Stream:	Internet Engineering Task Force (IETF)
RFC:	9774
Obsoletes:	6472
Updates:	4271, 5065
Category:	Standards Track
Published:	May 2025
ISSN:	2070-1721

Authors:	W. Kumari	K. Sriram	L. Hannachi	J. Haas
	Google, Inc.	*USA NIST*	*USA NIST*	*Juniper Networks, Inc.*

RFC 9774
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 AS_PATH segment types in the Border Gateway Protocol (BGP). This document advances that recommendation to a standards requirement in BGP; it prohibits the use of the AS_SET and AS_CONFED_SET path segment types 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 BGP route clearer. This will also simplify the design, implementation, and deployment of various BGP security mechanisms. This document updates RFC 4271 by deprecating the origination of BGP routes with AS_SET (Type 1 AS_PATH segment) and updates RFC 5065 by deprecating the origination of BGP routes with AS_CONFED_SET (Type 4 AS_PATH segment). Finally, it obsoletes RFC 6472.

Status of This Memo

This is an Internet Standards Track document.

This document is a product of the Internet Engineering Task Force (IETF). It represents the consensus of the IETF community. It has received public review and has been approved for publication by the Internet Engineering Steering Group (IESG). Further information on Internet Standards is available in Section 2 of RFC 7841.

Information about the current status of this document, any errata, and how to provide feedback on it may be obtained at https://www.rfc-editor.org/info/rfc9774.

Copyright Notice

Table of Contents

1. Introduction

[BCP172] recommends not using AS_SET [RFC4271] and AS_CONFED_SET [RFC5065] AS_PATH path segment types in the Border Gateway Protocol (BGP). This document advances the BCP recommendation to a standards requirement in BGP; it prohibits the use of the AS_SET and AS_CONFED_SET types of path segments in the AS_PATH. The purpose is to simplify the design and implementation of BGP and to make the semantics of the originator of a BGP route clearer. This will also simplify the design, implementation, and deployment of various BGP security mechanisms. In particular, the prohibition of AS_SETs and AS_CONFED_SETs removes any ambiguity about the origin AS in RPKI-based Route Origin Validation (RPKI-ROV) [RFC6811] [RFC6907] [RFC9319].

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 [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 BGP routes for more specific destinations into a new route for a less specific destination (see [RFC4271], Section 9.1.2.2). Aggregation may blur the semantics of the origin AS for the prefix being announced by producing an AS_SET or AS_CONFED_SET. Such sets 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 (see [RFC6480], [RFC6811], [RFC6907], [RFC8205], and [RFC9319]). This could result in reachability problems for the destinations covered by the aggregated prefix.

From analysis of historical Internet routing data, it is apparent that aggregation that involves AS_SETs is very seldom used in practice on the public Internet (see [Analysis]). When it is used, it is often used incorrectly; only a single AS in the AS_SET is the most common case [Analysis]. 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.

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. Updates to the Requirements of RFCs 4271 and 5065

Unless explicitly configured by a network operator to do otherwise (e.g., during a transition phase), BGP speakers:

- **MUST NOT** advertise BGP UPDATE messages containing AS_SETs or AS_CONFED_SETs and
- **MUST** use the "treat-as-withdraw" error handling behavior per [RFC7606] upon reception of BGP UPDATE messages containing AS_SETs or AS_CONFED_SETs in the AS_PATH or AS4_PATH [RFC6793].

Per the above specifications, this document updates [RFC4271] and [RFC5065] by deprecating AS_SET (see [RFC4271], Section 4.3) and AS_CONFED_SET (see [RFC5065], Section 3), respectively.

4. Treatment of Routes with AS_SET in RPKI-Based BGP Security

Resource Public Key Infrastructure (RPKI) [RFC6480] uses X.509 extensions for IP addresses and AS identifiers [RFC3779]. RPKI-ROV [RFC6811] [RFC6907] is a BGP security technology that never allows a route with AS_SET to be considered Valid. BGPsec [RFC8205] and Autonomous System Provider Authorization (ASPA) [ASPA-VERIFICATION] are also BGP security technologies based on RPKI. BGPsec does not support AS_SETs. In ASPA-based AS_PATH verification, a route with AS_SET is always considered Invalid and hence ineligible for route selection.

