

Internet Engineering Task Force (IETF)  
Request for Comments: 7705  
Updates: 4271  
Category: Standards Track  
ISSN: 2070-1721

W. George  
Time Warner Cable  
S. Amante  
Apple, Inc.  
November 2015

Autonomous System Migration Mechanisms  
and Their Effects on the BGP AS\_PATH Attribute

Abstract

This document discusses some existing commonly used BGP mechanisms for Autonomous System Number (ASN) migration that are not formally part of the BGP4 protocol specification. It is necessary to document these de facto standards to ensure that they are properly supported in future BGP protocol work.

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/rfc7705>.

Copyright Notice

Copyright (c) 2015 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 . . . . . 2
  - 1.1. Requirements Language . . . . . 3
  - 1.2. Documentation Note . . . . . 3
- 2. ASN Migration Scenario Overview . . . . . 3
- 3. External BGP Autonomous System Migration Mechanisms . . . . . 5
  - 3.1. Modify Inbound BGP AS\_PATH Attribute . . . . . 5
  - 3.2. Modify Outbound BGP AS\_PATH Attribute . . . . . 7
  - 3.3. Implementation . . . . . 8
- 4. Internal BGP Autonomous System Migration Mechanisms . . . . . 9
  - 4.1. Internal BGP AS Migration . . . . . 10
  - 4.2. Implementation . . . . . 12
- 5. Additional Operational Considerations . . . . . 13
- 6. Security Considerations . . . . . 14
- 7. References . . . . . 14
  - 7.1. Normative References . . . . . 14
  - 7.2. Informative References . . . . . 14
- Appendix A. Implementation Report . . . . . 16
- Acknowledgements . . . . . 16
- Authors' Addresses . . . . . 16

1. Introduction

This document discusses some existing commonly used BGP mechanisms for Autonomous System Number (ASN) migration that are not formally part of the BGP4 [RFC4271] protocol specification. These mechanisms are local to a given BGP speaker and do not require negotiation with or cooperation of BGP neighbors. The deployment of these mechanisms do not need to interwork with one another to accomplish the desired results, so slight variations between existing vendor implementations exist and will not necessarily be harmonized due to this document. However, it is necessary to document these de facto standards to ensure that new implementations can be successful, and 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 mechanisms.

The migration mechanisms discussed here are useful to ISPs and organizations of all sizes, but it is important to understand the business need for these mechanisms and illustrate why they are so critical for ISPs' operations. During a merger, acquisition, or divestiture involving two organizations it is necessary to seamlessly migrate both internal and external BGP speakers from one ASN to a second ASN. The overall goal in doing so is to simplify operations through consistent configurations across all BGP speakers in the combined network. In addition, given that the BGP Path Selection algorithm selects routes with the shortest AS\_PATH attribute, it is

critical that the ISP does not increase AS\_PATH length during or after ASN migration, because an increased AS\_PATH length would likely result in sudden, undesirable changes in traffic-patterns in the network.

By default, the BGP protocol requires an operator to configure a router to use a single remote ASN for the BGP neighbor, and the ASN must match on both ends of the peering in order to successfully negotiate and establish a BGP session. Prior to the existence of these migration mechanisms, it would have required an ISP to coordinate an ASN change with, in some cases, tens of thousands of customers. In particular, as each router is migrated to the new ASN, to avoid an outage due to ASN mismatch, the ISP would have to force all customers on that router to change their router configurations to use the new ASN immediately after the ASN change. Thus, it was critical to define a more asymmetric migration method so that the ISP could seamlessly change the ASN within its network(s) but allow the customers to gradually migrate to the ISP's new ASN when convenient. These customer migrations could occur either by coordinating individual session reconfigurations or, to allow for truly asymmetric migration, by accepting sessions using either the old or new ASN.

