

# Diagnosis of Route Leaks Among Autonomous Systems In The Internet

M. S. Siddiqui<sup>†‡</sup>, D. Montero<sup>†</sup>, M. Yannuzzi<sup>†</sup>, R. Serral-Gracià<sup>†</sup>, X. Masip-Bruin<sup>‡</sup>

<sup>†</sup>Networking and Information Technology Lab (NetIT Lab)

<sup>‡</sup>Advanced Network Architectures Lab (CRAAX)

Technical University of Catalonia

Vilanova i la Geltrú, Spain

Email: {siddiqui, dmontero, yannuzzi, rserral, xmasip}@ac.upc.edu

**Abstract**—Border Gateway Protocol (BGP) is the defacto inter-domain routing protocol in the Internet. It was designed without an inherent security mechanism and hence is prone to a number of vulnerabilities which can cause large scale disruption in the Internet. Route leak is one such inter-domain routing security problem which has the potential to cause wide-scale Internet service failure. Route leaks occur when Autonomous systems violate export policies while exporting routes. As BGP security has been an active research area for over a decade now, several security strategies were proposed, some of which either advocated complete replacement of the BGP or addition of new features in BGP, but they failed to achieve global acceptance. Even the most recent effort in this regard, lead by the Secure Inter-Domain Routing (SIDR) working group (WG) of IETF fails to counter all the BGP anomalies, especially route leaks. In this paper we look at the efforts in countering the policy related BGP problems and provide an analytical insights into why they are ineffective. We contend a new direction for future research in managing the broader security issues in the inter-domain routing. In that light, we propose a naive approach for countering the route leak problem by analyzing the information available at hand, such as the RIB of the router. The main purpose of this paper was to position and highlight the autonomous smart analytical approach for tackling policy related BGP security issues.

## I. INTRODUCTION

The BGP protocol [1] has been the default protocol for exchanging reachability information in the inter-domain arena of the Internet. The current version of BGP, the BGPv4 protocol, does not provide any performance or security guarantees. The intrinsic assumption of trust on the information exchanged between the ASes through the BGP protocol does not stand realistic in today's Internet. The heavy reliance of mission critical applications on the Internet have played a vital role in motivating the increased interest in improving the security of the Internet and its infrastructure. Although the BGP always remained part of the research community due to its performance issues but the recent global outages in the Internet acted as a catalyst for shifting the research focus towards securing it.

This implicit trust model among BGP speakers paves the way for a number of vulnerabilities and attacks which have led to paralysis of Internet services globally as well as in particular regions. There are several known attacks on the BGP but the main security problems include false IP prefix origination and

false route propagation, which can result in Denial of Service (DoS) or in traffic sniffing, in the worst case.

An deceptively simple but complex security problem of inter-domain routing that has caused large scale disruption in Internet is Route Leak. A route leak occurs when a route is not advertised according to the business relationship or the link classification. In February 2012, a misconfiguration at a multi-homed ISP leaked all its internal routes to one of its providers, including the routes from other providers, causing a national level disruption in Internet service in Australia [2].

The previous attempts in securing the BGP include several fault detection mechanisms and protocols [3]–[10], suggesting minor to major changes or replacement of the BGP protocol, but none succeeded practically. More recently, an IETF WG, Secure Inter-Domain Routing (SIDR) has put forward a set of comprehensive recommendations for securing the BGP. A couple of these recommendations have already been implemented and are in the testing phase. However, even the huge security infrastructure proposed by SIDR WG for securing the BGP, does not secure the BGP so as to avoid all the BGP anomalies, specifically the policy related attacks. This is because most of the security proposals, including SIDR WG, approach BGP security from an operational perspective and neglect the business policies among the ASes.

In this paper, we focus on the route leak problem only and highlight our research direction for resolving it. We identify the occurrence scenarios of route leaks and diagnose why they succeed in causing large scale service failures. Furthermore, we provide research direction for possible pragmatic ways to avoid, detect, or counter route leaks among ASes in inter-domain routing. The paper is organized into five sections: firstly, Section II provides a brief overview about the BGP protocol and the policies in the inter-domain routing. Secondly, Section II describes the main policy related security problem, namely Route Leak. Section III briefly describes the conventional and recently developed mechanisms that could be used to counter route leak problem. Next, we present a new direction for research in solving the policy related BGP security problems in Section IV and finally the Conclusions in Section V.