5. BGP AS_PATH "Brief" Aggregation

Sections 9.1.4 and 9.2.2.2 of [RFC4271] describe BGP aggregation procedures. Appendix F.6 of [RFC4271] describes a generally less utilized "Complex AS_PATH Aggregation" procedure.

[RFC4271], Section 5.1.6 describes the ATOMIC_AGGREGATE Path Attribute and notes that:

> When a BGP speaker aggregates several routes for the purpose of advertisement to a particular peer, the AS_PATH of the aggregated route normally includes an AS_SET formed from the set of ASes from which the aggregate was formed. In many cases, the network administrator can determine if the aggregate can safely be advertised without the AS_SET, and without forming route loops.

> If an aggregate excludes at least some of the AS numbers present in the AS_PATH of the routes that are aggregated as a result of dropping the AS_SET, the aggregated route, when advertised to the peer, **SHOULD** include the ATOMIC_AGGREGATE attribute.

When BGP AS_PATH aggregation is done according to the procedures in [RFC4271], Section 9.2.2.2, and any resulting AS_SETs are discarded, it is typically referred to as "brief" aggregation in implementations. Brief aggregation results in an AS_PATH that has the following property (from [RFC4271], Section 9.2.2.2):

> [D]etermine the longest leading sequence of tuples (as defined above) common to all the AS_PATH attributes of the routes to be aggregated. Make this sequence the leading sequence of the aggregated AS_PATH attribute.

The ATOMIC_AGGREGATE Path Attribute is subsequently attached to the BGP route, if AS_SETs are dropped.

5.1. Issues with "Brief" AS_PATH Aggregation and RPKI-ROV

While brief AS_PATH aggregation has the desirable property of not containing AS_SETs, the resulting aggregated AS_PATH may contain an unpredictable origin AS. This is because the aggregating AS may be different from the purported origin AS (for the aggregate), which may vary as explained below. Such unpredictable origin ASes may result in RPKI-ROV validation issues:

- Depending on the contributing routes to the aggregate route, the resulting origin AS may vary.
- The presence of expected contributing routes may be unpredictable due to route availability from BGP neighbors.
- In the presence of such varying origin ASes, it would be necessary for the resource holder to register ROAs [RFC9582] for each potential origin AS that may result from the expected aggregated AS_PATHs.

5.2. Recommendations to Mitigate Unpredictable AS_PATH Origins for RPKI-ROV Purposes

To ensure a consistent BGP origin AS is announced for aggregate BGP routes for implementations of "brief" BGP aggregation, the implementation **MUST** be configured to truncate the AS_PATH after the right-most instance of the desired origin AS for the aggregate. The desired origin AS could be the aggregating AS itself. A ROA would be necessary for the aggregate prefix with the desired origin AS.

This form of brief aggregation is referred to as "consistent brief" BGP aggregation.

If the resulting AS_PATH would be truncated from the otherwise expected result of BGP AS_PATH aggregation (an AS_SET would not be generated and possibly some ASes are removed from the "longest leading sequence" of ASes), the ATOMIC_AGGREGATE Path Attribute **SHOULD** be attached. This is consistent with the intent of [RFC4271], Section 5.1.6.

6. Operational Considerations

This section provides advice to operators regarding deployment and configuration.

6.1. Implementing Consistent Brief Aggregation

When aggregating prefixes, network operators **MUST** use consistent brief aggregation as described in Section 5.2. In consistent brief aggregation, the AGGREGATOR and ATOMIC_AGGREGATE Path Attributes are included, but the AS_PATH does not have AS_SET or AS_CONFED_SET path segment types. See Appendix B for examples of brief aggregation while keeping the origin AS unambiguous and generating appropriate ROAs.

6.2. Not Advertising Aggregate Routes to Contributing ASes

An aggregate prefix **SHOULD NOT** be announced to the contributing ASes. Instead, more specific prefixes (from the aggregate) **SHOULD** be announced to each contributing AS, excluding any that were learned from the contributing AS in consideration. See Appendix A for an example of this filtering policy.

