

Evaluation of the Deployment Status of RPKI and Route Filtering

Johannes Deger, Frank Kargl

Institute of Distributed Systems, Ulm University, Ulm, Germany

{johannes.deger, frank.kargl}@uni-ulm.de

I. INTRODUCTION

The Border Gateway Protocol (BGP) is an essential infrastructure element, often termed “the glue that keeps the Internet together”. Even in its current version 4 [1], BGP misses essential security mechanisms that would allow to validate routing information distributed through BGP in terms of its authenticity and integrity. While mechanisms like BGPsec [2] have been proposed many years ago, so far they have not found widespread adoption and many experts believe they never will due to their inherent complexity [3].

Incidents happening as early as 1997 like AS7007 [4] or the more recent Pakistan YouTube hijack [5] illustrate the problems stemming from BGP route information not being integrity protected and authenticated. In today’s Internet, BGP routing regularly gets manipulated by criminals or state actors with the goal of seizing control of certain portions of address space¹ for criminal or other purposes.

To ensure a minimal level of protection, most Internet service providers (ISPs) rely on heuristic filtering of routing information advertised from neighboring autonomous systems (AS). One approach is called *Path Origin Validation* where an ISP tries to verify whether the AS advertising a certain IP prefix is actually the legitimate owner of this prefix.

II. IRR AND RPKI

Currently, the goto-solution is filtering based on route-objects in so-called Internet Routing Registries (IRR). IRRs can be described as a decentralized and distributed database of text objects regarding information on Internet resources. Different organizations like regional registrars but also private actors run such IRRs and allow participating ISPs to add so-called *route-objects* to them. Essentially, a route-object is a text record that includes an IPv4 or v6 prefix, an origin AS number, and a creator plus optional descriptions. It expresses that the listed AS can act as originator of this prefix in BGP routing, thereby stating a routing-intention. In addition, IRRs can also hold other objects to express routing policies of autonomous systems. However, in most cases no authentication to create objects is required. As we have investigated in our work, data in such IRRs is often badly maintained and thereby IRRs only hold data of questionable quality.

The *Resource Public Key Infrastructure (RPKI)* [6, 7] offers an alternative that is expected to provide better data quality.

It is based on X.509 certificates and a public-key infrastructure run by registrars and offers a cryptographically secured way of distributing routing-intentions from a prefix-owner to all participants in BGP-routing. *Route Origin Authorizations (ROAs)* are the equivalent to IRR’s route objects and contain information that links a prefix to an AS. Additionally they contain the max-length that a prefix can be announced with. In contrast to the IRRs, ROAs are only valid if they are signed with the resource certificate of a party that is authorized to advertise this prefix and thereby form a tangible link between a resource-owner and the routing-intention. It is notable that RPKI’s root of trust is hereby located within the *Regional Internet Registries (RIRs)* that offer RPKI as part of their services. In a future increment, IANA shall hold the root certificate for the PKI. This separates RPKI from other PKIs such as Web certification with its decentralized CAs. The distribution of keys happens through the so called *Trust Anchor Locators (TALs)* where the RIRs publish their key repositories. ISPs can use information from route objects or ROAs to determine their import policies, governing which announcements they are willing to accept from their peers.

III. A COMPARATIVE STUDY FOR IRR AND RPKI

In our work, we asked the question about data quality and coverage of the IP address space of the rather new RPKI compared to the older IRR approach. Earlier works from 2015 and 2016 have come to differing conclusions regarding actual RPKI deployment [8, 9].

To assess this, live Internet-routing is validated via RIPE’s Routing Information Service (RIS) through a newly developed analytics framework (see Figure 1) for routing data and the results are then compared to IRRs yielding a metric of how credible the information in either RPKI and IRRs is. As a side-product it yields a framework capable of processing up to 30.000 routing-messages per second and which we intend to open-source soon²

Additionally the adoption of Route Origin Validation is assessed by injecting own prefix information through a dedicated test AS and verifying its dissemination through the Internet using probes from the RIPE Atlas project.

