BGP with BGPsec: Attacks and Countermeasures

Qi Li, Jiajia Liu, Yih-Chun Hu, Mingwei Xu, Jianping Wu

Abstract

The BGP suffers from numerous security vulnerabilities, for example, fake routing updates incurring traffic hijacking and interception. The BGPsec protocol is supposed to fix these vulnerabilities by attesting routing updates. Although the BGP security problem has been extensively studied, the security of BGP with BGPsec is not well studied yet. We argue that even secured with BGPsec, BGP still has inherent security vulnerabilities. In particular, traffic can still be hijacked. In this article, we systematically study the vulnerabilities of BGP with BGPsec. We find that the protocol still cannot achieve the desired security guarantee of inter-domain routing. In particular, it is unable to ensure correct packet delivery on the Internet. We measure the impacts of the vulnerabilities by using a real data trace, and discuss enhancements to the design and the implementation of the secure BGP protocol, which allows BGP to achieve strong secure inter-domain routing.

Introduction

The Border Gateway Protocol (BGP) is the de-facto inter-domain routing protocol that enables Internet connectivity for various autonomous systems (ASes), that is, networks operated by different organizations. It has been deployed since the Internet was built, and the current version 4 is widely used on the Internet. The protocol in each AS exchanges reachability information with neighbors and selects one of many paths that it learned from the neighbors to transmit packets. In general, it works well on the current Internet in ensuring the connectivity of the global Internet, in spite of the inability to provide any security mechanisms for connectivity.

Since BGP does not have built-in security mechanisms to verify if a route update generated by an AS is genuine, it suffers from serious security vulnerabilities. Therefore, any AS (or any BGP router) can announce any arbitrary route, that is, any routing path, via BGP. For instance, on February 24, 2008, Pakistan Telecom (AS17557) announced an unauthorized routing path for prefix 208.65.153.0/24, and PCCW Global (AS3491), that is, Pakistan Telecom's provider forwarded this fake announcement to the rest of the Internet, resulting in hijacking YouTube traffic on a global scale for over two hours. Many similar traffic hijacking and interceptions with BGP attacks and misconfigurations have been reported. Therefore, such vulnerabilities seriously impact the security of inter-domain routing.

To prevent fake (or false) routing announcements, a variety of secure BGP schemes has been proposed [1–4]. However, most of these schemes cannot be deployed in practice due to the complexity of computation and deployment. Prefix filtering [2] could be deployed to prevent attacks such as the YouTube traffic hijacking if it can be correctly applied by ISPs. However, in practice, it is difficult to achieve by ISPs because of the difficulty in configuring the filters. Among these, BGPsec [4] is the most promising secure BGP scheme that has recently been proposed by the IETF. It allows ASes to perform verification of legitimacy and authenticity of BGP route advertisements.

In this article, we systemically study the vulnerabilities of the current secure inter-domain routing protocol, that is, BGP with BGPsec. We aim to study if the design can achieve the goal of secure inter-domain routing, that is, correct packet delivery among ASes. Hence, we investigate the security of BGP with BGPsec instead of the security of BGPsec only. The existing studies [5–7] show that despite the significant vulnerability fix effort, serious vulnerabilities in the secure BGP schemes, for example, BGPsec, still exist. Several BGP attacks can be constructed to illustrate that the secure BGP design still has fundamental security weaknesses. We evaluate the impacts of the vulnerabilities by using a real trace, and find that the vulnerabilities can be easily exploited to construct attacks. In order to fix the vulnerabilities and mitigate their impacts on packet delivery on the Internet, we discuss enhancements to the design and implementation of BGPsec as well.

Background: BGPsec

Desirable Properties for BGP Security

The goal of inter-domain routing is to ensure correctness of packet forwarding among various ASes by computing and enforcing correct paths to correct destinations. In particular, it should be able to achieve the following properties by preserving correct packet forwarding paths even under attacks.

