A Multi-Tiered Anomaly Detection Framework for BGP Routing Security

1st Swastik Kar
Amrita Vishwa Vidyapeetham, Amritapuri)
Amritapuri, India https://orcid.org/0009-0000-0058-039X

***Abstract*—BGP, the protocol responsible for interdomain routing across the global internet, continues to face critical vulnerabilities, particularly in the form of routing anomalies and prefix hijacking attacks. Despite the existence of various anomaly detection systems, most current solutions suffer from poor scalability and require frequent, resource-intensive retraining to accommodate the continuous addition of new Autonomous Systems (ASes) and the emergence of novel attack patterns. In this work, we propose a multi-tiered anomaly detection framework that allows network operators and service providers to select different tiers of security based on their operational priorities and available computational resources. This will allow the network to respond quickly, in case of events where new ASes are added to the global network or new attacks are discovered. This tier-based approach ensures that lightweight, rapidly deployable detection mechanisms are available as a first line of defense, while more comprehensive, resource-heavy models continue training in the background.**

I. Introduction

Overview The border gateway protocol was created in the late 1980s, and since then, it has become one of the most, if not the most, widely used internet protocol to route traffic across the internet. It works by exchanging route announcements amongst organisations or Autonomous Systems (ASes) - which contains path information to reach the ASes. However, even though it has been widely adopted as the leading internet routing protocol, it doesn't have any built in authentication or security feature. Due to this, it is susceptible to various kinds of threats and attacks. An AS that has been compromised can be forced to announce wrong or random routes, which might cause it route traffic through and to unintended ASes or locations, which are often malicious, and trying to intrude upon information or data that is not meant for them, or just trying obstruct the global routing process. This process is called BGP hijacking, which is forcing an AS to form a route with a malicious AS. There are many other types of attacks as well, such as route leaks, which causes the rerouting of traffic across unwanted paths. Sometimes unintentionally, caused due to accidentally incorrect routing information or human error, these anomalies can occur. However, intentionally or unintentionally, these anomalies or hijacks cause a disruption in the global routing flow of the internet. In the past decades there have been many significant BGP attacks, and these are continuing to happen. For example, in 2008 [1], there was a BGP hijacking incident involving Pakistan Telecom (AS17557), which caused access to YouTube to be blocked for a major part of the world. Similarly, in 2014 [2], there was a route leak involving an AS in Malaysia, called Telekom Malaysia (AS4788). This caused significant problems in the networks for the global routing systems, causing Internet slowdown. The users of a cryptocurrency platform in South Korea, called KLAYswap [3], lost USD 1.9 million, after malicious hackers conducted a BGP Hijack attack on the infrastructure of the server of one of the providers of the platform, in 2022. In 2019 [4], AS 21217, SafeHost propagated announcements to another provider, AS 4134, China Telecom, which caused the redirection of large amounts of traffic that was originally meant for European mobile networks to be routed through China Telecom, which would cause loss of private data or invasion of privacy. Several solutions have been proposed in literature in order to counter these problems with BGP, such as BGPsec [5], psBGP [6], KC-x [7], s-BGP [8], among others - however these don't see widespread use. This can be attributed to high cost of implementation, as this would involve an infrastructure change, and network providers are often hesitant to incur these costs. Additionally, since routing protocols need to be standardized, unless the same protocol is implemented by a number of number providers, changing the routing protocol is often redundant. Therefore, to make the internet safer to navigate, a different approach needs to be taken. All of these approaches can be broadly divided into two approaches, preventive and reactive. Preventive approaches are those that try to modify the underlying infrastructure of the Border Gateway Protocol to make it safer, by altering the way BGP functions. Reactive approaches use detection systems in order to detect anomalies that might have occurred in a network, and then identify them as suspicious or not suspicious.

Autonomous System The internet is a collection of many different networks, clustered together, sharing direct or indirect connections between each other, forming an interconnected web-like structure, or a mesh like structure. Each of these individual networks is called an AS, short for Autonomous System. These Autonomous Systems are a collection of Internet Protocol networks or IP networks and routes which are under the control or governance of a single organization, and this organization follows a unified routing policy, which is decided by the AS themselves. These ASes could be internet service providers, which provide services to other ASes or networks, domain providing services, or big multinational corporations, perhaps even the network of colleges or universities big enough to have their own network. Each of these autonomous systems is provided with an unique identification number, which is called the autonomous system number or ASN, which is assigned by Internet Assigned Numbers Authority (IANA), which is then used to identify the AS globally. As of June 2025, there are approximately 120,000 autonomous systems [9] present globally. Without any governing body, all of these ASes might have used a different routing algorithm in order to communicate with each other, which would make intercommunications very difficult and tedious to conduct, and for a single chain of communication multiple routing protocols might be used. Hence it became increasingly important to make use of a unified routing protocol, so that communications could be standardized, and hence the border gateway protocol was introduced.

