

Internet Engineering Task Force (IETF)
Request for Comments: 6793
Obsoletes: 4893
Updates: 4271
Category: Standards Track
ISSN: 2070-1721

Q. Vohra
Juniper Networks
E. Chen
Cisco Systems
December 2012

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

Abstract

The Autonomous System number is encoded as a two-octet entity in the base BGP specification. This document describes extensions to BGP to carry the Autonomous System numbers as four-octet entities. This document obsoletes RFC 4893 and updates RFC 4271.

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 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/rfc6793>.

Copyright Notice

Copyright (c) 2012 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.

1. Introduction

In the base BGP specification [RFC4271], the Autonomous System (AS) number is encoded as a two-octet entity. To prepare for the anticipated exhaustion of the two-octet AS numbers, this document describes extensions to BGP to carry the AS numbers as four-octet entities.

More specifically, this document defines a BGP capability code, "support for four-octet AS number capability", to be used by a BGP speaker to indicate its support for four-octet AS numbers. Two attributes, AS4_PATH and AS4_AGGREGATOR, are introduced that can be used to propagate four-octet-based AS path information across BGP speakers that do not support four-octet AS numbers. This document also specifies mechanisms for constructing the AS path information from the AS_PATH attribute and the AS4_PATH attribute.

The extensions specified in this document allow a gradual transition from two-octet AS numbers to four-octet AS numbers.

This document obsoletes RFC 4893 and updates RFC 4271. It includes several clarifications and editorial changes, and it specifies the error handling for the new attributes.

2. Specification of Requirements

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. Protocol Extensions

For the purpose of this document, we define a BGP speaker that does not support the new four-octet AS number extensions as an OLD BGP speaker, and a BGP speaker that supports the new four-octet AS number extensions as a NEW BGP speaker.

BGP carries the AS numbers 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 AS numbers in the BGP Communities attribute.

A NEW BGP speaker uses BGP Capabilities Advertisements [RFC5492] to advertise to its neighbors (either internal or external) that it supports four-octet AS number extensions, as specified in this document.

The capability that is used by a BGP speaker to convey to its BGP peer the four-octet Autonomous System number capability also carries the AS number (encoded as a four-octet entity) of the speaker in the Capability Value field of the capability. The Capability Length field of the capability is set to 4.

The AS path information exchanged between NEW BGP speakers is carried in the existing AS_PATH attribute, except that each AS number in the attribute is encoded as a four-octet entity (instead of a two-octet entity). The same applies to the AGGREGATOR attribute -- the same attribute is used between NEW BGP speakers, except that the AS number carried in the attribute is encoded as a four-octet entity.

The AS_PATH attribute and the AGGREGATOR attribute carried between a NEW BGP speaker and an OLD BGP speaker will continue to contain two-octet AS numbers.

To preserve the AS path information with four-octet AS numbers across OLD BGP speakers, this document defines a new BGP path attribute called AS4_PATH. This is an optional transitive attribute that contains the AS path encoded with four-octet AS numbers. The AS4_PATH attribute has the same semantics and the same encoding as the AS_PATH attribute, except that it is "optional transitive", and it carries four-octet AS numbers.

To prevent the possible propagation of Confederation-related path segments outside of a Confederation, the path segment types AS_CONFED_SEQUENCE and AS_CONFED_SET [RFC5065] are declared invalid for the AS4_PATH attribute and MUST NOT be included in the AS4_PATH attribute of an UPDATE message.

Similarly, this document defines a new BGP path attribute called AS4_AGGREGATOR, which is optional transitive. The AS4_AGGREGATOR attribute has the same semantics and the same encoding as the AGGREGATOR attribute, except that it carries a four-octet AS number.

Currently assigned two-octet AS numbers are converted into four-octet AS numbers by setting the two high-order octets of the four-octet field to zero. Such a four-octet AS number is said to be mappable to a two-octet AS number.

This document reserves a two-octet AS number called "AS_TRANS". AS_TRANS can be used to represent non-mappable four-octet AS numbers as two-octet AS numbers in AS path information that is encoded with two-octet AS numbers. (In this context, four-octet AS numbers that are not mapped from two-octet AS numbers are referred to as "non-mappable".) We 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 and only if the speaker does not have a (globally unique) two-octet AS number.

4. Operations

4.1. Interaction between NEW BGP Speakers

A BGP speaker that supports four-octet AS numbers SHALL advertise this to its peers using BGP Capabilities Advertisements. The AS number of the BGP speaker MUST be carried in the Capability Value field of the "support for four-octet AS number capability".