Blackhole-Resistant Routing: Any AS cannot hijack network traffic. Typically, a blackhole is used to attract traffic to an AS that normally would not traverse that AS. This security property will prevent the following two types of prefix hijacking.

- **Traffic Hijacking:** The traffic of hijacked prefixes will be completely dropped and cannot be returned to the original destinations.

Digital Object Identifier:
10.1109/MNET.2018.1800171

Qi Li, Mingwei Xu, and Jianping Wu are with Tsinghua University; Jiajia Liu (Corresponding author) is with Xidian University; Yih-Chun Hu is with the University of Illinois at Urbana-Champaign.

- **Traffic Interception:** The traffic of hijacked prefixes can be returned to the original destinations. Note that the attack does not impact network availability (i.e., reachability between any two network pairs). However, the traffic will be redirected into known adversarial networks [1].

Loop-Free Routing: Any traffic will not enter a forwarding loop incurred by false (or inaccurate) routing updates. A forwarding loop serves as an attack amplification mechanism, and can impact network connectivity, overload links, or even disrupt the network. In particular, the impact of routing loops includes significant packet loss and delay for packets caught in the loop, and increased link utilization and corresponding delay and jitter for packets that traverse the link but are not caught in the loop [8]. Therefore, violation of this property will significantly reduce network availability.

In this article, we examine the above necessary properties of secure BGP.

Securing BGP by BGPsec

Prior secure BGP schemes, such as Secure BGP (S-BGP), Secure Origin BGP (SoBGP), and SPV, focus on verifying the authenticity of BGP routing updates and authorization of ASes (or BGP routers) [1, 3]. For example, S-BGP provides both prefix origin and routing path validation to secure BGP. However, S-BGP introduces significant computation and communication overhead. In particular, it exacerbates BGP convergence performance. Recently, the IETF has been working on standardizing a new secure BGP protocol called BGPsec [4], which aims to reduce overhead while enabling similar security guarantees to S-BGP, that is, authenticating prefix origin and validating routing paths.

BGPsec leverages Resource Public Key Infrastructure (RPKI) to authenticate prefix origins [4]. The RPKI service is provided by different regional Internet registries (RIRs), for example, RIPE, APNIC, and ARIN, each of which issues certificates for the prefixes it allocates [9]. An issued certificate called a route origination authorization (ROA) for an AS that is authorized to advertise a given prefix. An ROA specifies prefixes, the maximum length of more specific prefixes that the AS is also authorized to advertise, which allows the AS to perform compact authorization and advertise a set of prefixes contained in a given length of the prefix, and the AS that is allowed to announce the prefixes. Each AS receiving a routing update verifies the ROA encoded in the update, and rejects unauthorized prefix announcements. Figure 1 illustrates an ROA that specifies that AS_x is authorized to announce prefix {10.0.0.0/16}. By verifying this ROA, ASes, for example, AS_z, can successfully validate that AS_x is indeed the owner of the prefix.

Similar to the existing secure BGP schemes, for example, S-BGP, BGPsec attempts to ensure that an AS (or a BGP router) inserts the correct AS number, that is, its own AS number, into the routing paths it announced such that the announced routing paths correctly represent real AS paths used for packet forwarding. Meanwhile, BGPsec relies on the RPKI and allows different registries to issue resource certificates specifying

