       Recommendation for Not Using AS_SET and AS_CONFED_SET in BGP
                  draft-ietf-idr-deprecate-as-sets-06

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 the BGP protocol 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 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 http://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 April 9, 2012.

Table of Contents

1. Introduction

The AS_SET path segment type of the AS_PATH attribute ([RFC4271], Section 4.3) is created by a router that is performing route aggregation and contains an unordered set of ASs 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 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 the BGP protocol. 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 which fell in line with the original intended use. Without 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 more 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. IANA Considerations

This document requires no IANA actions.

5. 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 will remove complexity and code that is not exercised very
often, which decreases the attack surface. This will also simplify

the design and implementation of the RPKI and systems that will rely
on it.

6. Acknowledgements

The authors would like to thank Tony Li, Randy Bush, John Scudder,
Curtis Villamizar, Danny McPherson, Chris Morrow, Tom Petch, 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 Hones, Alvaro Retana, everyone in IDR and everyone else who
provided input

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

7. Informative References

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

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

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

[RFC4271] Rekhter, Y., Li, T., and S. Hares, "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.

[analysis]
           Sriram, K., "Measurement Data on AS_SET and AGGREGATOR,
           SIDR WG presentation, IETF-78", July 2010, < http://
           www.antd.nist.gov/~ksriram/AS_SET_Aggregator_Stats.pdf>.

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