### 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].

### 1.2. Documentation Note

This document uses Autonomous System Numbers (ASNs) from the range reserved for documentation as described in RFC 5398 [RFC5398]. In the examples used here, they are intended to represent Globally Unique ASNs, not private use ASNs as documented in Section 5 of RFC 6996 [RFC6996].

## 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 prior to the ASN migration use AS 64500 and 64510, respectively. AS 64500 will be the permanently retained ASN used across the consolidated set of both ISPs network equipment, and AS 64510 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 64500. To all external BGP speakers, the AS\_PATH length will not be increased.

In this same scenario, AS 64496 and AS 64499 represent two separate customer networks: C and D, respectively. Originally, customer C (AS 64496) is attached to ISP B, which will undergo ASN migration from AS 64510 to AS 64500. Furthermore, customer D (AS 64499) is attached to ISP A, which does not undergo ASN migration since the ASN for ISP A will remain constant, (AS 64500). Although this example refers to AS 64496 and 64499 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.

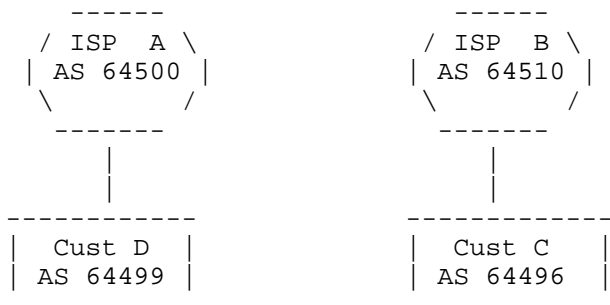


Figure 1: Before Migration

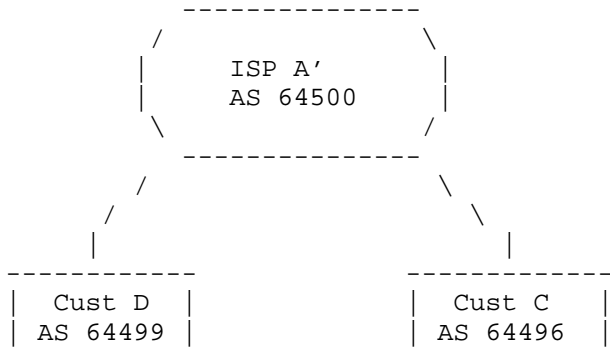


Figure 2: After Migration

The general order of operations, typically carried out in a single maintenance window by the network undergoing ASN migration (ISP B), is as follows. First, ISP B will change the global BGP ASN used by a Provider Edge (PE) router, from ASN 64510 to 64500. At this point, the router will no longer be able to establish External BGP (eBGP) sessions toward the existing Customer Edge (CE) devices that are attached to it and still using AS 64510. Second, since ISP B needs to do this without coordinating the simultaneous change of its ASN

with all of its eBGP peers, ISP B will configure two separate, but related ASN migration mechanisms discussed in this document on all eBGP sessions toward all CE devices. These mechanisms enable the router to establish BGP neighbors using the legacy ASN, modify the AS\_PATH attribute received from a CE device when advertising it further, and modify AS\_PATH when 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 CE devices at the edge of the network are not aware of the fact that their upstream router is now in a new ASN and do not observe any change in the length of the AS\_PATH attribute. However, after the changes discussed in this document are put in place by ISP A', there is a change to the contents of the AS\_PATH attribute to ensure the AS\_PATH is not artificially lengthened while these AS migration parameters are used.

In this use case, neither ISP is using BGP Confederations [RFC5065] internally.

### 3. External BGP Autonomous System Migration Mechanisms

The following section addresses optional capabilities that are specific to modifying the AS\_PATH attribute at the Autonomous System Border Routers (ASBRs) of an organization (typically a single Service Provider). This ensures that external BGP customers/peers are not forced to make any configuration changes on their CE routers before or during the exact time the Service Provider wishes to migrate to a new, permanently retained ASN. Furthermore, these mechanisms eliminate the artificial lengthening of the AS\_PATH both transmitted from and received by the Service Provider that is undergoing AS Migration, which would have negative implications on path selection by external networks.

