The Policy-Based AS_PATH Verification to Prevent 1-Hop AS Path Hijacking By Monitoring BGP Live Streams

Je-Kuk Yun, Beomseok Hong, Yanggon Kim
Information Technology
Towson University
Towson, U.S.A.
jyun4, bhong1@students.towson.edu, and 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, the SIDR working group 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 BGP 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 1.67% of the AS_PATH attributes is invalid and 98.33% of the AS_PATH attributes is valid based on original data including duplication from the ninth of February in 2014 to the fifth of February in 2015. In addition, our results state that 94.41% of the AS_PATH attributes is invalid and 94.41% of the AS_PATH attributes is valid after removing duplicated the AS_PATH attributes. We conducted the performance test and it verified 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 a BGP router completely trusts other BGP routers. This lack of consideration of BGP vulnerabilities often causes severe failures of Internet service provision [3]. The most well-known threat of the failures is the YouTube hijacking by Pakistan Telecom (AS17557) on the 24th of February in 2008 [4]. 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) [5], BGPmon [6], Argus [7], and a Prefix Hijack Alert System (PHAS) [8].

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 [9]. In addition, The Dell SecureWorks Counter Threat Unit (CTU) research team discovered a repeated traffic hijacking to Bitcoin mining sites between February and May 2014. Compromised networks belong to Amazon, Digital Ocean, OVH, etc. The attacker hijacked cryptocurrency miners' traffic and earned an estimated $83,000 [10]. Furthermore, AS 23274, owned by China Telecom, announced approximately 50,000 prefixes, which are registered to other ASes in 2010. The reason the incident was being magnified is because China Telecom is the 11th largest Internet provider. If small ISPs hijacks a large part of the Internet, they do not have the capacity to deal with a huge amount of traffic. China Telecom, however, has the capability to operate under such traffic, and redirect its desired destination. The incident was not recognized for 18 minutes[11]. A root cause of BGP hijacking can be discovered by empirical data analysis using BGP updates from Routeviews, RIB from iPlane project, paths from traceroute, etc. However, proving a malicious intent is hardly possible. According to this research, China Telecom incident is most likely caused by a routing table leak [9].

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. However, BGPsec is currently a work in progress [11]. In addition, a study propounds that BGP armed with BGPsec cannot be secured because of BGP's fundamental design [13].

We proposed 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. We extended the previous study to conduct experiments with increased period of collecting the BGP routing policy data [1]. For the validation test with the real BGP messages, the SAPBGP receives live BGP streams from the BGPmon project [14]. 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, we introduce current active studies on BGP security in Section II. With the fact that BGP is vulnerable to MITM attack, we describe an attack scenario and a solution in Section III. In Section IV, we introduce and explain the SAPBGP in detail. We discuss the SAPBGP environment and analyze the result of the SAPBGP validation and the performance test in Section V. Lastly, we conclude the paper in Section VI.

II. RELATED RESEARCH

A. *Origin validation*

The origin validation is to verify whether the originator of update message has been authorized to announce its prefixes. In order to validate originators, the Resource Public Key Infrastructure (RPKI) is implemented by SIDR working group on January in 2013 and is currently used for origin validation. RPKI is a Public Key Infrastructure (PKI) [15], [16] where an organization called IANA manages officially verifiable Internet resources that are the allocation of hierarchy of IP addresses, Autonomous System Numbers (ASN), and signed objects for routing security. IANA is the trust anchor who allows third party to officially validate assertions according to resource allocations. The authorization is hierarchically assigned from IANA to the Regional Internet Registries (RIRs), Local Internet Registries (LIRs), National Internet Registries (NIRs), and Internet Service Providers (ISPs) as shown in Figure 1.

