RPKI-based BGP Origin Attestation: From Reactive Validation to Proactive Security

draft-huang-idr-bgp-origin-attestation-00

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

Mar. 2026
IETF 125@Shenzhen

Contents

- **Problem Statement**
  - Progress of RPKI ROV
  - Long-Standing Challenges
  - Underlying Systemic Issues
- **Solution Approach – Two Complementary Layers**
  - Layer 1: Pre-validation at the Origin
  - Layer 2: In-band Attestation (BGP-OA)
- **Summary & Call to Action**

RPKI ROV – Significant Progress

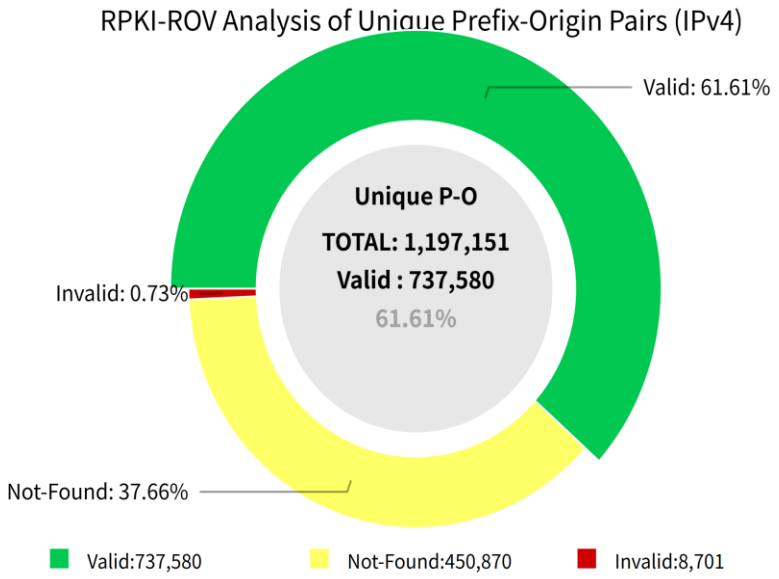

NIST RPKI Monitor: RPKI-ROV Analysis as of 2026-03-11

- **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

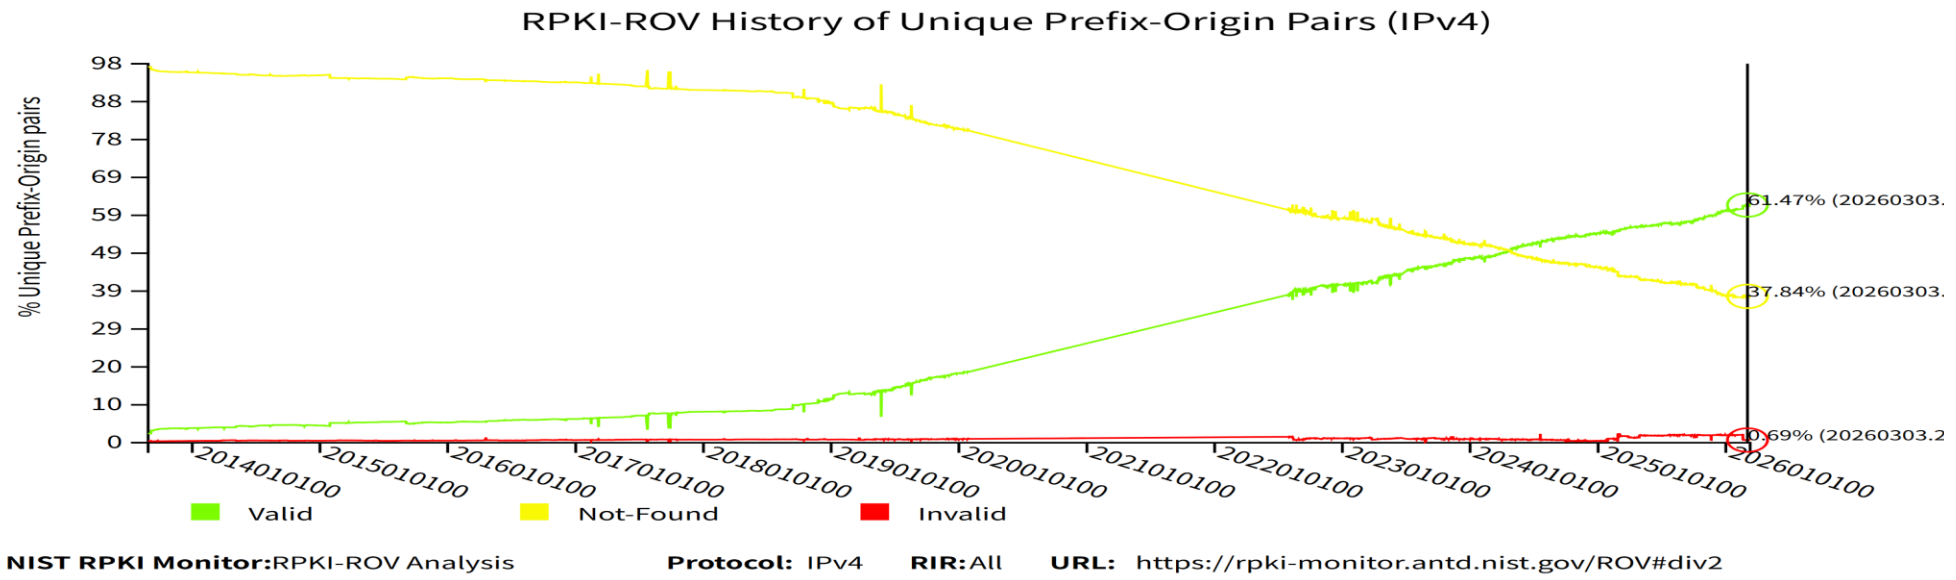

- NIST RPKI Monitor: RPKI-invalid prefixes (IPv4) have **remained around 1% for years**, often thousands of routes.
- Studies (NDSS 2026, "Demystifying RPKI-Invalid Prefixes"): **>96% of these are false positives** caused by misconfiguration, not hijacks.

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.**

Bottom Line

- **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.**

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 VRP 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.

Contents

- **Problem Statement**
  - Progress of RPKI ROV
  - Long-Standing Challenges
  - Underlying Systemic Issues
- **Solution Approach – Two Complementary Layers**
  - Layer 1: Pre-validation at the Origin
  - Layer 2: In-band Attestation (BGP-OA)
- **Summary & Call to Action**

Layer 1: Pre-validation at the Origin

- **What**: Mandatory, protocol-integrated check of self-originated routes against local RPKI cache **before advertisement**.

- **How it helps**:
