

Internet Engineering Task Force
Internet-Draft
Intended status: Informational
Expires: December 15, 2014

W. George
Time Warner Cable
S. Amante
Apple, Inc.
June 13, 2014

Autonomous System (AS) Migration Features and Their Effects on the BGP
AS_PATH Attribute
draft-ietf-idr-as-migration-01

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 future BGP protocol work such as 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 December 15, 2014.

Copyright Notice

Copyright (c) 2014 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	4
3. External BGP Autonomous System Migration Features	6
3.1. Modify Inbound BGP AS_PATH Attribute	6
3.2. Modify Outbound BGP AS_PATH Attribute	8
3.3. Implementation	9
4. Internal BGP Autonomous System Migration Features	10
4.1. Internal BGP Alias	10
4.2. Implementation	13
5. Additional Operational Considerations	14
6. Conclusion	14
7. Acknowledgements	15
8. IANA Considerations	15
9. Security Considerations	15
10. Appendix: Implementation report	15
11. References	16
11.1. Normative References	16
11.2. Informative References	16
Authors' Addresses	17

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. These features are local to a given BGP Speaker 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, so slight variations between existing vendor implementations exist. 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 and illustrate why they are critical, particularly for ISPs' operations. However, these features are not limited to ISPs and organizations of all sizes use these features for similar reasons to

ISPs. 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.

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 becomes critical to allow the ISP to make this process a bit more asymmetric, so that it could seamlessly migrate the ASN within its network(s), but not disturb existing customers, and allow the customers 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].

1.2. Documentation note

This draft 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 RFC 6996 [RFC6996] section 10.

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 64500 and 64510, respectively, before the ASN migration is to occur. AS 64500 will be the permanently retained ASN used going forward 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 64496) is attached to ISP A, which does not undergo ASN migration since ISP A's ASN 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, are as follows. First, ISP B, will change the global BGP ASN used by a PE router, from ASN 64510 to 64500. 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 64510. 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 CE devices at the edge of the network are not aware of 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 for the duration of time that these AS migration parameters are used.

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

There are multiple implementations with equivalent features deployed and in use. Some documentation pointers to these implementations, as well as additional documentation on migration scenarios can be found in the appendix. The examples cited below use Cisco IOS CLI for ease of illustration purposes only.

3. External BGP Autonomous System Migration Features

The following section addresses features 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 features 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

ISP B needs to reconfigure its router(s) to participate as an internal BGP speaker in AS 64500, 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 64510 to AS 64500. The next step is for ISP B to establish iBGP sessions with ISP A's existing routers, thus consolidating ISP B into ISP A resulting in operating under a single AS: ISP A', (AS 64500).


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 a eBGP session toward the existing CE devices using the legacy AS, AS 64510, 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 64510. Thus, within ISP A' the AS_PATH toward customer C would appear as: 64510 64496, 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 64510, 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 ASN in length: 64496.

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 AS'es in the AS_PATH: 64500 64496. 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. Ultimately, it is considered mandatory by operators that both the "Local AS" and "No Prepend" configuration parameters always be used in conjunction with each other in order to ensure the AS_PATH length is not increased.

PE-1 is a PE that was originally in ISP B. PE-1 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-1 a member of ISP A'. PE-2 is a PE that was originally in ISP A. Although its global configuration ASN remains AS 64500, throughout this exercise we also consider PE-2 a member of ISP A'.



Note: Direction of BGP UPDATE as per the arrows.

Figure 3: 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 64500
  neighbor <CE-1_IP> remote-as 64496
  neighbor <CE-1_IP> local-as 64510 no-prepend
    
```


As a result of the "Local AS No Prepend" configuration, on PE-1, CE-2 will see an AS_PATH of: 64500 64496. CE-2 will not receive a BGP UPDATE containing AS 64510 in the AS_PATH. (If only the "local-as 64510" feature was configured without the keyword "no-prepend" on PE-1, then CE-2 would see an AS_PATH of: 64496 64510 64500, which is unacceptable).

3.2. Modify Outbound BGP AS_PATH Attribute

The previous feature, "Local AS No Prepend", was only designed to modify the AS_PATH Attribute received by the ISP in 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 No Prepend" does not concurrently modify the AS_PATH Attribute for BGP UPDATES that are transmitted by the ISP to CE devices. Specifically, with "Local AS No Prepend" enabled on ISP A's PE-1, it automatically causes a lengthening of the AS_PATH in outbound BGP UPDATES from ISP A' toward directly attached eBGP speakers, (Customer C in AS 64496). This is the result of the "Local AS No Prepend" feature automatically appending the new global configuration ASN, AS64500, after the legacy ASN, AS64510, on ISP A' PE-1 in BGP UPDATES that are transmitted by PE-1 to CE-1. The end 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, it will cause an increase in AS_PATH length within customer's networks directly attached to ISP A', which is unacceptable.

A second feature was designed to resolve this problem (continuing the use of Cisco CLI in the examples, it is called "Replace AS" in the examples below). This feature allows ISP A' to prevent routers configured with this feature from appending the global configured AS in outbound BGP UPDATES toward its 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 4: 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 64500
  neighbor <CE-1_IP> remote-as 64496
  neighbor <CE-1_IP> local-as 64510 no-prepend replace-as
  
```

By default, without "Replace AS" enabled, CE-1 would see an AS_PATH of: 64510 64500 64499, which is artificially lengthened by the ASN Migration. After ISP A' changes PE-1 to include the "Replace AS" feature, CE-1 would receive an AS_PATH of: 64510 64499, which is the same AS_PATH length pre-AS migration.