When a NEW BGP speaker processes an OPEN message from another NEW BGP speaker, it MUST use the AS number encoded in the Capability Value field of the "support for four-octet AS number capability" in lieu of the "My Autonomous System" field of the OPEN message.

A BGP speaker that advertises such a capability to a particular peer, and receives from that peer the advertisement of such a capability, MUST encode AS numbers as four-octet entities in both the AS_PATH attribute 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 encode AS numbers as four-octet entities.

The new attributes, AS4_PATH and AS4_AGGREGATOR, MUST NOT be carried in an UPDATE message between NEW BGP speakers. A NEW BGP speaker that receives the AS4_PATH attribute or the AS4_AGGREGATOR attribute in an UPDATE message from another NEW BGP speaker MUST discard the path attribute and continue processing the UPDATE message.

4.2. Interaction between NEW and OLD BGP Speakers

4.2.1. BGP Peering

Note that peering between a NEW BGP speaker and an OLD BGP speaker is possible only if the NEW BGP speaker has a two-octet AS number. However, this document does not assume that an Autonomous System with NEW BGP speakers has to have a globally unique two-octet AS number -- AS_TRANS MUST be used when the NEW BGP speaker does not have a two-octet AS number (even if multiple Autonomous Systems would use it).

4.2.2. Generating Updates

When communicating with an OLD BGP speaker, a NEW BGP speaker MUST send the AS path information in the AS_PATH attribute encoded with two-octet AS numbers. The NEW BGP speaker MUST also send the AS path information in the AS4_PATH attribute (encoded with four-octet AS numbers), except for the case where all of the AS path information is composed of mappable four-octet AS numbers only. In this case, the NEW BGP speaker MUST NOT send the AS4_PATH attribute.

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

The NEW BGP speaker constructs the AS4_PATH attribute from the AS path information. Whenever the AS path information contains the AS_CONFED_SEQUENCE or AS_CONFED_SET path segment, the NEW BGP speaker MUST exclude such path segments from the AS4_PATH attribute being constructed.

The AS4_PATH attribute, being optional transitive, will be carried across a series of OLD BGP speakers without modification and will help preserve the non-mappable four-octet AS numbers in the AS path information.

Similarly, if the NEW BGP speaker has to send the AGGREGATOR attribute, and if the aggregating Autonomous System's AS number is a non-mappable four-octet AS number, then the speaker MUST use the AS4_AGGREGATOR attribute and set the AS number field in the existing AGGREGATOR attribute to the reserved AS number, AS_TRANS. Note that if the AS number is mappable, then the AS4_AGGREGATOR attribute MUST NOT be sent.

4.2.3. Processing Received Updates

When a NEW BGP speaker receives an update from an OLD BGP speaker, it MUST be prepared to receive the AS4_PATH attribute along with the existing AS_PATH attribute. If the AS4_PATH attribute is also received, both of the attributes will be used to construct the exact AS path information, and therefore the information carried by both of the attributes will be considered for AS path loop detection.

Note that a route may have traversed a series of Autonomous Systems with two-octet AS numbers and OLD BGP speakers only. In that case, if the route carries the AS4_PATH attribute, this attribute would have remained unmodified since the route left the last NEW BGP speaker. The trailing AS path information (representing Autonomous Systems with two-octet 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).

Under certain conditions, it may not be possible to reconstruct all of the AS path information from the AS_PATH and the AS4_PATH attributes of a route. This occurs, for example, when two or more routes that carry the AS4_PATH attribute are aggregated by an OLD BGP speaker, and the AS4_PATH attribute of at least one of these routes carries at least one four-octet AS number (as opposed to a two-octet AS number that is encoded in 4 octets). Depending on the implementation, either the AS4_PATH attribute would be lost during route aggregation, or both the AS_PATH attribute and the AS4_PATH attribute would contain valid, partial information that cannot be combined seamlessly, resulting in incomplete AS path information in these cases.

A NEW BGP speaker MUST also be prepared to receive the AS4_AGGREGATOR attribute along with the AGGREGATOR attribute from an OLD BGP speaker. When both of the attributes are received, if the AS number in the AGGREGATOR attribute is not AS_TRANS, then:

- the AS4_AGGREGATOR attribute and the AS4_PATH attribute SHALL be ignored,
- the AGGREGATOR attribute SHALL be taken as the information about the aggregating node, and
- the AS_PATH attribute SHALL be taken as the AS path information.