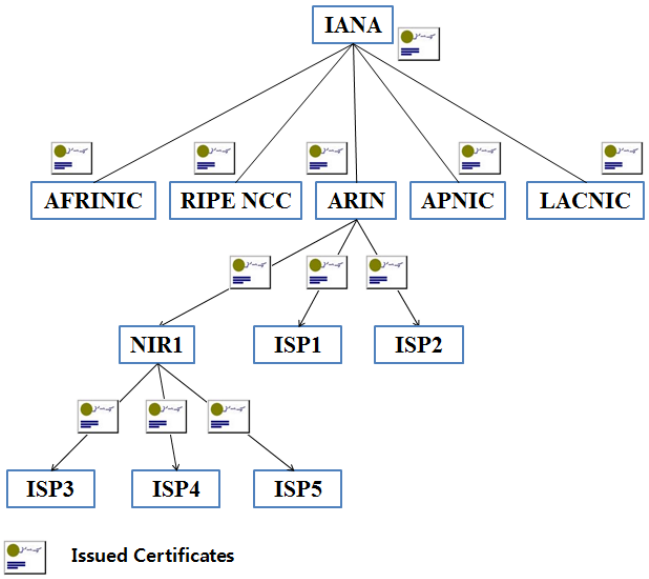

Figure 1. Hierarchy of the RPKI

There are five RIRs and they act as trust anchors like IANA. The RIR issues certificates to NIR, ISP and subscribers. NIR and ISP are allowed to issue certificates to downstream providers and to subscribers. IP address holders specify which ASes are authorized to announce their own IP address prefixes.

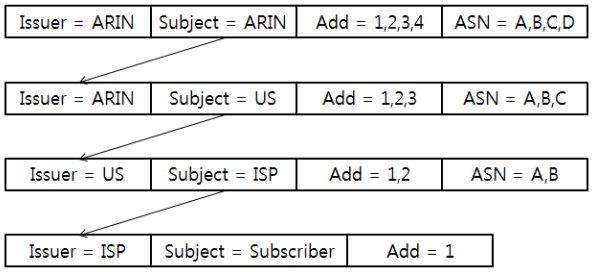

Figure 2. Certificate Chain

Figure 2 explains how a subscriber hierarchically gets certificates regarding their IP address. For example, ARIN issues certificates for US regarding addresses 1, 2, and 3 and ASN A, B, and C as shown in Figure 2. US issues certificates to ISP regarding addresses 1 and 2 and ASN A and B. Then, a subscriber can get a certificate from ISP regarding its IP addresses. As shown in Figure 3, the certificate, called Route Origin Authorizations (ROAs) [17] is a digital object formatted following the Cryptographic Message Syntax Specification (CMS) [18] and composes origin AS Number, validity date range, and one or more IP addresses with a CIDR block. If the address space holder needs to authorize multiple ASes and the IP prefixes are same, the holder should issues multiple ROAs.

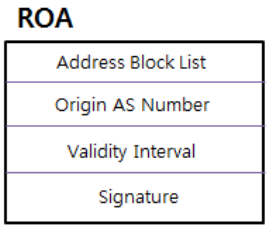

Figure 3. ROA Format

The value of Address Block List is more than one prefixes, corresponding to the NLRI that the ROA signer authorizes for prefix announcements by one or more ISPs. The value of origin AS number that is authorized to announce the prefixes indicated in the address block list. Validity interval indicates the start and end date for which the ROA is valid. Signature includes pairs of information that is used to verify the ROA. One is certificate pointer that directs its parent so that the certificate has been issued by CA. The other one is signature that is digitally signed hash data including address block list, origin as number, validity interval, hash algorithm, and digital signature algorithm. Therefore, if prefix hijacker announce other's prefixes, other network operators can check whether the announcement is invalid after comparing the IP prefixes, ASN included in the update message to the ROA

For example, as shown in Figure 4, there are five ASes. Towson University, AS 6059, announced its prefix 204.62.48.0/22. As the update message is transferred, each ASN is added to the AS_PATH attribute, and finally Verizon receives the update message and knows how to reach the prefix 204.62.48.0/22 through the AS_PATH attribute. However, if a hijacker router, AS 7922, sends the same prefix 204,62,48.0/22, then Verizon will choose AS 7922 as the final destination because the number of hops is shorter than the other as shown in Figure 4.

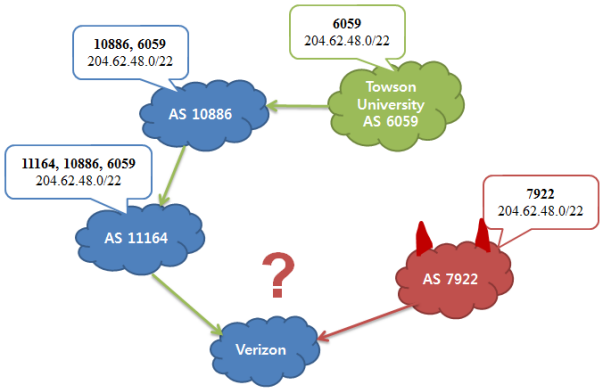

Figure 4. Scenario of IP hijacking

At this moment, if Verizon maintains ROAs and checks the ROAs then Verizon will realize that AS 7922 is not authorized to originate the prefix 204.62.48.0/22 because the ROA as shown in Table I indicates that AS 6059 has been authorized to announce the prefix 204.62.48.0/22. As a result, Verizon can choose the other route as the best path and Internet traffic will be transferred to AS 6059, which is the right destination.

