

Workgroup: Network Working Group
Internet-Draft:
draft-ietf-idr-deprecate-as-set-confed-set-07
Obsoletes: [6472](#) (if approved)
Updates: [4271](#) [5065](#) (if approved)
Published: 7 March 2022
Intended Status: Standards Track
Expires: 8 September 2022
Authors: W. Kumari K. Sriram L. Hannachi
 Google, Inc. USA NIST USA NIST
 J. Haas
 Juniper Networks, Inc.
 Deprecation of AS_SET and AS_CONFED_SET in BGP

Abstract

BCP 172 (i.e., RFC 6472) recommends not using AS_SET and AS_CONFED_SET in the Border Gateway Protocol. This document advances this recommendation to a standards requirement in BGP; it proscribes the use of the AS_SET and AS_CONFED_SET types of path segments in the AS_PATH. This is done to simplify the design and implementation of BGP and to make the semantics of the originator of a route clearer. This will also simplify the design, implementation, and deployment of various BGP security mechanisms. This document (if approved) updates RFC 4271 and RFC 5065 by eliminating AS_SET and AS_CONFED_SET types, and obsoletes RFC 6472.

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 8 September 2022.

Copyright Notice

Copyright (c) 2022 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. Requirements Language](#)
- [3. Recommendations](#)
- [4. Updates to Existing RFCs](#)
- [5. Operational Considerations](#)
- [6. Security Considerations](#)
- [7. IANA Considerations](#)
- [8. Acknowledgements](#)
- [9. References](#)
 - [9.1. Normative References](#)
 - [9.2. Informative References](#)
- [Authors' Addresses](#)

1. Introduction

BCP 172 [[RFC6472](#)] makes a recommendation for not using AS_SET (see [[RFC4271](#)]) and AS_CONFED_SET (see [[RFC5065](#)]) in the Border Gateway Protocol (BGP). This document advances the BCP recommendation to a standards requirement in BGP; it proscribes the use of the AS_SET and AS_CONFED_SET types of path segments in the AS_PATH.

The AS_SET path segment in the AS_PATH attribute (Sections 4.3 and 5.1.2 of [[RFC4271](#)]) is created by a router that is performing route aggregation and contains an unordered set of Autonomous Systems (ASes) that contributing prefixes in the aggregate have traversed. The AS_CONFED_SET path segment (see [[RFC5065](#)]) in the AS_PATH attribute is created by a router that is performing route aggregation and contains an unordered set of Member AS Numbers in the local confederation that contributing prefixes in the aggregate have traversed. It is very similar to an AS_SET but is used within a confederation.

By performing aggregation, a router is combining multiple existing routes into a single new route. The aggregation together with the use of AS_SET blurs the semantics of origin AS for the prefix being announced. Therefore, the aggregation with AS_SET (or AS_CONFED_SET) can cause operational issues, such as not being able to authenticate

a route origin for the aggregate prefix in new BGP security technologies such as those that take advantage of X.509 extensions for IP addresses and AS identifiers [[RFC3779](#)] [[RFC6480](#)] [[RFC6811](#)] [[RFC8205](#)]. This in turn could result in reachability problems for the aggregated prefix and its components (i.e., more specific prefixes).

From analysis of past Internet routing data, it is apparent that aggregation that involves AS_SETs is very seldom used in practice on the public Internet [[Analysis](#)] and when it is used, it is often used incorrectly -- only a single AS in the AS_SET are by far the most common cases. Also, very often the same AS appears in the AS_SEQUENCE and the AS_SET in the BGP update. The occurrence of reserved AS numbers ([[IANA-SP-ASN](#)]) is also somewhat frequent. Because the aggregation involving AS_SETs is very rarely used, the reduction in table size provided by this is extremely small, and any advantage thereof is outweighed by additional complexity in BGP. As noted above, AS_SETs also pose impediments to implementation of new BGP security technologies.

2. 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.

