Consolidated Review of
On Measuring the Client-Side DNS Infrastructure

1. Strengths:

The paper sketches the complexity of today's DNS infrastructure and shows a number of clever specific techniques to understand some aspects of DNS resolvers. The paper is written and presented well. It describes the complex infrastructure of DNS very well so that someone who is not keen on the system can also appreciate the paper. The reviewers also like the methodology to discover hidden entities in the DNS infrastructure. This methodology is a new contribution in the field. It's a clever exploration, which significantly contributes to the understanding of the current client-side of the DNS infrastructure.

2. Weaknesses

The paper doesn't give a clear "big picture" of what the research goals are, or how close it comes to reaching those goals. Although providing considerable data, the paper does not go very deep into understanding the causes of these data.

As someone who is not an expert in the topic, it was not clear how novel the contributions in this paper are. For instance, the authors mention their coverage of the DNS space covered is twice larger than the previous work. Is this a notable contribution? The paper is also the first to assess how various entities treat TTL. Given that DNS system has been around for a long time, why has no one looked at this issue? Does it mean this is an important topic to discuss? Some of these questions arise because most of the references are not recent and are from a few years back.

The negative side the reviewers see in this paper is that most of the findings are related to buggy software and unconscious deployment of open resolvers that really shouldn't be open. Part of the measurement findings will probably need frequent updates, as the default configuration of home routers that act as open resolvers and forwarders may possibly change in the near future.

3. Comments

This paper explores a lot of aspects of what appear to be a very complex part of the Internet that has not yet been fully studied.

On the way, the paper shows a number of clever specific techniques to understand some aspects of DNS resolvers. There is a lot of care put into developing specific queries that identify cache handling in at different levels of the DNS resolution chain, some hidden from direct observation. These techniques are really nice and the audience will appreciate them.

However, the paper as a whole doesn't give a clear "big picture" of what the research goals are, or how close it comes to reaching those goals. Part of the challenge here is that the paper is very good about saying how the work was done. But often the "how" comes at the cost of obscuring what the problem under consideration is and why it is important. As a specific example, the first paragraph of 7.1 is: "To investigate aggregate behavior of the DNS resolver infrastructure, we performed further probing of 2.4M FDNSes during the S5 scan." which makes it very clear the paper will do further investigation, but says nothing about what the question is. The heading of 7 adds some information (it's "caching behavior"). The actual reason for the investigation is buried in the middle of the 2nd paragraph: "recipients are known to disobey TTL [5,16]" and seems to hint that the reader should care because it has something to do with the Kaminisky attack. The authors could provide a clearer statement of what the question is.

Second, although providing considerable data, the paper does not go very deep into understanding the causes of these data. As a specific example, Section 7 shows that different groups of resolvers mishandle TTLs with different probabilities, with FDNSs often claiming TTL=10000, RDNSdis 3600 or 86400, etc. Why does this happen? Answering "why" is hard, but the authors could perhaps provide probable guesses (perhaps the Linksys firmware has hardcoded TTL, and that accounts for 60% of FDNS?). The reviewers expect a deeper analysis in a long IMC paper.

Finally, the reviewers raised two concerns about methodology:

1. The paper spends a lot of time on scanning methods, and accounting for potential bias. Complete scans of IPv4 seems fairly commonplace today. A reasonable computer should be able to easily probe 6k addresses per second (and probably well more than that with optimization)---eight of those should cover 2^32 IPv4 addresses in one day. Why risk bias when full coverage seems to be not that hard?
2. The statement: "We double, from 15 to 30 million, previous estimates of the number of open resolvers on the Internet." Seems somewhat inconsistent with other reports, and some assumptions behind the projections in sec. 6.1 seem incorrect. The paper compares to 15M as on 2010 data from The Measurement Factory. They actually publish daily reports: http://dns.measurement-factory.com/surveys/openresolvers/ASN-reports/ with current values reporting only 92k. This is fewer then the 15M the authors report from their data and the 30M estimated. Some of the overestimate may arise because the authors scale up sample to the 2^32 addresses, even though though 1/8th of IPv4 is either multicast or unallocated and so cannot host DNS. One might also consider observations that IPv4 is very unevenly used (see for example "Census and survey of the visible internet", http://dx.doi.org/10.1145/1452520.1452542), which would further reduce the scaleup factor.

