Inter-Domain Routing
Internet-Draft
Intended status: Informational
Expires: 17 August 2026

Y. Guo
Zhongguancun Laboratory
K. Xu
Tsinghua University
X. Wang
Capital Normal University

13 February 2026

BGP Communities for Security Policy Intent
draft-guo-idr-bgp-security-policy-community-01

Abstract

This document specifies a set of standardized BGP community to signal inter-AS routing security policy intent. The initial focus is on RPKI-based Route Origin Validation (ROV) using ROAs [RFC6482] [RFC6811] [RFC9582]. ROV produces validation outcomes such as "Valid", "Invalid", and "NotFound", but the operational treatment of "NotFound" and similar ambiguous cases is entirely a matter of local policy and often differs across networks.

This document defines transitive community that allows an Origin AS to explicitly express its security policy expectations regarding how its own originated routes SHOULD be treated when downstream Autonomous Systems (ASes) perform ROA-based origin validation. A typical example is an Origin AS indicating a preference for strict handling of ambiguous validation outcomes (e.g., NotFound) for its prefixes.

Unlike validation states, these community does not assert correctness, authorization, or RPKI deployment status, which is confront to [AVOID-RPKI-STATE-IN-BGP]. Instead, they communicate origin-declared policy intent to all downstream ASes, enabling them to correlate this intent with locally derived validation results. By enabling explicit signaling of security expectations without exporting validation state, this mechanism allows downstream ASes to make more informed policy decisions while reducing the risk of accidental outages caused by misalignment between origin expectations and downstream local policies.

The mechanism is orthogonal to existing routing security validation technologies and does not alter their semantics or deployment models.

Status of This Memo

This Internet-Draft is submitted in full conformance with the provisions of BCP 78 and BCP 79.

This Internet-Draft will expire on 17 August 2026.

Copyright Notice

Table of Contents

1. Introduction

Inter-domain routing security mechanisms intentionally separate validation from policy. While this separation improves robustness, it also creates persistent operational ambiguity.

Internet routing security relies on distributed validation mechanisms like RPKI ROA-based Route Origin Validation (ROV) [RFC6480] [RFC6482] [RFC6811] [RFC9582]. However, there is a functional gap between "knowing a route's validity state" and "knowing the origin's policy intent".

These security mechanisms are typically enforced locally only. No standardized method exists for an AS to signal its security policy expectations for its originated prefixes as they propagate through the inter-domain routing system. [AVOID-RPKI-STATE-IN-BGP] advises against carrying actual RPKI-derived validation state in BGP, in particular using transitive attributes such as BGP Communities. This is an important safeguard, but it leaves downstream ASes with only their local validation state and no explicit information about the origin' s security policy intent.

For example, when a Transit AS observes an RPKI "NotFound" state for a route, it cannot, based on RPKI state alone, distinguish between:

* an Origin AS that has intentionally not deployed ROAs (and may consider "NotFound" operationally acceptable for the time being);

* and an Origin AS that has deployed ROAs but is experiencing a configuration error or a hijack attempt (and may consider "NotFound" operationally undesirable or suspicious).

From the point of view of the validation algorithm, these cases all appear as "NotFound". Without additional information about the origin's expectations, downstream ASes must treat them according to their own local policies, which may or may not align with the origin's operational intent.

By allowing the Origin AS to signal a stricter security policy intent, downstream ASes can apply different local policies to otherwise identical validation outcomes based on the origin-declared preference. This signaling does not reveal the underlying cause of a "NotFound" state, and it does not carry validation results. Instead, it exposes the origin' s desired treatment of its routes under such ambiguity (for example, that the origin prefers its routes to be handled more strictly when validation is inconclusive).

Current mechanisms do not allow an Origin AS to express, in a standardized way:

* whether it prefers strict handling to be applied to ambiguous validation outcomes for its own routes;

* whether certain propagation behaviors are explicitly expected or operationally acceptable for its routes;

* and whether "NotFound" or similar outcomes are considered operationally acceptable for its routes.

As a result, downstream networks frequently face situations where routing information is technically acceptable according to their local validation policy, yet operationally unexpected from the perspective of the Origin AS. In large-scale deployments, this ambiguity may lead to:

* inconsistent treatment of identical prefixes;

* difficulty distinguishing misconfiguration from malicious behavior at an operational level.

By signaling security policy intent, an Origin AS can explicitly inform the network of its operational expectations regarding routing security for its own prefixes. For example, an Origin AS may indicate that it prefers downstream ASes to apply stricter handling for its prefixes when their local ROV results are ambiguous.

