

Network Working Group
Internet-Draft
Intended status: Standards Track
Expires: 28 May 2026

S. Zhuang
H. Wang
Huawei Technologies
24 November 2025

Enhanced BGP Resilience
draft-zwg-rtgwg-enhanced-bgp-resilience-00

Abstract

According to the base BGP specification, a BGP speaker that receives an UPDATE message containing a malformed attribute is required to reset the session over which the offending attribute was received. RFC7606 revises the error handling procedures for a number of existing attributes. The use of the "treat-as-withdraw" and "attribute discard" approaches significantly reduces the likelihood of BGP sessions being reset when receiving malformed BGP update messages, thereby greatly enhancing network stability. However, in practical applications, there are still numerous instances where BGP session oscillations occur due to the receipt of malformed BGP update messages, unrecognized attribute fields, or routing rules generated by a certain BGP AFI/SAFI that affect the forwarding of BGP messages.

This document introduces some approaches to enhance the stability of BGP sessions.

Requirements Language

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 RFC 2119 [RFC2119] [RFC8174] when, and only when, they appear in all capitals, as shown here.

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 May 2026.

Copyright Notice

Copyright (c) 2025 IETF Trust and the persons identified as the document authors. All rights reserved.

This document is subject to BCP 78 and the IETF Trust's Legal Provisions Relating to IETF Documents (<https://trustee.ietf.org/license-info>) in effect on the date of publication of this document. Please review these documents carefully, as they describe your rights and restrictions with respect to this document. Code Components extracted from this document must include Revised BSD License text as described in Section 4.e of the Trust Legal Provisions and are provided without warranty as described in the Revised BSD License.

Table of Contents

1. Introduction	2
2. Definitions and Acronyms	3
3. BGP Protection Mode	3
4. BGP Diagnostic Mode	5
5. Error Handling	6
6. IANA Considerations	6
7. Security Considerations	6
8. Contributors	6
9. Acknowledgements	7
10. References	7
10.1. Normative References	7
10.2. Informative References	9
Authors' Addresses	9

1. Introduction

As the internet carries an increasing number of services, its stability has become more and more important. The BGP protocol plays a crucial role in the internet, enabling different ISPs and enterprises to provide symmetrical and reliable connection services, ensuring that information and data on the internet can be transmitted and exchanged quickly and securely. The oscillation of BGP protocol sessions can have a significant negative impact on the internet. The method introduced in RFC7606 promotes the stability of BGP protocol sessions. However, in practical applications, there are still

numerous instances where BGP session oscillations occur due to the receipt of malformed BGP update messages, unrecognized attribute fields, or routing rules generated by a certain BGP AFI/SAFI that affect the forwarding of BGP messages.

This document introduces some approaches to enhance the stability of BGP sessions.

2. Definitions and Acronyms

- * AFI: Address Family Identifiers
- * FlowSpec: Flow Specification
- * ISP: Internet Service Provider
- * SAFI: Subsequent Address Family Identifiers

3. BGP Protection Mode

As shown in the figure below, the process of BGP protection mode operation is described.

