An Analysis of BGP Multiple Origin AS (MOAS) Conflicts

Xiaoliang Zhao, Dan Pei, Lan Wang, Dan Massey, Allison Mankin, S. Felix Wu,Lixia Zhang

Abstract—

This paper presents a detailed study of BGP Multiple Origin AS (MOAS) conflicts observed in the Internet. A MOAS conflict occurs when a particular prefix appears to originate from more than one AS. We analyzed data from archived BGP routing tables over 1279 days. Most of the conflicts were short-lived, lasting only a small number of days. The potential causes for the MOAS conflicts and impact on BGP fault-tolerance are discussed in detail.

I. Introduction

This paper presents a detailed analysis of BGP[12] routes that appear to originate from multiple Autonomous Systems. The Internet is made of thousands of Autonomous Systems, loosely defined as a connected group of one or more IP prefixes which have a single and clearly defined routing policy[11]. BGP[12] is the standard inter-AS routing protocol. A BGP route lists a particular prefix (destination) and the path of ASes used to reach that prefix. The last AS in an AS path should be the origin of the BGP routes. A Multiple Origin Autonomous System (MOAS) conflict occurs if a prefix appears to originate from more than one ASes. More precisely, suppose prefix d is associated with AS paths $asp_1 = (p_1, p_2, \ldots p_n)$ and $asp_2 = (q_1, q_2, \ldots q_m)$. We say a MOAS conflict occurs if $p_n \neq q_m$.

In an effort to improve the fault-tolerance and security properties of the BGP routing protocol, we have been measuring the behavior of BGP. The MOAS conflicts are interesting to us for a number of reasons. First, RFC 1930[11] recommends that a prefix should originate from a single AS, but MOAS conflict may occur for a limited number of valid reasons. Second, MOAS conflicts could be the result of a fault or an attack, where a BGP router falsely originates routes to some other organization's prefixes. We would like to understand the characteristics of valid MOAS conflicts due to operational needs and invalid MOAS conflicts caused by faults.

MOAS conflict data was obtained from Internet routers and was analyzed based on the total number of conflicts, duration of the conflicts, and the prefix length. Both the number of MOAS conflicts and distribution of the duration of MOAS conflicts were different than what we anticipated. This paper presents the measurement results and analysis of potential causes.

The remainder of the paper is organized as follows. Section II reviews the related work. Section III describes the methodology we used to collect and process the data. Section IV presents the MOAS conflict data. Section V and VI provide detailed analysis of the results and our explanations of the results. Section VII discusses the implications of this work and summarizes the paper.

II. Related work

MOAS conflicts have also been observed by other researchers, but no one has considered the problem in detail. The most relevant work comes from Geoff Huston's BGP Table Statistics website[9]. Starting on 2/18/2001, this site began tracking a daily count of MOAS conflicts[1] using data from some ISPs and from the Oregon Route Views Server. On 04/19/2001, the website switched to tracking MOAS conflicts on a bi-hourly instead of daily basis. However, the BGP Table Statistics work provides only a basic count of MOAS conflicts and no further explanations or analysis is offered.

The MOAS conflict issue has also been discussed within the IETF. RFC 1930[11] recommends that a prefix should belong to only one AS. If this recommendation was followed, MOAS conflicts would not occur, with the possible exception of a few unique cases discussed further in Section VI-D. Berkowitz[13] discussed the potential causes of

This material is based upon work supported by the Defense Advanced Research Projects Agency (DARPA) under Contract No DABT63-00-C-1027. Any opinions, findings and conclusions or recommendations expressed in this material are those of the authors and do not necessarily reflect the views of the DARPA.

xzhao@unity.ncsu.edu, peidan@cs.ucla.edu, lanw@cs.ucla.edu, masseyd@isi.edu, mankin@isi.edu, wu@cs.ucdavis.edu, lixia@cs.ucla..edu

[1] Huston uses the term "multiple-origin prefixes" in place of our term "MOAS conflicts"

MOAS conflicts. However, the discussion is not complete and no implications of MOAS conflicts are analyzed.

III. Methodology

The BGP route for a prefix (destination) includes an AS path. The last AS along the path to the prefix is considered to be the origin AS. We examined the AS paths that led to the same prefix but ended in different origin ASes.

