Automatic Analysis and Classification of Multiple Origin AS (MOAS) Conflicts

Uli Bornhauser, Jan Luehr, Michael Schaefer, Thomas Trimborn, Zia Ul Huda
University of Bonn – Institute of Computer Science 4 – Friedrich-Ebert-Allee 144 – D - 53113 Bonn
{ub, luehr, schaefer, trimborn}@cs.uni-bonn.de, zia_ul_huda@yahoo.com

***Abstract*—Global routing is consistently affected by minor and major disruption, so called *anomalies*. Reasons for anomalies are multilateral, as is the resulting impact. Even though the majority of anomalies typically causes negligible implications, incidents may have global impact. In the worst case, they culminate in worldwide unavailability of popular address spaces [1].**

A well known and very frequently observed type of anomaly are ASs that originate foreign address space without permission. Traffic for the corresponding destinations may be redirected to the wrong AS, thus services and users may become unreachable. Technically, such anomalies become visible as *Multiple Origin AS (MOAS) conflicts*: Routing information for the same destination exists that is simultaneously originated by different ASs. However, even though every destination should originate from a unique AS, reasons exist which render a MOAS conflict legitimate.

At LCN 2011, we present the *MOAS Analyzer*, a tool to identify, analyze, and automatically classify MOAS conflicts. As far as we know, it is the first tool that implements automatic classification of MOAS conflicts.

***Index Terms*—Routing Protocols, BGP, Security, Prefix Hijacking**

I. Introduction

The Internet is a compulsory part of our daily life. On a high level of abstraction, it consists of several tens of thousands of interconnected – so called – *Autonomous Systems (ASs)*. ASs are networks of networks operated by a common operational authority with a specific routing policy. Traffic between these systems is forwarded on the basis of routing information that is exchanged via the *Border Gateway Protocol (BGP)* [2].

The basis for providing global services and connecting users to the Internet is a unique global address space. To avoid that addresses are ambiguous, global address spaces are assigned to AS operators by *Regional Internet Registries* (RIRs). By advertising the own address space through BGP into the inter-AS routing, services and users become globally reachable.

A. Motivation

As long as every globally used address is only announced by a unique *origin*, traffic for corresponding destinations can be delivered uniquely. However, because BGP as it is productively used today does not provide any kind of origin authentication, ASs may announce arbitrary address spaces in practice. From different points in the Internet, traffic for the corresponding destinations may be delivered to differed ASs [1], cf. figure 1. As consequence, services and users that belong to the address space may become at least from certain systems unavailable. According to the concept behind ASs [3], such *Multiple Origin Autonomous System (MOAS) conflicts* should never happen.

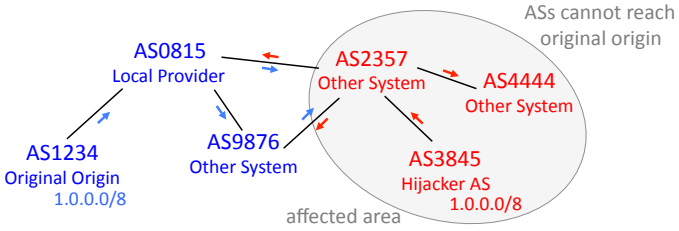

Fig. 1. If an AS hijacks the address space actively used by another system, network traffic from other systems may be redirected to the wrong destination.

Even if MOAS conflicts should never appear in theory, there are indeed some legitimate reasons in practice. For example, a conflict may be caused by two systems that keep a customer provider relationship, where the customer AS uses a dedicated sub-address space of its provider AS, cf. figure 2.a. Another simple and justifiable reason is anycast routing as applied for most of the DNS root servers [4], cf. figure 2.b. In these and similar scenarios, a conflict can be termed as *legitimate*. In contrast to legitimate conflicts, MOAS conflicts may also be *illegitimate*. Illegitimate MOAS conflicts usually appear if a system originates an address space as its own, it legally does not own. Even if purposeful attacks are an obvious explanation for such conflicts, the vast majority of illegal MOAS conflicts is caused by misconfigurations or typos.

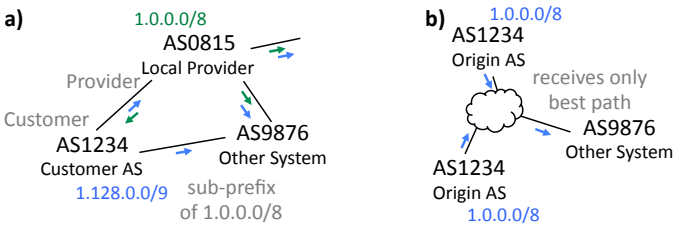

Fig. 2. MOAS conflicts may be caused by prefix delegation from a provider to a customer AS (a) or anycast as used for DNS root servers (b), for example.

B. Problem Description, Relevance, and Objectives

While desired MOAS conflict are only side effects of highly wanted and reasonable configurations, illegitimate conflicts are indications for routing disruptions. Even if most incidents only have a very limited scope and affect only a globally negligible

address space, they are highly problematic for the affected ASs and their services. In contrast to those rare incidents that affect a popular or large address space, cf. [1], for example, smaller events are not that obvious: As resulting connectivity problems often only concern some parts of the Internet, even an affected operator may not notice the phenomena directly. Nevertheless, especially in the long term, connectivity problems may have a significant impact on its business activities. This motivates to detect illegitimate MOAS conflicts as soon as possible.

