Network Working Group
Internet-Draft
Updates: draft-li-trustworthy-routing-discovery-00 (if approved)
Intended status: Standards Track
Expires: 28 August 2025

X. Li
X. Wei
Beijing University of Posts and Telecommunications

24 February 2025

Trustworthy Routing Discovery
draft-li-trustworthy-routing-discovery-00

Abstract

End users with high security and privacy requirements expect their data to be transmitted only through trusted devices to mitigate the risk of data leakage. Therefore, the development of trustworthy routing mechanisms is anticipated to be a key trend in the future evolution of the internet. This specification describes a network architecture that supports trustworthy routing and a scheme for carrying trustworthy routing information (TRI) using the BGP and BGPsec protocols.

Status of This Memo

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

Internet-Drafts are working documents of the Internet Engineering Task Force (IETF). Note that other groups may also distribute working documents as Internet-Drafts. The list of current Internet-Drafts is at https://datatracker.ietf.org/drafts/current/.

Internet-Drafts are draft documents valid for a maximum of six months and may be updated, replaced, or obsoleted by other documents at any time. It is inappropriate to use Internet-Drafts as reference material or to cite them other than as "work in progress."

This Internet-Draft will expire on 28 August 2025.

Copyright Notice

Table of Contents

1. Introduction

In the existing network system, data transmission security is almost entirely achieved through end-to-end encryption. However, users with high security and privacy requirements are no longer satisfied with this mechanism. They now demand that traffic be forwarded only by network devices that meet their expected security properties. For example, some clients may require their data to traverse through trusted devices and links only, to avoid data being exposed to insecure devices, causing leakage. By utilizing remote attestation technology [RFC9334], network Autonomous Systems (AS) can provide evidence of the level of security they support. This evidence can be exchanged between ASes through existing network protocols, such as BGP [RFC4271] or BGPsec [RFC8205], and parsed by ASes that support trustworthy routing. Through distributed management, ASes can perceive the trustworthiness of other ASes and plan appropriate transmission paths according to the security requirements of endpoints' traffic. Note that despite the extensive body of work on trust assessment for devices, there is still no standardized framework for trustworthiness evaluation among researchers. Therefore, in this draft, a more flexible solution-the Trust Assessment Profile (TAP) is proposed, which contains a set of evaluation rules designed to assess whether a device meets specific

trust requirements. Any network domain can maintain a set of TAPs, provided they are recognized by other domains. Within a domain, the orchestrator evaluates network devices based on specific TAPs and classifies them as either "trusted" or "untrusted" according to the rules of the current TAP. Additionally, we recommend advancing towards a multi-level trust classification in future research to enable more granular control over trust management. This document presents a network architecture that supports trustworthy routing and a scheme for carrying trustworthy routing information (TRI) using the BGP and BGPsec protocols.

2. Terminology

2.1. Terms

The following terms are imported from [RFC9334]: Attester, Evidence, Verifier.

Newly defined terms for this document:

Trust Assessment Profile --- a set of rules that specify how the verifier evaluates the trustworthiness of the attester based on the evidence provided during the remote attestation process.

Trustworthy Routing Information --- a series of messages that are exchanged between Autonomous Systems (AS) to propagate information regarding the trustworthiness of forwarding devices within the domain.

Trust Assessment Result --- whether a network device meets a specific trust assessment profile, with the result being either "trusted" or "untrusted".

2.2. Requirements Notation

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.