3. Recommendations

BGP speakers conforming to this document (i.e., conformant BGP speakers) MUST NOT locally generate BGP UPDATE messages containing AS_SET or AS_CONFED_SET. Conformant BGP speakers SHOULD NOT send BGP UPDATE messages containing AS_SET or AS_CONFED_SET. Upon receipt of such messages, conformant BGP speakers SHOULD use the "Treat-as-withdraw" error handling behavior as per [[RFC7606](#)].

If a network operator wishes to consider BGP UPDATE messages with AS_SET or AS_CONFED_SET (received from an external peer) for path selection, they MAY have a feature (knob) in their BGP speaker to opt to do so on a per peer basis. The operator should understand the full implications of choosing this option. There is no knob concerning locally generated BGP UPDATE messages, i.e., as stated before a conformant BGP speaker must not locally generate BGP UPDATE messages with AS_SET or AS_CONFED_SET.

Network operators MUST NOT locally generate any new announcements containing AS_SET or AS_CONFED_SET. If they have announced routes with AS_SET or AS_CONFED_SET in them, then they SHOULD withdraw

those routes and re-announce routes for the aggregate or component prefixes (i.e., the more specific routes subsumed by the previously aggregated route) without AS_SET or AS_CONFED_SET in the updates.

It is worth noting that new BGP security technologies (such as those that take advantage of X.509 extensions for IP addresses and AS identifiers [[RFC3779](#)] [[RFC6480](#)] [[RFC6811](#)] [[RFC8205](#)]) might not support routes with AS_SET or AS_CONFED_SET in them, and may treat routes containing them as infeasible even before the updated BGP in this document is implemented.

4. Updates to Existing RFCs

This document deprecates the AS_SET (type 1) AS_PATH segment type from [[RFC4271](#)]. BGP speakers conforming to this document (i.e., conformant BGP speakers) MUST NOT locally generate BGP UPDATE messages containing AS_SET. Conformant BGP speakers SHOULD NOT send BGP UPDATE messages containing AS_SET. Upon receipt of such messages, conformant BGP speakers SHOULD use the "Treat-as-withdraw" error handling behavior as per [[RFC7606](#)].

This document deprecates the AS_CONFED_SET (type 4) AS_PATH segment type from [[RFC5065](#)]. Conformant BGP speakers MUST NOT locally generate BGP UPDATE messages containing AS_CONFED_SET. Conformant BGP speakers SHOULD NOT send BGP UPDATE messages containing AS_CONFED_SET. Upon receipt of such messages, conformant BGP speakers SHOULD use the "Treat-as-withdraw" error handling behavior as per [[RFC7606](#)].

Wherever mentions of AS_SET or AS_CONFED_SET occur in [[RFC4271](#)] and [[RFC5065](#)], appropriate modification or elimination of the text must be made in future RFCs that would replace these RFCs, consistent with the deprecation of AS_SET and AS_CONFED_SET.

5. Operational Considerations

When aggregating prefixes, network operators MUST use brief aggregation. In brief aggregation, the AGGREGATOR attribute is included but the AS_SET or AS_CONFED_SET attribute is not included.

When doing the above, operators MUST form the aggregate at the border in the outbound BGP policy and omit any prefixes from the AS that the aggregate is being advertised to. In other words, an aggregate prefix MUST NOT be announced to the contributing ASes. Instead, more specific prefixes (from the aggregate) MUST be announced to each contributing AS, excluding any that were learned from the contributing AS in consideration. For illustration, if p1/24 (from AS1), p2/24 (from AS2), p3/24 (from AS3) and p4/24 (from AS4) are aggregated to p/22, then p/22 will not be announced to AS1, AS2, AS3, or AS4. Instead, as further illustration, p1/24, p2/24 and

p4/24 are announced to AS3. Or, possibly q/23 (aggregate of p1/24 and p2/24) and p4/24 are announced to AS3.

