# RPKI-based Validation with Prioritized Resource Data

## draft-zhang-sidrops-prioritized-route-validation-00


| | |
|---|---|
| **Jia Zhang** | ZGC Lab |
| Mingwei Xu | CERNET |
| Nan Geng | Huawei |
| Chongfeng Xie | China Telecom |
| Yangyang Wang | Tsinghua University |



IETF 125 SIDROPS

March 2026

# Problem Statement

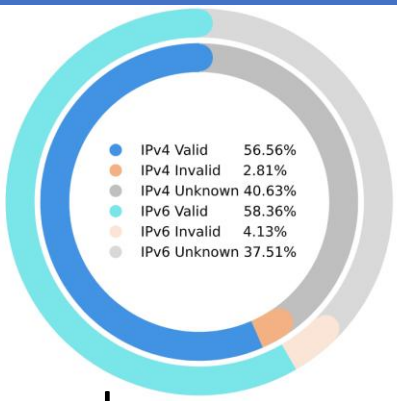


- Complex Operational Reality:
  - **Incomplete ROA Coverage**: A significant portion of routes remain "Unknown" due to low RPKI adoption.
  - **Erroneous ROA Registrations**: Incorrect ROAs cause legitimate routes to be flagged as "Invalid."
  - **Heterogeneous Data Exists**: Operators have valuable local data (e.g., AI inference, local policies, historical analysis) that could enhance validation (through SLURM).
- Rigid Current System:
  - An "Invalid" result triggers the same action (e.g., discard) regardless of whether it comes from a highly-trusted RPKI ROA or a less-trusted local supplement.
  - This forces a binary choice: either fully trust or completely ignore supplementary data.

# Core Idea: Introducing Priority

- Key Insight: Assign a Priority Level to every piece of validation data.
  - *Example:* RPKI-signed ROA -> **Priority 1 (High)**
  - *Example:* AI-inferred or local policy -> **Priority 2 (Medium)**
- Fundamental Shift:
  - Validation states evolve from Valid/Invalid/Unknown to Valid/Invalid-1/Invalid-2/Unknown.
  - The result now carries metadata about its own credibility.
  - **Routing policy can be fine-tuned based on this credibility level.**

# Use Case

- Scenario: A router holds two ROA records:
  - *192.0.1.0/24 -> AS 65001* (Priority 1, from Global RPKI)
  - *192.0.2.0/24 -> AS 65002* (Priority 2, from AI Inference)
- Validation & Action:
  - Route A: *192.0.1.0/24* from *AS 65003* -> Invalid-1
    - Action: DISCARD. High confidence of hijack. Strict policy applied.
  - Route B: *192.0.2.0/24* from *AS 65003* -> Invalid-2
    - Action: DEMOTE or WARN. Lower confidence. Apply a conservative policy (e.g., lower Local-Pref) to maintain connectivity safely.
- Conclusion: Enables **granular**, **risk-aware** routing security policies.

