Reputation Rating Algorithm for BGP Links

Hospice Alfred Arouna[1(✉)], Lionel Metongnon[1,2], and Marc Lobelle[2]

[1] Université d'Abomey-Calavi, Abomey-Calavi, Benin
hospice.arouna@uac.bj
[2] Université Catholique de Louvain, Ottignies-Louvain-la-Neuve, Belgium
{lionel.metongnon,marc.lobelle}@uclouvain.be

Abstract. BGP is a dynamic protocol used by Autonomous Systems (AS) constituting the Internet to exchange information in order to set up or remove links between AS. It takes into account the status of existing links and the internal policy of the AS. New links can be either legitimate or malicious. Having an objective way to detect route-leaks and/or route-hijacks could be a good starting point for deciding to accept or reject newly advertised links. In this work, an algorithm has been developed to evaluate link reputation on the basis of metrics. The work proceeded in three steps: first, BGPStream is used to overcome difficulties related to the collection of BGP record files from various collectors and projects. In the analysis phase (second phase), the algorithm is applied on collected data. The final phase is to visualize the results with a modified version of BGPlayJs to display the links reputation by coloring them from green to red. This algorithm could be used for baseline leak/hijack detection.

Keywords: BGP · Algorithm · Link · Reputation · Visualization

1 Introduction

BGP is the *de facto* inter-domains routing protocol used to maintain and exchange routing information on Internet. AS_PATH is one of the most important attributes of BGP [13]. BGP peers implicitly trust each other [2,9]. Any AS may announce and/or update any prefix (*i.e.*, IP address blocks) even if this prefix has already been assigned and/or announced by another AS or is a bogon (illegal prefix) [7]. This implicit trust is the fertile breeding ground for malicious activities, censorship and configuration errors. Various solutions have been proposed to improve the security of the protocol, but most of them remained at the project stage [2,9]. Since then, visualization solutions with their user-friendly analysis approach stand out and get popular. However most of those graphical solutions are not algorithm-based and are useless to detect if a link is legitimate or malicious [1]. The goal of this paper is to develop a BGP link reputation algorithm. For each link on each AS_PATH, a reputation will be computed based on other metrics. The rest of the paper is organized as follows.

V. Odumuyiwa et al. (Eds.): AFRICOMM 2017, LNICST 250, pp. 352–357, 2018.
https://doi.org/10.1007/978-3-319-98827-6_33

Section 2 summarizes related work. Section 3 explain our methodology. Section 4 give details about our approach. Section 5 presents tests cases and the results. Section 6 introduce African perspective while Sect. 7 concludes the paper.

2 Related Work

Studies [4,5,15] are related to ASes trust. Konte et al., [10] focused on malicious ASes fully dedicated to supporting cybercrime while Sankar et al., [14] developed a framework to detect suspicious deviation in the `AS_PATH` between a source and destination. The linkrank project from UCLA was using only one metric. However, Lad et al. show in [11] the need to have more than one metric to evaluate link reputation. The linkrank project from UCLA is a rare example of graphical tool that is algorithm based and allows BGP routing dynamics observation [11]. Although, this project was abandoned since July 2011, the idea has been taken up during the linkrank challenge from the Center for Applied Internet Data Analysis (CAIDA) BGP hackathon [3]. The objective of CAIDA linkrank, was to find a solution to help discriminate legitimate from malicious link information. The algorithm described in this paper combines different metrics to determine link reputation and clear the CAIDA linkrank challenge.

3 Methodology

This paper presents a proof of concept for link reputation computation. With the huge amount of BGP data available, *BGPStream* [12] is an efficient tool for processing large amounts of distributed and/or live BGP measurement data, enabling rapid prototyping and large-scale monitoring applications building. The JavaScript version of *BGPlay* [6]; *BGPlayJS* is used to visualize the algorithm results. With those 2 tools already available, our solution can focus on the algorithm part. In phase 1, *BGPStream* is used to collect BGP data during a defined time interval. The last phase is about results where a modified version of *BGPlayJS* is used to display the reputation of links by coloring them from green to red. Red means that link has bad reputation while green means the opposite. Phase 2 is the principal part where our intuition has been implemented in mathematical functions and algorithms. Our implementation has been evaluate with existing tests cases.

