

Firewalling Scenic Routes: Preventing Data Exfiltration via Political and Geographic Routing Policies

Kevin Benton
School of Informatics and Computing
Indiana University
Bloomington, Indiana, USA
Email: KTBenton@Indiana.edu

Dr. L. Jean Camp
School of Informatics and Computing
Indiana University
Bloomington, Indiana, USA
Email: LJCamp@Indiana.edu

ABSTRACT

In this paper we describe a system that allows the real time creation of firewall rules in response to geographic and political changes in the control-plane. This allows an organization to mitigate data exfiltration threats by analyzing Border Gateway Protocol (BGP) updates and blocking packets from being routed through problematic jurisdictions. By inspecting the autonomous system paths and referencing external data sources about the autonomous systems, a BGP participant can infer the countries that traffic to a particular destination address will traverse. Based on this information, an organization can then define constraints on its egress traffic to prevent sensitive data from being sent via an untrusted region. In light of the many route leaks and BGP hijacks that occur today, this offers a new option to organizations willing to accept reduced availability over the risk to confidentiality. Similar to firewalls that allow organizations to block traffic originating from specific countries, our approach allows blocking outbound traffic from transiting specific jurisdictions.

To illustrate the efficacy of this approach, we provide an analysis of paths to various financial services IP addresses over the course of a month from a single BGP vantage point that quantifies the frequency of path alterations resulting in the traversal of new countries. We conclude with an argument for the utility of country-based egress policies that do not require the cooperation of upstream providers.

1. INTRODUCTION

The Border Gateway Protocol (BGP) is responsible for the distribution of routes between autonomous systems on the Internet. BGPv4, the latest version of the protocol, was released in 1995 [25] the newest revisions were released in 2006 [28]. Due to the trusted nature of the Internet at the time of the protocol's creation, BGP participants are assumed to behave well and only advertise routes to networks to which they had connectivity. Subsequently, the protocol itself contains no protection from participants advertising false routes.

Permission to make digital or hard copies of all or part of this work for personal or classroom use is granted without fee provided that copies are not made or distributed for profit or commercial advantage and that copies bear this notice and the full citation on the first page. Copyrights for components of this work owned by others than ACM must be honored. Abstracting with credit is permitted. To copy otherwise, or republish, to post on servers or to redistribute to lists, requires prior specific permission and/or a fee. Request permissions from permissions@acm.org.

SafeConfig'16, October 24, 2016, Vienna, Austria

© 2016 ACM. ISBN 978-1-4503-4566-8/16/10...\$15.00

DOI: <http://dx.doi.org/10.1145/2994475.2994477>

The trust architecture of BGP means that simple misconfigurations can take down large chunks of the Internet. In a particularly large-scale example, in 2015 the ISP *Telekom Malaysia* advertised short paths to 179,000 prefixes to its peers which redistributed them globally. The result was a major portion of the Internet traffic in Asia saturating Telekom Malaysia's network [2] as it followed the shortest path. While these route leaks are quickly identified as errors by network operators, the standard BGP protocol does nothing to stop it.

In addition to these route leaks, there are focused outages caused by ISPs advertising prefixes they do not own (a.k.a prefix hijacking). A well-known example of this is when Pakistan Telecom attempted to block YouTube in the country by setting up a black-hole route¹ for YouTube's prefixes. While this is a common method of blocking traffic, they made the mistake of redistributing the route to their BGP peers outside of the country, which resulted in a global YouTube outage. [19]

In both cases, the error was immediately noticeable to operators so they could quickly take steps to intervene (e.g. setup route filters, stop peering with bad neighbors, etc). However, when prefixes are hijacked in a targeted manner, the result can go unnoticed for weeks and/or until well after the attacker stops the attack. For example, an attacker used an ISP to hijack prefixes containing control-nodes of a group of Bitcoin² miners just long to enough command them to connect to a different control-node. Under the control of the new node, the reward for all of the miners' work was directed to the attacker, which was about \$83,000 worth of Bitcoin. Due to the short, targeted nature of the attacks, none of the operators of the hijacked prefixes (mainly hosting providers) even noticed the event. [3]

