

Towards a Testbed for Naming and Internet Protocol Experimentation (abstract)

John Heidemann Wes Hardaker Terry Benzel (USC/Information Sciences Institute)

1. BROADENING THE COMMUNITY OF DNS RESEARCHERS

DNS is today a critical part of nearly every Internet use—every web page resolution and e-mail message, and it is often a component of the Content Distribution Networks that provide video, images and other content.

Unfortunately, DNS’ ubiquity is one factor that has made it increasingly difficult to evolve. At a technical level, DNS has an installed base of “on everything everywhere”, much of which cannot or will not be upgraded. Changes to the DNS require great care to ensure no clients are left behind.

Unfortunately, DNS evolution is also hindered by a research community that is too small. While DNS began in academia, today it is often challenging for researchers to fully test new ideas about DNS and Internet naming because the lack of access to meaningful datasets and realistic experimental setups.

A broader research community for DNS is essential for the continued health and growth of the Internet. DNS has a role to play in improving the Internet in new ways: with the strong integrity provided by DNSSEC, DNS has the potential to significantly increase Internet security. DNS and DNSSEC may also have a role to play as embedded, Internet-of-Things devices need to securely acquire updates and communicate with the world. In addition to these new opportunities, DNS *must* evolve to continue doing what it does today (name resolution) in the face of a changing Internet. DNS infrastructure is being confronted with increasingly large Denial-of-Service attacks [3, 4], and today’s internet places much higher expectations on data privacy and security than when DNS was first conceived [1, 5].

We believe new *Naming and Internet Protocol Experimentation Testbed* (NIPET) can provide a new kind of research infrastructure, allowing a much broader community of researchers to contribute to DNS evolution and new applications. NIPET will be shared research infrastructure that supports long-term data analysis, experimentation under real-world conditions, and a smoother path from idea, to test, and to deployment than is possible today. This research infrastructure will be informed by the real-world challenges and realistic data of a DNS root service.

2. GOALS AND REQUIREMENTS

We see NIPET as marrying real-world data from a Root Server with new components to support research:

Longitudinal data collection: Some studies require *continuous, long-term data collection*. While DITL events happen regularly [2], each DITL collection is brief (typically two days), with long gaps in between. Analysis of trends is difficult with sparse data, and questions that require longer measurement periods (such as evaluation of caching effects, or trends in client population) cannot be answered.

Experimental comparisons: Many questions require *experimental evaluation of real-world traffic*. While off-line testing with synthetic data can answer some questions, real-world traffic provides greater diversity and intensity. Our goal is to support evaluation of new technologies against real-world traffic in real-time, allowing exploration of the timing and caching constraints that often matter in DNS. Such testing will also support A/B comparisons of new approaches with existing software and designs.

Smooth transition to practice: It must be easy to move experimental protocols and software in and out of service. While off-line evaluation of protocol changes give some confidence in their applicability, deployment must eventually place experimental software into operation. We need such transition to be seamless, allowing experimental software to be run in parallel with production software, eventually shifted into service (perhaps answering a fraction of traffic), and also shifted out of service should (or when!) problems arise.

An open community: We seek to open up these tools to a large researcher community. We see our proposed research infrastructure being shared across multiple researchers in academia and industry. We envision both a central facility that many can access, and software that others can download and use locally. Our ultimate goal is to grow the size of the research community and their output, for all to benefit.

While this research infrastructure benefits from the experiences of an operational Root DNS service, we also recognize the responsibility that operational networks require. First, research performed within NIPET must not impede correct operation, the primary mission of a Root DNS service. Second, the infrastructure must reflect the sensitivity of real-world data, balancing research value against privacy concerns. While DNS data at the root is typically filtered through recursive resolvers, we see policy constraints, agreements on how data is to be used, and anonymization (where possible) as important tools. Development of better DNS-specific anonymization methods is critical.

3. REFERENCES

- [1] S. Bortzmeyer. DNS privacy considerations. RFC 7626, Internet Request For Comments, August 2015.
- [2] DNS-OARC. Day In The Life of the internet (DITL) 2014. <https://www.dns-oarc.net/oarc/data/ditl>, April 2014.
- [3] Root Server Operators. Events of 2015-11-30. Technical report, Root Server Operators, Dec. 4 2015.
- [4] Root Server Operators. Events of 2016-06-25. Technical report, Root Server Operators, June 29 2016.
- [5] Liang Zhu, Zi Hu, John Heidemann, Duane Wessels, Allison Mankin, and Nikita Somaiya. Connection-oriented DNS to improve privacy and security. In *Proceedings of the 36th IEEE Symposium on Security and Privacy*, pages 171–186, San Jose, California, USA, May 2015. IEEE.