

Internet Engineering Task Force
Internet-Draft
Intended status: Informational
Expires: March 28, 2013

W. George
Time Warner Cable
S. Amante
Level 3 Communications
September 24, 2012

**Autonomous System (AS) Migration Features and Their Effects on the BGP
AS_PATH Attribute
draft-ga-idr-as-migration-00**

Abstract

This draft discusses common methods of managing an ASN migration using some BGP features that while commonly-used are not formally part of the BGP4 protocol specification and may be vendor-specific in exact implementation. It is necessary to document these de facto standards to ensure that they are properly supported in BGPSec.

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 March 28, 2013.

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.

Table of Contents

1.	Introduction	3
1.1.	Requirements Language	4
2.	ASN Migration Scenario Overview	4
3.	Local AS: Modify Inbound BGP AS_PATH Attribute	5
4.	Replace AS: Modify Outbound BGP AS_PATH Attribute	7
5.	Additional Operational Considerations	8
6.	Conclusion	8
7.	Acknowledgements	9
8.	IANA Considerations	9
9.	Security Considerations	9
10.	References	9
10.1.	Normative References	9
10.2.	Informative References	9
	Authors' Addresses	10

1. Introduction

This draft discusses common methods of managing an ASN migration using some BGP features that while commonly-used are not formally part of the BGP4 [\[RFC4271\]](#) protocol specification and may be vendor-specific in exact implementation. This draft does not attempt to standardize these features, because they are local to given implementation and do not require negotiation with or cooperation of BGP neighbors. The deployment of these features do not need to interwork with one another to accomplish the desired results. However, it is necessary to document these de facto standards to ensure that any future protocol enhancements to BGP that propose to read, copy, manipulate or compare the AS_PATH attribute can do so without inhibiting the use of these very widely used ASN migration features.

It is important to understand the business need for these features, as well, to illustrate why they are critical, particularly for ISP's operations. (It should be noted that these features are not limited to ISP's and that organizations of all sizes use these features for similar reasons to ISP's). During a merger, acquisition or divestiture involving two organizations it is necessary to seamlessly migrate BGP speakers from one ASN to a second ASN. The overall goal in doing so, particularly in the case of a merger or acquisition, is to achieve a uniform operational model through consistent configurations across all BGP speakers in the combined network. In addition, and perhaps more importantly, it is common practice in the industry for ISPs to bill customers based on utilization. ISPs bill customers based on the 95th percentile of the greater of the traffic sent or received, over the course of a 1-month period, on the customer's PE-CE access circuit. Given that the BGP Path Selection algorithm selects routes with the shortest AS_PATH attribute, it is critical for the ISP to not increase AS_PATH length during or after ASN migration, toward both downstream transit customers as well as settlement-free peers, who are likely sending or receiving traffic from those transit customers. This would not only result in sudden changes in traffic patterns in the network, but also (substantially) decrease utilization driven revenue at the ISP. It can also be bad for the end customers, as they may use AS prepends as a way to manage the AS_PATH to ensure optimal traffic engineering from their multiple upstream SPs, and a sudden change in the AS_PATH may cause an undesirable shift in traffic.

Lastly, it is important to note that, by default, the BGP protocol requires an operator to configure a single remote ASN for the eBGP neighbor inside a router, in order to successfully negotiate and establish an eBGP session. Prior to the existence of these features, it would have required an ISP to work with, in some cases, tens of

thousands of customers. In particular, the ISP would have to encourage those customers to change their CE router configs to use the new ASN, in a very short period of time, when the customer has no business incentive to do so. Thus, it became critical to allow the ISP to seamlessly migrate the ASN within its network(s), not disturb existing customers and allow the customer's to gradually migrate to the ISP's new ASN at their leisure.

1.1. Requirements Language

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 [RFC 2119](#) [[RFC2119](#)].