TABLE I. AS 6059'S ROA

ROA
204.62.48.0/22
AS 6059

If every address spaces are authorized by its address holders, then IP prefix hijacking will be fully prevented by RPKI.

B. *BGPsec*

A SIDR working group is designing BGPsec to cryptographically prevent the AS-PATH hijacking [19]. In BGPsec, an optional and non-transitive path attribute, BGPsec_Path attribute, is included in BGP update messages. BGPsec depends on RPKI certificates and BGP router, which wants to send a BGP update messages that including the BGPsec_Path should have a private key associated with the BGP router's AS number. When the BGP router originates IP prefixes, the BGP router signs the update message with its private key so that any BGP router that receives the update message can check that the update message has been originated by the right BGP router by verifying the signature with the public key corresponding to the private key. In addition, BGP routers that receive the BGP update message sign the BGP update message with their private key and forward the BGP update message to neighbors. If every router that receives and forwards the BGP update messages signs the BGP update message, the BGP update message can be considered as the message that has not been synthesized by hijackers.

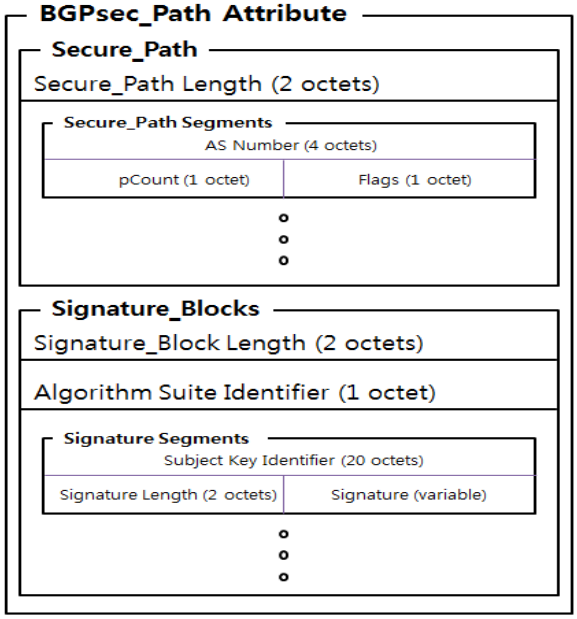

Figure 5. BGPsec_Path Attribute

In order to protect BGP update message, especially to protect AS_PATH attributes, the BGP update message should carry the secured information such as digital signature.

We call the BGP update message including a BGPsec_Path attribute BGPsec update messages as shown in Figure 5. The AS_PATH attribute in BGP update messages is replaced with BGPsec_Path attribute in the BGPsec update messages

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 ROA to ASes that are authorized to advertise a specific prefix. 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.

C. 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 [14]. Any BGP router can be a source that offers real-time update messages if the BGP router is connected to BGPmon. Currently, 9 organizations participate in the BGPmon project as a source router. In addition, BGPmon collects Multi-threaded Routing Toolkit (MRT) format [20] live streams 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.

D. 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 original purpose of the BGP policy is to filter incoming BGP messages and to choose BGP peers that will receive the BGP messages using BGP import and export policies. BGP router operators voluntarily upload their BGP policies to Internet Route Registries (IRR) through a predefined format, called Routing Policy Specification Language (RPSL) [20] that is provided by IRR. RIPE NCC database has been part of IRR and is composed of a set of online databases that is available for research purposes. In addition, RIPE NCC monitors Internet routing data and stores links between the routing data that has been seen by RIPE NCC. 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 IRR. RIS has collected and stored Internet routing data from several locations all over the world since 2001. The organizations or individuals who currently hold Internet resources are responsible for updating information in the database.

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 6 demonstrates an example of manipulating data in BGP update messages. Suppose AS 400 has a connection to AS 500 and creates a fake 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 do not have a BGP connection.