  - **Eliminates self-inflicted invalids at the source** – prevents an AS's own configuration errors from ever entering the global table.
  - Studies (NDSS 2026) show that **>96% of today's RPKI-invalids are caused by misconfiguration** – meaning this approach targets the root cause.
  - **Immediate "selfish" benefit for the deploying AS**:
    - No accidental de-peering or self-inflicted outages
    - The AS becomes the guardian of its own reachability
  - **Ensures cross-plane (management-plane & control-plane) alignment:**
    - Verifies that what the router advertises matches what the ROA authorizes
    - Catches discrepancies before they become public routing problems

    Note: cross-router synchronization is addressed by the next layer – BGP-OA

Relationship to RFC 8893 and 9324

- **RFC 8893 is a standard – but it's a standard for *how* to do egress validation, not a mandate that you *must* do it.**
  - It requires that **if** you implement egress validation, you use the correct "effective origin AS". It does **not** require every router to perform egress validation, nor does it define what to do with routes that fail – drop? cache? log?.

- **RFC 9324 fills one gap**: it requires routers to cache ROV-dropped routes in Adj-RIB-In, so they can be re-evaluated without refresh storms. But it says nothing about Adj-RIB-Out – the routes an AS *intends to send* but discards at egress.

- **BGP-OA completes the picture**:
  - It makes pre-validation **mandatory** at the origin – not optional.
  - It requires caching of failed outbound routes (closing the Adj-RIB-Out gap).
  - And most importantly, when a validated route is finally announced, it carries a **cryptographic attestation** – turning a local check into a globally visible proof.

