UNIVERSITY OF CALIFORNIA, SAN DIEGO UCSD

BERKELEY • DAVIS • IRVINE • LOS ANGELES • MERCED • RIVERSIDE • SAN DIEGO • SAN FRANCISCO SANTA BARBARA • SANTA CRUZ

CENTER FOR APPLIED INTERNET DATA ANALYSIS
(858) 534-8333

9500 GILMAN DR # 0505
LA JOLLA, CA 92093-0505

Marlene H. Dortch Secretary
Federal Communications Commission
45 L Street NE
Washington, DC 20554
Subject: **Re: Notice of *Ex Parte* Meeting, *Secure Internet Routing*, PS Docket No. 22-90**

Dear Ms. Dortch,

On January 20, 2023, David Clark (CSAIL/MIT) and I (KC Claffy) met via videoconference with Mr. Ken Carlberg, Chief Technologist in FCC's Public Safety and Homeland Security Bureau, and Padma Krishnaswamy of FCC's Office Of Engineering and Technology to discuss the FCC's Notice of Inquiry regarding *Secure Internet Routing.*

We discussed some of the many thoughtful comments in the FCC proceeding [1], which are an indication of the broad understanding of the severity of the problem, and which reflect a wide range of opinion on the most beneficial role for the FCC and other stakeholders in addressing the problem. Several U.S. government agencies, including the DHS and a joint filing by the DOD and DOJ, have urged the FCC to take action [42, 43]. Other commenters, including from other U.S. government agencies, emphasized the challenges of regulation in this domain, concerns about the U.S. government's own rate of adoption of routing security practices, and the importance of non-regulatory approaches including increased engagement with multistakeholder fora such as the Communications Security, Reliability, and Interoperability Council (CSRIC) to develop and promote deployment of enhanced BGP security practices [1]. Tension is clearly increasing on this topic, as multistakeholder activities have continued for over a decade, with the most optimistic predictions for deployment of protocol-based solutions to path hijacks estimating at least another decade. In the meantime, the risk and prevalence of both accidental and malicious BGP hijacks grows, rendering even the largest companies in the world victims of hijacks [8].

We commented in this proceeding last year, responding to several questions in the Notice of Inquiry. We offer additional perspectives as academic researchers who have studied Internet architecture, routing, and interconnection topology for decades. We summarize and elucidate our discussion topics and include related references below.

1 Potential harms of BGP hijacks

We discussed the serious harms that can result from hijacks. Depending on the attack objective and strategy, the harm can affect the user of a service, the service itself, the owner of the hijacked addresses (whether users or servers), or a third party. The simplest harm that can result from a hijack is that traffic goes to the wrong part of the Internet, where it is then discarded, resulting in **loss of availability.** A more pernicious harm is **server impersonation**, where a rogue endpoint carries out an exchange which seems to the victim to be with a legitimate party. Recent sophisticated attacks to steal crypto currencies [2], including from wallets hosted in the AWS cloud [13, 8] have demonstrated the use of targeted hijacks to intercept and deviate traffic in ways that allow theft of user credentials. A third potential harm is **analysis, inspection, or modification of traffic** that an attacker intercepts via a hijack. A fourth way to exploit hijacking is **source impersonation**. A common example is to send spam [32, 44], in which case the harm (aside from the harms that arise from spam) is to the owner of the address block, which may acquire a poor reputation and thus suffer blocklisting, reducing its value. Another impersonation use of hijacking is to generate fake user activity, for example to conduct ad fraud [47]. The attacker creates a web page, fills it with advertisements, and then creates fake users ostensibly located in the hijacked address block to click on these ads. In addition to the owner suffering loss of reputation of the address block, the advertisers suffer the financial harm of having to pay for the fraudulent clicks.

Although there have been anecdotal examples of such harms, there is no systematic source of such data, which presents a challenge when trying to evaluate the relative costs and benefits of action and inaction.

