

Errors Announcing 32-bit ASNs in BGP Routes

Riad Mazloun
UPMC Sorbonne Universités

Jordan Augé
Telecom ParisTech

Dario Rossi
Telecom ParisTech

Timur Friedman
UPMC Sorbonne Universités

Abstract—We present evidence of a phenomenon where there is large scale misconfigurations of BGP routers starting in 2007 and continuing to the present date. This has to do with the change from 16-bit to 32-bit autonomous system (AS) numbering and the employment of the fictitious AS, AS23456. This error strongly biases estimations of the number of ASes performing a type of inter-domain routing called multi-exit routing. In data from 2010, over half of apparent cases are in fact false. We show how to detect this error and obtain a truer picture of the extent of multi-exit routing.

I. INTRODUCTION

BGP (Border Gateway Protocol) is the globally implemented inter-domain routing protocol. It glues together the thousands of networks (autonomous systems or ASes) that form the Internet. An AS learns how to route its traffic toward destinations in other ASes through BGP. Some BGP routers publicly share their routing tables and route updates. Projects such as Route Views [1] and RIPE-RIS [2] collect this information and make it available to the community. BGP feeds are the main source of data for most scientific research on inter-domain routing and AS relationships. One of the attributes carried by BGP routes is the *AS Path*, which describes the sequence of ASes to be traversed toward the destination *prefix*.

There are some filtering and sanitizing steps on BGP feeds that the scientific community performs before using the datasets, such as considering only routes that seem to be stable ([3], [4] for instance), or by removing routes that contain loops [4]. Recent research on spurious BGP routes states that one of the public BGP routers has announced hundreds of thousands of BGP routes that do not exist in reality [5].

We point in this work at another source of bias that to our knowledge has not been reported previously in the scientific literature, where we identify BGP routes that carry malformed AS Paths because of the employment of AS23456 (also called AS_TRANS) which is used solve the transition from 16-bit ASNs to 32-bit ASNs. We show how this phenomenon highly impacts our observations of what we call multi-exit routing. Multi-exit routing happens when an AS uses multiple *next-hop* ASes toward a given destination prefix. Figure 1 illustrates the general case when there are two next-hops: routers in X announce two different routes ρ and ρ' toward a destination prefix p , each route having a different next-hop AS, Y and Y' respectively. We define and then quantify multi-exit routing in [6], where we also show that ASes that do such routing seem to violate some commonly used routing assumptions. We applied in this work our methodology to snapshots over ten years to




Fig. 1. Multi-exit routing general case

quantify ASes that are doing multi-exit routing, which allowed us to discover yet another source of artifacts in the BGP feeds that is caused by the employment of AS_TRANS. We show how it highly impacts quantification of multi-exit routing where the number reaches its peak in 2010 with a more than doubled number of multi-exits caused by spurious instances. We also explain how this phenomenon affects research making use of BGP feeds.

II. TRANSITION TO 32-BIT ASNs

Each AS is identified by a *unique AS Number*, or ASN. ASNs used to be 16-bits numbers. Then, as the number of ASes has increased, this numbering is no longer enough to identify all of the ASes. In 2007, the move was made to 32-bit ASNs (RFC 4893 [7]). In this representation, all ASNs are encoded on 32-bits, even those that are 16-bit numbers. However, not all BGP routers can handle AS Paths when they contain 32-bit ASNs. RFC 4893 introduces AS23456 (also called AS_TRANS) to provide a solution that assures compatibility between old BGP routers and AS Paths containing 32-bit ASNs. Whenever a new BGP router forwards routes to an old one it rewrites the AS Path using only 16-bit numbers. ASNs that can be represented on 16-bits are directly written as a 16-bit number, but all those that cannot are replaced by the number 23456 at each occurrence, which is a 16-bit number. In order to preserve the original AS Path that contains the real ASNs, the router writes it down into a new BGP attribute, which is the AS4-PATH. AS4-PATH is defined as an *optional, transit* BGP attribute, which means that when a router does not recognize it, it forwards it to its neighbors as it is, guarantying that it will not be lost. The old BGP router treats the AS Path as any normal one and may forward the route to its neighbors. Whenever this route reaches a new BGP router, this router will extract the original AS Path from the AS4-PATH attribute, rewrites the AS Path in 32-bit encoding and replaces each occurrence of 23456 with the corresponding 32-bit ASN.

Collaboration through the LINC laboratory. Full institutional affiliation of UPMC Sorbonne Universités authors: Sorbonne Universités, UPMC Univ Paris 06, UMR 7606, LIP6, F-75005, Paris, France.

