topDNS Best Practice Series Part 5: DNSAI Compass Dashboards – Benchmarking the prevalence and persistence of malware and phishing for registries and registrars

The topDNS webinar on "DNSAI Compass Dashboards" was held on 28 November 2023. This was the fifth in a series of topDNS best practice webinars which showcase what the domain name industry is doing to fight DNS abuse. The webinar was moderated by Lars Steffen, Head of International, Digital Infrastructures & Resilience at eco – Association of the Internet Industry.

The specific aim of the webinar was to learn about the DNSAI Compass, which is an initiative to measure phishing and malware across the DNS ecosystem. The details on this initiative were presented by **Rowena Schoo, Director of Programs and Policy at the DNS Abuse Institute**.

In commencing the presentation, Rowena explained that, in 2021, the DNS Abuse Institute (DNSAI) had been created and funded by the US-based Public Interest Registry (**.org**) in pursuit of a non-profit mission. At the institute, DNS abuse is regarded as entailing technical abuse relating to phishing, pharming, malware, botnets and spam. In combatting this abuse, the work covers all top-level domains (TLDs) and registrars. The institute therefore does not just focus on the **.org**, but works outwardly on a free-of-charge basis to support the entire industry and to engage more broadly with stakeholders. In addition, the institute has a multi-stakeholder advisory council that helps with the institute's strategic direction.

As Rowena proceeded to explain, one of the institute's main projects for reducing DNS abuse at the DNS level is the DNSAI Compass ("Compass"), with this project set out to empower the DNS community with sufficiently granular, accurate and comparable data at the registrar and registry level. The principles set by the institute include transparency, credibility, independence, accuracy and reliability. As Rowena highlighted, in order to guarantee a level of independence and academic rigour to the project, the institute contracted an external party, KOR Labs, who facilitated the best possible way to measure DNS abuse. An online public methodology explains how Compass consequently accesses the required numbers.

Emerging from KOR Lab's facilitation, Rowena indicated how Compass now optimises accuracy rather than coverage of all harm, thereby measuring a very specific subset. Compass commences by zooming in on what is referred to as the "iceberg" principle, and then further narrows down to focus exclusively on DNS abuse in terms of phishing and malware, given that these forms of DNS abuse are deemed to be the most reliably evidenced at this point in time. What is therefore notified through the Compass reports is what has occurred in phishing and malware across the ecosystem, and whether or not mitigation has occurred, with this bearing a particular importance among registrars. Compass also gains an understanding of the type of registration, whether it's compromised or malicious. In this regard, as spelt out by Rowena, a malicious domain is one that is registered for the purpose of phishing or malware, while a compromised domain is a benign approach which is set out for negative things to occur.

In pursuing this specific subset, Rowena illuminated how Compass starts by viewing the input into reputation blocklists (RBLs), including APWG, PhishTank, OpenPhish and URLHaus, with these being measured at the unique domain level rather than the URL level. Secondly, Compass goes through a process of removing types of domain names where it would be inappropriate to act at the DNS level; with this, for example, leading to the removal of "special domain" names (such as URL shorteners, subdomain providers known to the institute, and file sharing services such as docs.google.com). Subsequently, there is a process of analysis that includes taking a fingerprint of the site – i.e., recording the DNS records, the content and the screenshot. The follow-on process then determines whether the domain name has been registered maliciously or whether it has been compromised, with the relevant domain name then visited at very consistent intervals (starting with five, 15 and 30 minutes; moving onwards to one, two, three, four, five, six to 12 hours; and proceeding for a period of 30 days). After the 30-day period, the measuring draws to a close, with the complete information then summarised into outputs that go into three main categories: interactive charts available on the institute's website, PDF reports that are also available on the website, and dashboards which are available for the registry and registrars to look out for their own data.

In this measurement of DNS abuse, Rowena acknowledged that there are some existing challenges. For example, there are obstacles such as false positives, involving phishing simulations or grey areas, with the latter supplying insufficient evidence to affirm that there is abuse among suspicious domains (see Rowena's blog on this topic). A further challenge concerns the terms of getting data for ccTLDs in terms of zone size and new registrations. Furthermore, in presenting overall data, this can be seriously skewed – for example, one bad month endured by a registry or registrar may appear very negatively. All in all, as Rowena reported, Compass naturally does its utmost to minimise all of these challenges and is very thoughtful in terms of how information is presented.

Public and Private Reporting

In presenting specific details on Compass's **public reporting**, Rowena noted how Compass mirrors their dashboards, with all of this data available publicly on their website. In Rowena's set of slides on these interactive charts, she displayed how people can zoom into aggregate trends on either phishing or malware, as well as being able to see how much of this has been perceived to have been mitigated. They can also learn whether or not the domain has been maliciously registered or compromised, with this adding an understanding of the speed of mitigation in terms of registrars' median mitigation time. A further aspect of the public reporting was outlined as involving PDF reports which are also available on Compass's website, with such reports presented online on a monthly basis. These reports do not just include aggregate reporting, but, since June 2023 also entail "specific reporting" on registrars and registries, listed in four different metrics and focused exclusively on malicious registrations.

In the final element of her presentation, Rowena reported on Compass's dashboards for **private reporting**. Such reports are also available free-of-charge, and can be of interest to those who, having seen the public releases, would like to see how close they are to the public tables or how they compare with their peers. Rowena displayed an example of these dashboards, which can be seen to serve as a valuable tool for monitoring the impact of policy changes over time, offering an independent and academically rigorous metric. This system centres on gTLD registrations, with the dashboard presenting high-level data on domain registrations, newly registered domains, and observed abuse, with a focus on phishing and malware. Users can toggle between months, view abuse trends, and assess mitigation rates. Peer group comparisons based on domains under management help users

benchmark their performance, although specific peers are redacted to allow users to only share or publish the data as desired. The dashboard also provides insights into registration types and mitigation strategies, empowering users to gauge their performance within their size category. Overall, the dashboards offer a comprehensive and user-friendly interface for assessing domain abuse metrics and fostering data-driven decision-making.

Feedback

Following on from her presentation, Rowena responded to a number of questions from various participants. For example, in addressing a question regarding reports on phishing, Rowena pointed out that DNSAI also has a project called NetBeacon which has been created as an intermediary to serve between reporters and registrars, with its related attempt being to standardise the evidence and to enrich the report as it goes through the institute's system.

In response to a final query regarding feedback from registries and registrars, Rowena stated that the provided data had been overwhelmingly positive, with users expressing gratitude for the opportunity to track abuse metrics for themselves as registries or registrars. As had become apparent, responsible entities investing in zone safety find value in gauging whether their efforts are effective, and the feedback indicates a welcomed tool for understanding team performance, with some adopting it as a key performance indicator. Even for those whose dashboards may not showcase stellar results, the information becomes crucial for making internal cases for resource allocation or procedural changes. As Rowena accentuated, users appreciate the uniqueness of the data, which delves into phishing and malware at the registry and registrar level, filling a gap often present in publicly available information primarily focused on network protection.