

BGP support for four-octet AS number space

[draft-ietf-idr-as4bytes-07.txt](#)

1. Status of this Memo

This document is an Internet-Draft and is in full conformance with all provisions of [Section 10 of RFC2026](#).

Internet-Drafts are working documents of the Internet Engineering Task Force (IETF), its areas, and its working groups. Note that other groups may also distribute working documents as Internet-Drafts.

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.''

The list of current Internet-Drafts can be accessed at <http://www.ietf.org/ietf/1id-abstracts.txt>

The list of Internet-Draft Shadow Directories can be accessed at <http://www.ietf.org/shadow.html>.

2. Abstract

Currently the Autonomous System number is encoded in BGP [[BGP](#)] as a two-octets field. This document describes extensions to BGP to carry the Autonomous System number as a four-octets field.

3. Protocol Extensions

For the purpose of this document let's define a BGP speaker which doesn't support the new 4-octets AS number extensions as an OLD BGP speaker, and a BGP speaker which supports the new 4-octets AS number extensions as a NEW BGP speaker.

BGP carries the Autonomous System number in the My Autonomous System field of the OPEN message, in the AS_PATH attribute of the UPDATE message, and in the AGGREGATOR attribute of the UPDATE message. BGP also carries the Autonomous System number in the BGP Communities attribute.

A NEW BGP speaker uses BGP Capability Advertisements [[RFC2842](#)] to advertise to its neighbors (either internal or external) that it supports 4-octets AS number extensions, as specified in this document.

The Capability that is used by a BGP speaker to convey to its BGP peer the 4-octets Autonomous System number capability, also carries the 4-octets Autonomous System number of the speaker in the Capability Value field of the Capability Optional Parameter. The Capability Length field of the Capability is set to 4.

NEW BGP speakers carry AS path information expressed in terms of 4-octets Autonomous Systems numbers by using the existing AS_PATH attribute, except that each AS number in this attribute is encoded not as a 2-octets, but as a 4-octets entity. The same applies to the AGGREGATOR attribute - NEW BGP speakers use the same attribute, except that the AS carried in this attribute is encoded as a 4-octets entity.

To preserve AS path information with 4-octets AS numbers across OLD BGP speakers, this document defines a new AS path attribute, called NEW_AS_PATH. This attribute is optional transitive and contains the AS path encoded with 4-octets AS numbers. The NEW_AS_PATH attribute has the same semantics as the AS_PATH attribute, except that it is optional transitive, and it carries 4-octets AS numbers.

To prevent the possible propagation of confederation path segments outside of a confederation, the path segment types AS_CONFED_SEQUENCE and AS_CONFED_SET [[RFC3065](#)] are declared invalid for the NEW_AS_PATH attribute.

Similarly, this document defines a new aggregator attribute called NEW_AGGREGATOR, which is optional transitive. The NEW_AGGREGATOR attribute has the same semantics as the AGGREGATOR attribute, except that it carries 4-octets AS numbers.

Currently assigned 2-octets Autonomous System numbers are converted into 4-octets Autonomous System numbers by setting the high-order 2 octets of the 4-octets field to zero. Such a 4-octets AS number is said to be mappable to a 2-octets AS number.

To represent 4-octets AS numbers (which are not mapped from 2-octets) as 2-octets AS numbers in the AS path information encoded with 2-octets AS numbers, this document reserves a 2-octets AS number. Lets denote this special AS number as AS_TRANS for ease of description in the rest of this specification. This AS number is also placed in the My Autonomous System field of the OPEN message originated by a NEW BGP speaker if the speaker doesn't have a (globally unique) 2-octets AS number.

4. Operations

4.1. Interaction between NEW BGP speakers

A BGP speaker that supports 4-octets Autonomous System numbers may advertise this to its peers using the BGP Capability Advertisements. A BGP speaker that advertises such capability to a particular peer, and receives from that peer the advertisement of such capability **MUST** encode Autonomous System numbers as 4-octets entities in both the AS_PATH and the AGGREGATOR attributes in the updates it sends to the peer, and **MUST** assume that these attributes in the updates received from the peer encode Autonomous System numbers as 4-octets entities.

