

DNS Trace Replay at Scale (abstract)

Liang Zhu John Heidemann
USC/Information Sciences Institute

1. THE NEED FOR DNS EXPERIMENTS

The Domain Name System (DNS) has grown to play various of broader roles in the Internet, beyond name-to-address mapping. It provides query engine for anti-spam [2] and replica selection for content delivery networks (CDNs) [3]. DANE [1] provides additional source of trust by leveraging the integrity verification of DNSSEC. The wide use and critical role of DNS prompt its continuous evolution.

However, DNS protocol evolution and expansion of its use has been slow because advances must consider a huge and diverse installed base: a complex ecosystem of many implementations, archaic deployments, and interfering middleboxes.

DNS performance issues are also a concern, both for choices about protocol changes, and for managing inevitable changes in use. There are a number of important open questions: How does current server operate under the stress of a Denial-of-Service (DoS) attack? What is the server and client performance when protocol or architecture changes? What if all DNS requests were made over QUIC, TCP or TLS? What about changes in DNSSEC key sizes?

Ideally models would guide these questions, but DNS is extraordinarily difficult to model because of interactions of caching and implementation optimizations across levels of the DNS hierarchy and between clients and servers.

We believe accurate, high-speed *trace replay* is essential to study many open questions in DNS, because DNS performance can be very sensitive to query timing and caching, and interactions across levels of the DNS hierarchy and multiple servers. These interactions seem impossible to model, and difficult to capture with a naive set of servers.

2. CONFIGURABLE DNS TESTBED

DNS experiments face the challenges of modeling as well as additional practical constraints. The distributed nature of DNS makes it hard to recreate a global hierarchy in a controlled experiment, because the hierarchy of DNS involves millions of authoritative servers.

Our goal is to build a configurable, general-purpose DNS testbed that enables DNS experiments at scale in several dimensions: many zones, numerous levels of DNS hierarchy, large query rates, and diverse query sources with the following design requirements.

Emulate complete DNS hierarchy efficiently: it must emulate multiple independent levels of the DNS

hierarchy and provides correct response, using minimal commodity hardware in a lab environment.

Minimal traffic to the Internet: experimental traffic must stay inside the testbed, without polluting the Internet. Otherwise large replay traffic and repeated experimental runs would stress the real DNS.

Manipulate queries arbitrarily: Replay must be able to manipulate traces to answer “what if” questions with variations of real traffic.

Support multiple protocols: As a special case of manipulating queries, it should be possible to replay queries with TCP and TLS to evaluate potential shifts in traffic mix.

Support high query rates accurately: Replay must can replay queries at fast rates, while preserving correct timing. to reproduce interesting real-world traffic patterns, both regular and under attack.

3. APPLICATIONS

We expect our trace replay testbed will support answering a number of research questions, such as:

Impact of Changes in DNSSEC Usage: Longer Zone Signing Key (ZSK) and more queries DNSSEC enabled (the DO bit set) will increase reply traffic. We expect to evaluate scenarios with different key sizes, and different mixes (up to 100%) of DNSSEC-enabled traffic. Testing against real-world traces allows evaluation of additional fragmentation, and increases in reply bitrates.

Performance of DNS over TCP/TLS: The use the TCP and TLS improves the security and privacy of DNS. While studies have suggested increased use of TCP and TLS has only modest cost, trace replay can provide a more complete evaluation. Important open questions include evaluation of connection-based DNS across multiple levels of the DNS hierarchy, and careful evaluation of memory requirements on actual server implementations.

4. SOFTWARE RELEASE

The software of our system will be publicly available at: <https://ant.isi.edu/software/ldplayer/index.html>

5. REFERENCES

- [1] P. Hoffman and J. Schlyter. The DNS-Based Authentication of Named Entities (DANE) Transport Layer Security (TLS) Protocol: TLSA. RFC 6698, Aug. 2012.
- [2] C. Lewis and M. Sergeant. Overview of Best Email DNS-Based List (DNSBL) Operational Practices. RFC 6471, Jan. 2012.
- [3] A.-J. Su, D. R. Choffnes, A. Kuzmanovic, and F. E. Bustamante. Drafting behind akamai (travelocity-based detouring). SIGCOMM '06.