Registries Stakeholder Group Statement

Issue: **RySG comments on the DAAR Methodology Paper and Methodology Reviews**

Date statement submitted: **24 August 2018**

Background[1]

On 20 July ICANN published a paper describing the methodology used in the Domain Abuse Reporting (DAAR) system and two reviews of that methodology.

The Community was invited to submit comments on the reports and reviews by 24 August that will be considered in the final drafting of the DAAR methodology paper.

Announcement and documents at https://www.icann.org/news/announcement-2018-07-20-en .

RySG Comments on the DAAR Methodology Paper and Methodology Reviews

The Registries Stakeholder Group (RySG) appreciates the opportunity to comment on the Domain Abuse Activity Reporting (DAAR) system and recent opinions ICANN requested from its selected experts on its methodology. We have carefully reviewed ICANN's methodology paper, the Statement of Work, and the papers prepared by Mr. Marcus Ranum ("Ranum Paper") and Mr. John Bambenek ("Bambenek Paper").

Transparency and Public Comments

- The paper asserts that the DAAR initiative has been undertaken in response to requests from the community but provides very little specificity as to what those requests were or from where in the community they generated (beyond a mention of the GAC's Abu Dhabi Communiqué). Without understanding the specific concerns that the DAAR initiative is trying to address, it is difficult to properly assess the methodology and whether it is meeting those goals.

- ICANN Org has so far failed to justify its reason for creating DAAR. It is still a solution in search of a problem. The DAAR may prove to be a highly useful tool for researchers to utilize, but ICANN Org has not yet justified to the ICANN Community why it spent the money to build a secret tool for which it now seeks post hoc approval and endorsement. The RySG objects to and has grave concerns about both the lack of transparency about when and why ICANN chose to build DAAR and why ICANN continues to hunt for legitimate use cases. These are answers experts cannot provide. They can only validate after the fact "how" DAAR

[1] *Background: intended to give a brief context for the comment and to highlight what is most relevant for RO's in the subject document – it is not a summary of the subject document.*

was built. The statement of work makes it clear that ICANN has only partially understood RySG concerns about the feeds selected for DAAR and the inclusion of spam.

- We are disappointed that ICANN did not publish the DAAR methodology paper and attendant reviews through the standard public comment process. Many RySG members missed the announcement about these materials when it was published. Furthermore, collecting comments through a single email address is not transparent and does not allow community members to review and respond to each other's input.

About the DAAR Initiative

- The "Purposes of the DAAR System" (page 4) of the Methodology paper highlights a key concern that the RySG has with the DAAR initiative: Each of the purposes are couched in terminology that overtly suggests that DAAR can be of direct aid to the Registry Operator or Registrar in assisting with anti-abuse investigations. While DAAR outputs may be capable of indicating a likelihood or a trend of potential abuse, the system does not provide actionable evidence of that abuse. DAAR outputs are merely indicators, and not specific enough to enable a registry or registrar to pinpoint abuse, which would be essential in order for it to aid in any anti-abuse investigations that registries conduct, as the "Purposes" section appears to claim.

- Furthermore, the paper notes that the DAAR does not attempt to measure mitigation activity or otherwise provide tools for registries and registrars to address or resolve instances of abuse. The results are also only accessible by ICANN staff. As such, we have some concerns about the potential utility of this initiative, especially given the resources that have been devoted to developing it.

- Additionally, the Methodology paper suggests that ICANN Compliance may be assisted by DAAR in the case of complaints filed against a registry or registrar. Members of ICANN's Contractual Compliance staff are already on the record as noting that they cannot use the DAAR statistics in isolation as an enforcement tool. It would be a worrying precedent should ICANN Compliance actually use such data, in its current form, to ground any enforcement action, especially considering the stated lack of actionable evidence and the decision not to perform any actual quality review or verification of the Data Feeds (RBLs) in use.

- The RySG recognizes that ICANN has sought to alleviate the concerns of, particularly, ICANN's contracted parties, but is disappointed to see that those efforts were largely unsuccessful. As the RySG has not yet seen a demonstration of the DAAR tool, we have not yet formulated an opinion about DAAR as a neutral piece of technology. It probably does what it's designed to do very well. We challenge what it was designed to do.

- It is unclear whether an appropriate review has been conducted of how DAAR fits into ICANN's narrow remit as defined in the Bylaws. Security threats that relate to content and don't directly impact the DNS are likely out of scope for ICANN and the list of reputation data feeds should be reviewed against these criteria.

- The paper also raises the question of what kind of legal obligations or liability ICANN may have in identifying and collecting data on domain names associated with different forms of abuse.

- The ICANN community does not have a common, agreed-upon understanding of domain abuse overall, which raises the question of whether the DAAR initiative can really achieve the lofty goals laid out in the introduction to the paper. The DAAR initiative focuses on certain types of abuse related to domain names, e.g., phishing, malware, botnet command-and-control, and spam. While these are certainly serious issues that merit attention, they are not the only forms of abuse that impact domain names. In the past, the ICANN community has raised concerns over and worked to address other types of abuse like domain tasting and front running, which the DAAR initiative does not cover.