2 Persistent barriers to secure BGP solutions

We discussed barriers to improving BGP security against hijacks, and we agreed to include some related references as followup, which we hereby do. Over the last 30 years, over 20 proposals to secure BGP have come out of academia, industry and the Internet Engineering Task Force (IETF) [37, 20, 14, 48, 46, 12, 16, 38, 45, 28, 19, 33, 31, 18, 7, 15, 41, 21, 22]. As early as 2009 researchers began to survey the array of efforts and analyze why they failed to gain traction [27, 5] and have continued to undertake such surveys [36, 26, 39]. Researchers have also explored how to overcome the counter-incentives to deployment of protocol-based approaches to routing security [10, 11], and whether partial deployment of such approaches can be sufficiently effective to justify their promotion [24]. The OECD [29], ICANN [35], and BITAG [4] have all recently published reports with extensive references related to routing security research and challenges.

The recognized best current practice in routing security is a process called *Route Origin Validation* (ROV), which the IETF specified in 2013 as a mechanism to mitigate the risk of *origin hijacks* (the simplest form of hijack) [34], ROV uses a Resource Public Key Infrastructure (RPKI) (i.e., authoritative database maintained outside of BGP to support cryptographic signatures (certificates) that authorize designated ASes to originate address blocks; routers drop BGP announcements that do not match these certificates. RFC 6811 [34] specifies the ROV protocol with important caveats: its dependence on the integrity of the database used to validate routes, and its inability to prevent path hijacks. In particular, an attack can impersonate the valid source AS by appending it to a forged BGP announcement (recently observed in the wild [30]). RFC 6811 cautioned: "*..this system should be thought of more as a protection against misconfiguration than as true 'security' in the strong sense.*"

Aiming for a more rigorous approach to protecting against both origin and path hijacks, for many years the IETF's Secure Interdomain Routing Working Group discussed, debated, and designed a new variant of BGP called BGPsec [22], finally documented in 2017 in RFC 8205. Cryptographic attestation of paths requires propagation of a new layer of cryptographic transaction at each hop, which is computationally expensive but also poses a router-level (rather than AS-level or prefix-level) key distribution challenge, since every router must have its own public key signed by a certificate authority.

In addition to the standardization challenges that protocols face, other barriers arise from the design and operation of BGP, and the larger Internet ecosystem in which it operates. Several such barriers relate to misalignment of economic and operational incentives that prioritize other aspirations in tension with routing security.

Collectively, network operators have an interest in a well-performing, reliable public Internet. However, endpoints rather than transit ISPs are the most common target of hijacks and thus primarily benefit from reduced hijacks, while transit ISPs bear the cost of deployment of mitigations. Investment in the mechanism and inevitable errors associated with the learning curve in deployment will increase transit ISPs' costs, making them less competitive. There is thus a *last mover advantage*, i.e., the first ISPs to deploy today's BGP security mechanism may see no real benefit, either to themselves or their customers.

Operators are aware of the risk that (at least early) authoritative references will contain errors. Whether via the RPKI or other methods, attempts to discern malicious (hijacks) from benign behavior in real-time run the risk of impairing some legitimate activities that will cause some reduction in availability of the basic packet carriage service – the ISP's primary goal. Although hijacks themselves can hinder availability, this risk of false positives leaves ISPs in an intractable position trying to balance integrity against availability. For many operators, unintended impairment to availability that arise from attempting to prevent hijacks is a larger threat than the occasional route hijack—at least to the ISP if not their customers.[1]

Finally, different parts of the world, and/or different business sectors, will have different counter-incentives to deployment of security improvements. A workable solution to a security problem must balance these considerations.

[1] A recent OECD report: "For a network operator whose main concern is maintaining connectivity for its customers, turning on a solution that may drop a customer's traffic by mistake poses a substantial risk" (p.34)."Taking the example of RPKI, once an AS decides to implement ROV filtering fully (e.g., by discarding invalid responses), any invalid route would be discarded. In a large ISP with many customer ASes, there is a chance of mistakenly dropping a customer's traffic when turning on ROV filtering (e.g., if they have an incorrect ROA, for example)." (p.42) [29]

3 Limitations of MANRS as voluntary initiative

We discussed the strengths and limitations of the MANRS initiative, which was mentioned in the NOI. MANRS specifies four practices for participating networks, two of which correspond to the RPKI/ROV steps of registering authoritative information about one's prefixes, and verifying BGP announcements against authoritative information. The exact wording of these two practices are: (1) *Prevent propagation of illegitimate routes from customer networks or one's own network.*; and (2) *Document in a public routing registry the prefixes that the AS will originate.*[2]

