Consolidated Review of

Speedtrap: Internet-Scale IPv6 Alias Resolution

1. Strengths:

The problem is timely and will be important in the future. The research was thorough (although of limited scope suitable for a short paper) and the paper is well written, with a fair assessment of the limitations as well as strengths.

An important problem, and little work in the area for IPv6. IPv6 alias resolution is an open problem. The authors did large-scale measurement on real IPv6 network.

The authors have/will incorporate this into Ark and regularly produce and make available the datasets, and the technique is part of scamper.

A bunch of neat techniques: tricking routers with the packet size, the overall technique, the fact that you can probe at a slow rate because of the lack of background velocity and the large ID space.

Nice extension of the Midar technique to IPv6, using induced fragmentation to force a router to generate IDs. With recent RFC update there may be no "natural velocity" in IPv4 IP ID generation, raising the importance of techniques such as the one proposed to be used for alias resolution.

2. Weaknesses

It is not clear how IPV6 is widely used in operational network and how representative the router implementation observed in authors' experiments—are and if the same observation hold after IPv6 is seriously and widely used.

A step in a good direction for IPv6 alias resolution, but not much (yet) in the way of scientific contribution. The delta over existing techniques is pretty low at this point. It seems to be at a too-early stage at this point and will need more work to really become an "internet-scale" technique. Even for small IPv6 Internet, technique is pretty slow. The amount of probing required is substantial. The methods described for potential speed-up will only increase speed by a constant factor. As IPv6 network expands, the proposed technique will likely not keep pace.

3. Comments

This paper is very interesting. The use of path-MTUs to induce IPIDs is creative. The measurement of the time for which fragmentation continues is informative, especially since it deviates from the standard. The term "velocity" should be explained when it is first used. The description "debugged speedtrap" on page 1 sounds odd. The description of Step 3 ("distill candidate routers") could be written more clearly.

Are there ways to associate IPv4 addresses on a router with IPv6 addresses?

Abstract: Consider concluding with some accuracy or coverage numbers. Intro: Aliasing is useful for much more than just resilience and robustness.

Intro: 11k aliases struck me as really small. Can you give some context? 2: When you talk about managing the load, I wasn't sure where you need to manage it / what this means. After reading the whole paper, I think you just mean that you need to rate limit your probes, which seems like a minor point.

- 3.1: You eventually talk about which behaviors map to which vendors, but I was already curious about it here and at other points. Consider moving that result earlier. In addition to confirming them in your lab, can you confirm with the operators who provided your ground truth datasets?
- 3.1: Can you speculate whether there is likely to eventually be background velocity?
- 3.2 Step 1: Explain where the 65535 comes from
- 4.1: Could some of the single interfaces ones be incrementing but non-shared IDs?

Table 2: The fractions for Interfaces are not clear. It is the number you observed in your probes? If so, can you take the complete ground truth data and probe them, to remove IP discovery from the variables.

4.3: Why did you use so few Ark nodes? 4.3: Can you contact operators to try to confirm the filters? 6: Have you seen the effects of the Feb 2013 spec update in MIDAR?

I have few concerns about this paper. First, how widely IPv6 is used in operation? If IPv6 is not seriously used in operations by major ISPs, then observation of router implementation how they response to probing packets may not be very representative. Hence, how will the alias resolution technique would work is still undetermined. Second, it seems to me that the speedtrap tool is based on techniques reported in authors' previous work. When applying it to IPv6, it is quite limited since it cannot handle interfaces returning random fragmentation identifiers (they are a large fraction of interfaces in authors data set). I would like to see a solution that can address this issue before claiming a solution for IPv6 alias resolution.

4. Summary from PC Discussion

This is flagged as quick accept. The reviewers who rated themselves as more experienced (3 vs. 2) were also more enthusiastic about the merit (4 vs. 3).

Overall, reviewers thought the paper did a good job adapting a slow IPv6 technique and a scalable v4 technique to arrive at a better v6 technique. However, the work is incremental and not really a full solution as the IPv6 network grows.

Strengths:

- ❖ V6 alias resolution is or will be important, and there has been little work in the area.
- ❖ Large scale measurement, and will be incorporated into Ark to keep fresh data available.
- Nice adaptation of v4 techniques to new challenges raised in v6, and some of the new techniques might end up being useful in v4 too (RFC update)
- * Thorough paper, including limitations

Weaknesses:

- Probing overhead, slow, and therefore not really Internetscale: will not keep pace as IPv6 network expands.
- ❖ As network expands, the technique might not have good coverage if new routers do not support the technique.

- Technique is incremental.
- ❖ Not yet really Internet-scale

5. Authors' Response

We added a section (4.3) which focuses on scalability of speedtrap in response to the concerns about the technique as the set of interfaces to resolve grows. We chose a relatively slow PPS rate in our paper but it is simple to use a faster PPS rate. We emphasized that our validation data suggests a single fragment ID counter is used per router and not per interface, i.e., we observed a single-interface on most routers, rather than a fewer number of routers with per-interface counters. We also fixed typos pointed out by the reviewers, and shifted some text to earlier in the paper as suggested.