The final type of attack is more subtle than short hijacks. Instead of advertising a false route to the entire Internet, the attacker only advertises a bad prefix to one of its peers while leaving the others unaffected. This forces traffic from some ISPs to travel to the attacker. However, the attacker can then forward it via the unaffected peers to its destination after recording or making modifications to the traffic. This method was used to induce traffic between two ISPs in Denver to first travel to the UK and Iceland before returning to Denver. [16, 13]

In this paper we analyze BGP topology changes that cause traffic to a given destination to traverse countries through

¹A black-hole route is a route that just drops packets which match it.

²Bitcoin is a digital currency

which it was not previously routed. We show that it is relatively simple to detect these changes as a BGP participant and that organizations can implement exfiltration policies based on this detection to prevent sensitive traffic from traversing undesirable administrative domains. Our results are agnostic as to the definition of undesirable, which is left to the network operator. The analysis we present here is a worse case. Even with the assumption that all changes in jurisdiction are undesirable, we show there would be no disruption in the majority of cases, and minimal disruption in few cases under normal operation.

In the next section we discuss related work on detecting and preventing bad routes. In section 3 we discuss how an organization can adopt policies that restrict which autonomous systems traffic to a particular destination should traverse. In section 4, we provide an analysis based on Route Views data which quantifies how frequently an observer may see changes in the countries in an AS path to destinations they deem sensitive. In section 5 we demonstrate a proof of concept implementation and then we conclude in section 8.

2. RELATED WORK

2.1 Detecting Bad Routes

Defending against hijacks and route leaks has long been a concern. One approach is to add explicit trust information using a PKI. A PKI, once adopted across ASes, can provide attestation to source, path, and right to advertise a particular IP. Resource Public Key Infrastructure (RPKI) is specified in RFC 6480 [17], which defines a public key hierarchy that can be used with the existing BGP protocol to sign BGP advertisements. Once RPKI is fully adopted, it will prevent prefix hijacking but it will not address path spoofing attacks since it only binds a prefix to a specific origin. There is ongoing work in RFC 7353 [4] to extend RPKI to sign each hop in the AS PATH to prevent path modification attacks as well; however, it is not finalized let alone adopted on current Internet infrastructure.

Given the current state of BGP, and lack of PKI, defenses have leveraged the control plane, the data plane, or both.

Zhang et al. created a framework to detect bogus routes by using traceroutes and observing both the paths and the timestamps. This work showed that bogus routes resulted in significant variations, even if these were not otherwise visible to the target path [30]. These measurements of data-plane round-trip-time (RTT) inherently include some geographical information, based simply on the sheer distance traveled that is a part of many hijacks. One strength of this approach is that individuals can adopt it at the edge of the network without coordinating with any peers. The weakness of this approach is that it requires near real-time data from probes as well as correct ICMP responses from all the nodes on the path. Thus malicious ASes have a path to attack the system by manipulating traceroute probes.

Solutions and detections from more than one network location shows promise for addressing the challenge of the single malicious AS in the path. Hiran et al. proposed a collaborative mechanism, with a crowd-sourced approach to RTT measurements.[11] Qui et al. proposed a system with distributed monitors, without explicit crowd-sourcing as one entity could control all these monitors [24]. The goal of their system is LOCating the prefix hijackKer, condensed into the name LOCK. The potential hijack candidates are identified

by clustering and ranking that takes into account the most popular routing policies.

However, the potential for distributed RTT in the data plane is complicated by the use of BGP anycast. The use of anycast is widespread, and results in correct RTT variance from different parts of the network. Although this does not require the level of coordination necessitated by the protocol-level PKI-based solutions, it would require large-scale coordinated adoption to be effective. Identifying and preventing malicious participants from engaging as defenders and polluting data is also a potential challenge.

