Towards security transparency of Autonomous Systems on the Internet

Shyam Krishna Khadka[1], Ralph Holz[1,2], and Cristian Hesselman[1,3]

[1] University of Twente, Drienerlolaan 5, 7522 NB Enschede, Netherlands
[2] University of Münster, Schlosspl. 2, 48149 Münster, Germany
[3] SIDN Labs, Meander 501, 6825 MD Arnhem, Netherlands
{s.k.khadka,r.holz,c.e.w.hesselman@utwente.nl}@utwente.nl

1 Introduction

The use of the Internet is not just limited to daily activities such as communication, entertainment, and shopping, but many critical services including finance, healthcare, and the modern versions of infrastructures (e.g., power grids, transportation systems, water, oil, and gas pipelines)[19] increasingly use the Internet for their operations as well. This makes Internet paths a factor that must be considered a part of supply chain security.

In the following, we use the term "Internet path" to refer to a sequence of ASes that appear in the "AS path" attribute in BGP. The problem we face when considering Internet paths as part of the supply chain is that the Internet is a "black box": it offers no real insight into the security aspects of Internet paths and no control for relying parties to influence which paths are chosen in routing.

In this work, we propose a research agenda to investigate *how to build an AS reputation mechanism that could be used for measuring the reliability of Internet paths and selecting paths based on this.* Transparency of AS reputation and Internet path selection form an extension to the current BGP-based Internet, as envisioned for the "Responsible Internet" [8], a concept inspired by the notion of "Responsible AI", which provides more insight into the inner workings of AI systems [7]. We assume the deployment of AS security transparency and path selection will occur through so-called "trust zones" [4], which are coalitions of ASes that collaboratively implement security policies.

In this extended abstract, we discuss four research questions (RQ1-RQ4) that we have identified and sketch how we aim to address them. The objective of our abstract is to facilitate and inspire discussion at the TAURIN 2023 workshop.

2 RQ1: What may be the security attributes of an AS?

Enterprises and AS operators may prefer to select ASes on Internet paths that are not linked to malicious or suspicious incidents, simply because they view such paths as safer and more reliable [19]. We aim to create a reputation mechanism for this purpose. To this end, we collect data and measure attributes linked to incidents.

Approach: We will begin with the following datasets to understand which ASes that occur on Internet paths have been involved in incidents or may lack in security. As some data sources are snapshots, we also explore measurement methods that can be carried out in a sustained way.

(a) We identify ASes that are misbehaving or misconfigured, as identifiable in the following sources.
   – Abuse: phishing, spamming, malware distribution, botnet command & control hosting. Datasets: abuse feeds [1, 21, 13, 5], OpenINTEL[17];
   – Routing: route leaks, route hijacks [23]. Datasets: MANRS observatory [12], Internet Health Report [9], NANOG mailing list [14];

(b) Routing and traffic security measures implemented: Route filtering, Route Origin Authorization (ROA), Route Origin Validation(ROV), protection against DDoS and spoofing, security frameworks in use. Datasets: RPKI repository [18], CAIDA spoofer [22], network telescope data [3, 13];;

(c) Security status of hardware and software infrastructures used by AS operator, e.g., border routers. Datasets: NVD [16], Shodan [20], Censys [3];

3 RQ2: How to rank ASes in terms of reputation?

Knowing which ASes rank higher regarding their reputation will help AS operators choose or avoid particular ASes, assuming they can choose between different Internet paths.

Approach: We will explore the use of Multi-Criteria Decision Analysis (MCDA) [10] where operators will provide their preferences for the values of the security attributes that we explored in RQ1. We will also consider taking an expert opinions and aim to make the mechanisms output scores in real-time.

4 RQ3: How to choose paths based on ASes' reputations?

The AS operator should be able to choose/avoid particular ASes on their paths based on our work of RQ2, particularly within a trust zone.

Approach: We will set up a simulated testbed (e.g., using Mininet or BIRD BGP [2]) to explore different path selection/avoidance techniques, like using BGP community attributes, source-based routing, and traffic engineering techniques within a trust zone. We will study how this might affect the operational management of ASes.

5 RQ4: What may be a way to verify a path?

Users should be able to verify if a data path is the one they chose based on the methods from RQ3.

Approach: We will explore cryptographic principles [15, 11] as well as the concept of "BGP lies" [6] and reverse traceroute [24].

