

Global Routing Operations  
Internet-Draft  
Obsoletes: 7454 (if approved)  
Intended status: Best Current Practice  
Expires: 9 October 2026

T. Fiebig  
MPI-INF  
N. Hilliard  
INEX  
7 April 2026

BGP Operations and Security  
draft-ietf-grow-bgpopssecupd-15

## Abstract

The Border Gateway Protocol (BGP) is a critical component in the Internet to exchange routing information between network domains. It is important to understand the security and reliability requirements that can and should be met to prevent accidental or intentional routing disturbances.

Previously, security considerations for BGP have been described in RFC7454 / BCP194. Since the publication of RFC7454, changes in operational practice have taken place, which are partially conflicting with the advice given in RFC7454. This document obsoletes RFC7454, and provides less implementation-specific best practices, with the goal of being less prone to becoming outdated or conflicting with changed operational practices.

## 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 9 October 2026.

## Copyright Notice

Copyright (c) 2026 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
1.1. Requirements Language . . . . .	3
2. Scope of the Document . . . . .	3
3. Protection of the BGP Speaker and Session . . . . .	3
3.1. Protecting the BGP Session . . . . .	3
3.2. Protecting the Management Interface of the BGP Speaker . . . . .	4
4. Route Filtering . . . . .	4
4.1. Importing BGP Routes . . . . .	4
4.2. Originating and Propagating BGP Routes . . . . .	5
4.3. Altering BGP Routes . . . . .	6
5. IANA Considerations . . . . .	6
6. Security Considerations . . . . .	6
7. References . . . . .	6
7.1. Normative References . . . . .	6
7.2. Informative References . . . . .	7
Acknowledgements . . . . .	8
Authors' Addresses . . . . .	9

## 1. Introduction

The Border Gateway Protocol (BGP), specified in [RFC4271], is the protocol used in the Internet to exchange routing information between network domains. BGP does not directly include mechanisms that control whether the routes exchanged conform to the various guidelines defined by the Internet community. Furthermore, the BGP protocol itself, by its design, does not have any direct way to protect itself against threats to confidentiality, integrity, and availability.

This document summarizes security properties and requirements when operating BGP for securing the infrastructure as well as security considerations regarding the exchanged routing information. Operators are advised to consult documentation and contemporary informational documents concerning methods to ensure that these properties are sufficiently ensured in their network.

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

## 2. Scope of the Document

The guidelines defined in this document are intended for BGP when used to exchange generic Internet routing information within the Default-Free Zone (DFZ). It specifically does not cover other uses of BGP, e.g., when using BGP for exchanging routes in a data-center context, or other use-cases when using BGP without globally unique identifiers between networks. This document does not specify how the outlined requirements and properties can be technically realized at a specific point in time. Instead, operators are advised to consult applicable documentation and contemporary informational documents describing implementation specifics (e.g., [I-D.ietf-grow-routing-ops-sec-inform] and [I-D.ietf-grow-routing-ops-terms]).

## 3. Protection of the BGP Speaker and Session

The BGP speaker, i.e., the node running BGP to exchange routing information, needs to be protected from external attempts to impact integrity or availability of the BGP session and the node.

### 3.1. Protecting the BGP Session

To protect a BGP speaker on the network layer, an operator MUST ensure the following properties using technical or organizational measures:

- \* Prevent off-path attackers from injecting BGP messages into existing sessions.
- \* Prevent off-path attackers from interrupting existing sessions.
- \* Prevent off-path attackers from preventing the establishment of new sessions.
- \* Prevent remote systems from overwhelming the BGP speaker by sending large volumes of unsolicited packets or BGP messages.
- \* Ensure that unstable sessions do not threaten the availability of BGP speakers within the network.

Example technologies to accomplish this include GTSM/TTL-security [RFC5082], BGP-MD5 / TCP-AO [RFC5925], limiting traffic to the control plane via Control Plane Policing (CoPP), and setting maximum prefix limits for the number of prefixes a neighbor may send. When implementing prefix limits, operators SHOULD be aware of the operational implications of exceeding prefix limits, i.e., a loss of an established session. Hence, operators SHOULD appropriately weigh this impact within the specific operational circumstances, and ensure appropriate prefix limits to not cause outages under normal operations.