III. AS_TRANS ARTIFACTS

If the method described earlier is implemented correctly then any new BGP router should be reporting AS Paths that do not contain AS23456¹. What we have found is that some of the routers that provide public BGP feeds clearly understand 32-bit ASNs as they report paths that contain such ASNs, however some of those same routes carry also AS23456². Then we looked at AS Paths of different routes toward the same destination. In some cases, one of the AS Paths has an extra AS hop compared to the other, which is AS23456, located between the AS that originated the route (which has a 32-bit ASN) and the AS after it. In other cases, we saw two AS Paths, one having AS23456 as the origin AS followed by an AS (which also has a 32-bit ASN), and one having this AS as the origin AS followed by AS23456. Both of the previous cases make us think that there is a misconfiguration somewhere in the AS that originated the route. In consequence, we may observe false multi-exits where there are two next-hops, one is the real AS, as seen from a BGP route, and the other is AS23456, as seen from another BGP route.

We considered BGP routes from Route Views and RIPE-RIS, the two commonly used public BGP data sources. A snapshot at a given instant is obtained by extracting all of the BGP routes that were used at that instant. We took snapshots on the first of August of each year from 2005 to 2014.

The total height of the bars in Figure 2 represent the number of ASes doing multi-exit routing that are observed after applying the method in [6] on the ten snapshots. The first remark on this figure is the spike in number of ASes in 2010. After investigation, we found first that almost all ASes are observed only in the 2010 snapshot. Then we remarked that most of them seem to be small ASes that have a couple of prefixes. Those ASes are distributed all around the world (we looked at the countries where they exist) and they are the origin of the routes that form the multi-exits (those routes concern their own prefixes). At last, we found that those cases are formed by two next-hop ASes, one of them is always AS23456 and the other one is the same AS that comes right after 23456 in the first AS Path, which means that they are false multi-exits.

We remove all routes that contain AS23456 in their AS Path from the dataset, then we do again the quantification of multi-exit ASes (which is represented by hashed bars in Figure 2). We can see, looking at white bars, that there were no observations related to AS23456 before 2007, then increase reaching a peak in 2010, and fall after that. We recall ASN assignment policy to understand what happened. If an network wants to become an autonomous system it has to do a request to obtain an ASN. The ASN assignment policies³ state that 32-bit ASNs were available since 1st January 2007 upon request, which explains that no observations were before this year. Then from 1st January 2009, 32-bit ASNs were assigned by default (except for networks that ask for a 16-bit ASN). After




Fig. 2. Fraction of ASes falsely observed to do multi-exit routing

that, since 1st January 2010, there were no distinction and ASNs were assigned from 32-bit pool. The peak in 2010 can be explained that those ASes were new ASes that were assigned 32-bit ASNs, and perhaps, since the AS is a new AS, some misconfiguration can appear, but do not last long time as they are solved. The number of ASes that are observed to do multi-exit routing in 2010 decreases to *less than the half* after removing routes containing AS23456. Interestingly, there are still false multi-exits in 2014 indicating that there are still misconfigured BGP routers. For the first look, such artifact seem to affect only multi-exit quantifications, but a second look shows that these routes can affect research where the inexistent link between an AS and AS23456 can be seen as an alternative route. Such data is employed by research on identifying the origin of routing events for instance, or on inferring alternative routes between two ASes in cases of network changes (we omit citing examples of such research due to lack of space), which means that indeed research making use of BGP feeds should be aware of those routes.

IV. CONCLUSION

We described in this paper one of the artifacts of 32-bit representation of AS Numbers that can be observed in public BGP feeds. Through observations of multi-exit routing, we show how greatly it biases quantification results (53% of the observations in 2010 are false). Finally, we explain how those routes can bias also research making use of BGP feeds.

REFERENCES

- [1] University of Oregon Route Views Project. <http://www.routeviews.org/>.
- [2] RIPE-RIS. <http://www.ripe.net/data-tools/stats/ris/>.
- [3] X. Dimitropoulos, D. Krioukov, M. Fomenkov, B. Huffaker, Y. Hyun, k. claffy, and G. Riley, "AS relationships: inference and validation," *ACM SIGCOMM CCR*, 2007.
- [4] M. Luckie, B. Huffaker, A. Dhamdhere, V. Giotsas, and k. claffy, "AS relationships, customer cones, and validation," in *Proc. IMC*, 2013.
- [5] M. Luckie, "Spurious routes in public bgp data," *ACM SIGCOMM CCR*, 2014.
- [6] R. Mazloum, M.-O. Buob, J. Auge, B. Baynat, D. Rossi, and T. Friedman, "Violation of interdomain routing assumptions," in *Proc. PAM*, 2014.
- [7] Q. Vohra and E. Chen, "BGP support for four-octet Autonomous System (AS) number space," Internet Engineering Task Force, RFC 4893, 2007.

¹Unless the route has been aggregated.

²To see how frequent this ASN is observed, we refer the reader to Hurricane Electric website (<http://bgp.he.net/AS23456>). We can see in this web page a list of hundreds of prefixes that seem to be originated by this AS and a list of ASes that seem to be interconnected with it.

³We refer the reader to <http://www.ripe.net/ripe/docs/ripe-525> and <http://www.apnic.net/policy/asn-policy#6.3> for more details.