNS traffic should be conducted in a well-scoped and limited capacity to reduce costs, computational demands, and analytical efforts. To that end, Study Two data questions should be scoped, when appropriate, to only the following strings:

1. Detailed case studies of the CORP, MAIL, and HOME strings that highlight changes over time of the properties of DNS queries, traffic alterations as a result of DNS evolution, and mitigation mechanisms tailored to these strings.

2. Top N strings from the current 2020 data based on a threshold of a string receiving more than 100 million queries per day at the root. Using this threshold and DNS query data from A and J root servers results in six strings:.local, .home, .internal, .lan, .corp, and .mail.

Additional Details on Sensitivity Analysis

The goal of the Impact and Data Sensitivity analysis is to (1) identify the minimum data requirements for analysis to allow the Board to make decisions about Collision Strings per its questions to the SSAC; and (2) make recommendations regarding best practices for data handling and processing.

Among the questions to be considered regarding minimum data requirements are the following:

- Is root server data sufficient?
- Is DITL data sufficient?
- Is resolver data required?
- Is there data an application for new gTLD could provide that would be helpful?
- What is the time window of data needed?
- What is the quantity of data needed?

Among the questions to be considered regarding data handling and processing requirements are the following:

- What are the privacy considerations?
- What are the access considerations?
- What are the data retention considerations?

General Questions to be answered by the sensitivity analysis

- To what extent is root and resolver data generally available for analysis?
- What constraints from a terms-of-use or Personally Identifiable Information data anonymization impact data fidelity?
- How sensitive is risk analysis when using subsets of data for future assessments (e.g. using only one root vs a full DITL collection)?

DNS Traffic Characteristics Questions

- To what extent have query volumes changed?
    a. Overall query volume, NX Domain percentage, qtypes.
- Has traffic expanded to a broader distribution or consolidated to major recursive players?
- Have the names being queried changed significantly?
    a. Distinct TLD strings, number of labels, first labels and contextual identifiers with TLD string, etc.
- Is there apparent misuse or gaming of name collision data?

Additional Details Relative to Answering Board Questions

Below are questions, thoughts, and research ideas presented by the Discussion Group that will form the foundation of answering the Board's questions.

Board Question 2: the role that negative answers currently returned from queries to the root for these strings play in the experience of the end user, including in the operation of existing end systems;

- How has application logic evolved to depend on DNS (while being cognizant many legacy systems that are not well understood still persist)?
- Are there examples of new technologies that take advantage of NXDOMAIN?
- Can specific systems or trends be identified by looking into the data to find new software that relies on non-delegated strings?
- Why are people and systems still explicitly relying on non-delegated strings?

- What advice can be given to people so that maybe they'll behave better?

Board Question 3: the harm to existing users that may occur if Collision Strings were to be delegated, including harm due to end systems no longer receiving a negative response and additional potential harm if the delegated registry accidentally or purposely exploited subsequent queries from these end systems, and any other types of harm;

- Likely to be a tabletop exercise on our part, extrapolating from what we know and name collisions that have occurred.
- Might be dependent on types of collision and mitigation, i.e., perhaps there is a mitigation framework that would be helpful.
- What is "harm"? Does it imply physical? Cyber? Reputational? Or is it compromised credentials, systems, or data? The connotation of "harm" may include numerous things making it difficult to appropriately apply scale and context to this otherwise broad term within the scope of name collisions.
- We propose the following broad categories based on our analysis of the literature and data reported:
    - Interception and Manipulation: Private DNS, or similar, queries leaking into the public DNS that were previously answered by the root servers can be subsequently received and answered by various parties, either purposefully or unknowingly, after the delegation of a TLD string. In such a scenario, an attacker's exploitation of name collisions will allow them to intercept and manipulate DNS queries. Through these name collision events, attackers may capitalize on a variety of passive and active attack vectors including reconnaissance/enumeration, MitM attacks, internal or personal document leakage, malicious code injection, and credential theft. Some of these attack vectors and corresponding risks stem from DNS-SD or zero-configuration protocols that utilize the DNS as a bootstrapping mechanism. Coupling those protocols with either intentional rooting of a namespace in an undelegated TLD or through unintended consequences of suffix search lists, these types of queries are often the most exploitable attack vector in a name collision scenario.
    - Signaling Interruption: This is likely a spillover of Board question #2 that discusses the role played by negative answers currently returned from queries to the root. Some things that come to my mind would be breakage of applications that utilize the DNS as a signaling tool rather than as a directory (e.g. Chrome startup, Mozilla DoH, etc.). These situations again are likely due to search list processing. Do we want to talk about the impacts of signal changes when controlled interruption is deployed or the TLD is delegated (with registrations)? For example, how a browser would change its user displayed error message from something like "Domain not found: NXDOMAIN" to something around "Cannot connect to…." Another scenario is one in which conditional logic of the returned DNS answer is baked into the application and can be handled in many different ways making it difficult/impossible to assess/track/remediate/etc. (e.g., Mozilla encoding of 127.0.53.53 into their DoH logic within the application).

Board Question 4: possible courses of action that might mitigate harm.

