Longitudinal Insights into Root DNS Traffic Patterns

Matthew Thomas & Danny McPherson

During its lifetime, the DNS system has evolved considerably to support new features such as DNSSEC, the expansion of top-level domains, and fundamental protocol changes. In lieu of readily available DNS usage metrics, various ad hoc studies have been required to evaluate the preparedness and performance before and throughout these major DNS transitions. Since 2013, Verisign has made concerted efforts to preserve high fidelity A and J root traffic data as well as other aggregated DNS traffic measurements to inform various security, stability, and resiliency objectives. To date, Verisign's data repository and analytic platforms are processing over 4PB of uncompressed (~640TB uncompressed) root resolution data, and we continue to preserve and collect this data for longitudinal study. While ensuring the data is in compliance with the fluid privacy landscape, our DNS telemetry has been essential for ensuring the availability of our A and J root servers.

Longitudinal measurements and usage trends are increasingly difficult to rely upon as the underlying system evolves or bifurcation within the system occurs. Since the 2012 round of new gTLDs and name collisions assessments, the DNS has evolved to include privacy enhancing techniques such qname-minimization (RFC 7816), aggressive negative caching (RFC 8198), local root techniques (RFC 7706), and the introduction of forwarders and recursive intermediaries, changes in recursive operator synthesis behavior and applications that attempt to evade this (RFC 8484), etc. These changes and their resulting implications, some profound, on DNS metrics directly affect programs and efforts that rely on such data.

Recent measurements based on A and J root data reveal a dynamic ecosystem of name leakage patterns. Persistent strings such as CORP, MAIL, and HOME continue to see increased query volume and sources since 2013. Meanwhile, we see numerous transient anomalies and new emerging strings associated with networking equipment and software systems becoming some of the most prevalent non-existent TLDs.

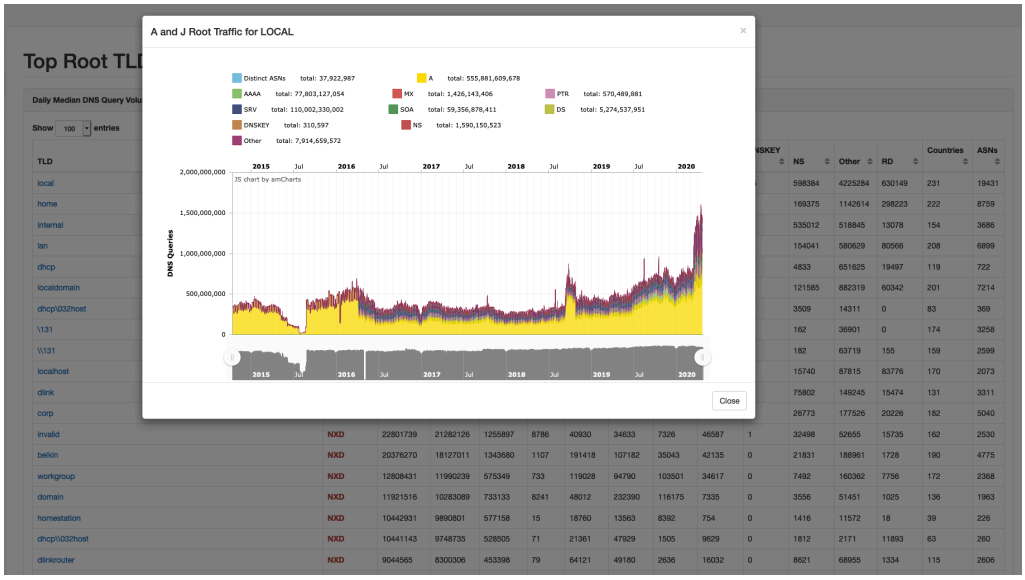

In support of NCAP-related efforts and a common SSR mission, Verisign continues[1] to be committed to making access to this data available to ICANN org for analysis in assessing these changes and riskiness for currently undelegated strings at the root. We continue to believe that the data will help address the data gaps identified by the NCAP DG to better understand and quantify the impacts to root server data that ultimately influences *proactive* name collision risk assessments and mitigation efforts, as well as new long-term longitudinal insights related to systemic changes and implications of various characteristics of and pressures on both the root server system and the global DNS.

[1]https://community.icann.org/display/NCAP/8+May+2019?preview=/109480152/109482829/NCAP%20Discussion%20Group%20Teleconference%2008%20May%202019.pdf