### 3.2. Protecting the Management Interface of the BGP Speaker

In addition to the control plane / exchange of BGP protocol messages, the management plane of BGP speakers must be appropriately secured. Hence, operators MUST ensure that:

- \* No unauthorized third-parties can obtain access or connect to the management interface of a BGP speaker in a way that impacts confidentiality, integrity, or availability.
- \* External activity towards the management interface does not interfere with the integrity or availability of BGP sessions.

## 4. Route Filtering

The purpose of BGP is to exchange routing information. Importing or exporting incorrect or malicious routes is a security risk for receiving networks and may threaten connected and/or remote networks. As such, operators MUST ensure the following properties when importing or exporting routing information from their neighbors.

### 4.1. Importing BGP Routes

When importing BGP routes from a neighbor, an operator MUST ensure that all imported routes conform to the following properties by implementing technical or organizational measures:

- \* The AS originating BGP routes for a prefix MUST be globally authorized to originate that prefix. Operators MAY deviate from this for default routes (::/0 and 0.0.0.0/0), if they granted the specific neighbor permission to announce default routes towards them. Operators are cautioned to evaluate carefully how accepting a default route affects their network, as this occludes limitations in forwarding coverage by the upstream from which the default route was received.

- \* For received BGP routes with an AS\_PATH = {AS1, AS2, ..., ASn}, where AS1 is the neighbor that sent the UPDATE and ASn is the originator, for each k in 1..n-1, AS(k+1) MUST be authorized to export the BGP routes to ASk according to their bilateral routing policy (e.g., providercustomer, peer, or lateral-peer).
- \* The AS\_PATH MUST NOT contain AS numbers reserved for private [RFC6996] or special-use cases, except for those AS numbers explicitly dedicated to a special-use that requires their presence in the global routing table [IANAASNSpec].
- \* The length of the AS\_PATH for a received BGP routes MUST NOT exceed the maximum length supported by the local router.
- \* The number of transitive BGP attributes, e.g., BGP communities [RFC1997], extended BGP communities [RFC4360], or Large BGP communities [RFC8092] attached to a received NLRI MUST NOT exceed the maximum supported by the local router.
- \* The number of BGP routes received from a neighbor MUST NOT exceed the resources of the local router.

#### 4.2. Originating and Propagating BGP Routes

When originating or propagating BGP routes, an operator MUST ensure that all BGP routes they export conform to the following properties by implementing technical or organizational measures:

- \* The propagating AS MUST be authorized to propagate BGP routes for the specific prefix when received from the AS directly to its right in the AS\_PATH. Additionally, each AS in the AS\_PATH not originating the prefix MUST be authorized to propagate the prefix when receiving it from the next AS to its right.
- \* The AS originating BGP routes for a prefix MUST be globally authorized to originate that prefix. Operators MAY deviate from this for default routes (::/0 and 0.0.0.0/0), if they originate the default route and the specific neighbor granted them permission to announce default routes towards them.
- \* The AS\_PATH MUST NOT contain AS numbers reserved for private [RFC6996] or special-use cases, except for those AS numbers explicitly dedicated to a special-use that requires their presence in the global routing table [IANAASNSpec].

#### 4.3. Altering BGP Routes

When processing BGP routes, an operator MUST ensure that the basic properties of these routes are not altered:

- \* An operator SHOULD NOT change or remove immutable transitive BGP attributes, e.g., ORIGIN as per [RFC4271]. In selected cases, if a specific attribute is known to be malicious, an operator MAY either remove that specific attribute from the BGP routes when importing them or filter the routes carrying the attribute. Transitivity of attributes unknown to an operator cannot be established. Treating such attributes as immutable enables incremental deployment of new BGP features, while processing unknown attributes may harm availability if the eBGP speakers used by an operator are unable to handle the attribute safely. Hence, operators SHOULD carefully assess the tradeoff between incremental deployment and BGP security for their network.
- \* Routes carried by BGP MUST NOT carry RPKI validation states in transitive BGP path attributes [I-D.ietf-sidrops-avoid-rpki-state-in-bgp].

