Securing BGP Networks using Consistent Check Algorithm

C. K. Man, K.Y. Wong, and K. H. Yeung

Abstract—**The Border Gateway Protocol (BGP) is the critical routing protocol in the Internet infrastructure. However, there is no security concern in the original design of BGP, which suffers from various kinds of threats for attacks. To secure the BGP operation, this paper proposes an algorithm called consistent check. The algorithm is to verify the correctness of AS path in an incoming BGP update message by consulting the knowledge of other autonomous systems in the network. Unlike existing solution, this proposed algorithm does not require the need of cryptography calculation.**

I. Introduction

A. BGP basic

The Border Gateway Protocol (BGP) [1][2] is a path vector routing protocol. The basic idea is that each router exchanges network reachability information with its directly connected neighbors. BGP treats each autonomous system as a single point on the path to any given destination.

In BGP, a set of address (or an IP clock) (e.g., 1.2.3.0/24) that is being routed is called a *prefix*, and the list of autonomous system (AS) that the packet must pass through to reach the prefix (e.g., {20 30 40}) is called *AS path*. A BGP message contains the information about the AS path for a given prefix. For simplicity, we use something like [{20, 30, 40} 1.2.3.0/24] to represent the information.

The operation of BGP can be briefly described as follows. Initially, each router originates the prefixes it can reaches in addition to its own AS number, and advertises this information to its neighbors using UPDATE messages.

When a router receives an advertisement, it adds new routes to its own local routing table based on the contents of the advertisement. On the other hand, it adds itself to the AS path before advertising this information to the next router.

Manuscript received November 17, 2008. The work described in this paper is supported by Hong Kong Government General Research Fund numbered CityU-123608.

C. K. Man is with the Department of Electronic Engineering, City University of Hong Kong (phone: +852-27889865; fax: +852-27887791; e-mail: 50584642@ee.cityu.edu.hk).

K. Y. Wong is with the Computer Studies Program, Macao Polytechnic Institute (e-mail: kywong@ipm.edu.mo).

K. H. Yeung is with the Department of Electronic Engineering, City University of Hong Kong (e-mail: eeayeung@cityu.edu.hk).

When a router finds that its AS number has already been included in the AS path, it rejects the route information in order to prevent the formulation of a routing loop. However, BGP can only detect loops among autonomous systems and it cannot ensure there is no loop within an AS.

In an update message, apart from prefixes, BGP attributes are included to provide additional information about those prefixes, such as path preference and aggregation information.

B. BGP Threats

There are three primary security vulnerabilities in BGP: 1) data integrity, 2) origin authentication and 3) path validation.

Firstly, data integrity is not provided in BGP routing. Therefore, the received information has a risk that it is modified during data transfer. On the other hand, when a BGP router receives a prefix advertisement, it only considers the AS path length but not validates ownership of the prefix. Finally, BGP routers do not validate the correctness of the received AS path. They simply check the length of the path to make routing decision (the shorter one will be used). Without path validation, AS path could be modified in the middle or maliciously created, which can redirect traffic to an unexpected direction.

As a result, the BGP routing networks suffer from various kinds of attacks, achieving one or more objectives. The attack objectives include:

Black Hole refers to the situation that packets go to a router or network but do not come out.

Traffic Redirection redirects packets, using the path that is not supposed to use, to an incorrect destination that can drop or modify the packets..

Traffic Subversion is a special case of traffic redirection. In traffic subversion, after eavesdropping or modifying the redirected packets, the router will forward the packer to the actual destination eventually.

Instability refers to the situation that the routing structure changes frequently or the routing nodes have inconsistent views of the network, causing a very long convergence delay.

There are various kinds of attacks [3][4][5] to achieve the objectives. Fig. 1 shows a simple example of the prefix hijacking attack. Considering the network shown in Fig. 1, in normal operation, AS 6 advertises [{6 5 1} 172.1.1.0/24] and AS 4 advertises [{4 3 2 1} 172.1.1.0/24] to the Internet. Since the advertisement of AS6 shows a shorter AS path

than that of AS4, the packet destined for 172.1.1.0/24 from the Internet will be sent to AS 6. However, after being compromised by an attacker, AS 4 advertises to the Internet a fake AS path {4, 1}. As this fake path is shorter than the correct one from AS 6, the Internet forwards the packets destined for the prefix A to AS 4 (not AS 6). Now, AS 4 has hijacked the prefix A.