3. Architecture

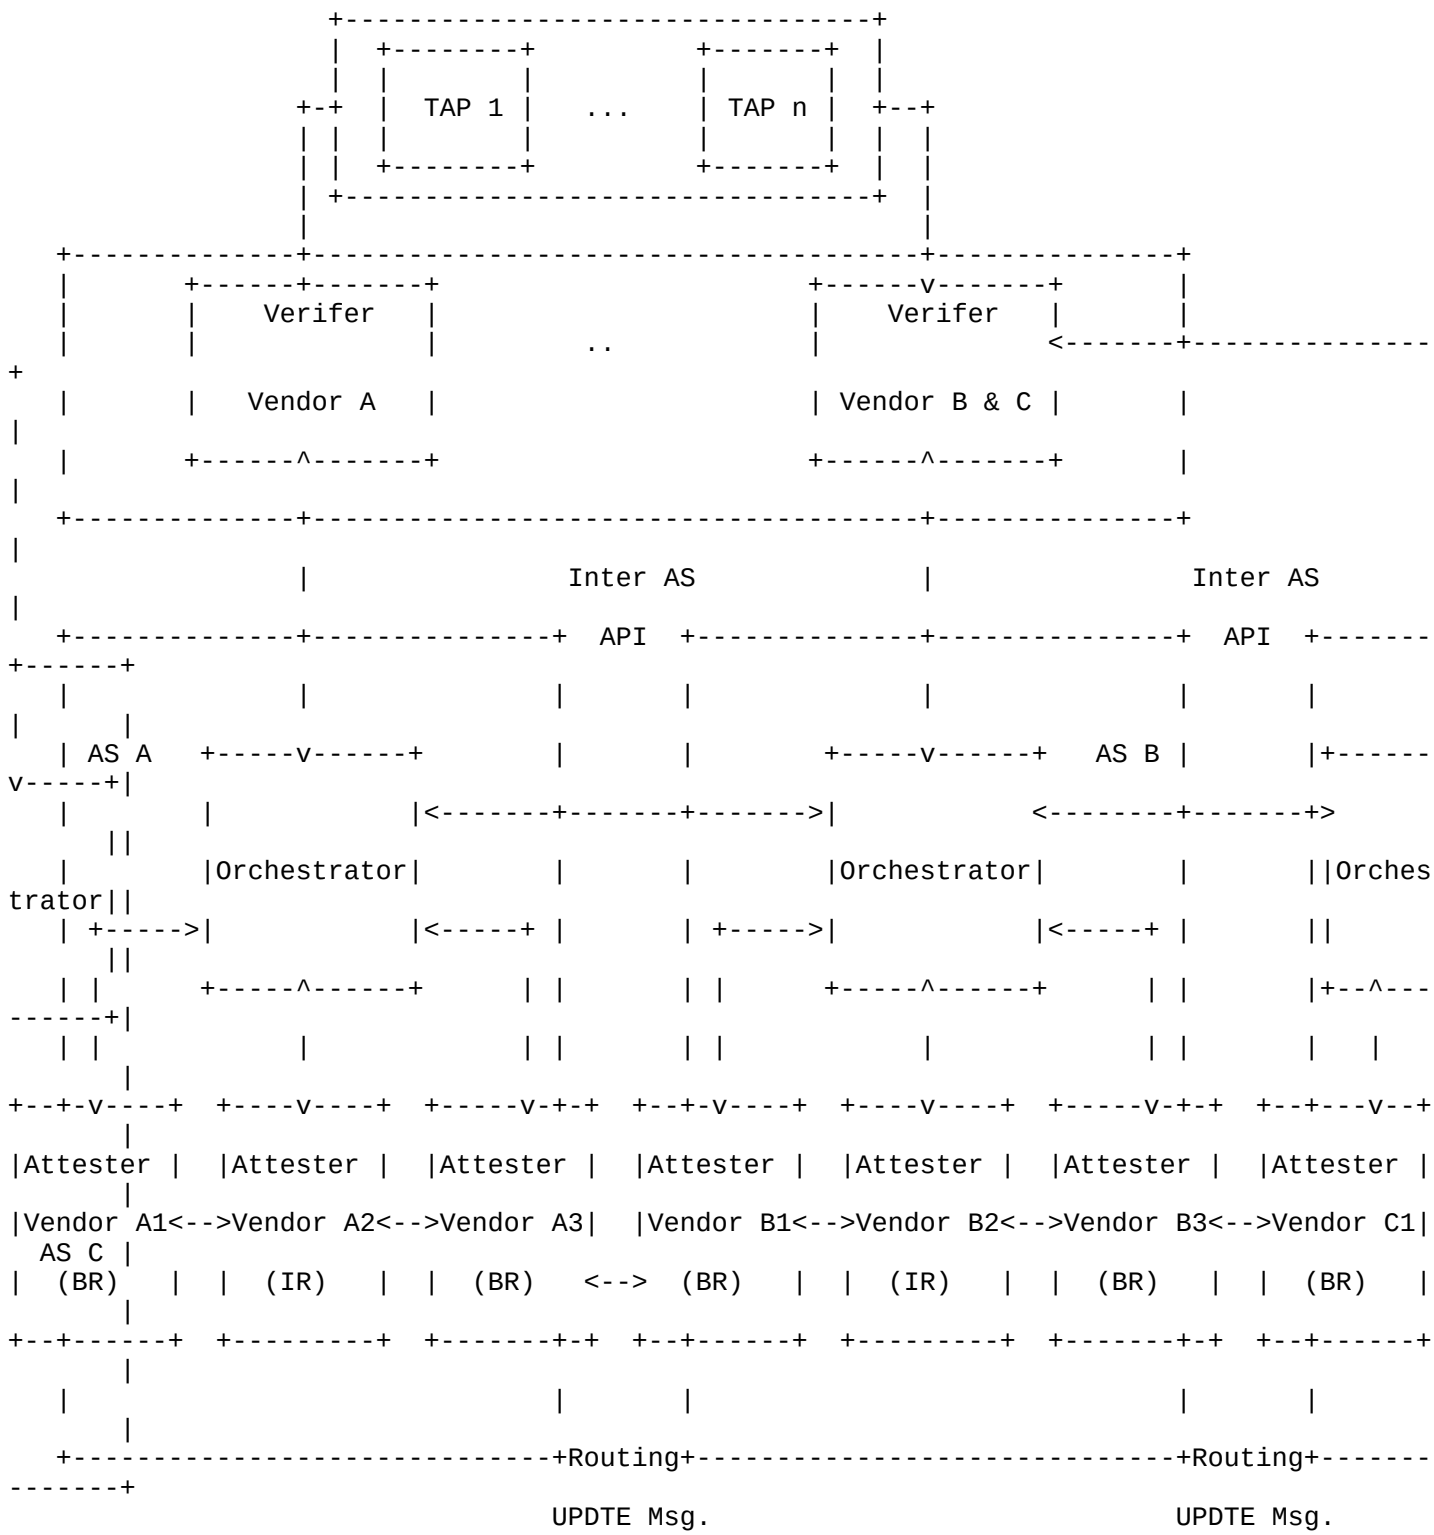

Figure 1: Example architecture of trustworthy-routing supported network

Figure 1 illustrates an example architecture of a trustworthy-routing supported network in a scenario involving multiple ASes. Please note that the proposed architecture is based on the "multi client-multi operator" framework from [nasr-arch]. With this architecture, each orchestrator is operated by the AS’s operator and is responsible for managing devices (Attesters) within its domain, providing TRI to

other ASes, and communicating with other ASes' orchestrators via the Inter-AS API to negotiate trustworthy routing policy flexibly. Each attester holds a remote attestation report issued by the verifier of its respective domain, indicating its trustworthiness assessment results (TAR). The TAR of an attester is determined by a set of remote attestation rules, which are described in a trust assessment

profile (TAP) and identified by a globally unique identifier, i.e., the trust assessment profile ID. All verifiers maintain consensus on the TAP. Note that within this framework, the TAR of a device is relative, determined by whether it is considered trusted under the specific rules of a given TAP.

TRI is propagated among ASes through BGP UPDATE or BGPsec UPDATE messages. How the BGP and BGPsec protocol stacks are extended to carry TRI will be explained later. With the assistance of BGP or BGPsec, each AS can perceive the TAR of other ASes and select appropriate paths for transmission according to the service requirements of the client. It is important to note that the design of this architecture also allows non-adjacent ASes to establish an L3 logical link through orchestrator negotiation, thereby shielding the trustworthiness of intermediate domains from affecting the entire transmission chain. For example, in Fig 