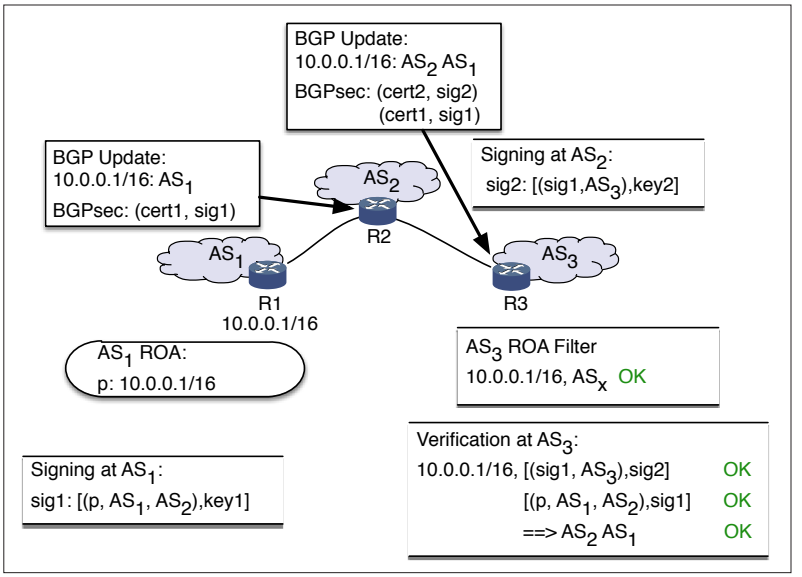

FIGURE 1. Securing BGP with BGPsec.

AS number allocation for routing path verification in each AS. For simplicity, in this article we use ASes as the entities to sign and verify routing paths, which is achieved by the signing and verification operations performed in BGP routes in the AS. Each AS signs a routing path specified in the routing update before sending the update to the neighbor AS. Different from S-BGP, BGPsec only signs the verified signature encoded in the corresponding received routing update and the AS number to which the update is sent. Figure 1 shows an example of the routing path validation of BGPsec. AS_1 signs the prefix p, its own AS number, that is, AS_1, and the AS number of the peer AS to which the update is being sent, that is, AS_2, and embeds the signature in the routing update sent to AS_2. AS_2 first verifies the signature before adopting the received update such that it can validate the authenticity of the routing path announced by the update. If the verification succeeds, it signs the previously verified signature and the AS number of the neighbor, that is, AS_3, to which the update is being sent, and then embeds the newly generated signature in the routing update sent to AS_3.

Note that in practice, ASes can leverage relying parties [4], for example, RPKI cache servers, to verify AS origins and validate routing paths in the route updates, and then distribute the verified records to all BGPsec routers within the AS. A verified route record specifies the prefixes encoded in the update (including the maximum lengths of the prefixes), the origin ASes, and the routing paths. Thereby, the BGPsec routers can directly verify if received routing updates are valid by comparing them to the stored records without performing the verification operations [4].

Vulnerabilities in BGP with BGPsec

It has been claimed that BGPsec is secure and provides authenticated prefix origin and routing paths announcements in routing updates. Unfortunately, BGPsec cannot provide the security properties we list above. In particular, BGP with BGPsec still has the following vulnerabilities.

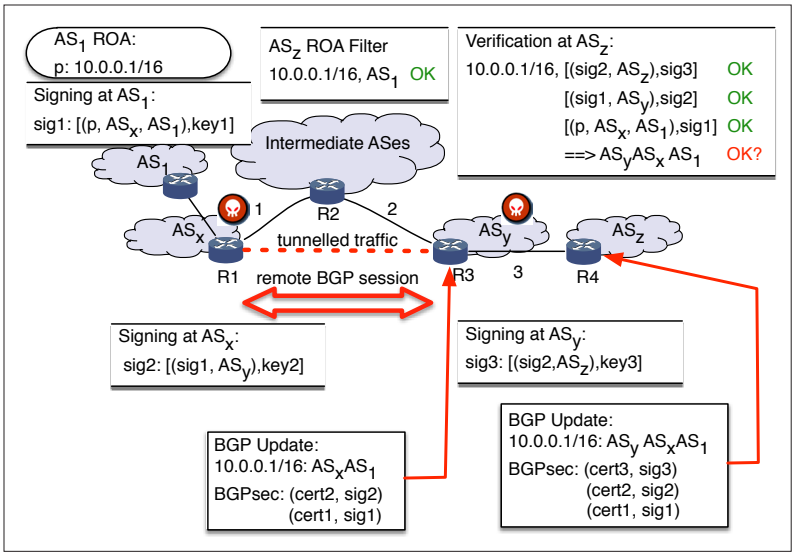

FIGURE 2. The wormhole attack to AS_z is constructed by collusion between AS_x and AS_y.

