# A Data-Driven Approach to Understanding the State of Internet Routing Security

Cecilia Testart and David Clark

February 1, 2020

## 1 Introduction

This paper is concerned with the global routing system of the Internet, the serious security vulnerabilities in this system, the persistence of these vulnerabilities and the reasons for this persistence, the multi-disciplinary considerations that shape this space, and some recent trends that may finally improve the security of this critical system. In short, the forces that have shaped this space, and have contributed to its persistent insecurity, do not center on lack of technical proposals, but on economic issues (incentives, externalities, and first mover disadvantage), the international character of the Internet (which precludes simple domestic regulatory interventions and makes consensus building more challenging) and the presence of actors in the ecosystem with adverse interests.

Until recently there had been little real-world evidence of the levels of malicious activity enabled by BGP flaws and of how operational practices can indeed limit the impact of malicious activity. This data-driven perspective provides a strong justification to push for better security and for the scheme proposed in this paper.

This paper is organized as follows: section 2 introduces routing, the flaws in the protocol used for routing and why routing security is a difficult and persistent problem. Section 3 presents evidence of malicious activity from technical measurement work. Section 4 examines a path to partially solve the problem and the evidence available of its benefits. In Section 5 we explain issues not covered by this path, discuss a solution formulated by the IETF and its challenges, and finish by proposing a new scheme based on mutual agreements of ISPs to form a routing *zone of trust*. We conclude in section 6.

## 2 An Introduction to Global Routing and its Flaws

The Internet is composed of *Autonomous Systems* (ASs), which interconnect to form the global Internet. These parts are called "autonomous" because they are realized by independent entities that make their own decisions about business plans, investment, interconnection and other policies that in total affect how the resulting interconnected global Internet works.

There are about 70K active ASs in the Internet today, which range in size from large broadband access providers like Comcast, ATT, Verizon and the like in the U.S., global networks that interconnect ASs in different parts of the world (the so-called *Tier-1* providers), and many thousands of enterprises that connect to the edge of the Internet for access. The operators of the ASs are often direct competitors, but must still decide whether to establish interconnection points, and on what terms.

Blocks of Interned addresses have been allocated to these various ASs, which they assign to the devices connected to their networks.[1] These addresses, which are in the header of the packets that cross the Internet, are used by the routers to determine the path along which to send the packets toward their destination.

For the process of packet forwarding to work, it is necessary for routers in the Internet to know which ASs are using which address blocks. This information is propagated across all the ASs in the Internet using

[1] In the technical literature, these address blocks are called *prefixes*, because an address block is defined by the first part (the prefix) of the address. Address blocks can be of various sizes, so a prefix is specified by the first address in the block and the number of bits of the address that define the prefix. Internet addresses (in version 4) are usually written as four decimal digits separated by periods. A typical prefix looks like this: 76.96.217.0/24, where 76.96.217.0 is the first address in the block and the /24 means that the all the addresses with the same first 24 bits are in the block. There are 256 addresses in this prefix.

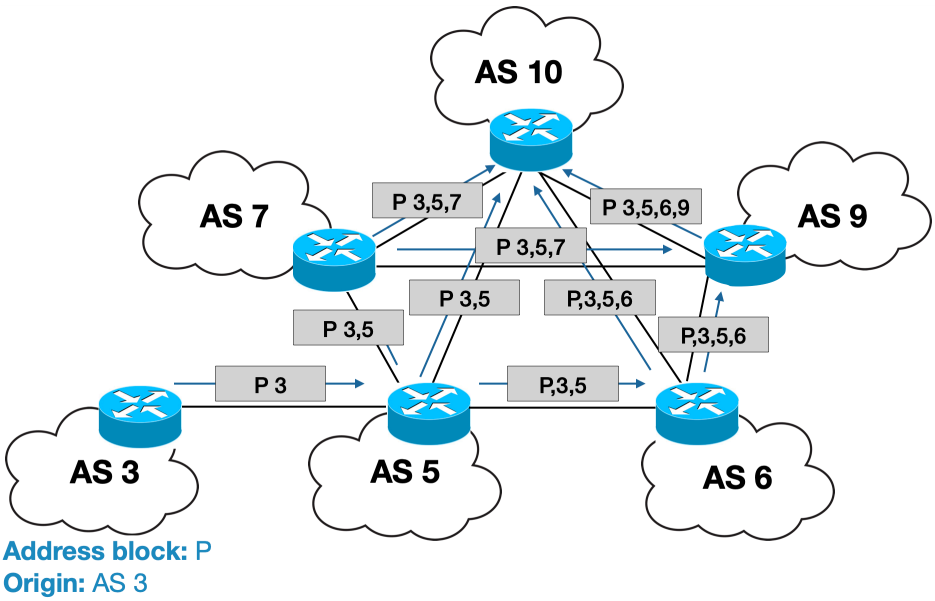


Figure 1: *AS 3, which uses address block P originates a BGP announcement to his neighbor AS5, which in turn forwards the announcement to its neighbors AS 6,and 7, and so on. At each hop, the announcement lists the sequence of ASs that provide the path back to AS 3 and address block P.*

a global routing protocol called the *Border Gateway Protocol*, or BGP. BGP (and it persistent security vulnerabilities) is the topic of this paper.

## 2.1 How BGP works

At a simplified level, the operation of BGP is straight-forward. Each AS tells its directly connected neighbor ASs the address blocks that it is using. This step is called *originating* a BGP announcement. That neighbor in turn passes that information to *its* neighbors, and so on, until the information has propagated everywhere. As each AS forwards this information along, it appends its AS number to the message, so at any point the message is the series of ASs that define the path back to that address block. Figure 1 illustrates the process of announcement propagation. AS3, which is using address block P, originates the BGP announcement P,3 to its neighbor AS5. This AS adds its AS number and forwards the announcement P,3,5 to its neighbor, and so on.

Figure 1 also illustrates how BGP route selection process operates when a network receives multiple announcements for the same IP address block. AS9 receives two announcements of the address block P with routes back to AS3, one via AS6 and one via AS7. AS9 must make two decisions when it receives these messages. It must select among them the one it will use and it must decide to which neighbors it will forward the announcement. ASs must pick the "best" announcement to forward onward, and there are a set of route selection rules to decide which one to use. An AS also has the option of dropping the announcement and not forwarding it at all.

This explanation sweeps a large number of details under the rug. For example, ASs can connect to each other as *peers* or in a *subscriber-provider* relationship. The customary rules for which announcement to forward are different in those two cases, and there are lots of rules like this. But these rules are not essential at this point to understand the basics of BGP.

## 2.2 Abuse of BGP

The critical security flaw with BGP is well-known: a rogue Autonomous System can announce a falsehood into the global routing system, i.e., a false assertion that it uses or is the path to a block of addresses that it does not in fact have the authority to announce. Traffic to addresses in that block may then travel to that rogue AS, which can drop, inspect, or manipulate that traffic.

Figure 2 illustrates this case. AS 666 announces that it holds address block P, even though it does not.

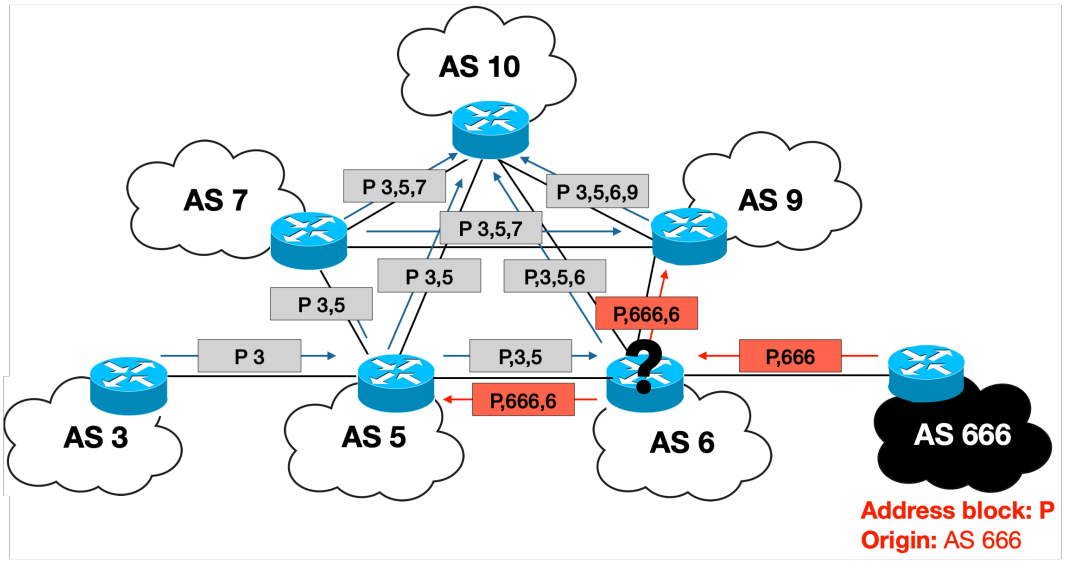


Figure 2: *AS 666 attempts to hijack address block P by falsely asserting that it has been allocated those addresses. When AS 6 receives this false announcement, it must pick whether to forward this one or the one it received from AS 5. Since the announcement from AS 666 is shorter, it might choose to forward it. Any AS that receives one of these announcements (in red) and acts on it will forward traffic intended for address block P to AS 666 as opposed to AS 3.*

When AS 6 receives this announcement, it must pick between the two competing announcements–the one from its neighbor AS666 or the longer one from AS3 via AS5. If it picks the shorter path and forwards it, a number of ASs may send traffic intended for address block P to AS666.

Harms that can result from a hijack can be grouped in three types:

- **Loss of availability:** The most simple harm that can result from a hijack is that traffic goes to the wrong part of the Internet, where it is then discarded. This outcome leads to a loss of availability between the intended communicating parties.
- **Fake end-point:** A more pernicious kind of harm is that the hijacked traffic goes to a rogue end-point that mimics the behavior of the intended end-point and carries out an exchange which seems to the victim to be with a legitimate party. This attack can lead to theft of information such as user credentials, which can then be exploited by the malicious actor. It can also lead to the download of malicious software, or malware, onto the victim's computer. For instance, in April 2018, hackers stole crypto currencies from a crypto currency wallet housed in the AWS cloud by using targeted BGP announcements to deviate traffic into a fake replica of the wallet where users credentials where stolen.[2]
- **Abuse IP address block:** A more indirect kind of harm happens when the hijackers send traffic using the hijacked address block to send spam or fake users activity. Spammers have been know to use hijacks for spam campaigns.[3] And hijacks have also been used in an online ad fraud scheme that faked user activity from hijacked addresses.[4]

## 2.3 The challenge of validating information in BGP

Creating a database of records that specifies what information in BGP announcements should be valid might seem as an obvious solution. An AS that receives an announcement would check the announcement against the database and reject it if it does not match.

