Summary Report of Public Comment Proceeding

Statistical Analysis of DNS Abuse in gTLDs (SADAG) Study

Publication Date:	9 August 2017
Prepared By:	Brian Aitchison, ICANN Operations and Policy Research

Public Comment Proceeding	
Open Date:	9 August 2017
Close Date:	27 September 2017
Summary Report Due Date:	13 October 2017

Important Information Links
Announcement
Public Comment Proceeding
View Comments Submitted

Staff Contact:	Brian Aitchison	**Email:**	brian.aitchison@icann.org

Section I: General Overview and Next Steps

The SADAG study was requested by the Competition, Consumer Trust and Consumer Choice Review Team (CCTRT) as part of their mandate from Section 9.3 of ICANN's Affirmation of Commitments to examine "malicious abuse issues" and the "effectiveness of…safeguards put in place to mitigate issues involved in…the expansion [of the top-level domain space]" as a result of ICANN's New Generic Top-Level Domain (gTLD) Program.

The CCTRT determined that measuring levels of abuse in new and legacy gTLDs serves as a proxy for consumer trust, positing that any increase in abuse levels in new gTLDs would likely affect consumer trust of those gTLDs.

In defining the parameters of the study, the CCTRT sought to measure rates of common forms of abusive activities in the domain name system, such as spam, phishing, and malware. The study compares rates of these activities between new and legacy gTLDs, as well as employs inferential statistical analysis to measure the effects of DNSSEC, domain parking, and registration restrictions on abuse rates using historical data covering the first three full years of the New gTLD Program (2014 – 2016).

The CCTRT will review public comments on the study's findings and incorporate them into their final report as they deem appropriate.

Section II: Contributors

*At the time this report was prepared, a total of **ten** community submissions had been posted to the forum. The contributors, both individuals and organizations/groups, are listed below in chronological order by posting date with initials noted. To the extent that quotations are used in the following narrative (Section III), such citations will reference the contributor's initials.*

Organizations and Groups:

Name	Submitted by	Initials
Registry Stakeholder Group	Stephane Van Gelder	RySG
International Trademark Association	Lori Schulman	INTA
Governmental Advisory Committee	Fabien Betremieux	GAC
Google Registry	Ben McIlwain	GR
Generic Names Supporting Organization (GNSO) Business Constituency	Steve DelBianco	BC
Internet Service Provider and Connectivity Provider Constituency	Mark McFadden	ISPCP
Intellectual Property Constituency	Gregory Shatan	IPC
Generic Names Supporting Organization (GNSO) Non-Commercial Stakeholder Group		NCSG

Individuals:

Name	Affiliation (if provided)	Initials
Greg Aaron and Rod Rasmussen	iThreat Cyber Group and R2Cyber	AR
John Poole	Registrant, Editor DomainMondo.com	JP

Section III: Summary of Comments

General Disclaimer: This section intends to summarize broadly and comprehensively the comments submitted to this public comment proceeding but does not address every specific position stated by each contributor. The preparer recommends that readers interested in specific aspects of any of the summarized comments, or the full context of others, refer directly to the specific contributions available at https://www.icann.org/public-comments/sadag-final-2017-08-09-en.

Note that throughout this document, the Statistical Analysis of DNS Abuse in gTLDs study will be referred to as "the report" or "SADAG."

The summary of comments below is arranged according to comment submitter:

Aaron and Rasumussen (AR)

- Credit SADAG for contribution to understanding of DNS abuse issues
- Cite noteworthy findings, including those showing concentrations of abuse in relatively small number of registrars and registries
- Provide constructive critique of study's limitations, and emphasize a need for more analysis of "*why* abuse tends to be concentrated at a few registrars and in a few registries" [emphasis in original]
- Note the report's finding that "abuse counts primarily correlate with stricter registration policies" (SADAG, p.1), but suggest that the most critical driving factor of abuse rates is *price*

    - Remark on need to further research the relationship between domain pricing and abuse
- Suggest more in-depth analysis of particular business practices that may contribute to abuse (e.g. examining "portfolio players" who run multiple TLDs, and thus have similar abuse profiles across those TLDs)
- Emphasize shortcomings in methodology that undercounts numbers of maliciously registered domains

Business Constituency (BC)

