CommunityWatch: The Swiss-Army Knife of BGP Anomaly Detection

Vasileios Giotsas
Lancaster University
v.giotsas@lancaster.ac.uk

1 MOTIVATION

The design of the Internet as a network of independent networks, or Autonomous Systems (ASes), allowed it to spontaneously evolve to the core communications technology for contemporary society, but also resulted in the ossification of its core protocols, and the opacity of its structure. The current version of BGP [16], the de-facto inter-domain routing protocol, is over two decades old, and despite various revisions since then, a number of serious problems have been building over time, including weak security, non-deterministic behavior, and proneness to misconfiguration [10].

While the vulnerabilities inherent in the Internets architecture have been known for decades, and there has been a great extent of research to address them [14], the proposed solutions have not been widely deployed due to the costs and risks involved in replacing the existing network equipment [12]. As a result, most networks rely on reactive defense mechanisms [21]. Nonetheless, the highly distributed ownership of the Internet infrastructure and its highly dynamic nature make the development of the appropriate anomaly detection mechanisms far from trivial. Operators have full control over their own infrastructure, but little knowledge of what happens beyond their network perimeter. Third-party services can extend the detection capabilities for some classes of anomalies beyond an AS's domain, but the costs involved and concerns with data sharing make many operators reluctant to outsource such functionalities [20]. Instad, operators often resort in social media and mailing lists in an effort to crowdsource the debugging of their routing issues [2], an approach that can be error-prone and inefficient.

We take steps toward remedying this situation by developing *CommunityWatch*, an open-source system that enables timely and accurate detection of BGP routing anomalies, by leveraging meta-data encoded by AS operators directly on their BGP messages through the use of the BGP Communities attribute.

2 HOW CommunityWatch WORKS

The key insight of our approach is that BGP is no longer purely an information hiding protocol [18]. The flattening of the Internet hierarchy [6], has led to very dynamic and continuously growing peering clusters, and complex peering practices [9]. Consequently, operators require increasing flexibility and expressiveness in defining their routing policies and communicating them to their neighbors. The optional BGP Communities attribute [3] offers this flexibility by allowing operators to encode arbitrary information on their prefix announcements, including business relationship types, route redistribution policies, location data, and traffic blackholing requests to mitigate attacks [4]. Their use has become increasingly popular, allowing us to use them as an automated crowdsourcing mechanism for acquiring accurate operator-provided information for about 50% of IPv4 and 30% of IPv6 updates. Between 2010 and 2016, the visible ASes using BGP Communities more than doubled, and the number of unique community values tripled to more than 50,000.

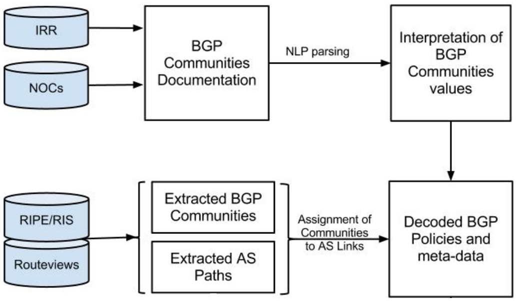
IRR
NOCs
BGP Communities Documentation
NLP parsing
Interpretation of BGP Communities values
RIPE/RIS
Routeviews
Extracted BGP Communities
Extracted AS Paths
Assignment of Communities to AS Links
Decoded BGP Policies and meta-data

Figure 1: Data collection methodology.

BGP Communities have the format X:Y, where X, Y are two 16-bit values (extended communities use four octets [19]). By convention, the first two octets encode the ASN of the operator that sets the community, while the next two octets encode denote the specific information carried by the Community, as the ingress location of a route. Importantly, Communities is a *transitive* attribute, which means that they can be propagated through multiple AS hops, and we can mine their values through publicly available BGP collectors.

Successful interpretation of the attached Communities values allows monitoring of BGP routes and gathering of routing intelligence based on authoritative data instead of heuristics. However, the Communities attribute lacks standardized values and semantics. Many operators document their Communities values in Internet Routing Registry (IRR) records, or their webpages, but typically not in machine-parsable format.

