sidrops
Internet-Draft
Intended status: Informational
Expires: 22 October 2026

J. Zhang
Zhongguancun Laboratory
M. Xu
CERNET
N. Geng
Huawei
C. Xie
China Telecom
Y. Wang
Tsinghua University

20 April 2026

RPKI-based Validation with Prioritized Resource Data
draft-zhang-sidrops-prioritized-route-validation-01

Abstract

RPKI ROAs and other digitally signed objects provide a practical solution to validate BGP routes for preventing route hijacks and leaks. During ROV operations, validation data may be sourced not only from issued ROAs but also from other local sources. As these data sources vary in credibility, ROV operations may therefore require different response actions for invalid or unknown routes. This document takes ROV as an example to describe a flexible RPKI-based route validation mechanism with multi-priority data, and outlines relevant use cases, framework, and requirements for ROV operations that involve multi-priority data.

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 22 October 2026.

Copyright Notice

Copyright (c) 2026 IETF Trust and the persons identified as the document authors. All rights reserved.

This document is subject to BCP 78 and the IETF Trust's Legal Provisions Relating to IETF Documents (https://trustee.ietf.org/license-info) in effect on the date of publication of this document. Please review these documents carefully, as they describe your rights and restrictions with respect to this document. Code Components extracted from this document must include Revised BSD License text as described in Section 4.e of the Trust Legal Provisions and are provided without warranty as described in the Revised BSD License.

Table of Contents

1. Introduction 2
1.1. Requirements Language 3
2. Gap analysis 3
3. Framework 3
4. Requirements for Multi-Priority RPKI ROV 5
5. Security Considerations 7
6. IANA Considerations 7
7. References 7
7.1. Normative References 7
7.2. Informative References 8
Authors' Addresses 8

1. Introduction

RPKI Route Origin Authorizations (ROAs) and other digitally signed objects provide a practical solution to validate BGP routes for preventing route hijacks and leaks. For example, Route Origin Validation (ROV) built on ROAs stands as a practical and effective approach to combat prefix origin hijacking. In ROV operations, the validating data utilized is not limited to ROAs alone; it may also include various types of local data from other sources. These additional data sources can exhibit varying levels of credibility, with some being highly reliable due to their authoritative origins while others being less credible due to potential inconsistencies or lack of verification. Correspondingly, ROV operations require sufficient flexibility to adopt distinct actions even for the same type of invalid routes, as well as for unknown routes. This document takes ROV as an example to describe a flexible RPKI-based route mechanism with multi-priority data, and elaborates the gap analysis, framework, and specify the key requirements for implementing ROV with multi-priority data in current RPKI infrastructure.

1.1. 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.

2. Gap analysis

RPKI-based Route Origin Validation (ROV) fundamentally relies on Route Origin Authorization (ROA) data. However, it is recognized that ROA deployment has the following deployment issues: incomplete coverage of routes in the global routing table, the existence of ROAs with erroneous registrations, and the common practice where network operators locally filter certain Validated ROA Payload (VRP) data or supplement it with external sources (e.g., data inferred through machine learning). Technologies like SLURM (Simplified Local Internet Number Resource Management) can be used to apply such local modifications.

Due to such mixed data sources, the credibility of the data varies to different degrees and is no longer uniform. Accordingly, network operators require the ability to apply different actions to validation results based on the credibility of the underlying data. This approach offers benefits such as the following use case: an ISP (Internet Service Provider) might supplement RPKI ROAs with data derived from its own operational experience, but assign a lower credibility to this supplementary data. In this scenario, when a route is validated as "invalid" by RPKI ROAs, the router can discard the route; if a route is validated as "invalid" by supplementary data with medium credibility, the router can be configured to trigger an alert.

However, current RPKI technology does not support such operations. Regardless of the source of the validation data, the same type of validation result triggers the same operation. This fails to meet the need for customized processing of different scenarios in network operations. This document describes a RPKI validation framework and requirements with multi-priority data to make RPKI-based validation processing more flexible and operable.

3. Framework

This document proposes a framework shown as the following figure. It sopports RPKI-based validation with prioritized resource data.