The description of DNS infrastructure on Section 3 is very informative. Yet, because there are so many abbreviations in the paper (e.g., ODNS, ADNS, RDNS, FDNS), the later analysis part becomes a bit confusing to understand. It would help the readers understand the paper better if there are several conceptual figures that depict how the measurement was done.

It seems there are more things that could be investigated using the dataset from these measurements. E.g.: Sections 6.2 and 6.3 Figure 9 and 10: it would be nice to see if these distributions change when broke down per AS, country, AS type, etc. Similar observation for Figure 11.

Would it be possible to match the observed misbehavior in managing TTL with source code of specific software by inspecting for the corresponding bugs?

End of section 5.2: "We detect if a record is in the cache by sending a DNS re- quest for the hostname to the RDNSd and comparing the re- turned time-to-live (TTL) value with the TTL we expect to be set by the Website's ADNS—which we established separately. DNS record TTLs begin counting down when the record enters the cache; once the TTL reaches 0, the record is removed." --> don't you show later in the paper that TTLs from resolvers are unreliable?

Referencing a paper as generically "under submission" is not very useful. Hope in case of acceptance the camera ready gets fixed.

Section 6.3 "The differences with our experiment could be due to different vantage points." -> can you be less generic? On how many observations are their analyses based?

Footnote 5: you scanned the whole Internet. Such accuracy values refer to the USA

Captions explaining the main meaning/take away message of the figures would improve the paper presentation

4. Summary from PC Discussion

In PC discussion, http://openresolverproject.org came up as another pre-existing summary of open resolvers. Their data overlaps with the IMC submission, starting in March 2013. We encourage the paper to cite this work as concurrent work that verifies the 30 million estimate based on sampling with a full scan.

5. Authors' Response

The reviewers expressed concern that the "big picture" goals of our research were not clear. We attempted to clarify the big picture in the paper. In particular we explained why partial scans--and hence effective strategies for undertaking partial scans--are important (sec. 5), and why understanding the caching behavior of resolver infrastructure is important (sec. 7).

The reviewers expressed concern that our estimate of the number of open resolvers may be inaccurate. This stems from a misunderstanding of our probing strategy that we have attempted to clear up in the final paper. In particular, while the reviewer is correct that not all 4B IPv4 addresses are in use, we did not try to distinguish between those in use and those not in use in our probing. Therefore, the right basis for our estimates is in fact the full IPv4 address space. Further, we are happy to see that the PC was able to recognize as inaccurate the review claim that the previous estimate of the number of open resolvers was 92K. Finally, while we were previously unaware of the Open Resolver Project, we added a reference based on the PC's feedback and note that the project's full scan of the Internet closely agrees with our estimates with respect to the overall number of open resolvers.

As reviewers are prone to do, the reviewers desired a deeper analysis of our data. We do not disagree. Our data is indeed rich and there is no doubt additional insight to draw from our various datasets. Obviously, it is impractical to explore all possibilities within the scope of a single conference paper. We intend to continue to mine the data. In addition, the final version of the paper includes a reference researchers can follow to obtain our datasets for their own use.

The reviewers note, correctly, that our conclusion at the end of Sec. 5.2 utilizes TTL responses from RDNSes, which we later find unreliable. We have modified our conclusion (Fig. 7) to consider only the difference in behavior, not absolute numbers. Whatever lies RDNSes use in reporting TTL will manifest in these behaviors.

Finally, we addressed various minor issues raised by the reviewers, which improved the presentation and for which we are grateful.