Projects like BGPmon [5] and Cyclops [6] have proven that MOAS conflicts are relatively easy to detect in general. Being able to recognize MOAS conflicts, the next conceptual step is classification. Our preliminary studies have shown that manually classifying MOAS conflicts is no alternative. Exemplary analyses of the first quarter of 2010 have revealed over 200,000 MOAS conflicts during that period. Even if around 60,000 of these conflicts are virtually stable and thus most likely legitimate, the remaining amount of MOAS conflicts is still too big for manual analyses and classification. An automatic analysis and classification process for MOAS conflicts is needed. This is the main aspect where the *MOAS Analyzer* presented at LCN 2011 comes into play.

II. The MOAS Analyzer

The MOAS Analyzer is a detection, analysis, and classification tool for MOAS conflicts that is developed at the University of Bonn. Currently being under development, it is available for a restricted user group at *https://moas.net.cs.uni-bonn.de/*. The main objective is to implement an automatic, realtime capable analysis of MOAS conflicts for scientific purposes. The MOAS Analyzer is work in process. Access to the web-interface may be requested by researchers and other interested parties.

A. MOAS Conflict Detection

As already mentioned above, the detection of MOAS conflicts is a functionality that is already public available [5], [6]. Similar to these tools, the MOAS Analyzer also gathers data on MOAS conflicts from publically available routing information. Technically, our tool is capable to use raw data provided by the *RIPE NCC Routing Information Service (RIS)* and the Route Views projects, cf. [7], [8]. For the presentation, we only made use of RIS data to avoid possible inconsistencies, for example due to the times zones used for timestamps.

The objective to provide data for scientific purposes borders the conflict detection of the MOAS Analyzer significantly from concepts like BGPmon and Cyclops: In principle, both tools focus to report new MOAS conflicts that affect certain prefixes in real time. In contrast to this, the MOAS Analyzer provides a full overview on conflicts that appeared in the Internet since the beginning of data acquisition. Instead of grouping them by address prefixes, it allows for grouping them by aspects like time, origin AS, and participating groups of origin ASs, for example. Based on the different aggregation criteria, an in-depth understanding for the frequency and characteristics of MOAS conflicts in the Internet can be gained. The basic architecture of the MOAS analyzer is illustrated in figure 3.

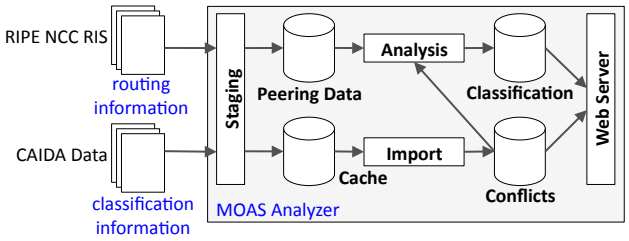

Fig. 3. The basic system architecture of the MOAS Analyzer.

B. Automatic MOAS Conflict Classification

Even if a global overview on MOAS conflicts may already be a helpful feature, the main strength on the MOAS Analyzer is the automatic analysis and classification of conflicts. As already explained above, MOAS conflicts can be legitimate and illegitimate. Common properties unify conflicts that depend on the same underlying reason. This observation provides a basis for classifying conflicts automatically.

Our research focus is the development of indicators that allow for classifying incidents causing MOAS conflicts as legitimate and illegitimate. We use prototypic implementations and analyses in order to refine the classification process. The indicator development process in realized in the MOAS Analyzer. The concepts and their implementations are currently evolved and extended. Results we already achieved reveal that a considerable number of conflicts can already be classified automatically. The interface of our tool allows users to evaluate results of the automatic analyses and assess the relevance of many conflicts. Our final goal is to classify as many MOAS conflicts as possible at least in such a way that the reason for the conflict are revealed or – if this is not possible – that at least plausible explanations are assessed. The current intermediate state on this way is what is demonstrated at LCN 2011.

III. Significance and Scope of the Demo

A. Significance

Incidents in the past have impressively shown that illegitimate announcements of routing information may significantly disturb the global routing, cf. [1]. As this problem has also be understood by router manufacturers and network operators, the IETF has also begun to standardize BGP extensions to identify illegitimate path announcements [9]. However, currently only first prototype implementations are available. Also the basic infrastructure, the *Resource Public Key Infrastructure* [10], is not globally deployed yet. Experiences with other extensions show that it will certainly require at least two to five years [11] until a comprehensive origin validation is implemented. For at least this period, identifying illegitimate MOAS conflicts through an analysis of global BGP information seems to be an important facet to improve the robustness of the inter-domain routing. The MOAS Analyzer is, for all that we know, the first tool that allows an automatic classification of MOAS conflicts as illegitimate or legitimate ones. As data is not only provided for ongoing conflicts but also for conflicts observed in the past, a significant relevance for practice and research is given.

