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.

Draft Answer (workspace below)

Some distinct types of harm which can be identified at this stage include the following:
DEFINE THESE FOR THE BOARD AND THE COMMUNITY

- Reconnaissance/enumeration
- MitM attacks (Man in the Middle attacks may need subgroups)
- Internal document leakage
- Personal document leakage
- Malicious Code Injection
- Credential Theft

Workspace: for data and documentation

:

Applicable notes from Study 1
* ICANN has collected some incident reports
* Current literature will have some information
* Likely will need to be a thought exercise for us

Proposed Gap: * Likely to be a thought exercise on our part, extrapolating from what we know

Questions from Study 2 Proposal, Appendix 3

- Likely to be a thought exercise on our part, extrapolating from what we know and name collisions that have occurred..
- Might be dependent on classes of collision types and mitigation types, i.e., perhaps there's 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 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).

Think of this around the consequences for what happens if collision happens:

- signaling interruption (how will apps and/or program logic change)

- interception and manipulation (the name is resolving in an unintended manner, opening the door to a MitM attack, data leakage, or other issues)
    - there is research on wpad and other similar things that we can use
    - two sorts of categories here
        - disclosure of information
        - security compromise
    - are there harms that can be prevented through contracted relationships (i.e., ICANN contracted parties)? would be good to separate out the issues of incompetence from inherent problems with new gTLDs
        - harms you can do something about
            - a MitM attack by a malicious registry
            - Unclear if harms can be prevented by contracts; this doesn't directly help the impacted party.
        - harms that can't be controlled
            - a MitM attack by another registrant
    - The above can be categorized as we dive into some of the details. Example: think of previous case studies, which identified some systemic vs specific issues.
    - We'll need to be careful to stay focused on name collision issues; vendors not recognizing some TLDs is not a name collision problem.
        - in the case where that leads to collisions or is a result of collisions, that may well be in our remit. Think: .crypto
    - The harms in the case of blockchain projects are where those projects are establishing mappings between 'names' and wallet or contract addresses. A gTLD delegated to the root name system might create a MIM. I don't want to expand the z-axis of this topic, but it might be worth having this on the radar.
    - Are we giving people incentive to work around the concept of one authoritative root? If we take the view that we have to look out for the potential victim instead of focusing on the party causing the issue, ICANN would never delegate anything again and other parties would create their own systems. We can't really address the "what if" questions.
        - when we talk about mitigation, we'll have to talk about how the different elements impact the delegation decision
        - need to distinguish between causes, effects, and remedies. Also, there is more than one remedy besides "delegate/not delegate"
    - As we iterate on the harms, need to be clear on "harms to whom"

  - .crypto example: people who set that up were aware of real root and how it works. They have chosen to use that string. If it gets delegated, they get harmed, but so does everyone who has registered in that namespace and who doesn't know about ICANN. The potential operator is also potentially harmed since .crypto won't work as well as other TLDs
    - Though the users/registrants who are ultimately harmed are also protected by the law (e.g., pyramid scheme protections)
    - When we're looking at victimization, and what our response is to this problem, we need to consider the principle of "buyer beware". We can't exclude the fact we'll have to call out the problem exists, even if we don't have an answer to it.
    - Laws may provide remedies, but we need to consider order at the root. The remedies may not be "let the collision live because registrants have signed up to it"
- would be helpful, if only for our own purposes, to be clear what are known (explicit, active today) harms vs what are theoretical/future looking

From Ann's comment (4/21 call)
It is important for us to be clear about the definitions around the different types of harm. The report doesn't have the necessary contextual definition that the board will need. This will impact all the other questions as well.

We may not be able to enumerate all types of harm. We will probably start by creating lists and looking at specific examples. This could go towards finding categories of harm that matter and explaining through the use of examples. (see notes above re: consequences)

Tom: "there is an underlying assumption here that the app generating traffic that would collide with an icann tld is occurring at the root. what about collisions that do not occur at the root but were introduce by ICANN adding a new TLD?" Collisions at level below the TLD level are not within scope for us. TLD collisions are managed by registrations being first come/first serve. See: https://community.icann.org/display/NCAP/NCAP+Working+Documents?preview=/79437474/111387704/Definition%20of%20Name%20Collision%20and%20Scope%20of%20Work%20for%20the%20NCAP.pdf

- taking the example of .crypto. They are showing their customer base how to alter DNS locally so it goes to an alternate DNS. So, no collision today because there is no ICANN .crypto. What happens if ICANN delegates a .crypto at the root, and there are people that aren't initially harmed because they're pointing at a different DNS? What if their browser vendor unilaterally switched over to the ICANN root? Then the users are suddenly harmed.
    - Not within our scope.
    - .crypto is squatting on a name; we're only worried about our namespace. The alternate root is their problem, and their customers being impacted by the (potential) ICANN root is their problem.
- Are we going to address harm that gets exposed because of alternate namespaces? Need to determine if we're going to say something here in our work product.
- We need to acknowledge the issue of alternate roots, but agree if we're talking about harm re: ICANN harming .crypto if ICANN delegates .crypto is the wrong perspective. This may functionally look like name collision, but it doesn't have the same merits to the issue that some of the previous issues (.home, .web, .onion, etc) had. Different problem, same result.
    - This could touch on how we handle Board Question 9
    - Some of the domains intended as internal are already spilling out into the public internet
    - .onion was approved as a special use domain by the IETF and the ICANN Board.
- We have yet to see where a DoH or DoT might introduce some fragmentary namespaces within their own systems that we'd NEVER measure
- We have a definition of name collision of what's in our scope, and is what we used in Study 1. We can change it, if we have a compelling reason to do so.
- Should the IETF special use domain list be reserved in future gTLD rounds?
    - note there is concern that the IETF should not be in the middle of defining the special use list; .onion should be treated as a special case and not as a precedent
    - SubPro also considered this issue
    - See also Note: There is also this statement: https://www.iab.org/documents/correspondence-reports-documents/2017-2/iab-statement-on-the-registration-of-special-use-names-in-the-arpa-domain/
    - Here is the original MoU https://tools.ietf.org/html/rfc2860
    - .SPA was going to be a canary for collision with a handshake TLD but it appears to be delayed

- See the letter from Göran Marby asking for a meeting with the IAB/IETF on the topic: https://www.icann.org/en/system/files/correspondence/marby-to-cooper-kuhlewind-22oct20-en.pdf
    - IAB response: https://datatracker.ietf.org/liaison/1706/
    - Here is the ICANN posting of that response: https://www.icann.org/en/system/files/correspondence/cooper-kuhlewind-to-marby-12nov20-en.pdf