To conform with the first practice, a MANRS member must verify two aspects of an announcement from a customer: it must confirm that the customer has used an ASN that it is legitimately allowed to use, and for any prefix originated by that customer, that the ASN is allowed to announce that prefix. This step blocks simple hijacks based on an invalid origin announcement. It does not prevent more complicated hijacks based on illegitimate *path* announcements. We call this test *verification* of the correctness of a BGP announcement from a customer, in contrast to the *validation* performed by the ROV protocol against the RPKI. MANRS does not specify how a member AS should verify the assertions of its customers, and in particular *does not require the use of RPKI/ROV (ROAs) in this verification*. The AS can use ROAs, or can verify against information in the Internet Routing Registry (IRR), or rely on a private arrangement with its customer. (The MANRS requirements do currently specify that network operators must *encourage* their customer network operators to register ROAs.)

The MANRS initiative has a key strength: it illustrates that ISPs can institutionalize their recognition of the need for a collective commitment to operational practices to reduce threats to the routing system. However, as the FCC observed in its NOI, the MANRS program has had limited success, with only a few hundred ASes joining the program. Most of these ASes use the IRR rather than the more authoritative (but more complex) RPKI. Many of the largest ISPs do not participate in MANRS.

We believe the limited success of MANRS is rooted in misaligned incentives that manifest in three ways. First, although if consistently implemented, the MANRS practices will reduce the incidence of invalid origin hijacks, there is no direct relationship between the action of any given MANRS member and the overall security of the Internet, or even the security of any customer of a MANRS member.

Second, the current MANRS practices, even the cryptographically protected RPKI/ROV options, only aim to prevent origin hijacks rather than path hijacks, which many providers believe does not justify the cost and complexity of participation.

Third, there is currently insufficient auditing of conformance to lend confidence to the assumption of consistent implementation. Our independent assessment in 2022 found a significant fraction (16%) of members do not actually conform with the practices [9, 3]. Instead the MANRS Observatory [25] reports trends of all ASes on the Internet, without focusing on or identifying the conformance of individual members. More rigorous auditing would be expensive and further reduce the incentive to participate.

The essential point is that the MANRS practices are not structured to provide incremental benefit, thus achieving participation still relies on the good will of participants. Some ISPs are not convinced the MANRS practices are worth the investment given that they do not offer protection against path hijacks, which are arguably not much harder to perpetrate than origin hijacks. Furthermore, a member ISP cannot leverage its participation in MANRS to advertise to its customers that it provides additional security benefits that non-MANRS participants cannot provide. The current set of MANRS practices thus fails to align the incentives of ISPs and their customers with improved routing security.

The varying deployment of the MANRS practices illustrates that some operators are more motivated than others to invest in security. An ideal set of operational practices would bring security benefits to the networks that choose to deploy them. To accommodate the reality that some operators will not have the resources to invest in stronger security, an ideal set of operational practices would also enable customer ASes *who are not conformant to them* to improve their own protection from hijacks by obtaining transit service from a network who is conformant. This aspiration drives our thinking, and our final topic of discussion at this meeting.

[2]The other two practices are: Ensure correct contact information for addresses is in public databases; and prevent traffic with spoofed source IP address from leaving one's network. In 2019 the MANRS program expanded to include a slightly different set of actions specific to CDNs and IXPs, accounting for the different characteristics of such networks. For example, the anti-spoofing filters can be challenging and thus risky to implement at large IXPs, so this action is not required for IXP members of MANRS.