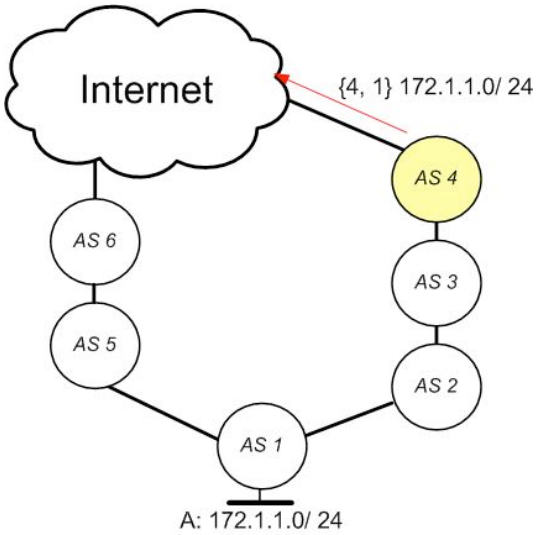

Fig. 1 Prefix Hijacking Attack

II. BGP EXISTING SOLUTIONS

There are a number of existing solutions to tackle the vulnerabilities of BGP. In this section, the three most commonly known will be discussed. SBGP, SoBGP and IRV will be discussed.

A. A. Secure BGP (S-BGP)

Secure BGP (S-BGP) is an extended and enhanced version of BGP, which uses strong security algorithm to tackle BGP security vulnerabilities. Four major additions are introduced in S-BGP. They are Public Key Infrastructure (PKI), Address Attestations (AA), Route Attestations (RA), and Internet Protocol Security (IPSec). The detail operations of them are referred to [6], [7].

AA is a digital signature issued by a trusted certificate authority, which is used to provide origin authentication to BGP. That is, by using AA, the prefix ownership for a particular AS can be verified.

RA is to provide path validation to BGP. It is used to authorize a neighbor AS to advertise the specified prefix.

On the other hand, IPSec is used to provide data integrity. That is to ensure all the data exchange is protected between routers.

B. Secure Origin BGP (soBGP)

soBGP [8] is another solution to secure the BGP network, which highly depends on the use of digital certificates. In soBGP, three certificates are introduced, and they are called EntityCert, AuthCert and PolicyCert. All the certificates are signed by the private key of the distributing AS.

EntityCert and AuthCert are used to achieve the origin authentication. EntityCert is used for validating the identity of an AS, which contains public key for the given AS. AuthCert is used to verify prefix ownership. It contains the advertised prefix information, the authorizing AS (itself) and the authorized AS (neighbor).

On the other hand, PolicyCert is used for path validation. In a PolicyCert, the peers of an AS are included.

C. Interdomain Routing Validation (IRV)

Unlike S-BGP and soBGP that use digital certificates, IRV [9] uses a new type of application-level server, called IRV server, to secure the BGP network. It requires at least one IRV server in each AS.

When a BGP receives an update message and finds it suspicious, it will ask the IRV server in its own AS to validate the message. The IRV server will in turns ask another IRV server in the predecessor AS, which is listed in the AS path of that update message. Eventually, the origin AS of the prefix will be consulted.

Therefore, path authentication is achieved as all corresponding IRV server listed in the AS path are queried. Origin authentication is also achieved as the IRV server in the origin AS is queried.

On the other hand, as IRV servers use IPSec or TLS to transfer the queries and response on a secure layer, data integrity can be achieved as well.

An important design issue in the IRV architecture is the way to locate the IRV server corresponding to a particular AS. One way to do that is the establishment of a well-known registry that records the location information (e.g., IP addresses) for all ASs.

D. Problems of existing solutions

Existing solutions for securing BGP commonly require the implementation of a new secure protocol, which is difficult to be widely deployed in practice. For example, the successful deployment of S-BGP and soBGP requires the presence of PKI and a common certificate authority (CA) which is trusted by all participating routers. As the new protocols heavily use the digital signature, it requires them to have much higher computation and memory requirement than the current BGP.

