DNS
ABUSE
INSTITUTE

2023-02-24

Dear Sebastien Ducos,
Chair, GNSO

Thank you for your letter dated January 6th, 2023 on the topic of Bulk Domain Registration (BDR). In that letter, three specific questions were asked, and while we answer those questions to the best of our knowledge below, we've also taken this opportunity to share thoughts and insights on this topic in general, where appropriate. The DNS Abuse Institute appreciates the opportunity to provide input, and is very pleased to act as a resource to the GNSO, and the wider ICANN community.

Q1: What information, evidence or complaint statistics can you share that can shed further light on the potential role of bulk registrations in DNS Abuse?

The DNSAI does not currently have any statistics or evidence on bulk registrations. The DNS Abuse Institute does however collect a substantial amount of data as part of our DNSAI: Compass™ initiative. We'll conduct some exploratory research to see if we can identify bulk domain registrations, and would plan to share any relevant results. If others have data, or a proposal on how to study this issue, we're happy to engage there as well.

In response to your letter, the DNSAI conducted a modest survey of available academic or commercial research on the issue of bulk domain registration, with the hope of finding insights into how others have defined the issue, and what they have learned. Our survey uncovered very little research specifically into issues of BDR, though a number of works touch on the topic. References to BDR often reference the same single source of research. The papers we examined are included in the appendix to this letter.

The two papers (1, 2) primarily concerned with BDR offer some interesting insights, but appear to be fundamentally inappropriate as a starting point for community discussion as they do not examine, or attempt to examine, bulk registered domains unrelated to blocklists. They each take as their starting place domains from blocklists, leaving the exploration of benign but bulk registered domains untouched. It is nearly certain that some domain names acquired in bulk are also malicious registrations, but any real data-driven understanding of bulk domain registration needs to start from an attempt to discover *all* BDRs, to then understand what proportion is harmful.

Interestingly, and perhaps more significantly, from our research not only is there no consistent definition of bulk registration, there is no particular effort to define the concept clearly. Bulk Domain Registration has been referred to as both the mechanism by which the domains were acquired, such as via API, or by the amount of domains acquired within an arbitrary amount of time. A general categorization of the work would be that most appear to see bulk domain registration where a single actor is acquiring more than one domain, from a single registrar, over an undefined amount of time, though in the research above it was measured in days.

Q2: Are you of the view that further work may be beneficial to address potential issues with bulk registrations in the context of DNS Abuse? If yes, please provide further details.

Following on from our first answer, we are of the view that research would need to be conducted to determine the scale of any issues related to BDR prior to any policy work from the community. Without any insights into the role that BDR plays in general, it is difficult to determine definitively that work needs to be conducted with regards to the role BDR plays in DNS Abuse.

Definitional Issues

As discussed in our response to Q1, there is no accepted, or even meaningfully proposed definition of what Bulk Registration means. Further, if you define "Bulk" as a specific number of domains registered by an individual registrant within a specific amount of time, you've created a clear recipe for bad actors to avoid detection. There is a tremendous amount of diversity in the marketplace, with the largest registrars adding tens of thousands of domains a day, and hundreds of registrars that add less than a 100 domains a day; "Bulk" registration might be very different from one registrar to another.

In order to avoid providing a threshold to exploit, defining bulk registrations as relative to the number of domains per transaction[1] *for each registrar* is a potential avenue forward. Where the number of domains in a purchase is above some defined percentile (recognizing the threshold would be quite a high percentile) of domains per transaction for that registrar. However, we note that even relative standards for defining bulk registrations could have challenges within individual registrars, as registrar transactional volume and transaction characteristics (e.g. domains per transaction) can vary significantly due to marketing campaigns and promotions.

Absolute definitions are exploitable and could present an unfair burden on both very large and very small registrars or those who serve particular markets, and relative definitions are complicated and require sophistication to implement. There is no easy answer; a thoughtful and nuanced approach is required.

Q3: What measures, if any, do registrars and/or registries have in place in relation to bulk registrations (examples might include, but are not limited to, additional checks adopted where registrations go above a certain threshold, and restrictions on bulk registrations from new accounts)? Are these found to be effective in constraining malicious actors?