2. ASN Migration Scenario Overview

The use case being discussed here is an ISP merging two or more ASNs, where eventually one ASN subsumes the other(s). In this use case, we will assume the most common case where there are two ISPs, A and B, that use AS 200 and 300, respectively, before the ASN migration is to occur. AS 200 will be the "keep" ASN used across the consolidated set of both ISPs' network equipment and AS 300 will be retired. Thus, at the conclusion of the ASN migration, there will be a single ISP A' with all internal BGP speakers configured to use AS 200. To all external BGP speakers, the AS_PATH length will not be increased.

In this same scenario, AS 100 and AS 400 represent two, separate customer networks: C and D, respectively. Originally, customer C (AS 100) is attached to ISP B, which will undergo ASN migration from AS 300 to AS 200. Furthermore, customer D (AS 100) is attached to ISP A, which does not undergo ASN migration since ISP A's ASN will remain constant, (AS 200). Although this example refers to AS 100 and 400 as customer networks, either or both may be settlement-free or other types of peers. In this use case they are referred to as "customers" merely for convenience.

The general order of operations, typically carried out in a single maintenance window by the network undergoing ASN migration, ISP B, are as follows. First, ISP B, will change the global BGP ASN used by a PE router, from AS 300 to 200. At this point, the router will no longer be able to establish eBGP sessions toward the existing CE devices that are attached to it and still using AS 300. Second, ISP B will configure two separate, but related ASN migration features discussed in this document on all eBGP sessions toward all CE devices. These features modify the AS_PATH attribute received from and transmitted toward CE devices to achieve the desired effect of not increasing the length of the AS_PATH.

At the conclusion of the ASN migration, The subset of eBGP speakers, AS 400, that were formerly configured to establish a BGP session with ISP B's legacy AS 300 will still receive AS 300 in the AS_PATH received from ISP A'. To be clear, that same subset of CE devices, in AS 400, will not see the new AS 200 in the AS_PATH, until the time those CE devices have changed their configuration to establish an eBGP session to the new AS, AS 200.

In this use case, neither ISP is using BGP Confederations [RFC 5065](#) [[RFC5065](#)] internally.

Additional information about this scenario, including vendor-specific implementation details can be found here: Cisco [[CISCO](#)] and here: Juniper [[JUNIPER](#)]. Equivalent features do exist in several implementations, however publicly available documentation is not available. Finally, the examples cited below use Cisco IOS CLI for ease of illustration purposes only.

3. Local AS: Modify Inbound BGP AS_PATH Attribute

ISP B needs to reconfigure its router(s) to participate as an internal BGP speaker in AS 200, to realize the business goal of becoming a single Service Provider: ISP A'. ISP B needs to do this without coordinating the change of its ASN with all of its eBGP peers, simultaneously. The first step is for ISP B to change the global AS in its router configuration, used by the local BGP process as the system-wide Autonomous System ID, from AS 300 to AS 200. The next step is for ISP B to establish iBGP sessions with ISP A's existing routers, thus consolidating ISP B into ISP A operating under a single AS: ISP A', (AS 200).

The next step is for ISP B to reconfigure its PE router(s) so that each of its eBGP sessions toward all eBGP speakers with a feature called "Local AS". This feature allows ISP B's PE router to re-establish an eBGP session toward the existing CE devices using the legacy AS, AS 300, in the eBGP session establishment. Ultimately, the CE devices, (i.e.: customer C), are completely unaware that ISP B has reconfigured its router to participate as a member of a new AS. Within the context of ISP B's PE router, the second effect this feature has is that, by default, it prepends all received BGP UPDATE's with the legacy AS of ISP B: AS 300. Thus, within ISP A' the AS_PATH toward customer C would appear as: 300 100, which is an increase in AS_PATH length from previously. Therefore, a secondary feature "No Prepend" is required to be added to the "Local AS" configuration toward every eBGP neighbor on ISP B's PE router. The "No Prepend" feature causes ISP B's PE router to not prepend the legacy AS, AS 300, on all received eBGP UPDATE's from customer C.

