Workgroup: Network Working Group
Internet-Draft:
draft-ietf-sidrops-rpki-rov-timing-06
Published: 7 February 2022
Intended Status: Informational
Expires: 11 August 2022
Authors: R. Bush
Internet Initiative Japan & Arrcus, Inc.
J. Borkenhagen T. Bruijnzeels J. Snijders
AT&T NLnet Labs Fastly

Timing Parameters in the RPKI based Route Origin Validation Supply Chain

Abstract

This document explores, and makes recommendations for, timing of Resource Public Key Infrastructure publication of ROV data, their propagation, and their use in Relying Parties, caches, and routers.

Requirements Language

The key words "MUST", "MUST NOT", "REQUIRED", "SHALL", "SHALL NOT", "SHOULD", "SHOULD NOT", "RECOMMENDED", "NOT RECOMMENDED", "MAY", and "OPTIONAL" in this document are to be interpreted as described in BCP 14 [RFC2119] [RFC8174] when, and only when, they appear in all capitals, as shown here.

Status of This Memo

This Internet-Draft is submitted in full conformance with the provisions of BCP 78 and BCP 79.

Internet-Drafts are working documents of the Internet Engineering Task Force (IETF). Note that other groups may also distribute working documents as Internet-Drafts. The list of current Internet-Drafts is at https://datatracker.ietf.org/drafts/current/.

Internet-Drafts are draft documents valid for a maximum of six months and may be updated, replaced, or obsoleted by other documents at any time. It is inappropriate to use Internet-Drafts as reference material or to cite them other than as "work in progress."

This Internet-Draft will expire on 11 August 2022.

Copyright Notice

Table of Contents

1. Introduction
2. Related Work
3. Certification Authority Publishing
4. Relying Party Fetching
5. Router Updating
6. Effect on Routing
7. Alternative Technologies
8. Operational Expectations
9. Security Considerations
10. IANA Considerations
11. References
  11.1. Normative References
  11.2. Informative References
Appendix A. Acknowledgements
Authors' Addresses

1. Introduction

This document explores, and makes recommendations for, timing of Resource Public Key Infrastructure (RPKI) publication of ROV data, their propagation, and their use in Relying Parties (RP), caches, and routers.

The RPKI ROA supply chain from CAs to when they reach routers has the following structure:

Cerification Authorities: The authoritative data of the RPKI are meant to be published by a distributed set of Certification Authorities (CAs) at the IANA, RIRs, NIRs, and ISPs (see [RFC6481]).

Publication Points: The CAs publish their authoritative data in publicly accessible repositories which have a Publication Point

(PP) for each CA. A CA may publish directly at a PP or may use the RPKI Publication Protocol [RFC8181].

Relying Parties: Relying Parties are a local (to the routers) set of one or more collected and verified caches of RPKI data which the RPs collect from the PPs.

Currently RPs can pull from other RPs, thereby allowing a somewhat complex topology.

Routers: Validating routers fetch data from local RP caches using the RPKI to Router Protocol, [RFC8210] and [I-D.ietf-sidrops-8210bis]. Routers are clients of the caches. Validating routers MUST have a trust relationship with, and a trusted transport channel to, any RP(s) they use. [I-D.ietf-sidrops-8210bis] specifies mechanisms for a router to assure itself of the authenticity of the cache(s) and to authenticate itself to cache(s).

As Resource Public Key Infrastructure based Route Origin Validation (ROV) becomes deployed in the Internet, the quality of the routing control plane, and hence timely and accurate delivery of packets in the data plane, increasingly depend on prompt and accurate propagation of the RPKI data from the originating Certification Authorities (CAs), to Relying Parties (RPs), to Border Gateway Protocol (BGP) speaking routers.

Origin Validation based on stale ROAs allows accidental or intentional mis-origination; announcement of a prefix by an AS which does not have the authority to do so. Delays in ROA propagation to ROV in routers might cause loss of good traffic. Therefore minimizing propagation time of data from CAs to routers is important.