#### 3.1. Modify Inbound BGP AS\_PATH Attribute

The first instrument used in the process described above is called "Local AS". This allows the router to supersede the globally configured ASN in the "My Autonomous System" field of the BGP OPEN [RFC4271] with a locally defined AS value for a specific BGP neighbor or group of neighbors. This mechanism allows the PE router that was formerly in ISP B to establish an eBGP session toward the existing CE devices using the legacy AS -- AS 64510. 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 the former ISP B PE router, the second effect this specific mechanism has on AS\_PATH is that, by default, it prepends all received BGP UPDATES with the legacy AS of ISP B -- AS 64510 --

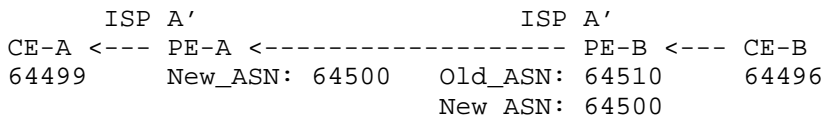
while advertising it (Adj-RIB-Out) to other BGP speakers (A'). Within the Loc-RIB on ISP B prior to the migration, the AS\_PATH of route announcements received from customer C would appear as 64496, whereas the same RIB on ISP A' (ISP B routers post-migration) would contain AS\_PATH 64510 64496.

The second instrument, referred to as "No Prepend Inbound", is enabled on PE routers migrating from ISP B. The "No Prepend Inbound" capability causes ISP B's routers to not prepend the legacy AS (AS 64510), when advertising UPDATES received from customer C. This restores the AS\_PATH within ISP A' for route announcements received from customer C so that it is just one ASN in length: 64496.

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

1. "Local AS": Allows the local BGP router to generate a BGP OPEN to an eBGP neighbor with the old, legacy ASN value in the "My Autonomous System" field. When this capability is activated, it also causes the local router to prepend the <old\_ASN> value to the AS\_PATH when installing or advertising routes received from a CE to Internal BGP (iBGP) neighbors inside the Autonomous System.
2. "No Prepend Inbound (of Local AS)": The local BGP router does not prepend the <old\_ASN> value to the AS\_PATH when installing or advertising routes received from the CE to iBGP neighbors inside the Autonomous System

PE-B is a PE that was originally in ISP B, and has a customer eBGP session to CE-B. PE-B has had its global configuration ASN changed from AS 64510 to AS 64500 to make it part of the permanently retained ASN. This now makes PE-B a member of ISP A'. PE-A is a PE that was originally in ISP A, and has a customer peer CE-A. Although its global configuration ASN remains AS 64500, throughout this exercise we also consider PE-A a member of ISP A'.



Note: Direction of BGP UPDATE as per the arrows.

Figure 3: Local AS and No Prepend BGP UPDATE Diagram

As a result, using both the "Local AS" and "No Prepend Inbound" capabilities on PE-B, CE-A will see an AS\_PATH of 64500 64496. CE-A will not receive a BGP UPDATE containing AS 64510 in the AS\_PATH. (If only the "Local AS" mechanism was configured without "No Prepend

Inbound" on PE-B, then CE-A would have seen an AS\_PATH of 64500 64510 64496, which results in an unacceptable lengthening of the AS\_PATH). Note: If there are still routers in the old ASN (64510), it is possible for them to accept these manipulated routes (i.e., those with 64510 removed from the AS\_PATH by this command) as if they have not already passed through their ASN, potentially causing a loop, since BGP's normal loop-prevention behavior of rejecting routes that include its ASN in the path will not catch these. Careful filtering between routers remaining in the old ASN and routers migrated to the new ASN is necessary to minimize the risk of routing loops.