6.3. Mitigating Forwarding Loops

When both less specific and more specific destinations are present, it's possible to create forwarding loops between networks, as discussed in Section 5.1 of [RFC4632].

As a reminder, Rule #2 in Section 5.1 of [RFC4632] requires that BGP implementations performing aggregation discard packets that match the aggregate route but do not match any of the more specific routes.

Further discussion of forwarding loops and their relationship to AS_SETs can be found in Appendix C.

7. Security Considerations

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

8. IANA Considerations

This document has no IANA actions.

9. References

9.1. Normative References

[BCP172] Best Current Practice 172, <https://www.rfc-editor.org/info/bcp172>.
At the time of writing, this BCP comprises the following:

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>.

[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>.

[RFC4632] Fuller, V. and T. Li, "Classless Inter-domain Routing (CIDR): The Internet Address Assignment and Aggregation Plan", BCP 122, RFC 4632, DOI 10.17487/RFC4632, August 2006, <https://www.rfc-editor.org/info/rfc4632>.

[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>.

[RFC6793] Vohra, Q. and E. Chen, "BGP Support for Four-Octet Autonomous System (AS) Number Space", RFC 6793, DOI 10.17487/RFC6793, December 2012, <https://www.rfc-editor.org/info/rfc6793>.

[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>.

9.2. Informative References

[Analysis]

"Detailed analysis of AS_SETs in BGP updates", commit eb0fc22, March 2022, <https://github.com/ksriram25/IETF/blob/main/Detailed-AS_SET-analysis.txt>.

[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-22, 23 March 2025, <https://datatracker.ietf.org/doc/html/draft-ietf-sidrops-aspa-verification-22>.

[IANA-SP-ASN] IANA, "Special-Purpose Autonomous System (AS) Numbers", <https://www.iana.org/assignments/iana-as-numbers-special-registry>.

[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>.

[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>.

[RFC6907] Manderson, T., Sriram, K., and R. White, "Use Cases and Interpretations of Resource Public Key Infrastructure (RPKI) Objects for Issuers and Relying Parties", RFC 6907, DOI 10.17487/RFC6907, March 2013, <https://www.rfc-editor.org/info/rfc6907>.

[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>.

[RFC9319] Gilad, Y., Goldberg, S., Sriram, K., Snijders, J., and B. Maddison, "The Use of maxLength in the Resource Public Key Infrastructure (RPKI)", BCP 185, RFC 9319, DOI 10.17487/RFC9319, October 2022, <https://www.rfc-editor.org/info/rfc9319>.

[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/info/rfc9582>.

Appendix A. Example of Route Filtering for Aggregate Routes and Their Contributors

The illustration presented below shows how an AS_SET is not used when aggregating and how data plane route loops are avoided. Consider that p1/24 (from AS 64501), p2/24 (from AS 64502), p3/24 (from AS 64503), and p4/24 (from AS 64504) are aggregated by AS 64505 to p/22. AS_SET is not used with the aggregate p/22 but AGGREGATOR and ATOMIC AGGREGATE are used. Data

plane route loops are avoided by not announcing the aggregate p/22 to the contributing ASes, i.e., AS 64501, AS 64502, AS 64503, and AS 64504. Instead, as further illustrated, p1/24, p2/24, and p4/24 are announced to AS 64503. The routing tables (post aggregation) of each of the ASes are depicted in the diagram below.