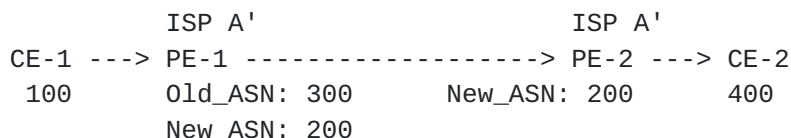
This restores the AS_PATH within ISP A' toward customer C so that it is just one AS long: 100.

In the direction of CE -> PE (inbound):

1. 'local-as <old_ASN>': appends the <old_ASN> value to the AS_PATH of routes received from the CE
2. 'local-as <old_ASN> no-prepend': does not prepend <old_ASN> value to the AS_PATH of routes received from the CE

As stated previously, local-as <old_ASN> no-prepend, (configuration #2), is critical because it does not increase the AS_PATH length. Ultimately, this ensures that routes learned from ISP B's legacy customers will be transmitted through legacy eBGP sessions of ISP A, toward both customers and peers, will contain only two ASes in the AS_PATH: 200 100. Thus, the legacy customers and peers of ISP A will not see an increase in the AS_PATH length to reach ISP B's legacy customers.

PE-1 is a PE that was originally in ISP B. PE-1 has had its global configuration AS changed from AS 300 to AS 200 to make it part of the permanently retained ASN. This now makes PE-1 a member of ISP A'. PE-2 is a PE that was originally in ISP A. Although its global configuration AS remains AS 200 throughout this exercise we also consider PE-2 a member of ISP A'.



Note: Direction of BGP UPDATE as per the arrows.

Figure 1: Local AS BGP UPDATE Diagram

The final configuration on PE-1 after completing the "Local AS" portion of the AS migration is as follows:

```

router bgp 200
  neighbor <CE-1_IP> remote-as 100
  neighbor <CE-1_IP> local-as 300 no-prepend

```

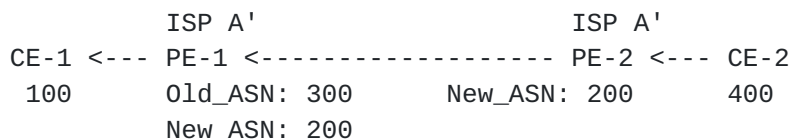
As a result of the "Local AS No Prepend" configuration, on PE-1, CE-2 will see an AS_PATH of: 200 100. CE-2 will not receive a BGP UPDATE containing AS 300 in the AS_PATH. (If only the "local-as 300" feature was configured without the keyword "no-prepend" on PE-1, then CE-2 would see an AS_PATH of: 100 300 200, which is not desirable).

4. Replace AS: Modify Outbound BGP AS_PATH Attribute

The previous feature, "Local AS No Prepend", modifies not only the AS_PATH Attribute received from CE devices, but also the AS_PATH Attribute transmitted to CE devices as well. Specifically, with "Local AS No Prepend" enabled on ISP B's PE, it causes outbound BGP UPDATES toward its directly attached eBGP speakers, customer C in AS 100, to have the AS_PATH Attribute prepended with the new global configuration AS on ISP B's PE. The result will be that customer C, in AS 100, will receive the following AS_PATH: 300 200 400. Therefore, if ISP B takes no further action, it will cause an increase in AS_PATH length within customer's networks directly attached to ISP B, which is unacceptable.

A second feature, called "Replace AS", was designed to resolve this problem. This feature allows the ISP to not prepend the global configured AS in outbound BGP UPDATE's toward customer's networks configured with the "Local AS" feature. Instead, only the historical (or, legacy) AS will be prepended in the outbound BGP UPDATE toward customer's network, restoring the AS_PATH length to what it was before AS Migration occurred.

To re-use the above diagram, but in the opposite direction, we have:



Note: Direction of BGP UPDATE as per the arrows.