[1] Or, per customer

Would there be value in promoting the adoption of such measures on a voluntary basis, or should adoption through policy development be considered? Is there potential harm in the adoption of such measures?

On the Registry side, there has been considerable work done on a subset of BDR that does clearly fall into the category of DNS Abuse; Domain Generating Algorithms (DGAs). The RySG DNS Abuse Working Group and the GAC Public Safety Working Group published a Framework on Domain Generating Algorithms Associated with Malware and Botnets that is useful in identifying Registry practices in addressing DGAs. Similarly, the Internet and Jurisdiction Policy Network has a helpful Framing Brief - Improving the Workflow of Fighting Botnets: Handling Algorithmically Generated Domains on the topic of DGAs. While each of these works are helpful, DGAs are just one small subset of the possible universe of what could be “Bulk” registrations. And we note that, depending on the definition, registrations from a DGA might not be classified as “Bulk” registrations.

The experience of the DNSAI is that there is substantial diversity within the Registrar ecosystem on how registrars address issues of abuse, and how each architects its registration flow. Most retail registrars have access to some anti-fraud tools as part of processing payments. These tools are generally able to use an array of transactional attributes, including number of items, transaction amount, and length of customer relationship to flag potentially fraudulent transactions. The DNSAI published a best practice on the use of these tools to reduce abuse in December of 2022. The financial risk associated with fraud and abuse increases with the transaction amount, which can be impacted by the number of domains (and price of those domains) increasing the importance of these practices.

We are supportive of payment-based approaches to fighting DNS abuse for a few reasons:

- Payment processors are going to have state of the art methods for detecting fraud and malicious actors
- Most retail registrars are already going to be integrated to these tools, minimizing development time and effort
- Registrars are incentivized to detect fraud, and reduce credit card chargebacks. We can leverage this self-interest to simultaneously reduce abuse and fraud.
- The attributes of abusive registrations may already be included in detection algorithms regardless of whether they are “bulk”

While we are very vocal in our support of leveraging payment-based techniques and technologies, in our view, mandating their use via policy development goes too far in dictating registrar operational processes. We are already advocating their voluntary adoption, and will be happy to continue spreading that message.

Internet users, registrants, and registrars all need the freedom to be able to develop new business models, operational approaches, and uses for domain names, potentially including the bulk registration of domains. Thresholds and algorithms are always going to be imperfect and evolving, especially in the context of an ever-changing threat landscape. The attributes of abuse that are relevant for detection are not static, and committing to particular technological approaches risks becoming problematic, overbroad, redundant, exploitable, or harmful, and can have unintended consequences.

Friction at the time of registration is not the only potential approach for limiting harms from bulk registrations. An avenue potentially worth exploring is to encourage Registrars to investigate all of the domains in a customer account where one is identified as malicious[2]. This could have the effect of identifying other malicious domains not acquired in bulk, and the expense of such investigations should incentivize Registrars to implement controls they deem efficient.

There are sensible and practical options available to registrars that will reduce DNS Abuse regardless of transaction volume right now. We should encourage their adoption, and endeavour to understand the landscape of bulk domain registration before committing to any particular solutions.

Regards,
Graeme Bunton,
Executive Director,
DNS Abuse Institute

[2] As opposed to compromised

References:

1. Vissers, Thomas, Jan Spooren, Pieter Agten, Dirk Jumpertz, Peter Janssen, Marc Van Wesemael, Frank Piessens, Wouter Joosen, and Lieven Desmet. “Exploring the Ecosystem of Malicious Domain Registrations in the .Eu TLD.” In *Research in Attacks, Intrusions, and Defenses*, edited by Marc Dacier, Michael Bailey, Michalis Polychronakis, and Manos Antonakakis, 472–93. Lecture Notes in Computer Science. Cham: Springer International Publishing, 2017. https://doi.org/10.1007/978-3-319-66332-6_21.

2. Dave Piscatello and Dr. Colin Strutt. “Criminal Abuse of Domain Names: Bulk Registration and Contact Information Access.” Interisle, October 17, 2019. https://interisle.net/sub/CriminalDomainAbuse.pdf.

Appendix 1: Further Reading

Affinito, Antonia, Raffaele Sommese, Gautam Akiwate, Stefan Savage, KC Claffy, Geoffrey M Voelker, Alessio Botta, and Mattijs Jonker. “Domain Name Lifetimes: Baseline and Threats,” n.d.

