

Anomaly Detection on D-root (Abstract)

Zhihao Li, Dave Levin, Bobby Bhattacharjee, Neil Spring

University of Maryland, College Park

DNS root name servers play a crucial role in the Internet operation. Detecting and identifying anomalous activities around root servers is a critical task for network operators. It is not hard to “detect” the huge attacks [1], but how do we detect more than just the strongest, most extreme signals? How can we go about extracting, studying and understanding the smaller (but still nontrivial) anomalous events? These events might be from leakage traffic from botnet activities, throw-away traffic from misconfigured resolvers, or traffic load changes due to route issues, etc. To detect all these events requires one to effectively extract anomalous patterns from massive multi-dimensional measurements.

We present initial work towards detecting and identifying anomalous activities on DNS root servers. The method is based on Principal Component Analysis (PCA) to separate DNS traffic measurements into disjoint subspaces corresponding to normal and anomalous network behaviors. We have performed a preliminary analysis using data from D-root servers operated by UMD, and identified the detected anomalies by manual inspection.

Many techniques have been proposed to detect anomalies in network traffic, but they mainly focus on volume anomalies [2–4]. However, DNS traffic provides us more features, like query names, query types and DNSSEC queries. There has been extensive works on DNS traffic analysis focus on specific anomalous activities, such as botnet’s traffic [5–7] (domain-flux or fast-flux traffic) or DoS traffic [8, 9]. We aim to detect general anomalous traffic on root servers, that potentially relates to attacks, misconfigured resolvers, routing issues, and so on.

The dataset used in our study consists of sampled DNS queries and responses collected from the 98 anycast sites of D-root. On each site, we aggregate both queries and responses for each one-hour time period, and compute desired measurements. The one-hour time period is used as a tradeoff between the amount of data to be processed in each period and the granularity of anomalies to be detected.

We focus on the following measurements: (1) Query number per second (QPS); (2) Source address number per second; (3) Query diversity: number of unique query names over number of queries; (4) Source address entropy: the entropy of the query number distribution among all source addresses, which describes the degree of concentration of the query distribution.

Given one of the measurements and a time interval T , we construct a $T \times p$ measurement matrix \mathbf{X} , where p is the number of anycast sites. We then apply our detection method to the measurement matrix.

The idea of PCA-based anomaly detection is to iden-

tify typical variations among measurements and detect anomalous deviation from the typical variations [2]. Given the measurement matrix \mathbf{X} , we apply PCA to the covariance matrix of \mathbf{X} to compute a set of principal components $\{v_i\}_{i=1}^p$ that captures the variance among \mathbf{X} . Then we select the first $m \ll p$ principal components $\{v_i\}_{i=1}^m$ to construct the normal subspace, in which the majority of the variation is captured; the rest of the components constructs the residual subspace. When a new observation \mathbf{y} comes in, it is then decomposed onto normal ($\hat{\mathbf{y}}$) and residual ($\tilde{\mathbf{y}}$) subspaces, i.e., $\mathbf{y} = \hat{\mathbf{y}} + \tilde{\mathbf{y}}$. The energy of $\tilde{\mathbf{y}}$ (i.e., $\|\tilde{\mathbf{y}}\|^2$) describes the degree of deviation from normal variation, thus statistical tests [10] can be applied to this energy to test if the observation is anomalous.

To handle time-dependent (diurnal and weekly) patterns in the DNS traffic, we construct matrix \mathbf{X} using measurements from the one week-long time window prior to the new observation. We update \mathbf{X} only if the new observation is not anomalous, and rerun the principal components analysis accordingly.

We applied our method to all the measurements mentioned above, but here we show initial results of anomalies based on QPS. Among QPS measurements from all of D-root’s anycast sites throughout 2015, we detected 136 hour-long time periods when anomalies happened. In order to verify and identify these anomalies, we manually inspected them. We focused on identifying significant patterns in QPS, query diversity and source address entropy during the anomalies.