1, supposing AS A and AS C are considered trusted under a specific TAP T, while AS B is not. Without additional measures, the transmission path from A to C cannot support the TAP T due to the trust bottleneck at B. In this case, orchestrators in AS A and AS C can negotiate and establish a logical link between the edge routers (A3 - C1) of the two domains using L3 security technologies (such as IPsec), thereby satisfying the requirement of TAP T in the data transmission from A to C.

4. Protocol Design

This draft presents the implementation of inter-domain propagation of TRI using the BGP or BGPsec protocols. Specifically, these protocols are extended to carry TRI segments. Each segment contains the TRI of an AS that supports trustworthy routing along the path.

The format of the TRI segment is shown in Figure 2. The first three fields represent the AS identifier that publishes the TRI, the identifier of the verifier that conducts the remote attestation for the AS, and the remote attestation report identifier, respectively. The attestation report ID can be provided as a URL, allowing the administrators of the BGP router’s domain to directly access the report and verify its authenticity. The trust assessment profile ID indicates the TAP identifier adopted by the AS. The TAR denotes whether this AS can be considered trusted under the given TAP. The timestamp denotes the publication time of the remote attestation result, serving to ensure the freshness of this information so that it consistently reflects the attester’s latest state. Additionally, each TRI segment contains a signature field, where the AS sign the entire segment using its private key to ensure the authenticity of the TRI issuer’s identity and prevents tampering of the TRI field during transmission. The selection of algorithms and key formats in the signing process can be guided by [RFC8205].