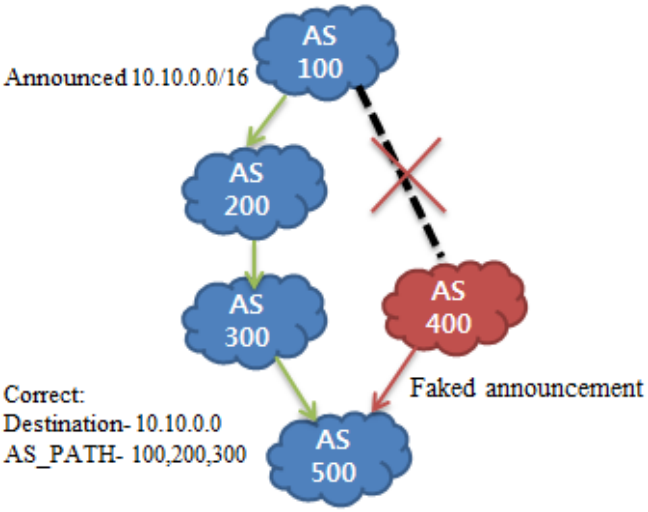

Figure 6. Manipulating a BGP message

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.

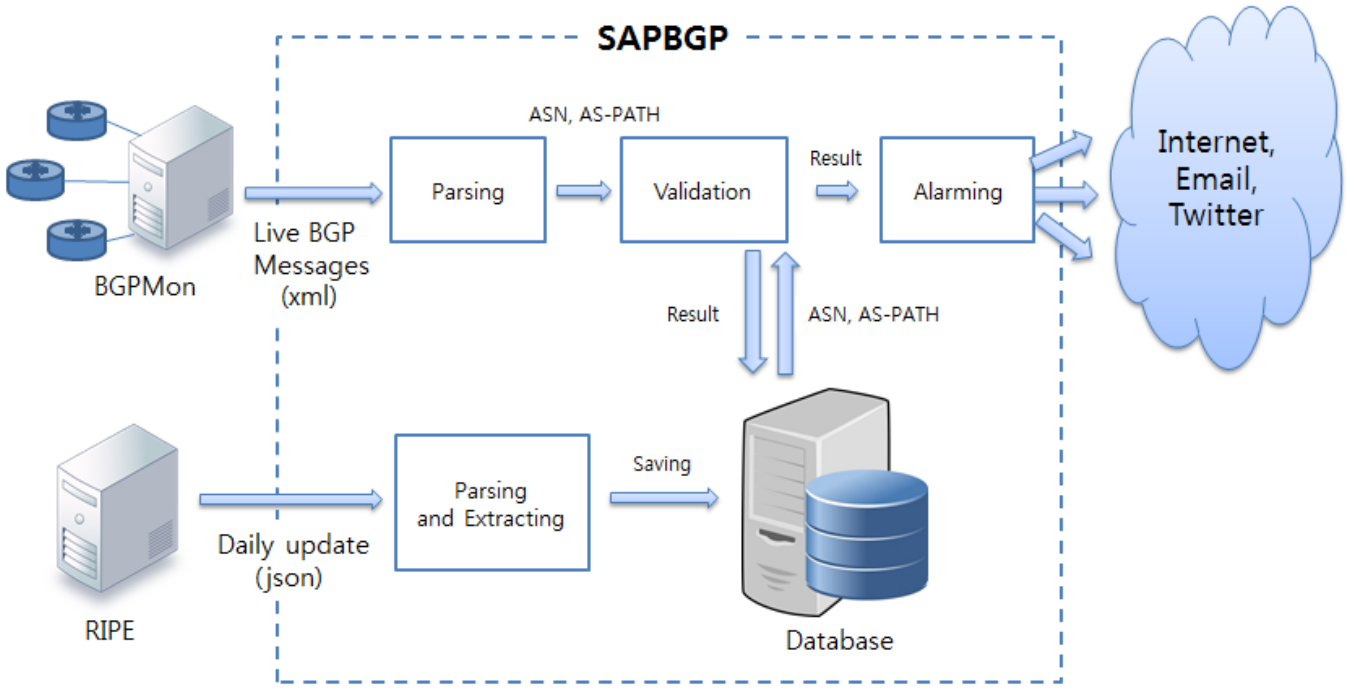

Figure 7. The architecture of the SAPBGP

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 the entire internet and they inform the targeted networks of hijacking incidents. With the support from LINX(London Internet Exchange) and other IXPs(Internet Exchange Point), they can make a more correct judgment over the hijacking. Renesys found 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 (ISPs), financial institutions, etc. [9]. Renesys detected several AS path hijacking attempts: Beltelecom (AS 6697) and Siminn (AS 6677) [9] . Victims whose traffic was diverted varied by day, and included major financial institutions, governments, and network service providers. Affected countries included the US, South Korea, Germany, the Czech Republic, Lithuania, Libya, and Iran.