Rather than focusing on a range of sources, Hu et al. proposed a system based on route update collection and data-plane fingerprinting [12] for a range of destinations. These authors used multiple BGP feeds to fingerprint paths based on analyses of to probes sent to endpoints in each prefix. This approach allows for the detection of anomalies using geographic approaches, as well as previous topological patterns (which they refer to as relationship and edge constraints). Again this solution has the advantage that potential targets need not cooperate with others for adoptions. The negative is the requirements for traffic and real time processes, requirements which increase as the number of potential legitimate traffic destinations increases.

Similar to reputation approaches and with a focus on topology, Qiu et al. examines changes in BGP over time[23]. The system examines previously unseen routes, with the assumption that such routes are likely to be hijacks. With the published analysis the result would be roughly .02% of updates (they estimate this as 20 per day) being false positives. Even with updated heuristics, a significant number of legitimate updates were being incorrectly classified.

As opposed to the use of cryptographic assertions, Chang et al. proposed a reputation system. With this system, called AS-CRED, route leaks and rapid route announcement withdrawals are tracked over time to provide an indicator of the trustworthiness of a particular AS. Their analysis shows that such behaviors are not uniformly distributed but rather that ASes vary enough that reputation is a useful indicator [5]. The power of reputation-based analysis can be shown by the enduring power of Spamhaus, whose listing of malicious IP addresses over time can be used to create AS reputation. Macroeconomic analyses of the distribution of malicious IP addresses has shown that both jurisdiction and AS are significant in likelihood of a particular IP address being malicious [8, 26]. While our work focuses on geographic anomalies, AS-CRED and other reputation mechanisms are potentially complementary sources of information that could be used to identify, and thus avoid, undesirable ASes.

The existence of hijacks is not currently disputed. However, the number of hijacks varies by an order of magnitude depending on the method of detection. For example, Argus indicates that there are tens of thousands of anomalies which may be hijacks every year, and this number is growing. With a more detailed investigation in search of ground truth indicated 2000 in 18 months [27]. McArthur et al. showed that, even with the proposed detection systems, targeted hijacks remain undetected [20]. Thus the ability to make higher level policies, such as identifying ASes or regions to avoid, can be useful.

2.2 Preventing Propagation of Bad Routes

Once suspicious routes have been identified, there has also

been significant research focused on determining what actions can be taken to prevent these from harming the network.

Zhang et al. showed that just 20 well-connected ISPs can reduce a hijack’s impact by 50% by ignoring the bad route [29]. However, getting these strongly-connected Tier 1 ISPs to adopt these schemes is difficult due to poor economic incentives. The customers of Tier 1 ISPs are mainly other ISPs, who are not normally the target of route hijacks.

Karlin et al. proposed a system that delays the acceptance of origin-altering routes if the original route is still being announced [14]. One of the issues with this historical approach is that it can delay a time-sensitive legitimate update (e.g. adding a new CDN to handle a DDoS).

There are also proposals that require widespread adoption of software changes to become effective. Gersch et al. proposed a system that checks for prefix information stored in DNS records to validate it [9]. Qi et al. proposed that all routers perform attestations on a neighbor’s routing software before accepting their routes [18]. However, the invasive nature of these changes makes their adoption by router vendors unlikely.

Other work focuses on particular domains or solutions. In Tan et al. the defenders use the techniques identified in [30] in the context of the Tor network [10]. The proposed defense for the Tor network is to create a set of guard nodes who maintain a list of trustworthy routes, primarily depending on the dataplane (i.e. traceroute). It combines this with control plane information to defend against shortest path and longest-prefix attacks.

Anderson et. al. argued that detection of bad routes and protection from them in multitenancy environments was a core value proposition of SDN, although the work is agnostic about detection. [1]

All of this work attempts to fundamentally prevent route hijacks and route leaks. The key difference between our work and these is that we assume hijacks/leaks will happen and examine what a BGP end-user (i.e. companies operating networks receiving BGP updates from upstream ISP(s)) can do to react to these events. Even if prefix hijacking and path-length attacks were to be completely eliminated, an organization may want to cease sending data in reaction to a topology change that would send sensitive data to an undesired geographical region.

