

Deprecation of AS_SET and AS_CONFED_SET

<https://tools.ietf.org/html/draft-ietf-idr-deprecate-as-set-confed-set-02>

K. Sriram, Warren Kumari, Lilia Hannachi, Jeff Haas

IETF IDR WG Meeting

IETF 106

November 2019

Interest in deprecation of AS_SET and AS_CONFED_SET

- WG seems to have strong motivation to eliminate the use of these Attributes

Analysis of AS_SETs in BGP (IPv4)

Unique prefixes (with or without AS_SET) : 826535

Total # routes with AS_SETs : 477

routes with only one AS in AS_SET : 383

routes that are /24 prefix (aggregate) announcements : 239

Total # routes that seem meaningless or malformed : 456

Total # routes that seem meaningful : 21

Details: https://www.nist.gov/sites/default/files/documents/2019/10/23/detailed-as_set-analysis.txt

Analysis of AGGREGATOR, ATOMIC_AGGREGATE

*** When there is AGGREGATOR without AS_SET ***

Unique prefixes (with or without AS_SET) : 826535

Unique prefixes without AS_SET but with AGGREGATOR: 75698 (9.2%)

Unique prefixes with ATOMIC_AGGREGATE: 47258

Unique prefixes with AGGREGATOR and ATOMIC_AGGREGATE: 44971

Unique prefixes with AGGREGATOR and without ATOMIC_AGGREGATE: 31769

https://www.nist.gov/sites/default/files/documents/2019/10/23/detailed-as_set-analysis.txt

Source of some unusual aggregated AS_PATHs (jhaas)

41.196.34.0/23 701 174 8452 24863 {37069}

RFC 4271 compliant implementations of aggregation can yield an AS_SET of length one under the following conditions:

1. One or more contributing routes that are completely internal. (NULL AS_PATH.)
2. One or more contributing routes with the same single AS number.

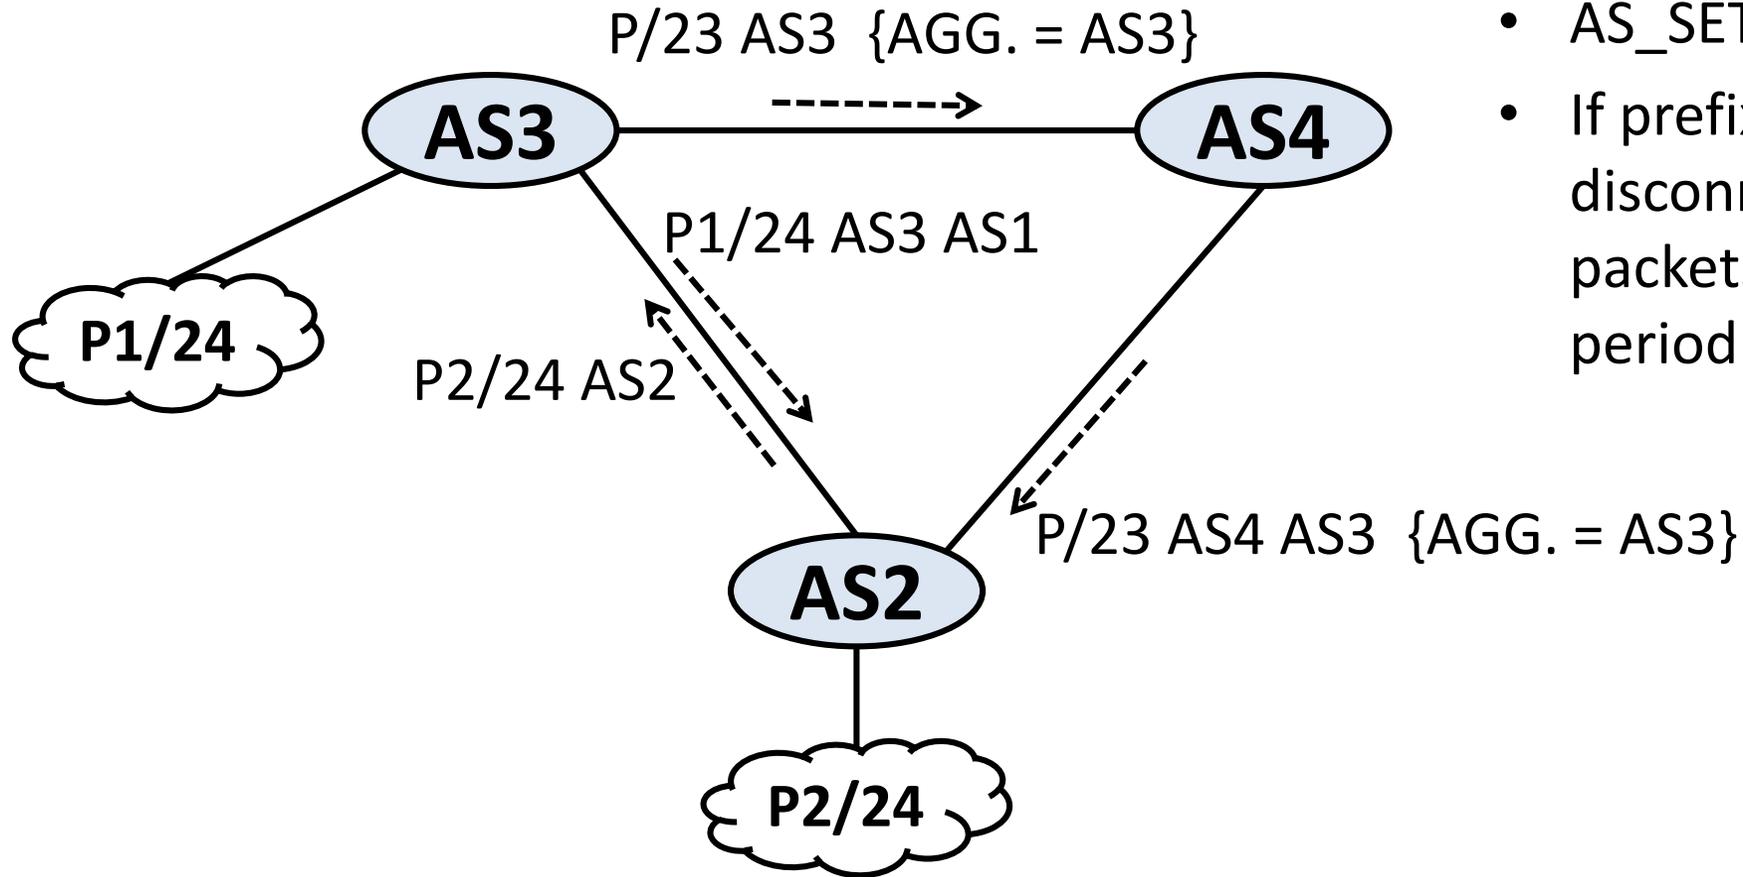
The longest common AS_PATH per the rules is NULL. Putting all additional ASes in a set yields an AS_SET of length one.