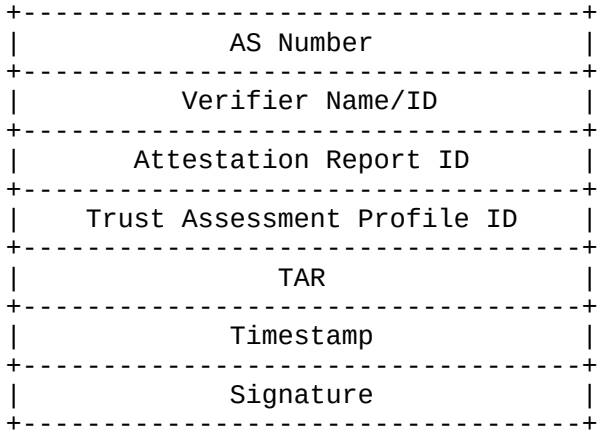

Figure 2: TRI segment format

This draft follows the passport model outlined in [RFC9334] to construct the data flow diagram between network entities. As shown in Figure 3, when BGP router 1 seeks to prove the TAR of its AS to a router in another AS, such as BGP router 2, it first conveys its state information, known as "evidence", to a verifier that compares this evidence against a specific TAP. The verifier then returns the attestation result to BGP router 1. BGP router 1 does not consume the attestation result, but caches it and further encapsulates it into the TRI following the format outlined in Fig. 2. BGP router 1 then presents this TRI to BGP router 2 via BGP or BGPsec protocol, and BGP router 2 checks this information to determine whether to accept it. Specifically, BGP router 2 first compares the TAP ID contained in this message to determine whether the TAP is supported within BGP router 2’s AS. It then verifies the signature field to ensure the authenticity of the TRI and subsequently updates the TAR associated with the AS of BGP router 1.

