ICANN NCAP Meeting #38 (2021-01-20) Report Recap

Introduction. The meeting opened with a brief status report on Study 2, which is currently in the process of being formally written up. Matt Thomas then briefly discussed details behind the VeriSign blog post on name collision query remediation published the week prior. The remainder of this meeting was taken up by two technical presentations. The first, led by Jeff Schmidt, was a review of a five-year-old JAS report based on DNS-OARC DITL data. The second, led by Matt Thomas, was an analysis of current `.internal` query name activity seen by VeriSign at root servers A and J. What follows is a technical summary of those two presentations.

JAS Name Collisions Report. The 3700+ page JAS report is in three parts. The first 40 pages constitute the narrative, highlighting the findings and recommendations, while two large appendixes are a compilation of regular expressions on NXDOMAIN responses across the entire data set or per TLD. This study helped inform decisions about the future state of `.corp`, `.home`, and `.mail` TLDs, which have all been seen in use in the wild even though they have never been officially assigned or delegated TLDs. The JAS report, in cooperation with a data science partner firm, helped sift through the enormous amount of query string data to construct regular expressions, visualization patterns, and features not easily detected otherwise. An example pattern widely seen in the data set was the query string started with `adroot`, which provided hints to dig further and uncover vendors and software responsible. This example led to the discovery of a critical Microsoft Active Directory vulnerability. Another example was how A

root receives uniquely different query traffic, particularly in volume from others likely due to hard coded configurations and shortcuts taken in embedded devices.

> *[Editor's Note: The bulk of the DNS-OARC data comes from authoritative servers at the upper hierarchy of the name space where NXDOMAIN responses are returned for names that are either not registered or not provisioned in the name space. It is possible, such as with WPAD or ISATAP technology, that collisions from these names will be hidden when names associated with these technologies exist or are registered below the TLDs.]*

Follow up questions led to discussions about how JAS was able to track down the source of various query behavior, and the generally positive remediation outcomes even if not necessarily glamorous or easily automated. Much like the recent VeriSign name collision mitigation project, JAS performed a great deal of outreach to operators and vendors for remediation. However, it was suggested that a similar study conducted today may be less effective with the advent of qname minimization.

.INTERNAL Analysis. Matt Thomas provided a handful of graphs from A and J root showing query traffic patterns and trends over the past few years. The `corp`, `home`, `local`, and `mail` labels under the `.internal` TLD seen at the root exhibited a bias in query type behavior for SRV RRs when compared to other, allocated TLDs. There was also a noticeable increase in query volume and distinct source IP addresses beginning in March 2020 for `.internal` names, which was hypothesized to be related to the changing work patterns as a result of COVID-19. However, when examining the source autonomous systems of these queries, approximately 85% are coming from Amazon. The cause of which is currently under investigation.

The difference between a prior *.mail* analysis and this *.internal* analysis showed some clear distinctions. The *.mail* queries were more diverse and largely concentrated within the first few labels, whereas *.internal* tended to have a distribution of label count centered around 4 or 5.

The remaining 15% of *.internal* queries still represent a significant amount of query volume. A common label pattern of rancher.internal accounted for a large proportion of them. These are queries related to old versions of software used with Kubernetes management.

.SVC Analysis. Included as part of the *.internal* analysis was an aside into the *.svc* TLD seen in the wild, which has very recently begun to increase dramatically. This label is also related to Kubernetes, but appears to be coming from Microsoft's cloud infrastructure.

References

ICANN NCAP Meeting #38 page