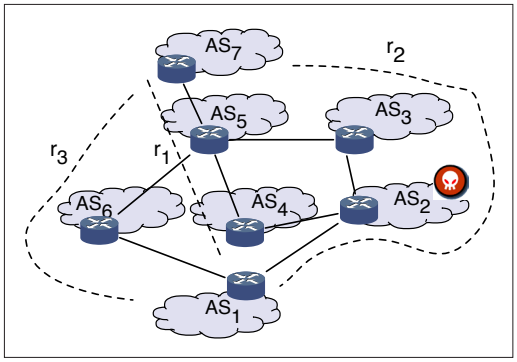

FIGURE 3. The protocol manipulation attack incurs traffic interception by leveraging the MRAI timer. AS_2 periodically announces and withdraws r_1 and r_2 such that r_3 is falsely damped forever.

Vulnerabilities in the Control Plane: BGPsec aims to secure the routing control plane to prevent blackhole attacks caused by route hijacking and propagation of forged routes. However, Li *et al.* [5] show that route hijacking attacks are still possible on the Internet even with full deployment of BGPsec, by employing wormhole attacks. Moreover, the protocol manipulation attacks in the data plane can also incur blackhole attacks [6].

Vulnerabilities in the Data Plane: Loop free routing is an important property for any routing protocol. Li *et al.* [5] show that attackers can generate forwarding loops and easily overload network links by launching a *mole attack* in the Internet. Mole attacks violate the loop-free routing property. Note that, although BGPsec does not aim to address data plane vulnerabilities, it is still necessary to fix the issue in BGP since it violates the correctness of packet forwarding. We will show that the vulnerability can possibly be fixed if the protocol is carefully operated.

Vulnerabilities Incurred by the Inconsistency between the Control and Data Planes: BGPsec does not have any considerations to verify the consistency between the control and data planes, which makes BGP vulnerable to protocol manipulation attacks [7]. By exploiting the vulnerability, traffic can be easily hijacked, even though the control and data planes are correctly verified, which violates the blackhole-resistant routing property.

In the following sections, we will elaborate on these vulnerabilities by illustrating the attacks.

Vulnerabilities in the Control Plane

Wormhole Attack: A wormhole attack can be constructed by any colluding ASes to generate fake links so as to hijack traffic, which does not require any modifications to the BGP protocol nor its implementation [5]. A wormhole attack can be easily launched by simple network configurations in two colluding ASes. Note that, although this attack is not addressed by BGPsec, it is important to fix this issue to ensure blackhole-resistant routing on the Internet. Figure 2 shows a basic wormhole attack. We assume that AS_x and AS_y want to attract and hijack traffic sent from AS_z. To achieve this, these two ASes need to collaborate and generate a routing path concealing the intermediate ASes between them, that is, AS_k and AS_l, in the announced routing path so that the length of the fake routing path is shorter than the real routing path from AS_z's point of view, which can be easily achieved by setting tunnels between them. Note that the updates via the link are transparent to the intermediate ASes since they can be encrypted.