3.3. Implementation

While multiple implementations already exist, the following should document the expected behavior such that a new implementation of this feature could be done on other platforms.

These features MUST be configurable on a per-neighbor or per peer-group basis to allow for maximum flexibility. When this feature set is invoked, an ASN that is different from the globally-configured ASN is provided as a part of the command as exemplified above. To implement this feature, a BGP speaker MUST send BGP OPEN messages to the configured eBGP peer using the ASN configured for this session as the value sent in MY ASN. The speaker MUST NOT use the ASN configured globally within the BGP process as the value sent in MY ASN in the OPEN message. This will avoid the BGP OPEN Error message BAD PEER AS, and is typically used to re-establish eBGP sessions with peers expecting the legacy ASN after a router has been moved to a new ASN. Additionally, when the BGP speaker configured with this feature receives updates from its neighbor, it MUST append the configured ASN in the AS_PATH attribute before processing the update as normal. Note that processing the update as normal will include appending the globally configured ASN to the AS_PATH, thus processing this update will result in the addition of two ASNs to the AS_PATH attribute. Similarly, for outbound updates sent by the configured BGP speaker to its neighbor, the speaker MUST append the configured ASN to the AS_PATH attribute, adding to the existing global ASN in the AS_PATH, for a total of two ASNs added to the AS_PATH.

Two options exist to manipulate the behavior of this feature. They modify the behavior as described below:

No prepend inbound - When the BGP speaker configured with this option receives inbound updates from its neighbor, it MUST NOT append the configured ASN in the AS_PATH attribute and instead MUST append only the globally configured ASN.

No prepend outbound - When the BGP speaker configured with this option generates outbound BGP updates to the configured peer, the BGP speaker MUST remove the globally configured ASN from the AS_PATH attribute, and MUST append the locally configured ASN to the AS_PATH attribute before sending outbound BGP updates to the configured peer.

While the exact command syntax is an implementation detail beyond the scope of this document, the following consideration may be helpful for implementers: Implementations MAY integrate the behavior of the options described above into a single command that addresses both inbound and outbound updates, but if this is done, implementations MUST provide a method to select its applicability to inbound updates, outbound updates, or updates in both directions. Several existing implementations use separate commands (e.g. local-as no-prepend vs local-as replace-as) for maximum flexibility in controlling the behavior on the session to address the widest range of possible migration scenarios.

4. Internal BGP Autonomous System Migration Features

The following section describes features that are specific to performing an ASN migration within medium to large networks in order to realize the business and operational benefits of a single network using one, globally unique Autonomous System. These features assist with a gradual and least service impacting migration of Internal BGP sessions from a legacy ASN to the permanently retained ASN. It should be noted that the following feature is very valuable to networks undergoing AS migration, but its use does not cause changes to the AS_PATH attribute.

4.1. Internal BGP Alias

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 of iBGP neighbor's 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 with more than a few dozen internal BGP speakers. There is also the concomitant service interruptions as these changes are made to routers within the network, resulting in a reset of iBGP sessions and subsequent reconvergence times to reestablish optimal routing paths. Operators do not, and in some cases, cannot make such changes given the associated risks and 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 Alias" [JUNIPER] feature, 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 feature 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. NB: Cisco doesn't have an exact equivalent to "Internal BGP Alias", but the combination of the Cisco features iBGP local-AS and dual-as provides similar functionality.

When the "Internal BGP Alias" feature is enabled, typically just on one side of a iBGP session, it allows that iBGP speaker to establish a single iBGP session with either the legacy ASN or the new, permanently retained ASN, depending on which one it receives in the "My Autonomous System" field of the BGP OPEN message from its iBGP session neighbor. It is important to recognize that enablement of the "Internal BGP Alias" feature preserves the semantics of a regular iBGP session, (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 Alias" are no different than those exchanged across an iBGP session without "Internal BGP Alias" configured, as defined by [RFC4271] and [RFC4456].