3. EXFILTRATION POLICIES BASED ON BGP PATHS

One of the primary benefits offered by BGP and other routing protocols is a high resiliency to link failures in networks. As long as routes to a prefix are being advertised by any participant in the network, everyone will be able to send traffic to that prefix. This property has made the Internet very reliable in spite of accidents cutting fiber links, natural disasters, and the lack of a centralized entity to dictate how entities should peer.

However, this global connectivity at all costs can conflict with the security requirements of an organization. Having servers exposed to the entire Internet can expose an entity to attacks from locations it does not intend to serve. To address this, some organizations will leverage IP geolocation, which attempts to identify the geographic location of an IP address. By correlating incoming traffic with a source

location, an organization can set filtering policies to drop traffic from regions it does not want to expose services to. For example, a credit card transaction processor that only serves businesses in Switzerland may decide to block any traffic that does not originate from a Swiss IP address.

There is ongoing research studying the efficacy of different IP Geolocation methods [15, 22, 6]. But even with the known limitations there is enough demand for it that many firewalls, including open source firewalls like IPTables³ and PFSense⁴, include the ability to block traffic based on country of origin.

Given the demand for filtering traffic from specific regions, we believe there are many use cases for preventing traffic from transiting specific regions. For example, a company that maintains a site-to-site VPN connection with another office in the same country to exchange sensitive information may want to adopt a policy that blocks the VPN connection if the traffic were to transit another country. Choosing such a policy would be exchanging availability for a higher degree of confidentiality, which may be acceptable for many applications without always-on requirements (e.g. periodic off-site backups).

To allow these types of policies, we propose a system that inspects the BGP update stream that an organization receives from its upstream provider. For any given IP address that an organization is going to send traffic to, it can perform a longest-prefix-match on the BGP routing information base to get the matching advertisements and extract the autonomous system (AS) paths from them. Then by cross-referencing each AS with its registration info, the system can determine if traffic to that IP will violate a regional constraint. If a constraint is violated, the system will generate access control lists to block traffic to that IP. Figure 1 shows this process.




Figure 1: Exfiltration Prevention Architecture

An important assumption we are making is that the upstream ISP is trustworthy. If the upstream ISP or an ISP in the normal path within a country wanted to, it could lie about the AS path and route the traffic via another country or perform malicious activities itself. While this is a valid threat, it is not the one we are addressing in this work. Our threat model is BGP hijacking by a BGP participant in a region the organization does not trust.

While organizations should always encrypt data in transit over networks, a defense-in-depth approach to security dictates that organizations should have other protection mechanisms should there be a flaw in the encryption protocol/software. Adding the ability to restrict communications with sensitive

³<http://iptables-addons.sourceforge.net/geoip.php>

⁴https://doc.pfsense.org/index.php/Country_Block

IP addresses to transiting specific regions reduces the points at which traffic recording or manipulation can take place.

The primary concern with a scheme like this is false positives. We do not want to frequently block traffic due to the normal churn of the global BGP topology. In the next section we analyze the churn of transiting countries over the course of a month from a single BGP vantage point.

4. DATA ANALYSIS

To examine how frequently routes to specific IP addresses change, we chose the top 50 banking/financial websites from Alexa⁵ as a rough approximation for sensitive addresses to which an organization may want to apply transit restrictions. We then examined the BGP state from a single ISP based in San Francisco over the course of April, 2016 and identified each time the country associations would change in the AS path to each banking site.

We chose jurisdiction and banking for the example implementation since avoiding and recovering from fraud is made more difficult when such crime crosses jurisdictions. It is also the case that different types of fraud are more or less prevalent in different jurisdictions.