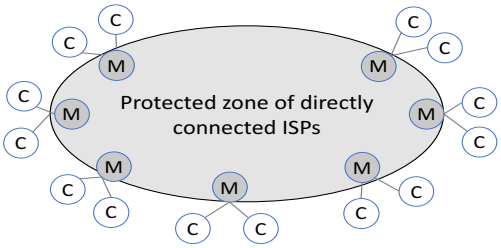

Figure 1: *If there were an operational practice that prevented path hijacks from effectively entering the zone, then for an AS directly connected to the zone, or an AS in the zone, any prefix they originate and announce to the zone will be protected from both invalid origin and path hijacks in the zone. As a result, an AS directly connected to the zone will not receive from the zone a BGP announcement that constitutes a hijack.*

4 Toward a Routing Zone of Trust

We believe that the network operator community could develop a set of enhanced practices, such that accountable conformance to these practices within a *connected* set of ISPs could overcome the essential barriers of previous routing security frameworks, as well as achieve protection against path hijacks. We have discussed this idea at a high-level in previous work [6, 40], and are exploring the application of these ideas to interdomain routing security. Success would require multistakeholder participation, and the FCC could play a constructive coordination role.

Imagine a connected region of the Internet composed of ISPs that commit to perform Route Origin Validation on all announcements coming into that region. Then, within that region there will be no propagation of origin hijacks. Furthermore, any AS that directly attaches to that region (i.e., the AS is a customer of an ISP in that region, not necessarily in the region itself) cannot be the victim of an origin hijack inside that region (Figure 1). Similarly, if a new set of operational practices prevents the propagation of *path hijacks* in the region, then customers directly attached to the region will not be the victim of a path hijack inside that region. We call this region a *zone of trust* because the protection arises at the perimeter of the zone. This protection requires that ASes in the zone be able to trust that the routers at the perimeter function correctly, which will require some degree of transparency and accountability.

The idea of a coherent perimeter around a zone is missing from today's interdomain routing system. Global deployment has always been a routing security protocol design assumption. Yet, recognition that ASes themselves can be the threat actors sheds doubt on any aspiration to make BGP *globally* secure. Our premise is that creating a zone of trust through perimeter protection (a trust-but-verify regime) offers a more pragmatic approach for today's routing system. Most important is the alignment of incentives. A zone of trust approach would be able to clearly articulate the benefit that the practices of that zone are bringing to their customers.

We define such a zone of trust as including: (1) any tier 1 provider that commits to the trusted operational practices; and (2) any AS that commits to the practices and has at least one transit provider that is already in the zone.[3] We elaborated on (and updated) an analysis in our earlier comment to this NOI that such a coherent topological region exists today, in the context of the MANRS initiative. On 1 December 2022, MANRS had 747 ISP members (as well as 22 MANRS CDN member organizations) participating with 906 and 25 ASNs, respectively [17]. (Some ISPs participate with more than one AS.) To quantify the coherent region, we start with the Tier 1 providers that are MANRS members, and recursively examine their customers' ASes to identify which are also members of MANRS. Using CAIDA's AS2Org and CAIDA's AS Relationship data sets to infer customer relationships [23], we find a connected region with 452 members operating 563 ASNs. There are 25,330 customer ASes (owned by 22,854 organizations) directly connected to this region.[4] These numbers illustrate the potential of an effort that leverages existing practices (ROV), existing institutions (MANRS), and existing capabilities (Internet topology analysis) to provide enhanced protections against routing hijacks. Further discussion and analysis is required to understand the implications of various interconnection scenarios, including multi-homed customers, and peering between MANRS and non-MANRS members.

[3] A more restrictive definition would require that *all* transit providers of a zone member be in the zone. Residual risks for both definitions is an important topic of analysis.

[4] This 25,330 number uses a loose definition. It considers all customers of this core set, even if the customer also has another provider not in the core. The stricter definition includes the (15,793) customers (14,259 organizations) who only have providers in the core region.

5 Summary

There is currently no consensus as to the next step to secure BGP beyond the simplest type of hijacks. As of 2023, BGPsec has no production deployment, and arouses significant controversy over the operational feasibility of its key management aspects. Although the cost of deployment of any given solution is hard to quantify, the cost of doing nothing is even harder to quantify. What we know is that incentives are badly misaligned: ISPs must bear the cost of action, but users bear the cost of inaction. And even those who are willing to invest their way to increased routing security cannot effectively do so due to the collective-action characteristics of the problem. We are reminded of Coase's theory that the market can solve any problem so long as all the externalities are internalized. The challenge is how to internalize them.

