

ASN VULNERABILITY METRICS — THEORY & SECURITY RATIONALE

Why These 6 Metrics Are Sufficient for Modeling Origin-Side Routing Environment Vulnerability & Attack Propagation Success

1. Introduction

ASN vulnerability, as used throughout this platform, is **not an intrinsic weakness of an autonomous system as an isolated entity**.

Instead, it represents the **abuse potential of the origin-side routing environment associated with that ASN**, and the ability of announcements originated from that environment to propagate across the global Internet under current routing conditions.

This model reflects a realistic attacker scenario in which a malicious actor:

- legitimately controls or operates an ASN, and
- originates misleading, incorrect, or policy-violating BGP announcements from within the routing environment surrounding that ASN.

The resulting risk arise from **insufficient constraint imposed by the surrounding routing ecosystem**.

The platform adopts a **dual-layer vulnerability model**:

(1) Origin-Side Environment Permissiveness

To what extent does the routing environment surrounding an ASN allow unsafe or weakly validated announcements to be originated and released?

(2) Propagation Success Potential (PSP)

Once originated, how widely and reliably would such announcements propagate across the Internet?

The second dimension — **Propagation Success Potential (PSP)** — elevates the system beyond simple origin-validation checks into a **global attack-surface and impact-prediction engine**.

The purpose of this document is to justify why the selected **ASN-level metrics** are sufficient to model both:

- **origin-side environment vulnerability**, and
- **real-world attack propagation success**.

Naturally, the actual impact of any routing attack also depends on the **specific prefix being used**.

For this reason, the platform combines **prefix-level structural vulnerability metrics** with **ASN-level origin-environment and propagation metrics**.

Only the **joint analysis of origin environment and prefix characteristics** can accurately predict real-world routing risk.

The six ASN-level metrics used are:

- `rpki_coverage_pct`
- `rov_enforcement_score`
- `irr_routeobject_coverage`
- `asn_filters_strict`
- `reachability_propagation`
- `as_path_len`

Each metric captures a **fundamental structural aspect of the routing environment surrounding an ASN**.

Together, they form a complete, multi-dimensional model of **origin-side routing environment vulnerability**.

2. What Is ASN Vulnerability? (Conceptually)

An ASN is considered **vulnerable as an origin environment** when:

- it operates within a routing ecosystem permissive enough to allow unsafe or weakly validated announcements to be originated,
- filtering and validation practices along origin-adjacent paths are weak or inconsistent,
- the surrounding topology allows originated announcements to propagate widely,
- upstreams and peers exhibit filtering gaps that can be leveraged.

In other words, **ASN vulnerability describes abuse potential at the origin environment level.**

Two Core Pillars of Origin-Side ASN Vulnerability

(A) Origin-Side Environment Permissiveness

Origin-side permissiveness reflects **how permissive the routing environment surrounding an ASN is.**

Each ASN operates within an **origin environment** composed of:

- upstream providers,
- peers,
- their respective filtering and validation practices,
- and the effective policy surface exposed to originated announcements.

An origin environment is considered permissive when, collectively:

- RPKI / ROA validation is weakly enforced along origin-adjacent paths,
- IRR-based prefix filtering is incomplete or inconsistently applied,
- abnormal or weakly validated AS paths are accepted by adjacent networks,

- import/export and adjacency policies allow unsafe announcements to exit the origin environment.

Under such conditions, **unsafe or malformed announcements can be originated and released into the global routing system**, not because the ASN is impersonated, but because **the surrounding routing environment fails to sufficiently constrain origin behavior**.

(B) Propagation Success Potential (PSP)

Even when unsafe announcements can be originated, **not all origin environments are equally capable of amplifying them**.

Propagation success depends on **structural properties of the ASN's position within the Internet topology**, including:

- upstream and peer diversity,
- position in the global routing graph,
- empirically observed propagation strength,
- average AS-path length,
- visibility across global BGP collectors.

These factors determine **how widely and how rapidly announcements originating from a given origin environment are likely to spread**.