Border Gateway Protocol The Border Gateway Protocol or BGP is one of the most widely used routing protocols used globally for routing information between different ASes on the internet. It can be looked at like the postal system of the internet - it ensures that packets reach their destination after leaving the source by managing how packets get routed from AS network to network It creates paths for communication between the ASes by exchanging route announcements in the network. These route announcements contain information about the path that is to be followed in order to reach certain networks or prefixes, such as an IP address or a block of IP addresses. Typically, these announcements are processed by edge routers of the network - these are essentially those routes that are present on the very edge or border of a network, and these are really the ones that directly connect to the other external networks - as internally within a network there other routing protocol being followed; but for global routing BGP is used. These edge routers maintain a routing table and based on other attributes such as AS-PATH length, local preference, MED and origin type, BGP decides the path that is to be followed. However, in spite of it being adopted globally adopted for use, it suffers from a major drawback - it does not have any built in security or authentication feature - which makes it particularly susceptible to malicious attacks, that could result in major privacy concerns or data loss from, which are designed to reroute the network traffic, such as prefix hijacking, route leaking and route hijacking.

BGP Attacks A few of the main types of BGP attacks are prefix hijacking, route leaking, and route hijacking. Prefix hijacking itself can be divided into two main categories: regular prefix hijacking and sub prefix hijacking.

Regular Prefix Hijacking - In this attack, the hijacker forges the prefix of the target AS, causing traffic originally destined for the target AS to be rerouted through the hijacker's AS, allowing them to intercept or steal the data.

Sub-prefix Hijacking - This is a more targeted form of prefix hijacking, where the hijacker not only forges the prefix but advertises a more specific (longer) prefix than the legitimate AS. Due to BGP's Longest Prefix Match (LPM) rule, this causes traffic to be rerouted through the hijacker's AS instead of the intended destination.

Route Leaking - In a Route Leaking attack, a misbehaving or anomalous AS improperly propagates route announcements to another AS, violating the target AS's routing policies. This causes traffic to be forwarded through unintended paths, exposing it to security, privacy, and confidentiality risks, as malicious ASes along the path can intercept, monitor, or modify the data.

Route Hijacking - Route hijacking typically occurs in two forms: (i) a fake path attack, where the hijacker advertises a forged AS-PATH that appears more attractive than the legitimate path, and (ii) false ownership of a prefix, where the hijacker falsely claims ownership of an IP prefix. In the fake path scenario, the attacker simply needs to announce a fabricated AS-PATH leading to the target AS to mislead upstream routers.

II. Related Works

There is research that has been done towards making BGP more secure and equipping it with security features. These approaches can be usually divided into two types - preventive approaches and reactive approaches. BGP hijack prevention approaches usually are based on cryptographic authentication. KC-x [7] provides a keychain based signature scheme, which has low CPU costs and low memory overheads, and while providing an incentive for incremental deployment. They have two versions of this scheme, based on different signature algorithms, one based on RSA called KC-RSA, and the other based on the Merkel hash tree, called KC-MT. s-BGP[8] made use of Public Key Infrastructure to assign the ownership of an IP prefix to an AS in order to provide authentication and verify identities of BGP speakers. Each BGP speaker verifies route announcements by previous speakers via the public key, while adding their own private key to the route. However, this makes it expensive in both computation and storage. psBGP [6] made use of a decentralized authentication system a prefix assertion list which contains the route announcements of local ASes and their peers. The consistency of this information is checked around the origin of the local AS and its peers and thus is validated. soBGP [10] proposes a system which allows a BGP speaker to check the authenticity of an originating AS and the validity of the path by providing a secure mechanism. PGBGP [11] offers protection by detecting suspicious advertisements by making use of historical hits and then delaying their announcements and assigning them a low preference and ignoring them temporarily. The other approach, reactive solutions monitor a network and observe the route announcements made in order to detect anomalies and flag routes as anomalous routes and ASes as hijacked. BGP2Vec [12] made use of NLP by treating route announcements as sentences for an NLP model in order to detect anomalous routes. AP2Vec [13] extended on BGP2Vec's approach and compared their embedding that they obtained based on the routes to the new routes in order to identify large differences. Moriano et al. [14] noted that BGP anomalies tend to occur in groups of bursts, which they have used in order to detect anomalous routing incidents. Dong et al. [15] proposed a self attention based LSTM model to obtain the difference between BGP anomalies including feature and time dimensions. BEAM [16] used similarities in hierarchy and proximity between ASes by comparing new routes to the old ones in order to detect anomalies. These solutions often rely on outdated, faulty or incomplete data - which means that incase of some error in the data, the same will be reflected in the training as well, which will result in faulty results. Furthermore, these solutions often require a change to the network framework which needs modifications to the network infrastructure, which most operators aren't willing to or able to take, as these incur in additional cost on the operators; additionally, since all the networks need the same protocol for standardization in communication, modifying the infrastructure would bear no benefits, until a significant number of operators simultaneously agree to make changes too. On top of that, the internet's routing environment is constantly changing - new attacks appear regularly, new networks or ASes are added frequently, and it's not feasible to retrain the data every time a new change occurs. Based on the operational limitations and research gaps identified in current BGP anomaly detection systems, our objective is to develop a tier-based, adaptable anomaly detection framework that provides varying levels of protection tailored to the specific needs and resources of individual customers or network operators; with the various tiers corresponding with different levels of security. The idea behind this approach is to have quick training and deployment in case of discovery of a new attack or addition of a new AS. This will allow the deployment of a first line of defence while the other tougher models are still training, and the detection would still be functional when a new attack or AS is introduced.