[2] Dan Goodin. *Suspicious event hijacks Amazon traffic for 2 hours, steals cryptocurrency.* en-us. Apr. 2018. URL: `https://arstechnica.com/information-technology/2018/04/suspicious-event-hijacks-amazon-traffic-for-2-hours-steals-cryptocurrency/` (visited on 11/01/2019).

[3] Anirudh Ramachandran and Nick Feamster. "Understanding the network-level behavior of spammers". In: *ACM SIGCOMM Computer Communication Review*. Vol. 36. ACM, 2006, pp. 291–302.

[4] *The Hunt for 3ve: Taking down a major ad fraud operation through industry collaboration.* Tech. rep. Nov. 2018. URL: `https://services.google.com/fh/files/blogs/3ve_google_whiteops_whitepaper_final_nov_2018.pdf?__hstc=&__hssc=&hsCtaTracking=c7b87c5c-1676-4d53-99fb-927a07720b17%7C9d63bf77-0926-4d08-b5ec-46b1a06846bc` (visited on 10/29/2019).

This idea is fine in principle but raises several critical issues in practice. The first, and perhaps most critical, is that a database of that sort gets out of date. Owners of addresses make a change to how they are using them and forget to update the database. When this happens, valid assertions will be dropped, and the resulting losses of connectivity are hard to debug.

Since the era of the telephone system, designers have incorporated databases that are supposed to be the source of authority as to the state of the network, and the results have usually been problematic. The insight that emerged is that as the elements of the network become more sophisticated and able to communicate about their state, the network itself should be the ground truth about the network.[5]

The idea that the network itself is the authoritative record of what the network is (and should be) is actually very empowering and can reduce management costs and complexity. But it opens the door to malice, because elements can now lie about their state. So this approach carries with it its built-in Achilles heel.

If the network itself cannot be the authoritative source of what should go in BGP, then the question is what should be the source. One option is to have network operators (or a selected group of them) be the authoritative source. Still operators could lie about their resources. Another option would be to have Regional Internet Registries (RIRs), which are already in charge of regional IP address block delegation, to be the authoritative source. In this case, the trust is transferred to RIRs and their infrastructure. Currently in the Internet, solutions using both these options co-exist in partial deployment.

There are other problems with a global database of record. One is that operators may not want to announce to the whole world all their operating state, for business reasons. A second is that the database itself now becomes an attractive target for malicious attack. If an attacker can change the database the scope of mischief is broad.

As the next sections will discuss, current proposals to improve the security of BGP do depend on some sort of independent record of what *should* be that can be used to validate what *is*. All the issues above apply to the solutions now being put forward. In addition, there is a general observation that informs all the approaches to improved security. *Improved security is contrary to the primary goal of the Internet*, which is availability—to deliver data. Security mechanisms keep people from doing things, so by definition they interfere with availability.

## 2.4 Why insecurity is persistent

There have been over 20 proposals to secure BGP coming from academia, industry and the Internet Engineering Task Force (IETF) [6]. Each of these proposals tackle BGP security in different ways. The main difference are:

- **The threat model:** solutions are focused on protecting from distinct threat models and as such they differ in the changes proposed to BGP or BGP operation.
- **Authoritative source:** as discussed in Section 2.3, some proposals have the ISPs (or a subset) be the authoritative source whereas in others it is the RIRs. Some propose to use historical BGP data.
- **Information to validate:** some proposals aim to validate IP address blocks and the AS that originated them; others look to validate the whole AS path reported in BGP messages.

[5] As a simple example, in the era of telephone service based on copper pairs, the phone companies tried to maintain a database of which pairs were in service, but workers in the field would make changes and not update the database, which then led to future confusion and failed installs. There is no way to "query" a copper pair and ask what its state is. However, with current access technology such as fiber to the home or hybrid fiber coax, the elements in the network are active, and so it is possible (for example) to send a control message down a fiber to confirm what is at the other end and whether it is working. In this way, the network itself can be the authoritative record.

[6] For technical reviews of BGP security proposals see (Martin O. Nicholes and Biswanath Mukherjee. "A survey of security techniques for the border gateway protocol (BGP)". in: *IEEE Communications Surveys Tutorials* 11.1 [2009], pp. 52–65. ISSN: 1553-877X. DOI: 10.1109/SURV.2009.090105; Kevin Butler et al. "A Survey of BGP Security Issues and Solutions". In: *Proceedings of the IEEE* 98.1 [Jan. 2010], pp. 100–122. ISSN: 0018-9219, 1558-2256. DOI: 10.1109/JPROC.2009.2034031. URL: http://ieeexplore.ieee.org/document/5357585/ [visited on 08/04/2017]; Muhammad S. Siddiqui et al. "A survey on the recent efforts of the Internet Standardization Body for securing inter-domain routing". en. In: *Computer Networks* 80 [Apr. 2015], pp. 1–26. ISSN: 13891286. DOI: 10.1016/j.comnet.2015.01.017. URL: http://linkinghub.elsevier.com/retrieve/pii/S1389128615000286 [visited on 02/15/2018]). For the life-cycle of ideas behind security proposals, see (Cecilia Testart. "Reviewing a Historical Internet Vulnerability: Why Isn't BGP More Secure and What Can We Do About it?" en. In: Washinton, DC: Social Science Research Network, Aug. 2018. URL: https://papers.ssrn.com/abstract=3141666 [visited on 09/29/2019])

- **Validation strategy:** One approach is that every router should always validate each BGP announcement; another is that a router need only validate an announcement if there are conflicting announcements.[7]

After more than 20 years from the first security proposal, there still is no consensus about all the different options explored by security proposals. Some are starting to emerge and will be discussed in section 4. A number of specific challenges involving disagreements, incentives and governance that have impeded the selection and deployment of a scheme from all the proposed options:

- Persistent disagreement as to which BGP vulnerabilities are the most important and should be prioritized. There are several points in the design of BGP that represent potential security vulnerabilities.
- Lack of agreement as to which proposals are actually practical, taking into account the issues of deployment and operation.
- Lack of framework that drives toward a consensus. Reaching agreement requires advocates to give ground, and there is no award for doing so.
- Misaligned incentives for actors. ISPs will almost certainly bear the major cost and complexity of deploying a change to BGP, but they are not the beneficiaries of the changes. It is primarily the end points that benefit, not the ISPs, that benefit from reduced hijacks.
- First-mover disadvantage. The first ISPs to deploy a mechanism may see no real benefit, either to themselves or their customers. The investment in the mechanism might increase their costs, making them less competitive.
- Global problem not easily be shaped by domestic regulation.

A final challenge to deploying schemes to improve security is the more or less inevitable imprecision of most approaches. In an idealized world, designers would invent security mechanisms that only prevent *bad* things from happening, but that is an unrealistic expectation. First, there is no precise definition of "bad". And second, it is very hard to design a mechanism that draws the line in exactly the right place, even if there is agreement as to where the line should be. So proposals either undershoot or overshoot. If they undershoot, they leave options for the malicious actors, who are quick to find and exploit them. Critics will argue that this sort of intervention is not worth doing. If proposals overshoot, they prevent (or inconvenience) too many legitimate efforts, and critics say that the harm is greater than the benefit, so they are not worth doing. This balance is the fundamental challenge for the design of any security mechanism and again, proposals to improve BGP must deal with this dilemma.

Additionally, it is very difficult to understand the boundaries of under- and over-shooting without knowing much about the real level of malicious activity in BGP. How "bad" is the hijacking problem? Would things improve by tackling only hijacks involving the IP address block and origin violations or is it too much of an undershoot? And how effective would solutions be in partial deployment? Are there some benefits at the start of the deployment or only close to full deployment?

The full answers to the questions above are still unknown. In the next two sections, we report results from technical work looking at misbehaving networks and the impact of operational security practice that provide partial answers and are helpful to grasp the full extent of hijacking activity in BGP and the current and potential effectiveness of operational security practices.

# 3 BGP hijacks

BGP hijacks have been studied for a long time but they are still a prevalent threat and concern for network operators.[8] There have been many efforts in the research community to characterize BGP hijacking events[9]

[7] We note that hijacks of IP address blocks that are not supposed to be routed in BGP can still be harmful. Indeed, it has been shown that many address blocks that are reserved according to RFCs are still routed in BGP. They are called *bogons* and unfortunately not all network filter them out.(*Cymru BGP Bogon Refence*. `https://team-cymru.com/community-services/bogon-reference/`)

[8] Pavlos Sermpezis et al. "A Survey among Network Operators on BGP Prefix Hijacking". In: *ACM SIGCOMM Computer Communication Review* 48.1 (Apr. 2018), pp. 64–69. ISSN: 01464833. DOI: `10.1145/3211852.3211862`. URL: `http://dl.acm.org/citation.cfm?doid=3211852.3211862` (visited on 09/28/2018).

[9] Mohit Lad et al. "Understanding Resiliency of Internet Topology against Prefix Hijack Attacks". In: *37th Annual IEEE/IFIP International Conference on Dependable Systems and Networks (DSN'07)*. June 2007, pp. 368–377. DOI: `10.1109/`

and to develop hijack detection systems. Most available hijacking detection system start by detecting a route change in BGP by contrasting current BGP announcements with historical announcements. The challenging part is to distinguish whether the new route is an acceptable dynamic change, an operator mistake or an intentional malicious event.

Hijacking detection systems use different approaches, heuristics, and vantage points to discard changes that seem acceptable or unintended.[10] Unfortunately, the high rate of false positive has limited the adoption of these system. On a local scale, an input ground truth can be used to detect unacceptable route changes with high confidence.[11] Unfortunately, this approach only works when ground truth is available and up to date.

Even when a hijack is properly detected, mitigating it might not be possible. Indeed the perpetrator can be connected anywhere in the Internet and to fully stop the hijack, either that network or all of its transit providers need to stop or drop the hijacking announcement. Currently, there is no established process between networks operators for such a request. Operators usually start by contacting the perpetrator network or its transit providers but they may not get an answer. Some operators disclose the event in network operators mailing lists or forum in the hope that peer-pressure would make networks answer. The victim network may chose to take an active approach and reduce the hijack impact by changing its own routing announcements in ways that will it more likely to be preferred. However, this option is not always possible.