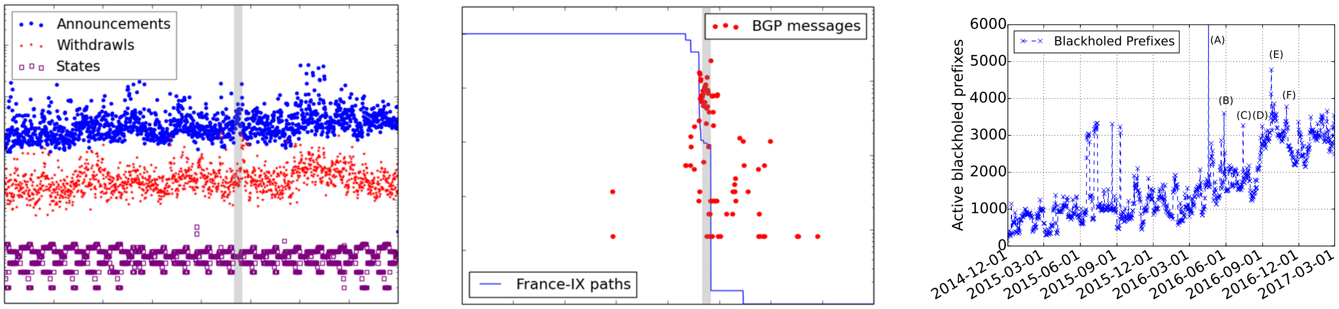

(a) Aggregated routing activity.

(b) Routing activity filtered based on BGP Communities.

(c) Longitudinal growth of blackholed prefixes

Figure 2

To decipher the Communities values in an automated manner, we combine a web-mining tool with Google's Google Cloud Natural Language API [11] to achieve the automatic compilation of a Communities dictionary, as explained in [7]. As of March 2018, the dictionary included 11,830 interpreted Communities, 48% of which encode geolocation data, 21% encode relationship type, and the rest encode different types of routing policies (selective advertisement, blackholing, local preference tuning, path prepending).

Figure 1 illustrates the data extraction methodology. The system combines the interpreted Communities with a live stream of BGP data, obtained through through BGPStream [15], to extract BGP updates annotated with the corresponding Communities. During the initialization phase, we continuously monitor the incoming BGP messages to establish a baseline of paths that are consistently tagged with a stable set of Communities. Then, CommunityWatch monitors the baseline of annotated paths to capture changes through explicit BGP withdrawals, or through changes in the attached Community values. Routing updates are binned in time intervals to correlate path changes with routing incidents. The system uses a binning interval of 60 seconds (twice the default MRAI time [17]). Whenever we detect a binning interval for which the paths deviating from the baseline exceed a minimum threshold, we trigger a signal of potential routing anomaly. Depending on the type of Communities, the corresponding signal investigation module analyzes the affected paths to determine the root cause of the observed change.

3 ANOMALY DETECTION RESULTS

The fact that BGP Communities encode different categories of routing meta-data, means that CommunityWatch can detect a wide range of routing anomalies. In this section we illustrate three such cases.

Infrastructure Outages In the past, the AS-path data have been used to study changes in prefix availability and reachability, and reveal the occurrence of outages due to country-level censorship, attacks, or natural disasters [1]. However, the coarse granularity of AS-paths has hindered detailed analysis of infrastructure outages, since many failures may change the infrastructure-level path, e.g., switching to another PoP, but the AS-path remains the same. Thus, an AS-path level and prefix-level analysis cannot show such failures. Communities are often used to encode location information at fine granularities, such as IXP-level, and facility-level Points-of-Presence. For instance, figure Figure 2b illustrates how using BGP Communities to filter the routing activity can reveal an outage at the France-IX IXP [5], by effectively de-noising the aggregated routing activity 2a that obscures the impact of localized events on the dynamics of BGP. We provide more details on the detection of infrastructure-level outages in [7].

Detection of Blackholed Prefixes Blackholing is a popular DDoS mitigation strategy inside a single network or among multiple networks. BGP enables blackholing by leveraging the BGP communities attribute. Networks trigger blackholing requests by sending BGP announcements to their BGP neighbors for specific destination prefixes with the appropriate blackhole community. Parsing of these values enable CommunityWatch to differentiate blackholing requests from normal BGP announcements, and characterize prefixes under attack, as we explain in [8].

Route leaks and policy violations Detection and analysis of export policy violations, such as the violations of the valley-free rule, is of particular importance for understudying BGP misconfiguration and characterizing misbehaving networks. Detection of such violations requires accurate AS relationship data, but the universality of valley-free rule is a fundamental assumption of relationship inference algorithms. CommunityWatch detects such violations by parsing relationship-tagging communities which are free from inference heuristics biases. We analyzed BGP data in March 2018 to find that over 3% of BGP paths violate the valley-free rule, which cannot be capture by the AS-Rank algorithm [13].