Before the data can start on the CA to router supply chain, the resource holder (operator) MUST create, modify, or delete the relevant ROA(s) through the CA's operational interface(s). The operator is responsible for anticipating their future needs for ROAs, be aware of the propagation time from creating ROAs to effect on routing, and SHOULD create, delete, or modify ROAs sufficiently in advance of any needs in the routing system.

There are questions of how frewwww3quently a CA publishes, how often an RP pulls, and how often routers pull from their RP(s). Overall, the router(s) SHOULD react within an hour of ROA publication. In pessimistic circumstances, it could be two hours.

For CAs publishing to PPs, a few seconds to a minute seems easily achieved with reasonable software. See Section 3.

Relying Party validating caches periodically retrieve data from CA publication points. RPs using rsync to poll publication points every ten minutes would be a burden today, given the load it would put on publication services, cf. one notorious repository which was structured against specification. RPs using RRDP impose less load. As the infrastructure moves from rsync to RRDP [I-D.ietf-sidrops-prefer-rrdp], RRDP is designed for quite frequent polling as long as Relying Parties use the If-Modified-Since (see [RFC7232]) header and there is a caching infrastructure. For rsync, an hour would be the longest acceptable window and half an hour the shortest. See Section 4.

For BGP speaking router(s) pulling from the RP(s), five minutes to an hour is a wide window. But, the RPKI-Rtr protocol does have the Serial Notify PDU, the equivalent of DNS Notify [RFC1996], where the cache tells the router that it has new data. See Section 5.

We discuss each of these in more detail below.

2. Related Work

It is assumed that the reader understands BGP, [RFC4271], the RPKI [RFC6480], RPKI Manifests [RFC6486], Route Origin Authorizations (ROAs), [RFC6482], the RPKI Repository Delta Protocol (RRDP) [RFC8182], The Resource Public Key Infrastructure (RPKI) to Router Protocol [I-D.ietf-sidrops-8210bis], RPKI-based Prefix Validation, [RFC6811], and Origin Validation Clarifications, [RFC8481].

3. Certification Authority Publishing

One constraint on publication timing can be ensuring the CRL and Manifest ([RFC6486]) are consistent with each other and with respect to the other repository data. With both rsync and RRDP protocols, the publication point MUST be consistent before it becomes current and is published.

Operators should beware that there may be implementation dependent delays between instructing their CAs to create and/or update ROAs and the publication of these changes in the PPs.

4. Relying Party Fetching

rsync puts a load on RPKI publication point servers. Therefore relying party caches have been discouraged from fetching more frequently than on the order of a half hour. Times as long as a day were even suggested. We specify that RPs using rsync SHOULD pull from CA publication points every 30 to 60 minutes.

With RRDP ([RFC8182]), such constraints can be less relevant. [RFC8182] makes clear that polling as frequently as once a minute is acceptable if and only if Relying Parties use the If-Modified-Since

header and there is caching. Absent use of the If-Modified-Since header, the RRDP polling interval MUST NOT be more frequent than ten minutes. Use of the If-Modified-Since header is strongly RECOMMENDED.

Migration from rsync to RRDP in [I-D.ietf-sidrops-prefer-rrdp] is recommended. During dual RRDP/rsync operation, should an RP need to fall over from RRDP to rsync, a uniformly distributed jittered delay with a mean of half the rsync interval SHOULD be used; so clients falling over to rsync are as spread out as they would be if they used rsync initallly.

A number of timers are embedded in the X.509 RPKI data which should also be considered. E.g., CRL publication commitments, expiration of EE certificates pointing to Manifests, and the Manifests themselves. Some CA operators commonly indicate new CRL information should be available in the next 24 hours. These 24 hour sliding timers, when combined with fetching RPKI data once a day, would expose failure windows, especially in the face of transient network issues between the CA and RP. To ameliorate this, RPs SHOULD update from CAs at least as frequently as once an hour.

In summary, the following timing constraints SHOULD be applied to data update: RPs SHOULD update from CAs at least once an hour. To avoid excess load, RPs SHOULD NOT update via rsync more frequently than every 30 minutes. RPs using RRDP SHOULD NOT need to update more frequently than every 10 minutes. Some form of timing jitter MUST be applied to ensure load distribution across the community. RPs SHOULD NOT force data fetch to be on the hour or similar times. Publication Points SHOULD deploy RRDP services which honor If-Modified-Since.