The anomalies are classified into four types based on the patterns. **62** anomalies are classified as “botnet activities”, as they included huge traffic volume increase and clear malicious query name patterns that are related to DoS attacks or algorithmically generated domains, such as [nonce] + “.ts8899.net(20)”. **18** anomalies had high volume traffic with query names potentially relate to bugs or faults in resolvers, such as “www.”, “http.” and “.”. There are **22** “traffic switch” anomalies, when the traffic volume decreased on several replicas, but increased on others. And **21** “traffic drop” anomalies showed significant volume decrease on some replicas, but had no corresponding increase on others. The rest **13** anomalies could not be classified into any of the above.

For the 62 periods with “botnet activities” detected in our dataset, D-root observed about 7 billion suspicious queries in total; by comparison, approximately 18 billion attack queries were observed during the widely publicized DDoS attack on A-root on Dec.1, 2015 [11]. With the series of botnet activities, we can profile their behaviors and track their evolution. Anomalies with buggy queries

reveal widespread faults among DNS resolvers; and we can infer routing issues related to anycast operations from “traffic switch” and “traffic decrease” anomalies.

We are working on applying clustering techniques to automatically cluster anomalies, and developing classification schemes that classify the anomaly clusters into types with distinct behavior patterns and causes. We hope that a discussion at DINR can help us better understand how to identify anomalies, what other measurements are indicative of anomalies, and whether we can, as a community, collectively construct an “anomalies dataset” to serve as a ground truth for future studies.

References

- [1] Giovane CM Moura, Ricardo de O Schmidt, John Heidemann, Wouter B de Vries, Moritz Müller, Lan Wei, and Cristian Hesselman. Anycast vs. ddos: Evaluating the november 2015 root dns event (extended).
- [2] Anukool Lakhina, Mark Crovella, and Christophe Diot. Diagnosing network-wide traffic anomalies. In *ACM SIGCOMM Computer Communication Review*, volume 34, pages 219–230. ACM, 2004.
- [3] Daniela Brauckhoff, Xenofontas Dimitropoulos, Arno Wagner, and Kavè Salamatian. Anomaly extraction in backbone networks using association rules. *IEEE/ACM Transactions on Networking (TON)*, 20(6):1788–1799, 2012.
- [4] Paul Barford, Jeffery Kline, David Plonka, and Amos Ron. A signal analysis of network traffic anomalies. In *Proceedings of the 2nd ACM SIGCOMM Workshop on Internet measurement*, pages 71–82. ACM, 2002.
- [5] Ricardo Villamarín-Salomón and José Carlos Brustoloni. Identifying botnets using anomaly detection techniques applied to dns traffic. In *2008 5th IEEE Consumer Communications and Networking Conference*, pages 476–481. IEEE, 2008.
- [6] Manos Antonakakis, Roberto Perdisci, Yacin Nadji, Nikolaos Vasiloglou, Saeed Abu-Nimeh, Wenke Lee, and David Dagon. From throw-away traffic to bots: detecting the rise of dga-based malware. In *Presented as part of the 21st USENIX Security Symposium (USENIX Security 12)*, pages 491–506, 2012.
- [7] Hyunsang Choi, Hanwoo Lee, Heejo Lee, and Hyogon Kim. Botnet detection by monitoring group activities in dns traffic. In *Computer and Information Technology, 2007. CIT 2007. 7th IEEE International Conference on*, pages 715–720. IEEE, 2007.
- [8] Keisuke Ishibashi, Tsuyoshi Toyono, Hirotaka Matsuoka, Katsuyasu Toyama, Masahiro Ishino, Chika Yoshimura, Takehiro Ozaki, Yuichi Sakamoto, and Ichiro Mizukoshi. Measurement of dns traffic caused by ddos attacks. In *2005 Symposium on Applications and the Internet Workshops (SAINT 2005 Workshops)*, pages 118–121. IEEE, 2005.
- [9] Laura Feinstein, Dan Schnackenberg, Ravindra Balupari, and Darrell Kindred. Statistical approaches to ddos attack detection and response. In *DARPA Information Survivability Conference and Exposition, 2003. Proceedings*, volume 1, pages 303–314. IEEE, 2003.
- [10] J Edward Jackson and Govind S Mudholkar. Control procedures for residuals associated with principal component analysis. *Technometrics*, 21(3):341–349, 1979.
- [11] Matt Weinberg and Duane Wessels. Review and analysis of anomalous traffic to a-root and j-root (nov/dec 2015). 24th DNS-OARC Workshop (presentation), 2016.