Use of WHOIS Data

- The Methodology paper acknowledges that rate limiting sometimes makes gathering data very difficult but does not explain why real-time queries are necessary. Many registries and registrars have WHOIS terms of service that expressly prohibit automated, high-volume queries as the DAAR methodology describes using. Further, the paper states this challenge but does not provide any suggestions to work around it. Registries provide ICANN with access to bulk registration data files, but the paper does not consider this mechanism and why it may not be sufficient to meet the DAAR's needs. It also does not provide a strategy for how often to query to deal with transfers or updates.

- We also have some concerns about whether the practice of having the DAAR collection system query registration data is compliant with GDPR and other privacy regulations. Such queries would require a disclosure of personal data that has not previously been contemplated and could have significant repercussions for both registry operators and ICANN.

Data Sources and RBLs

- ICANN should provide more information about how it selects the sources for the DAAR collection system, including selection criteria and how the quality of the sources is assessed and measured over time.

- While the paper emphasizes the quality of the RBLs that the DAAR system relies on and asserts that they have low rates of false positives, the fact is that the DAAR system has the potential to amplify the negative effects of RBLs. Although these RBLs may have "well-defined" processes for removing false positive domains, real-world experience has shown that "well-defined" does not mean "well-supported" or "well-used." In Marcus Ranum's review, he points out that this is not the DAAR's problem: it is an issue between the RBL and the domain registrant. The DAAR, however, is amplifying and anonymizing the RBL's potential mistakes in an unaccountable fashion.

- Moreover, while RBLs are theoretically correctable, the DAAR's historical reports are not. Although the DAAR may correctly state the historical opinion of the Internet community, this opinion is not correctable by the falsely accused.

- The RySG has previously made clear statements expressing our concerns about the use of Spamhaus[2], but ICANN has chosen to continue to use this source for the DAAR initiative. The

[2] For example in its comment on the 'Statistical Analysis of DNS Abuse in gTLDs' (september 2017) the RySG cautioned against relying on statistical information provided by Spamhaus and others, and noted that "Some previous data supplied

RySG has significant concerns about the methods and practices Spamhaus uses, such as adding registrar infrastructure like SRS and mail servers to its RBLs, which could create security and stability issues.

Aggregate Statistics

- The DAAR does not provide aggregate statistics over individual RBLs, which prevents registries and registrars from identifying the lists that flag the most problematic domains and, by extension, addressing the underlying issues. Because the RBLs used have little overlap, registries or registrars seeking to act on these accusations have no recourse but to subscribe to the same set of RBLs (at least for particular security concerns) and build a similar DAAR system to identify accused bad actors. It is not clear what corrective steps registries and registrars might take, either to assist in clearing the good name of their customers or to restrict bad actors' usage of their systems.

- Relying on aggregate statistics also creates a risk of gaming the system by adding abuse domains from one's competitor registries or registrars to the RBLs used. The paper lists various disincentives against adding abuse, e.g., cost, discovery of bad actors, but if these disincentives were truly reliable, then the DAAR would not be necessary. In addition, damaging competitors' DAAR scores and reputations can also be accomplished by reporting additional domains (legitimate or not) from one's competitors to RBLs.

- The paper also fails to address the matter of miscounting domain names that may have been legitimately sink holed by law enforcement or security practitioners. Such domains have the potential to skew overall counts and thus the results displayed in the DAAR reports.

Comments on the Ranum Paper

The RySG urges ICANN to completely disregard the Ranum Paper. It contains numerous condescending, conclusory statements that are not justified by a single reference or even a logical argument.[3] Furthermore, it inaccurately characterizes many of our concerns by casting them aside as hyperbole:

- The paper repeatedly and excitedly observes how helpful this DAAR product will be to researchers, leading the RySG to wonder if ICANN has committed to making this tool public or otherwise promised researchers access to the DAAR. We would like ICANN to clarify whether any such commitments have been made.

- Ranum states, "Any complaints about the RBL scoring are not ICANN's problem, they are the RBL providers', or the registrars [sic]" and says "…if there are charges of inaccuracy, they are deflected over to the maintainers and producers of the blacklists." This illustrates a key problem for the RySG: ICANN has chosen RBLs, but has left the contracted parties holding the bag if there are problems with the feeds. Blacklists may refuse to talk to us or to

to the industry by these organizations was later found to include significant errors." - https://docs.wixstatic.com/ugd/ec8e4c_31c358b409d14b3ba80ad1ce454e6460.pdf

[3] Example of the many completely conclusory statement that are unsupported:
"We are satisfied that there are incentives to ensure accuracy in the data upon which DAAR is built." [No rationale as to why they are satisfied or why the community should be.]

whitelist our domains even if we take corrective action or decide after an investigation that there is no abuse under our respective policies.