Propagation success is therefore governed by:

- surrounding routing hygiene and validation behavior,
- IRR and filtering consistency along propagation paths,
- AS-path credibility,
- global reachability and visibility,
- observed propagation behavior.

This is exactly what the selected metrics measure — enabling the platform to predict **not only whether unsafe announcements can be originated**, but **whether they would meaningfully propagate at Internet scale**.

3. Rationale for Each Metric

3.1 RPKI Coverage Percentage — **rpki_coverage_pct**

What it measures

The percentage of prefixes originated within the ASN's routing environment that have valid ROAs.

Why it matters

RPKI coverage reflects how strongly the origin environment anchors routing behavior in cryptographic authorization.

Security relevance

Low RPKI coverage → **high origin-side environment permissiveness**.

Prefixes without ROAs appear as *Unknown*, which many networks still accept.

High RPKI coverage → **lower permissiveness**, indicating stronger routing hygiene.

3.2 ROV Enforcement Score — **rov_enforcement_score**

What it measures

Whether RPKI origin validation is actually enforced within the origin environment.

Why it matters

ROV enforcement is the most effective operational control for constraining unsafe origin behavior.

Security relevance

Consistent enforcement significantly reduces abuse potential at the origin environment level.

3.3 IRR Route Object Coverage — **irr_routeobject_coverage**

What it measures

Completeness and correctness of IRR route objects associated with prefixes originated from the environment.

Why it matters

IRR data remains widely used for prefix filtering by transit providers.

Security relevance

Weak or incorrect IRR objects allow unsafe announcements to escape the origin environment.

3.4 BGP Filters Strictness — **asn_filters_strict**

What it measures

Observed strictness of import/export filtering based on real routing behavior, including:

- prefix acceptance rules,
- customer-to-provider filtering patterns,
- adjacency permissiveness.

Security relevance

Strict filters reduce origin-side environment abuse potential.

Weak filters increase permissiveness.

3.5 Reachability Propagation — **reachability_propagation**

What it measures

How widely and fast announcements originating from the environment propagate across the global routing system.

Security relevance

Strong propagation characteristics enable large-scale amplification of unsafe announcements.

Propagation relevance

This metric directly models **Propagation Success Potential (PSP)**.

3.6 AS Path Length Average — **as_path_len**

What it measures

Average AS-path length of announcements originated from the environment.

Why it matters

AS-path length influences credibility and acceptance by upstream filters.

Propagation relevance

Shorter, consistent paths propagate more effectively.

4. Why These 6 Metrics Are Fully Sufficient

Origin-side ASN vulnerability depends on both:

- **environment permissiveness**, and
- **propagation success potential**.

These six metrics fully cover both dimensions:

Security Domain	Metrics
Origin Authorization	<code>rpki_coverage_pct</code> , <code>rov_enforcement_score</code>
Policy & Routing Hygiene	<code>irr_routeobject_coverage</code> , <code>asn_filters_strict</code>
Propagation & Reachability	<code>reachability_propagation</code> , <code>as_path_len</code>
Combined Attack Viability	all metrics jointly

There is no missing structural dimension.

5. Integration With ML

The ML engine learns nonlinear interactions between metrics, for example:

- Low IRR + weak ROV + high reachability → **extreme propagation risk**

- High RPKI + strict filters → **very low abuse potential**
- Weak filters with low reachability → **moderate risk with limited impact**

The system therefore predicts:

“If unsafe announcements originate from the routing environment associated with ASN X, how likely are they to propagate globally?”

This goes beyond origin validation — it models **real-world impact**.

6. Conclusion

These six ASN-level metrics:

- accurately capture **origin-side routing environment vulnerability**,
- model **real-world propagation success**,
- are grounded in observable BGP behavior,
- provide a complete foundation for ML-based routing-risk modeling.

Together, they form a **rigorous, technically correct, and enterprise-grade framework** for predicting **origin-environment abuse potential and global attack propagation impact**.