This signaling enables downstream ASes, if they choose to honor it, to better align their local policies with the origin's expectations, while still deriving and using their own validation results locally. The mechanism defined in this document is explicitly designed to

follow the guidance in [AVOID-RPKI-STATE-IN-BGP] by avoiding the carriage of RPKI-derived validation state in BGP and instead signaling only origin-declared policy intent.

1.1. Goals and Non-Goals

This document is scoped to signaling origin-declared policy intent for ROA-based origin validation only. It does not attempt to define new validation mechanisms or to standardize local routing policies.

The goal of this document is to provide a mechanism for an Origin AS to explicitly express a routing security policy intent to downstream ASes: namely, that routes originated by this AS SHOULD be subject to a stricter local policy when downstream ASes perform ROA-based origin validation.

The community defined in this document is intended to convey only the Origin AS’s intent concerning the desired treatment of its own routes. They do not:

* assert or reveal whether the Origin AS has actually deployed RPKI, ROAs, or ROV locally;

* export or encode any RPKI validation state (e.g., "Valid", "Invalid", "NotFound");

* and guarantee that downstream ASes will enforce or even interpret the signaled intent in a particular way.

Enforcement of any stricter policy remains entirely a local decision of each downstream AS. An Origin AS can request stricter handling via this community, but it cannot enforce that request on other ASes.

1.2. Conventions and Definitions

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.

Strict:

In this document, the term "Strict" as used in community names (namely, "ROA-Strict") refers solely to the Origin AS's stated policy preference regarding the handling of ambiguous or unfavorable validation outcomes for its own originated routes. It does not imply any mandated filtering, dropping, or preference change by downstream ASes, and it MUST NOT be interpreted as a remote instruction to suppress routes.

Security Policy Community:
A BGP Large Community defined by this document to convey origin-declared routing security policy intent for the origin's own prefixes.

Validation State:
A locally derived outcome of a security validation mechanism, such as RPKI Prefix Origin Validation ("Valid", "Invalid", "NotFound"). Validation state is explicitly out of scope for the community defined in this document and MUST NOT be encoded or inferred from them.

2. Architecture and Operations

2.1. Overview

The mechanism defined in this document operates entirely within the BGP control plane and does not alter protocol message formats or path selection procedures.

This Security Policy Community uses BGP Large Community [RFC8092]. Its format is "Global-Administrator:Action-ID:Parameter".

An Origin AS attaches one Security Policy Community when originating a route. This community is propagated unchanged unless explicitly removed or modified by policy.

Intermediate and receiving ASes on the propagation path may choose to:

* ignore the community;

* log or monitor their presence;

* correlate them with local validation results;

* and incorporate them into local policy decisions.

No mandatory processing behavior is defined, and no interoperability dependency is introduced.

2.2. Policy Signaling Versus Validation State

This document makes a strict distinction between:

* validation state, which is derived locally from cryptographic or registry-based mechanisms; and

* policy intent, which reflects the Origin AS's operational expectations for its own routes.

The community defined herein exclusively signal policy intent. They do not encode validation outcomes, confidence levels, or security posture of any AS.

Downstream ASes MUST NOT interpret these communities as an indication that validation has been successfully performed, nor as a substitute for local validation.

Implementations and operators MUST NOT configure policies that set, clear, or modify the Security Policy Community defined in this document based solely on per-route validation outcomes (for example, if validation is Valid, then attach ROA-Strict). Doing so would effectively re-export validation state in BGP Communities, contrary to the guidance in [AVOID-RPKI-STATE-IN-BGP]. This design explicitly aligns with the guidance in [AVOID-RPKI-STATE-IN-BGP], and avoids propagating dynamic security state within BGP.

2.3. Transitivity Considerations

All communities defined in this document are specified as transitive, with the intent that the origin-declared policy intent can be observed by all ASes along the AS-PATH.

Intermediate ASes may apply local policy that removes or modifies communities, such behavior is outside the scope of this specification. This specification does not impose any requirement on intermediate ASes to preserve these communities. Preservation is an operational choice intended to maximize the visibility of origin-declared policy intent.

2.4. Origin AS Behavior

An Origin AS that chooses to signal security policy intent SHOULD attach the appropriate Security Policy Community when originating a route. The communities are intended to reflect a relatively stable per-origin policy posture, not per-route or per-event state.

