

DRAFT SSAC/NCAP Study 2 Report

A Report from the Security and Stability Advisory Committee

Preface

In this document the Security and Stability Advisory Committee (SSAC) {TBD}.

The SSAC focuses on matters relating to the security and integrity of the Internet's naming and address allocation systems. This includes operational matters (e.g., pertaining to the correct and reliable operation of the root zone publication system), technical administration matters (e.g., pertaining to address allocation and Internet number assignment), and registration matters (e.g., pertaining to registry and registrar services). SSAC engages in ongoing threat assessment and risk analysis of the Internet naming and address allocation services to assess where the principal threats to stability and security lie, and advises the ICANN community accordingly. The SSAC has no authority to regulate, enforce, or adjudicate. Those functions belong to other parties, and the advice offered here should be evaluated on its merits. SSAC members participate as individuals, not as representatives of their employers or other organizations. SSAC consensus on a document occurs when the listed authors agree on the content and recommendations with no final objections from the remainder of the SSAC, with the exception of any dissenting opinions or alternative views that are included at the end of the document.

Table of Contents

Executive Summary	4
1 Introduction	5
1.1 Background	5
1.2 Methodology	5
1.3 Terminology	5
2 Advice to the ICANN Board	5
3 Study Reports	6
3.1 Case Study of Collision Strings	6
3.2 Data Sensitivity Analysis	6
3.3 Root Cause Analysis	7
4 Board Questions	7
5 General Recommendations	8
5.1 Decision Tree	8
6 Conclusion	9
7 Acknowledgments, Statements of Interest, and Dissents, Alternative Views and Withdrawals	10
7.1 Acknowledgments	10
7.2 Statements of Interest	10
7.3 Dissents and Alternative Views	10
Appendix A: Case Study of Collision Strings	11
Appendix B: Data Sensitivity Analysis	11
Appendix C: Root Cause Analysis	11

Executive Summary

TBA

1 Introduction

The Name Collision Analysis Project (NCAP) Study 2 final report brings together the research and analysis of several studies that touch key facets around the issue of name collision. The report is structured in such a way as to walk the reader through the methodology and findings of the three research studies. The conclusions from those studies in turn supports the answers to the questions the ICANN Board has asked the NCAP to respond to and the recommendations on how to proceed when handling the potential for name collision in the DNS.

This first section describes the background of the NCAP and the mandate set forth by the ICANN Board in 2017. It goes on to describe the methodology of the study group as a whole, including the timeline of research, community outreach, and study group consensus. The study group found several points where they needed to establish a working group definition for several terms; the results of those discussions is a terminology section that makes clear how those terms are used in this report.

While this report is primarily intended as input to the ICANN Board, all parties interested in the future expansion of the gTLD space, from applicants to community groups, may find the material relevant to their efforts.

1.1 Background

The work behind the NCAP began a decade ago with [SAC057: SSAC Advisory on Internal Name Certificates](#), wherein the SSAC referred to the issue of "name collision" and provided the ICANN Board with steps for mitigating the issue. On 18 May 2013, the ICANN Board adopted resolution [2013.05.18.09 – 2013.05.18.11](#), regarding SAC057, commissioning a study on the use of TLDs that are not currently delegated at the root level of the public DNS in enterprises. From there, the work continued to evolve as the understanding regarding the depth and breadth of the issue grew.¹

Moving ahead to 2017, the ICANN Board requested via [resolutions](#) (2017.11.02.29 - 2017.11.02.31) that the ICANN Security and Stability Advisory Committee (SSAC) conduct studies to present data, analysis and points of view, and provide advice to the Board on the topics around DNS name collision. In response, SSAC formed the Name Collision Analysis Project.² This project is organized into three studies. The [first study](#), which provided a primer on the topic of name collision and a list of datasets that either existed at the time of the study or would need to be generated to support further analysis, was finalized on 19 June 2020 after submission to the Board and a public consultation period.

Whereas Study 1 established baseline documentation, the scope of Study 2 was revised based on new information regarding the DNS (e.g., the use of new transport protocols for the DNS, the existence of new data sets). The revised scope was described in "SSAC2021-02: Revised Study

¹ "History of the Name Collision Analysis Project," <https://community.icann.org/display/NCAP/History+of+the+Name++Collision+Analysis+Project>.

² "Invitation: Name Collision Analysis Project (NCAP) Discussion Group," <https://www.icann.org/en/announcements/details/invitation-name-collision-analysis-project-ncap-discussion-group-17-4-2019-en>.