#### 5. IANA Considerations

This document does not require any IANA actions.

#### 6. Security Considerations

This document is entirely about BGP operational security. It lists requirements and properties operators MUST ensure using technical or organizational measures when operating BGP routers in the DFZ.

#### 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/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>>.

- [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>>.
- [RFC6996] Mitchell, J., "Autonomous System (AS) Reservation for Private Use", BCP 6, RFC 6996, DOI 10.17487/RFC6996, July 2013, <<https://www.rfc-editor.org/info/rfc6996>>.
- [IANAASNSpec]  
IANA, "Special-Purpose Autonomous System (AS) Numbers", <<https://www.iana.org/assignments/iana-as-numbers-special-registry/iana-as-numbers-special-registry.xhtml>>.

## 7.2. Informative References

- [RFC1997] Chandra, R., Traina, P., and T. Li, "BGP Communities Attribute", RFC 1997, DOI 10.17487/RFC1997, August 1996, <<https://www.rfc-editor.org/info/rfc1997>>.
- [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>>.
- [RFC5082] Gill, V., Heasley, J., Meyer, D., Savola, P., Ed., and C. Pignataro, "The Generalized TTL Security Mechanism (GTSM)", RFC 5082, DOI 10.17487/RFC5082, October 2007, <<https://www.rfc-editor.org/info/rfc5082>>.
- [RFC5925] Touch, J., Mankin, A., and R. Bonica, "The TCP Authentication Option", RFC 5925, DOI 10.17487/RFC5925, June 2010, <<https://www.rfc-editor.org/info/rfc5925>>.
- [RFC7454] Durand, J., Pepelnjak, I., and G. Doering, "BGP Operations and Security", BCP 194, RFC 7454, DOI 10.17487/RFC7454, February 2015, <<https://www.rfc-editor.org/info/rfc7454>>.
- [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/info/rfc8092>>.
- [I-D.ietf-grow-routing-ops-sec-inform]  
Fiebig, T., "Current Options for Securing Global Routing", Work in Progress, Internet-Draft, draft-ietf-grow-routing-ops-sec-inform, 9 April 2025, <<https://datatracker.ietf.org/doc/draft-ietf-grow-routing-ops-sec-inform/>>.

`[I-D.ietf-grow-routing-ops-terms]`

Fiebig, T., "Currently Used Terminology in Global Routing Operations", Work in Progress, Internet-Draft, draft-ietf-grow-routing-ops-terms, 9 April 2025, <<https://datatracker.ietf.org/doc/draft-ietf-grow-routing-ops-terms/>>.

`[I-D.ietf-sidrops-avoid-rpki-state-in-bgp]`

Snijders, J., Fiebig, T., and M. A. 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, 3 October 2024, <<https://datatracker.ietf.org/doc/draft-ietf-sidrops-avoid-rpki-state-in-bgp/>>.

## Acknowledgements

This document has been originally based on [RFC7454] and we thank the original authors for their work.

We thank the following people for reviewing this draft and suggesting changes:

- \* Gert Doerring
- \* Jeff Haas
- \* Geng Nan
- \* Martin Pels
- \* Job Snijders
- \* Berislav Todorovic
- \* Linda Dunbar
- \* Wolfgang Tremmel
- \* Florian Obser
- \* Ben Maddison
- \* Mohamed Boucadair
- \* Tom Petch

Authors' Addresses

Tobias Fiebig  
Max-Planck-Institut fuer Informatik  
Campus E14  
66123 Saarbruecken  
Germany  
Phone: +49 681 9325 3527  
Email: tfiebig@mpi-inf.mpg.de

Nick Hilliard  
Internet Neutral Exchange Association  
4027 Kingswood Road  
Citywest, Dublin  
D24 AX96  
Ireland  
Phone: +353 1 433 205 2  
Email: nick@inex.ie