

Internet Engineering Task Force (IETF)
Request for Comments: 6472
BCP: 172
Category: Best Current Practice
ISSN: 2070-1721

W. Kumari
Google, Inc.
K. Sriram
U.S. NIST
December 2011

Recommendation for Not Using AS_SET and AS_CONFED_SET in BGP

Abstract

This document recommends against the use of the AS_SET and AS_CONFED_SET types of the AS_PATH in BGPv4. This is done to simplify the design and implementation of BGP and to make the semantics of the originator of a route more clear. This will also simplify the design, implementation, and deployment of ongoing work in the Secure Inter-Domain Routing Working Group.

Status of This Memo

This memo documents an Internet Best Current Practice.

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 BCPS is available in Section 2 of RFC 5741.

Information about the current status of this document, any errata, and how to provide feedback on it may be obtained at <http://www.rfc-editor.org/info/rfc6472>.

Copyright Notice

Copyright (c) 2011 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 (<http://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 Simplified BSD License text as described in Section 4.e of the Trust Legal Provisions and are provided without warranty as described in the Simplified BSD License.

Table of Contents

1. Introduction	2
2. Requirements Notation	3
3. Recommendation to Network Operators	3
4. Security Considerations	4
5. Acknowledgements	4
6. References	4
6.1. Normative References	4
6.2. Informative References	4

1. Introduction

The AS_SET path segment type of 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 the update has traversed. The AS_CONFED_SET path type ([RFC5065]) of 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 the update has traversed. It is very similar to AS_SETs but is used within a confederation.

By performing aggregation, a router is, in essence, combining multiple existing routes into a single new route. This type of aggregation blurs the semantics of what it means to originate a route. Said aggregation can therefore 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 the "X.509 Extensions for IP Addresses and AS Identifiers" [RFC3779]). This in turn would result in reachability problems for the aggregated prefix and its components (i.e., more specifics). Said aggregation also creates traffic engineering issues, because the precise path information for the component prefixes is not preserved.

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 network [Analysis] and, when it is used, it is usually used incorrectly -- reserved AS numbers ([RFC1930]) and/or only a single AS in the AS_SET are by far the most common case. Because the aggregation involving AS_SETs is very rarely used, the reduction in table size provided by said aggregation is extremely small, and any advantage thereof is outweighed by additional complexity in BGP. As noted above, said aggregation also poses impediments to implementation of said new BGP security technologies.

In the past, AS_SET had been used in a few rare cases to allow route aggregation where two or more providers could form the same prefix, using the exact match of the other's prefix in some advertisement and configuring the aggregation differently elsewhere. The key to configuring this correctly was to form the aggregate at the border in the outbound BGP policy and omit prefixes from the AS that the aggregate was being advertised to. The AS_SET therefore allowed this practice without the loss of BGP's AS_PATH loop protection. This use of AS_SET served a purpose that fell in line with the original intended use. Without the use of AS_SET, aggregates must always contain only less specific prefixes (not less than or equal to), and must never aggregate an exact match.

2. Requirements Notation

The key words "MUST", "MUST NOT", "REQUIRED", "SHALL", "SHALL NOT", "SHOULD", "SHOULD NOT", "RECOMMENDED", "MAY", and "OPTIONAL" in this document are to be interpreted as described in [RFC2119].

3. Recommendation to Network Operators

It is RECOMMENDED that operators not generate any new announcements containing AS_SETs or AS_CONFED_SETs. If they have already announced routes with AS_SETs or AS_CONFED_SETs in them, then they SHOULD withdraw those routes and re-announce routes for the component prefixes (i.e., the additional specifics of the previously aggregated prefix) without AS_SETs in the updates. This involves undoing the aggregation that was previously performed (with AS_SETs), and announcing more specifics (without AS_SETs). Route aggregation that was previously performed by proxy aggregation (i.e., without the use of AS_SETs) is still possible under some conditions. As with any change, the operator should understand the full implications of the change.

It is worth noting that new technologies (such as those that take advantage of the "X.509 Extensions for IP Addresses and AS Identifiers" [RFC3779]) might not support routes with AS_SETs/AS_CONFED_SETs in them, and may treat as infeasible routes containing them. Future BGP implementations may also do the same. It is expected that, even before the deployment of these new or future technologies, operators may filter routes with AS_SETs/AS_CONFED_SETs in them. Other than making that observation, this document is not intended to make any recommendation for how an operator should behave when receiving a route with AS_SET or AS_CONFED_SET in it. This document's focus is entirely on the sender side, as discussed in the preceding paragraph.

4. Security Considerations

This document discourages the use of aggregation techniques that create AS_SETs. Future work may update the protocol to remove support for the AS_SET path segment type of the AS_PATH attribute. This future work will remove complexity and code that are not exercised very often, thereby decreasing the attack surface. This future work will also simplify the design and implementation of the Resource Public Key Infrastructure (RPKI) and systems that will rely on it.

5. Acknowledgements

The authors would like to thank Tony Li, Randy Bush, John Scudder, Curtis Villamizar, Danny McPherson, Chris Morrow, Tom Petch, and Ilya Varlashkin, as well as Douglas Montgomery, Enke Chen, Florian Weimer, Jakob Heitz, John Leslie, Keyur Patel, Paul Jakma, Rob Austein, Russ Housley, Sandra Murphy, Steve Bellovin, Steve Kent, Steve Padgett, Alfred Hoenes, Alvaro Retana, everyone in the IDR working group, and everyone else who provided input.

Apologies to those who we may have missed; it was not intentional.

6. References

6.1. Normative References

[RFC2119] Bradner, S., "Key words for use in RFCs to Indicate Requirement Levels", BCP 14, RFC 2119, March 1997.

6.2. Informative References

[Analysis] Sriram, K. and D. Montgomery, "Measurement Data on AS_SET and AGGREGATOR: Implications for {Prefix, Origin} Validation Algorithms", SIDR WG presentation, IETF 78, July 2010, <www.antd.nist.gov/~ksriram/AS_SET_Aggregator_Stats.pdf>.

[RFC1930] Hawkinson, J. and T. Bates, "Guidelines for creation, selection, and registration of an Autonomous System (AS)", BCP 6, RFC 1930, March 1996.

[RFC3779] Lynn, C., Kent, S., and K. Seo, "X.509 Extensions for IP Addresses and AS Identifiers", RFC 3779, June 2004.

- [RFC4271] Rekhter, Y., Ed., Li, T., Ed., and S. Hares, Ed., "A Border Gateway Protocol 4 (BGP-4)", RFC 4271, January 2006.
- [RFC5065] Traina, P., McPherson, D., and J. Scudder, "Autonomous System Confederations for BGP", RFC 5065, August 2007.

Authors' Addresses

Warren Kumari
Google, Inc.
1600 Amphitheatre Parkway
Mountain View, CA 94043
US

Phone: +1 571 748 4373
EMail: warren@kumari.net

Kotikalapudi Sriram
U.S. NIST
100 Bureau Drive
Gaithersburg, MD 20899
US

Phone: +1 301 975 3973
EMail: ksriram@nist.gov