For example, an AS 65001 that wishes to indicate a strict policy posture with respect to ROA-based validation ambiguity for a given prefix may attach the following Large Community:

* 65001:1000:1 (ROA-Strict, default strict policy posture)

An Origin AS MUST NOT attach the Security Policy Communities defined in this document as a function of per-route validation outcomes. Instead, the communities are intended to describe the Origin AS’s policy intent independent of the current dynamic validation state of any one route.

This mechanism does not require the Origin AS to disclose whether it has deployed RPKI, ROAs, or ROV, and the presence or absence of these communities MUST NOT be interpreted as such disclosure.

2.5. Intermediate AS Behavior

A receiving AS that chooses to process these communities SHOULD, when using Large Communities, verify that the Global Administrator ASN in the Large Community matches the rightmost (Origin) AS in the AS_PATH. If they do not match, the community MUST be ignored for the purpose of interpreting origin policy intent, in order to limit unauthorized policy signaling.

If this plausibility check succeeds, a receiving AS MAY correlate the presence of a ROA-Strict community with its locally derived validation results as part of its local policy framework.

For example, a local policy may, if configured by the operator, treat a route carrying a ROA-Strict community as less acceptable when the local RPKI validation state is NotFound. Such behavior is illustrative only and is not mandated by this specification.

A receiving AS MUST NOT treat the presence of a ROA-Strict community as evidence that validation has already been performed, or that a particular validation outcome exists.

2.6. Parameter Field Extensibility

The "Parameter" field in the BGP Security Policy Community is explicitly designed for extensibility. Currently, a value of "1" conveys the default strict policy posture for the associated security action (e.g., ROA-Strict). Future assignments may introduce further parameters to support nuanced policy signaling, such as variant handling levels, time-limited policies, or security requirements specific to more recent routing security enhancements.

Example encodings (illustrative only, not assigned in this document):

* 65001:1000:1 (ROA-Strict, default strict policy posture)

* 65001:1000:2 (hypothetical future refinement of ROA-Strict behavior)

* 65001:1001:1 (hypothetical new Action-ID for a future policy, e.g., ASPA-Strict)

The Action-ID and Parameter range is managed through IANA for orderly growth as the community adopts richer security policies.

The extensibility of the Parameter field is intentionally limited to policy refinement and does not introduce conditional logic or dynamic state signaling.

3. Applicability and Deployment

3.1. Applicability to Route Leak and Hijack Detection

Security policy communities may serve as additional context for routing analysis systems.

For example, a route that violates an origin-authorized export constraint, while also exhibiting abnormal AS path patterns, may be flagged as anomalous with higher confidence when it also carries a strict policy community from the expected origin.

Such signals can reduce false positives in detection systems by providing operator-declared intent, without asserting correctness. This document does not define detection algorithms or mitigation procedures.

3.2. Deployment Considerations

The proposed mechanism is compatible with existing routing policies and does not require changes to BGP implementations. Deployment of this mechanism is expected to be incremental and partial.

The proposed mechanism is intended for gradual deployment and interoperability with existing BGP technologies. Origin ASes may selectively signal policy intent for specific prefixes. In hybrid networks where both supporting and non-supporting ASes are present, policy communities will simply be ignored by legacy BGP speakers, providing backward compatibility. No coordination between ASes is required.

Operators are encouraged to monitor for loss or modification of policy communities due to intermediate ASes that filter or rewrite BGP Community attributes, so as to ensure policy expectations are properly signaled end-to-end.

4. Security Considerations

This document defines a policy signaling mechanism using BGP communities. It does not define a security mechanism and does not provide independent security guarantees. It is not intended for real-time attack mitigation or automated incident response.

This document follows the guidance in [AVOID-RPKI-STATE-IN-BGP] by not carrying any RPKI-derived validation state in BGP. The communities defined here do not encode or imply specific validation outcomes (such as "Valid", "Invalid", or "NotFound"). Instead, they allow an Origin AS to express a relatively stable policy posture regarding how downstream ASes may treat ambiguous validation results for its routes, if they choose to do so.

4.1. Authenticity and Integrity

The communities defined in this document are not cryptographically protected and may be modified, removed, or added in transit. This is consistent with existing BGP community usage and with the design goals of this document.

No attempt is made to ensure integrity or authenticity of community propagation. Accordingly, these communities MUST NOT be treated as authoritative security assertions and MUST NOT be used as a basis for accepting otherwise invalid routes.