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 streams from BGPmon. For example, in Figure 6, 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 SAPBGP works and Figure 7 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.

When the BGP was designed for the first time, the initial number of bits for the AS number was 16 bits, so AS number ranged from 0 to 65535. However, the number of bits for the AS number was changed to 32 bits. After that, each RIR reserves AS numbers as indicated Table II. We collected policy information from AS 1 to AS 394239 and skipped unallocated AS numbers that are not indicated in Table II.

TABLE II. 32 BITS AS NUMBER ALLOCATION ABOVE 65535

	Allocation	***The number of ASes***
APNIC	131,072-135,580	4,509
RIPE NCC	196,608-202,239	5,632
LACNIC	262,144-265,628	3,485
AFRINIC	327,680-328,703	1,024
ARIN	393,216-394,239	1,024

We sent queries to RIPE NCC 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 8 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 8. A screen capture of the policy table

B. Monitoring Live BGP Stream

BGPmon provides live BGP streams through telnet to the public. So, whenever the routers that are connected to 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 as shown in Table III.

TABLE III. 9 ORGANIZATIONS THAT PARTICIPANTE IN THE BGPMON PROJECT

AS number	***Organization name***
812	Rogers Cable Communication Inc.
3303	Swisscom (Switzerland) Ltd
3257	Tinet SpA (RIPE NCC)
5568	ROSNIIROS Russian Institute for Public Networks
6447	University of Oregon
10876	MAOZ.com
14041	University Corporation for Atmospheric Research
12145	Colorado State University
28289	Americana Digital Ltda.

We measured the number of update messages that BGPmon propagates for 1 hour on the twenty sixth of February in 2014. Table III shows the minimum, maximum, and average number of update messages per 10 seconds.

TABLE IV. THE NUMBER OF UPDATE MESSAGES FROM BGPMON

	The number of update messages per 10 seconds
Minimum	38
Maximum	1,672
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

In order to monitor AS path hijacking in the real world, we collected BGP live stream from the BGPmon project and compared the AS_PATH attribute to our policy-based database. The policy-based database is updated daily because BGP policy information changed whenever network operators wanted to change their BGP policies. A new BGP policy table is created every day, so we used the BGP policy table that is collected one day before the day we conducted the experiment. The number of BGP routing policies that are registered by AS holders is 55,395 on February in 2015, which means only 68% of AS holders registered their BGP routing policies as shown in Figure 9.

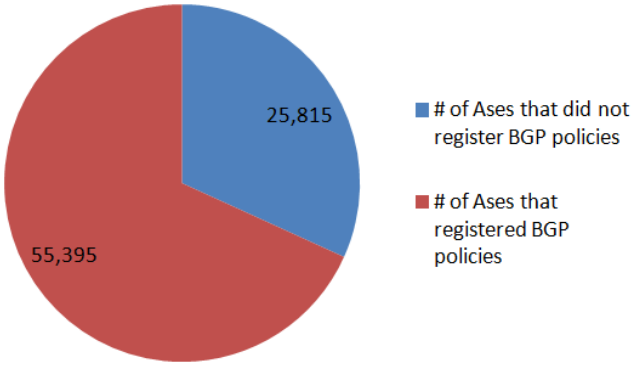

Figure 9. Ratio of ASes that registered BGP routing policies

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 streams from BGPmon.

TABLE V. THE COMPARISON OF THE RESULTS

	Duplication included	*No duplication*
Valid	1,950,904	83,636
Invalid	34,938	5,271
Valid by the accumulated records	107,795	5,463

Table V 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 duplicated more than others.