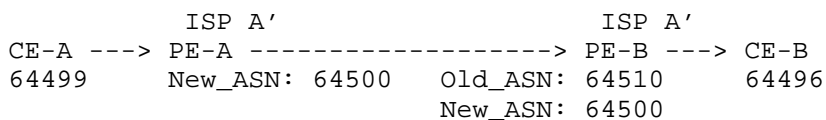
### 3.2. Modify Outbound BGP AS\_PATH Attribute

The two aforementioned mechanisms, "Local AS" and "No Prepend Inbound", only modify the AS\_PATH attribute received by the ISP's PEs in the course of processing BGP UPDATEs from CE devices when CE devices still have an eBGP session established with the ISPs legacy AS (AS64510).

In some existing implementations, "Local AS" and "No Prepend Inbound" do not concurrently modify the AS\_PATH attribute for BGP UPDATEs that are transmitted by the ISP's PEs to CE devices. In these implementations, with "Local AS" and "No Prepend Inbound" used on PE-B, there is an automatic lengthening of the AS\_PATH in outbound BGP UPDATEs from ISP A' toward directly attached eBGP speakers (customer C in AS 64496). The externally observed result is that customer C in AS 64496 will receive the following AS\_PATH: 64510 64500 64499. Therefore, if ISP A' takes no further action, there will be an unacceptable increase in the AS\_PATH length within the customer's networks directly attached to ISP A'.

A tertiary mechanism, referred to as "Replace Old AS", is used to resolve this problem. This capability allows ISP A' to prevent routers from appending the globally configured ASN in outbound BGP UPDATEs toward directly attached eBGP neighbors that are using the "Local AS" mechanism. Instead, only the old (or previously used) AS will be prepended in the outbound BGP UPDATE toward the customer's network, restoring the AS\_PATH length to what it was before AS Migration occurred.

To reuse the above diagram, but in the opposite direction, we have:



Note: Direction of BGP UPDATE as per the arrows.

Figure 4: Replace AS BGP UPDATE Diagram

By default, without the use of "Replace Old AS", CE-B would see an AS\_PATH of 64510 64500 64499. After ISP A' changes PE-B to use "Replace Old AS", CE-B would receive an AS\_PATH of 64510 64499, which is the same AS\_PATH length pre-AS migration.

### 3.3. Implementation

The mechanisms introduced in this section MUST be configurable on a per-neighbor or per-neighbor-group basis to allow for maximum flexibility. (Here, "neighbor group" refers to a group of similar BGP neighbor statements that reuse some common configuration to simplify provisioning.) When the "Local AS" capability is used, a local ASN will be provided in the configuration that is different from the globally configured ASN of the BGP router. To implement this mechanism, a BGP speaker SHOULD send BGP OPEN [RFC4271] (see Section 4.2) messages to the configured eBGP peer(s) using the local ASN configured for this session as the value sent in "My Autonomous System". The BGP router SHOULD NOT use the ASN configured globally within the BGP process as the value sent in "My Autonomous System" in the OPEN message. This prevents causing the eBGP neighbor to unnecessarily generate a BGP OPEN Error message "Bad Peer AS". This method is typically used to re-establish eBGP sessions with peers expecting the legacy ASN after a router has been moved to a new ASN.

Implementations MAY support a more flexible model where the eBGP speaker attempts to open the BGP session using either the ASN configured as "Local AS" or the globally configured AS as discussed in BGP Alias (Section 4.2). If the session is successfully established to the globally configured ASN, then the modifications to AS\_PATH described in this document SHOULD NOT be performed, as they are unnecessary. The benefit to this more flexible model is that it allows the remote neighbor to reconfigure to the new ASN without direct coordination between the ISP and the customer.

Note that this procedure will vary slightly if the locally or globally configured ASN is a 4-octet ASN. See Section 3 of [RFC6793].



