RySG
gTLD Registries
Stakeholder Group

David Conrad
Chief Technology Officer
Office of the Chief Technology Office
ICANN

Dear David

As you are aware the RySG has had a longstanding interest in ICANN’s Domain Activity Abuse Reporting (DAAR) Tool since its inception in 2017. In March 2019, the RySG formed an internal working group (RySG DAAR WG) to focus on two particular areas of the DAAR:

(1) Community misuse or misunderstanding of the data (potentially including Compliance); and
(2) Inaccuracy or, at least, lack of granularity of the data.

The WG has recently completed its work and I am pleased to provide you with a copy of the Final Report that has consensus support from the RySG. The report contains a number of findings and recommendations that we believe, if implemented, will improve the community’s understanding and utilization of DAAR.

The RySG DAAR WG is sincerely appreciative and grateful for the assistance provided to this effort by OCTO’s John Crain and Samaneh Tajalizadehkhoob, who have shared their expertise and knowledge with the WG on a regular basis. John and Samaneh’s willingness to work collaboratively with the WG on this effort has ultimately resulted in, what we believe is, a well-considered and informed Final Report.

We look forward to continuing this collaboration as we engage in discussions associated with implementation of the recommendations.

Kind regards

D Austin

Donna Austin
Chair, RySG

9 September 2020

RySG DAAR Working Group Report

Background

ICANN first referenced creation of its Domain Activity Abuse Reporting (DAAR) Tool in May 2017 when the Board mentioned the project in its response to the GAC's Copenhagen Communique. The RySG has taken multiple opportunities to comment on DAAR both formally[1]and informally, during face to face meetings with staff. The 2018 Compliance Audit of registry operators compliance with Specification 11, 3(b) of the Registry Agreement manifested in questions about the purpose and informative nature of DAAR.   In March 2019, the RySG formed a Working Group to focus on two areas: (1) Community misuse or misunderstanding of the data (potentially including Compliance), and (2) Inaccuracy or, at least, lack of granularity of the data.

Summary of work

- Studied the available DAAR reports and compiled a list of concerns.
- Refined the list of concerns into an Issues + Suggestions list which we discussed with OCTO's John Crain and Samaneh Tajalizadehkhoob over several meetings.
- Drafted suggested new text for OCTO to use in the DAAR reports.
- Suggested ways for OCTO to look at DAAR data and reviewed the results with them.
- OCTO noted that in 2020, the new DAAR reports would: (1) not present data as gTLDs vs legacy TLDs, (2) break out reports of spam into a separate graphic, (3) use a new metric for tracking abuse, and (4) would provide a monthly average, not a point-in-time snapshot.

[1] Comment on DAAR Methodology paper: https://docs.wixstatic.com/ugd/ec8e4c_f35b6c497f6346d2acfcbecd2251a5bb.pdf and comment on CCT-RT final report: https://docs.wixstatic.com/ugd/ec8e4c_ff7158c0470f40bc89c8c6ae2f8f5e03.pdf.

Findings

Finding #1: Reputation lists represent varying views of reported abuse activity that is correlated with domain names.

Reputation list providers are dependent on reports of abusive activity (including the use of honeypots as a type of report). This limits the types of abuse for which they can provide a view. Some of the types that are currently presented by DAAR do not align with the types that registries and registrars regularly seek to mitigate, for example spam.

Reputation lists are formulated from a variety of different sources: many are based on crowd-sourced data (i.e., based on reports from ordinary internet users) and some are based on data reviewed and confirmed by the reputation list provider itself or private groups of professional security researchers. Reputation list providers sometimes combine data from those sources, and sometimes offer different lists based on the source or type of reported abuse, among other options. Reputation list providers may change their formulation over time without notice.

The Domain Abuse Activity Reporting System combines these varying views of reported abuse activity to create a set of numbers that is intended to reflect reports of abuse activity.

Ordinary internet users and many internet-based enterprises use reputation lists to inform their decisions about what is or is not a domain name engaged in abusive behavior. Typically, these lists are used according to a weighted scale[2] set by the user of the list based on what works best in that user's environment, thus acknowledging that some reputation lists are more applicable in some circumstances than others. DAAR does not use a weighted scale instead applying all chosen reputation lists equally.

Finding #2: DAAR shows only reported abuse activity, i.e., it does not show any abuse mitigation nor the speed at which abuse is actively mitigated.

An essential point here is that no part of the registration system, i.e., neither registries nor registrars, controls whether or not abuse is present. A great deal of abuse is the result of

[2] A weighted scale permits a user to apply greater preference to one reputation list versus another. Often lists are assigned a quality value by a user or a user's tool and the combined values of a domain name on multiple lists determines whether or not a threshold for mitigation is crossed. Such values are determined by empirically observed quality by the user or user's tool regarding whether or not the list is fit for the purpose at hand.

malefactors who take advantage of weak security at the web site or other service hosted at a particular domain name. This type of abuse has nothing to do with the registration system and cannot be controlled by the registration system. Thus, abusive behavior will always be present and DAAR will always show the presence of reported abuse activity.

DAAR shows the count of domain names that have correlated reports of abusive behavior. However, when mitigation is applied, a domain name counted in one month may not be present in the next month or may not have been removed from the reputation list. Thus, although a given TLD may show "X" number of reports of abusive activity every month for several months in a row, the domain names listed in the "X" number of reports each month may be completely different from month to month. In such a case, the aforementioned TLD is not a "hot bed" of abuse activity, but rather an exemplary TLD actively mitigating abuse. The feeds on which the DAAR is based are designed for use in various types of filters, where "stickiness" is a feature. They do not attempt to tell the whole story of registry, registrar, or service provider responsiveness.