In general, CAs should have Manifest, CRL, ... timers of a few days to allow relying party operators to go away for the weekend and not fear for their control plane.

5. Router Updating

The rate of change of ROA data is estimated to remain small, on the order of a few ROAs a minute, but with bursts. Therefore, the routers may update from the (presumed local) relying party cache(s) quite frequently. Note that [I-D.ietf-sidrops-8210bis] recommends a polling interval of one hour. This polling timing is conservative because caches can send a Serial Notify PDU to tell routers when there are new data to be fetched. As the RP cache and the router belong to the same operator, routers are free to hammer the RPs as frequently as they wish.

A router SHOULD respond with a Serial Query when it receives a Serial Notify from a cache. If a router can not respond appropriately to a

Serial Notify, then it MUST send a periodic Serial Query no less frequently than once an hour.

6. Effect on Routing

Once a router has received an End of Data PDU from a cache, the effect on Route Origin Validation MUST be a matter of seconds to a minute. The router MAY allow incoming VRPs to affect Origin Validation as they arrive instead of waiting for the End of Data PDU. See [I-D.ietf-sidrops-8210bis] for some cautions regarding the arrival and processing sequence of VRPs.

7. Alternative Technologies

Should the supply chain include components or technologies other than those in IETF documents, the end effect SHOULD be the same; the router(s) SHOULD react to invalid AS origins within the same overall time constraint, one hour, two at most, from ROA creation at the CA publication point to effect in the router.

8. Operational Expectations

Assuming the above recommendations, in worst conditions such as an RPKI-rtr Notify PDU being ignored, it may take up to two hours for a new ROA to propagate from creation at the CA to BGP speaking routers. Therefore it is RECOMMENDED that planned changes in ROAs take this propagation time into consideration. E.g. if a new route is to be announced in BGP, the operators SHOULD create the ROA around three hours before BGP announcement, or it may not propagate globally.

9. Security Considerations

Despite common misconceptions and marketing, Route Origin Validation is not a magic security protocol. It is intended to catch operational errors, and is easily gamed and attacked through, for example, AS Path manipulation. It is one tool in the prudent operator's kit, and a good one.

If an attacker can add, delete, or modify RPKI data, either in repositories or in flight, they can affect routing and thereby steer or damage traffic. The RPKI system design does much to deter these attacks. But the 'last mile' from the cache to the router uses transport, as opposed to object, security and is vulnerable. This is discussed in [I-D.ietf-sidrops-8210bis].

Similarly, if an attacker can delay prompt propagation of RPKI data on the supply chain described in this document, they can affect routing, and therefore traffic flow, to their advantage.

10. IANA Considerations

None

11. References

11.1. Normative References

[I-D.ietf-sidrops-8210bis] Bush, R. and R. Austein, "The Resource Public Key Infrastructure (RPKI) to Router Protocol, Version 2", Work in Progress, Internet-Draft, draft-ietf-sidrops-8210bis-05, 22 December 2021, <https://www.ietf.org/archive/id/draft-ietf-sidrops-8210bis-05.txt>.

[RFC1996] Vixie, P., "A Mechanism for Prompt Notification of Zone Changes (DNS NOTIFY)", RFC 1996, DOI 10.17487/RFC1996, August 1996, <https://www.rfc-editor.org/info/rfc1996>.

[RFC2119] Bradner, S., "Key words for use in RFCs to Indicate Requirement Levels", BCP 14, RFC 2119, DOI 10.17487/RFC2119, March 1997, <https://www.rfc-editor.org/info/rfc2119>.

[RFC4271] Rekhter, Y., Ed., Li, T., Ed., and S. Hares, Ed., "A Border Gateway Protocol 4 (BGP-4)", RFC 4271, DOI 10.17487/RFC4271, January 2006, <https://www.rfc-editor.org/info/rfc4271>.