Acknowledgements: This work was conducted as part of the project CATRIN (https://www.catrin.nl), which received funding from the Dutch Research Council (NWO).

References

1. Abuse.ch — Fighting malware and botnets, https://abuse.ch, Last accessed 10-July-2023.
2. The BIRD Internet Routing Daemon Project, https://bird.network.cz/, Last accessed 09-July-2023.
3. Exposure Management and Threat Hunting Solutions — Censys, https://censys.io/, Last accessed 05-July-2023.
4. Clark, D., Claffy, K.: Trust zones: A path to a more secure internet infrastructure. Journal of Information Policy **11**, 26–62 (2021). https://doi.org/10.2139/ssrn.3746071
5. Cybercrime Information Center, https://cybercrimeinfocenter.squarespace.com/, Last accessed 05-July-2023.
6. Del Fiore, J.M.: Detecting hidden broken pieces of the Internet: BGP lies, forwarding detours and failed IXPs. Ph.D. thesis, Université de Strasbourg (2021)
7. Dignum, V.: Responsible artificial intelligence: designing AI for human values. ITU Journal: ICT Discoveries (September 2017)
8. Hesselman, C., Grosso, P., Holz, R., Kuipers, F., Xue, J.H., Jonker, M., de Ruiter, J., Sperotto, A., van Rijswijk-Deij, R., Moura, G.C., et al.: A responsible internet to increase trust in the digital world. Journal of Network and Systems Management **28**, 882–922 (2020). https://doi.org/10.1007/s10922-020-09564-7
9. Internet Health Report — Monitoring networks health, https://ihr.iijlab.net/ihr/en-us, Last accessed 02-July-2023.
10. Ishizaka, A., Nemery, P.: Multi-criteria decision analysis: methods and software. John Wiley & Sons (2013)
11. Kim, T.H.J., Basescu, C., Jia, L., Lee, S.B., Hu, Y.C., Perrig, A.: Lightweight source authentication and path validation. In: Proceedings of the 2014 ACM Conference on SIGCOMM. pp. 271–282 (2014). https://doi.org/10.1145/2740070.2626323
12. MANRS Observatory, https://observatory.manrs.org/, Last accessed 12-July-2023.
13. MISP Default Feeds, https://www.misp-project.org/feeds/, Last accessed 05-July-2023.
14. Welcome to Mailman.NANOG.Org, https://mailman.nanog.org/mailman/listinfo, Last accessed 03-July-2023.
15. Naous, J., Walfish, M., Nicolosi, A., Mazieres, D., Miller, M., Seehra, A.: Verifying and enforcing network paths with ICING. In: Proceedings of the Seventh Conference on Emerging Networking Experiments and Technologies. pp. 1–12 (2011). https://doi.org/10.1145/2079296.2079326
16. NIST:National Vulnerability Database, https://nvd.nist.gov/, Last accessed 03-July-2023.
17. OpenINTEL: Active DNS Measurement Project, https://openintel.nl/, Last accessed 05-July-2023.
18. FTP repos of RPKI in RIPE, https://ftp.ripe.net/rpki/, Last accessed 03-July-2023.

19. Schulzrinne, H.: Networking: The Newest Civil Engineering Challenge (August 2022), https://www.youtube.com/watch?v=5lvXIqI_mQ4, SIGCOMM Lifetime Achievement Award keynote, SIGCOMM 2022, Amsterdam
20. Shodan Search Engine, https://www.shodan.io/, Last accessed 05-July-2023.
21. DROP - Don't Route or Peer lists - The Spamhaus Project, https://www.spamhaus.org/drop/, Last accessed 01-July-2023.
22. Spoofer - CAIDA, https://www.caida.org/projects/spoofer/, Last accessed 03-July-2023.
23. Sriram, K., Montgomery, D., McPherson, D.R., Osterweil, E., Dickson, B.: Problem Definition and Classification of BGP Route Leaks. RFC 7908 (Jun 2016). https://doi.org/10.17487/RFC7908, https://www.rfc-editor.org/info/rfc7908
24. Vermeulen, K., Gurmericliler, E., Cunha, I., Choffnes, D., Katz-Bassett, E.: Internet scale reverse traceroute. In: Proceedings of the 22nd ACM Internet Measurement Conference. pp. 694–715 (2022). https://doi.org/10.1145/3517745.3561422