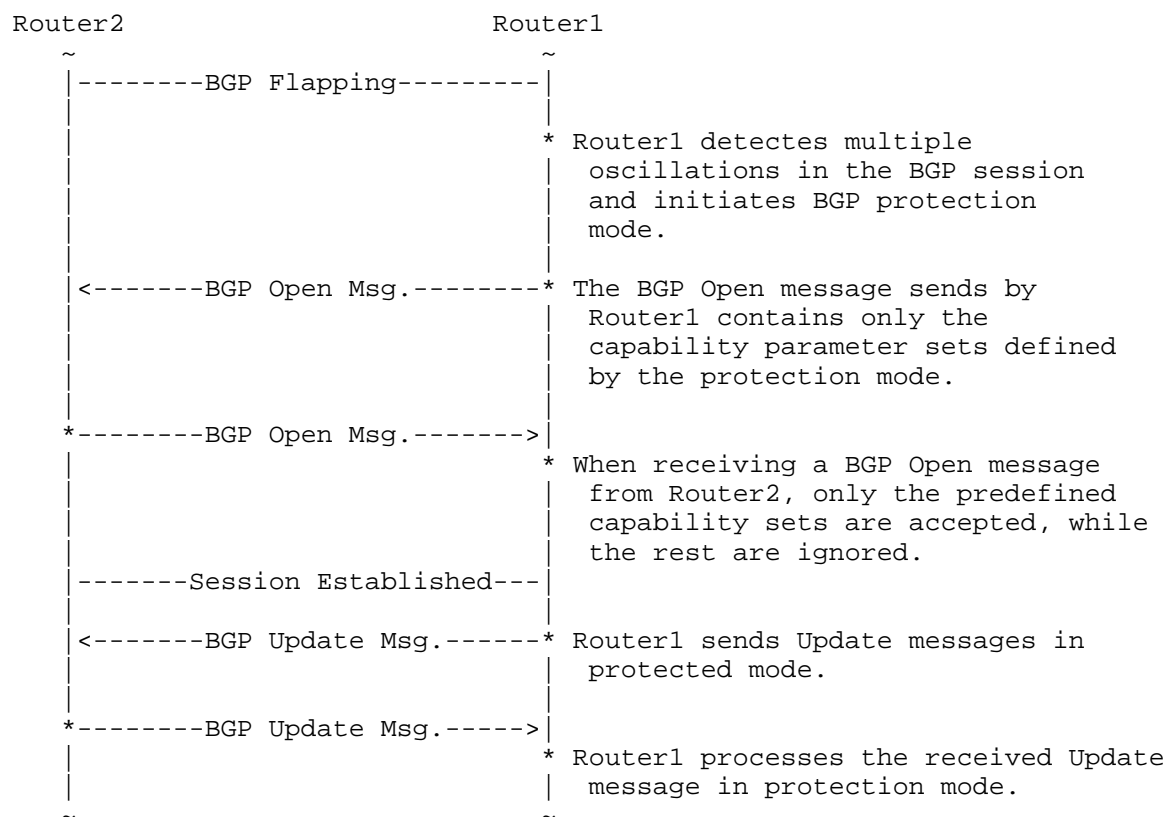


Figure 1: Schematic Diagram of BGP Protection Mode Operation

The processing procedure is as follows:

- 1) Router1 detectes multiple oscillations in the BGP session and initiates BGP protection mode.
- 2) The BGP Open message sent by Router1 contains only the capability parameter sets defined by the protection mode.

For example:

Before implementing this solution: Send IPv4 unicast address family, IPv4 Flowspec address family, Route Refresh capability, etc. In some erroneous operations, the rules published by the Flowspec address family can filter out protocol messages in the BGP session, leading to a prolonged absence of protocol message exchanges and ultimately causing the BGP session to be interrupted.

After implementing this solution: Only send IPv4 unicast address family and Route Refresh capability, and no longer send IPv4 Flowspec address family and other capability parameters that may cause oscillation.

3) When receiving a BGP Open message from Router2, only the predefined capability sets are accepted, while the rest are ignored. After this operational step, Router1 does not have the capability to filter protocol packets for the new session. This eliminates the issue of repeated BGP session flaps caused by problematic Flowspec routes.

4) After the BGP session is established, when Router1 sends a BGP Update message to Router2, the BGP Update message contains only the set of attributes customized for protection mode.

5) When Router1 receives a BGP Update message from Router2, it only accepts the set of attributes in the BGP update that are configured for protection mode, while ignoring other BGP path attributes.

4. BGP Diagnostic Mode

As shown in the figure below, the process of BGP Diagnostic mode operation is described.