Two Proposal for the Name Collision Analysis Project" and accepted as per [resolutions](#) 2021.03.25.11 – 2021.03.25.14. The revised scope focused on four key areas:

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

The third study, to understand the ramifications of name collisions with an ultimate goal of developing concrete guidance on how to avoid negative consequences, will be informed by the work from studies one and two.

1.2 Methodology

With the acceptance of the revised Study 2 proposal, the discussion group kicked off the study reports as described in Section 3. Study Reports and settled into a regular meeting cadence. While the discussion group considered the questions assigned by the ICANN Board, the researchers collected and analyzed available data relevant to understanding both how to observe and how to measure the impact of name collision. Each effort required coordination to make sure the Board questions were supported by the findings of the study reports, and that the study reports stayed in scope with the overall mandate for the group.

ICANN org provided administrative support for the NCAP, including project management and technical writing resources. [ICANN org also funded the research analyst(s)]

The discussion group chairs called for consensus on the responses to the Board questions, the study reports, and any special terminology after discussion on each item was concluded during the regular conference calls.

Throughout the NCAP study efforts, the study group has considered collisions at level below the TLD level and alternate roots as out of scope.

[All study reports went out for public consultation prior to their being used in this report to finalize the findings and recommendations to the ICANN Board.]

Item	Final Consensus Call Date	Result
Case Study of Collision Strings		[Full Partial None]
Data Sensitivity Analysis		[Full Partial None]
Root Cause Analysis		[Full Partial None]
Final Report		[Full Partial None]

1.3 Terminology

- Collision String - Board term; need to use or map to our preferred term
- Day-In-The-Life (DITL) - a large-scale data collection project initially undertaken every year since 2006. This data has historically been the primary measurement asset for name collision studies.
- Name Collision - (used in Study 1 and RFP) Name collision “refers to the situation where a name that is defined and used in one namespace may also appear in another. Users and applications intending to use a name in one namespace may attempt to use it in a different one, and unexpected behavior may result where the intended use of the name is not the same in both namespaces. The circumstances that lead to a name collision could be accidental or malicious.

Study 1 concerns name collisions in the context of top-level domains (TLDs), where the conflicting namespaces are:

- the global Internet Domain Name System (DNS) namespace reflected in the root zone overseen by the Internet Assigned Numbers Authority (IANA) Function; and
- any other namespace, regardless of whether that other namespace is intended for use with the DNS or any other protocol.”
- Root Server Identity (RSI) - thirteen identities, each of which is named with the letters ‘a’ to ‘m’, collectively administered by twelve root server operators. They are authoritative for the ‘root-server.net’ domain.
- Controlled Interruption - (From [FAQ](#)) “Controlled interruption is a method of notifying system administrators who have configured their networks incorrectly (knowingly or unknowingly) of the namespace collision issue, and helping them mitigate potential issues.”
- Enhanced Controlled Interruption
- Critical Diagnostic Measurement
- Source Diversity
- Query Volume
- Impact???
- Data Sensitivity
- Reconnaissance/enumeration
- MitM attacks (Man in the Middle attacks may need subgroups)
- DocumentInternal document leakage
- Personal document leakage
- Malicious Code Injection
- Credential Theft

2 Advice to the ICANN Board

[“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.” – Does our advice go beyond this original mandate re: just those gTLDs? Are we saying that our advice is generally applicable?]
[What exactly will ICANN need to do? Tool development to support enhanced controlled interruption? Process changes? Policy changes?]

3 Study Reports

As described in its revised scope, the NCAP Study Group 2 conducted three studies:

- Case Study of Collision Strings
- Impact and Data Sensitivity Analysis
- Root Cause Analysis

Each study offered several insights into how to look for and understand the impact of name collisions.

The first study report, the Case Study of Collision Strings, helped define all the critical diagnostic measurements required to identify name collisions and further, how to assess the impact of a name collision.

The second study report, the Data Sensitivity Analysis, considered if and how the available data sets from both individual root servers and global public resolvers were representative of the overall picture of the DNS queries that would help identify name collisions.

The third and last report, the Root Cause Analysis, considered known name collisions and evaluated what mitigation or remediation happened, particularly what they did and why. The following sections describe the results of those studies in greater detail; the full report for each is available in Annexes A through C at the conclusion of this report.