The semantics defined herein apply only to communities that are plausibly originated by the Origin AS, as determined, in the case of Large Communities, by the Global Administrator field matching the rightmost AS in the AS_PATH. This check is intended solely to limit semantic impersonation and does not constitute a security guarantee.

An adversary capable of hijacking a route may also attach, modify, or remove communities. It is worthing note that the mechanism is fail-closed with respect to adversarial injection of Security Policy Communities. An intermediate AS, or an active attacker on the path, could unilaterally attach a strict Security Policy Community to a route that did not originate it. However, this cannot weaken security posture: at worst, it causes the route to be treated under stricter assumptions (e.g., making certain ambiguous states appear more suspicious to detection systems) than the true origin may have intended. Because the communities do not grant additional

reachability or override existing validation results, but only bias analysis and filtering towards more conservative handling, unauthorized addition of such communities can at most increase false positives, not reduce protections.

4.2. Relationship to Validation Mechanisms

The communities defined in this document do not represent validation results, security states, or assertions of route correctness, legitimacy, or authorization.

In particular, the presence or absence of a Security Policy Community MUST NOT be interpreted as indicating whether a route is valid or invalid under RPKI, BGPsec, or any other validation mechanism.

These communities MUST NOT override locally derived validation results, including a "Valid" RPKI state. They may be correlated with validation outcomes as part of local policy or analysis, but they do not alter the semantics of those outcomes.

Implementations and operators MUST NOT use these communities as a reason to skip or short-circuit local validation.

4.3. Policy Semantics and Downstream Behavior

The communities defined in this document express origin-authorized policy intent only. They do not define required actions.

Downstream Autonomous Systems MAY ignore these communities entirely without violating this specification. Any routing, filtering, or preference decisions remain solely a matter of local policy.

The absence of a policy community MUST NOT be interpreted as expressing the opposite intent.

This document does not require or expect consistent interpretation or uniform behavior across Autonomous Systems. Differences in interpretation, deployment, and operational use are expected and acceptable.

4.4. Denial-of-Service and Abuse Considerations

This document intentionally avoids defining communities that directly request route suppression or traffic dropping. As a result, it reduces the risk that a malicious actor could trigger network-wide disruption through abuse of policy signaling.

Nevertheless, operators should be aware that misconfiguration or abuse of these communities may influence local policy decisions if such decisions are explicitly configured to consider them. Operators are encouraged to avoid automated hard actions based solely on the presence of these communities, and to combine them with independently derived validation results and operational context.

4.5. Threat Model

Potential abuse scenarios include, but are not limited to:

* false or misleading signaling of policy intent;

* removal or modification of policy signals during propagation;

* inconsistent signaling across multiple origin points.

These risks are inherent to existing uses of BGP communities and do not introduce new attack vectors. Operators SHOULD correlate these signals with independently verifiable information when making security-related decisions.

5. IANA Considerations

his document requests the creation of one new sub-registry under the "BGP Large Communities" registry. No other IANA actions are required.

5.1. BGP Security Policy Action IDs

This document defines new BGP Large Community values for signaling security policy intent. Large Communities are used, rather than Extended Communities, to avoid ASN exhaustion and ambiguity associated with 16-bit Global Administrator fields.

IANA is requested to create a sub-registry titled "BGP Security Policy Action IDs" under the "BGP Large Communities" registry.

Large Communities used for this purpose have the following format:

```
Global-Administrator:Action-ID:Parameter
```

The Global-Administrator field MUST be set to the ASN of the Origin AS. The Action-ID field is an integer whose semantics are defined in the "BGP Security Policy Action IDs" sub-registry created by this document. The Parameter field is left to future documents or operator-specific conventions.

The Action-ID space is globally coordinated by IANA so that the same Action-ID has the same semantics regardless of the Origin AS using it. The "Strict" qualifier expresses only an origin-declared preference and does not define any required downstream behavior.

The initial contents of the "BGP Security Policy Action IDs" sub-registry are:

Action ID	Name	Policy Intent Description
1000	ROA-Strict	Origin expresses a preference for strict handling of its originated routes when downstream RPKI validation results are ambiguous or unfavorable (e.g., NotFound or Invalid).

6. Relationship to Existing and Future Mechanisms

This section clarifies the relationship between the mechanism defined in this document and existing routing policy, traffic engineering, and routing security mechanisms. The goal is to explicitly delineate scope and avoid overlap or semantic ambiguity.

6.1. Relationship to RPKI and ROA

RPKI-based mechanisms such as ROA provide cryptographic or registry-backed validation outcomes for routing information. These mechanisms answer the question of whether a route is consistent with registered authorization data.