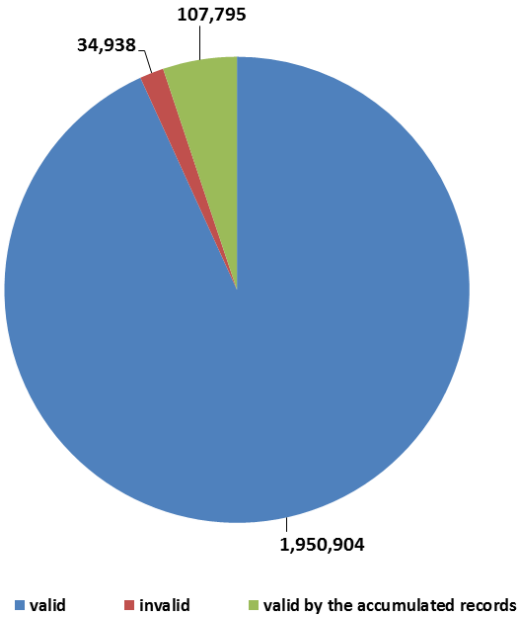

Figure 10. The result of the AS_PATH monitoring experiment that includes duplications

Figure 10 shows the result of the AS_PATH monitoring experiment through the SAPBGP from the ninth of February in 2014 to the fifth of February in 2015. We conducted the experiment randomly twice a month during that period. Figure 10 shows the original data that contains many duplicated results. Our result indicates 1.67% of the AS_PATH attributes are invalid and 98.33% of the AS_PATH attributes is valid.

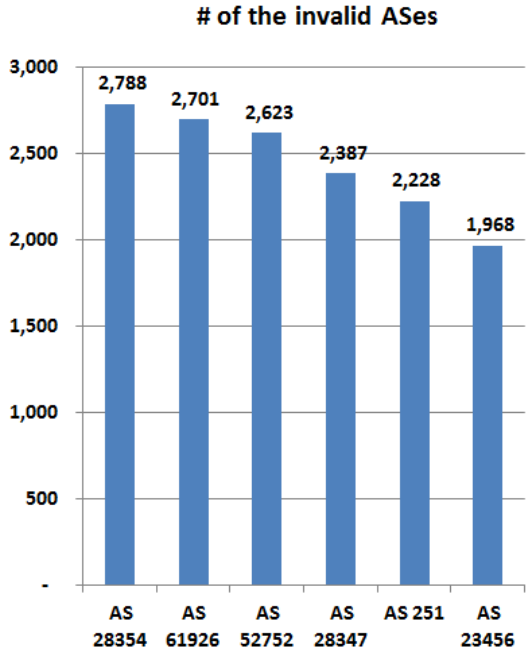

Figure 11. A portion of the policy table of the invalid ASes that includes duplications

Figure 11 illustrates a portion of the policy table of the invalid ASes that the SAPBGP detected in the experiment and this result contains duplications. The invalid ASes could signify either the AS holder does not configure policies or the AS_PATH attribute was manipulated by hijackers.

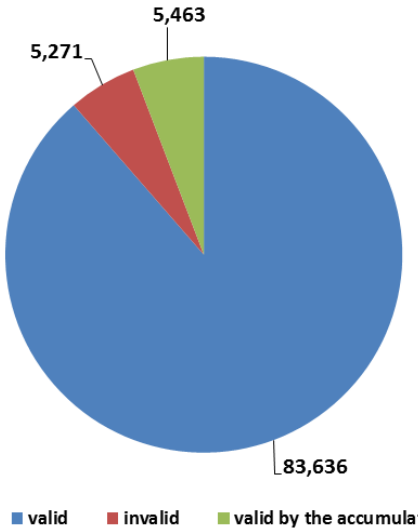

Figure 12. The result of the AS_PATH monitoring experiment that includes duplications

Since original data contains many duplicated information, we analyzed the result that does not contain duplications as well. Figure 12 shows the result of AS_PATH that does not contain the duplications. Our result shows 5.57% of the AS_PATH attribute are invalid and 95.43% of the AS_PATH attribute are valid.

Figure 13 illustrates a portion of the policy table of the invalid ASes that the SAPBGP detected in the experiment. The result does not contain duplications from the original results.