## 3.1 How serious is the problem?

The more damaging cases of route hijacks are often reported in the press and on mailing lists. If a hijack impacts a highly visible infrastructure or service such as a Domain Name Server or a large cloud provider, it might end up being reported in blogs from routing security experts and companies who detect the event while it is ongoing. In addition, the victim may disclose the hijack in operators mailing lists or forums in an attempt to contact involved networks. However, the overall level of hijacking has not been clear. Cisco's BGPStream monitoring tool detected about $1,300$ possible hijacks between May and November 2020,[12] but many of them are known to be false positive.

To understand the current level of abuse, and the importance of seeking ways to mitigate it, we, together with collaborators at CAIDA developed a scheme to identify malicious routing announcements based on their intrinsic characteristics. Working with five years of data curated by CAIDA, we demonstrated that there are Autonomous Systems that persist as malicious players in the Internet for years, issuing malicious routing assertions and deflecting ("hijacking") traffic from its intended destination. There are about 70K active Autonomous Systems in the Internet today. Using routing data and some machine learning tools, we identified about 400 of the ASs as highly likely serial hijackers, and another 400 that are probable hijackers. This work illustrates the value of data collection and analysis, both to understand the extent of the problem, and to provide support for proposals as to how to mitigate it.[13]

---

DSN.2007.95; Pierre-Antoine Vervier, Olivier Thonnard, and Marc Dacier. "Mind Your Blocks: On the Stealthiness of Malicious BGP Hijacks". en. In: *Proceedings 2015 Network and Distributed System Security Symposium*. San Diego, CA: Internet Society, 2015. ISBN: 978-1-891562-38-9. DOI: 10.14722/ndss.2015.23035. URL: https://www.ndss-symposium.org/ndss2015/ndss-2015-programme/mind-your-blocks-stealthiness-malicious-bgp-hijacks/ (visited on 10/02/2018).

[10] Jian Qiu and Lixin Gao. *Hi-BGP: A Lightweight Hijack-proof Inter-domain Routing Protocol*. Tech. rep. 2006, p. 12; Mohit Lad et al. "PHAS: A Prefix Hijack Alert System". In: *Proceedings of the 15th Conference on USENIX Security Symposium - Volume 15*. USENIX-SS'06. Berkeley, CA, USA: USENIX Association, 2006. URL: http://dl.acm.org/citation.cfm?id=1267336.1267347 (visited on 10/02/2018); Xin Hu and Z. Morley Mao. "Accurate Real-time Identification of IP Prefix Hijacking". In: *2007 IEEE Symposium on Security and Privacy (SP '07)*. May 2007, pp. 3–17. DOI: 10.1109/SP.2007.7; Z. Zhang et al. "iSPY: Detecting IP Prefix Hijacking on My Own". In: *IEEE/ACM Transactions on Networking* 18.6 (Dec. 2010), pp. 1815–1828. DOI: 10.1109/TNET.2010.2066284; Tongqing Qiu et al. "Locating Prefix Hijackers Using LOCK". in: *Proceedings of the 18th Conference on USENIX Security Symposium*. SSYM'09. Berkeley, CA, USA: USENIX Association, 2009, pp. 135–150. URL: http://dl.acm.org/citation.cfm?id=1855768.1855777 (visited on 10/02/2018); Xingang Shi et al. "Detecting Prefix Hijackings in the Internet with Argus". In: *ACM IMC*. 2012; Johann Schlamp et al. "HEAP: Reliable Assessment of BGP Hijacking Attacks". In: *IEEE Journal on Selected Areas in Communications* 34.6 (June 2016), pp. 1849–1861. ISSN: 0733-8716. DOI: 10.1109/JSAC.2016.2558978.

[11] Pavlos Sermpezis et al. "ARTEMIS: Neutralizing BGP Hijacking within a Minute". In: *arXiv:1801.01085 [cs]* (Jan. 2018). URL: http://arxiv.org/abs/1801.01085 (visited on 10/01/2018).

[12] *Cisco BGPStream CrossworkCloud*. https://bgpstream.com/.

[13] The next sections describe relevant aspects of the work. For more technical detail, please refer to (Cecilia Testart et al. "Profiling BGP Serial Hijackers: Capturing Persistent Misbehavior in the Global Routing Table". en. In: *Proceedings of the Internet Measurement Conference on - IMC '19*. Amsterdam, Netherlands: ACM Press, 2019, pp. 420–434. ISBN: 978-1-4503-6948-0. URL: https://dl.acm.org/doi/10.1145/3355369.3355581 [visited on 10/27/2019])

### 3.2 Profiling BGP serial hijackers

In this study, instead of focusing on detecting individual hijacking events, we study networks that repeatedly perform hijacks in BGP at different points in time. We call these networks *serial hijackers*. By extracting the AS numbers of perpetrators of hijacks disclosed in an operator mailing list, we are able to find 23 networks for which we have evidence of repeated hijacks over time. As an example, one of the serial hijackers is a network called BitCanal, for which there have been repeated report of malicious activity starting in 2014[14] and which was finally disconnected by its transit provider in July 2018 after a lengthy thread in an operator mailing list exposing some of its repetitive malicious behavior.[15]

To contrast the behavior of serial hijackers, we select about 200 benign networks we do not expect to be repeatedly engaging in malicious activities. Most networks we select are subscribers to the Mutually Agreed Norms of Routing Security (MANRS) initiative.[16] Since this networks voluntarily agreed to take steps to improve routing security, we believe they would not willingly engage in repetitive malicious behavior. To this group, we manually add a few more network we can manually verify to increase the diversity of the set of benign networks.

Then we study the overall routing behavior of both serial hijackers and benign networks to characterize how serial hijacker's BGP behavior differs from that of benign networks. Through CAIDA's BGPStream library,[17] we have access to BGP data coming from over 1,400 networks in the Internet, including some of the largest providers. Using this data, we extract the IP address blocks that are being advertised in BGP and the origin Autonomous System (AS) of those announcements, and compute their visibility by counting the number of networks that have each address block and origin AS in their routing table. With this information, we build timelines of address blocks being originated by a given network and we are then able to analyze the BGP behavior of that network from the timelines of IP address blocks it originates.

Figure 3 shows the visualization of the overall BGP behavior of a benign network AS5400 and a serial hijacker network AS197426, based on the timelines of individual IP address block advertised with these networks as origin. Over the course of 5 years, the hijacker originated almost 1,500 unique IP address blocks. And we also observe a highly irregular pattern of short-lived advertisement of address blocks. To put this behavior in contrast, Figure 3a shows the equivalent behavior of a large British residential and mobile ISP, where a more steady and overall constant pattern of longer announcement times can be seen. We note, however, that legitimate networks can also exhibit irregular patterns (see the white space between lines indicating an address block was not originated at that time), often due to configuration issues of the network in question or of third-party networks. From Figure 3 it becomes clear that these two networks show wildly different long-term behavior in the global routing table.

Studying the overall BGP behavior of serial hijackers and benign networks, we demonstrate that serial hijacker behavior differs in 5 main characteristics:

1. **Intermittent presence in routing:** Serial hijacker activity in BGP is often intermittent and has *offline* periods where those networks are not advertising IP address blocks in BGP, effectively disappearing from the global routing system. In contrast, once a benign network starts advertising IP address blocks in BGP, it usually never goes offline, unless the entity is ended.

2. **Short advertisements of IP address blocks:** A large portion of the IP address blocks serial hijackers originate are advertised for short periods of time. In contrast, benign networks tend to constantly advertise an address block once they start doing so. This difference can be seen in the amount of white space in figure 3b *vs.* figure 3a.

3. **Constant changes in the number of IP address block advertised:** As a result of short IP block advertisements, serial hijackers have a high variability of the total number of IP address block being advertised over time. In figure 3b, if we sum up the IP address blocks being originated by this

[14] Doug Madory. *Sprint, Windstream: Latest ISPs to hijack foreign networks — Dyn Blog*. en-US. `https://dyn.com/blog/latest-isps-to-hijack/`. Sept. 2014. (Visited on 05/12/2019); Doug Madory. *The Vast World of Fraudulent Routing — Dyn Blog*. en-US. `https://dyn.com/blog/vast-world-of-fraudulent-routing/`. Jan. 2015. (Visited on 05/12/2019).

[15] Doug Madory. *Shutting down the BGP Hijack Factory — Dyn Blog*. en-US. `https://dyn.com/blog/shutting-down-the-bgp-hijack-factory/`. July 2018. (Visited on 05/12/2019); Krebs on Security. *Notorious 'Hijack Factory' Shunned from Web*. `https://krebsonsecurity.com/tag/bitcanal/`.

[16] *Mutually Agreed Norms for Routing Security (MANRS)*. `https://www.manrs.org/`.

[17] Chiara Orsini et al. "BGPStream: A Software Framework for Live and Historical BGP Data Analysis". In: *Proceedings of the 2016 Internet Measurement Conference*. IMC '16. Santa Monica, California, USA: Association for Computing Machinery, Nov. 2016, pp. 429–444. ISBN: 978-1-4503-4526-2. URL: `https://doi.org/10.1145/2987443.2987482` (visited on 01/25/2020).

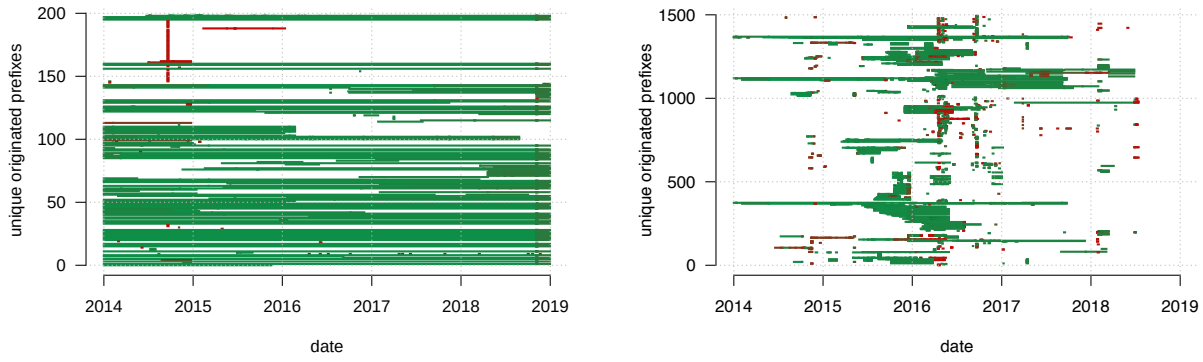


(a) Benign network: IP address blocks advertised with AS5400 (British Telecom) as origin over the course of 5 years. This network originates address blocks consistently over long time periods.