We primarily used data from the Oregon Route Views server [8] to obtain the BGP routes and AS paths used in this study. Currently, the Oregon Route Views server peers with 54 BGP routers in 43 different ASes. Each peer exports its BGP routing table to the Route Views server.

The Oregon Route Views data is particularly attractive because it provides data from a number of different vantage points. The data obtained from a particular local point, such as in an individual ISP, may show a smaller number of MOAS conflicts since fewer potential AS paths may be visible at that point in the network. For example, at a randomly selected time, the Oregon Route Views server observed 1364 MOAS conflicts, but three other individual ISPs observed 30, 12, and 228 MOAS conflicts during the same period. This only means that fewer MOAS conflicts were visible to these ISPs and even the number of MOAS conflicts observed from the Oregon Route Views Server may underestimate the total number of MOAS conflicts.

To obtain a relatively complete view, we used archived Oregon Route Views data from both NLANR[2] and PCH.net [3]. NLANR archived the Oregon Route Views data on a daily basis from 11/08/1998 to 03/16/2001. PCH.net archived the Oregon Route Views data on a daily basis from 03/16/2001 to the present. The MOAS conflicts are identified by prefixes only no matter whether a MOAS conflict was conflicted by same set of origin ASes or the conflict was continuous.

Note that AS sets did not play any meaningful role in our study. An AS path typically consists of the sequences of AS numbers used to reach prefix, but due to factors such as aggregation, the AS path may also contain AS sets as well as AS sequences. Out of of over 100K prefixes observed, roughly 12 routes ended in AS sets and these 12 routes were not included in the study.

IV. Results

The total number and durations of MOAS conflicts deviated substantially from our expectation. Based on these results, we believe the nature of these conflicts differs from what one might expect based on documents such as [11].

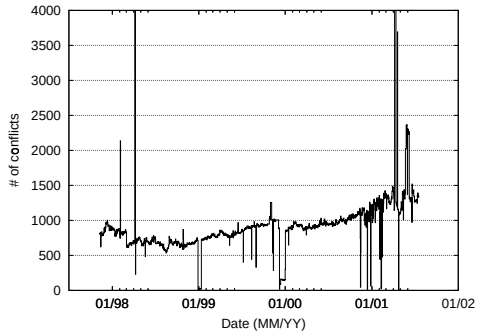

Fig. 1. The number of MOAS conflicts from 11/1997 to 07/200

Year	Median of MOAS conflicts	Increasing rate
1998	683	
1999	810.5	18.7%
2000	951	17.3%
2001	1294	36.1%

Fig. 2. Median of MOAS conflicts per year

A. Total Number of MOAS Conflicts

Figure 1 shows the total number of conflicts from 11/08/1997 to 07/18/2001 [2]. Overall 38225 conflicts were observed over 1279 days. The median number of MOAS conflicts for each year are listed in Figure 2. There is an increase from 683 conflicts in 1998 to 1294 conflicts in 2001.

B. Duration of MOAS Conflicts

Figure 3 shows the duration of MOAS conflicts, based on the data (Figure 1). Figure 3 shows that most of the conflicts are short-lived. 13730 out of 38225 conflicts appeared only once and lasted less than one day. 11358 of these one-time conflicts can be attributed to a configuration fault that occurred on April 7th, 1998. Excluding the one-time conflicts, the expectation of the duration is 30.9 days. Taking into account that many other short-lived conflicts might also be due to faults, we considered the data set which contains only conflicts whose duration is greater than 9 days (a total of 10177 conflicts). For these conflicts, the expectation of the duration is 107.5 days with 1002 conflicts lasted longer than 300 days. Figure 4 lists the expectation of the duration from the different data sets. The longest duration was 1246 days out of a possible 1279 days and 1326 conflicts were still ongoing as of the date the paper was written.

The duration of an individual conflict counts the total number days of the conflict was in existence, regardless of

[2]The number of conflicts reached its peaks of 11842 on 04/07/1998 and 10226 on 04/06/2001.