As shown in Fig. 2, assume AS_x and AS_y are two colluding ASes. AS_1 signs prefix *10.0.0.1/16*, the AS number of AS_1, and the AS number of the AS that is going to receive the announcement, that is, AS_x, together, and embeds the signature and its ROA certificate in the route update. It sends the route update to AS_x. After verifying the signature, AS_x signs the signature generated by AS_1 together and the AS number of the fake peer AS, that is, AS_y, and sends the route update to AS_y through the built wormhole session between AS_x and AS_y. Thus, AS_y obtains "authentic" signatures of the fake routing paths from AS_x, though the session is built upon the fake link AS_x-AS_y. To further propagate the fake routing update, AS_y only needs to sign the signature from AS_x together and the AS number of the victim AS, that is, AS_z, if it wants to attract the traffic from AS_z. Upon receiving the update, AS_z can successfully verify the prefix origin of 10.0.0.1/16 by using ROA filters, and verify the forged routing path {AS_y, AS_x, AS_1} by verifying the signatures generated by AS_1, AS_x, and AS_y, respectively. In this setting, AS_z will select the forged path as the preferred the routing path instead of the real path if these paths are set with equal preference values since the fake path has the shortest path length among all learned routing paths.

In summary, colluding ASes can generate fake links by constructing wormhole attacks such that routing updates including fake routing paths also have valid signatures from victim ASes' point view. Any victim ASes deployed with BGPsec receiving forged routing paths cannot identify if the announced paths include fake links (i.e., tunneled links). Thus, these fake routing paths can be successfully verified and adopted by the victim ASes. Existing schemes (e.g., soBGP) [1] can possibly detect the attack by certificating links. Unfor-

tunately, it cannot prevent prefix origin hijacking. Therefore, wormhole attacks can still easily raise routing blackholes on the Internet, which cannot be prevented if all BGP routers are equipped with BGPsec.

Protocol Manipulation Attack (I): The protocol manipulation attack (I) developed by Song *et al.* [6] constructs traffic interception (including traffic blockholes) by leveraging the MRAI or RFD timer. MRAI is the minimum amount of delay between consecutive announcements (including route announcement and withdrawal) of a route and limits the frequency of route announcements sent to neighbors [10], while RFD is a mechanism designed to damp unstable routes that frequently change. An attacker can manipulate the preference of the routes such that the victim AS falsely chooses less preferred routes to forward the traffic.

Let us take an example of the attack that is constructed by using the MRAI timer For example, as shown in in Fig. 3, a malicious AS, that is, AS_2, controls the two most preferred routes, that is, r_1: {AS_1, AS_2, AS_4, AS_5} and r_2: {AS_1, AS_2, AS_3, AS_5}, while there exists a good route that is not controlled by the malicious node, that is, r_3: {AS_1, AS_6, AS_5, AS_8}, where the preference order of AS_1 is $r_1 > r_2 > r_3$. Here, we assume RFD is not enabled. AS_2 announces the path from AS_1 to AS_4 and then withdraws it immediately. After an MRAI interval, AS_2 announces the path from AS_1 to AS_3 and then withdraws it. AS_2 periodically repeats the above route announcement and withdrawal after each MRAI interval. Since the route withdrawal is delayed by an MRAI, r_1 and r_2 will still be adopted even though they are withdrawn. As a result, AS_7 finds that AS_1 is unreachable though r_3 exists. If there is one additional good route r_4 from AS_1 and AS_7, which is less preferred over r_3, AS_7 will prefer r_4 over r_3 since r_3 is damped permanently. Similarly, we can construct the attack by using the RFD timer.

In a nutshell, the protocol manipulation attack is launched in the routing control plane by manipulating all routing paths having a higher preference over the benign paths. The attack cannot be detected and prevented by BGP enabled with BGPsec since all routes can be verified.

Vulnerabilities in the Data Plane

Mole Attack: A mole attack can be launched to construct a forwarding loop if an ROA for a prefix is not issued on the basis of the prefix usage [5], that is, an announced prefix allocated to an AS is not to be fully consumed by the AS. The mole attack can easily leverage such prefixes to violate the property of loop-free routing. On the current Internet, larger ASes (or organizations that joined the Internet in the early stage) normally have been assigned large blocks of prefixes by Regional Internet Registries (RIRs), but they only use part of prefixes (i.e., smaller blocks of the prefixes) and assign them to their customers and sub-organizations [5, 11]. Note that, a customer (or sub-organization) network could be with or without an AS number. In this article, for simplicity, we do not differentiate between "networks" and "ASes." Existing studies [5, 11] show that such unused or unallocated prefixes can be easily abused if they are not correctly announced in BGP.