When the BGP router receives UPDATES from its eBGP neighbor configured with the "Local AS" mechanism, it processes the UPDATE as described in RFC4271 Section 5.1.2 [RFC4271]. However, the presence of a second ASN due to "Local AS" adds the following behavior to processing UPDATES received from an eBGP neighbor configured with this mechanism:

1. Internal: The router SHOULD append the configured "Local AS" ASN in the AS\_PATH attribute before installing the route or advertising the UPDATE to an iBGP neighbor. The decision of when to append the ASN is an implementation detail outside the scope of this document. Some considerations factoring into this decision include consistency in the AS\_PATH throughout the AS, and implementation of the loop detection mechanism.
2. External: The BGP router SHOULD first append the globally configured ASN to the AS\_PATH immediately followed by the "Local AS" value before advertising the UPDATE to an eBGP neighbor.

Two options exist to manipulate the behavior of the basic "Local AS" mechanism. They modify the behavior as described below:

1. "No Prepend Inbound": When the BGP router receives inbound BGP UPDATES from its eBGP neighbor configured with this option, it MUST NOT append the "Local AS" ASN value in the AS\_PATH attribute when installing the route or advertising that UPDATE to iBGP neighbors, but it MUST still append the globally configured ASN as normal when advertising the UPDATE to other local eBGP neighbors (i.e., those natively peering with the globally configured ASN).
2. "Replace Old AS" (outbound): When the BGP router generates outbound BGP UPDATES toward an eBGP neighbor configured with this option, the BGP speaker MUST NOT append the globally configured ASN from the AS\_PATH attribute. The BGP router MUST append only the configured "Local AS" ASN value to the AS\_PATH attribute before sending the BGP UPDATES outbound to the eBGP neighbor.

#### 4. Internal BGP Autonomous System Migration Mechanisms

The following section describes mechanisms that assist with a gradual and least service-impacting migration of Internal BGP sessions from a legacy ASN to the permanently retained ASN. The following mechanism is very valuable to networks undergoing AS migration, but its use does not cause changes to the AS\_PATH attribute.

#### 4.1. Internal BGP AS Migration

In this case, all of the routers to be consolidated into a single, permanently retained ASN are under the administrative control of a single entity. Unfortunately, the traditional method of migrating all Internal BGP speakers, particularly within larger networks, is both time-consuming and widely service impacting.

The traditional method to migrate Internal BGP sessions was strictly limited to reconfiguration of the global configuration ASN and, concurrently, changing all iBGP neighbors' remote ASN from the legacy ASN to the new, permanently retained ASN on each router within the legacy AS. These changes can be challenging to swiftly execute in networks with more than a few dozen internal BGP routers. There are also the concomitant service interruptions as these changes are made to routers within the network, resulting in a reset of iBGP sessions and subsequent route reconvergence to reestablish optimal routing paths. Operators often cannot make such sweeping changes given the associated risks of a highly visible service interruption; rather, they require a more gradual method to migrate Internal BGP sessions, from one ASN to a second, permanently retained ASN, that is not visibly service impacting to its customers.

With the "Internal BGP AS Migration" mechanism described herein, it allows an Internal BGP speaker to form a single iBGP session using either the old, legacy ASN or the new, permanently retained ASN. The benefits of using this mechanism are several fold. First, it allows for a more gradual and less service-impacting migration away from the legacy ASN to the permanently retained ASN. Second, it (temporarily) permits the coexistence of the legacy and permanently retained ASN within a single network, allowing for uniform BGP path selection among all routers within the consolidated network.

The iBGP router with the "Internal BGP AS Migration" capability enabled allows the receipt of a BGP OPEN message with either the legacy ASN value or the new, globally configured ASN value in the "My Autonomous System" field of the BGP OPEN message from iBGP neighbors. It is important to recognize that enablement of the "Internal BGP AS Migration" mechanism preserves the semantics of a regular iBGP session (i.e., using identical ASNs). Thus, the BGP attributes transmitted by and the acceptable methods of operation on BGP attributes received from iBGP sessions configured with "Internal BGP AS Migration" capability are no different than those exchanged across an iBGP session without "Internal BGP AS Migration" configured, as defined by [RFC4271] and [RFC4456].