Layer 2: In-band Attestation (BGP-OA)

☐ Why Pre-validation Alone Isn't Enough

Pre-validation ensures the **origin's own management-plane and control-plane are aligned** – a critical first step. But the downstream still faces two fundamental challenges:

- **Challenge 1: Cross-Router Information Asymmetry Persists** – in Two Dimensions
  1) Horizontal inconsistency - across different paths
  2) vertical inconsistency - along the same forwarding path
  - **Result:** The same route can be treated inconsistently across the network, causing route flapping, unpredictable path selection, blackhole, invisible hijack, and operational uncertainty etc..
- **Challenge 2: Pre-validation Itself Is Partially Deployed.** Pre-validation cleans the routes from ASes that deploy it. But downstream **cannot assume** that every route it receives has undergone this cross-plane check.
  - **The result:** When a route arrives without an attestation, the downstream faces the same ambiguity as today – was it pre-validated? Was it not? It has no information to go on.

◆ What Downstream Needs – and What BGP-OA Provides

Downstream needs **visibility into the origin's validation state** to overcome both challenges:

- ✓ Cross-router asynchrony -- horizontal inconsistency and vertical inconsistency.
- ✓ Cross-plane uncertainty (did the origin actually validate?)

Pre-validation aligns the source. BGP-OA makes that alignment visible and verifiable to the rest of the Internet – bridging the gaps that a partially deployed, asynchronous system cannot avoid, and giving the long tail a reason to join.

Relationship with BGPSec

Some may ask: is this just BGPSec for origins?

	BGPSec	BGP-OA (this proposal)
What it protects	**Path integrity** – ensures the AS_PATH hasn't been modified	**Origin validity** – proves the source performed pre-validation
Scope	Every AS along the path must sign	Only the **origin AS** signs
Deployment model	All-or-nothing – if one AS doesn't sign, path security breaks	**Incremental** – works even if only the origin deploys; non-supporting routers just ignore the attribute
Operational overhead	High – every router must sign and verify every update	**Lower** – only origin signs; downstream verifies once
State in BGP	Replaces the AS_PATH with a signed object	Adds an **optional, transitive attribute** – existing ROV still works even if stripped

- **BGP-OA inherits BGPSec's trust model – but simplifies it**
  - Uses the same RPKI router certificates (RFC 8209) – no new PKI needed
  - But only the origin needs a certificate and signing capability
  - Downstream only needs to verify the origin's signature, not maintain a per-hop trust chain
- Enhancement to ROV, not replacement/new category of routing security. If the attribute is not included or even it is stripped, **ROV still works** – no downgrade vulnerability

Contents

- **Problem Statement**
  - Progress of RPKI ROV
  - Long-Standing Challenges
  - Underlying Systemic Issues
- **Solution Approach – Two Complementary Layers**
  - Layer 1: Pre-validation at the Origin
  - Layer 2: In-band Attestation (BGP-OA)
- **Summary & Call to Action**

Summary & Call to Action

- **Problem recap**:
  - ROV has come far, but false positives and uneven deployment persist.
  - Underlying issues: asynchrony, misconfiguration, incentive gap.
- **Our solution**:
  - Pre-validation + in-band attestation = proactive prevention + cooperative visibility.
  - **Not a replacement, but an enhancement** to ROV.

- **What we ask from the WG**:
  - Consensus on the problem statement.
  - Feedback on whether the two-layer approach addresses the root causes.
  - Guidance on next steps in IDR and SIDROPS.

Q&A / Thank You

Thanks!

Draft: draft-huang-idr-bgp-origin-attestation-00
Contact: huangmq@zgclab.edu.cn
Repository Link: *draft-huang-idr-bgp-origin-attestation-00*