B. Methodology and Technical Details

The MOAS analyzer and its automatic conflict classification presented at LCN 2011 is based of the following indicators:

1) Prefix Delegation: A typical reason for MOAS conflicts is that provider ASs *delegate* a sub-range of their address space to certain customer systems. As the result of such a delegation, provider ASs originate certain IP prefixes while their customer ASs originate equally or more specific (sub-)prefixes. Identified MOAS conflicts can be classified as *probably legitimate* if the involved ASs keep provider-customer relationships that indicate simple or multi prefix delegation. For this indicator, we make use of public AS-relationship data provided by *The Cooperative Association for Internet Data Analysis (CAIDA)*, cf. [12]. Our primality studies showed that around one third of the conflicts can be classified by this indicator.

2) AS Sibling Relationship: MOAS conflicts that are caused by ASs under the same administration are either intended or a local matter of the operator. ASs under the same administrative control are called *siblings*. To trace back a MOAS conflict to siblings and classify them as *probably legitimate*, our tool uses registry information provided by the RIRs. We compare data like registered organization names, contact persons, addresses, and similar information to identify siblings. In addition to this, we locally keep data on manually identified siblings that do not result from whois information such as sub- and supersidiaries.

3) Prefix Aggregation: BGP as specified in [2] also allows ASs to aggregate routing information. Aggregation also results in MOAS conflicts if parts of the original prefixes are globally visible. However, if the systems that originate the sub-address spaces are all visible in the origin AS-set of the super-address space, cf. [2], conflicts are most probably legitimate. To verify this property, no external data is needed.

4) Announcements of Private ASs: Network operators may use private ASs to border an administrative zone even if no public AS number is registered. As private AS numbers are not globally unique, they should never be announced globally [2]. Having a private AS, the easiest solution to avoid private AS numbers is to remove it from the AS-path at the upstream provider. If this functionality is not correctly implemented at any upstream provider, MOAS conflicts where a private AS number participates appear. Such conflicts are illegitimate. Checking conflicts for private AS numbers realizes a simple indicator to identify this kind of misconfiguration.

5) Typos: Simple typos, for example in the number of the AS a router belongs to, are another reason for illegitimate MOAS conflicts. Checking AS numbers and advertised address spaces for permutations and missing digits indicate such misconfigurations. By comparing the address spaces different ASs advertise, such typos can be identified.

IV. Technical Requirements and Concept

Detection as well as classification results are presented in a web-based front-end. It allows for navigating through detected MOAS conflicts, examining their properties, and reviewing their classifications. As the underlying processing takes place in real-time, it is performed on an appropriate lab system. This system is hosted at the University of Bonn. The system also provides the web-based front-end.

Presentation at the LCN demo session shall be realized via client systems. A client system is a common notebook that has access to our lab server and its services. We set up two or three client systems and allow attendees of the LCN 2011 to individually explore the data. In addition to this, we set up a client system connected to an external display to explain the concepts and our tool.

V. Conclusion

MOAS conflicts arise as a result of legitimate configurations. However, they are also a basic indication for routing problems caused by systems that originate an address space they are not authorized to: MOAS conflicts may be legitimate and illegitimate. Being able to differ illegitimate from legitimate MOAS conflicts could help to identify routing disruptions timely and improve the robustness of BGP.

As MOAS conflicts affect a high number of prefixes today, automatic detection and classification is necessary. Besides the real-time detection of MOAS conflicts on the basis of globally available routing information, this functionality is provided by the MOAS Analyzer. Being under development, our tool prototypically implements classification indicators. These indicators presented at the LCN 2011 provide the basis for further developments and refinements. In the long term, we seek for being able to give a reasonable explanation for most MOAS conflicts that can be observed.

References

[1] U. Bornhauser and P. Martini, "About Prefix Hijacking in the Internet," in *Accepted for the Proceedings of the 36th IEEE Conference on Local Computer Networks (LCN 2011)*. IEEE Computer Society, October 2011.

[2] Y. Rekhter, T. Li, and S. Hares, "A Border Gateway Protocol 4," January 2006, RFC 4271.

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

[4] Internet Corporation for Assigned Names and Numbers (Coordination), "DNS Root Servers," http://www.root-servers.org/.

[5] "BGPmon." [Online]. Available: http://bgpmon.net/

[6] "Cyclops." [Online]. Available: http://cyclops.cs.ucla.edu/

[7] RIPE NCC, "Routing Information Service (RIS)," http://ripe.net/data-tools/stats/ris/.

[8] University of Oregon, "Route Views," http://www.routeviews.org/.

[9] M. Lepinski and S. Kent, "An Infrastructure to Support Secure Internet Routing," February 2011, draft-ietf-sidr-arch-12.

[10] R. Bush, "RPKI-Based Origin Validation Operation," July 2011, draft-ietf-sidr-origin-ops-10.

[11] ——, "BGPsec Operational Considerations," March 2011, draft-ietf-sidr-arch-12.

[12] The Cooperative Association for Internet Data Analysis (CAIDA), "AS Rank: Autonomous System Ranking," http://as-rank.caida.org/.