The Policy-Based AS_PATH Verification to Monitor AS Path Hijacking

Je-Kuk Yun, Beomseok Hong
Information Technology
Towson University
Towson, U.S.A.
jyun4, bhong1@students.towson.edu

Yanggon Kim
Information Technology
Towson University
Towson, U.S.A.
ykim@towson.edu

***Abstract*— As the number of IP prefix hijacking incidents has increased, many solutions are proposed to prevent IP prefix hijacking, such as RPKI, BGPmon, Argus, and PHAS. Except RPKI, all of the solutions proposed so far can protect ASes only through the origin validation. However, the origin validation cannot detect specified attacks that alter the AS_PATH attribute, such as AS Insertion attack and Invalid AS_PATH Data Insertion attack. In order to solve these problems, SIDR proposed the RPKI using BGPSEC, but BGPSEC is currently a work in progress. So, we propose Secure AS_PATH BGP (SAPBGP) in which we monitor the AS_PATH attribute in update messages whether each AS in the AS_PATH attribute are connected to each other based on our policy database collected from RIPE NCC repository. Our analysis shows 4.57% of the AS_PATH attribute is invalid and 95.43% of the AS_PATH attribute is valid from the fifteenth of April in 2014 to the eighth of June in 2014. In addition, the performance test verifies that the SAPBGP can process all of the live BGP messages coming from BGPmon in real time.**

Keywords- border gateway protocol; interdomain routing; network security; networks; AS path hijacking.

I. INTRODUCTION

The Border Gateway Protocol (BGP) is the de-facto protocol to enable large IP networks to form a single Internet [1]. The main objective of BGP is to exchange Network Layer Reachability Information (NLRI) among Autonomous Systems (ASes) so that BGP routers can transfer their traffic to the destination.

However, BGP itself does not have mechanisms to verify if a route is valid because BGP speaker completely trusts other BGP speakers. This lack of consideration of BGP vulnerabilities often causes severe failures of Internet service provision [2]. The most well-known threat of the failures is the YouTube hijacking by Pakistan Telecom (AS17557) on the 24th of February in 2008 [3]. In response to the government's order to block YouTube access within their ASes, Pakistan Telecom announced a more specific prefix than YouTube prefix. Then, one of Pakistan Telecom's upstream providers, PCCW Global (AS3491), forwarded the announcement to other neighbors. As a result of this, YouTube traffic from all over the world was misled to Pakistan Telecom (AS17557) for two hours. In order to solve these problems, many studies were conducted, such as Resource Public Key Infrastructure (RPKI) [4], BGPmon [5], Argus [6], and a Prefix Hijack Alert System (PHAS) [7].

While there are many studies to IP prefix hijacking, few studies have been researched about AS path hijacking. There was some misdirected network traffic suspected of the man-in-the-middle (MITM) attack in 2013 observed by Renesys. In February 2013, global traffic was redirected to Belarusian ISP GlobalOneBel before its intended destination and it occurred on an almost daily basis. Major financial institutions, governments, and network service providers were affected by this traffic diversion in several countries including the U.S. From the thirty first of July to the nineteenth of August, Icelandic provider Opin Kerfi announced origination routes for 597 IP networks owned by a large VoIP provider in the U.S through Siminn, which is one of the two ISPs that Opin Kerfi has. However, this announcement was never propagated through Fjarskipti which is the other one of the two ISPs. As a result, network traffic was sent to Siminn in London and redirected back to its intended destination. Several different countries in some Icelandic autonomous systems and belonging to the Siminn were affected. However, Opin Kerfi said that the problem was the result of a bug in software and had been resolved [8].

In order to protect the AS path hijacking, the AS_PATH attribute should not be manipulated. However, the BGP itself cannot check whether the AS_PATH attribute has been changed or not. If a routing hijacker manipulates the AS_PATH attribute in a BGP message that is sent by another router and forwards the manipulated BGP message to other neighbors, the neighbors who receive the manipulated BGP message can be a victim of AS path hijacking. Only Secure Inter-Domain Routing (SIDR) working group proposed the RPKI using BGPSEC to validate the AS_PATH attribute, but BGPSEC is currently a work in progress [9]. In addition, a study propounds that BGP armed with BGPSEC cannot be secured because of BGP's fundamental design [10].