BGP-OA Attribute Extension

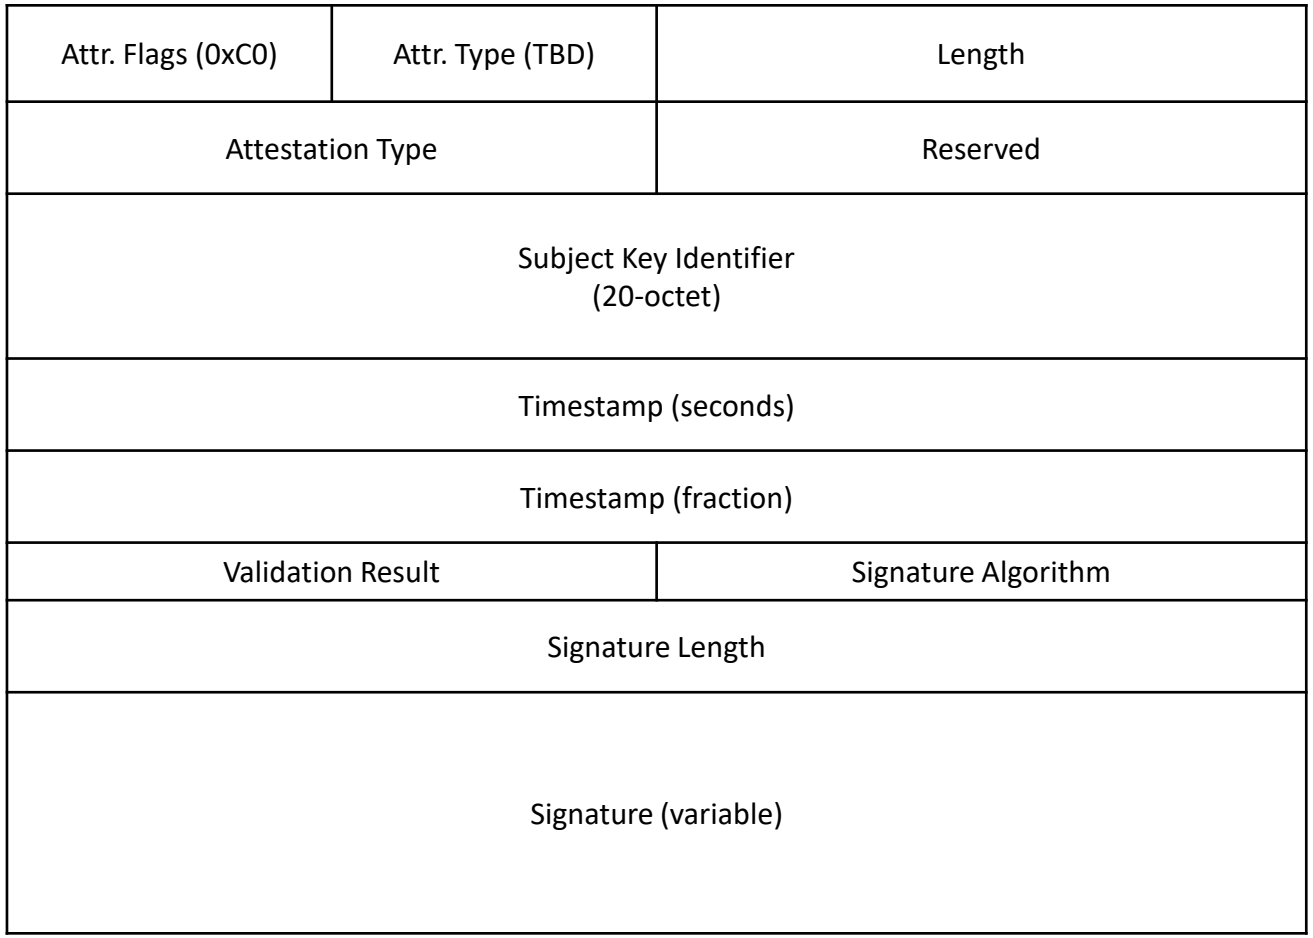

- **Optional, Transitive, Extended-Length BGP path attribute.**
- **Key Semantics & Trust:**
  - **Who:** (AS Number, SKI) identifies the router's key (from RPKI cert).
  - **What & Result:** The attested prefix and its Validation_Result (Valid / Not-Found).
  - **Proof:** Cryptographic signature over the prefix, AS, result, and timestamp.
  - **Attestation Type:** Carry enablement status (O-bit) and policy flag (V-bit).
- **Trust Anchor:** The signature is verified using the **RPKI Router Certificate** (RFC8209), linking the key to the AS and router.