4.1 Dataset and Methodology

The Route Views project [21] makes an archive available containing snapshots of the BGP routing information bases (RIB) every 2 hours of various BGP speakers around the world that peer with many other routers. We downloaded all of the Route Views San Francisco Internet Exchange RIB snapshots from April, 2016 to perform our analysis⁶.

Since the Route Views BGP speaker peers with several routers, we selected a single peer and filtered out the others to simulate the view a smaller organization would have (one single ISP BGP peer). The AS we selected was 32354, which belongs to a regional ISP in the San Francisco Bay area. After this filtering, the BGP RIBs reflected what any organization receiving BGP service from the same ISP would have had during that same time period.

For each of the 50 banking IPs, we performed the following steps on each snapshot:

1. Parse the snapshot into a list of prefixes with AS paths.
2. Perform a longest-prefix-match to find the preferred prefix and AS path for the target IP.
3. Convert the set of AS numbers in the path into country codes.
4. Store the result as a set of transited countries for the IP at that point in time.

Steps 2 is a standard operation every router performs to match an IP address to a routing table entry. We have to match that logic to know which AS path will be selected.

Once we collected each IP's country sets for every 2 hours over the month, we compared each set to see if it changed; and, if it did, we examined how long the changes lasted. We then analyzed each change to see what types of exfiltration policies they may have violated.

4.2 Results

Out of the 50 financial services IP addresses we analyzed over the 1-month period, only 4 of them experienced country

changes in the path to their destination and 1 experienced an outage. That means that for 46 of the IP addresses, an extremely strict exfiltration policy restricting to no AS path country changes would not have caused any outages during this time period from our vantage point in San Francisco.

For the remaining 4, we examined each individual case to understand how an exfiltration policy would have affected communication with the target IP. The country codes for the 4 IPs that did experience changes are highlighted in table 1.

Online.citibank.co.in is the website for Citi India. During the 1 month window, the IP address was being advertised from a prefix in the Netherlands (NL). For a period of 24 hours ranging from April 5th to April 6th, routes to that IP were completely withdrawn from the routing table. So a country-based exfiltration policy would not have negatively impacted communication with this site.

www.nbg.gr is the website for the National Bank of Greece. Its path changed on April 9th to a direct peering between a US provider and a Greek provider (GR), eliminating an intermediary peer in the EU. Due to the elimination of a country from the path, the route became more direct and therefore would not have been impacted by any country-based exfiltration policies.

Hsbc.com belongs to HSBC, which is a large international bank based in London. The path to their site alternated between being advertised directly in the US and being advertised by an ISP in Great Britain (GB) twice in the one month period. If an exfiltration policy had restricted it to US-only paths, it would have caused an 8 day outage and then a 3 day outage. However, when communicating with a bank headquartered in London, we believe it would be reasonable to have Great Britain in the permitted list, in which case there would have been no outages.

Icbc.com.cn belongs to the Industrial and Commercial Bank of China. At the start of the period, the path was from the US to China (CN) via Hong Kong (HK). For a brief period on April 5th, a direct route to China was advertised that removed Honk Kong from the path. However, the path via Honk Kong was restored <4 hours later. Because the path became more direct during the change, any country-based exfiltration policies would not have been violated.

Bsi.ir belongs to Bank Saderat Iran, headquartered in Tehran, Iran. From the start of the period until April 17th, the path traversed the US, Russia (RU), Azerbaijan (AZ), and Iran (IR). For the next 8 days after that, it traversed Oman (OM), the US, and Iran. Then for a short period on April 25th (<4 hours), it traversed Germany (DE), the US, and Iran; after which it switched back to the Oman route. Then again on the 28th it traversed an EU ISP for < 4 hours before switching back to the Oman route. It is difficult to speculate what an exfiltration policy might look like for a company in San Francisco communicating with a bank in Tehran; but, for the sake of illustration, let us assume it had adopted a policy to refuse any paths via EU countries. This would have caused two small outages during the month (<4 hours each) — one on the 25th when it transited Germany and one on the 28th with it transited another EU ISP.