Figure 2: Replace AS BGP UPDATE Diagram

The final configuration on PE-1 after completing the "Replace AS" portion of the AS migration is as follows:

```

router bgp 200
  neighbor <CE-1_IP> remote-as 100
  neighbor <CE-1_IP> local-as 300 no-prepend replace-as
  
```

By default, without "Replace AS" enabled, CE-1 would see an AS_PATH of: 300 200 400, which is artificially lengthened by the AS Migration. After ISP A' changes PE-1 to include the "Replace AS" feature, CE-1 would receive an AS_PATH of: 200 400, which is the same AS_PATH length pre-AS migration.

5. Additional Operational Considerations

This document describes just two variations of implementation-specific features to support ISP's and other organizations that need to perform ASN migrations. Other variations of these features may exist, for example, in legacy router software that has not been upgraded or reached End of Life, but continues to operate in the network. Such variations are beyond the scope of this document.

Companies routinely go through periods of mergers, acquisitions and divestitures, which in the case of the former cause them to accumulate several legacy ASNs over time. Since ISPs often do not have control over the configuration of customer's devices, (i.e.: the ISPs are often not providing a managed CE router service, particularly to medium and large customers that require eBGP), and since the ISP is using methods to perform ASN migration that do not require coordination with customers, there is not a finite period of time that the legacy ASNs will be completely expunged from the ISP's network. In fact, it is common that legacy ASNs along with the features discussed in this document can and do persist for several years if not longer. Thus, it is prudent to plan that legacy ASNs and associated features will persist in a operational network indefinitely.

6. Conclusion

Although the features discussed in this document are not formally recognized as part of the BGP4 specification, they have been in existence in commercial implementations for well over a decade. These features are widely known by the operational community and will continue to be a critical necessity in the support of network integration activities going forward. Therefore, these features are extremely unlikely to be deprecated by vendors. As a result, these features must be acknowledged by protocol designers, particularly when there are proposals to modify BGP's behavior with respect to handling or manipulation of the AS_PATH Attribute. More specifically, assumptions should not be made with respect to the preservation or consistency of the AS_PATH Attribute as it is transmitted along a sequence of ASNs. In addition, proposals to manipulate the AS_PATH that would gratuitously increase AS_PATH length or remove the capability to use these features described in this document will not be accepted by the operational community.

7. Acknowledgements

Thanks to Kotikalapudi Sriram for his comments.

8. IANA Considerations

This memo includes no request to IANA.

9. Security Considerations

This draft discusses a process by which one ASN is migrated into and subsumed by another. This involves manipulating the AS_PATH Attribute with the intent of not increasing the AS_PATH length, which would typically cause the BGP route to no longer be selected by BGP's Path Selection Algorithm in other's networks. This could result in a loss of revenue if the ISP is billing based on measured utilization of traffic sent to/from entities attached to its network. This could also result in sudden, and unexpected shifts in traffic patterns in the network, potentially resulting in congestion, in the most extreme cases. While this I-D does not introduce any new security considerations, as it is simply documenting current functionality, the reader should be aware that manipulation of the AS_PATH attribute has several known attack vectors that can result in the same sort of undesired traffic shifts.

Given that these features can only be enabled through configuration of router's within a single network, standard security measures should be taken to restrict access to the management interface(s) of routers that implement these features.

10. References

10.1. Normative References

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

10.2. Informative References

[CISCO] Cisco Systems, Inc., "BGP Support for Dual AS Configuration for Network AS Migrations", 2003, <http://www.cisco.com/en/US/docs/ios/12_3t/12_3t11/feature/guide/gtbgpdas.html>.

[JUNIPER] Juniper Networks, Inc., "Configuring the BGP Local

Autonomous System Attribute", 2012, <http://www.juniper.net/techpubs/en_US/junos12.2/information-products/topic-collections/nce/bgp-local-as/bgp-local-as-configuring.pdf>.

[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.

Authors' Addresses

Wesley George
Time Warner Cable
13820 Sunrise Valley Drive
Herndon, VA 20171
US

Phone: +1 703-561-2540
Email: wesley.george@twcable.com

Shane Amante
Level 3 Communications
1025 Eldorado Blvd
Broomfield, CO 80021
US

Phone:
Email: shane@level3.net