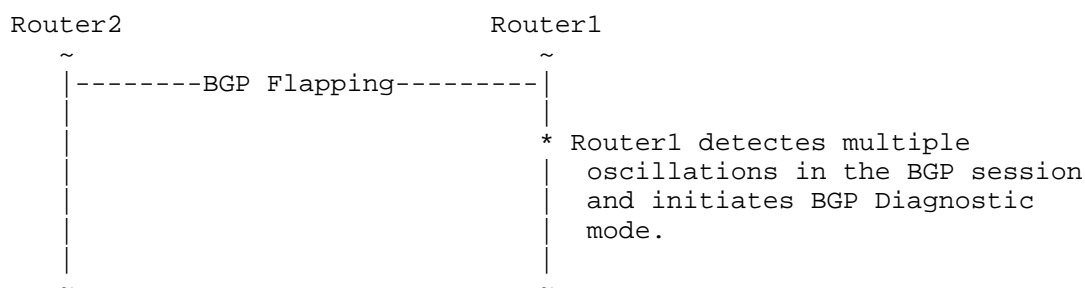


Figure 2: Schematic Diagram of BGP Diagnostic Mode Operation

After the BGP session has flapped multiple times (determine a threshold, for example, 5 times), the router implementing this solution enters diagnostic mode the next time the BGP session starts:

1) Upon entering diagnostic mode, the router allocates additional storage resources to the BGP module and optionally activates some diagnostic modules that are typically disabled in normal mode. These diagnostic modules may usually impact the operational performance of BGP.

2) The router records the BGP messages received and sent to the peer and stores them.

3) The diagnostic module on the router performs a diagnostic analysis of the legality of the BGP messages.

4) If the diagnostic module identifies some issues, when the session restarts or it actively restarts the session: at this point, it excludes the information causing the fault from the messages sent and ignores the information causing the fault when processing the information received from the peer.

5) If the session continues to flap after starting diagnostic mode (either because the diagnostic module did not promptly identify the issue or because there is no diagnostic module), the router initiates the protection mode; the subsequent process is the same as in section 3.

5. Error Handling

TBD

6. IANA Considerations

No IANA actions are required for this document.

7. Security Considerations

This document does not change the security properties of BGP.

8. Contributors

TBD

9. Acknowledgements

The authors would like to acknowledge the review and inputs from ... (TBD)

10. References

10.1. Normative References

- [I-D.ietf-idr-flowspec-redirect-ip]
Haas, J., Henderickx, W., and A. Simpson, "BGP Flow-Spec Redirect-to-IP Action", Work in Progress, Internet-Draft, draft-ietf-idr-flowspec-redirect-ip-04, 2 September 2025, <<https://datatracker.ietf.org/doc/html/draft-ietf-idr-flowspec-redirect-ip-04>>.
- [I-D.ietf-idr-fsv2-ip-basic]
Hares, S., Eastlake, D. E., Dong, J., Yadlapalli, C., and S. Maduschke, "BGP Flow Specification Version 2 - for Basic IP", Work in Progress, Internet-Draft, draft-ietf-idr-fsv2-ip-basic-03, 3 March 2025, <<https://datatracker.ietf.org/doc/html/draft-ietf-idr-fsv2-ip-basic-03>>.
- [I-D.ietf-idr-sr-policy-safi]
Previdi, S., Filsfils, C., Talaulikar, K., Mattes, P., and D. Jain, "Advertising Segment Routing Policies in BGP", Work in Progress, Internet-Draft, draft-ietf-idr-sr-policy-safi-13, 6 February 2025, <<https://datatracker.ietf.org/doc/html/draft-ietf-idr-sr-policy-safi-13>>.
- [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>>.
- [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>>.
- [RFC4360] Sangli, S., Tappan, D., and Y. Rekhter, "BGP Extended Communities Attribute", RFC 4360, DOI 10.17487/RFC4360, February 2006, <<https://www.rfc-editor.org/info/rfc4360>>.