## II. ROUTE LEAKS

Before describing route leak problem in detail, it is essential to understand at least the basic fabric of inter-domain routing. Currently, there are around 45,000 ASes in the Internet and reachability among them is achieved when each AS informs its neighbors (*directly connected ASes*) about its available IP prefixes according to the particular business relationship it has with each of them. The business relation between any two ASes dictates the kind of policies that would be implemented on that particular link. The relation between two ASes can be broadly classified into either Provider-Customer or Peer-Peer relation. In the latter case, both ASes advertise subsets of their routes, i.e., they only advertise their own or their customer's routes to each other. A Peer-Peer relation may not involve financial settlement up to a certain agreed traffic ratio. For the Provider-Customer relation, Provider and Customer are the opposite ends of the same link. That is, from Customer AS's perspective, it will only advertise its own routes and the routes of its Customers (i.e., *Customer's Customer routes*) toward its Provider link. From Provider AS's perspective, it will advertise all routes toward its customer hence providing it transit to rest of the Internet. The Provider charges its Customer for forwarding Customer's traffic to and fro.

In terms of maximizing revenues, an AS prefers a Customer route over a Peer or Provider route, and prefers a Peer route over a Provider route. In a nut-shell, the following Gao-Rexford model [11] is adopted by an AS to minimize expenses and maximize earnings:

- 1) Routes learned from Customers are further advertised to other Customers, Peers and Providers.
- 2) Routes learned from Peers are further advertised to Customers only.
- 3) Routes learned from Providers are further advertised to Customer's only.

The protocol used to exchange the routing information is the BGP protocol. The BGP enabled routers, called BGP speakers, setup a BGP session with their direct neighbors using TCP. The setting up of the BGP session among BGP speakers is called BGP peering. Two BGP peers inform each other about their routing information through BGP updates. The BGP updates are exchanged on regular basis to inform the neighbors about the recently learned new routes or the withdrawn routes along with their attributes. A BGP route has multiple attributes associated to it which assist the receiving BGP speaker in selecting the best route. The ASes implement their business policies by tuning the attributes of a route before considering it for route selection or advertising it further. The BGP protocol executes route selection and propagation process based on each AS's individual routing policy.

As mentioned earlier, the lack of any mechanism to verify the reachability information exchanged between two BGP speakers gives rise to vulnerabilities which can be exploited. Given the complex operation of the BGP, a number of BGP anomalies can occur due to misconfigurations or malicious intent but in this paper we focus on the Route Leak problem

and describe its causes.

### A. Diagnosis of Route Leaks

A route leak occurs when a route gets advertised over a link by an AS, which does not coincide with the link classification. The AS advertising this route can either be the owner of this route or had received this route from another AS where it was in accordance with the link classification [12].

In this regard, the valley-free rules can be used as basis for providing an initial definition of the route leak problem, i.e., if a route is advertised by an AS toward its neighbor AS such that it is in violation of any of the three valley-free rules, then it is called a route leak. Furthermore, we notice that first rule in Gao-Rexford model cannot be infringed, since an AS can always export customer routes independently of the business relationship it has with the neighbor to which it is exporting the route to. This implies that a route leak can not occur while exporting routes toward customer ASes i.e., a provider AS can not leak a route toward its customer ASes. We can further deduce that route leaks can occur while exporting routes either to a provider AS or a peer AS.

For better understanding the route leaks, let us represent the Internet as an undirected graph  $G = (V, E)$ , where  $V$  corresponds to set of all ASes in the Internet and  $E$  corresponds to set of all links between the ASes, then for a particular autonomous system  $v$ , let us define

$\mathbb{P}_v$ : set of all the providers of  $v$

$\mathbb{J}_v$ : set of all the peers of  $v$

$\mathbb{C}_v$ : set of all the customers of  $v$

and for  $x$  and  $y$  representing an AS or set of ASes, we define

$\mathbb{R}_{x,y}^I$ : set of all the routes imported by  $x$  from  $y$

$\mathbb{R}_{x,y}^E$ : set of all the routes exported by  $x$  to  $y$

$\mathbb{O}_v$ : the set of all the routes owned by  $v$

Mindful of that fact that route leak can occur while exporting routes toward a peer AS or a provider, we define:

$$\mathbb{R}_{v,j}^E \subseteq \mathbb{R}_{v,\mathbb{C}_v}^I \cup \mathbb{O}_v \quad \text{where, } j \in \mathbb{J}_v \quad (1)$$

$$\mathbb{R}_{v,p}^E \subseteq \mathbb{R}_{v,\mathbb{C}_v}^I \cup \mathbb{O}_v \quad \text{where, } p \in \mathbb{P}_v \quad (2)$$

as sets of routes that  $v$  can export to a peer or a provider, respectively. That is an AS can only export its own routes and its customers' routes toward its peer ASes or provider ASes.

