Consolidated Review of
D-mystifying the D-root Address Change

1. Strengths

The paper is interesting and well written.

The paper is based on a unique dataset that looks at interesting phenomena occurring as a result of the D-root DNS server changing IP addresses, and does a very good job in pinpointing potential causes for these phenomena.

Careful measurement and some interesting initial explanations for a variety of surprising discoveries

While previous studies have tackled the topic of root IP changes, this paper adds a level of depth I have no previously seen. I like the notion of writing down information about these once-in-a-long-while changes. I liked the idea of changing root IPs as a way of "garbage collection". That is both cute and probably actually a reasonable idea to cull chud. This is the quintessential IMC short paper. It isn't a huge study, but it is interesting and it is useful information.

2. Weaknesses

It is short on context. We get some insights from the J-root change but almost nothing about the B-root change. We also hear nothing of the experience moving the early root servers c. 1990-1991. Paul Mockapetris gave a number of talks about how long people kept querying A.ISI.EDU after it was shut off and also about the percentage of bogus queries. It would have been useful to do a bit more long-term analysis here.

I thought the history window before the change could have been much longer. Seemingly the authors have the data. However, we're supposed to believe that the day before the new address went live is supposed to capture the "before" behavior. Given that the behavior of most networking phenomena is not greatly stable I think this is a shaky notion.

3. Comments

While I think more history would have been useful, I don't think it is a showstopper for an IMC short.

Any evidence for the distribution of the PowerDNS around the globe, or its growing popularity over time, would strengthen the argument in Section 4.2.

In Section 5.2, I can understand that a faulty (misconfigured) resolver might continue to use the old IP address, but in the case of a resolver being used for attack, why shouldn't it switch to the new IP address? The explanation presented doesn't seem to address this.

Q3 (page 3) is left unanswered.

First paragraph of section 4: What is the timeframe over which these numbers are calculated?

4. Summary from PC Discussion

This paper was accepted without discussion.

5. Authors' Response

The reviewers raised the concern of a lack of context. When we performed our literature searches (before and after seeing the reviews), we were unable to find any analyses of the root causes of the generally well-known anomalies. One reviewer mentioned talks that Paul Mockapetris had given in the early 90s, but neither we nor the reviewer were able to find these materials (and the reviewer did not recall analysis of the root causes). In preparing the camera-ready, we contacted the J-root operator, Bill Manning, who gave the talk cited in our paper, but he informed us that they too had not had the opportunity to perform such an analysis. From these conversations, we draw two conclusions. First, our paper's analysis into root causes for changeover anomalies is the first of its kind; the context of our results is the set of prior observations of the anomalies, which we have tried to capture throughout the paper. Second, while our classification of hosts and our observation of PowerDNS's behavior describe the behavior we observed in *this* changeover, it is unclear to what extent they apply to previous changeovers, particularly those from 20 years ago. We believe this strengthens our (admittedly somewhat outlandish) suggestion to use periodic changeovers as a crude form of garbage collection. In summary, our paper was somewhat short on prior quantitative analysis into root causes because, unfortunately, there was none. We hope this work is the first of many that investigates anomalous behavior at root DNS servers.

One reviewer asked why we had not extended our analysis to more than a single day before the changeover. Our high fidelity data (captures of 200k packets every minute) does not extend to more than 24 hours before the changeover. We had collected coarser-grained data prior to the changeover. Our analysis of this historic data shows that the data collected from the changeover-eve was typical. In particular, the data from the prior week and the changeover-eve had nearly identical distributions of query types and return codes, as well as reasonably consistent query volume.