As a first step, general statistics have been established to describe the development of RPKI’s repository size and a exponential increase in the amount of ROAs was noticeable.

¹<https://mailman.nanog.org/pipermail/nanog/2020-January/105672.html>

²Live routing validation using RIS-Live and IRRs: <https://bit.ly/2UqdCIH>

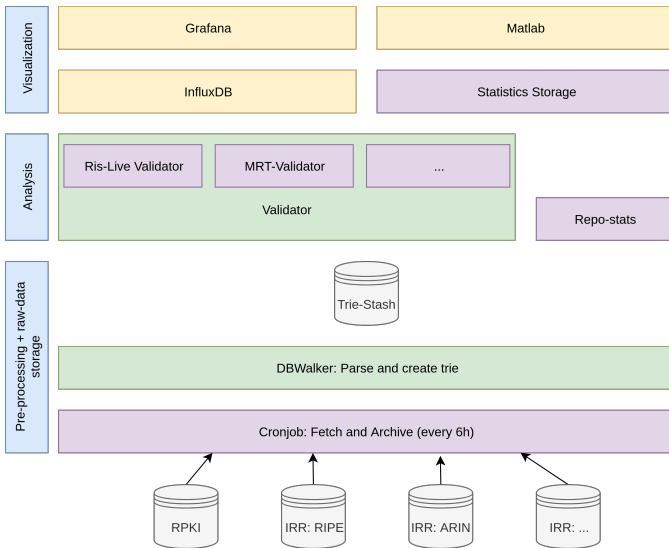


Fig. 1: Architecture of the analysis-framework. The majority of the work is done in the pre-processing consisting of a library to parse route-objects and ROAs and storing them in a Patricia-trie for efficient lookup [10]. Multiple tools use the Validation/Analysis-library to validate routing-data from different sources.

Interpolations show that RPKI will catch up in terms of prefixes covered within approximately 10 years from now as address-space holders adopt the technology.

To assess the quality of ROAs, we looked at the validation of routing-tables using RPKI route origin validation (ROV).

What can be seen is that invalid routes are virtually non-existent and the reason for them mainly were route-leaks and typos in ROAs that eventually got fixed quickly while our study was running. When looking at the valid announcements, we see that most of the ROAs cover the announcement directly, meaning that the received prefix is equal to the prefix in the ROA. Since filtering needs to be as specific as possible to provide adequate immunity to hijacks using de-aggregation attacks [11], this gives rise to the assumptions that network-operators are aware of that and create their ROAs according to their routing-reality.

Having this in mind, IRRs were inspected for their quality by using ROAs and comparing them to the route-objects and to actual BGP routing tables. The findings are very clear: IRRs have substantial quality-problems. Many route-objects contradict ROAs: While the majority of unique prefix/origin-as tuples are equally covered in both datasets (around 65%), many route-objects contradict information in ROAs. This is mostly due to outdated information present in the IRRs, most prominently caused by RADB³ that caused 84% of all conflicts. When validating the routing-tables against route-objects, we see that almost 10% of all routes are invalid according to route-objects. Additionally many of the routes are simply not covered by route-objects (around 13%). So not

³One of the major independent IRRs.

only are there many quality-problems in IRRs, but they don't even cover the whole announced address-space.

Wrapping up all results regarding the quality of ROAs vs the quality of route-objects, it is clear that ROAs are by far the better source of information.

Currently ROAs do not cover as much address-space as route-objects, but even now roughly 20% of all routes seen on the Internet could be successfully validated. Looking at invalids, we see far less RPKI-invalids than IRR-invalids, even though IRRs are the older technology and are currently considered as the goto-solution. The fear of many network-operators of losing routes by implementing ROV seems unaccounted for.

