Towards a Rigorous Methodology for Measuring Adoption of RPKI Route Validation and Filtering

Andreas Reuter
Freie Universität Berlin
andreas.reuter@fu-berlin.de

Randy Bush
IIJ Research Lab / Dragon Research
randy@psg.com

Italo Cunha
Universidade Federal de Minas Gerais
cunha@dcc.ufmg.br

Ethan Katz-Bassett
Columbia University
ethan@ee.columbia.edu

Thomas C. Schmidt
HAW Hamburg
t.schmidt@haw-hamburg.de

Matthias Wählisch
Freie Universität Berlin
m.waehlisch@fu-berlin.de

ABSTRACT

In this talk, we will report on our recent article "Towards a Rigorous Methodology for Measuring Adoption of RPKI Route Validation and Filtering", published in ACM Computer Communication Review, January 2018. We will also present new results that arise from the ongoing deployment of RPKI route origin validation (*e.g.,* default filtering at IXP route servers), and introduce a publicly available measurement platform for daily monitoring of the state of deployment.

1 INTRODUCTION

The Border Gateway Protocol (BGP) [9] is responsible for establishing Internet routes, yet it does not check that routes are valid. An autonomous system (AS) can hijack destinations it does not control by announcing invalid routes to them, either intentionally or unintentionally.

Because this critical aspect of the Internet is vulnerable, there are proposals to improve routing security [3], and one—the RPKI—is standardized and is in early adoption. The Resource Public Key Infrastructure (RPKI) [6] publishes Route Origin Authorization (ROA) objects, each specifying which AS is allowed to announce an IP prefix. Using ROA data, a BGP router can perform RPKI-based origin validation (ROV) verifying whether the AS originating an IP prefix announcement in BGP is authorized to do so [7] and labeling the route as valid or invalid. The validity of a route can be used as part of the router's local BGP policy decisions, *e.g.,* filtering routes that reflect invalid announcements or preferring valid ones. While the RPKI is fairly populated with ROAs and growing [5, 8, 13, 14], adoption of ROV and filtering has been negligible, according to operator gossip.

To measure RPKI route origin validation, two methods have been proposed. *Uncontrolled experiments* [1] are based on passive observation of routes. *Controlled experiments* [10] overcome limitations of uncontrolled experiments by deploying dedicated experiments that manipulate both BGP announcements and the ROAs that apply to them.

In this talk, we present our verifiable methodology for measuring ROV and report about lessons learned when replicating other ROV measurements [10]. In detail, we show that an uncontrolled experiment to identify ROV adoption has three major limitations. First, its characterizations of some networks change depending on which set of BGP collectors is used, inferring ROV adoption in some cases when it definitely has not been deployed and not inferring it in some cases when it may have been deployed. Second, the approach relies on invalid routes that happen to be announced, and so its coverage is limited by their rare nature [4]. Third, our supplemental measurements suggest that most networks flagged by uncontrolled experiments as using ROV are actually avoiding invalid routes for unrelated (non-security) traffic engineering purposes, without checking ROV status.

2 UNCONTROLLED EXPERIMENTS

Uncontrolled experiments use available BGP dumps and RPKI data to estimate a lower bound for ROV non-adoption and identify ROV filtering [1]. It compares AS paths taken by known ROV valid and known ROV invalid announcements from a single AS to a single vantage point. If the paths differ, it assumes that the invalid announcement was filtered by ROV on the path taken by the valid announcement, causing the divergence. This approach does not distinguish between a single router or an entire AS using ROV-based filtering, ignoring that AS are not atomic [11]. The method analyzes routes exported by vantage points as follows: (*i*) exclude AS observed to use invalid routes, (*ii*) identify AS that may be performing ROV filtering, and (*iii*) select filtering AS when there were seen filtering from at least three vantage points. We now discuss common challenges that have been not addressed.

Impact of Limited Vantage Point (VP) Sets. To quantify the impact of vantage point selection, we choose 44 Routeviews vantage points (the number used in previous work [1]) and calculate the number of AS identified in each step of the method. For each group, results can vary widely depending on which vantage points are used. This clearly illustrates that using BGP RIB dumps as a basis for uncontrolled measurements of ROV filtering (or non-filtering) is problematic. It makes inferences based on routes visible in the selected dumps, but lacks complete visibility of the Internet, leading to misclassification.