3.1 Case Study of Collision Strings

The NCAP discussion group met over the course of approximately two years to evaluate and consider questions posed by the ICANN Board on the delegation of currently reserved TLDs such as .CORP, .HOME, and .MAIL. The group undertook a review of past studies and literature, and conducted its own analysis from two root server identities. The result of that review is a modern picture of the impact and potential harm due to name collisions with the undelegated names under study. The analysis provides a sufficient basis from which to draw a number of important findings. Among these include the observation that queries for these undelegated names are both increasing in volume and diversity. These facts suggest that challenges relating to impact and mitigation are also increasing. The group also identified a number of critical diagnostic measurements that help determine the scope, impact, and potential harm of name collisions.

[Hypothesis/Goal: 1) defined all the critical diagnostic measurements, gave them a name and created the set. For us, name collisions are defined by the existence of those measurements. 2) to shift from focusing on harm to focusing on the impact of name collisions. We came up with guidance on how impact can be assessed (two characteristics of the CDM: volume as an absolute number, diversity of that number). More volume in CDMs and more diversity within those CDMs drives the impact.]

3.2 Data Sensitivity Analysis

[Hypotheses/Goals: two things: any root server operator is representative of root server operations. That makes the fact that ICANN is already publishing visible name collisions out of L-root and making that publicly available. Second, resolver operators see a different view of the DNS than root server operators. Right now, Matt only has one global resolver, would like to have others, but we've already learned the most important thing (that they are different from RSOs). That they are different drives the controlled interruption and the requirement for it. Because you need controlled interruption to pull data out of resolvers so root servers can see it. You have to do the delegation to do the controlled interruption and see any conflicts. Also, will never see the resolver data again, it's not public. It's not so much about the spread as it is about the pockets of individuals. It comes with search suffixes wrt how organizations are configured. So many places do this in-house that you won't see it at the global level. There can be further analysis on this, and it's not analysis we're exactly doing right now. The root server data will never provide you a guaranteed view of all name collisions.]

3.3 Root Cause Analysis

In October 2017, ICANN began receiving reports through its Web form of collisions associated with the domain name wpad.domain.name. The reports indicated that HTTP traffic for users in various countries around the world was being proxied through a third party. This man-in-the-middle (MITM) attack violated users' privacy and left them vulnerable to theft of credentials or even identity. The attacks reported resulted from 1) home router software that had a default network configuration, 2) a protocol that made use of that domain to determine where traffic should be directed, and 3) malicious entities that exploited that vulnerability by redirecting traffic to them.

This report was written in direct response to those reports submitted to ICANN. In it we discuss the attack itself and the reports submitted to ICANN. Using artifacts and inferences from historical and recent Internet data, we also create a timeline of events that collectively tell the story of how the network changed over time to create an unsafe environment for vulnerable clients and end users. We also discuss the implications of the circumstances leading to the attack and summarize the key takeaways to be applied to related studies.

4 Board Questions

Part of Board resolution 2017.11.02.30 were a set of topics that helped define much of the scope for the discussion group. While the resolution referred to these as questions, they were not in question format. The discussion group found it a valuable exercise to reconsider each topic as a question; the responses to each are below.

4.1 Defining Name Collision

(1) a proper definition for name collision and the underlying reasons why strings that manifest name collisions are so heavily used;

4.2 Negative Answers

Board topic	Question as understood by the NCAP DG
(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;	What is 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?

As noted in the SSAC Report, “Redirection in the Com and Net Domains,” uninstantiated names that result in negative answers might occur for a variety of reasons: “A name might not exist because it had been misspelled, had lapsed or had never been registered. A name might also be registered or reserved but not included in the lookup database used for domain name queries.”³ Regardless of the reason, the errors received when returning a negative answer are in and of themselves useful to systems and end users. For example, systems such as spam filtering services may rely on the error to help determine if a message is spam by checking whether the domain name of the sender exists.

Any interruption or intervention in the path that results in a negative answer has the potential to intrude upon end-user privacy by allowing the intervening system to collect data on the user’s behavior and the path attempted.⁴ From a system perspective, interruption or intervention in the flow by a third party could result in increased network charges for some classes of users, a reduction in performance, or the creation of work required to compensate for the consequent failure.⁵

³ pg 3

⁴ pg 22

⁵ pg 23

4.3 Harm

Board topic	Question as understood by the NCAP DG
(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;	What are the types of harm and their likelihood to existing users if Collision Strings were to be delegated? This should include considerations around 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, as well as any other types of harm.
(4) possible courses of action that might mitigate harm;	What possible courses of action can ICANN org take that might mitigate harm?
(5) factors that affect potential success of the courses of actions to mitigate harm;	

Regarding the question of "harm," the study group focused on the potential for harm. The connotation of "harm" may include numerous things, from cybersecurity risks to reputational damage to physical impacts, making it difficult to appropriately apply scale and context to this otherwise broad term within the scope of name collisions. Real-world, demonstrable harm has been difficult to identify due to the limitations in both available data and the lack of clarity around the definition of harm. Where the study group found evidence of actual, clear harm, we have called that out specifically.