Though IRV does not modify the current routing protocols, it requires additional IRV servers to store prefix information. Most importantly, it has to build another network that interconnects all IRV servers for ASs all over the world. Since the authentication of a routing update takes long latency (by consulting a number of IRV servers), as stated in the IRV proposal [9], it cannot be used for authenticating all routing updates; it should only authenticate updates at a random basis. Finally, the maintenance of the IRV servers is one of the deployment issues.

III. Consistent Check Algorithm

In this paper, Consistent Check (CC) algorithm is proposed to tackle the BGP vulnerabilities. The CC algorithm does not require the modification of existing routing protocols, and does not use cryptography calculation, making it both deployment and resource friendly. The algorithm is based on the common views from the ASs in the network to verify the validation of a routing update. That is, if one receives an update about the path P from AS X, it will do some calculation to see if other ASs (except AS X) in the network have the same information about P as that just received from AS X.

There are several steps to perform CC:

1. Each AS constructs a CC table based on its received BGP routing update messages.
2. The consistence of a received path can be checked by the CC table.
3. If conflict exists, the CC algorithm launches another process to find out the actual path.

The consistence check in step 2 is based on the predecessor information stored in the CC table. Predecessor refers to the second-to-last hop (predecessor) in the path to a destination. As can be seen in step 3, the CC algorithm not only provides the validation of a path to a destination, but also be able to identify the actual connection of that path.

A. AS-Path Pool and CC Table

Fig. 2 (a) shows an example of BGP network connecting nine autonomous systems. Considering the network A in the example, as BGP routers advertise their path to network A by BGP update messages, after receiving the messages from AS 4 and AS 6, AS 5 will have the AS-path pool of network A (see Fig. 1 (b)). In this case, there are two paths for network A, and the one [4, 3, 2, 1] is selected (indicated by ">") as it is shorter.

In the CC algorithm, CC tables have to be built, which are based on the AS-path pools. A CC table contains three components:

Destination: It indicates the destination ASs in the network.
Length: It indicates the number of hop from the origin AS to the destination ASs. For example, in Fig. 2(c), AS 7 is two hops away from the origin AS (AS 5).
Predecessor: It indicates the predecessor of the destination AS in the path to the origin AS. For example, look at the network topology in Fig. 2(a), the predecessor of AS 7 is AS 6 in the path to the origin AS (AS 5).

The following describe how AS 5 constructs the CC table for network A, as shown in Fig. 2(c). First, since AS 5 is the owner of the table, on the entry of AS 5, the length and predecessor values are zero and itself respectively. Second, since ASs 4 and 6 are the neighbors of AS 5, the length from them to AS 5 is 1, and their predecessor is AS 5. Third, base on the selected path in the AS-path pool of network A (see Fig. 1(b)), ASs 3 and 4 are in the middle of the path from AS 2 to AS 5. Therefore, for the entry of AS 2, the length is 3, and the predecessor is AS 3; whereas for the entry of AS 3, the length is 2 and the predecessor is AS 4. Note that the selected (best) path is used to calculate the length value, even though there is more than one path available. Similarly, based on another AS path {6, 7, 8, 9, 1} shown in the AS-path pool of network A, the entries of ASs 7, 8, and 9 in the CC table can be filled.

By using the CC tables from other ASs, AS 5 is able to check the correctness of a received routing update. During the consistency check, path will be traced back to the each destination in the CC tables. If conflict exists, AS 5 can discard the update and label the advertiser as suspicious. After that, AS 5 can find out the actual path to the claimed prefix in that update, by checking more CC tables from other ASs.

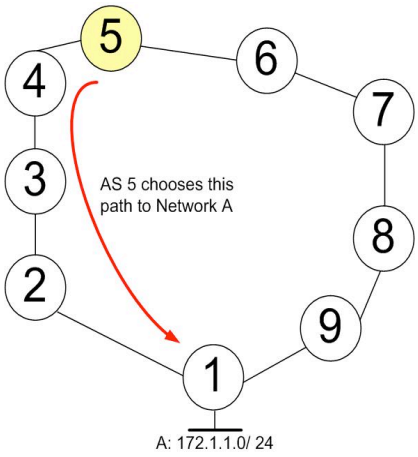

(a) The BGP network.

Network	Path	Next Hop	
172.1.1.0	>4 3 2 1	172.4.1.1	*(from AS 4)*
172.1.1.0	6 7 8 9 1	172.3.3.1	*(from AS 6)*

(b) The AS-path pool of network A

Destination	Length	Predecessor
5	0	5
6	1	5
4	1	5
3	2	4
7	2	6
2	3	3
8	3	7
1	4	2
9	4	8
A	5	1

(c) The CC table for prefix A.