Otherwise,

- the AGGREGATOR attribute SHALL be ignored,
- the AS4_AGGREGATOR attribute SHALL be taken as the information about the aggregating node, and
- the AS path information would need to be constructed, as in all other cases.

In order to construct the AS path information, it is necessary to first calculate the number of AS numbers in the AS_PATH and AS4_PATH attributes using the method specified in Section 9.1.2.2 of [RFC4271] and in [RFC5065] for route selection.

If the number of AS numbers in the AS_PATH attribute is less than the number of AS numbers in the AS4_PATH attribute, then the AS4_PATH attribute SHALL be ignored, and the AS_PATH attribute SHALL be taken as the AS path information.

If the number of AS numbers in the AS_PATH attribute is larger than or equal to the number of AS numbers in the AS4_PATH attribute, then the AS path information SHALL be constructed by taking as many AS numbers and path segments as necessary from the leading part of the AS_PATH attribute, and then prepending them to the AS4_PATH attribute so that the AS path information has a number of AS numbers identical to that of the AS_PATH attribute. Note that a valid AS_CONFED_SEQUENCE or AS_CONFED_SET path segment SHALL be prepended if it is either the leading path segment or is adjacent to a path segment that is prepended.

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 AS number. Quite clearly, this would not work for a NEW BGP speaker with a non-mappable four-octet AS number. Such BGP speakers should use four-octet AS specific extended communities [RFC5668] instead.

6. Error Handling

This section provides an update to RFC 4271 [RFC4271] with respect to the error conditions noted here and their handling.

Given that the two-octet AS numbers dominate during the transition and are carried in the AS_PATH attribute by an OLD BGP speaker, in this document the "attribute discard" approach is chosen to handle a malformed AS4_PATH attribute.

Similarly, as the AS4_AGGREGATOR is just informational, the "attribute discard" approach is chosen to handle a malformed AS4_AGGREGATOR attribute.

The AS4_PATH attribute and AS4_AGGREGATOR attribute MUST NOT be carried in an UPDATE message between NEW BGP speakers. A NEW BGP speaker that receives the AS4_PATH attribute or the AS4_AGGREGATOR attribute in an UPDATE message from another NEW BGP speaker MUST discard the path attribute and continue processing the UPDATE message. This case SHOULD be logged locally for analysis.

In addition, the path segment types AS_CONFED_SEQUENCE and AS_CONFED_SET [RFC5065] MUST NOT be carried in the AS4_PATH attribute of an UPDATE message. A NEW BGP speaker that receives these path segment types in the AS4_PATH attribute of an UPDATE message from an OLD BGP speaker MUST discard these path segments, adjust the relevant attribute fields accordingly, and continue processing the UPDATE message. This case SHOULD be logged locally for analysis.

The AS4_PATH attribute in an UPDATE message SHALL be considered malformed under the following conditions:

- the attribute length is not a multiple of two or is too small (i.e., less than 6) for the attribute to carry at least one AS number, or
- the path segment length in the attribute is either zero or is inconsistent with the attribute length, or
- the path segment type in the attribute is not one of the types defined: AS_SEQUENCE, AS_SET, AS_CONFED_SEQUENCE, and AS_CONFED_SET.

A NEW BGP speaker that receives a malformed AS4_PATH attribute in an UPDATE message from an OLD BGP speaker MUST discard the attribute and continue processing the UPDATE message. The error SHOULD be logged locally for analysis.

The AS4_AGGREGATOR attribute in an UPDATE message SHALL be considered malformed if the attribute length is not 8.

A NEW BGP speaker that receives a malformed AS4_AGGREGATOR attribute in an UPDATE message from an OLD BGP speaker MUST discard the attribute and continue processing the UPDATE message. The error SHOULD be logged locally for analysis.

7. Transition

When an Autonomous System is using a two-octet AS number, then the BGP speakers within that Autonomous System MAY be upgraded to support the four-octet AS number extensions on a piecemeal basis. There is no requirement for a coordinated upgrade of the four-octet AS number capability in this case. However, if an Autonomous System wishes to use a four-octet AS number as its own AS number, then this document assumes that an Autonomous System can use a four-octet AS number only after all the BGP speakers within that Autonomous System have been upgraded to support four-octet AS numbers.