III. Methodology

A. Overview

In this paper propose a multi-tiered detection system; in which multiple layers would be present, and each later more secure than the previous. We also propose a light weight model, which we have called Vanguard, which will act as the first tier of our detection system.

B. Dataset

For this research, we require access to a reliable, large-scale dataset of BGP routing information, both historical and live, to train and evaluate our anomaly detection models. To achieve this, we have made use of BGPstream [17], an open-source software framework developed by CAIDA (Center for Applied Internet Data Analysis). BGPstream is specifically designed for the collection, processing, and analysis of BGP data, offering a flexible and efficient means to gather routing information from multiple sources. BGPstream connects to public BGP data archives such as Route Views and RIPE RIS, as well as real-time data streams like RIS Live and BMP-enabled feeds, enabling us to capture a comprehensive view of internet routing activities over time. Once collected, this data serves as the basis for training our detection models. To evaluate the system's performance and reliability, we test the trained models against historical BGP anomaly events, verifying their capability to correctly identify known incidents while assessing detection accuracy and false positive rates. Due to practical constraints on computational resources and processing capacity, we have deliberately chosen to narrow the geographic and operational scope of our study. Specifically, for this paper, we have focused exclusively on ASes registered under the Asia Pacific region, obtaining data through APNIC (Asia Pacific Network Information Centre). This regional restriction reduces the total number of Autonomous Systems considered from approximately 200,000 globally to around 30,000 ASes in Asia. Likewise, the associated IP prefixes, spanning both IPv4 and IPv6, have been scaled down from about 36 billion globally to approximately 10 billion within the Asia-Pacific region.

C. Model

We first aggregate anomalous events detected by existing anomaly detection systems. Using the list of these detected events, we use BGPstream to retrieve the original BGP update packets associated with each incident. These BGP packets serve as the primary traffic input for our model. For anomaly detection, we implement a self-attention-based Long Short-Term Memory (LSTM) Recurrent Neural Network (RNN). The model architecture consists of an input layer with a size of 30,000, followed by a window size of 2, and an output layer of 4 × 30,000. The network includes five layers beyond the input layer, featuring a 1D convolutional layer with 32 filters of length 3, amounting to 3,072 trainable weights and 32 bias parameters. We employ the ReLU activation function and introduce a Max Pooling layer producing 32 feature maps of size 2 to reduce dimensionality. The final output layer uses a sigmoid activation function for binary classification. The model is trained using the Adam optimizer over 16 epochs with a batch size of 64.

IV. Experiments

As said earlier, for our scope, we constrained out dataset to include only Autonomous Systems (ASes) originating from the Asia-Pacific region. While the global AS count currently is almost 200,000, this restriction reduced our working dataset to approximately 30,000 ASes. Correspondingly, the total number of associated IP prefixes - encompassing both IPv4 and IPv6 allocations - was scaled down from an estimated 36 billion entries worldwide to around 10 billion within the regional scope. We tested our model against 7 real world anomaly events, and we managed to detect 4/7.

V. Future Scope and Conclusion

In the future, this model could be extended to incorporate Autonomous Systems (ASes) from across the globe, rather than being limited to a single region. Expanding the scope in this way would enable broader, more comprehensive protection against BGP routing anomalies at a global scale. Additionally,

Anomalous Event	Vanguard
Pakistan YouTube Route Hijack (2008)	✓
China Hijack (2010)	x
Indosat Prefix Hijack (2014)	✓
Bharti Airtel Route Leak (2015)	✓
China Route Leak (2019)	x
Malaysia Route Leak (2015)	✓
Myanmar Route Hijack (2021)	x
Total Count	4/7

deploying the system in a live operational network environment would allow for valuable real-world testing, helping to assess the model's ability to generalize and adapt to previously unseen anomalies or emergent routing events. To conclude, in this work we proposed and evaluated a tier-based anomaly detection framework for BGP route monitoring. The multi-tiered design enables network operators to select varying levels of security and resource allocation based on their operational demands and financial constraints. For the scope of this study, we restricted our dataset to ASes from the Asia-Pacific region to manage resource limitations and focus on a representative subset of the global AS topology. The proposed system demonstrated promising results by successfully detecting 4 out of 7 historical BGP anomaly events in the region. These findings highlight the viability of a lightweight, scalable, and flexible tier-based detection approach for improving BGP security.