[RFC6481] Huston, G., Loomans, R., and G. Michaelson, "A Profile for Resource Certificate Repository Structure", RFC 6481, DOI 10.17487/RFC6481, February 2012, <https://www.rfc-editor.org/info/rfc6481>.

[RFC6482] Lepinski, M., Kent, S., and D. Kong, "A Profile for Route Origin Authorizations (ROAs)", RFC 6482, DOI 10.17487/RFC6482, February 2012, <https://www.rfc-editor.org/info/rfc6482>.

[RFC6486] Austein, R., Huston, G., Kent, S., and M. Lepinski, "Manifests for the Resource Public Key Infrastructure (RPKI)", RFC 6486, DOI 10.17487/RFC6486, February 2012, <https://www.rfc-editor.org/info/rfc6486>.

[RFC6811] Mohapatra, P., Scudder, J., Ward, D., Bush, R., and R. Austein, "BGP Prefix Origin Validation", RFC 6811, DOI

10.17487/RFC6811, January 2013, <https://www.rfc-editor.org/info/rfc6811>.

[RFC7232] Fielding, R., Ed. and J. Reschke, Ed., "Hypertext Transfer Protocol (HTTP/1.1): Conditional Requests", RFC 7232, DOI 10.17487/RFC7232, June 2014, <https://www.rfc-editor.org/info/rfc7232>.

[RFC8174] Leiba, B., "Ambiguity of Uppercase vs Lowercase in RFC 2119 Key Words", BCP 14, RFC 8174, DOI 10.17487/RFC8174, May 2017, <https://www.rfc-editor.org/info/rfc8174>.

[RFC8181] Weiler, S., Sonalker, A., and R. Austein, "A Publication Protocol for the Resource Public Key Infrastructure (RPKI)", RFC 8181, DOI 10.17487/RFC8181, July 2017, <https://www.rfc-editor.org/info/rfc8181>.

[RFC8182] Bruijnzeels, T., Muravskiy, O., Weber, B., and R. Austein, "The RPKI Repository Delta Protocol (RRDP)", RFC 8182, DOI 10.17487/RFC8182, July 2017, <https://www.rfc-editor.org/info/rfc8182>.

[RFC8210] Bush, R. and R. Austein, "The Resource Public Key Infrastructure (RPKI) to Router Protocol, Version 1", RFC 8210, DOI 10.17487/RFC8210, September 2017, <https://www.rfc-editor.org/info/rfc8210>.

[RFC8481] Bush, R., "Clarifications to BGP Origin Validation Based on Resource Public Key Infrastructure (RPKI)", RFC 8481, DOI 10.17487/RFC8481, September 2018, <https://www.rfc-editor.org/info/rfc8481>.

11.2. Informative References

[I-D.ietf-sidrops-prefer-rrdp] Bruijnzeels, T., Bush, R., and G. Michaelson, "Resource Public Key Infrastructure (RPKI) Repository Requirements", Work in Progress, Internet-Draft, draft-ietf-sidrops-prefer-rrdp-01, 22 October 2021, <https://www.ietf.org/archive/id/draft-ietf-sidrops-prefer-rrdp-01.txt>.

[RFC6480] Lepinski, M. and S. Kent, "An Infrastructure to Support Secure Internet Routing", RFC 6480, DOI 10.17487/RFC6480, February 2012, <https://www.rfc-editor.org/info/rfc6480>.

Appendix A. Acknowledgements

The authors wish to thank George Michaelson, Massimiliano Stucchi and Ties de Kock.

Authors' Addresses

Randy Bush
Internet Initiative Japan & Arrcus, Inc.
5147 Crystal Springs
Bainbridge Island, Washington 98110
United States of America

Email: randy@psg.com

Jay Borkenhagen
AT&T
200 Laurel Avenue South
Middletown, NJ 07748
United States of America

Email: jayb@att.com

Tim Bruijnzeels
NLnet Labs

Email: tim@nlnetlabs.nl
URI: https://www.nlnetlabs.nl/

Job Snijders
Fastly
Amsterdam
Netherlands

Email: job@fastly.com