Source of some unusual aggregated AS_PATHs (jhaas) (2)

It is possible to alter the code for such cases of a single AS in the AS_SET.

- In such a case, it could be merged into the adjacent AS_SEQ.
- The path length is preserved.
- However, does this properly preserve the origin intent? This may be arguable. “brief” style aggregation discarding the AS_SET may be the right thing to do here.

Common Scenario: AGGREGATOR without AS_SET



- AS_SET not really required
- If prefix P2/24 gets disconnected, then data packets loop for a brief period

$$P/23 = P1/24 + P2/24$$

AGG. = AGGREGATOR

MUST / SHOULD Question

- Conformant BGP speakers MUST NOT locally generate BGP UPDATE messages containing *SET
- Upon receipt of messages with *SET, conformant BGP speakers SHOULD use the "Treat-as-withdraw" error handling behavior
 - SHOULD → MUST ?

Updating RFCs 4271, 5065, 6793 – Level of Detail?

- RFC 4271 has 26 mentions of AS_SET
- RFC 5065 has 11 mentions of AS_CONFED_SET
- RFC 6793 has 1 mention of AS_SET and 10 mentions of AS_CONFED_SET

Strategy for making necessary updates to these RFCs?

RFC 5065: Autonomous System Confederations for BGP

RFC 6793: BGP Support for Four-Octet Autonomous System (AS) Number Space

Alternative path for standardization (jhaas)

- Most (all?) implementations of BGP should be able to support “brief” style aggregation already. No new code need be deployed to change how aggregation works.
 - And no need to intrusively change several RFCs.
 - RFC 6472 already covers this requirement.
- Implementations should be asked to add a policy element that permits AS_SETS to be detected.
 - Having done so, it is possible to implement policy to discard routes having AS_SETs.
 - In the absence of operators cleaning up routes that have sets, RPKI filtering will eventually provide them “incentive” to clean up.

Alternative path for standardization (jhaas) (2)

- An Operational Considerations section for this document should be added that covers the issues with not using sets:
 - The aggregator must supply the more specific contributors to the contributing ASes.
 - The aggregator should not supply the aggregate route to the contributing ASes.
 - ASes that have reachability that is being aggregated should likely reject routes that contain their reachability to prevent forwarding loops.
 - Potentially enshrine the practice of internally advertising a discard route for the destination addresses belonging to one's subnet to prevent in-AS traffic from being sent off-AS. (However, see AS-bridging scenarios.)

Questions / Discussion