REFERENCES

[1] Anonymous. "YouTube Hijacking: A RIPE NCC RIS case study." ripe,17 March 2008, https://www.ripe.net/about-us/news/youtube-hijacking-a-ripe-ncc-ris-case-study/ [2] Toonk Andre. "Massive route leak causes Internet slowdown." bgpmon, 12 June 2015, https://www.bgpmon.net/massive-route-leak-cause-internet-slowdown/ [3] Siddiqui Aftab. "KlaySwap – Another BGP Hijack Targeting Crypto Wallets." manrs, 17 Feb 2022, https://manrs.org/2022/02/klayswap-another-bgp-hijack-targeting-crypto-wallets/ [4] Madory Doug. "Large European routing leak sends traffic through China Telecom." apnic, 7 Jun 2019 https://blog.apnic.net/2019/06/07/large-european-routing-leak-sends-traffic-through-china-telecom/ [5] M. Lepinski, K. Sriram. "BGPsec Protocol Specification" rfc-editor, Sep 2017 https://www.rfc-editor.org/rfc/rfc8205 [6] Oorschot, Paul Wan, Tao Kranakis, Evangelos. (2007). On interdomain routing security and pretty secure BGP (psBGP). ACM Trans. Inf. Syst. Secur.. 10. 10.1145/1266977.1266980. [7] Tao Wan, Evangelos Kranakis, and P.C. van Oorschot, "Pretty secure BGP (psBGP)," in Proceedings of Network and Distributed System Security Symposium (NDSS 2004), February 2004. [8] S. Kent, C. Lynn and K. Seo, "Secure Border Gateway Protocol (S-BGP)," in IEEE Journal on Selected Areas in Communications, vol. 18, no. 4, pp. 582-592, April 2000, doi: 10.1109/49.839934. [9] Anonymous. "RIR Delegations and RIPE NCC Allocations" Regional Internet Registries Statistics, https://www-public.telecom-sudparis.eu/ maigron/rir-stats/rir-delegations/world/world-asn-by-number.html [10] R. White and B. Akyol, "Deployment considerations for Secure Origin BGP(soBGP)," Internet Draft, Internet Engineering Task Force, June 2003. [11] J. Karlin, S. Forrest and J. Rexford, "Pretty Good BGP: Improving BGP by Cautiously Adopting Routes," Proceedings of the 2006 IEEE International Conference on Network Protocols, Santa Barbara, CA, USA, 2006, pp. 290-299, doi: 10.1109/ICNP.2006.320179. [12] T. Shapira and Y. Shavitt, "A deep learning approach for IP hijack detection based on ASN embedding," in Proc. Workshop Netw. Meets AI ML, 2020, pp. 35–41. [13] T. Shapira and Y. Shavitt, "Ap2vec: an unsupervised approach for bgp hijacking detection," TNSM, vol. 19, no. 3, pp. 2255–2268, 2022. [14] P. Moriano, R. Hill, and L. J. Camp, "Using bursty announcements for detecting BGP routing anomalies," Comput. Netw., vol. 188, Apr. 2021, Art. no. 107835 [15] Y. Dong, Q. Li, R. O. Sinnott, Y. Jiang, and S. Xia, "Isp self-operated bgp anomaly detection based on weakly supervised learning," in ICNP. IEEE, 2021, pp. 1–11. [16] Yihao Chen, Qilei Yin, Qi Li, Zhuotao Liu, Ke Xu, Yi Xu, Mingwei Xu, Ziqian Liu, and Jianping Wu. 2024. Learning with semantics: towards a semantics-aware routing anomaly detection system. In Proceedings of the 33rd USENIX Conference on Security Symposium (SEC '24). USENIX Association, USA, Article 288, 5143–5160. [17] Chiara Orsini, Alistair King, Danilo Giordano, Vasileios Giotsas, and Alberto Dainotti. 2016. BGPStream: A Software Framework for Live and Historical BGP Data Analysis. In Proceedings of the 2016 Internet Measurement Conference (IMC '16). Association for Computing Machinery, New York, NY, USA, 429–444.