Fig. 2. An example of AS-path and CC table in the CC algorithm.

B. Consistent check for the claimed AS path

After receiving a new BGP, AS 5 can validate the AS path claimed in that update by checking the corresponding CC table.

Suppose that, in the network shown in Fig. 2(a), AS 6 has a malicious router that advertises to AS 5 a BGP update containing a fake AS path {6, 9, 1} to the prefix A, 172.1.1.0/24. See Fig. 3.

If the CC algorithm is not used, as this fake path is shorter than the one {4, 3, 2, 1} AS 5 already knows, AS 5 will regard this fake path is better and put it in the routing table. As a result, AS 5 will send all packets destined for the prefix A to AS 6 (not AS 4 that is the correct next hop). Now, AS 6 can perform various kinds of attack, such as black hole or hijacking, on these packets that it is not supposed to receive.

With the CC algorithm, when AS 5 receives the claimed AS path to the prefix A from AS 6, it asks for the "network A CC table" from other involved ASs to validate the claimed AS path in the update.

In this case, since the claimed path is {6, 9, 1}, excluding the announcer, other involved ASs are AS 1 and AS 9. Suppose that AS 5 first check the AS 1's CC table, as shown in Fig. 4. By tracing the destinations in the table, AS 5 finds that AS 1 is the neighbor of AS 9. This relationship (9-1) matches part of the claimed AS path (6-9-1). However, as AS 1's CC table does not contain any information of AS6, AS 5 checks AS 9' CC table. After tracing the destinations in the table (see Fig. 5), AS 6 finds that only AS 1 and AS 8 are directly connected to AS 9. It is inconsistent to what the AS path claimed in the BGP update (which claims AS 6 is directly to AS 9). Therefore, AS 5 identifies that this update is suspicious.

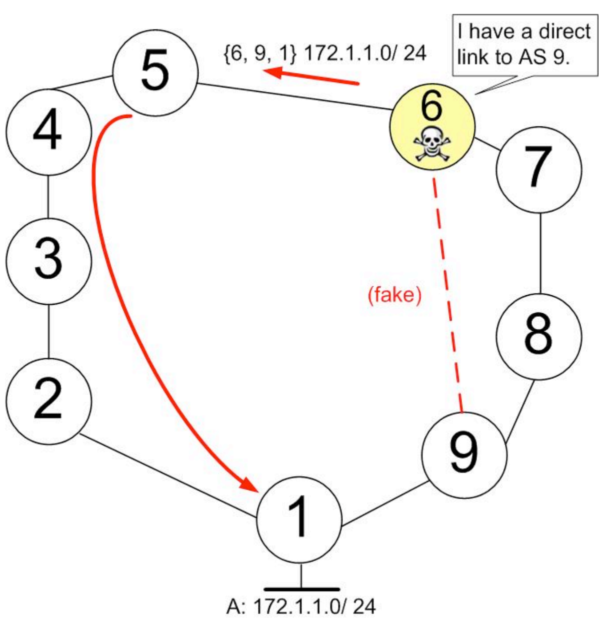

Fig. 3. The example of attack.

Destination	Length	Predecessor
1	0	1
2	1	1
9	1	1
A	1	1

Destination	Length	Predecessor
1	0	1
2	1	1
9	1	1
A	1	1

Fig. 4. Trace back process from AS 9 to AS 1 in AS 1's CC table.

Destination	Length	Predecessor
9	0	9
1	1	9
8	1	9
7	2	8
6	3	7
5	4	6
4	5	5
3	6	4
2	7	3
A	2	1

Destination	Length	Predecessor
9	0	9
1	1	9
8	1	9
7	2	8
6	3	7
5	4	6
4	5	5
3	6	4
2	7	3
A	2	1

Destination	Length	Predecessor
9	0	9
1	1	9
8	1	9
7	2	8
6	3	7
5	4	6
4	5	5
3	6	4
2	7	3
A	2	1

Destination	Length	Predecessor
9	0	9
1	1	9
8	1	9
7	2	8
6	3	7
5	4	6
4	5	5
3	6	4
2	7	3
A	2	1

Fig. 5. Trace back process from AS 6 to AS 9 in AS 9's CC table.