What the data from the time period and selected targets (top 50 banking sites) shows is that transited countries are relatively stable over time. The majority of the IP addresses incurred no country churn during the entire month,

⁵http://www.alexa.com/topsites/category/Business/Financial_Services/Banking_Services

⁶<http://archive.routeviews.org/route-views.sfmix/bgpdata/>

Site	Changes
Online.citibank.co.in	$NL, US \rightarrow N/A \rightarrow NL, US$
www.nbg.gr	$EU, GR, US \rightarrow GR, US$
Hsbc.com	$US \rightarrow GB, US \rightarrow US \rightarrow GB, US$
Icbc.com.cn	$CN, HK, US \rightarrow CN, US \rightarrow CN, HK, US$
Bsi.ir	$AZ, IR, RU, US \rightarrow IR, OM, US \rightarrow DE, IR, US \rightarrow IR, OM, US \rightarrow EU, IR, US \rightarrow IR, OM, US$

Table 1: Country Path Changes for top 50 Alexa Banking sites in April 2016

and the ones that did were hosted by companies headquartered in other countries. For the international sites, only two IPs would have experienced outages if exfiltration policies did not allow any new countries to show up in the path (*Hsbc.com* and *Bsi.ir*). Out of the two, *Hsbc.com* would likely have an exfiltration policy permitting Great Britain to begin with, leaving only *Bsi.ir* as the the only one that likely would have experienced outages.

With such a low churn rate, we believe that country-based transit policies would be very manageable by a security team at an organization. Especially if it is a policy that protects very few targets and only excludes a small set of untrusted countries.

5. PROOF OF CONCEPT CODE

To test our approach, we built a policy engine on top of ExaBGP [7], an open source BGP engine to peer with BGP speakers and output route information. The source code for our project is available under a BSD license⁷.

The policy engine takes a few simple configuration values:

- Peering info for BGP update feed.
- Path to ASN to country code CSV file.
- IP policies which take the form of (prefix, allowed country codes, blocked country codes). This allows a white-list or black-list approach.
- Output path for ACL entries.⁸

The engine never propagates information back to the organization’s BGP router and it just generates text files containing the blocked IPs so it is not invasive to test it to see what it would block under various constraints. This also means that performance of the system is not a concern because the only thing it will impact is the time it takes to generate block rules in reaction to country changes.

6. FUTURE WORK

Blocking traffic when a transit policy is violated is effective, but can cause potentially significant downtime depending on the lifetime of the undesired path. One thing we are examining is automating traffic redirection to a different

destination based on IP header rewriting. With the participation of a reflection service or a different remote office, transit violations could be avoided in an automated fashion at the cost of latency (and bandwidth for the reflector).

7. LIMITATIONS

One of the primary concerns with this approach is one of the intermediary ISPs sending traffic via another autonomous system that was not advertised in the BGP updates it was propagating. The expectation of BGP is that an AS will send traffic along to the next AS in the path that it is advertising. Violating this makes it difficult for other BGP participants to choose the shortest route and it risks creating routing loops. However, it is possible for an ISP to do this and it would not be visible in our BGP updates. We have not quantified how frequently this occurs.

The other limitation with this approach is that some traffic may be routed via a blocked AS immediately after a topology change before the BGP update is propagated back to the organization’s router and the policy engine. Depending on how disruptive the topology change was that caused it, it may take several minutes for BGP to converge and get the updated path to the policy engine. While this would leak some packets to the blocked AS, it still prevents any kind of long-term traffic analysis.

8. CONCLUSION

Defense in depth can include not sending sensitive information across the network where attackers may capture and analyze large amounts of information. Encrypting data transmissions ensures the security of content when all steps are implemented correctly. This includes selection of keys as well as certificate checking. Session keys may be weak, software may be compromised, certificates may be badly formed, based on incorrect information, or use weak algorithms. The threat of zero day vulnerabilities includes even improved factoring. Avoiding exfiltration of data in transit can offer another level of protection for organizations willing to exchange some availability for greater confidentiality.