4 Approach

This study is based on the intuition that links with good reputation have a majority of BGP announcements with little amount of bogon and that most peers use those links. From this intuition, three metrics has been defined and combine to compute the link reputation at a specific time: (i) `link stability` (number of announcements vs withdrawn), (ii) `bogon degree` (number of bogon per link) and finally (iii) `link sensibility` (number of prefixes - new and/or

modified - per link between t_2 and t_1). Those metrics have been converted first to mathematical functions and then as algorithms. For example, in an oriented graph $G = (V, E)$ where $V = \{v | v \in path_i\}$ (nodes are ASN on specific `AS_PATH`) and $E = \{\langle a, b\rangle | \langle a, b\rangle \in path_i\}$ (edges between ASN), the `bogon degree` is given by the following formula:

$$bogons_t(\langle a, b\rangle) = \begin{cases} \sum_{i=0}^{n} 1, & \text{if } prefix_{ti} \bigcap bogon_{set} \\ 0, & \text{otherwise} \end{cases} \quad (1)$$

where n is the total number of prefixes on link $\langle a, b\rangle$ for event t. For each event (BGP annoncement) at time t and for each prefix i seen on link $\langle a, b\rangle$ between ASes a and b, a positive counter is incremented by one if $prefix_{ti}$ is present in bogon database ($bogon_{set}$ from *Team Cymru*). With `link stability`, link behavior is observed on each event by removing the number of withdrawn prefixes from the number of announced prefixes. `link sensibility` metric is the one used by the linkrank project from UCLA. This metric helps us to observe the link utilization by peers and is taken as it is [11].

BGPStream helped us filter collected BGP events related to our tests cases. For each of those test cases, a specific list of BGP events is available. For the purpose of our study, a data structure called `record` has been added to each BGP event result. Each `record` contains a list of `links` which represents each link observed on the current AS_PATH. In each `link` data structure, elements like `prefixes_list` and `as_path_list` are present.

```
   Input: events list
   Output: link_reputation
 1 while we still have event do
 2     foreach record in event do
 3         foreach link in record do
 4             link_stab ← Stability(link)
 5             link_bogon ← Bogons(link)
 6             link_sens ← Sensibility(link)
 7             link_reputation
               ← 1 + link_stab − link_bogon + link_sens − link_stab^link_bogon
 8         end
 9     end
10 end
```

Algorithm 1. Link Reputation (Main algorithm)

In the main algorithm (Algorithm 1), while there are BGP events (line 1) and `record` (line 2) from *BGPStream*, link reputation can be compute for each link (line 3). `link stability` metric result is saved on `link_stab` variable (line 4). On line 5, `bogon degree` metric result from function `Bogon` apply formula

1 and the last metric `link sensibility` result is obtain on line 6. With all components, the link reputation is compute by combining those metrics using the formula on line 7. The last part this formula express exponential impact of correlation between announcement and the type of prefixes. The 1 is a fallback to make sure we have *positive cost* when others parts of the formula gives 0. We have evaluated other formulas to combines those metrics without conclusive results. Our tests shows that the current formula is quite acceptable as shown in Sect. 5.

5 Results

Due to resources limitation, our implementation has been tested on four cases: censorship (Youtube hijack), malicious activities (Link Telecom hijack), configuration error (Malaysia Telekom route leaks) and country-wide outage on Egypt. For each case, two states has been observed: one *incident* state and one control state, called *normal* state. For the purpose of this paper, only Egypt country-wide outage case will be discuss.