Despite the evolving nature of the Internet, there are fundamental architectural and economic constraints the prevent a perfect solution to routing security. There is definitely no zero-cost solution. We accept these limitations, and propose the development of empirically-grounded enhancements to a well-known set of routing security best practices, in hope of measurably improving routing security for any interested network.

The current MANRS practices describe what an individual member of MANRS is expected to do. Performing ROV on origin announcements of directly connected customers will reduce the propagation of origin hijacks, but it is difficult to quantify the benefit. However, the current MANRS framework has nothing to offer directly to ASes that want protection from having their addresses hijacked, or having their packets to a distant prefix hijacked. The insight we believe requires further exploration is that if there is a *region* of the Internet with a coherent perimeter, and with practices limiting the malicious traffic entering that region, then the operational practices, if enforced, can provide much stronger guarantees against abuse for those who join, and thus incentive to participate.

We propose an applied research direction that leverages existing capabilities and institutions to make measurable progress to prevent route hijacks. In particular, we propose to leverage the existing RPKI infrastructure and ROV capabilities, the existing MANRS framework, and current techniques for collecting and analyzing interdomain (BGP) topology data, to create a new framework for protection against not only BGP origin but also path hijacks. The MANRS initiative has already demonstrated that ISPs do recognize the need for a collective commitment to a set of operational practices to improve routing security. As steward of the MANRS initiative, the Internet Society has also launched a working group to explore enhancements to the MANRS practices. We proposed this enhancement directly to the MANRS steering committee in August 2022.

A key insight of our proposal is that *topology matters*, and that there is a coherent core of ISPs that emerges organically in the ecosystem, which we can leverage to create a *zone of trust*, a region that protects not only all networks in the region, but *all directly attached customers*. We consider it a worthy line of inquiry to extend this existing connected region into a BGP zone of trust. Leveraging the resulting connected zone of trust offers immediate benefit to potential targets of path hijacks as well as origin hijacks. The result is a virtuous circle, where customers benefit from choosing ISPs committed to the practices, and ISPs (thus) benefit from committing to the practices.

We suggested to the FCC that a multistakeholder workshop or otherwise structured conversation to discuss and rigorously evaluate such a direction might be prudent at this stage. Specific topics at such a workshop could include elaborations of the topics we discussed, e.g., how to get better data on the relative harms of action versus inaction; how to develop and evaluate a set of operational practices that could provide more incentive alignment than MANRS; what data sources and data science (e.g., topology analysis, conformance checking) are required to support a trust zone approach; subtleties of interconnection scenarios that may complicate conformance; and comparison of proposed approaches in terms of costs, operational complexity, and benefits. We are currently refining our own thinking about how to create a BGP trust zone and hope to submit a more wdetailed proposal to this proceeding (if it is still open) next month. In the meantime we wanted to thank the FCC for their efforts in this area and offer any assistance we could provide to inform their analysis of the challenges and opportunities to improving the security of global routing in the Internet.

KC Claffy	David Clark
Research Scientist	Research Scientist
CAIDA/UC San Diego	CSAIL/MIT

References

[1] Filings to No. 22-90. FCC Notice of Inquiry. In the Matter of Secure Internet Routing, 2022. https://www.fcc.gov/ecfs/search/search-filings/.

[2] M. Apostolaki, A. Zohar, and L. Vanbever. Hijacking Bitcoin: Routing Attacks on Cryptocurrencies. In *2017 IEEE Symposium on Security and Privacy (SP)*, May 2017.

[3] Ben Du. Studying Conformance of MANRS Members, January 2023. https://blog.caida.org/best_available_data/2023/01/21/studying-conformance-of-manrs-members/.

[4] Broadband Internet Technical Advisory Group. Security of the Internet's Routing Infrastructure, 2022. https://www.bitag.org/documents/BITAG_Routing_Security.pdf.

[5] K. Butler, T. R. Farley, P. McDaniel, and J. Rexford. A Survey of BGP Security Issues and Solutions. *Proceedings of the IEEE*, Jan 2010.

[6] D. Clark and k. claffy. Trust zones: A path to a more secure internet infrastructure. *Journal of Information Policy*, 11:38, 2021-08.