Finding #3: DAAR shows its calculated scores based on a point-in-time assessment.

As noted in Finding #2 above, the registration system does not control the presence or absence of malefactors. This particularly includes when malefactors choose to act. Thus, on any given day, a registry may or may not show any signs of reported abuse activity. This is likely to create an unbalanced, and even inaccurate, perception of the presence or absence of reports of abuse activity.

The working group believes a better approach would be to use an average of the reports of abuse activity. The count of the number of domain names reported with abuse activity for each day could be used to calculate an average for the calendar month. Other options may also be relevant, for example the median.

Finding #4: DAAR presents its reports as new gTLDs vs legacy gTLDs with no foundation for this distinction, nor do they identify any useful conclusion from the reports shown with this distinction present.

The working group believes all TLDs should be treated equally from the point of view of having domain names potentially exhibiting abusive behavior. Although there is some commonality between a large number of TLDs, generally each TLD has its own business model and its own procedures for addressing reports of suspected abusive behavior, and this is independent of whether the TLD is "new" or "legacy".

In addition, the question in the future will be to define what qualifies as a "new gTLD". ccTLDs are already being added and it is presumed that additional gTLDs will be added again soon. Are ccTLDs new or legacy, or yet another category of TLD? If gTLDs are added in rounds will each round be "new" or will all rounds combined, including the last one, be considered "new"? What if the next gTLDs are added as a continuous process rather than in rounds?

Finding #5: There is an inconsistent understanding within the ICANN community about what DAAR is and what it shows.

Anecdotally, it is clear there is a misunderstanding as to what DAAR actually shows and how what it shows can be used. The RySG DAAR working group repeatedly circled back to community perceptions of DAAR as being the biggest factor in our challenge to suggest and assess new metrics and variables.

One specific example worth noting is the broadly accepted community perception that DAAR shows[3] a systemic abuse problem in or related to the DNS. Since the vast majority of registries actively and consistently work to mitigate abuse we know this perception is false. This is the principal motivation for wanting to improve DAAR.

OCTO does take steps to provide context for the data presented in the reports. One observation is that the context is fairly broad thus suggesting that the data is applicable to all users equally. It is worth considering if it's possible to reconfigure the reports and the context to be applicable to specific user communities.

Recommendations

Recommendation #1: There should be a display of the volume of reported abuse activity by type of abuse activity based on the percentage of the total number of domain names with reported abuse activity.

[3] The CCT-RT's final recommendations assumed that the DAAR data could lead to policy development. The ATRT3's draft report indicates ICANN's accountability pages do not explain how DAAR shows accountability. The SSR2 report mistakes the purpose of DAAR throughout the draft, indicating the community misunderstands its purpose and scope.

Recommendation #1a: The percentages calculated in Recommendation #1 should be displayed over time, one display per type of abuse activity including the current month and the prior 12 months.

Recommendation #2: There should be a display of the volume of reported abuse activity by type of abuse activity based on the percentage of the total number of delegated domain names.

Recommendation #2a: The percentages calculated in Recommendation #2 should be displayed over time, one display per type of abuse activity including the current month and the prior 12 months.

Recommendation #3: Display a measure of the "persistence" of reported abusive activity.

The objective is to demonstrate that reports of abusive activity do change, for example by: (1) measuring: the average (mean) length of time a domain name exhibits reported abuse activity, and (2) describing characteristics of the population distribution such as standard deviation and shape (e.g., normal, bi-modal, or flat).

Recommendation #4: DAAR should consider discontinuing the distinction it makes between legacy gTLDs and new gTLDs, and rather focus on factors that more directly indicate the presence of abuse.

Recommendation #5: We recommend that ICANN update the DAAR report messaging to address the most prevalent community misperceptions about DAAR.

The DAAR report should clearly articulate: (i) feeds are subjective and often crowd-sourced and frequently show false positives (because the nature of the feeds prioritizes being over-inclusive), (ii) neither the feeds nor DAAR are capable of showing how reliably or quickly abuse is mitigated, and (iii) neither the feeds nor DAAR are capable of showing whether the domains were registered maliciously or have been compromised. We are willing to work with OCTO to ensure the DAAR report captures these concerns and recognizes the benefits of the new DAAR reports.

Recommendation #6: We recommend that ICANN conduct proactive community-wide webinars (perhaps linked to ICANN's smaller regional meetings) to introduce the community to the newly revised DAAR and its features and limitations.

The launch of the updated DAAR reports is a good reason to offer this sort of community-wide socialization of the data, which will provide a good opportunity to also educate the community on its usefulness and limitations. We are available to contribute to these webinars.

Recommendation #7: Collaborate with ICANN OCTO to create an infographic explaining the role, capabilities, and limitations of the various components in the abuse-reporting process: reporter/automated tool → 3P vendor→ cybersecurity folks, DAAR→ some ROs→ RRs → hosts → customers to the actual mitigation. Post the infographic on the DAAR page and provide a link in the DAAR report.

Conclusion

The RySG DAAR Working Group was formed to consider if there were changes or enhancements that could be suggested for ICANN's DAAR with the objective of improving the community's understanding and utilization of DAAR. We began with internal discussions and then reached out to work collaboratively with OCTO. We have facilitated and are continuing to facilitate changes to DAAR, with the cooperation of OCTO. We have prepared this final slate of recommendations to summarize our work.

We recommend the RySG review this report and its recommendations, and consider submitting this report to OCTO as a consensus position of the RySG. The working group will continue to engage with OCTO as the recommendations are implemented. We do not intend to develop additional recommendations.