The responses to the board questions are not intended to address (a) the probability of an end user being harmed in any of these manners, (b) the frequency with which these harms would occur, (c) the degree of harm (if any) that could be incurred by any particular end user, (d) whether the existing end system's intended to exploit negative answers in designing its systems, nor (e) whether such harm could be avoided or mitigated in ways other than refusing to delegate strings. Instead, we have focused on offering information that will help the board with their decisions regarding domains on, or even yet to be added, to the collision string list.

Some distinct types of harm which can be hypothesized at this stage include the following:

- Reconnaissance/enumeration
- MitM attacks (Man in the Middle attacks may need subgroups)
- DocumentInternal document leakage
- Personal document leakage

- Malicious Code Injection
- Credential Theft

In the case of harms related to the delegation of Collision Strings, the study group has considered several potential areas of harm. These have not been directly observed; data that would allow for such analysis is not available.

For the sake of this discussion, we have categorized potential harms as involving either Interception and Manipulation or Signaling Interruption.

Interception and Manipulation includes private queries leaking into the public DNS. These would be queries that were previously answered by the root servers and which can subsequently be 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, man-in-the-middle (MitM) attacks, 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. When 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, as mentioned in Board Question 2 (???), discusses the role played by negative answers currently returned from queries to the root. This could include breaking applications that utilize the DNS as a signaling tool rather than as a directory (e.g. Chrome startup, Mozilla DoH, etc.). These situations are likely due to search list processing. 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. Unfortunately, the sheer number of possible local variations makes it impossible with current technologies to measure, assess, or remediate this potential for harm.

Regarding mitigations for harm, it is commonly held that you cannot change what you cannot measure. Mitigating name collisions is particularly difficult as detecting and reporting name collisions is itself challenging. At the local level, organizations are unlikely to be able to see the problem (e.g. transient corporate devices used on corporate networks) or even be able to reliably trace the causes. On a broader level, registrars and registries are unlikely to detect name collisions until well after the fact. As described further in **Section 6. General Recommendations** for this report, there will need to be a workflow that supports collecting information and having

the ICANN Board evaluate each case individually to determine the appropriate course of action based on their analysis of the potential risk.

4.4 Risks of Delegation

Board topic	Question as understood by the NCAP DG
(6) potential residual risks of delegating Collision Strings even after taking actions to mitigate harm;	What are the potential residual risks of delegating Collision Strings even after taking the actions described in Board Question 4 to mitigate harm?

4.5 Undelegated Strings and Collision Strings

Board topic	Question as understood by the NCAP DG
(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;	What are the 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?
(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	What are the suggested criteria for determining when a collision has been sufficiently mitigated that a Collision String can be removed from the list.
(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.	What measures would be appropriate and effective to protect against intentional or unintentional creation of situations that might cause strings to be placed in a Collision String category? What are the potential negative effects of a collision string list?

By “Collision String”, the NCAP Discussion Group assumes the Board is asking for a method of identifying a Top Level string that, due to the high risk of name collisions associated with the string and the potential for harm (as described in response to Board Question #3), that particular

string should not be delegated and should be reserved. The DG notes that there is, in fact, a spectrum or “range” of risks of harm associated with delegation of each new string as was observed in connection with ICANN’s Alternate Path to Delegation process followed in the 2012 round, which permitted more rapid delegation of certain strings provided certain names were “blocked” at the second level. In this regard, the subsequent Name Collision Framework which adopted a system of 90-day “controlled interruption” is a system which lends itself to identification of name collisions, but does not, in and of itself, mitigate those collisions.

5 Analysis and Findings

The study reports described above provided a wealth of information that informed both the responses to the board questions and the general recommendations below.