Typically, in medium to large networks, BGP Route Reflectors (RRs) [RFC4456] are used to aid in reduction of configuration of iBGP sessions and scalability with respect to overall TCP (and BGP) session maintenance between adjacent iBGP routers. Furthermore, BGP Route Reflectors are typically deployed in pairs within a single Route Reflection cluster to ensure high reliability of the BGP Control Plane. As such, the following example will use Route Reflectors to aid in understanding the use of the "Internal BGP AS Migration" mechanism. Note that Route Reflectors are not a prerequisite to enable "Internal BGP AS Migration" and this mechanism can be enabled independent of the use of Route Reflectors.

The general order of operations is as follows:

1. Within the legacy network, (the routers comprising the set of devices that still have a globally configured legacy ASN), one member of a redundant pair of RRs has its global configuration ASN changed to the permanently retained ASN. Concurrently, the "Internal BGP AS Migration" capability is enabled on all iBGP sessions on that device. This will comprise Non-Client iBGP sessions to other RRs as well as Client iBGP sessions, typically to PE devices, both still utilizing the legacy ASN. Note that during this step there will be a reset and reconvergence event on all iBGP sessions on the RRs whose configuration was modified; however, this should not be service impacting due to the use of redundant RRs in each RR Cluster.
2. The above step is repeated for the other side of the redundant pair of RRs. The one alteration to the above procedure is that the "Internal BGP AS Migration" mechanism is now removed from the Non-Client iBGP sessions toward the other (previously reconfigured) RRs, since it is no longer needed. The "Internal BGP AS Migration" mechanism is still required on all RRs for all RR Client iBGP sessions. Also during this step, there will be a reset and reconvergence event on all iBGP sessions whose configuration was modified, but this should not be service impacting. At the conclusion of this step, all RRs will have their globally configured ASN set to the permanently retained ASN and "Internal BGP AS Migration" enabled and in use toward RR Clients.
3. At this point, the network administrators would then be able to establish iBGP sessions between all Route Reflectors in both the legacy and permanently retained networks. This would allow the network to appear to function, both internally and externally, as a single, consolidated network using the permanently retained network.

4. To complete the AS migration, each RR Client (PE) in the legacy network still utilizing the legacy ASN is now modified. Specifically, each legacy PE would have its globally configured ASN changed to use the permanently retained ASN. The ASN configured within the PE for the iBGP sessions toward each RR would be changed to use the permanently retained ASN. It is unnecessary to enable the "Internal BGP AS Migration" mechanism on these migrated iBGP sessions. During the same maintenance window, External BGP sessions would be modified to include the above "Local AS", "No Prepend", and "Replace Old AS" mechanisms described in Section 3 above, since all of the changes are service interrupting to the eBGP sessions of the PE. At this point, all PEs will have been migrated to the permanently retained ASN.
5. The final step is to excise the "Internal BGP AS Migration" configuration from the Router Reflectors in an orderly fashion. After this is complete, all routers in the network will be using the new, permanently retained ASN for all iBGP sessions with no vestiges of the legacy ASN on any iBGP sessions.

The benefit of using the aforementioned "Internal BGP AS Migration" capability is that it is a more gradual and less externally service impacting change to accomplish an AS migration. Previously, without "Internal BGP AS Migration", such an AS migration change would carry a high risk and need to be successfully accomplished in a very short time frame (e.g., at most several hours). In addition, it would likely cause substantial routing churn and rapid fluctuations in traffic carried -- potentially causing periods of congestion and resultant packet loss -- during the period when the configuration changes are underway to complete the AS Migration. On the other hand, with "Internal BGP AS Migration", the migration from the legacy ASN to the permanently retained ASN can occur over a period of days or weeks with reduced customer disruption. (The only observable service disruption should be when each PE undergoes the changes discussed in step 4 above.)