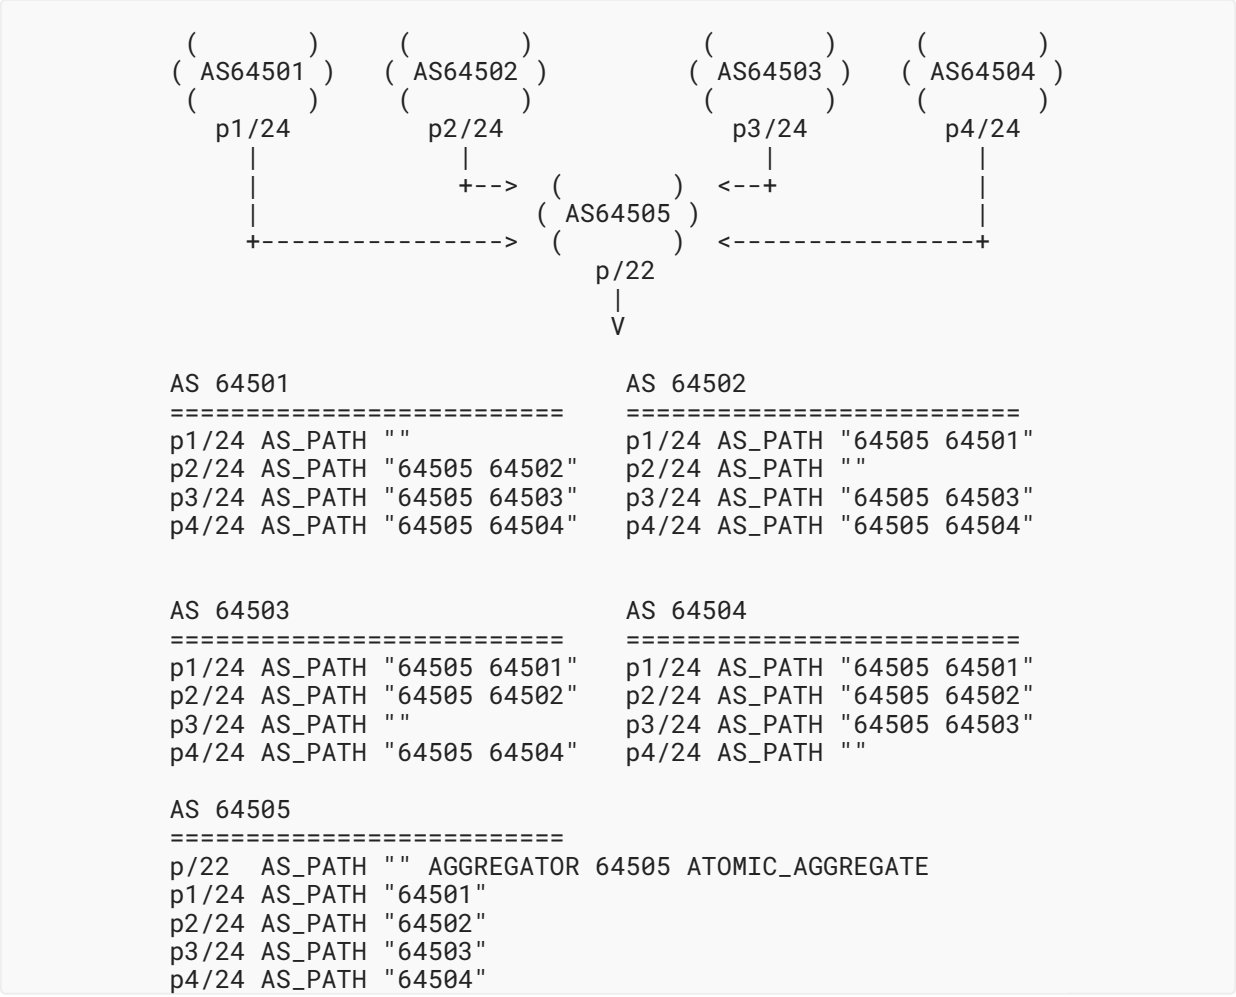

Appendix B. Examples of Consistent and Inconsistent BGP Origin AS Generated by Traditional Brief Aggregation

The examples below illustrate how traditional brief aggregation may result in an inconsistent origin AS.

AS 64500 aggregates more specific routes into 192.0.2.0/24.

Consider the following scenarios where brief aggregation is done by AS 64500 and what the resultant origin ASes would be.