With the Case Study of Collision Strings, we learned what information is needed to identify and define all the critical diagnostic measurements required to identify name collisions. In particular, we were able to understand how to evaluate the level of impact as determined by the diversity and volume of queries. With the data at hand, we can see that while name collisions remain an issue, there is data that will allow for the development of appropriate mitigation and/or remediation strategies. We can also show that the telemetry data used by an RSI will continue to diminish in its fidelity as DNS protocol changes (e.g., QNAME minimization) further get deployed in the DNS.

The Data Sensitivity Analysis demonstrated that the available data sets from both individual root servers and global public resolvers are representative of the overall picture of the DNS queries that would help identify name collisions. Specifically, that data from any one RSI is equivalent to any other RSI, and that data from global public resolvers is not equivalent to any RSI. In practice this means that ICANN org can rely on the data it has available via the L RSI. Global public resolvers will have a different view, but that data is not readily available. In order to be certain of the existence of name collision activity, that string must be delegated to log the associated DNS request.

[The third and last report, the Root Cause Analysis, considered known name collisions and evaluated what mitigation or remediation happened, particularly what they did and why. The following sections describe the results of those studies in greater detail; the full report for each is available in Annexes A through C at the conclusion of this report.]

With the conclusions coming from the Case Study and Data Sensitivity Analysis, we demonstrate that Controlled Interruption provides a way to capture whether or not a name collision exists by using the data from any of the RSIs. However, Controlled Interruption by itself does not allow for evaluating the impact of the name collision. It does not allow for capturing the diversity nor volume of queries when they come in through protocols outside the

DNS. Understanding the level of impact is a critical component when evaluating the risk of delegating a given string. To get at that level of information, Enhanced Controlled Interruption becomes necessary as it allows for reviewing data that comes through other protocols.

- Case study
 - CDMs [therefore we know what to measure; we know how to find name collisions]
 - Impact as Diversity and Volume [therefore we have a mechanism to understand potential impact; the telemetry data used by L RSI will continue to diminish in its fidelity as DNS protocol changes (e.g., QNAME minimization) further get deployed in the DNS (less data is sent to the various portions of the DNS hierarchy)]
 - Name collisions continue to exist [therefore but they can accept the risk or support mitigation/remediation efforts (outreach efforts. May be more difficult in certain scenarios (CPE devices)).]
- Data sensitivity Analysis
 - Any RSO is representative []
 - Resolvers have a different view than RSO [therefore ... in order to be sure you are seeing previously unobserved NC activity related with that string you would be required to do a delegation to “bring that data up” to a third party logging the NC DNS requests (this holds partially true via CI)?]
- Root Cause analysis
 - There is no standard template for mitigation [therefore each instance needs to be considered individually, and any resulting mitigation requires re-analyzing the data]

Analysis

- RSO are all the same therefore ICANN can use what it already has control over = L root
- Need published NC data open to all (snapshots are not reliable)
- Need CI to see if NC exist
- Need ECI to develop a mitigation/remediation plan
- Need a role to evaluate NCs - both TRT and Applicant
 - the data is exists
 - someone needs to evaluate the data
 - to support the transparency required by ICANN and the community, both the applicant and a third-party technical review team need to have access to the data in order to evaluate it.
- Need to be able to pull back from CI and ECI
- Audience for the mitigation/remediation recommendation is important
 - Applicant would take actions
 - Board may choose to accept the risk and not require actions

Findings

- Need to ensure accurate and complete CI results and availability
- Need a role to evaluate NCs - both TRT and Applicant
- Need to be able to pull back from CI and ECI

6 General Recommendations

[Paper will already have told ICANN to prepare to change their tools/process/etc. Here's where we walk them through the details of the workflow with a focus on what applicants will need to do.]

[Technical Review Team: there's a function that has to happen. Coming out of the case study we make the observation that the critical diagnostic measurements exist and you're looking for certain criteria. You're looking for a technical evaluation about what's coming out of controlled interruption. We label that role as a "technical review team" - this is a function that has to happen at this point in time. We describe all the things this function has to do, but we don't speak to how that function is implemented. It could be a third-party contract, it could be a group within OCTO, or some other model that the board needs to decide on. – maybe not have this here]

6.1 Application Workflow

[What purpose will a decision tree serve? Can we turn this into a diagram like an actual decision tree?]