We propose Secure AS_PATH BGP (SAPBGP) in which the SAPBGP constructs its own policy-based database by collecting RIPE NCC repository and checks the AS_PATH attribute in BGP update messages whether or not the ASes listed in the AS_PATH attribute are actually connected. For the validation test with the real BGP messages, the SAPBGP receives a live BGP stream from the BGPmon project [11]. In addition, we conduct the performance test of the SAPBGP to measure the duration of the validation with the live BGP messages.

In this paper, with the fact that BGP is vulnerable to MITM attack, we describe an attack scenario and a solution in Section 3. In Section 4, we introduce and explain the SAPBGP in detail. We discuss the SAPBGP environment and analyze the result of the SAPBGP validation and the

 ISBN: 978-1-61208-376-6

performance test in Section 5. Lastly, we conclude the paper in Section 6.

II. RELATED RESEARCH

A. BGPsec

BGPsec is a mechanism to provide routing path security for BGP route advertisements and a work in progress by SIDR [9]. BGPsec relies on RPKI where the root of trust consists of the Regional Internet Registries (RIRs), including ARIN, LACNIC, APNIC, RIPE, and AFRINIC. Each of the RIRs signs certificates to allocate their resources. RPKI provides Route Origination Authorization (ROA) to ASes that are authorized to advertise a specific prefix [12]. The ROA contains the prefix address, maxlength, and AS number, which certifies the specified AS has permission to announce the prefixes. For routing path validation, each AS receives a pair of keys, which are a private key and a public key, from its RIR. Each AS speaker signs the routing path before forwarding it to their neighbors.

B. BGPmon

BGPmon is a monitoring infrastructure, implemented by Colorado State University that collects BGP messages from various routers that are distributed and offers the BGP messages as the routes for destinations are changed in real-time [5]. Any BGP speaker can be a source that offers real-time update messages if the BGP speaker is connected to BGPmon. Currently, 9 BGP speakers are participated in the BGPmon project as a source router. In addition, BGPmon collects Multi-threaded Routing Toolkit (MRT) format [13] live stream from the RouteViews project through indirect peering. The MRT format defines a way to exchange and export routing information through which researchers can be provided BGP messages from any routers to analyze routing information. Clients can be connected to the BGPmon via telnet and receive the live BGP stream in real time.

C. RIPE NCC

RIPE NCC is one of the Regional Internet Registries (RIRs) in charge of the Europe/Middle-East region. RIPE NCC manages RIPE Data Repository that is a collection of datasets, such as IP address space allocations and assignments, routing policies, reverse delegations, and contacts for scientific Internet research. The organizations or individuals who currently hold Internet resources are responsible for updating information in the database. As a result, RIPE NCC can keep the Data Repository up to date and provide database APIs so that data users can access the RIPE data repository through web browsers or programs.

III. BGP THREATS AND SOLUTION

In this section we introduce a scenario of the AS path hijacking that leads to the MITM attack. In addition, we discuss how the routing policy-based AS_PATH validation is operated in order to prevent the AS path hijacking.

A. Manipulating data in BGP updates

A BGP router inserts its own ASN into the AS_PATH attribute in update messages when the BGP router receives the update message from neighbors. However, the BGP router can insert one or more ASNs into the AS_PATH attribute in update messages other than its own ASN. In addition, a BGP router might pretend as if the BGP router is connected to a certain BGP router by manipulating data contained in BGP updates.

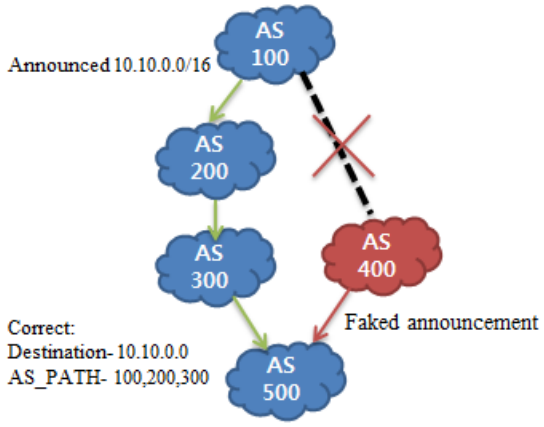

Figure 1. Manipulating a BGP message

Figure 1 demonstrates an example of manipulating data in BGP update messages. Suppose AS 400 has a connection to AS 500 and creates a faked\ BGP announcement to pretend that AS 400 received a BGP message originated by AS 100 and forwarded the update message to AS 500 even though AS 100 and AS 400 actually don't have a BGP connection. In terms of AS 500, the traffic heading for prefix 10.10.0.0/16 will choose AS 400 as the best path because AS 500 selects the shortest path and AS 400 is shorter than AS 300. Even if the AS 500 can conduct origin validation, the AS 500 cannot prevent this attack because prefix and ASN information is correct. As a result, AS 400 will have the traffic heading for prefix 10.10.0.0 and might start another attack using the traffic, such as a MITM attack.