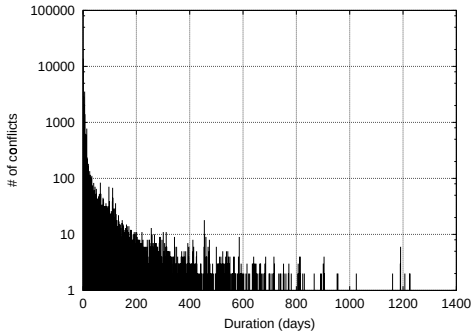

Fig. 3. Duration of MOAS conflicts

Expectation (days)	Measured data set
30.9	longer than 0 day
47.7	longer than 1 days
107.5	longer than 9 days
175.3	longer than 29 days
281.8	longer than 89 days

Fig. 4. Expectation of the duration of MOAS conflicts

whether the conflict was continuous and whether the same ASes were involved.

The results seem a little surprising if one assumes that multi-homing, discussed in Section VI-B, is the major reason for the MOAS conflicts. Multi-homing would seem to imply that the MOAS conflicts should last longer than what is observed here and this is discussed further in Section VI-F.

C. Distribution of MOAS Conflicts

Figure 5 shows the distribution of conflicts among prefix length. The /24 (netmask of 255.255.255.0) attracts most of conflicts. This is not unexpected since /24 prefixes make up the bulk of the BGP routing table.

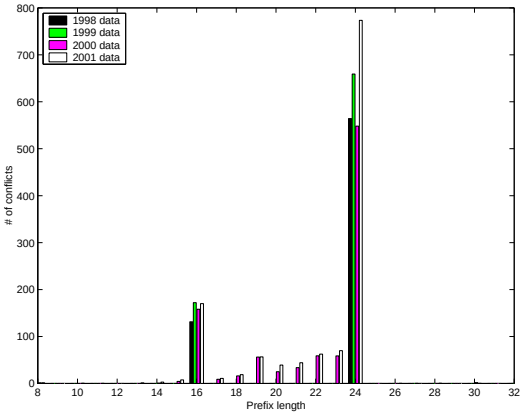

Fig. 5. Distribution among prefix length

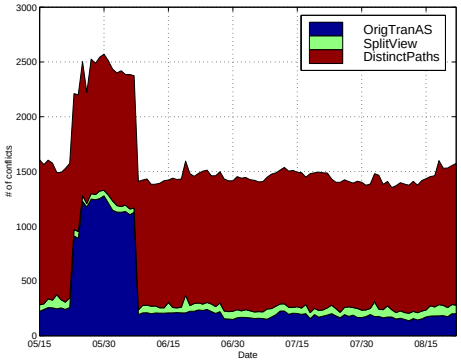

Fig. 6. Distribution of classes

V. Classification of MOAS conflicts

If a MOAS conflict occurs, prefix p will be associated with at least two different AS paths:

- $asp_1 = (p_1, p_2, \ldots p_n)$
- $asp_2 = (q_1, q_2, \ldots q_m)$

By definition, $p_n \neq q_m$ for a MOAS conflicts. In order to better understand the type of conflicts and the potential causes, we divided the MOAS conflicts into three classes based on relationships between the two AS paths.

OrigTranAS: $p_n = q_j \ (j < m)$.
In this case, AS p_n announces itself as the origin AS in asp_1 and announces itself as a transit AS in asp_2.
SplitView: $p_i = q_j \ (i < n, j < m)$.
In this case, AS p_i announces different routes to different neighbors.
DistinctPaths: $p_i \neq q_j \ (\forall i \in [1..n], j \in [1..m])$.
In this case, there are two totally different routes for the prefix d.

Instances of all three cases were observed and Figure 6 shows the number of conflicts for each class. In the OrigTranAS class, an AS acts as both the origin AS and a transit AS. In the SplitView class, a transit AS offers two different paths to the prefix and these paths end in different origin ASes.

The OrigTranAS and SplitView conflicts indicate that a single AS may advertise multiple paths to the same prefix. This is often because of the traffic engineering practices used at large ISPs. An AS might prefer that traffic to the same destination flow through different paths due to constraints such as geographical distances, link speed, or economic reasons.

In the DistinctPaths class, there are two completely disjoint AS paths for the same prefix. Figure 6 shows that the DistinctPaths class is dominant in the MOAS conflicts, which is not unexpected because BGP only choose one best route if no traffic engineering practice.

VI. Explanations and Implications

There are a number of possible explanations for MOAS conflicts. Unique cases such as exchange points, some forms of multi-homing, and faults all contribute to the MOAS conflicts. Each of these factors was observed in this study.