(b) Serial Hijacker network: IP address blocks advertised AS197426 (Bitcanal) as origin over the course of 5 years. This network originates a large number of address blocks over short time periods.

Figure 3: Long-term BGP behavior for a benign network and a serial hijacker network. We visualize each originated IP address block as a row on the $y$-axis and time on the $x$-axis. Red colored address block have low visibility (less than 15%).

hijacker network over time, that amount would constantly vary. Conversely, in figure 3a the sum of address blocks originated over time would be relatively stable, except for a short period of time in 2014 where there are many address blocks advertised for a short period of time. As mentioned earlier in this section, in a 5-year window, it is likely that networks had a configuration or infrastructure incident that modified their usual routing behavior for a limited time.

4. **Conflicts with other IP address blocks in the global routing table:** Serial hijackers have a larger share of the IP address blocks they advertise as being the origin having a conflict with the same IP address block being advertised with another network as origin. Although there are legitimate reason why an IP address blocks may have different networks originating them, those cases typically exhibit a different pattern. Usually legitimate address block conflicts tend to last for a long time whereas conflicts with serial hijackers are short in duration.

5. **Broad geographical distribution of address space originated:** IP address blocks used by networks in the Internet are delegated by Regional Internet Registries (RIRs). Each one of the 5 RIR covers a distinct geographical zone. As a result, most networks have all or most of the IP address block they host and advertise in BGP coming from one RIR, related to networks' main geographic location. However, serial hijacker do not only use in BGP address blocks delegated to them and tend to start advertisements in BGP for address block from multiple RIR.

## 3.3 Finding BGP serial hijackers

Using the key distinctions between serial hijacker and benign network behavior, we train a machine learning classifier to find networks in the Internet that exhibit a behavior in BGP similar to serial hijackers. Based on the extensive analysis of network's behavior in BGP, we select 52 features that capture the differences and the variability in behavior between the two categories of networks.[18]

We run our classifier on all networks that originate advertisements for at least 10 IP address blocks in our dataset. The classifier finds that 934 networks out of the 19,103 network in that group have similar

[18]For details of the machine learning classifier used and the training process refer to the technical paper. We note that the precision of the classifier is about 80% because a few benign networks in the training set are wrongly classified as exhibiting serial hijacker's behavior. This is a strong reminder that the behavior of networks selected by the classifier is not *necessarily illegitimate* (Testart et al., "Profiling BGP Serial Hijackers").

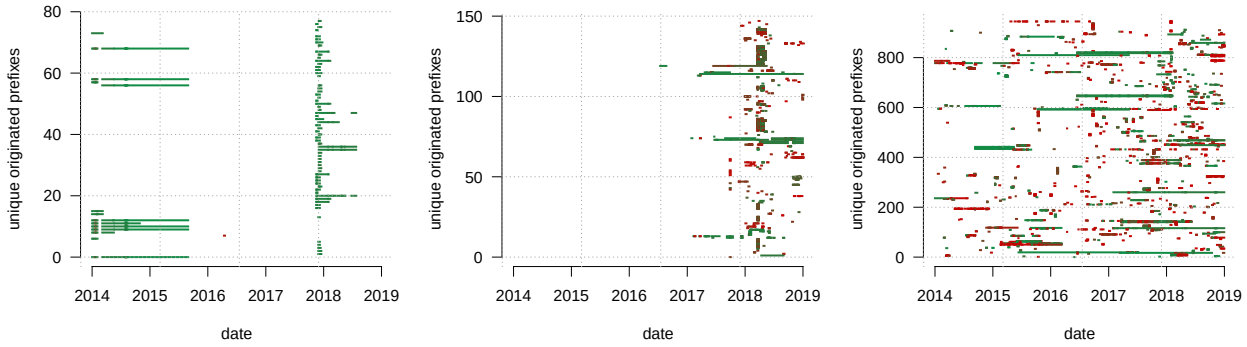


(a) AS19529, a hijacker identified by our classifier for which we found corroborating evidence of hijacking activity.

(b) AS134190, the most recent detected case of a potential serial hijacker.

(c) AS5, a *fat finger* and defunct network that is flagged by our classifier. The company that owns AS5 went out of business over 10 years ago.

Figure 4: Visualization of BGP behavior from 3 networks flagged by our classifier as having similar behavior to serial hijackers.

BGP behavior to serial hijackers. This represents 4.9% of networks in that group and 1.2% of networks in the Internet.

Networks flagged as having similar behavior to BGP serial hijackers are medium to large networks, with volatile BGP behavior and originating IP addresses coming from multiple RIRs. Figure 4 shows the visualization of the overall BGP behavior for 3 Autonomous Systems flagged by our classifier that illustrate these characteristics. Overall, flagged networks have originated more IP address blocks than non-flagged networks. Flagged networks have a median of 41 IP address blocks advertised in 5 years compared to 23 blocks for the others. And those address block tend to be advertised for much shorter periods. 50% of flagged networks advertise the address blocks they originate for less than 50 days, when only about 15% of the non-flagged network have such amount of short advertisements. Along the same lines, flagged networks tend to disappeared much more often from the global routing table, with over 25% of them being inactive at least 30% of the time when less than 0.1% of non-flagged networks have such amount of inactive time.

### 3.3.1 Categories of suspicious networks

To further scrutinize the networks flagged by our classifier as having similar behavior to serial hijackers, we look for indications of malicious behavior and misconfigurations, and investigate likely false positives.

- **Indications of malicious behavior:** Many networks flagged by our classifier can be found either in the Spamhaus Don't Route Or Peer (DROP) ASN list,[19] which contains ASN of networks controlled by "spammers, cyber criminals, and hijackers", or in the UCEPROTECT Spam blacklist [20] because of repetitive and substantial spam campaigns at the same time that the networks are active. 84 flagged networks are in the Spamhaus DROP list, making it 10 times more likely that a flagged network is in this DROP list than non-flagged networks. Additionally, we find at least one address block blacklisted for spam for 38% of flagged networks.

- **Indications of misconfiguration:** 114 networks flagged by our classifier have a private Autonomous System Number that should not be advertised in BGP according to RFC 6996 [21]. However, due to router misconfiguration, they appear at the origin of AS paths, and as the misconfigurations get corrected,

[19]This list contains Autnomous Systems Numbers that Spamhaus deemed controlled by professional cyber-criminal operators involved in "professional spam, [...], dissemination of malware, trojan downloaders, botnet controllers" (The Spamhaus Project. *DROP - Don't Route or Peer lists - The Spamhaus Project*. `https://www.spamhaus.org/drop/`).

[20]We use the UCEPROTECT Blacklist Policy Level 2 list, which includes IP address blocks where many IP address have been linked to spam campaigns in the past 7 days. (UCEPROTECT. *Blacklist Policy LEVEL 2*. `http://www.uceprotect.net/en/index.php?m=3&s=4`)

[21]This RFC provides a list of Autonomous System Numbers that are for private use only, which can be use within a network or in specific settings such as when a small network is connected to only one provider and all its traffic would go through that provider (Jon Mitchell. *RFC 6996: Autonomous System (AS) Reservation for Private Use*. en. July 2013. URL: `https://tools.ietf.org/html/rfc6996` [visited on 05/11/2019]).

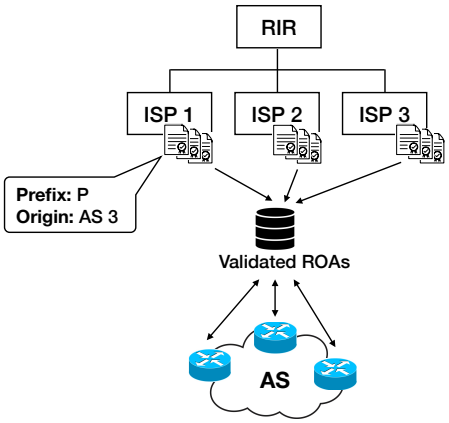


Figure 5: *The Resource Public Key Infrastructure (RPKI) is a system to secure BGP using cryptographic records. Route Origin Authorizations (ROAs) are cryptographically signed assertions specifying an address block and Autonomous System Number authorized to originate it in BGP. Using validated ROAs, ISPs drop announcements that have invalid information. This process is known as Route Origin validation (ROV). In this example, we consider ISP 1 to be the owner of AS 3 and address block P, and it issues a ROA to authorize AS 3 to originate P.*

these networks appear to have very volatile behavior. In addition, all single digit ASN networks are flagged by our classifier. Because of configuration errors linked to traffic engineering practices in some networks, these single-digit ASN networks appear to be originating many address block from other networks. A notable example of these so-called *fat finger errors* is AS5, whose overall BGP behavior is depicted in Figure 4c. AS5 is registered by a company that went out of business 20 years ago, but through fat-fingers errors, it is periodically revived.

- **False positive cases:** We find 18 networks flagged by our classifier that offer denial of service protection using BGP. These networks perform *benign hijacks* when advertising IP address blocks usually advertised by their clients that are under attack in order to attract the traffic. Additionally, 4 flagged networks are in the top 500 networks according to CAIDA AS-Rank. These very large networks are unlikely to perform serial hijacking.[22]

There are 441 networks flagged by our classifier as having similar BGP behavior to serial hijackers that do not fit in any of the categories described above. Figure 4b shows the behavior of AS134190, flagged by our classifier but not found to be in any block list. This network has originated about 150 prefixes in less than 2 years and half of them are advertised for less 12 hours.

Given that we use 5 years of BGP data for this study, many AS that are flagged by our classifier as having similar behavior to serial hijackers are not active anymore. Of the 820 flagged ASNs that are not for Private Use only, 30% are not active anymore. Of the 441 ASes that cannot be further classified in the categories mentioned above, 143 are not currently active.