B. Man-in-the-middle (MITM) attack

The man-in-the-middle attack is an active eavesdropping in which the attacker secretly creates connections to the victims and redirects large blocks of internet traffic between the sources and the destinations as if the sources and destinations communicate directly. In such a case, the victims can only notice a little enlarged latency time because the internet packets travel longer hops than normal. In the meantime, the attacker can monitor and manipulate the packets so that the attacker can create future chances to try another attack.

Renesys monitors MITM attacks and its clients were victims of route hijacking caused by MITM attacks for more than 60 days. The victims are governments, Internet Service Providers (ISP), financial institutions, etc.

 ISBN: 978-1-61208-376-6

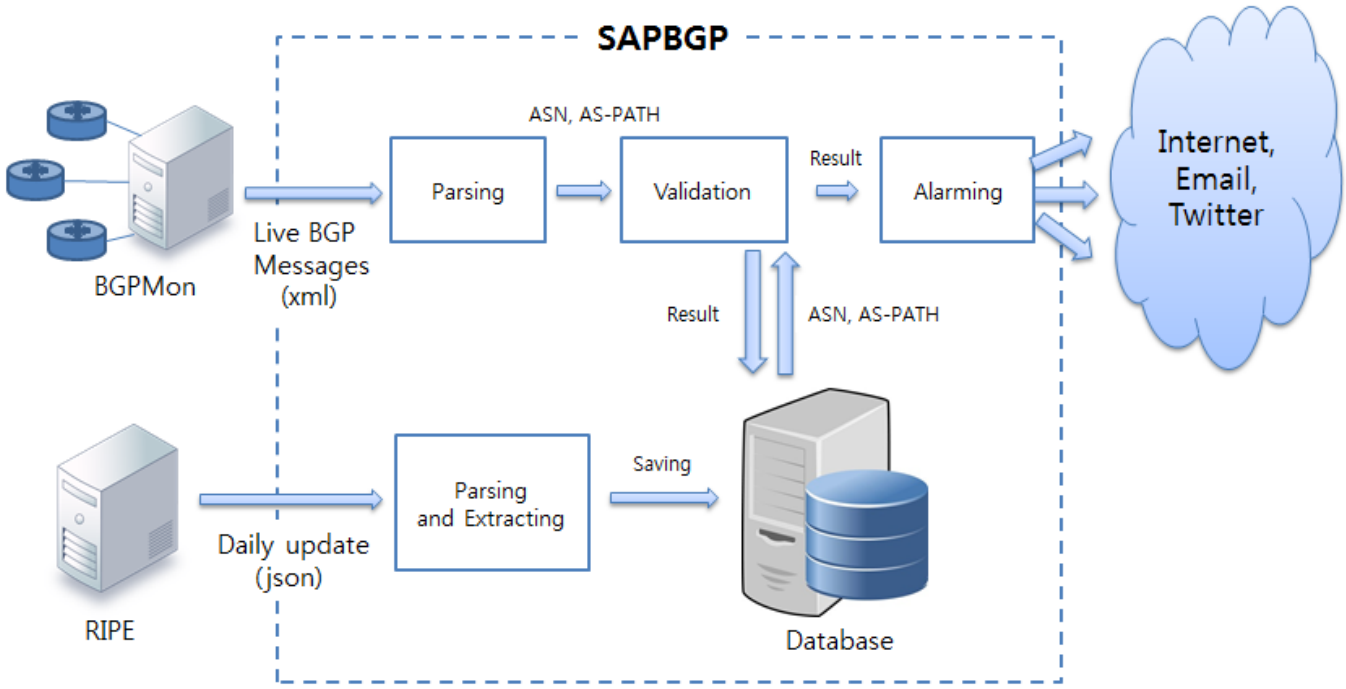

Figure 2. The architecture of the SAPBGP

C. *Routing policy based AS_PATH Validation*

RIPE NCC provides users with RIPE Data Repository that contains BGP peer information. Through this information, we can know if any ASes are connected to other ASes. This peer information has been collected by either Routing Information Service (RIS) or Internet Routing Registry (IRR). RIS has collected and stored Internet routing data from several locations all over the world since 2001.