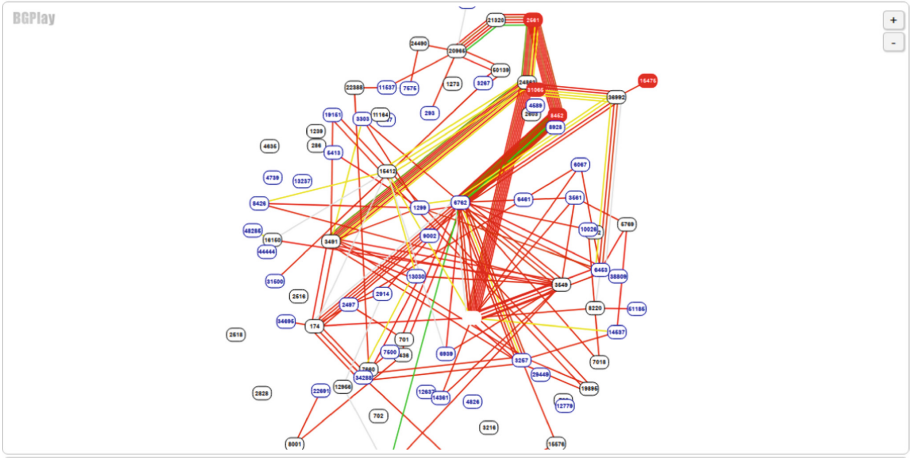

Fig. 1. Internet in Egypt offline (incident state)

Figure 1 presents the *incident* state where majority of the 3089 prefixes (196.219.246.0/24, 81.21.104.0/24 and 193.227.0.0/18 are monitored for this analysis) assigned to Egyptian Ases have been withdrawn on January 27th, 2011 [8]. We can notice a majority of red lines between `Origin AS` and transit/peers Ases, helping us verify our intuition: a majority of the links have *bad reputation*.

Figure 2 is the *normal* state of this case (February 2nd, 2011). Here we have mixed link color. Some have *good reputation* while other have *bad reputation*. Since Egyptian Ases started advertising theirs prefixes as before the censorship, reputation of those links is gradually migrating from *bad* to *good*.

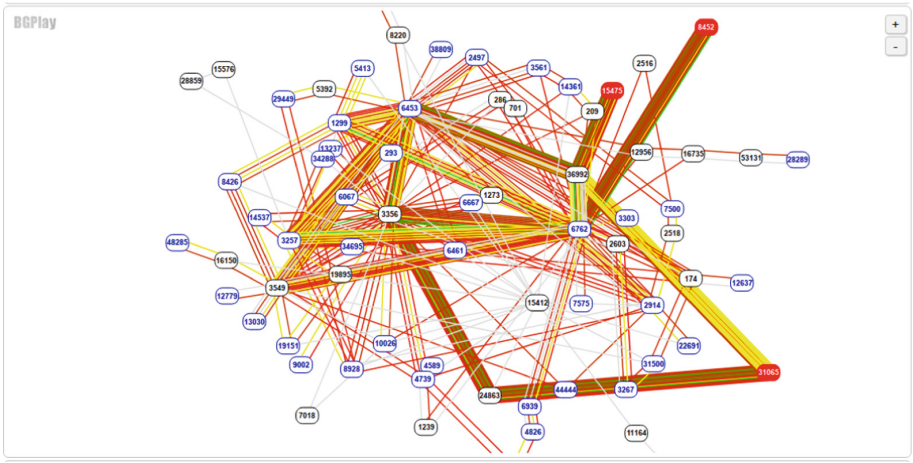

Fig. 2. Egypt back online (control state) (Color figure online)

6 African Perpective

The low density of Internet in Africa does not lower BGP threats in the continent. On the other hand, the lack of resources makes cheap security solutions a requirement. This paper is a step in this direction. The next step is a real time implementation of our solution that will be tested with operators we met during African Peering and Interconnection Forum.