C. Find the correct AS path

The purpose of last step is to check the consistency of a claimed AS path. However, if inconsistency occurs, it is only able to identify that the update is suspicious, but not able to figure out the correct AS path.

If a path to a given prefix is correct, the ASs in the network should have the consistent knowledge about that path. Therefore, by checking more CC tables from other ASs, it is able to figure out the correct AS path which is common to other ASs in the network.

Consider the attack example in Fig. 3. Since the received AS path is {6, 9, 1}, AS 5 will check the CC tables from all the ASs listed in the path except the malicious sender. Therefore, to figure out the correct AS path to the prefix A, AS 9 and AS 1 will be checked. Suppose that it firstly checks the CC table from AS 9. By tracing the table, as shown in Fig. 5, AS 5 finds that the path from AS 6 to AS 9 is "6 - 7 - 8 - 9." After that, AS 5 traces the CC table of AS 1. It then finds that the path from AS 9 to AS 1 is "9 - 1." Now AS 5 has two following two partial paths:

1. From AS 9: 6 - 7 - 8 - 9
2. From AS 1: 9 - 1 - A

By merging these two paths, AS 5 can figure out the complete path from itself to the prefix A should be "5 - 6 - 7 - 8 - 9 - 1."

D. Discussions

Since the number of normal routers is more than that of malicious routers in the network, by checking the knowledge about the path from the majority of ASs in the network, the proposed CC algorithm is able to figure out the correct path.

As the CC algorithm checks the CC tables from the ASs listed in the AS path. The longer the AS path is, the more reliable the CC algorithm is. Therefore, if a malicious router

claims it has a direct link to a prefix, the CC algorithm has difficulty to figure out the actual path. To do that, CC has to check other ASs that are in the network but not listed in the AS.

The CC algorithm can provide path and prefix validation. Though data integrity is not introduced provided, it can be achieved if a secure transport layer, such as Secure Sockets Layer (SSL) is used to transfer the CC tables.

IV. Conclusion

In this paper, the threats and attack objectives are pointed out. The idea and operation of the proposed Consistent Check (CC) algorithm are then described. CC provides an effective way for tackling three primary security vulnerabilities of BGP. The major merit of CC is that, unlike existing solutions, it does not require the need of cryptography calculation, or the need to modify the existing routing protocol. Therefore, it can be easily deployed in the existing networks.

Reference:

[1] Y. Rekhter, T. Li, and S. Hares, "A Border Gateway Protocol 4 (BGP-4)," RFC 4271, January 2006.

[2] Russ White, Danny McPherson and Shrihari Sangli, "Practical BGP", Addison-Wesley 2005

[3] Ola Nordstorm and Constantinos Dovrolis, "Beware of BGP Attacks" *ACM SIGCOMM Computer Communication Review*, Vol. 34, pp. 1 - 8, Apr. 2004.

[4] Meiyuan Zhao, Sean W. Smith and David M. Nicol, "The Performance Impact of BGP Security" *IEEE Network*, Vol. 19, pp.42 – 48, Nov. 2005.

[5] Hitesh Ballani, Paul Francis, and Xinyang Zhang, "A study of prefix hijacking and interception in the internet," *ACM SIGCOMM Computer Communication Review*, Vol. 37, pp. 265 – 276, 4 - Oct. 2007

[6] Stephen T. Kent, Charles Lynn and Karen Seo, "Secure Border Gateway Protocol (S-BGP)" *IEEE Journal On Selected Areas In Communications*, Vol. 18, No. 4, pp. 582 – 592, Apr. 2000

[7] Bradley R. Smith and J. J. Garcia-Luna-Aceves, "Securing the Border Gateway Routing Protocol" *Global Telecommunications Conference, 1996. GLOBECOM '96. 'Communications: The Key to Global Prosperity*, pp. 81-85, Nov 1996

[8] Russ White, "Securing BGP Through Secure Origin BGP" *The Internet Protocol Journal*, Vol. 6, Number 3, pp. 15 – 22, Sep. 2003.

[9] K. Butler, T. Farley, P. McDaniel, and J. Rexford. A Survey of BGP Security Issues and Solutions. Technical Report TD-5UGJ33, AT&T Labs - Research, Florham Park, NJ, Feb. 2004.