In short, this work demonstrate that there are networks—the BGP serial hijackers—that repeatedly abuse BGP using the flaw described in section 2.2 without consequences. Some of these networks have been misbehaving in BGP for years and are still somehow connected to the Internet. In the next section, we describe and discuss a scheme can limit the spread of illicit announcement such as the ones coming from serial hijackers.

# 4 A path to better security

As mentioned in section 2.3, an approach to improving security for BGP will almost certainly require the creation of a database of record that specified what *ought* to happen. The Secure Inter-domain Routing (SIDR) working group of the Internet Engineering Task Force (IETF) has developed a scheme that supports

[22]There are a number of reasons why very large network providers have a more volatile BGP behavior. For instance, they tend to have more configuration and infrastructure incidents impacting their routing behavior, they tend to use multiple Autonomous System Numbers—called sibling ASNs—and change address blocks origination between them, and in general they make a large number of changes to their networks over time.

cryptographically signed assertions about the ownership and authority to announce address blocks in BGP. Figure 5 illustrates how the scheme works from the address block delegation to networks validating routing information in BGP.

An assertion is called a *Route Origin Authorization* (ROA), and the system that is used to assure their legitimacy is called the Resource Public Key Infrastructure (RPKI). The RPKI is a hierarchically organized structure in which cryptographically signed ROAs can be stored. There are five *roots of trust* for the RPKI, under the control of the five Regional Internet Registries (RIRs) that manage allocation of addresses in the Internet[23]. RIRs delegate IP address blocks to ISPs. The RPKI involves many sub-processes to fetch, sync and validate data but the key functionality for ROAs is simple: the essential feature is that the owner of an address block is given the authority by an RIR to register one or more ROAs for the addresses in the block (or a subset of the addresses in that block). A ROA specifies the address block and the AS that is allowed to originate an announcement for that block [24]. In figure 5, ISP 1 issues a ROA authorizing AS 3 to originate the address block P. Then an AS can validate that assertion and use that information in its border routers to validate information in BGP messages it receives.

If the owner of an address block has registered a suitable ROA for that block, any ISP that receives a routing assertion can confirm that the origin is valid. Referring back to Figure 2, if the owner of address block P registers a ROA that says that a BGP announcement for address block P with AS3 as origin is legitimate, then when AS6 gets the announcement P:666 from AS666, it can use the ROA to determine that it should drop P:666 and accept P:3,5. This step is called Route Origin Validation (ROV).

There are two parts to making ROAs and the RPKI scheme work. First, the owners of address blocks must register their ROAs, keep them up to date, and so on[25]. Uptake of this technology is low, although growing. Globally, about 25% of the Internet address space is protected by ROAs today[26]. Second, ISPs must use validated ROAs to validate routes in BGP[27]. ROAs do no good unless ISPs actually use them. ROAs are defined by a protocol specification,[28] but a protocol specification does not necessarily define the ways that the protocol is utilized in practice. We use the term *operational practice* to describe how a mechanism is being used—the description of the current or preferred operational practices is a complement to the specification of the mechanism.

The basic operational practice with respect to ROAs is that any AS that receives a BGP announcement can check the prefix and AS in the origin of the announcement and see if it is validated or invalidated by a ROA. In the next section, we describe the measurements we have devised to assess the extent that this practice is happening and the benefits it translate into for networks properly registering their ROAs in the RPKI. In Section 5 we describe an enhanced practice that can further enhance the benefit of ROAs.

## 4.1 Tracking ISPs that drop invalid (suspicious) routes

Although ROAs and the RPKI have been around since 2013, until recently, there was little evidence of any ISP using the information from ROAs to validate announcements in BGP. However, starting in 2019, there had been increasing anecdotal evidence of ISPs, such as AT&T, KPN and Telia, implementing RPKI filtering (ROV) and effectively droppping invalid announcements.[29] Together with collaborators at CAIDA, we developed a method to track which ISPs are dropping invalid routes and measure how this impacts the

[23]The five RIRs are: RIPE NCC for Europe and Middle East, ARIN for North America, APNIC for the Asia-Pacific region, LACNIC for Latin America and the Caribbean, and AfriNIC for Africa.

[24]An IP address block owner can allow more than one AS to originate an announcement to that block by issuing multiple ROAs for the same address block. The process of issuing ROAs is also known as RPKI registration of IP address blocks.

[25]ROA issuance is represented by the top part of figure 5.

[26]NIST has been monitoring the share of IP address blocks advertised in BGP that are covered by ROAs and their status (valid or invalid) since 2013. ROAs adoption has slowly but consistently grown. Currently, ROA coverage by RIR zone is 41% for RIPE, 13% for ARIN, 28% for APNIC, 34% for LACNIC and 8% for AfriNIC (National Institute for Stantards and Technology. *RPKI Deployment Monitor*. `https://rpki-monitor.antd.nist.gov/`).

[27]The use of ROAs to validate routing information is represented by the bottom part of 5.

[28]M. Lepinski, S. Kent, and D. Kong. *A Profile for Route Origin Authorizations (ROAs)*. RFC 6482 (Proposed Standard). RFC. Fremont, CA, USA: RFC Editor, Feb. 2012. URL: `https://www.rfc-editor.org/rfc/rfc6482.txt`; Geoff Huston and George Michaelson. *RFC 6483: Validation of Route Origination Using the Resource Certificate Public Key Infrastructure (PKI) and Route Origin Authorizations (ROAs)*. en. Feb. 2012. URL: `https://tools.ietf.org/html/rfc6483` (visited on 01/27/2020).

[29]AT&T announced in February 2019 that is was filtering RPKI invalid announcements, KPN announced in June 2019 and Telia in February 2020 (*AT&T/as7018 now drops invalid prefixes from peers*. `https://mailman.nanog.org/pipermail/nanog/2019-February/099501.html`; *AS286 Routing Policy*. `https://as286.net/AS286-routing-policy.html`; *Dropping RPKI Invalid Prefixes*. `https://blog.teliacarrier.com/2020/02/05/dropping-rpki-invalid-prefixes/`).

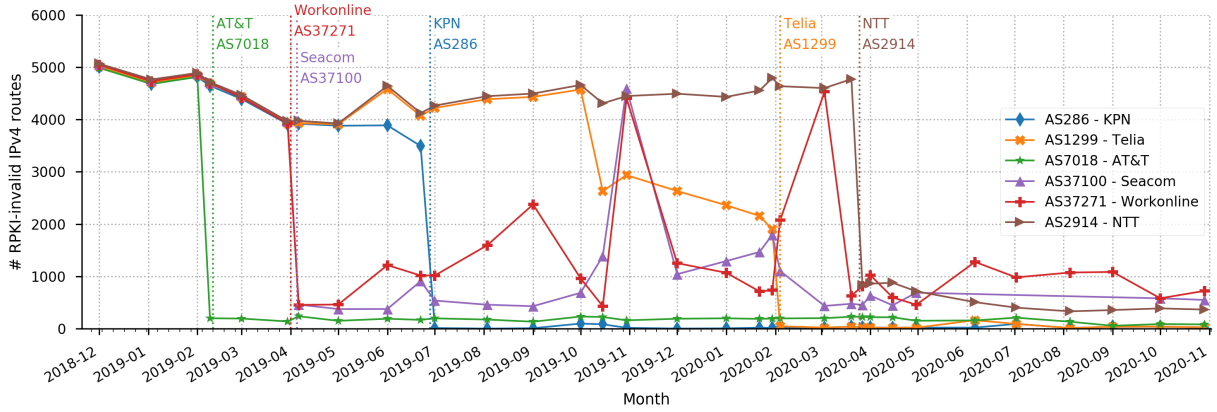


Figure 6: *Invalid IPv4 prefix assertions from ISPs that publicly announced they had started to discard invalid assertions with respect to RPKI ROAs in BGP, from January 2019 until November 2020. The vertical dotted lines correspond to the date of their public announcement.*

spread of illicit announcements. The next paragraphs are a high-level description of the methods and the main insights from that work.[30].

To track ISPs that are filtering invalid routing assertions, we leverage historical BGP data through CAIDA's BGPStream[31] and historical validated ROAs[32]. By coupling those datasets, we are able to measure the number of invalid assertions in ISPs that directly peer with the BGP collector infrastructure, which feeds CAIDA BGPStream. Note that invalid routes can come from misconfigurations in BGP, in ROAs and also hijacks. The volume of invalid routes does not represent the volume of hijacking.

Figure 6 shows the number of invalid route assertions from December 2018 to November 2020 for 6 large ISPs: AT&T (AS7018), KPN (AS286), Seacom (AS37100), Workonline Communications (AS37271), Telia (AS1299) and NTT (AS2914). These ISPs made public announcements when they started discarding routing assertions in BGP that are invalid with respect to ROAs in the RPKI during that period. For reference, in figure 6 the vertical dotted lines are placed on the date of ISPs' public announcement of RPKI filtering. All these ISPs forwarded over 4,000 invalid route announcements to collectors in early 2019, but after they finished deploying the invalid assertion filtering, they forward less than 800 invalid announcements. KPN and Telia also made announcements at the very beginning of their filtering deployment, May 2019 and September 2019 respectively, and the number of invalid assertion slowly declined, matching their deployment status.

In Figure 6 we can also see that Workonline Communication and Seacom, two major African ISPs, ave encountered operational issues when deploying RPKI filtering and have (partially) stopped filtering for some periods of time (intermittent upticks in the figure). Additionally, in our study we never find an ISP that filters *all* invalid route assertions. Besides synchronization delays between ROAs issuance and router filtering rules updates, we find that partial filtering can be related to operational issues, Autonomous Systems relationships (*e.g.* some networks do not filter invalid assertion when received from customers) and Regional Internet Registries policies (*e.g.* legal barriers have limited the availability of ARIN ROAs[33]).

To confidently track ISPs that filter invalid route announcements according to RPKI ROAs, we select ISPs that directly peer with route collectors and that share the equivalent of a full routing table. For these Autonomous Systems, we can conclusively track their filtering behavior over time by tracking the share of invalid assertions represent in their routing table.

[30]For technical details, please refer to (Cecilia Testart et al. "To Filter or Not to Filter: Measuring the Benefits of Registering in the RPKI Today". en. In: *Passive and Active Measurement*. Ed. by Anna Sperotto, Alberto Dainotti, and Burkhard Stiller. Vol. 12048. Series Title: Lecture Notes in Computer Science. Oregon, US: Springer International Publishing, 2020, pp. 71–87. ISBN: 978-3-030-44080-0 978-3-030-44081-7. DOI: 10.1007/978-3-030-44081-7_5. URL: http://link.springer.com/10.1007/978-3-030-44081-7_5).