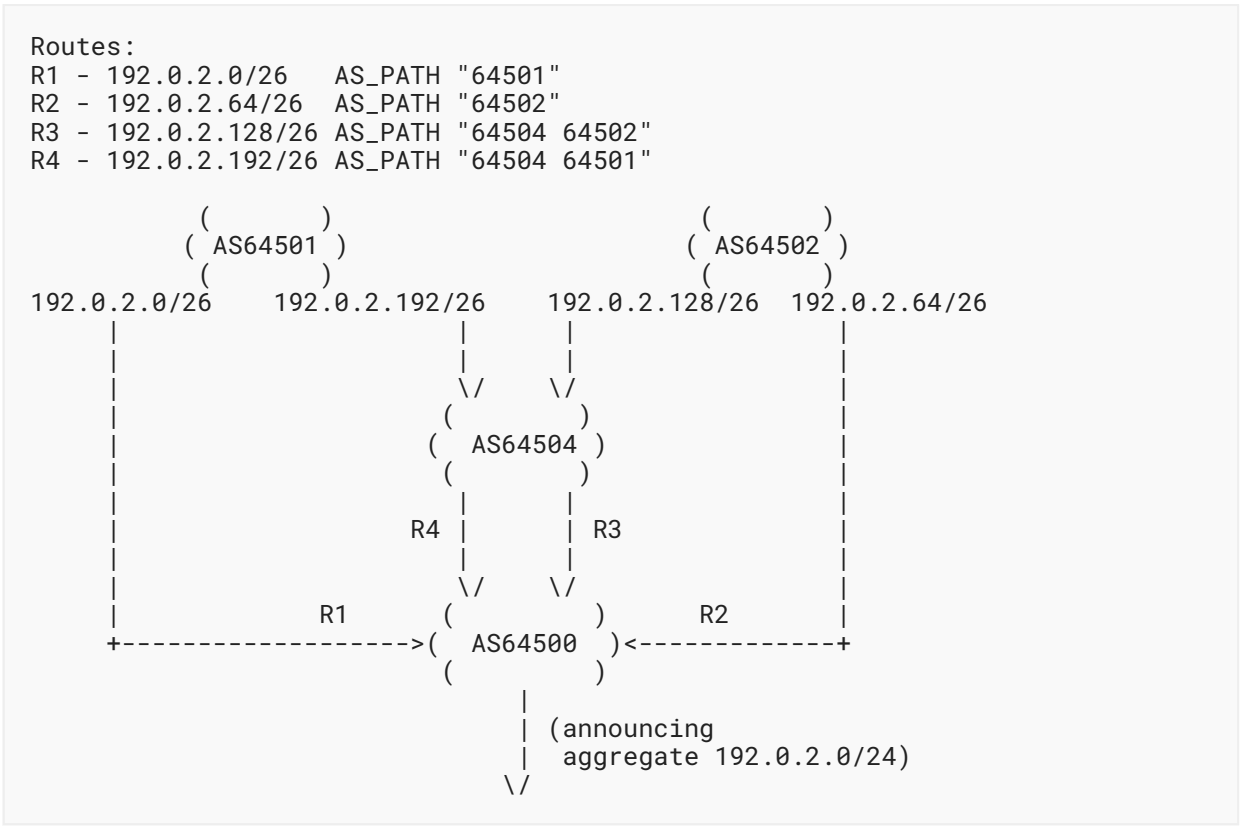

B.1. Scenario 1: First one route, then another, each with a fully disjoint AS_PATH

Receive R1. Aggregate 192.0.2.0/24 AS_PATH "64501"

Alternate "bug?": Aggregate 192.0.2.0/24 AS_PATH "[64501]"

(Note: AS numbers within square brackets represent an AS_SET.)

Receive R2. Aggregate 192.0.2.0/24 AS_PATH "[64501 64502]"

If brief aggregation is in use, the AS_PATH would be truncated to the empty AS_PATH, "".

The resulting AS_PATH is thus not stable and depends on the presence of specific routes.

B.2. Scenario 2: First one route, then another, and the AS_PATHs overlap at the origin AS

Receive R1. Aggregate 192.0.2.0/24 AS_PATH "64501"

Receive R4. Aggregate 192.0.2.0/24 AS_PATH "[64504 64501]"

If brief aggregation is in use, the AS_PATH is truncated to "".

The resulting AS_PATH is thus not stable and depends on the presence of specific routes.

B.3. Scenario 3: First one route, then another, and the AS_PATHs overlap at the neighbor AS

Receive R3. Aggregate 192.0.2.0/24 AS_PATH "64504 64501"

Receive R4. Aggregate 192.0.2.0/24 AS_PATH "64504 [64501 64502]"

If brief aggregation is in use, the AS_PATH is truncated to "64504".