The new attributes, NEW_AS_PATH and NEW_AGGREGATOR should not be carried in the UPDATE messages between NEW BGP peers. A NEW BGP speaker that receives the NEW_AS_PATH and NEW_AGGREGATOR path attributes in an UPDATE message from a NEW BGP speaker should discard these path attributes and continue processing the UPDATE message.

4.2. Interaction between NEW and OLD BGP speaker

4.2.1. BGP Peering

Note that peering between a NEW BGP speaker and an OLD one is possible only if the NEW BGP speaker has a 2-octets AS number. However, this document doesn't assume that an Autonomous System with NEW speakers has to have a globally unique 2-octets AS number - AS_TRANS could be used instead (even if multiple Autonomous System would use it).

4.2.2. Generating Updates

When communicating with an OLD BGP speaker, a NEW speaker MUST send the AS path information in the AS_PATH attribute encoded with 2-octets AS numbers. The NEW speaker also MUST send the AS path information in the NEW_AS_PATH attribute (encoded with 4-octets AS numbers), except for the case where the entire AS path information is composed of 2-octets AS numbers only. In this case the NEW speaker should not send the NEW_AS_PATH attribute.

In the AS_PATH attribute encoded with 2-octets AS numbers, non-mappable 4-octets AS numbers are represented by the well known 2-octets AS number, AS_TRANS. This will preserve the path length property of the AS path information; and will also help in updating the AS path information received on a NEW BGP speaker from an OLD speaker, as explained in the next section.

The NEW speaker constructs the NEW_AS_PATH attribute from the information carried in the AS_PATH attribute. In the case where the AS_PATH attribute contains either AS_CONFED_SEQUENCE or AS_CONFED_SET path segments, the NEW speaker, when constructing the NEW_AS_PATH attribute from the AS_PATH attribute, must exclude such path segments. The NEW_AS_PATH attribute will be carried across a series of OLD BGP speakers without modification and will help preserve the truly 4-octets AS numbers in the AS path information.

Similarly, if the NEW speaker has to send the AGGREGATOR attribute, and if the aggregating Autonomous System's AS number is truly 4-octets, the speaker constructs the NEW_AGGREGATOR attributes by taking the attribute length and attribute value from the AGGREGATOR attribute and placing them into the attribute length and attribute value of the NEW_AGGREGATOR attribute, and sets the AS number field in the existing AGGREGATOR attribute to the reserved AS number, AS_TRANS. Note that if the AS number is 2-octets only, then the NEW_AGGREGATE attribute should not be sent.

4.2.3. Processing Received Updates

When a NEW BGP speaker receives an update from an OLD one, it should be prepared to receive the NEW_AS_PATH attribute along with the existing AS_PATH attribute. If NEW_AS_PATH attribute is also received, both the attributes will be used to construct the exact AS path information, and therefore the information carried by both the attributes will be considered for AS path loop detection.

Note that a route may have traversed a series of autonomous systems with 2-octets AS numbers and OLD BGP speakers only. In that case, if

the route carries a NEW_AS_PATH attribute, this attribute may not have been updated since the route left the last NEW BGP speaker. The trailing AS path information (representing autonomous systems with 2-octets AS numbers and OLD BGP speakers only) is contained only in the current AS_PATH attribute (encoded in the leading part of the AS_PATH attribute). This AS path information should be prepended to the NEW_AS_PATH attribute to construct the exact AS path information.

Similarly, a NEW BGP speaker should be prepared to receive the NEW_AGGREGATOR attribute from an OLD BGP speaker. In that case, the AGGREGATOR attribute is ignored and the NEW_AGGREGATOR contains the exact information about the aggregating node.

4.3. Interaction between OLD BGP speakers

In all other cases the speaker MUST encode Autonomous System numbers as 2-octets entities in both the AS_PATH and the AGGREGATOR attribute in the updates it sends to the peer, and MUST assume that these attributes in the updates received from the peer encoded Autonomous System numbers as 2-octets entities.