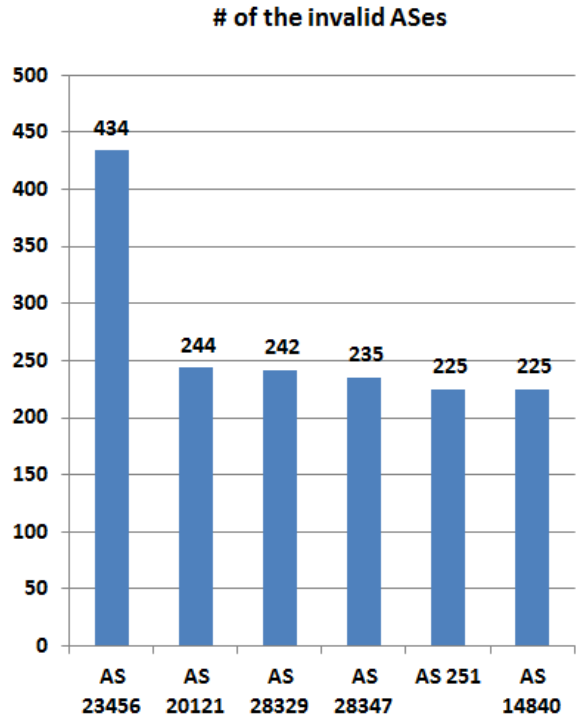

Figure 13. # of ASes that registered BGP policies that do not include duplications

The number of ASes that registered BGP routing policies are gradually increased according to our policy database. The total number of ASes is 81,210 and it will take a long time for every AS holder to register BGP policies. Figure 14 shows how many of ASes that registered BGP policies is increased for 1 year between March in 2014 and February in 2015. In order to check connections between two peers, BGP policy information from each BGP peer should contain the BGP policy against the other peer. However, we considered a BGP connection is valid if only one of two BGP peers has the BGP policy against the other peer because the number of ASes that registered BGP policy is still small. In addition, we considered a BGP message as valid message if one of an AS_PATH pair is the one of 9 organizations that participate in the BGPmon project.

We assumed that a pair of AS_PATH that is invalid and is placed at the second position in the AS_PATH attribute can be candidates of 1-hop hijacker because the number of hops should be shorter than others to draw Internet traffic to their AS. Since the first position is the destination AS, the second position AS can hijack Internet traffic heading for the first position AS.

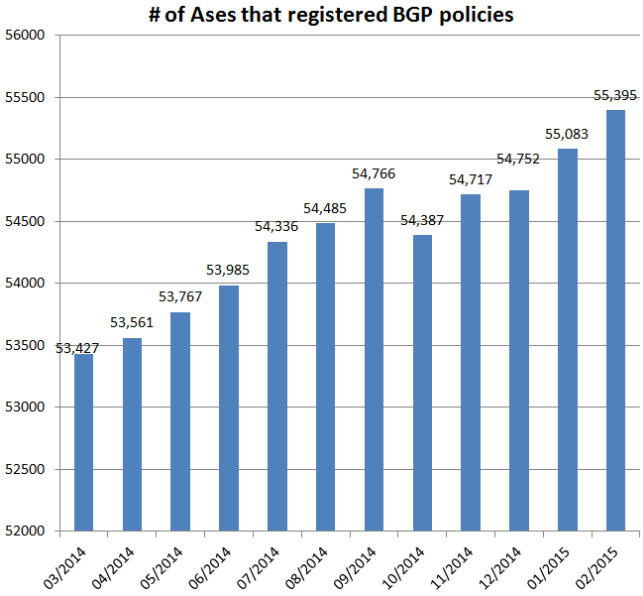

Figure 14. # of ASes that registered BGP policies

Table VI enumerates the top 20 1-hop hijacking candidates.

TABLE VI. TOP 20 1-HOP HIJACKING CANDIDATES

First position	*Second position*	*Frequency*
AS 4739	AS 3491	12
AS 4739	AS 1239	12
AS 4739	AS 1273	12
AS 4739	AS 1299	12
AS 3491	AS 7575	12
AS 4739	AS 209	12
AS 10026	AS 3491	11
AS 10026	AS 1273	11
AS 4739	AS 24115	11
AS 4739	AS 9488	11
AS 53237	AS 12956	11
AS 7575	AS 24490	11
AS 4739	AS 2914	11
AS 4826	AS 2828	11
AS 4739	AS 4635	10
AS 38809	AS 2914	10
AS 4826	AS 9498	10
AS 4739	AS 10026	10
AS 10026	AS 1299	10
AS 53237	AS 3549	10

We checked 94,370 invalid pairs of AS_PATH attributes that do not include duplications and we considered 1-hop hijacking candidates if the pair is located at first and second positions in the AS_PATH attribute.