Using peer information, the SAPBGP monitors live BGP stream from BGPmon. For example, in Figure 1, suppose that AS 400 pretends as if AS 400 is connected to AS 100, and AS 400 creates a BGP message as if the BGP message is coming from AS 100 and forwarding the BGP message. Then, AS 500 cannot check AS 400 and AS 100 are connected to each other even though the AS 500 can conduct the origin validation. However, suppose that either AS 500 or one of AS 500's neighbors is a BGPmon's participant and the SAPBGP can receive the live BGP stream related to AS 500. The AS_PATH attribute in the BGP stream should contain AS_PATH-100, 400, 500. Then, the SAPBGP can find that AS 100 and AS 400 are not connected to each other according to the peer information collected from RIPE NCC repository. As a result of this, an AS 500 administrator will be alerted by the SAPBGP and realize AS 400 might be trying the MITM attack to draw AS 500 traffic heading to AS 100.

IV. SECURE AS_PATH BGP

In this section, we introduce overall how the SASBGP works and Figure 2 describes the architecture of the SAPBGP.

A. *Constructing Database*

We construct our own database by using API provided by RIPE. We have collected, every day, all of the AS imports and exports policies information since the eighteenth of February in 2014. In addition, we have separated tables in the database to keep the daily information as well as the accumulated information by adding new exports and imports to the existing exports and imports.

As of the sixth of June in 2014, there are 77,776 ASes in the world. We sent queries to RIPE one by one. For example, if a query is related to AS 1 then the result includes AS 1's export policies, imports polices, and prefixes in the form of json. The SAPBGP parses the results so that the list of export policies and import policies can be stored to AS 1's record in the table. As a result, a new table is created every day to keep track of the daily policy information. In addition, the accumulated table is updated by adding new policies if AS 1 adds new policies against other ASes. Figure 3 shows the records from AS 10001 to AS 10005 in the policy table.

asn	export	import
10001	4680,2497,2516	
10002	2497,17224,9002,4716,251...	17225,4716,17232,45686,4732,10015
10003	4716,6939,2516,2497	4716,2516
10004	7682,4675,4732,4686,2519	7682,4732
10005		

Figure 3. A screen capture of the policy table

B. *Monitoring Live BGP Stream*

BGPmon provides live BGP stream through telnet to the public. So, whenever the routers that are connected to

 ISBN: 978-1-61208-376-6

BGPmon receives BGP update messages, BGPmon converts BGP update messages to XML format messages and propagates the XML format messages to their clients. Apart from the BGP update message, the XML format message includes timestamp, date time, BGPmon id, BGPmon sequence number, and so on.

Currently, there are 9 participants that are directly connected to BGPmon, such as AS 3043, AS 10876, AS 3257, AS 3303, AS 812, AS 5568, AS 14041, AS 28289, and AS 12145. We measured the number of update messages that BGPmon propagates for 1 hour on the twenty sixth of February in 2014. Table I shows the minimum, maximum, and average number of update messages per 10 seconds.

TABLE I. THE NUMBER OF UPDATE MESSAGES FROM BGPMON

	The number of update messages per 10 seconds
Minimum	38
Maximum	1672
Average	119.43

After parsing the live BGP message, the SAPBGP retrieves the ASN attribute and the AS_PATH attribute to check whether ASes in the AS_PATH attribute are connected to each other. Firstly, we compare the policy table in the database that is collected one day before. If we cannot find the pair, we compare the information from the accumulated table. If we cannot find the pair from the table, we consider the AS_PATH attribute as the suspicious AS_PATH attribute. If we find the suspicious AS_PATH attribute, we notify the AS network administrators of the suspicious AS_PATH attribute.

V. PERFORMACE TEST AND RESULT ANALYSIS

We explain the environment in which the SAPBGP constructs its own database by collecting RIPE repository and check the live BGP stream from BGPmon to check the invalid AS_PATH attribute in the BGP message. In addition, we conduct the performance test and analyze the result of the performance test in this section.