We have shown that the increase in confidentiality does not need to create significant disruptions in availability. We proposed a system to mitigate data exfiltration for traffic that crosses the Internet by generating real time firewall rules blocking connectivity to sensitive IP addresses when routes would cause transit through a BGP peer in an untrusted country. We provided an analysis based on historical BGP routing data that showed this method would not adversely impact communication with the top 50 banking sites even under strict country restrictions. We then introduced a simple proof-of-concept that can be used to experiment with this approach and build on this research in the future.

The proof of concept in this paper is achieved by inspecting paths in BGP updates and cross referencing the AS numbers with AS registration info. Other information for selecting ASes to avoid can be found in related work, using reputations as well as jurisdictions. The analysis implemented here can be repeated by any party on the network, to examine the possible impact of adopting rules which exclude potentially risky paths. We are seeking partners interested in limited data sharing for additional investigation into mitigating risks of exfiltration via hijacking or transmission across untrusted nodes.

⁷<https://github.com/kevinbenton/exabgp>

⁸Currently the policy engine just generates a list of the IPs that should be blocked and writes them to a file under the expectation that another tool will convert that into ACLs in the format required by the organizations filtering infrastructure (e.g. Router ACLs, Firewall Rules, OpenFlow rules, black-hole routes).

Acknowledgments

This material is based upon work supported, in part, by the DHS BAA 11-02-TTA 03-0107 Contract N66001-12-C-0137, Cisco Research Support Proposal 591000, Google Privacy & Security Focused Research Program and NSF CISE #1565252: Living in the Internet of Things. Any opinions, findings, and conclusions or recommendations expressed in this material are those of the author(s) and do not necessarily reflect the views of the DHS, NSF, DoD, Google, Cisco, or Indiana University.