Furthermore, (1) and (2) can also be expressed as:

$$\mathbb{R}_{v,j}^E \not\subseteq \mathbb{R}_{v,\mathbb{J}_v}^I \cup \mathbb{R}_{v,\mathbb{P}_v}^I \quad \text{where, } j \in \mathbb{J}_v \quad (3)$$

$$\mathbb{R}_{v,p}^E \not\subseteq \mathbb{R}_{v,\mathbb{J}_v}^I \cup \mathbb{R}_{v,\mathbb{P}_v}^I \quad \text{where, } p \in \mathbb{P}_v \quad (4)$$

That is to say that the set of routes exported by an AS toward a peer AS should not contain any routes learned from peer or provider ASes. Similarly, the set of routes exported by an AS toward a provider AS should not contain any routes learned from peer, provider ASes.

Using (3), we can define route leak as:

**Definition 1.** "For an AS  $v$ , if  $\mathbb{R}_{v,j}^E$ , where  $j \in \mathbb{J}_v$ , contains a route  $r$ , such that  $r \in \{\mathbb{R}_{v,\mathbb{J}_v}^I \cup \mathbb{R}_{v,\mathbb{P}_v}^I\}$ ,




Fig. 1. Customer Route Leak.

that is,  $\mathbb{R}_{v,j}^E \cap \{\mathbb{R}_{v,\mathbb{J}_v}^I \cup \mathbb{R}_{v,\mathbb{P}_v}^I\} \neq \emptyset$ , then it is a route leak”

Thus, an AS exporting a route toward its peer AS which it learned from another peer AS or provider AS falls in the category of route leaks. Similarly, using (4) we define route leak as,

**Definition 2.** “For an AS  $v$ , if  $\mathbb{R}_{v,p}^E$ , where  $p \in \mathbb{P}_v$ , contains a route  $r$ , such that  $r \in \{\mathbb{R}_{v,\mathbb{J}_v}^I \cup \mathbb{R}_{v,\mathbb{P}_v}^I\}$ , that is,  $\mathbb{R}_{v,p}^E \cap \{\mathbb{R}_{v,\mathbb{J}_v}^I \cup \mathbb{R}_{v,\mathbb{P}_v}^I\} \neq \emptyset$ , then it is a route leak”

Based on **Definition 1** and **Definition 2**, from a victim’s perspective, route leaks can be classified as either Peer Route Leaks (PRL) or Customer Router Leaks (CRL).

1) *Customer Route Leak*: The route leak occurring on a Customer-Provider link is termed as a Customer route leak. In Fig. 1,  $AS_3$  and  $AS_4$  have Peer-Peer relation.  $AS_1$  is customer of both,  $AS_3$  and  $AS_4$ , i.e., it is multi-homed. The  $AS_3$  has another customer  $AS_2$  which owns the IP prefix 10.1.1.0/24. The  $AS_2$  advertises the prefix 10.1.1.0/24 to its Provider  $AS_3$  (Step i). The  $AS_3$  being provider of  $AS_1$  and peer of  $AS_4$ , advertises it to both of them (Step ii a and ii b).

Now if  $AS_1$  advertises 10.1.1.0/24 to  $AS_4$  (Step iii), this falls in the category of route leak. In this case,  $AS_4$  gets advertisements of the same prefix from  $AS_1$  and  $AS_3$ . Usually customer’s routes are preferred over peer routes, so  $AS_4$  selects  $AS_1$  as its next hop for the IP prefix 10.1.1.0/24, which apart from being sub-optimal, also allows  $AS_1$  to sniff all the traffic between  $AS_3$  and  $AS_4$ . The increase in the length of AS-Path of the route advertised by  $AS_1$  does not help in detecting this problem because customer routes are preferred by setting the ‘local pref’ attribute, which is evaluated before the AS-Path attribute hence the route is decided before comparing the AS-Path length.