A. Experiment

We have constructed our database by daily collecting BGP policy records from the RIPE repository since the eighteenth of February in 2014. Based on our table, the SAPBGP checked the live BGP stream from BGPmon.

TABLE II. THE COMPARISON OF THE RESULTS

	Original results	***No duplication***
Valid	230575	13490
Invalid	3931	656
Valid by the accumulated records	4508	205

Table II shows the comparison between the original results and the result that does not contain duplications. Because of the difference of variation of BGP update periodic time, some pairs of ASes can be more duplicated than others.

Figure 4 shows the result of the AS_PATH monitoring experiment through the SAPBGP from the eighteenth of February in 2014 to the eighth of June in 2014. We conducted the experiment once a week during that period. The original data collected contains many duplicated results, but the outcome in Figure 4 does not contain the duplications. Our result shows 4.57% of the AS_PATH attribute is invalid and 95.43% of the AS_PATH attribute is valid.

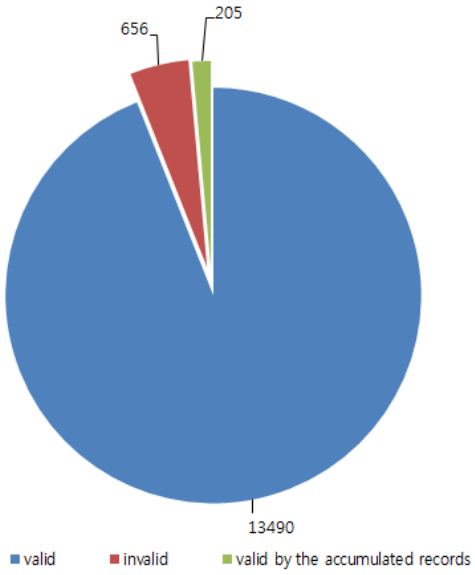

Figure 4. The result of the AS_PATH monitoring experiment

Figure 5 illustrates a portion of the policy table of the invalid ASes that the SAPBGP detected in the experiment. The invalid ASes could signify either the AS holder does not configure policies or the AS_PATH attribute was manipulated by hijackers.

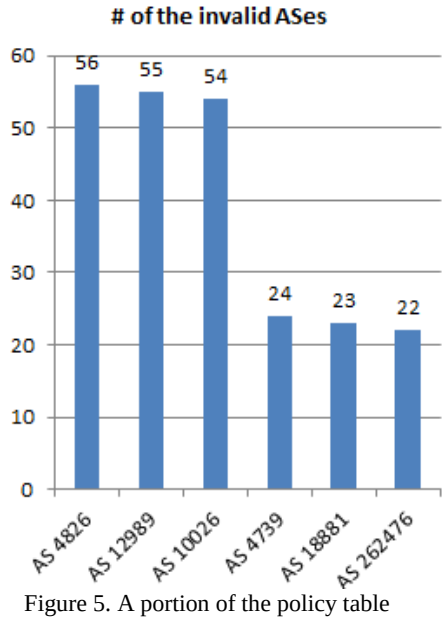

Figure 5. A portion of the policy table

B. *Performance Test*

The SAPBGP runs on a 3.40 GHz i5-3570 machine with 16 GB of memory running Windows 7. MySQL Ver. 14.14 Distrib 5.1.41 is used for the database. The SAPBGP is implemented by JAVA to collect daily updates from RIPE, to receive live BGP stream from BGPmon, and to validate the BGP stream by comparing the AS_PATH attribute to our database. The SAPBGP and database are located in the same machine to reduce the connection latency between them.