[31]Orsini et al., "BGPStream".

[32]Historical validated ROAs records are made available by (Taejoong Chung et al. "RPKI is Coming of Age: A Longitudinal Study of RPKI Deployment and Invalid Route Origins". In: *Proceedings of the Internet Measurement Conference*. IMC '19. Amsterdam, Netherlands: Association for Computing Machinery, Oct. 2019, pp. 406–419. ISBN: 978-1-4503-6948-0. URL: https://doi.org/10.1145/3355369.3355596 [visited on 01/25/2020])

[33]Christopher Yoo and David Wishnick. "Lowering Legal Barriers to RPKI Adoption". In: *Faculty Scholarship at Penn Law* (Jan. 2019). URL: https://scholarship.law.upenn.edu/faculty_scholarship/2035.

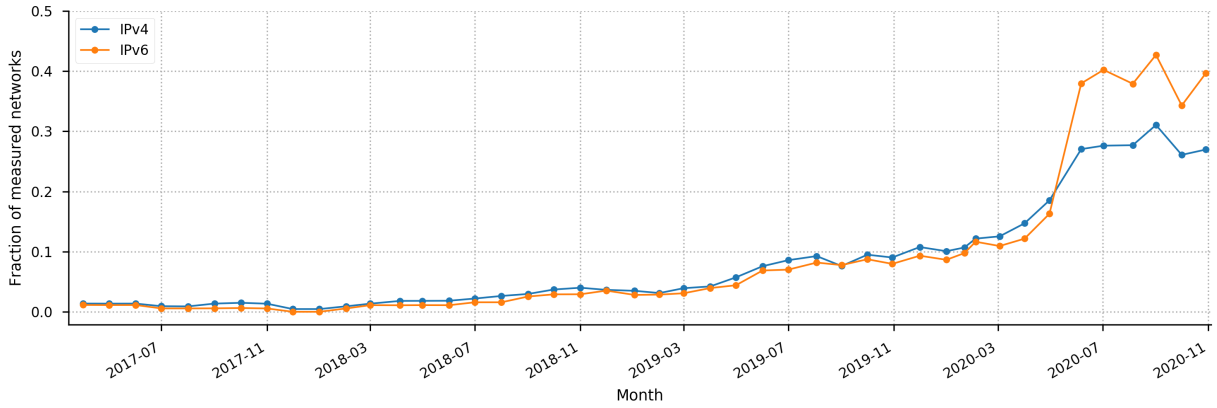


Figure 7: *Fraction of measured Autonomous Systems that filter invalid route assertions with respect to RPKI ROAs over time. In early November 2020, out of 330 measured ASes, 89 are filtering invalids in IPv4.*

The percentage of Autonomous Systems we can measure that are filtering invalid assertions in BGP using RPKI ROAs over time is in Figure 7. At the start of our measurement window in April 2017, less than 2% of networks were dropping invalid assertions in BGP. In 2019, many more networks started dropping invalids, reaching about 10% at the end of 2019. And in 2020 this practice gained even more traction, reaching close to 30% adoption.

## 4.2 The value of issuing ROAs

Encouraged by the uptake in networks dropping invalid routes, we studied the impact of that practice by looking at the visibility of IP address blocks in BGP collector's data. The visibility of IP address blocks in this data is a proxy of the reachability of address blocks in the Internet. If a hijacked address block has low visibility, the harm that can be done by the malicious actor is limited. For the month of September 2019, we process all BGP updates from RIPE and RouteViews collectors through CAIDA's BGPStream to build *timelines* of IP address blocks and the Autonomous System that originated the routing assertion of that block and the visibility over time. We then paired the IP address blocks and origin AS with the RPKI status depending on validated ROAs over time.

Studying the overall visibility of IP address block by their validation status, we find that overall invalid assertions have lower visibility than valid ones and never reach the highest level of visibility[34]. The fact that invalid assertions are never as visible as the valid or unknown state assertion is a direct effect of Autonomous Systems dropping invalid assertions.

In addition, we also look at the visibility of valid and invalid assertions in cases where there is conflict with other routing assertions in the global routing table, *i.e.* either the same IP address block is advertised with another ASN as origin or a part of the address block is advertised with the same or another ASN. In all cases, IP address block from valid route assertions had higher visibility than the ones from invalid route assertions. Issuing ROAs effectively reduced the visibility of illicit announcements by 10-15%, even though only 8% of measured networks were dropping invalid announcements at the time of measurement.

In summary, this study demonstrate that the operational practice of validating information in BGP using RPKI ROAs is gaining traction, and even in partial deployment, this practice effectively limits the spread of invalid information in BGP. The next section discusses accountability challenges of such operational practices.

## 4.3 Accountability of operational practices

The operational practice that is understood today with respect to ROAs is that ISPs should check *all* route assertions and drop those in which the announced route origin is inconsistent with a registered ROA. As we documented in the previous section, major ISPs are now starting to drop route announcements that are made invalid by a ROA, which stops invalid origin announcement from propagating across the Internet. But is this an effective practice? Is it reasonable to expect that most if not all large provider will eventually

[34]The maximum visibility reached by RPKI invalid announcements is about 80% instead of 100%.

adopt—and maintain—this practice? Which ASs should take on the role of checking BGP routes against registered ROAs? Given the lack of central authority and accountability for operational practices, enhanced security practices suffer from the tragedy of the common. An approach to solve this is to create group accountability.

**The Mutually Agreed Norms for Routing Security (MANRS):** In 2014, a group of network operators came together to define a set of operating practices that would enhance the security and resilience of the Internet routing system. This led to a set or practices called the Mutually Agreed Norms for Routing Security (MANRS), supported by the Internet Society, to which 500 ISPs have now subscribed.[35] One of the practices required of ISPs that are MANRS-compliant is that the ISP will check the routing assertions of its customers and verify that the announced AS and prefix are valid.[36]

What the MANRS requirement does is to specify *where* in the system the check of the route origin should be done. It is the responsibility of the immediate provider of a customer. If all ISPs checked the assertions of their customers, it would not be necessary for the large ISPs in the core of the Internet to perform this check, but since MANRS is not universally accepted, the checks we document in Section 4.1 is currently critical to prevent the propagation of invalid routes.

Additionally, although currently MANRS *requires* ISP to check routing assertions of their customers, there is no measurement of *compliance* with this practice, and therefore no clear consequences for networks not caring out the requirements. Since MANRS gives ISPs the freedom to choose on what the verification of customer announcements would be based (not necessarily ROAs), properly monitoring MANRS-compliance is challenging. Using ROAs to verify routing announcements allows to monitor compliance, making it a stronger case to adopt RPKI and ROAs.

# 5 Toward a stronger set of operational norms for routing security

In Section 2.4, we described the dilemma of designing a security mechanism that exactly divides good from bad behavior, and schemes in practice are either going to undershoot or overshoot. The scheme we have described in section 4 to validate address block and the network authorized to originate an announcement for that block in BGP undershoots in a serious way, which creates a easily exploitable path for hijackers to avoid detection. The next sections describe the possible attack and describe and compare two proposed defenses.

## 5.1 Invalid Path hijacks

In an *invalid path hijack*, a malicious Autonomous System (AS) asserts a route including connection between networks (ASs) that do not exist. Figure 8 illustrates an example where the malicious AS (AS666) asserts that it has a connection to AS3. AS666 announces a BGP path of the form P,3,666, *i.e.,* as if AS3 had originated an announcement for address block P and sent it to AS666 and AS666 is then forwarding that announcement to its neighbor (AS6). If there is a valid ROA for the address block P with origin AS AS3, the hijacked announcement will match the valid ROA, so the assertion is likely to be accepted and forwarded by neighbors of AS666. In BGP, any AS can send an assertion including connections between ASs that do not exist. This serious residual vulnerability has caused critics to argue that deploying a scheme based on ROAs is a waste of time. However, this expressed skepticism is based on an assumption that the operational practices associated with the use of ROAs are fixed. The next section presents two proposals to limit the propagation of invalid path hijacks

## 5.2 Proposals to defend against path hijacks

These section discusses and compares two proposals to prevent invalid path hijacks. One, being formulated by the IETF, calls for new mechanism, including in particular an additional global database. The other, which we are exploring, calls for enhanced operational practices.[37]

[35] *Mutually Agreed Norms for Routing Security (MANRS).*

[36] MANRS does not require the this verification be based on ROAs. It could also be based on an older system called IRRs, or a private sharing of data between the ISP and the customer. But this discussion will assume that the validation is based on the use of ROAs.

[37] There is another proposal, BGPsec, standardized by the IETF in 2017 to defend against path hijacks. However, mainly because of the scheme computational cost and operation complexity, BGPsec is not currently deployed (Matthew Lepinski and

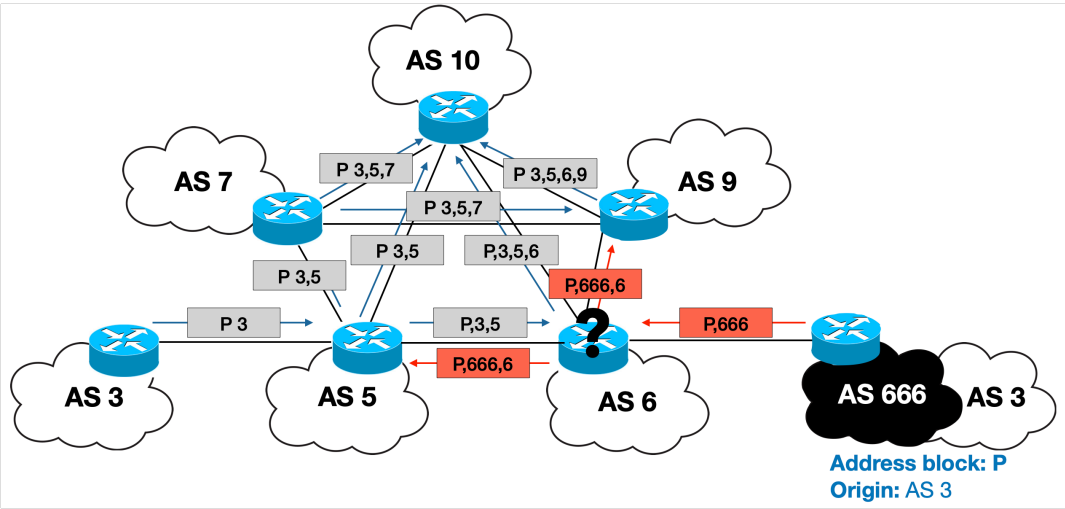


Figure 8: *AS 666 attempts to hijack address block P by falsely asserting that it has a phantom connection to AS3. It has no connection to AS3, but the announcement falsely asserts that it does. The resulting announcement has the form p,3,666. If AS6 checks the validity of the origin in this announcement, it will find that P,3 is confirmed by a valid ROA, and will accept the announcement.*

### 5.2.1 Autonomous System Provider Authorization (ASPA)

The new IETF proposal to defend against prefix hijacks is based on the creation of a new global database that would allow any AS to register a list of the ASs that are valid providers to the AS. This scheme is called Autonomous System Provider Authorization or ASPA.[38]

Using figure 8 as a reference of how this proposal would work, AS3 would register in an ASPA record that the only AS that is allowed to announce a path to it is AS5. Then when AS6 gets the announcement P,3,666, it could determine that the announcement is invalid because the connection between AS3 and AS666 is invalid with respect to the information in AS3 ASPA record.