4 ACKNOWLEDGMENTS

The source code of CommunitiesWatch and detailed documentation is publicly available at https://gitlab.com/vgiotsas/CommunitiesDetector.

REFERENCES

[1] Giuseppe Aceto, Alessio Botta, Pietro Marchetta, Valerio Persico, and Antonio Pescapé. 2018. A comprehensive survey on internet outages. *Journal of Network and Computer Applications* (2018).

[2] Ritwik Banerjee, Abbas Razaghpanah, Luis Chiang, Akassh Mishra, Vyas Sekar, Yejin Choi, and Phillipa Gill. 2015. Internet outages, the eyewitness accounts: Analysis of the outages mailing list. In *International Conference on Passive and Active Network Measurement*. Springer, 206–219.

[3] R. Chandra, P. Traina, and T. Li. 1996. BGP Communities Attribute. IETF RFC 1997. (1996).

[4] B. Donnet and O. Bonaventure. 2008. On BGP communities. 38, 2 (March 2008), 55–59.

[5] FranceIX. [n. d.]. Outage Notification. https://www.franceix.net/en/events-and-news/news/franceix-outage-notification/. ([n. d.]).

[6] Phillipa Gill, Martin Arlitt, Zongpeng Li, and Anirban Mahanti. 2008. The flattening internet topology: Natural evolution, unsightly barnacles or contrived collapse?. In *International Conference on Passive and Active Network Measurement*. Springer, 1–10.

[7] Vasileios Giotsas, Christoph Dietzel, Georgios Smaragdakis, Anja Feldmann, Arthur Berger, and Emile Aben. 2017. Detecting Peering Infrastructure Outages in the Wild. In *Proceedings of the Conference of the ACM Special Interest Group on Data Communication*. ACM, 446–459.

[8] Vasileios Giotsas, Philipp Richter, Georgios Smaragdakis, Anja Feldmann, Christoph Dietzel, and Arthur Berger. 2017. Inferring BGP blackholing activity in the internet. In *Proceedings of the 2017 Internet Measurement Conference*. ACM, 1–14.

[9] V. Giotsas, G. Smaragdakis, B. Huffaker, M. Luckie, and k. claffy. 2015. Mapping Peering Interconnections at the Facility Level.

[10] S. Goldberg. 2014. Why is it taking so long to secure internet routing? *Queue* 12, 8 (2014), 20.

[11] Google. [n. d.]. Cloud Natural Language. https://cloud.google.com/natural-language/. ([n. d.]).

[12] Mark Handley. 2004. Evolving the Internet: Changing the Engines in Mid Flight. Invited presentation at the ICSE 2004, Edingburgh, Scotland.. (May 2004). http://www.cs.ucl.ac.uk/staff/M.Handley/slides/icse.pdf.

[13] M. Luckie, B. Huffaker, A. Dhamdhere, V. Giotsas, and kc claffy. 2013. AS Relationships, Customers Cones, and Validations.

[14] A. Mitseva, A. Panchenko, and T. Engel. 2018. The State of Affairs in {BGP} Security: A Survey of Attacks and Defenses. *Computer Communications* (2018).

[15] C. Orsini, A. King, D. Giordano, V. Giotsas, and A. Dainotti. 2016. BGPStream: a software framework for live and historical BGP data analysis. In *Internet Measurement Conference (IMC)*.

[16] Y. Rekhter and T. Li. 1996. A Border Gateway Protocol 4 (BGP-4). IETF RFC 1654. (1996).

[17] Y. Rekhter, T. Li, and S. Hares. 2006. A Border Gateway Protocol 4 (BGP-4). IETF RFC 4271. (2006).

[18] M. Roughan, W. Willinger, O. Maennel, D. Pertouli, and R. Bush. 2011. 10 Lessons from 10 Years of Measuring and Modeling the Internet's Autonomous Systems. 29, 9 (2011).

[19] S. Sangli, D. Tappan, and Y. Rekhter. 2006. BGP Extended Communities Attribute. IETF RFC 4360. (2006).

[20] P. Sermpezis, V. Kotronis, A. Dainotti, and X. Dimitropoulos. 2018. A Survey among Network Operators on BGP Prefix Hijacking. *arXiv preprint arXiv:1801.02918* (2018).

[21] Zheng Zhang, Ying Zhang, Y Charlie Hu, and Z Morley Mao. 2007. Practical defenses against BGP prefix hijacking. In *Proceedings of the 2007 ACM CoNEXT conference*. ACM, 3.