5. Handling BGP Communities

As specified in [[RFC1997](#)], when the high-order two-octets of the community attribute is neither 0x0000 nor 0xffff, these two octets encode the Autonomous System number. Quite clearly this would not work for routers that use 4-octets Autonomous System numbers. Such routers should use the Extended Communities [EXT-COMM] attribute instead.

6. Transition

The scheme described in this document allows a gradual transition from 2-octets AS numbers to 4-octets AS numbers. One can upgrade one Autonomous system or one router at a time.

To simplify transition this document assumes that an Autonomous System could start using 4-octets AS number only after all the BGP speakers within that Autonomous System have been upgraded to support 4-octets AS numbers.

An OLD BGP speaker should not use AS_TRANS as its Autonomous System number.

A non-mappable 4-octets AS number can not be used as a "Member AS

Number" of a BGP Confederation until all the BGP speakers within the Confederation transitioned to support 4-octets AS numbers.

In an environment where an Autonomous System that has OLD BGP speakers peers with two or more Autonomous Systems that have NEW BGP speakers and use AS_TRANS (rather than having a globally unique AS number), use of Multi-Exit Discriminators by the Autonomous System with the OLD speakers may result in a situation where Multi-Exit Discriminator will influence route selection among the routes that were received from different neighboring Autonomous Systems.

Under certain conditions it may not be possible to reconstruct the entire AS path information from the AS_PATH and the NEW_AS_PATH attributes of a route. This occurs when two or more routes that carry the NEW_AS_PATH attribute are aggregated by an OLD BGP speaker, and the NEW_AS_PATH attribute of at least one of these routes carries at least one 4-octets AS number (as oppose to a 2-octets AS number that is encoded in 4 octets). When such aggregation results in creating a route that is less specific than any of the component routes, (route whose NLRI covers NLRI of all the component routes), loss of the AS path information does not create a risk of a routing loop. In all other cases loss of the AS path information does create a risk of a routing loop.

7. IANA Consideration

This document uses a BGP Capability code to indicate that a BGP speaker supports the 4-octets AS numbers. The Capability code has been assigned by IANA per [RFC 2842](#).

In addition, this document introduces two new BGP optional transitive attributes. The first is the NEW_AS_PATH attribute, which preserves the AS path information with 4-octet AS numbers across old BGP speakers. The second is the NEW_AGGREGATOR attribute, which is similar in use to the current AGGREGATOR attribute but it carries 4-octet AS numbers. The Type Codes for these attributes has been assigned by IANA.

Finally, this document introduces a reserved 2-octets AS number - AS_TRANS. The AS number for AS_TRANS has been assigned by the IANA.

8. Security Considerations

Security issues are not discussed in this document.

9. Acknowledgments

The authors would like to thank Yakov Rekhter, Chaitanya Kodeboyina, and Jeffrey Haas for the numerous discussions which went into the making of this draft.

10. References

[BGP] Rekhter, Y., Li, T., "Border Gateway Protocol 4", [draft-ietf-idr-bgp4-12.txt](#)

[EXT-COM] Ramachandra, S., Tappan, D., and Rekter Y. "BGP Extended Communities Attribute", [draft-ramachandra-bgp-ext-communities-02.txt](#)

[RFC1997] Chandra, R., Traina, P. and T. Li, "BGP Communities Attribute", [RFC 1997](#), August 1996.

[RFC2842] Chandra, R., and Scudder, J., "Capabilities Advertisement with BGP-4", [RFC 2842](#), May 2000.

[RFC3065] Traina, P., McPherson, D., Scudder, J., "Autonomous System Confederations for BGP", [RFC3065](#), February 2001.

11. Author Information

Quaizar Vohra
Juniper Networks
[1194 N.Mathilda Ave](#)
Sunnyvale, CA 94089
e-mail: qv@juniper.net

Enke Chen
Redback Networks, Inc.
[350 Holger Way](#)
San Jose, CA 95134
e-mail: enke@redback.com