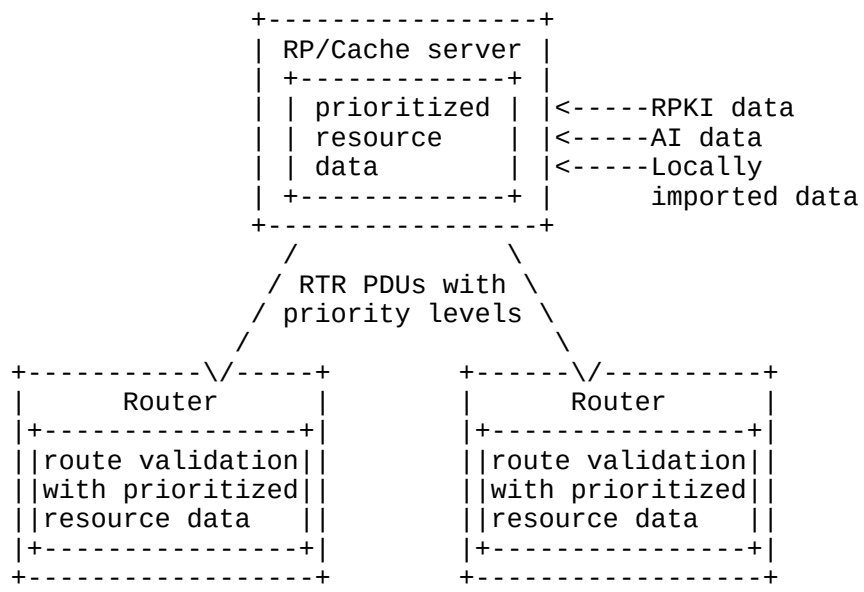

The RP/Cache server collects and manages resouce data (e.g., ROA/ other digitally signed objects, like ASPA) which are from different sources such as RPKI repository, AI inference, and local import. The data from different sources will be set to different priorities. The RP/Cache server needs to decide how to merge these data from different sources.

The data will be syncronized from the RP/Cache server to routers through tools like RTR. Routers will do BGP route validation with priorities being taken into consideration. Particularly, the validaiton output for a route can be Valid, Unknown, or Invalid-_level_. A route validated as invalid will be marked with Invalid as well as a credibility level of the validation result. For example, "Invalid-_1_" means the validation result of invalid is derived based on the source data with the priority of _1_ and thus has a credibility level of _1_.

Network operators can do configurations on routers to taken different policies on the invalid routes with different credibility levels. For example, suppose there are two route origin validation records on a router: (1) the prefix of 192.0.1.0/24 is originated from AS 65001, which has a high priority of _1_ and (2) the prefix of 192.0.2.0/24 is originated from AS 65002, which has a relatively low priority of _2_. For the route with the prefix of 192.0.1.0/24 and the origin of AS 65003, the validation result will be invalid-_1_. Operators can configure the router to discard the route because the validation result is with a high credibility level. For the route with the

prefix of 192.0.2.0/24 and the origin of AS 65003, the validation result will be invalid-_2_. Operators can choose to set a lower priority for this route to influence the route selection outcome. This is because the validation result is with a relatively low credibility level and adopting a conservative handling policy to the route may be safer.

The proposed framework brings two main benefits:

* Enhancing BGP route handling after route validation. When the resource data are from multiple data sources with different levels of credibility, operators can implement customized priority settings to the resource data and apply different handling policies. The multiple priority settings can be extended into a Generic Tagging Mechanism (conceptually similar to BGP Communities) to express flexible handling policies after route validation.

* Improving early deployment benefits and promoting the deployment of RPKI-based routing validation techniques. When the registration rate of RPKI data is not high, operators can supplement data with techniques like AI while still being able to take "discarding" action on invalid routes with high credibility levels, and take "alerting" action on invalid routes with lower credibility levels.

4. Requirements for Multi-Priority RPKI ROV

This section outlines the requirements for extending the RPKI architecture to support the processing and propagation of RPKI data with multiple priority levels. These requirements are necessary to enable differentiated handling of routing validation results based on their perceived credibility, such as those derived from authoritative sources (e.g., RPKI ROAs) versus inferred or supplemental sources (e.g., AI-inferred data).