Despite the clear advantages of RPKI over IRR-based filtering, the adoption of active ROV seems to be performed rather hesitantly, since only about 1.4% of networks seen by RIPE-Atlas filtered our invalid announcements. As we see from the results, big transit-providers have the most impact when deploying ROV as they implicitly cover and protect their customers. We have seen that 1.4% filtering networks results in about 6% of coverage for the whole Internet.

On the positive side, many big Tier 1 networks already announced the deployment of ROV, some of them are in the process of deploying ROV on all customer-links (Telia), others are already productive (Hurricane Electric). Not only is RPKI a viable alternative to IRR-based filtering, it is far superior.

With this work, we clearly showed that RPKI is a substantial enhancement over IRR in terms of data quality, that deployment is gaining traction even though large-scale coverage will still take years, but that active ROV based on RPKI ROAs starts to get deployed.

ACKNOWLEDGMENT

This work was supported in the bwNET100G+ project by the Ministry of Science, Research and the Arts Baden-Württemberg (MWK). The authors alone are responsible for the content of this paper.

REFERENCES

- [1] Y. Rekhter, T. Li, and S. Hares, "A border gateway protocol 4 (bgp-4)," Internet Requests for Comments, RFC Editor, RFC 4271, January 2006, <http://www.rfc-editor.org/rfc/rfc4271.txt>. [Online]. Available: <http://www.rfc-editor.org/rfc/rfc4271.txt>
- [2] M. Lepinski and K. Sriram, "Bgpsec protocol specification," Internet Requests for Comments, RFC Editor, RFC 8205, September 2017.
- [3] R. Lychev, S. Goldberg, and M. Schapira, "BGP security in partial deployment: Is the juice worth the squeeze?" *CoRR*, vol. abs/1307.2690, 2013. [Online]. Available: <http://arxiv.org/abs/1307.2690>
- [4] P. G. Neumann. Internet routing black hole. [Online]. Available: <http://catless.ncl.ac.uk/Risks/19.12.html#subj1>
- [5] DYN-Blog. Pakistan hijacks youtube — dyn blog. [Online]. Available: <https://dyn.com/blog/pakistan-hijacks-youtube-1/>
- [6] R. Bush and R. Austein, "The resource public key infrastructure (rpki) to router protocol," Internet Requests for Comments, RFC Editor, RFC 6810, January 2013.
- [7] R. Bush, "Origin validation operation based on the resource public key infrastructure (rpki)," Internet Requests for Comments, RFC Editor, BCP 185, January 2014.
- [8] M. Wählisch, R. Schmidt, T. C. Schmidt, O. Maennel, S. Uhlig, and G. Tyson, "Ripki: The tragic story of rpki deployment in the web ecosystem," in *Proceedings of the 14th ACM Workshop on Hot Topics in Networks*, ser. HotNets-XIV. New York, NY, USA: ACM, 2015, pp. 11:1–11:7. [Online]. Available: <http://doi.acm.org/10.1145/2834050.2834102>

- [9] A. Cohen, Y. Gilad, A. Herzberg, and M. Schapira, "Jumpstarting bgp security with path-end validation," in *Proceedings of the 2016 ACM SIGCOMM Conference*, ser. SIGCOMM '16. New York, NY, USA: ACM, 2016, pp. 342–355. [Online]. Available: <http://doi.acm.org/10.1145/2934872.2934883>
- [10] D. R. Morrison, "Patricia - practical algorithm to retrieve information coded in alphanumeric," *J. ACM*, vol. 15, no. 4, pp. 514–534, Oct. 1968. [Online]. Available: <http://doi.acm.org/10.1145/321479.321481>
- [11] O. Nordström and C. Dovrolis, "Beware of bgp attacks," *SIGCOMM Comput. Commun. Rev.*, vol. 34, no. 2, pp. 1–8, Apr. 2004. [Online]. Available: <http://doi.acm.org/10.1145/997150.997152>