[7] A. Cohen, Y. Gilad, A. Herzberg, and M. Schapira. Jumpstarting BGP Security with Path-End Validation. In *ACM SIGCOMM*, 2016.

[8] Dan Goodin. How 3 hours of inaction from Amazon cost cryptocurrency holders $235,000, Sep 2022. https://arstechnica.com/information-technology/2022/09/how-3-hours-of-inaction-from-amazon-cost-cryptocurrency-holders-235000/.

[9] B. Du, C. Testart, R. Fontugne, G. Akiwate, A. C. Snoeren, and k. claffy. Mind Your MANRS: Measuring the MANRS Ecosystem. In *ACM Internet Measurement Conference*, 2022.

[10] P. Gill, M. Schapira, and S. Goldberg. Let the market drive deployment: A strategy for transitioning to BGP security. In *ACM SIGCOMM Computer Communication Review*, 2011.

[11] S. Goldberg. Why is it taking so long to secure internet routing? *Communications of the ACM*, 57(10), 2014.

[12] G. Goodell, W. Aiello, T. Griffin, J. Ioannidis, P. D. McDaniel, and A. D. Rubin. Working around BGP: An Incremental Approach to Improving Security and Accuracy in Interdomain Routing. In *ISOC Symposium on Network and Distributed Systems Security*, 2003.

[13] D. Goodin. Suspicious event hijacks Amazon traffic for 2 hours, steals cryptocurrency, Apr 2018. https://arstechnica.com/information-technology/2018/04/suspicious-event-hijacks-amazon-traffic-for-2-hours-steals-cryptocurrency/.

[14] M. G. Gouda, E. N. Elnozahy, C.-T. Huang, and T. M. McGuire. Hop integrity in computer networks. *IEEE/ACM Transactions on Networking*, 10(3), Jun 2002.

[15] A. Heffernan. RFC 2385: Protection of BGP Sessions via the TCP MD5 Signature Option, Aug 1998.

[16] Y.-C. Hu, A. Perrig, and M. Sirbu. SPV: Secure path vector routing for securing BGP. *ACM SIGCOMM Computer Communication Review*, 34(4), 2004.

[17] Internet Society. Mutually Agreed Norms for Routing Security (MANRS) network operator participants. https://www.manrs.org/netops/participants/.

[18] J. Israr, M. Guennoun, and H. T. Mouftah. Credible BGP – Extensions to BGP for Secure Networking. In *Fourth International Conference on Systems and Networks Communications*, Sep 2009.

[19] J. Karlin, S. Forrest, and J. Rexford. Pretty Good BGP: Improving BGP by Cautiously Adopting Routes. In *IEEE International Conference on Network Protocols*, Nov 2006.

[20] S. Kent, C. Lynn, and K. Seo. Secure border gateway protocol (S-BGP). *IEEE Journal on Selected areas in Communications*, 18, 2000.

[21] M. Lepinski and S. Kent. An Infrastructure to Support Secure Internet Routing. RFC 6480 (Informational), Feb 2012.

[22] M. Lepinski and K. Sriram. RFC 8205: BGPsec Protocol Specification, Sep 2017.

[23] M. Luckie, B. Huffaker, A. Dhamdhere, V. Giotsas, and kc claffy. AS Relationships, Customers Cones, and Validations. In *ACM IMC*, 2013.

[24] R. Lychev, S. Goldberg, and M. Schapira. BGP security in partial deployment: is the juice worth the squeeze? In *ACM SIGCOMM*, pages 171–182, Aug 2013.

[25] MANRS. MANRS Observatory. https://www.manrs.org/manrs-observatory/.

[26] A. Mitseva, A. Panchenko, and T. Engel. The state of affairs in BGP security: A survey of attacks and defenses. *Computer Communications*, 124:45–60, Jun 2018.

[27] M. O. Nicholes and B. Mukherjee. A survey of security techniques for the border gateway protocol (BGP). *IEEE Communications Surveys Tutorials*, 11(1):52–65, 2009.

[28] O. Nordstrom and C. Dovrolis. Beware of BGP attacks. *ACM CCR*, 34(2), 2004.