The communities defined in this document do not provide validation and do not alter validation outcomes. They do not indicate that a route is valid, invalid, or authorized. Instead, they allow an Origin AS to express its operational expectations regarding how ambiguous or unfavorable validation outcomes (e.g., NotFound) _for its routes_ may be handled by downstream ASes.

This mechanism is therefore complementary to RPKI and related validation mechanisms. It operates strictly at the policy signaling layer and does not export validation state, consistent with the guidance in [AVOID-RPKI-STATE-IN-BGP].

6.2. Relationship to "Avoid RPKI State in BGP"

[AVOID-RPKI-STATE-IN-BGP] provides guidance against carrying RPKI-derived validation state in BGP, particularly via transitive attributes such as BGP communities, because doing so can leak local validation outcomes and policy decisions into the global routing system.

The mechanism described in this document has been explicitly designed to conform to that guidance:

* It does not encode, infer, or transport validation states (e.g., "Valid", "Invalid", "NotFound") in BGP attributes;

* it does not attempt to synchronize or standardize local validation policies across ASes; and

* it uses communities only to carry origin-declared policy intent for the treatment of the origin’s own routes.

As a result, downstream ASes remain responsible for performing their own validation and applying their own policies. The communities defined here are optional hints that can help align local policy decisions with the origin’s expressed expectations, without exporting dynamic validation state in BGP.

The work in [AVOID-RPKI-STATE-IN-BGP] argues that RPKI validation state (e.g., "Valid", "Invalid", "NotFound") and related information SHOULD NOT be propagated in BGP, in order to avoid leaking local validation outcomes and policy decisions into the global routing system.

This document is aligned with that principle:

* The communities defined here do not carry or reflect any RPKI validation state; and

* they do not reveal whether RPKI/ROV is deployed by the Origin AS or by any downstream AS.

Instead, this document only defines a way for the Origin AS to signal a policy intent to downstream ASes: "please apply stricter ROA-based validation policy to my originated routes, if you support such a policy." Whether and how a downstream AS uses this hint in its local policy is entirely at its own discretion.

6.3. Relationship to BGPsec

BGPsec [RFC8205] provides cryptographic protection of the AS_PATH to ensure path integrity and origin authentication. BGPsec is designed to assert and verify routing correctness.

The mechanism defined in this document does not provide cryptographic protection, path validation, or origin authentication. It does not attempt to replace or replicate BGPsec functionality. Instead, it provides an optional policy signal that may be used in conjunction with BGPsec or in environments where BGPsec is not deployed.

6.4. Relationship to BGP Color and Color-Aware Routing

BGP Color and Color-Aware Routing mechanisms [RFC9871] are primarily intended to support traffic engineering and transport-specific constraints, such as latency, bandwidth, or SR policy selection.

While both mechanisms use BGP communities as signaling vehicles, the semantics are fundamentally different. BGP Color expresses forwarding or transport preferences, whereas the communities defined in this document express origin-declared routing security policy intent.

This document does not define path selection behavior, traffic steering, or forwarding constraints, and does not overlap with the objectives of Color-Aware Routing.

6.5. Relationship to Only-To-Customer (OTC)

OTC [RFC9234] is a mechanism designed to assist in route leak prevention by signaling export intent at AS boundaries.

The communities defined in this document differ from OTC in scope and semantics. OTC signals propagation constraints related to business relationships, whereas this document signals security policy expectations related to validation ambiguity for an origin’s routes.

These mechanisms are complementary and may coexist on the same routes. Neither mechanism subsumes the other.

6.6. Relationship to Potential ASPA-Based Extensions

This document focuses exclusively on ROA-based origin validation and the associated ROA-Strict intent community. Other validation mechanisms such as ASPA [ASPA-Profile] [ASPA-Verification] may, in the future, benefit from similar origin-intent signaling constructs.

Any such ASPA-related communities, if defined, MUST follow the same core principles as this document:

* they MUST NOT export ASPA-derived validation state in BGP;

* they MUST signal only relatively stable origin policy intent;

* and they MUST leave all enforcement decisions to downstream local policy.

The specification of ASPA-based policy communities, including any additional Action IDs, is out of scope for this document and is expected to be covered in separate documents if there is interest.

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/rfc/rfc2119>.