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 streams 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.

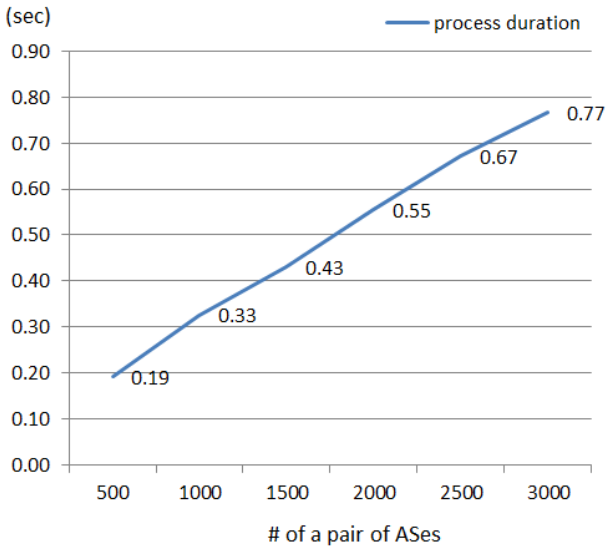

Figure 15. The result of the performance test for the AS_PATH validation

Figure 15 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. We conducted performance test for around 1,864,567 live BGP streams. As shown in Table VI, it takes 4.12 ms, on average, to validate a pair of ASes.

TABLE VII. AS_PATH VALIDATION TIME TO PROCESS ONE BGP UPDATE MESSAGE

	Duration for verifying a BGP message
Minimum	0.07ms
Maximum	9.86sec
Average	4.12ms

According to Table IV, the maximum number of live BGP messages for 10 seconds is 1,672. The SAPBGP can process 2,427.18 BGP messages for 10 seconds, on the average, based on the performance test as shown in Table VII. 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. Our analysis shows 1.67% of the AS_PATH attributes is invalid and 98.33% of the AS_PATH attributes is valid based on original data including duplication from the ninth of February in 2014 to the fifth of February in 2015. In addition, our results state that 94.41% of the AS_PATH attributes is invalid and 94.41% of the AS_PATH attributes is valid after removing duplicated the AS_PATH attributes. In addition, the result of the 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] J. Yun, B. Hong, and Y. Kim, "The Policy-Based AS_PATH Verification to Monitor AS Path Hijacking," The Eighth International Conference on Emerging Security Information, Systems and Technologies (SECURWARE 2014), Lisbon, Portugal, 16-20, November, 2014, pp.20-24.

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

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

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

[5] 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].

[6] 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].

[7] 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.

[8] 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.

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

[10] P. Litke and J. Steward,"BGP Hijacking for Cryptocurrency Profit" [Online]. Available: http://www.secureworks.com/cyber-threat-intelligence/threats/bgp-hijacking-for-cryptocurrency-profit [Accessed February 2015]

[11] J. Cowie. Renesys blog: China's 18-minute mystery [Online]. Available: http://www.renesys.com/blog/2010/11/chinas-18-minute-mystery.shtml[Accessed January 2015].

[12] M. Lepinski, Ed., and BBN, "BGPSEC Protocol Specification," Available: http://tools.ietf.org/html/draft-ietf-sidr-bgpsec-protocol-11, [Accessed January 2015].

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

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

[15] R. Housley, W. Ford, W. Polk, and D. Solo, Internet X.509 Public Key Infrastructure Certificate and CRL Profile, RFC2459, 1999.

[16] D. Cooper, S. Santesson, S. Farrell, S. Boeyen, R. Housley, and W. Polk, Internet X.509 Public Key Infrastructure Certificate, and Certificate Revocation List (CRL) Profile, RFC5280, 2008 .

[17] 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].

[18] R. Housley, "Cryptographic Message Syntax (CMS)," [Online]. Available: http://www.ietf.org/rfc/rfc3852.txt, [Accessed September 2014].

[19] IETF. "Secure Inter-Domain Routing (SIDR)". Online, Sep. 2010. Available from http://datatracker.ietf.org/wg/sidr/

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

[21] C. Alaettinoglu, C. Villamizar, E. Gerich, D. Kessens, D. Meyer, T. Bates, D. Karrenberg, and M. Terpstra, "Routing Policy Specification Language (RPSL)," RFC 2622, June 1999.