2) *Peer Route Leak*: Route leaks are also possible on a Peer-Peer link. In Fig. 2, the steps (a), (b)i, (b)ii, (c)i, (c)ii and (d) illustrate how route leaks can occur on a Peer-Peer link.  $AS_5$  advertises its IP prefix 11.1.1.0/24 to its provider.  $AS_3$  and  $AS_4$  learn about this prefix from their respective providers,  $AS_1$  and  $AS_2$ , respectively. Now, if  $AS_3$  advertises 11.1.1.0/24 to  $AS_4$ , then  $AS_4$  would prefer the peer route

over it’s provider route resulting in a route leak causing the traffic between  $AS_4$  and  $AS_5$  to go through  $AS_3$ , making it vulnerable to sniffing.

In subsequent sections, we briefly describe the available and newly proposed security measures which can be used to curtail route leak problem. Then we propose a novel direction for further research in this field to resolve route leaks. We discuss the apparent pros and cons of our proposed research direction and also provide explanation on how it can be integrated into BGP operations causing minimum entropy in the inter-domain routing.

### III. STATE-OF-THE-ART

In this section we describe the tools, techniques and methodologies, along with their pros and cons, that can be used to counter the Route Leak problem.

#### A. IETF Secure Inter-Domain Routing Working Group

The Secure Inter-Domain Routing (SIDR) working group (WG) of IETF consists of experts from industry and academia. The SIDR WG has put forward three proposals for securing the BGP. We present these mechanisms in a bit detail to have a better understanding of recent developments in securing the BGP.

The three proposals recommended by SIDR WG include a security infrastructure to support attestations and verifications related to the resources, a mechanism to secure origination of routing announcements and a scheme to secure route propagation in the BGP. The security infrastructure is called Resource Public Key Infrastructure (RPKI) [13]. The mechanism to counter false route origination is called Route Origin Authorization (ROA) [14]. The scheme to secure route propagation is called BGPSEC [15]. The ROA and the BGPSEC heavily rely on the RPKI to achieve their goals. A complete list of related RFCs and drafts can be found in [16].

1) *Resource Public Key Infrastructure (RPKI)*: The RPKI is a distributed publication infrastructure of cryptographically processed information to support third party resource verifications. The RPKI reflects an administrative resource allocation

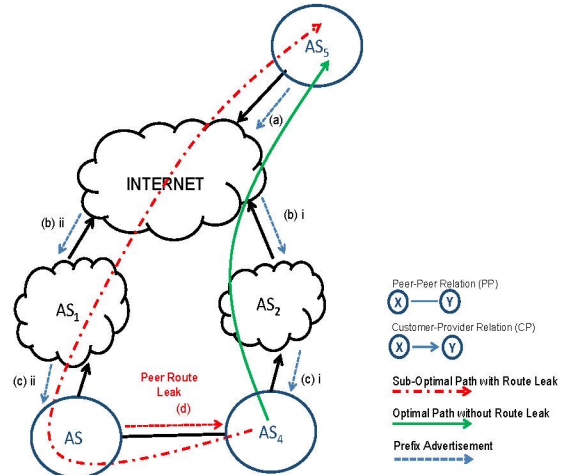


Fig. 2. Peer Route Leak.

hierarchy, where resources are distributed from Internet Assigned Numbers Authority (IANA) all the way down to End User with the use of digital certificates, forming a chain of verifiable trust.

Given such security skeleton, ASes can obtain certificates, called End Entity (EE) certificates, for the resources they own from the concerned resource allocation authorities. These certificates are regularly published in the respective RPKI repository of the AS. Each AS can have their own RPKI cache which is synchronized and updated regularly with the global RPKI. The ROA and the BGPSEC extensively rely on RPKI to achieve their goals, that is for verifying route origin advertisements and securing route propagation updates, respectively.

2) *Route Origin Authorization (ROA)*: The route origin authorization (ROA) proposal of SIDR WG targets the traffic hijacking problem due to false route origin advertisements. The ROA proposal makes use of RPKI credentials to assure integrity in the route origin announcements. It is achieved by the use of a particular signed authority, called Route Origination Authorization. The RPKI enables the legitimate owner of an IP prefix to produce an ROA with the help of EE certificate. This signed information binds the IP prefix resource with its owner's ASN while including the corresponding EE certificate.