#### 4.2. Implementation

The mechanism introduced in this section MUST be configurable on a per-neighbor or per-neighbor-group basis to allow for maximum flexibility. When configured with this mechanism, a BGP speaker MUST accept BGP OPEN and establish an iBGP session from configured iBGP peers if the ASN value in "My Autonomous System" is either the globally configured ASN or a locally configured ASN provided when this capability is utilized. Additionally, a BGP router configured with this mechanism MUST send its own BGP OPEN [RFC4271] (see Section 4.2) using either the globally configured or the locally

configured ASN in "My Autonomous System" as follows. To avoid potential deadlocks when two BGP speakers are attempting to establish a BGP peering session and are both configured with this mechanism, the speaker SHOULD send BGP OPEN using the globally configured ASN first, and only send a BGP OPEN using the locally configured ASN as a fallback if the remote neighbor responds with the BGP error "Bad Peer AS". In each case, the BGP speaker MUST treat UPDATES sent and received to this peer as if this was a natively configured iBGP session, as defined by [RFC4271] and [RFC4456].

Note that this procedure will vary slightly if the locally or globally configured ASN is a 4-octet ASN. See Section 3 of [RFC6793].

## 5. Additional Operational Considerations

This document describes several mechanisms to support ISPs and other organizations that need to perform ASN migrations. Other variations of these mechanisms 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. ISPs often do not have control over the configuration of customers' devices (i.e., the ISPs are often not providing a managed CE router service, particularly to medium and large customers that require eBGP). Furthermore, ISPs are using methods to perform ASN migration that do not require coordination with customers. Ultimately, this means there is not a finite period of time after which legacy ASNs will be completely expunged from the ISP's network. In fact, it is common that legacy ASNs and the associated External BGP AS Migration mechanisms 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 External BGP AS Migration mechanisms will persist in an operational network indefinitely.

With respect to the Internal BGP AS Migration mechanism, all of the routers to be consolidated into a single, permanently retained ASN are under the administrative control of a single entity. Thus, completing the migration from iBGP sessions using the legacy ASN to the permanently retained ASN is more straightforward and could be accomplished in a matter of days to months. Finally, good operational hygiene would dictate that it is good practice to avoid using "Internal BGP AS Migration" capability over a long period of time for reasons of not only operational simplicity of the network,

but also reduced reliance on that mechanism during the ongoing life cycle management of software, features, and configurations that are maintained on the network.

## 6. Security Considerations

This document 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 others' networks. This could result in sudden and unexpected shifts in traffic patterns in the network, potentially resulting in congestion.

Given that these mechanisms can only be enabled through configuration of routers within a single network, standard security measures should be taken to restrict access to the management interface(s) of routers that implement these mechanisms. Additionally, BGP sessions SHOULD be protected using TCP Authentication Option [RFC5925] and the Generalized TTL Security Mechanism [RFC5082]

## 7. References

### 7.1. Normative References

- [RFC2119] Bradner, S., "Key words for use in RFCs to Indicate Requirement Levels", BCP 14, RFC 2119, DOI 10.17487/RFC2119, March 1997, <<http://www.rfc-editor.org/info/rfc2119>>.
- [RFC4271] Rekhter, Y., Ed., Li, T., Ed., and S. Hares, Ed., "A Border Gateway Protocol 4 (BGP-4)", RFC 4271, DOI 10.17487/RFC4271, January 2006, <<http://www.rfc-editor.org/info/rfc4271>>.
- [RFC4456] Bates, T., Chen, E., and R. Chandra, "BGP Route Reflection: An Alternative to Full Mesh Internal BGP (IBGP)", RFC 4456, DOI 10.17487/RFC4456, April 2006, <<http://www.rfc-editor.org/info/rfc4456>>.