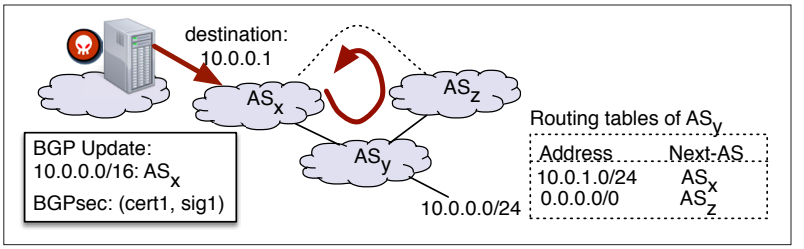

FIGURE 4. The mole attack generates a permanent forwarding loop that can be misused to overload the AS's links.

If equipped with BGPsec, ASes announce the prefix blocks with the correct signatures to the Internet, while they cannot understand the usage of the prefixes that are allocated to their customers. In the meantime, their customer ASes may set up a default route to one of their providers for simple network operations [12]. In this setting, if a customer AS does not completely consume the assigned prefixes, an attacker can easily leverage such unused prefixes and then construct mole attacks by generating traffic to the unused prefixes. The attack will significantly exacerbate the packet forwarding performance and even disrupt the connectivity of the Internet if the attacker generates a large amount of traffic to the prefix.

Figure 4 illustrates an example of a mole attack, where AS_y is multihomed to two provider ASes, that is, AS_x and AS_z, and sets up a default route to one of the provider ASes, that is, AS_z. We assume AS_y is authorized to announce prefix *10.0.0.0/16*. The routing update to announce the prefix origin and the corresponding routing paths is benign and can be verified by all BGP routers armed with BGPsec. We assume that AS_y does not fully consume *10.0.0.0/16* and the sub-prefix *10.0.0.0/24* is not assigned to any of its customers. Thus, traffic destined to the addresses in prefix *10.0.0.0/24* will be delivered among AS_x, AS_y, and AS_z permanently according to the adopted routing path and the default routing path. Therefore, the mole attack allows an attacker to easily construct forwarding loops. It can even cause congestion at the link loads among these ASes and flood the links by generating traffic to the unused prefix. Note that the attacker may not need to own a network to flood the link since the attacker can leverage botnets to generate attack traffic after probing the unused prefixes. Here, as a special case, if AS_x is the only provider AS that AS_y is attached to, only the link between AS_y and AS_x will be affected and flooded by an attacker constructing the mole attack. In order to construct a mole attack, as a prerequisite step, an attacker should identify a target prefix whose AS paths will traverse the target link and verify if the customer AS that the target link is attached to fully consumes the assigned prefixes, for example, using a network diagnosis tool. Here, a prefix not fully consumed is called the target prefix, while an AS owning the target prefix is called a target AS. If there is any target prefix, the attacker can successfully launch the mole attack.

Note that although the IPv4 prefixes have been completely allocated, there still exists a significant number of IP addresses that are not really used [13]. Therefore, it is easy to construct the mole attack on the current Internet. The situation

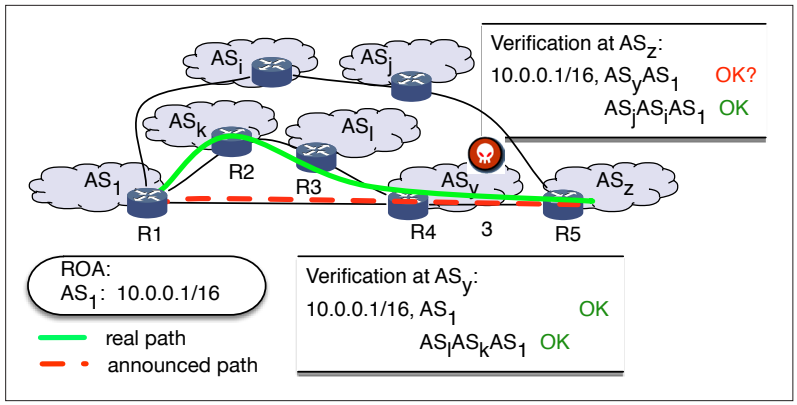

FIGURE 5. The manipulation attack allows a malicious AS to announce a shorter routing path that is not adopted in packet forwarding.