- Thought exercise on our part extrapolating from name collisions that have occurred.
- Dependent on classes of collision types and mitigation types, i.e., perhaps there is a mitigation framework that would be helpful.
- Why is mitigating name collisions difficult?
    - Are organizations even able to "see the problem" (e.g., transient corporate devices used on corporate networks) or even be able to reliably "trace the causes"
- Some reasonable mitigation plans:
    - Organizations using a private TLD, change it to use ones rooted in the global DNS.
    - If using a shortened name, ensure use of fully qualified domain names in various systems.
- Targeted Outreach
    - If the applied string has certain traffic properties direct outreach to the underlying manufacturer or ISP may be sufficient to remediate the issue (e.g., TELUS, CONSUL, CBA, etc.)
- SLD Blocklist (e.g., used from snapshots of DNS data)
    - Various research reports show that statistical sampling is flawed using this approach due to time, root server affinities, etc.
    - Provides blueprint to miscreants for domains with elevated traffic (and potentially higher risk profiles)
- Mitigation Strategies
    - Underlying causes of colliding strings likely requires various strategies to effectively mitigate (or inform) end systems/users.
    - Fail hard scenarios: Database connectivity, etc. Events in which an application explicitly requires a connection to one or more services and places corresponding exception handling processes to properly raise errors.
    - Systems designed to keep users unaware of actions: DNS-SD and Zero Configuration protocols. Service configuration is done via DNS and facilitates various MitM attacks if performed surreptitiously.
- Guidance to consider:
    - It's important to keep in mind that we probably won't be able to list all possible mitigation strategies, especially going forward. The advice that would be most helpful to the Board is how to evaluate mitigation strategies and considerations regarding who is responsible for the mitigation.
    - What are the parameters of a good mitigation strategy?
    - What are the parameters for measuring the success of a mitigation strategy?

Board Question 5: factors that affect potential success of the courses of actions to mitigate harm:

- This task will likely be highly dependent on the risk analysis of the items identified in Board question four.
- This will also be influenced by what we learn from name collisions that have occurred (e.g., an analysis of the ICANN name collision reports and retrospective analysis of DNS telemetry data).

Board Question 6: potential residual risks of delegating Collision Strings even after taking actions to mitigate harm:

- Risk analysis of mitigations taken as a result of known collisions is likely to influence our response.
- What is the effect of time on mitigation, i.e., does risk go up or down over time after mitigation has been applied?
- With most of the mitigation efforts, despite taking the mitigation actions, there is a risk that the collisions will still occur (lack of attention? Lack of realization that mitigation steps have been taken. Or Lack of caring).

Additional Details on Data Sets

The following data sets are suggested candidates to be used in the analysis tasks of Study Two. Their availability, licensing, and processing-capacity may render some or some subset of them as unusable. To that end, as part of Study Two's "conduct an impact analysis", the NCAP Discussion Group has recommended a sensitivity analysis be conducted on the data sets that are available. It is unclear at this time as to what data will readily be accessible and what frequency ICANN will utilize these types of data sets to evaluate strings in the subsequent rounds. Understanding the limitations of these data sets, specifically when used independently, is critical to help provide guidance to the Board's questions.

1. Root data: Root server traffic has served as the de facto standard for most name collision research. Continued use of this data, namely through the DNS-OARC data repository of DITL data, will likely be the primary data resource for Study Two.

2. Recursive resolver data: What additional telemetry or fidelity is gained when examining recursive resolver data. This will likely help provide insights to the differences of traffic with respect to FQDNs, inter-query timing, source diversity, when contrasted to the root. This data may better portray the end-user population and systems issuing queries and provide better insights into the root causes of name collisions and more effective mitigation strategies that could be deployed.

3. ICANN collision reports: While ICANN collision reports are not expected to be used in subsequent string evaluations, they are relevant inputs to Studies Two and Three. Retrospective analysis of DNS traffic data searching for "missed" signals within the 2012 data that aligns with characteristics of the reported collision (specific names, source IP addresses, etc.). This data will likely provide insights to the Board's second question but also inform other measurement tasks for Board questions three through six.

Appendix 4 - Original NCAP Project Study Two Proposal

The original NCAP Project proposal[6] had three goals for Study Two:

1. Build a data repository
2. Understand the root cause of most name collisions
3. Understand the impact of name collisions

The first goal, building the data repository, was pushed forward into Study Two from the original SSAC NCAP Project Proposal when the details of Study One were revised as part of launching that work.[7] Study One was revised to focus on creating a bibliography of all published works related to name collisions. A review of the published works was conducted to document data sets used and identify potential gaps in data sets or data providers that would be necessary to successfully complete Studies Two and Three.

The final SSAC proposal, before Study Two was revised, had the following 5 Study Tasks:

1. Conduct root cause analysis
2. Build a test system which can be used for impact analysis and to test possible mitigation strategies
3. Conduct impact analysis
4. Produce a report on the results of Study Two
5. Undertake an informal public consultation on the results of Study Two

These tasks were augmented with 10 additional tasks related to the building of the data repository, drawn from the final SSAC proposal's Study One. Here is the list of those 10 tasks for reference.

1. Develop rules regarding any datasets collected. This will need to consider:
    a. Anonymization of data to comply with privacy laws
    b. Protection of data submitted under confidentiality provisions
    c. Defining data retention policies
    d. Determining whether instrumentation for performing the data analysis should be made available for public use
2. Develop agreement for obtaining data.
3. Create a data register which logs the source of datasets, the date or period over which the data was collected and key identifying features
4. Create a common data repository where the data can be stored and processed efficiently and, if necessary, confidentially

[6] This is OCTO's version of the NCAP Project Proposal with all financial detail redacted. It is based on the SSAC version.

[7] Statement of work for Study One: https://community.icann.org/display/NCAP/NCAP+Working+Documents?preview=%2F79437474%2F111387705%2FNCAP+Study+1+30+May+Proposal.pdf.

5. Develop a set of guidelines for data depositors on how they can sanitize their data, removing all unnecessary information, while still allowing all the expected analysis. This may be in the form of levels of sanitization with guidelines for each level
6. Develop code to implement the data sanitization guidelines on common DNS data capture formats.
7. Gather data from past studies.
8. Confirm data gap analysis produced in Study One and documented in Study One final report
9. Define additional datasets or data providers that are needed
10. Gather new data