### 7.2. Informative References

- [ALU] Alcatel-Lucent, "BGP Local AS attribute", 2006-2012, <[https://infoproducts.alcatel-lucent.com/html/0\\_add-h-f/93-0074-10-01/7750\\_SR\\_OS\\_Routing\\_Protocols\\_Guide/BGP-CLI.html#709567](https://infoproducts.alcatel-lucent.com/html/0_add-h-f/93-0074-10-01/7750_SR_OS_Routing_Protocols_Guide/BGP-CLI.html#709567)>.

- [CISCO] Cisco Systems, Inc., "BGP Support for Dual AS Configuration for Network AS Migrations", 2013, <[http://www.cisco.com/c/en/us/td/docs/ios-xml/ios/iproute\\_bgp/configuration/xe-3s/asr1000/irg-xe-3s-asr1000-book/irg-dual-as.html](http://www.cisco.com/c/en/us/td/docs/ios-xml/ios/iproute_bgp/configuration/xe-3s/asr1000/irg-xe-3s-asr1000-book/irg-dual-as.html)>.
- [JUNIPER] Juniper Networks, Inc., "Understanding the BGP Local AS Attribute", December 2013, <[http://www.juniper.net/techpubs/en\\_US/junos13.3/topics/concept/bgp-local-as-introduction.html](http://www.juniper.net/techpubs/en_US/junos13.3/topics/concept/bgp-local-as-introduction.html)>.
- [RFC5065] Traina, P., McPherson, D., and J. Scudder, "Autonomous System Confederations for BGP", RFC 5065, DOI 10.17487/RFC5065, August 2007, <<http://www.rfc-editor.org/info/rfc5065>>.
- [RFC5082] Gill, V., Heasley, J., Meyer, D., Savola, P., Ed., and C. Pignataro, "The Generalized TTL Security Mechanism (GTSM)", RFC 5082, DOI 10.17487/RFC5082, October 2007, <<http://www.rfc-editor.org/info/rfc5082>>.
- [RFC5398] Huston, G., "Autonomous System (AS) Number Reservation for Documentation Use", RFC 5398, DOI 10.17487/RFC5398, December 2008, <<http://www.rfc-editor.org/info/rfc5398>>.
- [RFC5925] Touch, J., Mankin, A., and R. Bonica, "The TCP Authentication Option", RFC 5925, DOI 10.17487/RFC5925, June 2010, <<http://www.rfc-editor.org/info/rfc5925>>.
- [RFC6793] Vohra, Q. and E. Chen, "BGP Support for Four-Octet Autonomous System (AS) Number Space", RFC 6793, DOI 10.17487/RFC6793, December 2012, <<http://www.rfc-editor.org/info/rfc6793>>.
- [RFC6996] Mitchell, J., "Autonomous System (AS) Reservation for Private Use", BCP 6, RFC 6996, DOI 10.17487/RFC6996, July 2013, <<http://www.rfc-editor.org/info/rfc6996>>.

## Appendix A. Implementation Report

As noted elsewhere in this document, this set of migration mechanisms has multiple existing implementations in wide use.

- o Cisco [CISCO]
- o Juniper [JUNIPER]
- o Alcatel-Lucent [ALU]

This is not intended to be an exhaustive list, as equivalent features do exist in other implementations; however, the authors were unable to find publicly available documentation of the vendor-specific implementation to reference.

## Acknowledgements

Thanks to Kotikalapudi Sriram, Stephane Litkowski, Terry Manderson, David Farmer, Jaroslaw Adam Gralak, Gunter Van de Velde, Juan Alcaide, Jon Mitchell, Thomas Morin, Alia Atlas, Alvaro Retana, and John Scudder for their comments.

## Authors' Addresses

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

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

Shane Amante  
Apple, Inc.  
1 Infinite Loop  
Cupertino, CA 95014  
United States

Email: amante@apple.com