- Remarks on the study's lack of conclusiveness regarding the effectiveness of New gTLD Program safeguards. The BC suggests that if those safeguards were effective, a decrease in malicious registrations would have been observed.
- Remarks on a potential relationship between domain pricing and concentrations of abuse, including business models that "depend on low-cost, high-volume registrations" that alter "the distribution of new gTLD abuse away from legacy gTLDs and toward new gTLDs" (BC Comment, p. 2)
- Emphasizes a number of other key findings from the study related to abuse patterns and potential drivers of abuse
- Makes a number of recommendations, including:
    - Focus ICANN Contractual Compliance scrutiny on highly-abused registries
    - Conduct more segmented research into abuse by TLD and registry operator
    - Incentivize anti-abuse practices among ICANN contracted parties and reward low levels of abuse

Google Registry (GR)

- Notes a number of key findings from the study and comments on the dynamics of abuse rates, including:
    - Patterns of abuse vary with type of abuse
    - Variation in registry and registrar practices to combat abuse contributes to variation in abuse rates associated with those entities
    - Abuse patterns shifted due to "low-cost or high-volume" domain registration practices in new gTLDs
- Emphasizes relationship between pricing practices, registry policies, and abuse concentrations.
- Offers a number of considerations based on the above as a basis for further study

Governmental Advisory Committee (GAC)

- Highlights a number of findings from SADAG, including a particular emphasis on the relationship between pricing, relative strictness of registration restrictions within a TLD, and concentrations of abuse
- Suggests that findings from the study should inform policy development, in particular within the GNSO's New gTLD Subsequent Procedures Policy Development Process currently in progress and with regard to study findings on the correlation between stricter registration policies and lower abuse counts
- Welcomes the use of empirical, statistical analysis to inform policy development

International Trademark Association (INTA)

- Critiques the scope of the study, and argues that various forms of trademark abuse should have been included in the SADAG analysis (see "On Study Scope" in Section IV, "Analysis of Comments" below)

- Notes that the methodology, analysis, and conclusions of the study are useful, but that the limited scope affects overall understanding of DNS abuse, especially as it pertains to trademark abuse

Intellectual Property Constituency (IPC)

- Similar to INTA, critiques the study's limited scope and encourages more focus on abuse related to intellectual property, trademarks, and copyright infringement
- Offers a number of specific comments for consideration by the CCTRT pertaining to how abuse should be measured, researched, and reported, including, but not limited to:
    - Focusing on some forms of abuse related to content (e.g. intellectual property abuse)
    - Comparing abuse in ccTLDs to other types of TLDs
    - Making available and adding more granularity to ICANN Contractual Compliance complaint reporting and resolution data
    - Studying the relationship between Public Interest Commitments (PICs) and prevalence of abuse
    - Researching the persistence of abuse (i.e. how long an abusive domain remains active) in addition to prevalence of abuse
- Suggests report "underscores the need to create a draft framework for a high security zone verification program…to establish a set of criteria to assure trust in TLDs with higher risk of being abused…" (IPC Comment, p. 3)

Internet Service Provider and Connectivity Provider Constituency (ISPCP)

- Notes the report indicates ICANN enforcement actions and New gTLD Program safeguards are not having an effect on levels of abuse
- Suggests that criminals seem to prefer maliciously registering domains rather than compromising them, indicating registration safeguards need to be strengthened
- Notes the study indicates criminals now prefer to maliciously register domains rather than compromise third party domains, demonstrating that "safeguards surrounding registration of domains are failing to address stability, security and resilience issues" (ISPCP Comment, p. 2)
- Critiques domain blacklist data, noting the lack of overlap between these lists, which in turn suggests abuse may be under-reported (see "On Domain Blacklists" below)
- Recommends analysis of the relationship between domain pricing and abuse

Non-Commercial Stakeholder Group (NCSG)

- Suggests that any growth in abuse can also be attributed to the growth of the Internet expressed by the number of domain name registrations
- Warns that "any analysis associated with expansion [of the DNS] should also take into account the benefits of lower costs, wider access, and more choices" (NCSG Comment, p. 1).
- Comments on a number of findings, and hypothesizes that price/cost are a significant driver of abuse
- Cautions against drawing conclusions based on this early evidence, and emphasize possibility of confirming the null hypothesis, i.e. that no statistically significant difference between abuse levels in new and legacy gTLDs may be present
- Notes that abuse levels between new and legacy gTLDs appear to be converging
- Endorses GAC suggestion that ICANN continue to employ statistical analysis to measure DNS abuse