A non-mappable four-octet AS number cannot be used as a "Member AS Number" of a BGP Confederation until all the BGP speakers within the Confederation have transitioned to support four-octet 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 mappable AS number), the use of the MULTI_EXIT_DISC attribute [RFC4271] by the Autonomous System with the OLD BGP speakers may result in a situation where the MULTI_EXIT_DISC attribute 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 all of the AS path information from the AS_PATH and the AS4_PATH attributes of a route. This occurs when two or more routes that carry the AS4_PATH attribute are aggregated by an OLD BGP speaker, and the AS4_PATH attribute of at least one of these routes carries at least one four-octet AS number (as opposed to a two-octet 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 (routes whose Network Layer Reachability Information (NLRI) covers the NLRI of all the component routes), loss of the AS path information does not create the risk of a routing loop. In all other cases, loss of the AS path information does create the risk of a routing loop.

8. Manageability Considerations

If the BGP4-MIB [RFC4273] is supported, there are no additional manageability concerns that arise from the use of four-octet AS numbers, since the InetAutonomousSystemNumber textual convention [RFC4001] is defined as Unsigned32.

When IP Flow Information Export (IPFIX) [RFC5101] is supported, there are no additional manageability concerns that arise from the use of four-octet AS numbers. The `bgpSourceAsNumber` and `bgpDestinationAsNumber` information elements [IANA-IPFIX] can continue to be used, with a new template record, specifying the new length of 4 bytes.

9. IANA Considerations

This document expands the pool for AS numbers from 0-65535 to 0-4294967295. The AS numbers are managed by the IANA "Autonomous System Numbers" registry. Other than expanding the AS number pool, this document does not propose any modifications to the existing policies and procedures pertaining to the allocation of AS numbers.

This document uses a BGP capability code to indicate that a BGP speaker supports four-octet AS numbers. Capability Code 65 has been assigned by IANA per [RFC5492].

In addition, this document introduces two BGP optional transitive attributes, and their type codes have been assigned by IANA. The first one is the `AS4_PATH` attribute, value 17, which preserves the AS path information with four-octet AS numbers across old BGP speakers. The second one is the `AS4_AGGREGATOR` attribute, value 18, which is similar in use to the current `AGGREGATOR` attribute, but it carries a four-octet AS number.

Finally, IANA has replaced a reference to RFC 4893 with a reference to this document for a reserved two-octet AS number -- `AS_TRANS` (23456). Also, IANA has replaced a reference to RFC 4893 with a reference to this document for the "32-bit Autonomous System Numbers" registry.

10. Security Considerations

This extension to BGP does not change the underlying security issues inherent in the existing BGP, except for the following:

The inconsistency between the `AS_PATH` attribute and the `AS4_PATH` attribute can create loss of the AS path information, and potential routing loops in certain cases, as discussed in this document. This could be exploited by an attacker.

It is a misconfiguration to assign a non-mappable four-octet AS number as the "Member AS Number" in a BGP Confederation before all the BGP speakers within the Confederation have transitioned to support four-octet AS numbers. Such a misconfiguration would weaken AS path loop detection within a Confederation.

11. Acknowledgments

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

The authors would also like to thank members of the IDR Working Group for their review and comments.

12. References

12.1. Normative References

- [RFC1997] Chandra, R., Traina, P., and T. Li, "BGP Communities Attribute", RFC 1997, August 1996.
- [RFC2119] Bradner, S., "Key words for use in RFCs to Indicate Requirement Levels", BCP 14, RFC 2119, March 1997.
- [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.
- [RFC5492] Scudder, J. and R. Chandra, "Capabilities Advertisement with BGP-4", RFC 5492, February 2009.
- [RFC5668] Rekhter, Y., Sangli, S., and D. Tappan, "4-Octet AS Specific BGP Extended Community", RFC 5668, October 2009.

12.2. Informative References

- [IANA-IPFIX] IANA, "IP Flow Information Export (IPFIX) Entities", <<http://www.iana.org/assignments/ipfix>>.
- [RFC4001] Daniele, M., Haberman, B., Routhier, S., and J. Schoenwaelder, "Textual Conventions for Internet Network Addresses", RFC 4001, February 2005.
- [RFC4273] Haas, J., Ed., and S. Hares, Ed., "Definitions of Managed Objects for BGP-4", RFC 4273, January 2006.
- [RFC5101] Claise, B., Ed., "Specification of the IP Flow Information Export (IPFIX) Protocol for the Exchange of IP Traffic Flow Information", RFC 5101, January 2008.

Authors' Addresses

Quaizar Vohra
Juniper Networks
1194 N. Mathilda Ave.
Sunnyvale, CA 94089
USA

E-Mail: quaizar.vohra@gmail.com

Enke Chen
Cisco Systems, Inc.
170 W. Tasman Dr.
San Jose, CA 95134
USA

E-Mail: enkechen@cisco.com