Impact of Limited Prefix Visibility at VPs. As the approach uses pairs of non-invalid and invalid announcements, it relies on vantage points receiving such announcements from enough origins to reveal their policies. Combining all dumps from the RIPE RIS and Routeviews projects, we have data from 960 vantage points. But, not all vantage points provide routes to the same set of prefixes. Some vantage points have a near global view, while some have routes for only a very limited number of prefixes. Applying the method with only a subset of VPs as in the previous work [1], selecting vantage points with very limited prefix visibility misses a significant portion of origin AS, and thus underestimates the set of *ROV candidates* and can lead to misclassification.

Impact of Limited Control. Just because a vantage point uses different routes to reach a non-invalid and an invalid prefix from the same origin, it does not imply that the difference is caused by ROV-based filtering. We found traffic engineering as another possible explanation (unrelated to BGP security) for observed differences. For a multi-homed AS, a common technique to influence inbound traffic is to announce different (often overlapping) prefixes to different upstreams. Note that a major cause for invalid BGP announcements is issuing a ROA only for a prefix and then announcing subprefixes [2, 5, 13] which are not covered by the ROA. We found that the majority of AS paths of invalid routes either share the AS path of the covering non-invalid or diverge at the first hop, as would occur with traffic engineering. We also observed a router selecting different routes from the same origin AS due to route age (a BGP tiebreaker).

3 CONTROLLED EXPERIMENTS

Basic Approach. We use the PEERING testbed to make BGP announcements for prefixes we control from PEERING sites around the world to the hundreds of networks it peers with [12]. We use multiple /24 prefixes from the same /16 block. To control ROAs, we run a grandchild RPKI Certificate Authority (CA) in the RIPE region, enabling us to programmatically issue and revoke RPKI certificates and ROAs. To guard against uncommonly long ROA propagation delays, we conservatively keep every configuration (set of BGP announcements and ROA states) in place for eight hours.

In our basic approach, an AS must fulfill two assumptions to allow us to unambiguously determine whether the AS is using ROV-based filtering: (*i*) *connected-assumption.* The network peers with PEERING, either directly or using a route server. (*ii*) *visibility-assumption.* The network offers some means to check the BGP route it uses to reach an Internet destination, either via a Looking Glass or via a vantage point. While the connected assumption is limiting, it is necessary to maintain accuracy, relaxing it to allow networks that are not peers of PEERING introduces ambiguity.

We announce two prefixes via PEERING (AS47065), a *reference prefix* P_R and an *experiment prefix* P_E. We periodically change RPKI state for the experiment prefix, using an additional origin AS to alternate between the following configurations: (C1) P_R and P_E are valid. (C2) P_R is `valid` and P_E is invalid. We check the routes a vantage point chooses to both prefixes during both configurations. The reference prefix always has a *valid* RPKI state so should not be filtered via ROV, and so we omit any vantage points at which P_R is not visible. We expect both prefixes to be treated the same as long as both announcements are valid, and so we omit a vantage point if it uses different routes during configuration C1. Analysing only data from vantage points that pass both these requirements eliminates the problem of **limited visibility**, since there is no missing data anymore. We then check the routes a vantage point has chosen after the announcement of the expermiment prefix becomes *invalid.* Three observations might occur: *(O1) V* has the same route for both prefixes P_E and P_R. *(O2) V* has a different route for prefix P_E. *(O3) V* has no route to P_E.

In the cases of O2 and O3, we know that this route change *must* be because of the RPKI status change. Had it been for another reason we would expect a change in route for the reference prefix as well. The reference prefix combined with the ROA changes thus all but eliminates the problem of **limited control**. The experiments are repeated continuously to confirm the behaviour is consistent.

Results. Our original experiments were performed on February 20-27, May 11-17, and August 1-7, 2017. In our experiments in February and May 2017, we found AS8283, AS50300, and AS59715 were using ROV to filter invalid announcements. The experiments in August show AS50300 and AS59715 to be filtering, but not AS8283. It is worth noting that AS50300 only filtered routes learned via a route server at the Amsterdam exchange (AMSIX), which contradicts one of the assumption in uncontrolled experiments, whereas it is assumed that an AS found on the AS path of an invalid route does not use ROV based filtering. For all three AS we contacted the operators via email and they confirmed that they used ROV based filtering.