[Decision Tree: there are two things going on. We need to be careful that we are describing it as a workflow, not a decision tree. The decision tree phrase is something different. There are two broad recommendations to come out here: one, this is what the process looks like to capture what you need to know about name collisions and it has to be fit into the application process. The implementation review team will figure out exactly how that will work. The second thing we need to provide is "how is the board going to evaluate the technical package?" That would be a decision tree if it existed. Don't know that we'll be able to come up with a deterministic real flow diagram on how to evaluate name collisions. We'll only be able to talk about the issues to examine; every request will likely be unique. Each name collision situation is unique. There is no standard template strategy. They have to be investigated individually. That's what got us to enhanced controlled interruption so you could have enough data to draft a mitigation or remediation plan. The decision tree will offer whether or not to do enhanced control interruption. There may be visibility to the board to make that decision, rather than guidelines for the technical team to just do them. When it comes to Study 3, all we want to say is to motivate the need for ECI and the need for their to be an extra step on the part of the applicant to do the mitigation work and that will need to be evaluated on its own merit. Ultimately, we have to get to the point to demonstrate that a single template won't work.]

- Review risk of applying (Rec1: ICANN make as much NC data available to applicants (via ITHI, IMRS, etc.)
- Application submission (Rec2: ICANN use CDM to evaluate NC risks via available data (L-root IMRS) and the analysis be conducted by X to evaluate if the impact is low enough to move forward with temporary delegation for CI)

- Controlled Interruption (Rec3: ICANN has a platform/service for CI to collect NC data via the delegation of the TLD to a “third-party” [[should set out requirements not say this explicitly]] (not the applicant)) - needs to be defined process for discontinuing CI if needed
 - Board decision regarding Controlled Interruption - this will come from the Board’s approval of the application process
 - Capture a picture of the collisions
 - Applicant/Tech Review Team prepare addendum to application
- Enhanced Controlled Interruption - need to lay out requirements/guidelines for when to do this as well as exactly what this is
 - Board decision regarding Enhanced Controlled Interruption - need to lay out requirements - expecting that TRT will make recommendation to Board
 - Capture a picture of the collisions
 - Applicant/Tech Review Team prepare addendum to application
- Board gets final package for review
 - Should note that Board has opportunity to iterate CI or ECI
 - Should review mitigation/remediation plan and perhaps collaborate on what it really is with the applicant

[Ensure there is a "break glass" procedure if something goes wrong after the TLD is delegated. (see slide deck). There needs to be documentation - people will assert that the Board can always turn it off, but the detail that we need to document is it needs to be written down that it's possible. We won't examine in NCAP; this will result in a discussion with RSOs (RSSAC too). We need to recommend ICANN Board with with RSO and RSSAC on the process for this.]

6.2 Board Decision Process

[Board will consider the input and recommendation of the Technical Review Team and make a decision.]

7 Conclusion

[NCAP Study 3?]

We need controlled interruption

- Resolvers see collisions that roots don’t
- We won’t see resolver data again

We need enhanced controlled interruption (otherwise known as a honeypot)

- To investigate the origin of the name collision
- Protocol source - listen on all ports
- Protocol data - what’s on the inside to assess the activity
- QNAME minimization
- DoH/DoT/DoQ??
- Aggressive NSEC Caching, NXDomain Cut

8 Acknowledgments, Statements of Interest, and Dissents, Alternative Views and Withdrawals

In the interest of transparency, these sections provide the reader with information about aspects of the SSAC process. The Acknowledgments section lists the SSAC members, outside experts, and ICANN staff who contributed directly to this particular document. The Statements of Interest section points to the biographies of all SSAC members, which disclose any interests that might represent a conflict—real, apparent, or potential—with a member’s participation in the preparation of this Report. The Dissents and Alternative Views section provides a place for individual members to describe any disagreement with, or alternative view of, the content of this document or the process for preparing it. The Withdrawals section identifies individual members who have recused themselves from discussion of the topic with which this report is concerned. Except for members listed in either the Dissents and Alternative Views or Withdrawals sections, this document has the consensus approval of all of the members of SSAC.

8.1 Acknowledgments

The committee wishes to thank the following SSAC members and experts for their time, contributions, and review in producing this report.

SSAC Members

{TBD}

NCAP Discussion Group Members

{TBD}

ICANN staff

{TBD}

8.2 Statements of Interest

SSAC member biographical information and Statements of Interest are available at:
<https://www.icann.org/resources/pages/ssac-biographies-2021-10-21-en>

NCAP Discussion Group member Disclosure of Interest are available at:
<https://community.icann.org/display/NCAP/NCAP+Discussion+Group>

8.3 Dissents and Alternative Views

Appendix A: Case Study of Collision Strings

Appendix B: Data Sensitivity Analysis

Appendix C: Root Cause Analysis