Now, when an AS announces a particular IP prefix as its owner then with the help of RPKI a Relying Party (RP) can verify if this route origination announcement is legitimate or not. The RP queries the RPKI to confirm whether or not the announced IP resource belongs to the advertising AS. The response of the query from the RPKI can be used to take further decision according to the internal policy of the AS. However, the ROA requires changes to the BGP itself for performing IP prefix origin validation.

3) *Securing Route Propagation*: The BGPSEC protocol [15] provides a mechanism, based on public key cryptography, to secure the AS-Path information of an advertised route. In BGPSEC, a BGP speaker has additional router certificates and a pair of cryptographic keys which allow the BGPSEC router to sign BGP updates, making them BGPSEC updates. The BGPSEC is to be practiced only on the edge routers, i.e., among eBGP routers. This means the BGP updates are signed and validated on the boundaries of ASes only. The update originating AS signs the IP-Prefix, its ASN, ASN of next hop, and pCount field as shown in Fig. 3. The pCount enables optimized way of performing AS-Path prepending as well. The inclusion of Target ASN in the signature secures the forward direction of the update. Fig. 3 illustrates the use of signatures in a BGPSEC update origination and propagation. When a BGPSEC speaker receives a BGPSEC update, it verifies the update, using RPKI resources, and can propagate it to its peers either as a BGPSEC update or a BGP-4 update. This flexibility facilitates partial deployment scenarios.

The BGPSEC requires changes in the way BGP operates along with the requirement of a new BGP attribute, called BGPSEC\_Path. The introduction of changes in the BGP

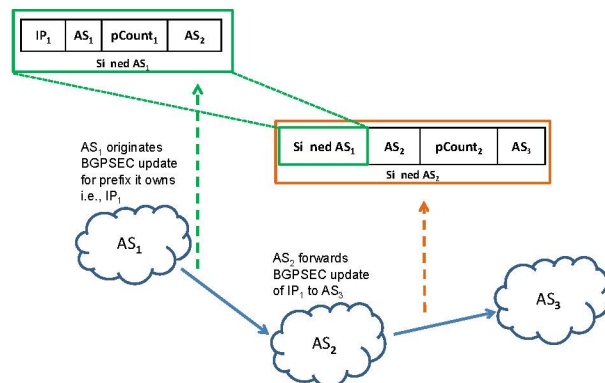


Fig. 3. BGPSEC Propagation with Forward Signing.

protocol makes this proposal susceptible to rejection, even though SIDR WG has given special considerations to partial deployment aspects and backward compatibility with BGP-4 protocol.

Even along with the huge excess baggage including, extra memory storage and high computational power for the BGP routers, new software and hardware equipments, changes to BGP protocol such as a new BGP attribute and path selection decision process, the SIDR proposals still fails to counter the route leaks. In Fig. 1, even if AS<sub>2</sub> had published an ROA for the IP prefix 10.1.1.0/24 in the RPKI and all the ASes in the scenario propagated BGPSEC updates, even then route leak would succeed if AS<sub>1</sub> advertises IP prefix 10.1.1.0/24 to AS<sub>4</sub>. The peer route leak would succeed in similar manner. This is because route leak exploits the default behavior of the BGP protocol where customer routes are preferred over peer routes and peer routes are preferred over provider routes. Furthermore, SIDR WG considers route leaks as out of the scope of its charter.

### B. Prefix Filters

Prefix filters, also called *route filters*, are a basic way of managing connections with other ASes and protecting BGP routers from accidental or intentional misconfigurations.

Prefix filters can be used both when receiving prefixes from a neighbor (i.e., ingress filter) and when advertising prefixes to a neighbor (i.e., egress filter).

The filters allow only those prefixes to pass through which have already been agreed upon by the two ASes. If an alien prefix is received or advertised, it would be filtered out at the ingress filter or egress filter, respectively. The egress filter are essential to counter 'fat finger' syndrome, that is mistakenly advertising more than the agreed set of prefixes. With proper ingress and egress filters, for all the prefix sets with the corresponding neighbors, an AS can counter route leak problem.