- Ranum calls out another one of the RySG's concerns perfectly when he says, "Some might consider that a complaint that the list was wrong [sic], but I consider it 'how reputation lists work.'" The reputations of registries (and registrars), many of which are publicly traded companies, are collateral damage in both errant threat data and in simple misuse/misapplication of it.[4]

- Ranum says, "…DAAR is not 'naming and shaming' anyone…" He is correct in that DAAR, as a piece of technology, does not name or shame anyone. DAAR itself is agnostic. However, we know from experience how ICANN and other parts of the community can react to misinterpreted information and we are not confident that this data will be used to help registries.

- We also take issue with Ranum's argument that spam is a threat simply because a spam monitor calls spam a threat -- a dizzyingly circular argument -- when he says, "Cisco's Talos 'threat detection center' specifically treats spam as a threat not an annoyance."

- Ranum says, "…if one is arguing against spam blocking, one is arguing for spam…," which fallacy requires no explanation. The RySG of course opposes spam. What the RySG challenges is both its own role, and ICANN's, in lumping spam in with security threats because it crosses the line into evaluating content, which ICANN cannot regulate under its current Bylaws. We are not – and should not be – contractually obligated to monitor for spam and we're now being evaluated against a system that includes it.

Comments on the Bambenek Paper

In contrast to the Ranum Paper, the Bambenek paper, while we didn't necessarily agree with all of it, was well-thought-out, well written, and contained useful suggestions. It was also not well-cited but approximated a better-constructed expert opinion that carefully considered other viewpoints rather than treating them as hysteria. Bambenek was able to see and appreciate some of the concerns articulated by the RySG and we appreciate his moderation and practical solutions.

- Bambenek, like Ranum, alludes to the future availability of DAAR as a research tool. This further alarms the RySG for the aforementioned reasons, as they heavily imply that ICANN has made commitments to provide researchers with access to the DAAR in the future, without providing any information or details to contracted parties or the rest of the ICANN community.

- While the RySG appreciates Bambenek's thoughtful review of which blocklists ICANN should purge from its Malware Patrol feed and recommends ICANN accept this suggestion, as well as his careful review of the feeds to identify which ones may be content-based, we are concerned that he fails to review the five other "major" sources. He seems to simply accept that these feeds are useful and valid simply because they are widely used.

- Bambenek says multiple times that "Not all abuse can directly be attributed to decisions made by registries or registrars." He then points to ways ICANN can mitigate the results and

[4] See also: "Someone may disagree with any particular RBL's scoring for a particular domain but that is a problem between them and the RBL maintainer."

interpret the data to take into account which actions are attributable to contracted party actions (see pages 24 & 25 and page 34). We urge ICANN to explicitly agree with Bambenek's statement and carefully review Bambenek's suggestions.

- Bambenek also says that"…making determinations in a programmatic way on whether a specific indicator is truly malicious, compromised, or simply a service provider being used by a criminal is well beyond the scope of DAAR" (see page 4). This goes to the heart of our concern and we urge ICANN to understand that even if ICANN Org itself does not intend to use the data against registries, other people will.

- The RySG supports Bambenek's helpful suggestions on pages 6 and 7 as ways to potentially mitigate harm to our members.[5]

- Bambenek partially identifies, then discounts, a key RySG concern at page 31: "The biggest risk is for small TLDs or registrars, but it is not likely this will be much of an issue…." Reputation is all many of us have to differentiate ourselves from our competitors. Furthermore, even large companies can be reputationally crippled by a single instance of inaccurate negative data. In today's fast-moving media, a hard-won reputation can be lost in an instant with no regard for the veracity of the claim. This risk is even greater for those registries and registrars that are or are Affiliates of publicly traded companies.

Outstanding Questions

- Both Bambenek and Ranum noted that most companies (both non-profit and commercial) provide filters and protections to consumers. The RySG is left wondering why ICANN doesn't leave these companies to solve for abuse-related problems as they exist and develop?

- The Registry Agreement currently requires registries to monitor for abuse. Before launching the DAAR initiative, we ask that ICANN demonstrate, with empirical data, that the majority of us are not doing this. Otherwise, it raises the question, what gap is ICANN trying to fill with DAAR?

- Has ICANN promised to make DAAR or its data available as a tool to researchers?

Recommendation

The RySG recommends that ICANN disregard the Ranum Paper and carefully review the Bambenek Paper for its many worthwhile suggestions. The RySG again reiterates its desire for an open dialogue with ICANN about ICANN's intentions for DAAR. The specifics of how DAAR works is only part of the many open questions that the RySG continues to have about this initiative.

---

[5] This suggestion is in the spirit of mitigation and the RySG retains all individual and collective rights to challenge ICANN's use of DAAR.