9. REFERENCES

- [1] R. Anderson and C. Hall. *Collaborating with the Enemy on Network Management (Transcript of Discussion)*, pages 163–171. Springer International Publishing, Cham, 2014.
- [2] B. Andree Toonk. Massive route leak causes internet slowdown. <http://www.bgpmon.net/massive-route-leak-cause-internet-slowdown/>, 2015.
- [3] W. Andy Greenberg. Hacker redirects traffic from 19 internet providers to steal bitcoins. <http://www.wired.com/2014/08/isp-bitcoin-theft/>, 2014.
- [4] S. Bellovin, R. Bush, and D. Ward. Rfc 7353: Security requirements for bgp path validation. Technical report, 2014.
- [5] J. Chang, K. K. Venkatasubramanian, A. G. West, S. Kannan, I. Lee, B. T. Loo, and O. Sokolsky. As-cred: Reputation and alert service for interdomain routing. *Systems Journal, IEEE*, 7(3):396–409, 2013.
- [6] B. Eriksson, P. Barford, J. Sommers, and R. Nowak. A learning-based approach for ip geolocation. In *Passive and Active Measurement*, pages 171–180. Springer, 2010.
- [7] Exa-Networks. exabgp. <https://github.com/Exa-Networks/exabgp>, 2015.
- [8] V. Garg and L. J. Camp. Macroeconomic analysis of malware. In *NDSS*, 2013.
- [9] J. Gersch and D. Massey. Rover: Route origin verification using dns. In *Computer Communications and Networks (ICCCN), 2013 22nd International Conference on*, pages 1–9. IEEE, 2013.
- [10] M. S. Henry Tan and W. Zhou. Data-plane defenses against routing attacks on tor. *Proceedings on Privacy Enhancing Technologies*, 2016, 2016.
- [11] R. Hiran, N. Carlsson, and N. Shahmehri. Crowd-based detection of routing anomalies on the internet. 2015.
- [12] X. Hu and Z. M. Mao. Accurate real-time identification of ip prefix hijacking. In *Security and Privacy, 2007. SP’07. IEEE Symposium on*, pages 3–17. IEEE, 2007.
- [13] R. Jim Cowie. The new threat: Targeted internet traffic misdirection - dyn research. <http://research.dyn.com/2013/11/mitm-internet-hijacking/>, 2013.
- [14] J. Karlin, S. Forrest, and J. Rexford. Pretty good bgp: Improving bgp by cautiously adopting routes. In *Network Protocols, 2006. ICNP’06. Proceedings of the 2006 14th IEEE International Conference on*, pages 290–299. IEEE, 2006.
- [15] E. Katz-Bassett, J. P. John, A. Krishnamurthy, D. Wetherall, T. Anderson, and Y. Chawathe. Towards ip geolocation using delay and topology measurements. In *Proceedings of the 6th ACM SIGCOMM conference on Internet measurement*, pages 71–84. ACM, 2006.
- [16] W. Kim Zetter. Someone’s been siphoning data through a huge security hole in the internet. <http://www.wired.com/2013/12/bgp-hijacking-belarus-iceland>, 2013.
- [17] M. Lepinski and S. Kent. Rfc 6480: an infrastructure to support secure internet routing. internet engineering task force (ietf), 2012.
- [18] Q. Li, M. Xu, J. Wu, X. Zhang, P. P. Lee, and K. Xu. Enhancing the trust of internet routing with lightweight route attestation. *Information Forensics and Security, IEEE Transactions on*, 7(2):691–703, 2012.
- [19] R. Martin Brown. Pakistan hijacks youtube. <http://research.dyn.com/2008/02/pakistan-hijacks-youtube-1/>, 2008.
- [20] C. McArthur and M. Guirguis. Stealthy ip prefix hijacking: don’t bite off more than you can chew. In *Global Telecommunications Conference, 2009. GLOBECOM 2009. IEEE*, pages 1–6. IEEE, 2009.
- [21] U. of Oregon. Route views project. <http://www.routeviews.org/>, 2016.
- [22] I. Poese, S. Uhlig, M. A. Kaafar, B. Donnet, and B. Gueye. Ip geolocation databases: Unreliable? *ACM SIGCOMM Computer Communication Review*, 41(2):53–56, 2011.
- [23] J. Qiu, L. Gao, S. Ranjan, and A. Nucci. Detecting bogus bgp route information: Going beyond prefix hijacking. In *Security and Privacy in Communications Networks and the Workshops, 2007. SecureComm 2007. Third International Conference on*, pages 381–390. IEEE, 2007.
- [24] T. Qiu, L. Ji, D. Pei, J. Wang, J. J. Xu, and H. Ballani. Locating prefix hijackers using lock. In *USENIX Security Symposium*, pages 135–150, 2009.
- [25] Y. Rekhter and T. Li. A Border Gateway Protocol 4 (BGP-4). RFC 1654, IETF, July 1995.
- [26] M. Van Eeten, J. M. Bauer, H. Asghari, S. Tabatabaie, and D. Rand. The role of internet service providers in botnet mitigation an empirical analysis based on spam data. TPRC, 2010.
- [27] P.-A. Vervier, O. Thonnard, and M. Dacier. Mind your blocks: On the stealthiness of malicious bgp hijacks. In *NDSS*, 2015.
- [28] Yakov Rekhter, Tony Li, and S. Hares. A Border Gateway Protocol 4 (BGP-4). RFC 4271, IETF, January 2006.
- [29] Z. Zhang, Y. Zhang, Y. C. Hu, and Z. M. Mao. Practical defenses against bgp prefix hijacking. In *Proceedings of the 2007 ACM CoNEXT conference*, page 3. ACM, 2007.
- [30] Z. Zhang, Y. Zhang, Y. C. Hu, Z. M. Mao, and R. Bush. Ispy: detecting ip prefix hijacking on my own. In *ACM SIGCOMM Computer Communication Review*, volume 38, pages 327–338. ACM, 2008.