In Fig. 1, if AS<sub>4</sub> has an ingress prefix filter, for link connecting to AS<sub>1</sub>, which would only allow the agreed prefixes between AS<sub>1</sub> and AS<sub>4</sub>, then the route leak could have been filtered out. Similarly, an appropriate egress filter at AS<sub>1</sub> on the link towards AS<sub>4</sub> could also have countered this route leak in case of a misconfiguration. Even in the case of an intentional

route leak by  $AS_1$ , the ingress filter at  $AS_4$  mitigates the route leak. In the case that  $AS_1$  has new routes to advertise, it should inform its provider  $AS_4$  and  $AS_3$  and in order to accommodate the new prefixes they should update their respective egress and ingress filters.

The prefix filter mechanism works well for ASes which have small number of prefixes to announce or manage but the prefix filters' maintenance becomes challenging as the number of prefixes go into thousands. The maintenance of thousands of prefix filters give rise to operational challenges stemming from the administrative overhead, expense of maintenance, and accuracy and timeliness of the prefix entries in the filters [2]. Thus, as the maintenance cost goes high, most of the ASes rely on trust and do not maintain prefix filters.

The reliance on trust or poorly maintained prefix filters may have been the reason why the route leak succeeded between Dodo and its provider Telstra in Australia in 2012 [2]. Therefore, the prefix filters solution for route leaks is practically unreliable in face of scalability.

### C. Internet Route Registry

The Internet Route Registries (*IRR*) are databases of route objects arranged in a programmable format using the Routing Policy Specification Language (*RPSL*). An AS can define the IP-Prefixes it originates along with the route policies in the IRR and other ASes can query this route registry for their purposes. The IRR and RPSL provides an another way to manage route filters and can be used to avoid route hijackings and route leaks as well.

For IRRs to succeed as a reliable defense mechanism for the BGP problems, all the ASes should maintain all their route objects up-to-date at their respective IRRs, so that other ASes can rely on the output of their filters which would query IRRs for the route registry contents of their interest.

The responsibility of keeping all their route registries record up-to-date at IRRs, pushes a huge overhead burden toward each AS administration. Plus the complexity of RPSL language requires professional resources in order to appropriately define the route objects along with their policies at IRRs. Furthermore, the IRR mechanism was adopted and encouraged only in some regions of the world and in other regions the network operators preferred their own customized tools for this purpose. Due to these reasons and few intrinsic shortcomings of the IRR model, the currently available IRRs suffer from duplicate, false and incomplete records which puts their reliability and integrity into question.

In [17], Steenbergen confirms the problems faced by route registries and quantifies the usage of IRRs as well as the validity and accuracy of IRRs' data as insufficient.

### D. BGP Monitoring Tools

The extra and costly administrative burden of maintaining IP-Prefix filters or keeping the IRR records up-to-date led many ASes to develop their own proprietary tools or rely on third party tools for automating the respective processes.

The BGP monitoring tools use registry data, i.e., data from RIRs or IRRs, as well as probe BGP data at different points in the network to look for inconsistencies against the already defined filters and send alarms to concerned party in case of a violation.

MyASN/IS Alarms is an example of a BGP monitoring tool provided by RIPE NCC [18]. It provides Prefix hijacking detection using its Routing Information Service (*RIS*) Alarm System. The RIS consists of 15 Remote Route Collectors (*RRC*) around the globe. The system inspects BGP routes and updates collected at RRCs to detect Prefix hijacking. It notifies the concerned party through email or Syslog.

Nemecis, Prefix Hijack Alert System (PHAS), Pretty Good BGP (PGBGP), BGPmon are examples of monitoring tools used for BGP anomaly detection.

A common shortcoming of the above mentioned monitoring tools is that they can only detect problems where they have probes or route collectors. A little strategic attack avoiding their vantage points can succeed without detection. On the other hand, the administrative burden of defining the policy filters in the monitoring tools is still required, which brings us to the same place as IP-Prefix filters and IRRs. Furthermore, the victim has to rely on information provided by third party entities in order to react to security attacks, i.e., the monitoring and the anomaly detection is performed by an entity which is not under victim's direct management.

## IV. RESEARCH DIRECTION FOR RESOLVING ROUTE LEAKS

The conventional methods for securing AS policies including IP-Prefix filters, IRRs and 3rd party BGP monitoring tools fall short of providing an efficient and effective solution for route leaks due to the huge administrative burden required for keeping the routes and policies up-to-date as the number of IP prefixes scale up into thousands and reliance on third party information for making critical decisions.