Registry Stakeholder Group (RYSG)

- Supports reliance on empirical research to understand DNS abuse trends

- Remarks on how abuse tends to be concentrated, noting that it affects a relatively limited number of new gTLDs
- Requests that future studies be segmented according to different business entities
- Notes that the study indicates the New gTLD Program did not result in a net increase in abuse
- Suggests exploration of more effective mechanisms to combat abuse outside of existing abuse mitigation mechanisms at the registry/TLD level
- Forwards the hypothesis that price can be a predictor of abuse, but cautions that any recommendations coming from bodies such as the CCTRT on pricing remain cognizant that low-priced registrations can also serve to bring more registrants online, particularly in developing areas
- Discusses benefits and defends the use of privacy and proxy services
- Make several suggestions to improve methodology and reporting, including:
    - Independent validation of blacklist data accuracy
    - More specificity with how findings are characterized and cited

John Poole (JP)

- Seconds comment offered by AR on the quality of the study, while emphasizing AR's point that the SADAG study's methodology under-counts malicious registrants
- Offers critiques directed at the New gTLD Program and ICANN in general

Section IV: Analysis of Comments

General Disclaimer: This section intends to provide an analysis and evaluation of the comments submitted along with explanations regarding the basis for any recommendations provided within the analysis.

Overall, the comments provided thoughtful and considered input regarding the scope, methodology, and findings of the SADAG study. The comments offer constructive critiques and note many areas for future research. Most comments welcomed the employment of scientifically rigorous empirical statistical analysis to research abuse rates, and recommended further such studies be conducted.

The most prominent theme expressed in the comments is a need for further research pertaining to the relationship between domain pricing and abuse rates (AR, BC, GAC, ISPCP, GR, NCSG, RYSG). RYSG and NCSG caution that any analysis of the relationship between pricing and abuse should also account for the benefits of low prices of domains in terms of increasing competition and access to the Internet (especially in the developing world).

Responses and Clarifications to Specific Issues Raised

AR and JP note that the methodology employed in the study likely under-counts maliciously registered domains, a point acknowledged on page 9 of the SADAG study. AR in particular note that phishers and spammers may “age” their domains in order to receive better reputation scores.

RYSG discusses the benefits of privacy and proxy services (RYSG Comment, p. 2). As the RYSG’s comment noted, the SADAG study found no relationship between the use of privacy/proxy services and disproportionately high levels of abuse associated with them.

The ISPCP expresses concerns about the comprehensiveness and accuracy of the data employed for the study, and a lack of clarity on “sample selection”:

“[SADAG] carefully states that only a sample of legacy gTLD data was used in comparison to scan

the whole zone of the new gTLDs. This would suggest the possibility of a skewed sample and misinterpreted, inaccurate results. In addition, this makes it impossible to compare results in future years where the analysis is repeated" (ISPCP Comment, p. 2)

The SADAG study had complete data for 18 legacy gTLDs for the 2014 – 2016 observation period: .aero, .asia, .biz, .cat, .com, .coop, .info, .jobs, .mobi, .museum, .name, .net, .org, .post, .pro, .tel, .travel, and .xxx. It also had data for the 1196 new gTLDs delegated during the study's observation period. Data for the .mil, .gov, .edu, and .arpa legacy TLDs were not available and were thus excluded from the analysis of legacy gTLDs. However, the study's authors believe the existing data provides an adequately large set to observe general abuse trends in new and legacy gTLDs from 2014 - 2016. As long as similar blacklists, WHOIS and zone data are analyzed, the results of future studies can be compared.

"…in terms of methodology, very little is said about the selection of samples (and the criteria used for this selection) and both whether and why those are believed to be representative of the whole 'population' of TLDs or data. An estimation of the error to be expected between different samples would have made it possible to ascertain the findings" (ISPCP Comment, p.3).

Sampling as a statistical method was not employed and not appropriate for a study of this type. The study utilized data for nearly the entire population of new and legacy gTLDs during the observation period (see comment above on data availability). Therefore, sampling the data was not appropriate given data for nearly the entire population of gTLDs was available for analysis.