The resulting AS_PATH is thus not stable and depends on the presence of specific routes.

B.4. Achieving Consistent Origin AS During Aggregation

In the three scenarios above, the aggregating AS 64500 is using traditional brief aggregation. This results in inconsistent origin ASes as the contributing routes are learned. This motivates the "consistent brief" BGP aggregation mentioned in Section 5.2 and discussed further with examples below.

The trivial solution to addressing the issue is to simply discard all of the ASes for the contributing routes. In simple BGP aggregation topologies, this is likely the correct thing to do. The AS originating the aggregate, 192.0.2.0/24 in this example, is likely the resource holder for the route in question. In such a case, simply originating the route to its BGP upstream neighbors in the Internet with its own AS, 64500, means that a consistent ROA could be registered in the RPKI for this prefix. This satisfies the need for a consistent (unambiguous) origin AS.

If the contributing ASes are themselves multihomed to the Internet outside of their connections to AS 64500, then additional ROAs would need to be created for each of the more specific prefixes.

In more complex proxy aggregation scenarios, there may be a desire to permit some stable (i.e., common) portion of the contributing AS_PATHs to be kept in the aggregate route. Consider the case for Scenario 3, where the neighbor AS is the same for both R3 and R4 -- AS 64504. In such a case, an implementation may permit the aggregate's brief AS_PATH to be "64504", and a ROA would be created for the aggregate prefix with 64504 as the origin AS.

Appendix C. Discussion on Forwarding Loops and AS_SETs

Although BGP-4 was designed to carry Classless Inter-Domain Routing (CIDR) routes, [RFC4271] does not discuss the installation of "discard" or "null" routes when implementing its aggregation procedures. Implementations could originate an aggregate prefix without a covering route for a more specific prefix (subsumed by the aggregate prefix) present in the local routing table.

When aggregating more specific routes according to the aggregation procedures of [RFC4271], the aggregating BGP speaker will place contributing routes into the generated AS_PATH, perhaps using AS_SETs. As a result, a contributing AS will not install the aggregated route into its RIB since the route is an AS_PATH loop. This provides a form of protection against forwarding loops created by BGP aggregation.

When brief aggregation methods are used, a BGP speaker may receive a route containing a less specific destination covering a local more specific destination and install it in its routing table since it is not prevented from doing so by BGP AS_PATH loop detection. This gives rise to the possibility of forwarding loops. To help prevent forwarding loops, it is critical to adhere to the following:

1. Rule #2 in Section 5.1 of [RFC4632]:

> A router that generates an aggregate route for multiple, more-specific routes must discard packets that match the aggregate route, but not any of the more-specific routes. In other words, the "next hop" for the aggregate route should be the null destination.

2. Not advertising aggregate routes to contributing ASes as specified in Section 6.2 of this document (also see Appendix A).

Acknowledgements

The authors would like to thank Alvaro Retana, John Scudder, Ketan Talaulikar, Keyur Patel, Susan Hares, Claudio Jeker, Nick Hilliard, Robert Raszuk, John Heasley, Job Snijders, Jared Mauch, Jakob Heitz, Tony Przygienda, Douglas Montgomery, Randy Bush, 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, Steven M. Bellovin, Steve Kent, Steve Padgett, and Alfred Hoenes for comments and suggestions. The comments and suggestions received from the IESG reviewers are also much appreciated.

Authors' Addresses

Warren Kumari
Google, Inc.
1600 Amphitheatre Parkway
Mountain View, CA 94043
United States of America
Phone: +1 571 748 4373
Email: warren@kumari.net

Kotikalapudi Sriram
USA NIST
100 Bureau Drive
Gaithersburg, MD 20899
United States of America
Phone: +1 301 975 3973
Email: ksriram@nist.gov

Lilia Hannachi
USA NIST
100 Bureau Drive
Gaithersburg, MD 20899
United States of America
Phone: +1 301 975 3259
Email: lilia.hannachi@nist.gov

Jeffrey Haas
Juniper Networks, Inc.
1133 Innovation Way
Sunnyvale, CA 94089
United States of America
Email: jhaas@juniper.net