in the IPv6 networks will be worse. As we discussed above, the root cause that the mole attack can be constructed in the Internet equipped with BGPsec is that RPKI certificates are not issued according to the prefix usage. However, such a prefix announcement practice similar to the example shown in Fig. 4 is very common on the Internet [14]. Although the mole attack is not incurred by route selections only, it still should be prevented by secure inter-domain routing so that different network-level attacks, for example, forwarding loops [5] and spams [11], can be ruled out on the Internet, which is an important goal of Internet routing.

Vulnerabilities Incurred by Inconsistent Control and the Data Plane

Protocol Manipulation Attack (II): The protocol manipulation attack (II) [7] allows an attacker to modify a route update for victim ASes so as to hijack traffic, even with deployment of BGPsec. The attack can be constructed by exploiting the inconsistency between the control plane and data plane [7]. Similar to the protocol manipulation attack (I), the protocol manipulation attack (II) is constructed by manipulating routing messages. The difference between protocol manipulation attack (I) and (II) is that the first attack manipulates messages by changing configurations of BGP, while the second attack manipulates messages by directly modifying the routing announcement messages.

The current BGPsec design assumes that routes computed by the control plane will be correctly enforced in the data plane, that is, routing control and data plane is consistent, and does not have any mechanism to verify the consistency between the control plane and data plane.

Figure 5 shows an example of the manipulation attack. Let us assume AS_y wants to attract traffic from AS_z. To achieve this, AS_y can sign a transit service contract with AS_1 such that it can receive route updates from AS_1 with correct signatures. AS_1 signs the prefix and AS numbers of AS_1 and AS_y, embeds the signature and its ROA certificate in a route update, and sends it to AS_y through the built BGP session. In the meanwhile, AS_y will receive a legitimate route update from AS_l with a routing path $\{AS_l, AS_k, AS_1\}$. AS_y can successfully verify the prefix origin by ROA filters and the two routing paths, that is, AS_1 and $\{AS_l, AS_k, AS_1\}$. AS_y adopts the later one as the best candidate routing path, for example, by assigning a high preference value to AS_l. However, it can still configure AS_y to announce routing path AS_1 that is not adopted by itself to AS_z.

In this setting, AS_z will receive two verifiable routing paths from AS_y and AS_j, for example, $\{AS_y, AS_1\}$ and $\{AS_j, AS_i, AS_1\}$, with the correct ROA certificate, respectively. Since the length of $\{AS_y, AS_1\}$ from AS_y is shorter than $\{AS_j, AS_i, AS_1\}$ received from AS_j, AS_z will select the routing path from AS_y. The traffic from AS_z will be hijacked to AS_y, and the actual routing path for traffic delivery is $\{AS_y, AS_l, AS_k, AS_1\}$, instead of the announced one, that is, $\{AS_y, AS_1\}$. The protocol manipulation attack allows a malicious AS to generate a routing path that can be verified with BGPsec but not really adopted by itself. Any victim ASes receiving the update cannot validate if the announced paths are used in traffic delivery.

Evaluation and Countermeasures

In this section, we evaluate the effectiveness and impact of the attacks above and then present the countermeasures to throttle the attacks.