Almashor, Mahathir, Ejaz Ahmed, Benjamin Pick, Sharif Abuadbba, Raj Gaire, Seyit Camtepe, and Surya Nepal. “Characterizing Malicious URL Campaigns.” arXiv, August 28, 2021. http://arxiv.org/abs/2108.12726.

Felegyhazi, Mark, Christian Kreibich, and Vern Paxson. “On the Potential of Proactive Domain Blacklisting,” n.d.

Hao, Shuang, Alex Kantchelian, Brad Miller, Vern Paxson, and Nick Feamster. “PREDATOR: Proactive Recognition and Elimination of Domain Abuse at Time-Of-Registration.” In *Proceedings of the 2016 ACM SIGSAC Conference on Computer and Communications Security*, 1568–79. CCS ’16. New York, NY, USA: Association for Computing Machinery, 2016. https://doi.org/10.1145/2976749.2978317.

Kidmose, Egon, Erwin Lansing, Søren Brandbyge, and Jens Myrup Pedersen. “Heuristic Methods for Efficient Identification of Abusive Domain Names.” *International Journal on Cyber Situational Awareness* 4, no. 1 (December 7, 2018): 121–42. https://doi.org/10.22619/IJCSA.2018.100123.

Korczynski, Maciej, Maarten Wullink, Samaneh Tajalizadehkhoob, Giovane C M Moura, and Cristian Hesselman. “Statistical Analysis of DNS Abuse in GTLDs Final Report,” n.d.

Lauinger, Tobias, Abdelberi Chaabane, Ahmet Salih Buyukkayhan, Kaan Onarlioglu, and William Robertson. “Game of Registrars: An Empirical Analysis of Post-Expiration Domain Name Takeovers,” n.d.

Lazar, David, Kobi Cohen, Alon Freund, Avishay Bartik, and Aviv Ron. “IMDoC: Identification of Malicious Domain Campaigns via DNS and Communicating Files.” *IEEE Access* PP (March 18, 2021): 1–1. https://doi.org/10.1109/ACCESS.2021.3066957.

Moura, Giovane C. M., Moritz Muller, Marco Davids, Maarten Wullink, and Cristian Hesselman. “Domain Names Abuse and TLDs: From Monetization towards Mitigation.” In *2017 IFIP/IEEE Symposium on Integrated Network and Service Management (IM)*, 1077–82. Lisbon, Portugal: IEEE, 2017. https://doi.org/10.23919/INM.2017.7987441.

Sabir, Bushra, M. Ali Babar, Raj Gaire, and Alsharif Abuadbba. “Reliability and Robustness Analysis of Machine Learning Based Phishing URL Detectors.” *IEEE Transactions on Dependable and Secure Computing*, 2022, 1–18. https://doi.org/10.1109/TDSC.2022.3218043.

Spooren, Jan, Thomas Vissers, Peter Janssen, Wouter Joosen, and Lieven Desmet. “Premadoma: An Operational Solution for DNS Registries to Prevent Malicious Domain Registrations.” In *Proceedings of the 35th Annual Computer Security Applications Conference*, 557–67. ACSAC ’19. New York, NY, USA: Association for Computing Machinery, 2019. https://doi.org/10.1145/3359789.3359836.

Szurdi, Janos, and Nicolas Christin. “Domain Registration Policy Strategies and the Fight against Online Crime,” n.d.

Weber, Michael, Jun Wang, and Yuchen Zhou. “Unsupervised Clustering for Identification of Malicious Domain Campaigns.” In *Proceedings of the First Workshop on Radical and Experiential Security*, 33–39. RESEC ’18. New York, NY, USA: Association for Computing Machinery, 2018. https://doi.org/10.1145/3203422.3203423.

Zhao, Chen, Yongzheng Zhang, and Yipeng Wang. “A Feature Ensemble-Based Approach to Malicious Domain Name Identification from Valid DNS Responses.” In *2020 International Joint Conference on Neural Networks (IJCNN)*, 1–7. Glasgow, United Kingdom: IEEE, 2020. https://doi.org/10.1109/IJCNN48605.2020.9207527.