The drawback to this scheme is that it requires the creation of a new database (or the addition of new sorts of records to the current RPKI databases used for ROAs), and that all the relevant ASs register ASPAs. It also requires that *all* ISPs take on the job of checking the routing assertions against the registered information, but does not define exactly which ISPs should do this. Just like the practice of validating BGP announcements with ROAs, ASPA validation suffers from the tragedy of the commons and the benefits it can bring to network registering ASPA records resides in the collective action of networks.

### 5.2.2 Establishing a routing *zone of trust*

We have been exploring an alternative approach that deals with invalid path hijacks by enhancing the practices that are required by the MANRS agreement described in section 4.3. The idea is to establish a routing *zone of trust* through the enhanced operational practices of MANRS participants, and have ISPs give priority to announcements coming from ISPs within that zone. We call our proposal *recursive MANRS*.

Our proposal requires that every MANRS-compliant ISP know which of its customers is also MANRS-compliant to establish the *zone of trust*. MANRS membership will not change rapidly, so it should not be a burden to track it.[39] If the customer of a MANRS-compliant provider is also MANRS-compliant, then the provider ISP can assume that the customer ISP has checked its own customers, *i.e.* it is within the *trust zone* established recursively between members of MANRS. Then, MANRS-compliant ISPs accept all announcements from MANRS-compliant customers. If the customer does not participate in MANRS, treat

Kotikalapudi Sriram, eds. *RFC 8205: BGPsec Protocol Specification*. en. Sept. 2017. URL: `https://tools.ietf.org/html/rfc8205`).

[38]The idea is that using ASPA, networks can verify the validity of implied AS connections in the paths of the announcements it receives in BGP and drop announcements that include a non-authorized connection (Alexander Aximov et al. *Verification of AS_PATH Using the Resource Certificate Public Key Infrastructure and Autonomous System Provider Authorization*. Nov. 2020).

[39]Any new member of MANRS could be required to contact all its providers about its new membership.

the BGP assertion as "suspicious" as it comes from outside the *trust zone*.[40]

Thus, in the case of Figure 8, if AS666 has committed to comply with the practices of recursive MANRS, AS6 would accept the announcement, since it is coming from the *trust zone*. Since routing assertions are public, BGP monitors would see this route, and if it were found to be malicious, AS666 would be ejected from the MANRS membership and not considered to be part of the trust zone. If AS666 is *not* MANRS-compliant, then AS6 should view the announcement as "suspicious", since it is not coming from the *trust zone*. To over-simplify, if the ISP receiving the assertion (for example, AS6) has another route to the origin which is within the trust zone, *i.e. not* "suspicious", it discards the suspicious one.[41]

This zone of trust proposal does not require a global data base. It depends on an AS knowing about its customers, but that knowledge can be private to those parties. However, the validation can only be done at the point where the announcement encounters the first MANRS-compliant AS . With ASPA, any AS can, if it chooses, inspect an announcement to detect an invalid path.

### 5.3 Comparing the two approaches

There are clear differences between the security guarantees of the two proposals described in the previous section. For instance, in the case of announcements coming from customers of customers outside the routing trust zone of recursive-MANRS, legitimate paths are indistinguishable from illegitimate paths. However, with ASPA, the illegitimate path could be identified *if* involved ASs had properly registered their providers in ASPA records. Nonetheless, there is also other type of malicious activity that would not be solved with ASPA. Conversely, in the case of an announcement involving unused or unassigned space for which there is no ROA, MANRS-compliance prevents any MANRS ISP or its customers from originating such announcements, whereas RPKI ROAs and ASPAs would not prevent those announcements from spreading.

Although we can identify differences in how effective the two proposal would be at stopping the spread of hijacks in different cases, there currently is no evidence to determine which sorts of hijack are more prevalent or more harmful. It is thus unclear what approach, or combination of approaches, would be most effective in creating a save zone free of harmful hijacks.

Still, the difference between the two proposals is less whether one blocks more types of potential hijacks than the other, but the complexity and incentives that would help or hinder the deployment of the two. In today's Internet, issues of real-world deployment of a scheme are probably more important than the technical elegance of a scheme. We argue that a global database may be technically elegant, but faces challenges in deployment and uptake similar to the ones discussed in section 2.4. In particular, the first networks implementing ASPA will experience first mover disadvantage and there is no framework to drive consensus on how to practically operationalize this solution.

Another big difference between the two proposal is the fact that ASPA centralizes the records needed to validate path in RIRs' infrastructure. This centralization leads to additional complexities that would need to be resolved:

- The ASPA database (which would physically be distributed) would probably be hosted by the five RIRs. How would their operating expenses be covered? What is their incentive to take this role on?
- The database, like the ROA database, would be an attractive target for attack, so it would have to be maintained at the highest level of security, which would further add to the operating costs.
- ASPA, like ROAs, create a potential point of control that could be exercised by a sovereign state to block certain connections in the Internet. To what extent are these systems a new point of vulnerability for state control?[42]

[40] Note that for the MANRS ISP to be compliant with MANRS requirements, it has to check all announcements originated by its customers. Therefore, announcements tagged as *suspicious* would never include the customer address space; those are either valid or filtered out by MANRS ISPs.

[41] There are many possible scenarios of how this would play out when legitimate and illegitimate announcements for the same address block come from customers of non MANRS-compliant ISPs. In some cases, legitimate announcements from non-MANRS ISPs would be unidentifiable from illegitimate ones. Although this is not ideal, this is the current state of routing. As more networks decide to join MANRS and the size of the routing zone of trust increases, the likelihood and impact of those cases would decrease.

[42] MANRS does not require that validation be based on ROAs—any valid arrangement between customer and provider, including a private contract, is acceptable. However, given *any* MANRS ISP could validate announcements with ROAs and not just the provider, issuing ROAs is beneficial even with MANRS. And, given that MANRS-compliant ISPs are required to verify *all* announcements from their customers, it provides security for address block for which a proper ROA cannot be issued such as unassigned and reserved address space and legacy address blocks distributed before RIR where consolidated.

Looking at how long it has taken for RPKI registration of ROAs and route validation in BGP using those ROAs to gain traction, it is reasonable to believe that ASPA will take a long time to reach a point where its benefits are significant for networks. In the meanwhile, a group of ISPs pushing for better routing security such as MANRS could implement the recursive MANRS proposal.

The proposal based on the MANRS agreement to establish a routing zone of trust can leverage MANRS as a framework to reach consensus on operational practice details (as members have already done for current requirements). This proposal also clarifies responsibilities—MANRS-compliant ISPs need to verify their customers announcements and give preference to announcements from other MANRS members—and with proper monitoring of compliance, it makes networks *accountable* for their practices.

# 6 Key points in the path to better security

In this paper, we revisit the long-studied problem of routing insecurity, showing real-world evidence of the state of routing security from recent technical works, with the goal of discussing possible options for better security. We start by describing the main vulnerability of BGP, which enables *BGP hijacking*, the deflection of traffic intended for one destination so that a malicious region of the Internet receives it instead. We review the differences and main barriers encountered by previous proposals to secure BGP. In addition, we discuss the challenge of designing a mechanism that perfectly draws the line between unwanted behavior and usual operation, especially without knowing much about the level of abuse in BGP.

To better understand malicious activity in BGP, we study BGP hijacking behavior and find the existence of *serial hijackers*, networks that persistently perform hijacks in BGP with almost no consequences. Switching to the defense side, we report the level of adoption and the resulting benefits of the IETF scheme to validate routing information using Route Origin Assertions (ROAs) published in the Resource Public Key Infrastructure (RPKI) framework. Even in partial adoption, we can measure that this operational practice reduces the spread of invalid information in BGP. However, in the RPKI framework, it is not clear which ISP should be responsible for undertaking which part of the solution. We describe the MANRS initiative, and point out that one of its benefits is that it clarifies where responsibility lies to implement the scheme. Only with clear allocation of responsibility is it realistic to put in place a scheme to monitor and track members for compliance.[43]

Finally, building on our study of past failed attempts to improve BGP security by the deployment of new protocols, we propose a new approach to tackle security issues that are not covered by RPKI and ROAs. This scheme proposes the creation of a routing *zone of trust* based on enhanced *operational practices* that will bring benefit to the members of the zone as it is incrementally deployed. We compare this approach to another proposal formulated by the IETF.

We offer two high-level conclusions. First, in designing an approach (or comparing proposed approaches) it is perhaps most important to understand the complexities and incentives for deployment. Second, any scheme that is picked for deployment will take time to have an effect, so it is important to have a realistic view of how quickly change can be accomplished. Quick fixes will usually be incomplete and inadequate fixes. But we cannot make progress until we start.

# References


*AS286 Routing Policy*. `https://as286.net/AS286-routing-policy.html`.

*AT&T/as7018 now drops invalid prefixes from peers*. `https://mailman.nanog.org/pipermail/nanog/2019-February/099501.html`.

Aximov, Alexander et al. *Verification of AS_PATH Using the Resource Certificate Public Key Infrastructure and Autonomous System Provider Authorization*. Nov. 2020.