[RFC4271] Rekhter, Y., Ed., Li, T., Ed., and S. Hares, Ed., "A Border Gateway Protocol 4 (BGP-4)", RFC 4271, DOI 10.17487/RFC4271, January 2006, <https://www.rfc-editor.org/rfc/rfc4271>.

[RFC8092] Heitz, J., Ed., Snijders, J., Ed., Patel, K., Bagdonas, I., and N. Hilliard, "BGP Large Communities Attribute", RFC 8092, DOI 10.17487/RFC8092, February 2017, <https://www.rfc-editor.org/rfc/rfc8092>.

[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/rfc/rfc8174>.

7.2. Informative References

[ASPA-Profile] Azimov, A., Uskov, E., Bush, R., Snijders, J., Housley, R., and B. Maddison, "A Profile for Autonomous System Provider Authorization", Work in Progress, Internet-Draft, draft-ietf-sidrops-aspa-profile-22, 6 February 2026, <https://datatracker.ietf.org/doc/html/draft-ietf-sidrops-aspa-profile-22>.

[ASPA-Verification]
Azimov, A., Bogomazov, E., Bush, R., Patel, K., Snijders, J., and K. Sriram, "BGP AS_PATH Verification Based on Autonomous System Provider Authorization (ASPA) Objects", Work in Progress, Internet-Draft, draft-ietf-sidrops-aspa-verification-24, 19 October 2025, <https://datatracker.ietf.org/doc/html/draft-ietf-sidrops-aspa-verification-24>.

[AVOID-RPKI-STATE-IN-BGP]
Snijders, J., Fiebig, T., and M. Stucchi, "Guidance to Avoid Carrying RPKI Validation States in Transitive BGP Path Attributes", Work in Progress, Internet-Draft, draft-ietf-sidrops-avoid-rpki-state-in-bgp-04, 12 February 2026, <https://datatracker.ietf.org/doc/html/draft-ietf-sidrops-avoid-rpki-state-in-bgp-04>.

[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/rfc/rfc6480>.

[RFC6482] Lepinski, M., Kent, S., and D. Kong, "A Profile for Route Origin Authorizations (ROAs)", RFC 6482, DOI 10.17487/RFC6482, February 2012, <https://www.rfc-editor.org/rfc/rfc6482>.

[RFC6811] Mohapatra, P., Scudder, J., Ward, D., Bush, R., and R. Austein, "BGP Prefix Origin Validation", RFC 6811, DOI 10.17487/RFC6811, January 2013, <https://www.rfc-editor.org/rfc/rfc6811>.

[RFC7999] King, T., Dietzel, C., Snijders, J., Doering, G., and G. Hankins, "BLACKHOLE Community", RFC 7999, DOI 10.17487/RFC7999, October 2016, <https://www.rfc-editor.org/rfc/rfc7999>.

[RFC8097] Mohapatra, P., Patel, K., Scudder, J., Ward, D., and R. Bush, "BGP Prefix Origin Validation State Extended Community", RFC 8097, DOI 10.17487/RFC8097, March 2017, <https://www.rfc-editor.org/rfc/rfc8097>.

[RFC8205] Lepinski, M., Ed. and K. Sriram, Ed., "BGPsec Protocol Specification", RFC 8205, DOI 10.17487/RFC8205, September 2017, <https://www.rfc-editor.org/rfc/rfc8205>.

[RFC9234] Azimov, A., Bogomazov, E., Bush, R., Patel, K., and K. Sriram, "Route Leak Prevention and Detection Using Roles in UPDATE and OPEN Messages", RFC 9234, DOI 10.17487/RFC9234, May 2022, <https://www.rfc-editor.org/rfc/rfc9234>.

[RFC9582] Snijders, J., Maddison, B., Lepinski, M., Kong, D., and S. Kent, "A Profile for Route Origin Authorizations (ROAs)", RFC 9582, DOI 10.17487/RFC9582, May 2024, <https://www.rfc-editor.org/rfc/rfc9582>.

[RFC9871] Rao, D., Ed. and S. Agrawal, Ed., "BGP Color-Aware Routing (CAR)", RFC 9871, DOI 10.17487/RFC9871, November 2025, <https://www.rfc-editor.org/rfc/rfc9871>.

Acknowledgments

TBD.

Authors' Addresses

Yangfei Guo
Zhongguancun Laboratory
Email: guoyangfei@zgclab.edu.cn

Ke Xu
Tsinghua University
Email: xuke@tsinghua.edu.cn

Xiaoliang Wang
Capital Normal University
Email: wangxiaoliang0623@foxmail.com