Typically, in medium to large networks, BGP Route Reflectors [RFC4456] (RRs) 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 speakers. 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 Alias" feature. Note that Route Reflectors are not a prerequisite to enable "Internal BGP Alias" and this feature 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), take one member of a redundant pair of RRs and change its global configuration ASN to the permanently retained ASN. Concurrently, enable use of "Internal BGP Alias" on all iBGP sessions. 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. Repeat the above step for the other side of the redundant pair of RRs. The one alteration to the above procedure is to disable use of "Internal BGP Alias" on the Non-Client iBGP sessions toward the other (previously reconfigured) RRs, since it is no longer needed. "Internal BGP Alias" 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 should now have their globally configured ASN set to the permanently retained ASN and "Internal BGP Alias" 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. The next steps to complete the AS migration are to gradually modify each RR Client, (PE), in the legacy network still utilizing the legacy ASN. Specifically, each legacy PE would have its globally configured ASN changed to use the permanently retained ASN. The ASN used by the PE for the iBGP sessions, toward each RR, would be changed to use the permanently retained ASN. (It is unnecessary to enable "Internal BGP Alias" on the migrated iBGP sessions). During the same maintenance window, External BGP sessions would be modified to include the above "Local AS No Prepend" and "Replace-AS" features, since all of the changes are service interrupting to the eBGP sessions of the PE. At this point, all PE's will have been migrated to the permanently retained ASN.
5. The final step is to excise the "Internal BGP Alias" configuration from the first half of the legacy RR Client pair -- this will expunge "Internal BGP Alias" configuration from all devices in the network. 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 "Internal BGP Alias" is that it is a more gradual and less externally visible, service-impacting change to accomplish an AS migration. Previously, without "Internal BGP Alias", such an AS migration change would carry a high risk and need to be successfully accomplished in a very short timeframe, (e.g.: at most several hours). In addition, it would cause substantial routing churn and, likely, rapid fluctuations in traffic carried -- potentially causing periods of congestion and resultant packet loss -- during the period the configuration changes are underway to complete the AS Migration. On the other hand, with "Internal BGP Alias", the migration from the legacy ASN to the permanently retained ASN can occur over a period of days or weeks with little disruption experienced by customers of the network undergoing AS migration. (The only observable service disruption should be when each PE undergoes the changes discussed in step 4 above.)

4.2. Implementation

When configured with this feature, a BGP speaker MUST accept BGP OPEN and establish an iBGP session from configured iBGP peers if the ASN value in MY ASN is either the globally configured ASN or the locally configured ASN provided in this command. Additionally, a BGP speaker configured with this feature MUST send its own BGP OPEN using both the globally configured and the locally configured ASN in MY ASN. To avoid potential deadlocks when two BGP speakers are attempting to establish a BGP peering session and are both configured with this feature, 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 ASN. 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].

Implementations of this feature MAY integrate the functionality from the eBGP features (Section 3) section as a part of this command in order to simplify support for eBGP migrations as well as iBGP migrations, such that an eBGP session to a configured neighbor could be established via either the global ASN or the locally configured ASN. If the eBGP session is established with the global ASN, no modifications to AS_PATH are required, but if the eBGP session is established with the locally configured ASN, the modifications discussed in eBGP features (Section 3) MUST be implemented to properly manipulate the AS_PATH.

5. Additional Operational Considerations

This document describes several features to support ISPs 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. 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). 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 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 External BGP AS Migration features will persist in an operational network indefinitely.

With respect to the Internal BGP AS Migration Features, 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 Alias" over a long period of time for reasons of not only operational simplicity of the network, but also reduced reliance on that feature during the ongoing lifecycle management of software, features and configurations that are maintained on the network.

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 ASN's. 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, Stephane Litkowski, Terry Manderson, David Farmer, Jaroslaw Adam Gralak, Gunter Van de Velde, and Juan Alcaide for their 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.

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. Appendix: Implementation report

As noted elsewhere in this document, this set of migration features 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.

11. References

11.1. Normative References

- [RFC2119] Bradner, S., "Key words for use in RFCs to Indicate Requirement Levels", BCP 14, RFC 2119, March 1997.
- [RFC5398] Huston, G., "Autonomous System (AS) Number Reservation for Documentation Use", RFC 5398, December 2008.

11.2. 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>.
- [CISCO] Cisco Systems, Inc., "BGP Support for Dual AS Configuration for Network AS Migrations", 2003, <http://www.cisco.com/c/en/us/td/docs/ios-xml/ios/iproute_bgp/configuration/xr-3s/asr1000/irg-xr-3s-asr1000-book/irg-dual-as.html>.
- [JUNIPER] Juniper Networks, Inc., "Configuring the BGP Local Autonomous System Attribute", 2012, <http://www.juniper.net/techpubs/en_US/junos13.3/topics/concept/bgp-local-as-introduction.html>.
- [RFC4271] Rekhter, Y., Li, T., and S. Hares, "A Border Gateway Protocol 4 (BGP-4)", RFC 4271, January 2006.
- [RFC4456] Bates, T., Chen, E., and R. Chandra, "BGP Route Reflection: An Alternative to Full Mesh Internal BGP (IBGP)", RFC 4456, April 2006.
- [RFC5065] Traina, P., McPherson, D., and J. Scudder, "Autonomous System Confederations for BGP", RFC 5065, August 2007.
- [RFC6996] Mitchell, J., "Autonomous System (AS) Reservation for Private Use", BCP 6, RFC 6996, July 2013.

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
Apple, Inc.
1 Infinite Loop
Cupertino, CA 95014
US

Email: samante@apple.com