
Operational issues - RPKI ROV

Bad data from Origin

M. Huang, N. Geng, D. Li, S. Yue
ZGCLab, Huawei, Tsinghua University, China Mobile

Mar. 2026
IETF 125@Shenzhen

Contents

□ Problem Statement

- ◆ if RPKI ROV is wrong from source ... What do I do?

□ Investigated the problem

- ◆ Progress of RPKI ROV
- ◆ Long-Standing Challenges
- ◆ Underlying Systemic Issues

Cause 1: Bad RPKI Data from Me

- “My boss holds me responsible if there is Bad RPKI data”

Cause 1: Bad data is originated from my network

- ◆ What do other people use to validate at the origin?
- ◆ What happens if you find that data in RPKI ROV you originated is bad?
- ◆ How quickly can it be fixed?
- ◆ Trade-offs between “no data vs “bad data” from leaf AS

Bad or no Data from Remote AS

“My boss does realize I get bad data from Leaf AS”

Cause 2: RPKI is 60% of routes, 20% of AS – Leaf AS have less.

□ Bad data originated from Leaf AS

- ◆ Must be fixed at origin
- ◆ Need suggestions on tools or ideas

If we get Bad Data from the remote AS, what happens

- ◆ False positives cause problems
- ◆ Impact of no RPKI ROV coverage
- ◆ 3 systemic reasons for Remote

Cost Remote AS sending Bad data

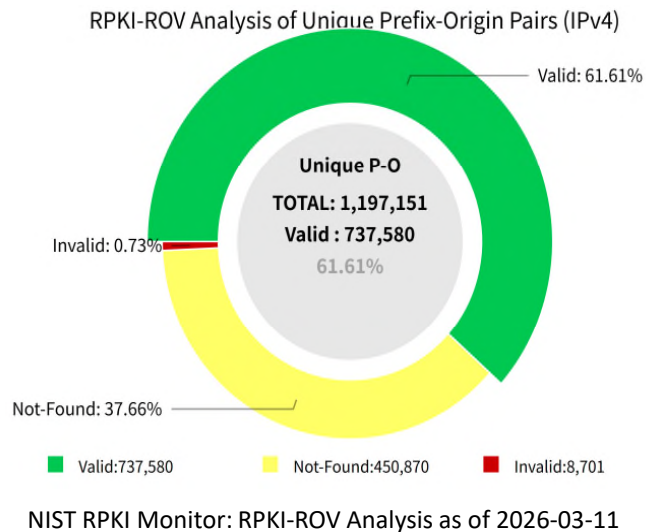
- Bad data from the origin (based on double-checking)
 - ◆ False positives cause problems in
 - ◆ Holes in RPKI ROV coverage impact

- Leaf AS issue
 - ◆ Lack of resources for some small networks – causes misconfiguration
 - ◆ Lack benefit for small networks

More Details

Details if we have time or you are interested

RPKI ROV – Significant Progress



- **ROA Coverage**


- NIST RPKI Monitor: **>60% of IPv4 routes** now have ROAs (as of Mar. 2026).
- Traffic-weighted coverage even higher: Kentik data shows >70% of Internet traffic flows to ROA-covered prefixes (As of May 2024, when ROA coverage was ~50%).

- **ROV Deployment**

- More and more tier-1 and large transit providers reject invalids (e.g., Sparkle, Deutsche Telekom, Bell Canada – IPinfo 2026).
- ATHENE (NDSS 2026): **~27% of ASes** observably deploy ROV in a way that affects routing decisions.

The foundation is strong. But coverage is only half the story, when we look beneath the surface, two long-standing challenges emerge.

Challenge 1: The Persistent False Positive Problem



- NIST RPKI Monitor: RPKI-invalid prefixes (IPv4) have **remained around 1% for years**, often thousands of routes.
- Studies (NDSS 2026, “Demystifying RPKI-Invalid Prefixes”): **>96% of these are false positives** caused by **misconfiguration**, not hijacks.

Despite increased ROV enforcement, the scale of invalid routes and the false positive rate have not improved. This indicates a deeper systemic issue.