1. Priority Setting. Implementations processing RPKI data on local cache (e.g., Relying Party software) SHOULD support the assignment of a priority level to each validated RPKI object. The priority MAY be configurable based on the data source (e.g., RPKIsigned, locally imported, or AI-inferred). The priority value SHOULD be represented in a standardized format to ensure interoperability.

2. Multi-Priority Data Merge. Implementations SHOULD merge data from multiple sources according to a defined algorithm. This algorithm SHOULD specify how to handle conflicts, including rules for merging data of the same priority and for superseding lower-priority data with higher-priority data.

3. SLURM Support for Priority Marking. The SLURM mechanism SHOULD be extended to allow local exceptions and additions to include a priority attribute. This enables network operators to override or supplement RPKI data with local policies that reflect differentiated credibility.

4. RTR Support for Priority Marking. The RPKI-to-Router (RTR) protocol MUST be extended to convey the priority of the validation data it delivers. This enables routers to apply appropriate local policy based on the credibility of the origin. To provide flexibility in deployment, two implementation models MAY be supported. Network operators SHOULD choose which implementation to deploy based on their specific operational preferences and infrastructure:

   4.1. One single local cache server transmits data of multiple priority levels. The protocol SHOULD be extended to include a new field within its Protocol Data Units (PDUs) to explicitly carry the priority level for each payload data item. This model allows a router to maintain a single, simple transport session with one cache server while receiving a mixed-priority data set.

   4.2. A router establishes transport sessions with multiple local cache servers, where each server is designated to provide data for a specific priority level (e.g., a primary server provides high-priority RPKI-validated data, and a secondary server provides low-priority supplemental data). The protocol itself remains unchanged, as the priority is derived from the configuration of the router-to-server association. When adopting this implementation, routers must be aware of data priority (at least data sources) to execute differentiated policies (e.g., Drop vs. Alert).

5. BGP route validation (e.g., ROV, ASPA, etc.) with Priority Awareness. BGP route validation processes SHOULD be enhanced to support a multi-priority data model. The validation process SHOULD annotate the resulting validation state (Valid, Invalid) with an indication of the priority level, which identifies the priority tier of the data that was used to reach that conclusion. For example, an "Invalid" result derived from a local override (high-priority) MUST be distinguishable from an "Invalid" result derived from inferred data (low-priority). This annotated

validation state SHOULD then be used to inform subsequent routing policy actions. Implementations SHOULD provide flexible policy mechanisms that allow network operators to define actions (e.g., reject, depreference, warn, accept) based on both the validation state (e.g., Invalid) and its associated assurance level (e.g., High or Low).

6. Router Handling of Priority-Based Invalid Routes. BGP speakers SHOULD support configurable policies to handle invalid routes based on the priority of the validation data. For example, routes invalidated by high-priority data MAY be deprioritized or discarded, while those invalidated by low-priority data MAY be retained with a warning.

7. BMP Support for Priority in Validation Reports. The BGP Monitoring Protocol (BMP) SHOULD be extended to include priority information in reports of BGP route validation results. This enables network operators to monitor and analyze routing decisions based on data credibility.

5. Security Considerations

This document defines a framework for handling RPKI data with multiple levels of priority (assurance), which introduces new considerations beyond those of the base RPKI system [RFC6480].

Amplification of RPKI Repository Failures: If a high-priority source (e.g., a primary RTR cache) becomes stale or unavailable, the system may fall back to low-priority data. This could lead to a mass re-evaluation of routes from a 'Unknown' state to a 'Valid' or 'Invalid' state based on less credible information, potentially causing widespread routing churn. Implementations should include mechanisms to detect such scenarios and allow operators to define appropriate fallback behaviors.

6. IANA Considerations

This document has no IANA action.

7. References

7.1. Normative References

[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>.

[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>.

7.2. Informative References

[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>.

Authors' Addresses

Jia Zhang
Zhongguancun Laboratory
Beijing
China
Email: zhangj@zgclab.edu.cn

Mingwei Xu
CERNET
Beijing
China
Email: xmw@cernet.edu.cn

Nan Geng
Huawei
Beijing
China
Email: gengnan@huawei.com

Chongfeng Xie
China Telecom
Beijing
China
Email: xiechf@chinatelecom.cn

Yangyang Wang
Tsinghua University
Beijing
China
Email: wangyy@cernet.edu.cn