The latest proposals from SIDR WG, including RPKI, ROA and BGPSEC do not counter all the BGP anomalies and attacks, especially the route leak problem. The RPKI, ROA and BGPSEC are being pushed forward by US government agencies like National Institute of Standards and Technology (NIST) as a solution for BGP security problems, however these proposals are facing resistance from the industry experts. Recently, engineers from Verisign Labs have exposed a synchronization delay problem in the distributed functionality of RPKI through an empirical study [19], which supports the hard feelings of network operators toward the SIDR proposals.

Most of the BGP security research was focused on the integrity, authentication, confidentiality, authorization, and validation of the BGP data, that is securing only the mechanics of the BGP. In order to counter security problems such as route leaks, it is important to explore other ways of thinking about securing the BGP which would include not only its operational aspects but business policies as well.

We propose an in-the-box approach for attempting to resolve the route leak problem. That is to develop methods and algo-

gorithms which utilize information available within the router, i.e., in the Routing Information Base (RIB), and route leak diagnosis presented in Section II-A to identify route leaks in real time. An example of in-the-box approach can be found in [20], where authors present a simple idea for countering route leaks using only the information available within the router.

In this regard, we propose another naive method for detecting route leaks following our new approach. As described in Section II-A, from the perspective of an AS which wants to detect route leaks, it can receive route leaks only from neighbor ASes with whom it has either customer or peer relationship. For example, in order for an  $AS_1$  to identify the route leaks it may receive from its neighbor customer  $AS_2$ , it computes customer cones for  $AS_2$  and maintains set of IP prefix belonging to the ASes in the customer cone  $IP - CC - AS_2$ . A customer cone is a set of ASes that can be reached through an AS following only provider-to-customer links. In case of the customer neighbor  $AS_2$ , for each advertised route,  $AS_1$  can check whether the advertised route is in  $IP - CC - AS_2$ . If the route advertised by  $AS_2$  does not belong to the  $IP - CC - AS_2$ , then it can suspect it as a route leak. This is because, given the relationship between  $AS_1$  and  $AS_2$ ,  $AS_2$  can only advertise the routes owned by itself or the routes of its customers and  $IP - CC - AS_2$  should contain all these routes. But there are valid exceptions when  $IP - CC - AS_2$  does not contain all the routes of  $AS_2$  or one of its customers due to policy routing or in case  $AS_2$  has a new customer. This suggests that the customer cone computation algorithm should be dynamic and self-learning. Initially, a customer cone can be computed using IP prefix filters during a training period. After the initial training period, IP prefix filters need not to be maintained and smart self-learning can take over for maintaining up-to-date IP prefix sets of customer cones for each neighbor customer or peer. In [21] authors mention different algorithms that can be used to compute customer cones, however more efficient and dynamic algorithms for computing customer cones in real time need to be further investigated.

## V. CONCLUSION

The conventional methodologies for avoiding route leaks, that are using IP-Prefix filters or IRRs, do not achieve the required goal practically. The latest security recommendations by the SIDR WG focus more on securing the operations of the BGP and fail to address the route leak problem as they do not take the AS policies into account. We believe that it is fundamental to change the way the route leak problem is approached and a new space of requirements and analytical tools, based on self-contained information including AS policies should be envisioned to attempt to resolve this problem. In this regard, we proposed a simple idea of applying smart analytics on BGP information available at hand, i.e., RIB and knowledge of direct neighbor AS relationships to detect route leaks without relying on third party information or alarms. The main purpose of this paper was to position and highlight the autonomous smart analytical approach for

tackling policy related BGP security issues. For future work, we intend to excel and refine the direction of research outlined in this paper with simulations and experimental results.

## ACKNOWLEDGMENT

This work was supported in part by an RFP granted by Cisco Systems, Inc., by the Spanish Ministry of Science and Innovation under contract TEC2012-34682, by the Catalan Government under contract 2009 SGR1508, and by IST Open-LAB Project under contract FP7-287581.