Operators MUST install egress filters to block data packets when the destination address belongs to an internal prefix. Similarly, any known single-homed customer prefix MUST also be included in the egress filters except on the interface for that customer. This mitigates looping in the data plane when connection to such an internal or customer prefix is lost. This mechanism effectively compensates for the lack of the additional loop detection capability accorded by AS_SETs (if they were allowed).

6. Security Considerations

This document obsoletes the use of aggregation techniques that create AS_SETs or AS_CONFED_SETs. Obsoleting these path segment types from BGP and removal of the related code from implementations would potentially decrease the attack surface for BGP. Deployments of new BGP security technologies [RFC6480] [RFC6811] [RFC8205] benefit greatly if AS_SET and AS_CONFED_SET are not used in BGP.

7. IANA Considerations

This document requires no IANA actions.

8. Acknowledgements

The authors would like to thank John Heasley, Job Snijders, Jared Mauch, Jakob Heitz, Keyur Patel, Douglas Montgomery, Randy Bush, Susan Hares, John Scudder, Curtis Villamizar, Danny McPherson, Chris Morrow, Tom Petch, Ilya Varlashkin, Enke Chen, Tony Li, Florian Weimer, John Leslie, Paul Jakma, Rob Austein, Russ Housley, Sandra Murphy, Steve Bellovin, Steve Kent, Steve Padgett, Alfred Hoenes, and Alvaro Retana for comments and suggestions.

9. References

9.1. 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>>.
- [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>>.

- [RFC5065] Traina, P., McPherson, D., and J. Scudder, "Autonomous System Confederations for BGP", RFC 5065, DOI 10.17487/RFC5065, August 2007, <<https://www.rfc-editor.org/info/rfc5065>>.

9.2. Informative References

- [Analysis] Hannachi, L. and K. Sriram, "Detailed analysis of AS_SETs in BGP updates", NIST Robust Inter-domain Routing Project Website, October 2019, <https://github.com/ksriram25/IETF/blob/main/Detailed-AS_SET-analysis.txt>.
- [IANA-SP-ASN] "Special-Purpose Autonomous System (AS) Numbers", <<https://www.iana.org/assignments/iana-as-numbers-special-registry/iana-as-numbers-special-registry.xhtml>>.
- [RFC3779] Lynn, C., Kent, S., and K. Seo, "X.509 Extensions for IP Addresses and AS Identifiers", RFC 3779, DOI 10.17487/RFC3779, June 2004, <<https://www.rfc-editor.org/info/rfc3779>>.
- [RFC6472] Kumari, W. and K. Sriram, "Recommendation for Not Using AS_SET and AS_CONFED_SET in BGP", BCP 172, RFC 6472, DOI 10.17487/RFC6472, December 2011, <<https://www.rfc-editor.org/info/rfc6472>>.
- [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/info/rfc6480>>.
- [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/info/rfc6811>>.
- [RFC7606] Chen, E., Ed., Scudder, J., Ed., Mohapatra, P., and K. Patel, "Revised Error Handling for BGP UPDATE Messages", RFC 7606, DOI 10.17487/RFC7606, August 2015, <<https://www.rfc-editor.org/info/rfc7606>>.
- [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>>.
- [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>>.

Authors' Addresses

Warren Kumari
Google, Inc.
1600 Amphitheatre Parkway
Mountain View, CA 94043
United States of America

Phone: [+1 571 748 4373](tel:+15717484373)
Email: warren@kumari.net

Kotikalapudi Sriram
USA NIST
100 Bureau Drive
Gaithersburg, MD 20899
United States of America

Phone: [+1 301 975 3973](tel:+13019753973)
Email: ksriram@nist.gov

Lilia Hannachi
USA NIST
100 Bureau Drive
Gaithersburg, MD 20899
United States of America

Phone: [+1 301 975 3259](tel:+13019753259)
Email: lilia.hannachi@nist.gov

Jeffrey Haas
Juniper Networks, Inc.
1133 Innovation Way
Sunnyvale, CA 94089
United States of America

Email: jhaas@juniper.net