A. Exchange Point Addresses

One potential cause of MOAS conflicts involves the prefixes associated with exchange points (or equivalently, links connecting ASes). A prefix associated with an exchange point is directly reachable from all the ASes at the exchange point and each AS at the exchange point might advertise the prefix as if it comes directly from that AS.

However, exchange point prefixes make up a small percentage of the MOAS conflicts observed in this study. In the examined BGP data, 30 out of 38225 prefixes could be definitively identified as exchange point prefixes. Our analysis of exchange point prefixes may underestimate the total number of exchange point prefixes, but the number of exchange point prefixes remains relatively small even if our estimate is off by two orders of magnitude. All of these exchange point prefix conflicts lasted for long periods, consisting of most or all of the observation periods. These MOAS conflicts do not present a problem for packet forwarding since each AS originating the route can directly reach the prefix.

B. Multi-homing Without BGP

In some cases, multi-homing can occur without the use of BGP and this can result in MOAS conflicts. Suppose there is a link between two ASes, but the routing across this link does not use BGP (and instead relies on static routing or some IGP). From a BGP perspective, it appears as if one AS can directly reach prefixes belonging to the other AS.

Again one would expect these conflicts to be long lasting since static routes are likely to have a long lifetime. These MOAS conflicts could present a problem for packet forwarding if the links necessary to support the static routes fail.

C. Multi-homing with Private AS Numbers

To prevent AS number exhaustion, Haas [10] suggests that a multi-homed customer uses a private AS number which is mutually agreeable to all providers. This technique is called AS number Substitution on Egress (ASE). If deployed, this approach could produce MOAS conflicts because the private AS number should be stripped off by the upstream providers and the real origin information will be lost.

Based on discussions with network operators, we do not believe this technique is used widely in practice. These MOAS conflicts would not present a problem for packet forwarding since all upstream providers can reach the private AS. Furthermore, if the link to the private AS is lost, the corresponding BGP route will also be withdrawn.

Because the links using non-BGP routing mechanisms or private AS numbers are "hidden" to BGP, the pure BGP data can not tell whether or not a MOAS conflict is due to multi-homing without BGP or multi-homing with private AS number. However, by contacting individual ASes, we did confirm such occurrences.

D. Theoretical Causes

Other factors have the potential to cause MOAS conflicts, but these factors did not occur during our study. In particular, RFC 1930[11] notes that aggregation could result in routes that end in AS sets. But overall, we typically observed 12 prefixes which ended in AS sets and these AS sets were consistent with each other.

Anycast address would also create MOAS conflicts since an anycast prefix is intended to originate from multiple ASes. No prefixes in our study were identified as anycast addresses.

E. Faulty or Malicious Configurations

MOAS conflicts can also occur when an AS incorrectly originates routes to some other organization's prefixes. This could occur due to configuration errors or even intentional attacks. Often, the faulty AS does not have a route to the incorrectly originated prefixes and packets that use the incorrectly originated route will reach the faulty AS and then be lost.

Figure 1 shows several notable examples of MOAS conflicts caused by faults. The graph shows a large spike on April 7th, 1998 and AS 8584 was involved in 11357 out of 11842 conflicts that occurred during that day. Discussions on a network operators mailing list[4] indicated that AS 8584 falsely originated routes to those conflicted prefixes. Consequently, some ASes selected the incorrectly originated route. Packets sent along this incorrectly originated route would reach AS 8584 and would then be lost.

The graph also shows a large spike on April 10th, 2001 and the sequence (AS 3561, AS 15412) was involved in 5532 out of 6627 MOAS conflicts that occurred during that day. Based on the archived data from RIPE RIS [1], AS 15412 normally originates only 5 prefixes. However, on April 6th, AS 15412 suddenly originated thousands prefixes due to a configuration error[5].

On April 25th, 1997, a severe Internet outage occurred[7] when one ISP falsely de-aggregated most of the Internet routing table and advertised the prefixes as if they originated from the faulty ISP[6]. The falsely originated prefixes resulted in MOAS conflicts. These examples show that invalid MOAS conflicts do occur and can have serious impacts on Internet routing.

Faulty aggregation could also cause MOAS conflicts. In faulty aggregation, an AS advertises an aggregated prefix, even though some of more specific prefixes are not reachable by the AS. A MOAS conflict occurs if an aggregate route is also generated by some other AS. Packets that use the faulty aggregated route will travel to the faulty AS and then may not be able to reach all the more specific prefixes.