7 Conclusion

BGP is a resilient protocol providing stable inter-domain routing since late 1993. BGP threats are well known but most proposed solution require significant changes to be applied worldwide for the protocol, which would require global agreement. However living with BGP threats is possible if these threats can be quickly identified. This work is a response to the need expressed by linkrank projects (UCLA and CAIDA). In this work a graphical link detection algorithm-based tool has been introduced, to easily provide baseline leak/hijack detection. However there is space for improvement. BGP is a large scale protocol, so massive data analysis tools like machine learning algorithms has to be used to increase the efficiency of the algorithm. Those technologies will also help to evaluate the number of links detected with correct/expected reputation. There is also the need to develop a more appropriate viewer, since *BGPlayJS* was not designed to display only one line between ASes.

References

1. Biersack, E., et al.: Visual analytics for BGP monitoring and prefix hijacking identification. IEEE Netw. **26**(6), 33–39 (2012)
2. Butler, K., Farley, T.R., McDaniel, P., Rexford, J.: A survey of BGP security issues and solutions. Proc. IEEE **98**(1), 100–122 (2010)
3. CAIDA: List-of-challenges. https://github.com/CAIDA/bgp-hackathon/wiki/List-of-Challenges#linkrank-1
4. Chang, J., et al.: AS-TRUST: a trust quantification scheme for autonomous systems in BGP. In: McCune, J.M., Balacheff, B., Perrig, A., Sadeghi, A.R., Sasse, A., Beres, Y. (eds.) Trust 2011. LNCS, vol. 6740, pp. 262–276. Springer, Heidelberg (2011). https://doi.org/10.1007/978-3-642-21599-5_20
5. Chang, J., et al.: AS-CRED: reputation and alert service for interdomain routing. IEEE Syst. J. **7**(3), 396–409 (2013)
6. Colitti, L., Di Battista, G., Mariani, F., Patrignani, M., Pizzonia, M.: Visualizing interdomain routing with BGPlay. J. Graph Algorithms Appl. **9**(1), 117–148 (2005)
7. Cymru, T.: The bogon reference. http://www.team-cymru.org/bogon-reference.html
8. Dainotti, A., et al.: Analysis of country-wide internet outages caused by censorship. In: Proceedings of the 2011 ACM SIGCOMM Conference on Internet Measurement Conference, IMC 2011, pp. 1–18. ACM, New York (2011). https://doi.org/10.1145/2068816.2068818
9. Huston, G., Rossi, M., Armitage, G.: Securing BGP - a literature survey. IEEE Commun. Surv. Tutor. **13**(2), 199–222 (2011)
10. Konte, M., Perdisci, R., Feamster, N.: ASwatch: an as reputation system to expose bulletproof hosting ASes. ACM SIGCOMM Comput. Commun. Rev. **45**(4), 625–638 (2015)
11. Lad, M., Zhang, L., Massey, D.: Link-Rank: a graphical tool for capturing BGP routing dynamics. In: 2004 IEEE/IFIP Network Operations and Management Symposium, NOMS 2004, vol. 1, pp. 627–640. IEEE (2004)
12. Orsini, C., King, A., Giordano, D., Giotsas, V., Dainotti, A.: BGPStream: a software framework for live and historical BGP data analysis. In: Proceedings of the 2016 ACM on Internet Measurement Conference, pp. 429–444. ACM (2016)
13. Rekhter, Y., Li, T., Hares, S.: A border gateway protocol 4 (BGP-4) RFC 4271. Technical report (2005)
14. Prem Sankar, A.U., Poornachandran, P., Ashok, A., Manu, R.K., Hrudya, P.: B-Secure: a dynamic reputation system for identifying anomalous BGP paths. In: Satapathy, S.C., Bhateja, V., Udgata, S.K., Pattnaik, P.K. (eds.) Proceedings of the 5th International Conference on Frontiers in Intelligent Computing: Theory and Applications. AISC, vol. 515, pp. 767–775. Springer, Singapore (2017). https://doi.org/10.1007/978-981-10-3153-3_76
15. Yu, H., Rexford, J., Felten, E.W.: A distributed reputation approach to cooperative internet routing protection. In: 1st IEEE ICNP Workshop on Secure Network Protocols (NPSec), pp. 73–78. IEEE (2005)