In response to a request from the GAC for further information on a reference to distribution of child abuse material on page 6 of the report (GAC Comment, p.2):

The SADAG study references previous research associated with the Dutch National Police in which the researchers employed a method to measure the amount of abuse associated with unique "fully qualified domain names (FQDNs)" distributing abusive imagery. This method to measure FQDNs was employed in the SADAG study as well. For further information please see: A. Noroozian, M. Korczynski, S. Tajalizadehkhoob, and M. van Eeten, "Developing security reputation metrics for hosting providers," in Proceedings of the 8th USENIX CSET, 2015, pp. 1–8, http://mkorczynski.com/UsenixCSETNoroozian.pdf

The following pages present a few points and clarifications in response to specific themes and issues discerned from the public comments:

On recommendations as a result of study findings:

The SADAG study was not intended to make recommendations on the effectiveness of New gTLD Program safeguards. Rather, it was intended to provide a baseline set of statistics on abuse rates and trends based on the best available data covering the observation period of 2014 – 2016 (i.e. the first three years of the New gTLD Program). The findings from SADAG are for the consideration of the CCTRT as they develop recommendations for their final report in Q4 2017.

On domain blacklists:[1]

[1] An extensive academic literature is devoted to measuring the accuracy and reliability of domain blacklists. For example, see: J. Ma, L. K. Saul, S. Savage, and G. M. Voelker, "Beyond Blacklists: Learning to Detect Malicious Web Sites from Suspicious URLs," in *Proceedings of the 15th ACM SIGKDD International Conference on Knowledge Discovery and Data Mining*, ser. KDD '09. ACM, 2009, pp. 1245–1254 and Kührer M., Rossow C., Holz T. (2014) Paint It Black: Evaluating the Effectiveness of Malware Blacklists. In: Stavrou A., Bos H., Portokalidis G. (eds) Research in Attacks, Intrusions and Defenses. RAID 2014. Lecture Notes in Computer Science, vol 8688. Springer

A number of critiques focused on the accuracy of domain blacklists in measuring abuse levels (IPC, ISPCP, RYSG). Indeed, blacklist providers' methods can generate false positives and negatives; "bad" domains are often reported by end-users rather than detected by technical means. Since most domain blacklists rely on end-user reporting of abuse, a useful way to think about what they measure is that they represent what *end-users* consider "abusive" domains. In this sense, any relative lack of accuracy is understood as part of intrinsic data collection issues.

Although, as several commenters noted, there is little overlap in domain blacklists between blacklist providers (a point corroborated by the academic literature on the subject cited in footnote 1, above), SADAG did show similar trends across those blacklists. For example, clear upward trends of abuse in new gTLDs were observed regardless of the blacklist provider. Comparing blacklists against each other—as the SADAG study did—is one means to ensure the *reliability* of the study's results.

Considering the known limitations of blacklist data, the study should be considered as a *reliable*—if not perfectly accurate—gauge of general abuse trends and dynamics based on end-user reporting, which must be corroborated by real-world experience and additional research.

On study scope:

Several comments critiqued the scope of the study, and suggested its expansion to include other forms of abuse (as characterized by the comment submitter). For example, both IPC and INTA suggested that trademark abuse be considered DNS abuse, and measured alongside those forms of abuse measured in the SADAG study.

The scope of the study was debated within the CCTRT. Ultimately, the team decided to focus on abuse types that affect all players in the DNS ecosystem, and was thus limited to broad categories of abuse, i.e. spam, phishing, and malware. In addition, these behaviors populate datasets (domain blacklists) that can be readily analyzed and cross-referenced with other data (such as the WHOIS and zone file data used for the SADAG study). The particular type or content of abuse beyond these broad categories was not considered in the scope of this study.

However, INTA and IPC in particular should note the potential for future studies that analyze large datasets of abuse that make use of novel methodologies to detect trademark and brand abuse. The methodology employed in the SADAG study distinguished compromised and maliciously registered domains: the researchers filtered data for domains that "contained a brand name or misspelled variant of a brand name" (p. 8, Section D). Such a methodology may be useful in future large dataset abuse studies focused on abuse associated with trademarks and brands.

The ISPCP comment is also indicative of comments to expand study scope. It suggests that "other, potentially better, regressions...could have provided a better fit" than those conducted in the study, and as with many other commenters hypothesizes price as a key potential driver of abuse.

Indeed, there may be other drivers that have more explanatory power in predicting an abuse rate. One such driver is price, which is mentioned extensively in the public comments on the report. Unfortunately, due to resource and time constraints, SADAG's researchers were unable to incorporate every potentially useful explanatory variable into their model. With this in mind, the SADAG study must be considered a first step toward more refined analyses based on community and peer feedback.