F. MOAS Conflict Durations and Potential Causes

With the exception of faults and intentional attacks, the possible explanations should have created long duration MOAS conflicts. MOAS conflicts for exchange point prefixes should remain as long as two or more ASes choose to advertise a route to the exchange point. The data confirmed this expected pattern and exchange point MOAS conflicts persisted for most, if not all, of the study. Multi-homing without BGP and multi-homing with Private AS numbers both require router policy configurations at two or more ASes and the resulting MOAS conflicts should persist for as long as the multi-homing policy remains in place. We expected that multi-homing policies (and the resulting MOAS conflicts) would occur over months, not days. But the data in Section IV shows a large number of short duration conflicts.

One possible reason for short-lived MOAS conflicts is that MOAS conflicts could occur during a transition period when a non-BGP customer switches from one provider to another. To guarantee the connectivity to the non-BGP customer, it is possible for both providers to originate the customer's prefix for a short period. Another possible and more likely reason for short-lived MOAS conflicts is router mis-configurations or other faults. These conflicts disappear when the faults are detected and corrected.

Overall, the duration can be a useful heuristic to distinguish between valid MOAS conflicts and invalid ones. However, such differentiation can not be accurate enough to be a solution to validate MOAS conflicts.

VII. Summary

The MOAS data presented in this paper can help in understanding the operational behavior of BGP. At a minimum, one would like to know if types of MOAS conflicts expected to occur actually match the type of conflicts actually occurring in the Internet. These results would indicate there are a large number of faults or large number of very short lived multi-homing policies.

From the standpoint of fault-tolerance and security, MOAS conflicts pose an interesting challenge. On the one hand, MOAS conflicts can occur for valid reasons, such as multi-homing without BGP and advertising routes to exchange points. On the other hand, router mis-configurations have also produced MOAS conflicts. Large scale network outages and other problems have been associated with MOAS conflicts. When a MOAS conflict is observed, we would like to be able to determine whether it is the result of a fault or a valid change in routing/multi-homing policy. Based on this MOAS data alone, we can not accurately differentiate a fault from a valid policy change, but we can utilize the MOAS analysis results as a valuable input to address BGP problems and we are investigating techniques for identifying invalid conflicts with a high degree of certainty.

VIII. Acknowledgments

We would like to thank a number of network operators and BGP routing experts for the advice. In particular, Randy Bush and Geoff Huston provided useful insights on the operation of large ISPs. Also, we would like to thank anonymous reviewers for their valuable comments.

References

[1] RIPE Routing Information Service, http://www.ripe.net/ripencc/pub-services/np/ris-index.html

[2] National Laboratory for Applied Network Research, http://moat.nlanr.net/Routing/rawdata/

[3] PCH.net, http://www.pch.net/documents/data/routing-tables/route-views.oregon-ix.net/

[4] B. Kroenung, "AS8584 taking over the internet", NANOG mailing list, msg00047, Apr. 7, 1998.

[5] J. Farrar, "C&W routing instability", NANOG mailing list, msg00209, Apr. 6, 2001.

[6] V. J. Bono, "7007 Explanation and Apology", NANOG mailing list, msg00444, Apr. 26, 1997.

[7] R. Barrett et al, "Routing Snafu Causes Internet Outage", ZDNet, April 25, 1997. http://www.zdnet.com/zdnn/content/inwo/0425/inwo0009.html

[8] University of Oregon Route Views Project, http://www.antc.uoregon.edu/route-views/

[9] G. Huston, "BGP table statistics", http://www.telstra.net/ops/bgp/as6447/bgp-multi-orgas.html.

[10] J. Haas, "Autonomous System Number Substitution on Egress", Internet Draft, Working in Progress, 2001.

[11] J. Hawkinson and T. Bates, "Guidelines for creation, selection, and registration of an Autonomous System (AS)", RFC 1930. 1996.

[12] Y. Rekhter and T. Li, "A Border Gateway Protocol 4 (BGP-4)', RFC 1771. 1995.

[13] H.Berkowitz, E.Davies and L. Andersson, "An Experimental Methodology for Analysis of Growth in the Global Routing Table", Internet Draft, Working in Progress, July, 2001.