

How to measure KINDNS? (abstract)

Raffaele Sommese

University of Twente - r.sommese@utwente.nl

Introduction

The Internet's naming system (DNS) is a hierarchically structured database, with hundreds of millions of domains. The DNS's distributed architecture is primarily responsible for its capability to scale to its current size, but it also poses stability and security problems. To increase DNS resilience, the Internet engineering task force (IETF) has published several operational best practices; nonetheless, operators must still make their own choices that balance security, cost, and complexity. Since these choices could affect the security of billions of Internet users, ICANN recently put out the Knowledge-Sharing and Instantiating Rules for DNS and Naming Security (KINDNS) initiative, which aims to codify best practices into a set of universal norms to enhance security.

This project was inspired by a similar effort for routing security, Mutually Agreed Norms for Routing Security. The MANRS program invites operators to voluntarily agree to a set of procedures that will improve collective routing security - a challenging task given the incentives to comply with these practices do not provide a clear return on investment for operators. Independent verification of compliance with the practices is a barrier for both initiatives. The KINDNS discussion has only recently begun, and stakeholders are currently arguing what should be included in the defined practices.

Our Contribution

At this early stage, we examine potential best practices in terms of third-party measurability, including a review of DNS measurement studies and available datasets.

The KINDNS group has proposed practices (P) for authoritative (A) and recursive (R) nameservers, and practices (H) for general infrastructure hardening. We concentrate on public-facing DNS infrastructure, such as open resolvers and authoritative nameservers.

We identify non-measurable practices and recommend alternative practices to include based on existing scientific research.

We hope to present this work to DINR2023 to gain community feedback on:

- What are the gaps in KINDNS best practices, and what should be added?
- How can we accurately assess and measure KINDNS adoption in a responsible way?
- What datasets can be used to assess KINDNS-readiness?
- How can operators (large and small) be invited to follow KINDNS best practices?