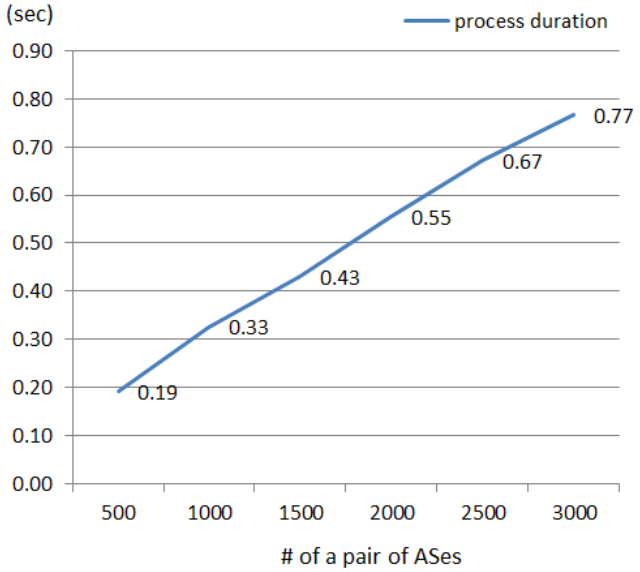

Figure 6. The result of the performance test for the AS_PATH validation

Figure 6 shows the AS_PATH validation time. The validation time includes accessing database, retrieving the specific AS record from a table, and comparing the AS_PATH attribute to the AS's record. It takes 256 microseconds, on average, to validate a pair of ASes. According to Table 1, the maximum number of live BGP messages for 10 seconds is 1672. So, the SAPBGP can process all of the live BGP messages coming from BGPmon in real time.

VI. CONCLUSION

Even though many solutions are proposed to prevent IP prefix hijacking, such as RPKI, BGPmon, Argus, and PHAS, these solutions cannot protect the AS path hijacking except RPKI. SIDR proposed the RPKI using BGPSEC but BGPSEC is currently a work in progress. In order to monitor the AS path hijacking, we propose Secure AS_PATH BGP (SAPBGP) in which we monitor the AS_PATH attribute in update messages whether each AS in the AS_PATH attribute are connected to each other based on our policy database collected from RIPE NCC repository. The result of the AS_PATH validation test shows 4.57% of the AS_PATH attribute is invalid and 95.43% of the AS_PATH attribute is valid from the fifteenth of April in 2014 to the eighth of June in 2014. In addition, the result of performance test verifies that the SAPBGP can process all of the live BGP messages coming from BGPmon in real time. In the result of the AS_PATH monitoring experiment, the ratio of invalid AS_PATH attribute is high because some AS routers still do not configure their policies. For the precise result of the policy based AS_PATH validation, every router needs to configure policies against its peers.

REFERENCES

[1] Y. Rekhter, "A Border Gateway Protocol 4 (BGP-4)," 2006, RFC 4271.

[2] S. Murphy, "BGP Security Vulnerabilities Analysis," 2006, RFC 4272.

[3] Rensys Blog, Pakistan hijacks YouTube [Online]. Available: http://www.renesys.com/blog/2008/02/pakistan_hijacks_youtube_1.shtml [Accessed February 2014].

[4] T. Manderson, L. Vegoda, and S. Kent, "Resource Public Key Infrastructure (RPKI) Objects Issued by IANA(Feb. 2012)," 2012, [Online]. Available: http://www.rfc-editor.org/rfc/rfc6491.txt [Accessed January 2014].

[5] BGPmon, Google's services redirected to Romania and Austria [Online]. Available: http://www.bgpmon.net/googles-services-redirected-to-romania-and-austria [Accessed October 2013].

[6] X. Shi, Y. Xiang, Z. Wang, X. Yin, and J. Wu., "Detecting Prefix Hijackings in the Internet with Argus", In Proc. of ACM IMC 2012.

[7] M. Lad, D. Massey, D. Pei, Y. Wu, B. Zhang, and L. Zhang, "PHAS: A prefix hijack alert system," 2006, In Proceedings of the 15th conference on USENIX Security Symposium - Volume 15 (USENIX-SS'06), Vol. 15, pp.153-166.

[8] Renesys Blog, Targeted Internet Traffic Misdirection [Online]. Available: http://www.renesys.com/2013/11/mitm-internet-hijacking [Accessed January 2014].

[9] M. Lepinski, Ed., and BBN, "BGPSEC Protocol Specification," Available: http://tools.ietf.org/html/draft-ietf-sidr-bgpsec-protocol-08.

[10] Q. Li, Y. Hu, and X. Zhang, "Even Rockets Cannot Make Pigs Fly Sustainably: Can BGP be Secured with BGPsec?," 2014.

[11] The BGPmon project, http://bgpmon.netsec.colostate.edu, [Accessed 6th July 2013].

[12] M. Lepinski, S. Kent, and D. Kong, "A Profile for Route Origin Authorizations (ROAs)," [Online]. Available: http://tools.ietf.org/html/rfc6482, [Accessed December 2012].

[13] L. Blunk, "Multi-Threaded Routing Toolkit (MRT) Routing Information Export Format," RFC 6396 , 2011.

 ISBN: 978-1-61208-376-6