## REFERENCES

- [1] Y. Rekhter, T. Li, and S. Hares, "A Border Gateway Protocol 4 (BGP-4)," IETF, RFC 4271, January 2006.
- [2] G. Huston, "Leaking routes," <http://www.potaroo.net/ispcol/2012-03/leaks.html>, March 2012.
- [3] S. Kent, C. Lynn, J. Mikkelsen, and K. Seo, "Secure Border Gateway Protocol (S-BGP)," *IEEE Journal on Selected Areas in Communications*, vol. 18, no. 4, pp. 103–116, 2000.
- [4] R. White, "Securing BGP through secure origin BGP (soBGP)," *Internet Protocol Journal*, vol. 6, no. 3, 2003.
- [5] P. v. Oorschot, T. Wan, and E. Kranakis, "On interdomain routing security and pretty secure BGP (psBGP)," *ACM Trans. Inf. Syst. Secur.*, vol. 10, no. 3, Jul. 2007.
- [6] G. Goodell, W. Aiello, T. Griffin, J. Ioannidis, P. McDaniel, and A. Rubin, "Working Around BGP: An Incremental Approach to Improving Security and Accuracy of Interdomain Routing," in *Proc. Internet Society Symp. Netw. Distributed Syst. Security (NDSS) (2003)*, 2003.
- [7] Y.-C. Hu, A. Perrig, and M. Sirbu, "SPV: Secure Path Vector routing for securing BGP," *SIGCOMM Comput. Commun. Rev.*, vol. 34, no. 4, pp. 179–192, Aug. 2004.
- [8] Q. Li, M. Xu, J. Wu, X. Zhang, P. P. C. Lee, and K. Xu, "Enhancing the trust of Internet Routing with Lightweight Route Attestation," in *Proceedings of the 6th ACM Symposium on Information, Computer and Communications Security*, ser. ASIACCS '11. ACM, 2011, pp. 92–101.
- [9] M. Suchara, I. Avramopoulos, and J. Rexford, "Securing BGP incrementally," in *Proceedings of the 2007 ACM CoNEXT conference*, ser. CoNEXT '07. New York, USA: ACM, 2007, pp. 52:1–52:2.
- [10] W. Aiello, J. Ioannidis, and P. McDaniel, "Origin authentication in interdomain routing," in *Proceedings of the 10th ACM conference on Computer and Communications Security*, ser. CCS '03. New York, USA: ACM, 2003, pp. 165–178.
- [11] L. Gao and J. Rexford, "Stable Internet Routing without Global Coordination," *SIGMETRICS Perform. Eval. Rev.*, vol. 28, no. 1, pp. 307–317, jun 2000.
- [12] B. Dickson, "Route Leaks – Requirements for Detection and Prevention thereof," draft-dickson-sidr-route-leak-reqts, March 2012.
- [13] M. Lepinski and S. Kent, "An Infrastructure to Support Secure Internet Routing," IETF, RFC 6480, 2012.
- [14] M. Lepinski, S. Kent, and D. Kong, "A Profile for Route Origin Authorizations (ROAs)," IETF, RFC 6482, 2012.
- [15] M. Lepinski, "BGPSEC Protocol Specification," draft-ietf-sidr-bgpsec-protocol, February 2013.
- [16] Secure InterDomain Routing (SIDR) Working Group IETF, "http://datatracker.ietf.org/wg/sidr/," 2013.
- [17] R. Steenberg. (2008, October) NANOG 44: Examining the validity of IRR data. [https://www.nanog.org/meetings/nanog44/presentations/Tuesday/RAS\\_irrdata\\_N44.pdf](https://www.nanog.org/meetings/nanog44/presentations/Tuesday/RAS_irrdata_N44.pdf).
- [18] RIPE Network Coordination Center, "http://www.ripe.net/," 2013.
- [19] E. Osterweil, T. Manderson, R. White, and D. McPherson, "Sizing Estimates for a Fully Deployed RPKI," Verisign, Tech. Rep. 1120005 version 2, 2012.
- [20] M. S. Siddiqui, D. Montero, M. Yannuzzi, R. Serral-Gracia, and X. Masip-Bruin, "Route leak identification: A step toward making Inter-Domain routing more reliable," in *2014 / 10th International Conference on Design of Reliable Communication Networks (DRCN 2014)*, Ghent, Belgium, April 2014.
- [21] M. Luckie, B. Huffaker, A. Dhamdhare, V. Giotsas, and k. claffy, "As relationships, customer cones, and validation," in *Proceedings of the 2013 Conference on Internet Measurement Conference*, ser. IMC '13. New York, NY, USA: ACM, 2013, pp. 243–256.