# Framework

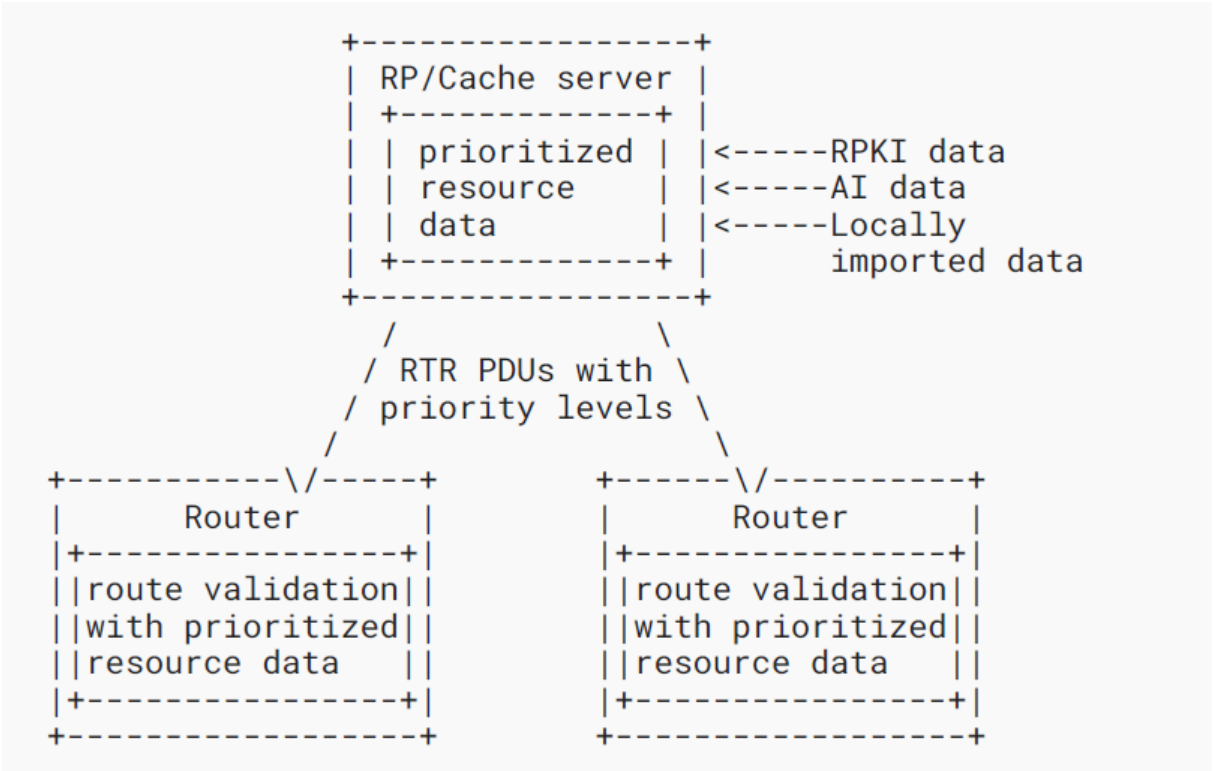


- Component Flow:
  - **RP/Cache Server**: Aggregates data from different sources(e.g., RPKI, AI, and local imports). Assigns priority levels and merges data based on a defined algorithm.
  - **RTR Protocol (Extended)**: Carries priority information within its PDUs to routers.
  - **Router**: Performs priority-aware BGP route validation. Produces annotated results (e.g., Invalid-1) and executes configured policies based on both state and credibility.

# Protocol Extension Requirements

- **Priority Setting & Data Merge:** RP SHOULD be able to assign and merge data from different sources, resolving conflicts.
- **SLURM Extension:** SLURM format SHOULD be extended with a priority attribute for local exceptions and additions.
- **RTR Extension (Core):** Two flexible implementation models are proposed:
  - Model A (Integrated): Extend RTR PDU with a new field to carry priority for each data item. Simple, one-session model.
  - Model B (Segregated): Router establishes multiple RTR sessions to different caches, each dedicated to a priority level. Protocol-agnostic, uses existing RTR.
- **Router Policy (Core):** Routers MUST use the annotated validation state (Invalid-Level) to apply configurable actions (discard, demote, warn, etc).

# Key Benefits

- **Enhanced Operational Flexibility**: Moves Route Validation from a binary to a multi-level decision process. Operators can define nuanced policies that match their risk tolerance.

- **Accelerates RPKI Deployment**: Lowers the barrier to entry. Operators can supplement low-coverage RPKI data with other sources without losing the ability to act decisively on high-confidence invalid routes.

# Key Feedbacks from the mailing list

- **Feedback1, from Maria Matejka**
  - "BIRD is already doing this by simply allowing multiple RTR, multiple tables and multiple ROA evaluations, treating different ROA valid/unknown/invalid results differently based by the source. "
  - "That approach allows for clear data source splits without updating the RTR protocol..."
  - "I would be very careful with creating more validation results."

- **Response1:**
  - Glad to see existing software (BIRD) validating the operational need for this!
  - Agree on keeping the base validation algorithm simple.

# Key Feedbacks from the mailing list

- **Feedback2, from Nick Hilliard**
  - "There is a cost associated with adding complexity..."
  - "If this prioritization mechanism were to be implemented solely in the local cache, it would provide most of the functionality..."

- **Response2:**
  - Zero Modification Option: Draft Sec 4.2 proposes a Multi-Session approach, fits existing implementations perfectly.
  - Why Routers: Routers must be aware of data confidence levels (at least data sources) to execute differentiated policies (e.g., Drop vs. Alert).

# Key Feedbacks from the mailing list

- **Feedback3, from William McCall**
  - "I'm not in favor of the narrow scope... For this to be worthwhile, I want much more flexibility/more generic."
  - "So think something more like communities than verdicts or priorities or whatever.“

- **Response3:**
  - Evolving this into a Generic RTR Tagging Mechanism (conceptually similar to BGP Communities) is a great idea!
  - Broader Community Benefit: not only solves our "mixed-trust" data scenario but also provides the broader community with a highly flexible tool for extensible routing policies.

# Key Feedbacks from the mailing list

- **Feedback4, from Ralph Covelli**
  - "Trust should be binary. Yes or No... RPKI should only be concerned with cryptographically validated information... "
  - "Imagine a worst case scenario where a rogue AI agent uses this to hijack traffic."

- **Response4:**
  - Cryptography remains authoritative, but deployment gaps exist (e.g., in early deployment). SLURM already sets a precedent for local data intervention.
  - Clarification: Prioritization actually protects the network. It ensures inferred/supplemental data (lower trust) cannot cause aggressive actions (like dropping routes) that signed ROAs can.

# Next Steps

- We propose a flexible framework that introduces fine-grained, priority-based control to the RPKI ecosystem.
- Feedback and comments are welcome.
  - We seek feedback and support from the SIDROPS WG.
  - We welcome discussion on the protocol extension details (e.g., RTR).
- Thank You & Q&A

# Thank you !

**Special thanks to Nick Hilliard, William McCall, Maria Matejka, and Ralph Covelli for their valuable suggestions and comments.**

RPKI-based Validation with Prioritized Resource Data

zhangj@zgclab.edu.cn

IETF 125 SIDROPS Meeting