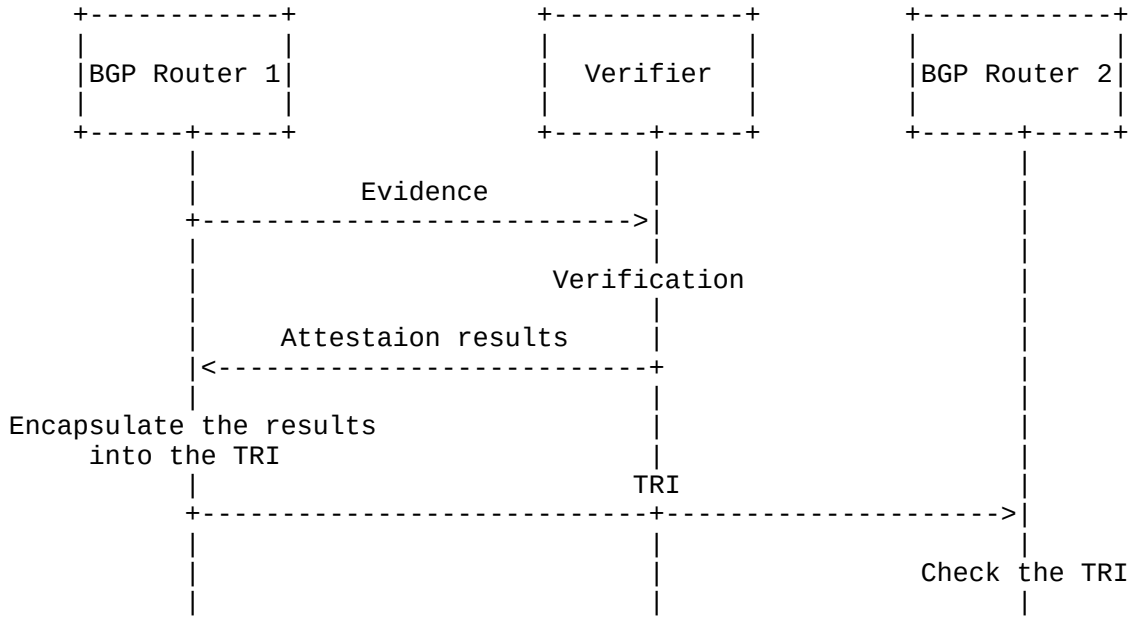

Figure 3: The data flow diagram between network entities

4.1. BGP

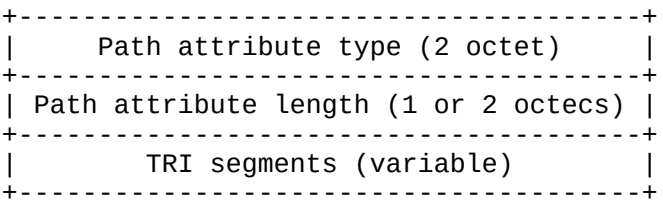

Figure 4: TRI path attribute format in BGP

When using the BGP protocol, an extended field called the trustworthy routing path attribute is introduced to carry TRI. This extended field follows the TLV format and [RFC4271], as shown in Figure 4. The path attribute type is a 2 octets field that consists of the attribute flags and attribute type code, as shown in Figure 5.

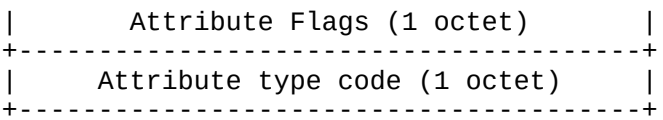

Figure 5: Path attribute type format

Since the trustworthy routing path attribute is an optional transitive BGP path attribute, the first two bits of its attribute flags are set to 11. The remaining 6 bits of the path flags are configured according to [RFC4271]. The attribute type code indicates that this path attribute carries TRI. The path attribute length is a 1-octet or 2-octet field that indicates the total length of the value field. The value field of the trustworthy routing path attribute contains a series of consecutively arranged TRI segments that represent the TRI along the path.

4.2. BGPsec

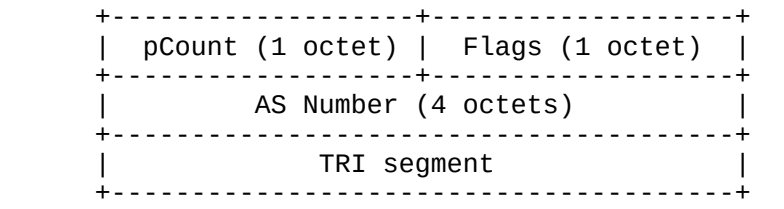

Figure 6: The secure_path segment format of trustworthy-routing supported BGPsec

When using the BGPsec protocol, the TRI segment for each AS is separately placed within the corresponding secure_path segment, as shown in Figure 6. One of the seven unassigned bits in the Flags field is used to indicate whether the AS supports trustworthy routing, referred to as the E_T_R Flag in Figure 7. If E_T_R Flag = 0, the AS does not support trustworthy routing, and the routing update message will be parsed according to the original BGPsec protocol. If E_T_R Flag = 1, the AS supports trustworthy routing, and its TRI will be included in the corresponding secure_path segment.

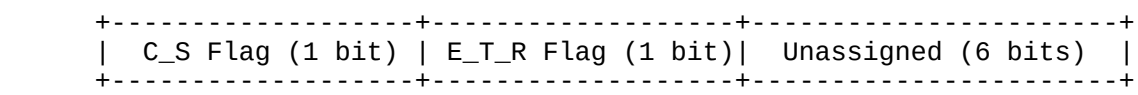

Figure 7: Flags format

It is important to note that the AS number is already included in the secure_path segment of BGPsec, so this information is omitted from the TRI. Additionally, since BGPsec already applies a signature mechanism to ensure the authenticity of the secure_path, the signature field in the TRI segment is also omitted.

5. Security Considerations

This document has no further security considerations.

6. References

6.1. Normative References

[RFC9334] Birkholz, H., Thaler, D., Richardson, M., Smith, N., and W. Pan, "Remote ATtestation procedureS (RATS) Architecture", RFC 9334, DOI 10.17487/RFC9334, January 2023, <https://www.rfc-editor.org/info/rfc9334>.

[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/info/rfc2119>.

[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/info/rfc8174>.

6.2. Informative References

[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/info/rfc4271>.

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

[nasr-arch] Liu, C., Chen, M., Richardson, M., and D. Lopez, "Network Attestation for Secure Routing (NASR) Architecture", October 2024, <https://datatracker.ietf.org/doc/html/draft-liu-nasr-architecture>.

Authors' Addresses

Xiaoyong Li
Beijing University of Posts and Telecommunications
No.10 Xitucheng Road
Beijing
100876
China
Email: lixiaoyong@bupt.edu.cn

Xinghai Wei
Beijing University of Posts and Telecommunications
No.10 Xitucheng Road
Beijing
100876
China
Email: junjuntvt@bupt.edu.cn