[29] OECD. BGP incidents, mitigation techniques and policy actions, 2022. https://www.oecd-ilibrary.org/science-and-technology/routing-security_40be69c8-en.

[30] L. Oliver, G. Akiwate, M. Luckie, B. Du, and k. claffy. Stop, DROP, and ROA: Effectiveness of Defenses through the Lens of DROP. In *ACM Internet Measurement Conference*, 2022.

[31] J. Qiu and L. Gao. Hi-BGP: A Lightweight Hijack-proof Inter-domain Routing Protocol, 2006.

[32] A. Ramachandran and N. Feamster. Understanding the network-level behavior of spammers. In *ACM SIGCOMM Computer Communication Review*, volume 36. ACM, 2006.

[33] P. Reynolds, O. Kennedy, E. G. Sirer, and F. B. Schneider. Using External Security Monitors to Secure BGP. *Transactions on Newtorking*, 2006.

[34] J. Scudder, R. Bush, P. Mohapatra, D. Ward, and R. Austein. RFC 6811: BGP Prefix Origin Validation, Jan 2013.

[35] Security and Stability Advisory Committee. SSAC Briefing on Routing Security, 2022. https://www.icann.org/en/system/files/files/sac-121-en.pdf.

[36] M. S. Siddiqui, D. Montero, R. Serral-Gracia, X. Masip-Bruin, and M. Yannuzzi. A survey on the recent efforts of the Internet Standardization Body for securing inter-domain routing. *Computer Networks*, 80:1–26, Apr 2015.

[37] B. R. Smith and J. J. Garcia-Luna-Aceves. Securing the border gateway routing protocol. In *Global Telecommunications Conference*, Nov 1996.

[38] L. Subramanian, V. Roth, I. Stoica, S. Shenker, and R. H. Katz. Listen and Whisper: Security Mechanisms for BGP. In *Symposium Networked System Design and Implementation,*, 2004.

[39] C. Testart. Reviewing a Historical Internet Vulnerability: Why Isn't BGP More Secure and What Can We Do About it? In *Telecommunications Policy Research Conference*. SSRN, Aug 2018.

[40] C. Testart and D. Clark. A Data-Driven Approach to Understanding the State of Internet Routing Security. In *Telecommunications Policy Research Conference*. SSRN, Sep 2021.

[41] J. Touch, A. Mankin, and R. P. Bonica. RFC 5925: The TCP Authentication Option, Jun 2010.

[42] U.S. Department of Homeland Security Cybersecurity and Infrastructure Security Agency (CISA). NOTICE OF INQUIRY. PS Docket No. 22-90. In the Matter of Secure Internet Routing, Feb 2022. https://www.fcc.gov/ecfs/document/1022806680214/1.

[43] U.S. Department of Justice National Security Division (Matthew G. Olsen) and U.S. Department of Defense Acquisition and Sustainment (William A. Laplante). Public Comment to NOTICE OF INQUIRY. PS Docket No. 22-90. In the Matter of Secure Internet Routing, Sep 2022. https://www.fcc.gov/ecfs/document/1091496862125/1.

[44] P.-A. Vervier, O. Thonnard, and M. Dacier. Mind Your Blocks: On the Stealthiness of Malicious BGP Hijacks. In *Proceedings 2015 Network and Distributed System Security Symposium*, San Diego, CA, 2015.

[45] T. Wan, E. Kranakis, and P. C. van Oorschot. Pretty Secure BGP, psBGP. In *Proceedings of the 2005 ISOC Symposium on Network and Distributed Systems Security*, San Diego, 2005.

[46] R. White. Securing BGP Through Secure Origin BGP - The Internet Protocol Journal - Volume 6, Number 3. *The Internet Protocol Journal*, 6(3), Sep 2003.

[47] White Ops and Google. The Hunt for 3ve: Taking down a major ad fraud operation through industry collaboration. Technical report, Google, Nov 2018.

[48] X. Zhao, D. Pei, L. Wang, D. Massey, A. Mankin, S. F. Wu, and L. Zhang. Detection of invalid routing announcement in the Internet. In *Proceedings International Conference on Dependable Systems and Networks*, 2002.