Since beginning of 2018, we run our experiments on a daily basis and publish our results on https://rov.rpki.net/. Currently, more than 35 AS deploy RPKI-based filtering; most of them do this implicitly by using ROV on route servers at Internet Exchange Points.

Acknowledgements

We would like to thank the authors of [1] for answering our questions regarding their methodology and data set. We gratefully acknowledge open discussions about RPKI filters with Job Snijders, Antonio Prado (AS59175), and operators of AS50300. Thanks to CAIDA for organizing the BGP hackathon in 2016, which was important to kick this collaboration off; and thanks to the RIPE NCC for providing us temporarily Internet ressources during the beginning of this journey.

REFERENCES

[1] Yossi Gilad, Avichai Cohen, Amir Herzberg, Michael Schapira, and Haya Shulman. 2017. Are We There Yet? On RPKI's Deployment and Security. In *Proc. of NDSS*. ISOC.

[2] Yossi Gilad, Sharon Goldberg, and Kotikalapudi Sriram. 2017. *The Use of Maxlength in the RPKI.* Internet-Draft – work in progress 00. IETF.

[3] Sharon Goldberg. 2014. Why is It Taking So Long to Secure Internet Routing? *Commun. ACM* 57, 10 (Sept. 2014), 56–63.

[4] Ethan Heilman, Danny Cooper, Leonid Reyzin, and Sharon Goldberg. 2014. From the Consent of the Routed: Improving the Transparency of the RPKI. In *Proc. of ACM SIGCOMM*. ACM, New York, NY, USA, 51–62.

[5] Daniele Iamartino, Cristel Pelsser, and Randy Bush. 2015. Measuring BGP route origin registration validation. In *Proc. of PAM (LNCS)*. Springer, Berlin, 28–40.

[6] M. Lepinski and S. Kent. 2012. *An Infrastructure to Support Secure Internet Routing*. RFC 6480. IETF.

[7] P. Mohapatra, J. Scudder, D. Ward, R. Bush, and R. Austein. 2013. *BGP Prefix Origin Validation*. RFC 6811. IETF.

[8] NIST. 2015. NIST RPKI Deployment Monitor. http://rpki-monitor.antd.nist.gov/. (2015).

[9] Y. Rekhter, T. Li, and S. Hares. 2006. *A Border Gateway Protocol 4 (BGP-4)*. RFC 4271. IETF.

[10] Andreas Reuter, Randy Bush, Ítalo Cunha, Ethan Katz-Bassett, Thomas C. Schmidt, and Matthias Wählisch. 2018. Towards a Rigorous Methodology for Measuring Adoption of RPKI Route Validation and Filtering. *ACM Computer Communication Review* 48, 1 (2018), 19–27.

[11] Matthew Roughan, Walter Willinger, Olaf Maennel, Debbie Perouli, and Randy Bush. 2011. 10 Lessons from 10 Years of Measuring and Modeling the Internet's Autonomous Systems. *IEEE Journal on Selected Areas in Communications* 29, 9 (2011), 1810–1821.

[12] Brandon Schlinker, Kyriakos Zarifis, Italo Cunha, Nick Feamster, and Ethan Katz-Bassett. 2014. PEERING: An AS for Us. In *Proceedings of the 13th ACM Workshop on Hot Topics in Networks (HotNets-XIII)*. ACM, New York, NY, USA, 18:1–18:7.

[13] Matthias Wählisch, Olaf Maennel, and Thomas C. Schmidt. 2012. Towards Detecting BGP Route Hijacking using the RPKI. *ACM CCR* 42, 4 (October 2012), 103–104.

[14] Matthias Wählisch, Robert Schmidt, Thomas C. Schmidt, Olaf Maennel, Steve Uhlig, and Gareth Tyson. 2015. RiPKI: The Tragic Story of RPKI Deployment in the Web Ecosystem. In *Proc. of 14th ACM Workshop on Hot Topics in Networks (HotNets)*. ACM, New York, 11:1–11:7.