Why False Positives Matter

Two Common Objections

Objection	Core Claim
1. "False positives don't affect reachability "	Traffic falls back to less-specifics or non-ROV paths → no real harm.
2. "False positives self-heal "	Transient sync delays auto-resolve; persistent errors trigger operator fixes.

Responses

- ◆ **Reachability ≠ Predictability**
 - Fallback paths are not designed for normal operation, they introduce: **unexpected latency, suboptimal routing, and route churn/instability.**
- ◆ **"Self-healing" Has a Hidden Cost**
 - **Transient delays:** The healing window is **not zero** – 10 to 100 minutes of instability is real operational damage, not "no harm."
 - **Persistent errors:** Fixes require manual intervention; **The damage—blackholes, customer complaints, degraded performance—happens first, before any repair begins.**
- ◆ **More importantly**
 - ◆ **No pain, no fix: If there's truly no pain, there's no fix. If there is pain, the damage already happened. Either way, the system loses.**
 - ◆ **The real danger: invisible hijacking – Tolerating false positives creates blind spots for real attacks.**
- **The design objective** of RPKI-based origin hijack protection should be: **Fast, accurate, and decisive – once a route is determined invalid, its reachability should be blocked immediately.**
 - ✓ Not "tolerated until proven harmful."
 - ✓ Not "self-healed after damage is done."
 - ✓ Not "masked by fallback paths that hide real attacks."
- ◆ **False positives undermine this objective.** Every tolerated invalid is a compromise: it either creates instability, delays true detection, or both.
- ◆ **Minimizing false positives is not optional – it's essential to making ROV work as designed: to detect and block hijacks, reliably, every time.**

Bottom Line

Challenge 2: The “Long Tail” of Networks Still Hesitate

◆ Deployment is uneven.

While precise statistics by network type are not formally published, the following pattern is widely observed in the community and supported by multiple data sources:

- **Tier-1 and large transit providers:** High ROV adoption. Most global transit-free networks now reject RPKI-invalids [IPinfo, 2026].
- **Content / eyeball networks:** High ROA coverage (driving traffic-weighted metrics), but lower incentive to reject others' invalids – their primary goal is protecting their own prefixes.
- **Regional ISPs / smaller operators:** Moderate adoption, often partial (e.g., filtering only at IXPs) [ATHENE / RoVista].
- **Stub ASes (the >50,000 long tail):** Very low deployment.

Data: RoVista (NANOG 90) found only 12.3% of ASes are fully protected--meaning even among 'deployed' ASes, most only filter partially."



◆ Why it matters:

- Even if the core drops invalids, **attackers can still reach victims through non-ROV paths.**
- **True protection against origin hijacking requires both ends of the connection:**
 1. **Source prevention**--clean at origin
 2. **Last-mile rejection**--drop invalids at the customer-facing edge.

Core deployment limits the blast radius – but it doesn't close the door.

The fence only works if the gate is closed at both ends – where the route starts, and where the traffic lands.

For true origin hijack protection, the long tail is not optional; it's the target.

Underlying Systemic Issues

These two challenges don't exist in isolation. They stem from three deeper systemic issues.

1. Synchronization Asymmetry

(RPKI Time-of-Flight, PAM 2023)

- RPKI management plane and router control plane are **decoupled**.
- **Delays:** average >10 minutes, maximum >100 minutes.
- **Consequences:**
 - **Transient false positives** due to lag.
 - Inconsistent VRRP snapshots across routers – **different decisions at the same time.**

2. ROA Misconfiguration

- ROA creation **involves human processes**, not just automated chains.
- Problems: stale ROAs, wrong ASN, incorrect maxLength etc.
- **Root cause:** Resource certificates bind IP prefixes to holders, but the **binding between prefix and authorized AS is a one-off signature** – no continuous verification of intent.

3. Incentive Gap for Small Networks

- For a stub AS, the **risk of a self-inflicted outage** outweighs the benefit of filtering others' routes.
- Without clear, low-risk signals, they stay in monitor mode – **This is precisely why the long tail remains long.**