Butler, Kevin et al. "A Survey of BGP Security Issues and Solutions". In: *Proceedings of the IEEE* 98.1 (Jan. 2010), pp. 100–122. ISSN: 0018-9219, 1558-2256. DOI: `10.1109/JPROC.2009.2034031`. URL: `http://ieeexplore.ieee.org/document/5357585/` (visited on 08/04/2017).


[43] We note that transparency does not imply accountability per se. The RPKI scheme is transparent but it leave it to ISPs to decide if/when/what they want to validate with the records. There is no consequences for not validating announcements in BGP.

Chung, Taejoong et al. “RPKI is Coming of Age: A Longitudinal Study of RPKI Deployment and Invalid Route Origins”. In: *Proceedings of the Internet Measurement Conference*. IMC ’19. Amsterdam, Netherlands: Association for Computing Machinery, Oct. 2019, pp. 406–419. ISBN: 978-1-4503-6948-0. URL: https://doi.org/10.1145/3355369.3355596 (visited on 01/25/2020).

*Cisco BGPStream CrossworkCloud*. https://bgpstream.com/.

*Cymru BGP Bogon Refence*. https://team-cymru.com/community-services/bogon-reference/.

*Dropping RPKI Invalid Prefixes*. https://blog.teliacarrier.com/2020/02/05/dropping-rpki-invalid-prefixes/.

Goodin, Dan. *Suspicious event hijacks Amazon traffic for 2 hours, steals cryptocurrency*. en-us. Apr. 2018. URL: https://arstechnica.com/information-technology/2018/04/suspicious-event-hijacks-amazon-traffic-for-2-hours-steals-cryptocurrency/ (visited on 11/01/2019).

Hu, Xin and Z. Morley Mao. “Accurate Real-time Identification of IP Prefix Hijacking”. In: *2007 IEEE Symposium on Security and Privacy (SP ’07)*. May 2007, pp. 3–17. DOI: 10.1109/SP.2007.7.

Huston, Geoff and George Michaelson. *RFC 6483: Validation of Route Origination Using the Resource Certificate Public Key Infrastructure (PKI) and Route Origin Authorizations (ROAs)*. en. Feb. 2012. URL: https://tools.ietf.org/html/rfc6483 (visited on 01/27/2020).

Krebs on Security. *Notorious ‘Hijack Factory’ Shunned from Web*. https://krebsonsecurity.com/tag/bitcanal/.

Lad, Mohit et al. “PHAS: A Prefix Hijack Alert System”. In: *Proceedings of the 15th Conference on USENIX Security Symposium - Volume 15*. USENIX-SS’06. Berkeley, CA, USA: USENIX Association, 2006. URL: http://dl.acm.org/citation.cfm?id=1267336.1267347 (visited on 10/02/2018).

Lad, Mohit et al. “Understanding Resiliency of Internet Topology against Prefix Hijack Attacks”. In: *37th Annual IEEE/IFIP International Conference on Dependable Systems and Networks (DSN’07)*. June 2007, pp. 368–377. DOI: 10.1109/DSN.2007.95.

Lepinski, Matthew and Kotikalapudi Sriram, eds. *RFC 8205: BGPsec Protocol Specification*. en. Sept. 2017. URL: https://tools.ietf.org/html/rfc8205.

Madory, Doug. *Shutting down the BGP Hijack Factory — Dyn Blog*. en-US. https://dyn.com/blog/shutting-down-the-bgp-hijack-factory/. July 2018. (Visited on 05/12/2019).

— *Sprint, Windstream: Latest ISPs to hijack foreign networks — Dyn Blog*. en-US. https://dyn.com/blog/latest-isps-to-hijack/. Sept. 2014. (Visited on 05/12/2019).

— *The Vast World of Fraudulent Routing — Dyn Blog*. en-US. https://dyn.com/blog/vast-world-of-fraudulent-routing/. Jan. 2015. (Visited on 05/12/2019).

Mitchell, Jon. *RFC 6996: Autonomous System (AS) Reservation for Private Use*. en. July 2013. URL: https://tools.ietf.org/html/rfc6996 (visited on 05/11/2019).

*Mutually Agreed Norms for Routing Security (MANRS)*. https://www.manrs.org/.

Nicholes, Martin O. and Biswanath Mukherjee. “A survey of security techniques for the border gateway protocol (BGP)”. In: *IEEE Communications Surveys Tutorials* 11.1 (2009), pp. 52–65. ISSN: 1553-877X. DOI: 10.1109/SURV.2009.090105.

*The Hunt for 3ve: Taking down a major ad fraud operation through industry collaboration*. Tech. rep. Nov. 2018. URL: https://services.google.com/fh/files/blogs/3ve_google_whiteops_whitepaper_final_nov_2018.pdf?__hstc=&__hssc=&hsCtaTracking=c7b87c5c-1676-4d53-99fb-927a07720b17%7C9d63bf77-0926-4d08-b5ec-46b1a06846bc (visited on 10/29/2019).

Orsini, Chiara et al. “BGPStream: A Software Framework for Live and Historical BGP Data Analysis”. In: *Proceedings of the 2016 Internet Measurement Conference*. IMC ’16. Santa Monica, California, USA: Association for Computing Machinery, Nov. 2016, pp. 429–444. ISBN: 978-1-4503-4526-2. URL: https://doi.org/10.1145/2987443.2987482 (visited on 01/25/2020).

Project, The Spamhaus. *DROP - Don’t Route or Peer lists - The Spamhaus Project*. https://www.spamhaus.org/drop/.

Qiu, Jian and Lixin Gao. *Hi-BGP: A Lightweight Hijack-proof Inter-domain Routing Protocol*. Tech. rep. 2006, p. 12.

Qiu, Tongqing et al. “Locating Prefix Hijackers Using LOCK”. In: *Proceedings of the 18th Conference on USENIX Security Symposium*. SSYM’09. Berkeley, CA, USA: USENIX Association, 2009, pp. 135–150. URL: http://dl.acm.org/citation.cfm?id=1855768.1855777 (visited on 10/02/2018).

Ramachandran, Anirudh and Nick Feamster. “Understanding the network-level behavior of spammers”. In: *ACM SIGCOMM Computer Communication Review*. Vol. 36. ACM, 2006, pp. 291–302.

Lepinski, M., S. Kent, and D. Kong. *A Profile for Route Origin Authorizations (ROAs)*. RFC 6482 (Proposed Standard). RFC. Fremont, CA, USA: RFC Editor, Feb. 2012. URL: https://www.rfc-editor.org/rfc/rfc6482.txt.

Schlamp, Johann et al. "HEAP: Reliable Assessment of BGP Hijacking Attacks". In: *IEEE Journal on Selected Areas in Communications* 34.6 (June 2016), pp. 1849–1861. ISSN: 0733-8716. DOI: 10.1109/JSAC.2016.2558978.

Sermpezis, Pavlos et al. "A Survey among Network Operators on BGP Prefix Hijacking". In: *ACM SIGCOMM Computer Communication Review* 48.1 (Apr. 2018), pp. 64–69. ISSN: 01464833. DOI: 10.1145/3211852.3211862. URL: http://dl.acm.org/citation.cfm?doid=3211852.3211862 (visited on 09/28/2018).

Sermpezis, Pavlos et al. "ARTEMIS: Neutralizing BGP Hijacking within a Minute". In: *arXiv:1801.01085 [cs]* (Jan. 2018). URL: http://arxiv.org/abs/1801.01085 (visited on 10/01/2018).

Siddiqui, Muhammad S. et al. "A survey on the recent efforts of the Internet Standardization Body for securing inter-domain routing". en. In: *Computer Networks* 80 (Apr. 2015), pp. 1–26. ISSN: 13891286. DOI: 10.1016/j.comnet.2015.01.017. URL: http://linkinghub.elsevier.com/retrieve/pii/S1389128615000286 (visited on 02/15/2018).

Stantards and Technology, National Institute for. *RPKI Deployment Monitor*. https://rpki-monitor.antd.nist.gov/.

Testart, Cecilia. "Reviewing a Historical Internet Vulnerability: Why Isn't BGP More Secure and What Can We Do About it?" en. In: Washinton, DC: Social Science Research Network, Aug. 2018. URL: https://papers.ssrn.com/abstract=3141666 (visited on 09/29/2019).

Testart, Cecilia et al. "Profiling BGP Serial Hijackers: Capturing Persistent Misbehavior in the Global Routing Table". en. In: *Proceedings of the Internet Measurement Conference on - IMC '19*. Amsterdam, Netherlands: ACM Press, 2019, pp. 420–434. ISBN: 978-1-4503-6948-0. URL: https://dl.acm.org/doi/10.1145/3355369.3355581 (visited on 10/27/2019).

— "To Filter or Not to Filter: Measuring the Benefits of Registering in the RPKI Today". en. In: *Passive and Active Measurement*. Ed. by Anna Sperotto, Alberto Dainotti, and Burkhard Stiller. Vol. 12048. Series Title: Lecture Notes in Computer Science. Oregon, US: Springer International Publishing, 2020, pp. 71–87. ISBN: 978-3-030-44080-0 978-3-030-44081-7. DOI: 10.1007/978-3-030-44081-7_5. URL: http://link.springer.com/10.1007/978-3-030-44081-7_5.

UCEPROTECT. *Blacklist Policy LEVEL 2*. http://www.uceprotect.net/en/index.php?m=3&s=4.

Vervier, Pierre-Antoine, Olivier Thonnard, and Marc Dacier. "Mind Your Blocks: On the Stealthiness of Malicious BGP Hijacks". en. In: *Proceedings 2015 Network and Distributed System Security Symposium*. San Diego, CA: Internet Society, 2015. ISBN: 978-1-891562-38-9. DOI: 10.14722/ndss.2015.23035. URL: https://www.ndss-symposium.org/ndss2015/ndss-2015-programme/mind-your-blocks-stealthiness-malicious-bgp-hijacks/ (visited on 10/02/2018).

Xingang Shi et al. "Detecting Prefix Hijackings in the Internet with Argus". In: *ACM IMC*. 2012.

Yoo, Christopher and David Wishnick. "Lowering Legal Barriers to RPKI Adoption". In: *Faculty Scholarship at Penn Law* (Jan. 2019). URL: https://scholarship.law.upenn.edu/faculty_scholarship/2035.

Zhang, Z. et al. "iSPY: Detecting IP Prefix Hijacking on My Own". In: *IEEE/ACM Transactions on Networking* 18.6 (Dec. 2010), pp. 1815–1828. DOI: 10.1109/TNET.2010.2066284.