- [RFC4760] Bates, T., Chandra, R., Katz, D., and Y. Rekhter, "Multiprotocol Extensions for BGP-4", RFC 4760, DOI 10.17487/RFC4760, January 2007, <<https://www.rfc-editor.org/info/rfc4760>>.
- [RFC5701] Rekhter, Y., "IPv6 Address Specific BGP Extended Community Attribute", RFC 5701, DOI 10.17487/RFC5701, November 2009, <<https://www.rfc-editor.org/info/rfc5701>>.
- [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>>.
- [RFC8660] Bashandy, A., Ed., Filsfils, C., Ed., Previdi, S., Decraene, B., Litkowski, S., and R. Shakir, "Segment Routing with the MPLS Data Plane", RFC 8660, DOI 10.17487/RFC8660, December 2019, <<https://www.rfc-editor.org/info/rfc8660>>.
- [RFC8669] Previdi, S., Filsfils, C., Lindem, A., Ed., Sreekantiah, A., and H. Gredler, "Segment Routing Prefix Segment Identifier Extensions for BGP", RFC 8669, DOI 10.17487/RFC8669, December 2019, <<https://www.rfc-editor.org/info/rfc8669>>.
- [RFC8955] Loibl, C., Hares, S., Raszuk, R., McPherson, D., and M. Bacher, "Dissemination of Flow Specification Rules", RFC 8955, DOI 10.17487/RFC8955, December 2020, <<https://www.rfc-editor.org/info/rfc8955>>.
- [RFC8956] Loibl, C., Ed., Raszuk, R., Ed., and S. Hares, Ed., "Dissemination of Flow Specification Rules for IPv6", RFC 8956, DOI 10.17487/RFC8956, December 2020, <<https://www.rfc-editor.org/info/rfc8956>>.
- [RFC9012] Patel, K., Van de Velde, G., Sangli, S., and J. Scudder, "The BGP Tunnel Encapsulation Attribute", RFC 9012, DOI 10.17487/RFC9012, April 2021, <<https://www.rfc-editor.org/info/rfc9012>>.
- [RFC9252] Dawra, G., Ed., Talaulikar, K., Ed., Raszuk, R., Decraene, B., Zhuang, S., and J. Rabadan, "BGP Overlay Services Based on Segment Routing over IPv6 (SRv6)", RFC 9252, DOI 10.17487/RFC9252, July 2022, <<https://www.rfc-editor.org/info/rfc9252>>.

- [RFC9256] Filsfils, C., Talaulikar, K., Ed., Voyer, D., Bogdanov, A., and P. Mattes, "Segment Routing Policy Architecture", RFC 9256, DOI 10.17487/RFC9256, July 2022, <<https://www.rfc-editor.org/info/rfc9256>>.

10.2. Informative References

- [I-D.ietf-pce-segment-routing-policy-cp] Koldychev, M., Sivabalan, S., Sidor, S., Barth, C., Peng, S., and H. Bidgoli, "Path Computation Element Communication Protocol (PCEP) Extensions for Segment Routing (SR) Policy Candidate Paths", Work in Progress, Internet-Draft, draft-ietf-pce-segment-routing-policy-cp-27, 4 April 2025, <<https://datatracker.ietf.org/doc/html/draft-ietf-pce-segment-routing-policy-cp-27>>.
- [RFC8754] Filsfils, C., Ed., Dukes, D., Ed., Previdi, S., Leddy, J., Matsushima, S., and D. Voyer, "IPv6 Segment Routing Header (SRH)", RFC 8754, DOI 10.17487/RFC8754, March 2020, <<https://www.rfc-editor.org/info/rfc8754>>.

Authors' Addresses

Shunwan Zhuang
Huawei Technologies
Huawei Bld., No.156 Beiqing Rd.
Beijing
100095
China
Email: zhuangshunwan@huawei.com

Haibo Wang
Huawei Technologies
Huawei Bld., No.156 Beiqing Rd.
Beijing
100095
China
Email: rainsword.wang@huawei.com