Vulnerability Evaluation

We measure the impacts of the attacks above by real traces. First, we use the measured AS topology from CAIDA (http://as-rank.caida.org/data/) to measure the impacts of wormhole attacks. Note that since protocol manipulation attacks have similar impacts on packet forwarding, for simplicity we do not present the results in this article. Here, we assume all ASes can be malicious. We randomly select 10 AS pairs to construct the attack and measure the number of routing paths hijacked by the attack. The topology includes 34 ASes with router-views monitors and their neighbor ASes. We directly use the CAIDA AS relationship report to set the relationships between these ASes, and uses Gao-Rexford conditions to compute routing policies for all ASes. In total, the topology contains 1425 ASes and 1405 links, and the number of routing paths is around 5510. We randomly choose ASes that have more than three neighbors to construct the attacks, and evaluate the number of routing paths of each AS that is impacted by the attacks. Figure 6a shows the number of routing paths hijacked by the wormhole attacks. We observe that around 72 percent of ASes have at least one routing path impacted by the attacks. Note that the attack can be constructed by malicious ISPs or attackers who compromised BGP routers. In particular, attackers can easily exploit the vulnerabilities of routers to compromise routers so that they can easily construct the attack by modifying the configuration of the routers. Therefore, we can conclude that the enhanced BGP protocol cannot really secure the inter-domain routing if there exist malicious ASes in practice.

Second, we evaluate the impact of mole attacks. We collect the information of unconsumed IP blocks and measure the number of links that can be impacted by the traffic delivered to these IP blocks. We use the traceroute tool to investigate the routing paths to all /24 prefixes. Also, we use the router-views data to map prefixes to corresponding ASes such that we can infer the AS links in the packet forwarding paths inferred

by traceroute. We identify various vulnerable AS links that can be exploited by the mole attacks. Figure 6b illustrates the distribution of vulnerable links that can be exploited by generating traffic to unused /24 prefixes. We find that most of the vulnerable links can be exploited by using more than five /24 prefix blocks. In particular, such attacks can be leveraged to construct stealthy DDoS attacks that can easily elude defense on the Internet. Note that the mole attack can be constructed by Internet users. Thus, the attack can be constructed stealthily to disrupt the victim network.

Possible Countermeasures

In this section, we discuss possible countermeasures against the attacks discussed above, which can be readily deployed on the Internet. Following the practice of IETF, we can possibly throttle the attacks by enhancing the BGPsec protocol and enabling a consolidated router design with the enhanced protocol.

Enhanced BGPsec Protocol: The root cause of the tiger and mole attacks is that ASes cannot verify the existence of an AS link in a routing path leading to a given prefix. To address this issue, we can extend BGPsec to generate ROA certificates only for used prefixes and generate certificates for links (instead of AS number). Thereby, a BGP router can verify the authenticity of real links in an announce path to the used prefix. Therefore, only routing paths containing real "physical" links will be adopted and announced, which prevents the tiger attack, and only the traffic to the used prefix will be forwarded and that to the unused prefix will be blackholed, which rules out the mole attack. In particular, we can develop a mechanism for BGPsec that can automatically detect the consistency between the announced prefixes and the used prefixes, and then block the unused prefixes. By leveraging the mechanism, an AS can automatically correlatively analyze the assigned prefixes and the prefixes in the routing table. To throttle the protocol manipulation attack that hijacks traffic, we encode the root cause of routing changes in the routing updates, which can avoid good routes being falsely damped [6].

Consolidated BGPsec Router Design: We can leverage trusted processors in routers, for example, Intel SGX and ARM TrustZone, to implement verifiable router implementation such that any BGP router can verify routing and packet forwarding engines of its neighbors and prevent malicious routing announcement behaviors. Any BGPsec router adopting trusted processors enabled instruction architectures, for example, Intel SGX enabled instruction architecture, can prevent system components from accessing and modifications in their protected memory regions. A trusted router design with trusted processors has been proposed in [15], which aims to attest used routing paths in neighbors to ensure that the neighbors correctly announce their adopted routing paths. Different from the design, we can allow a BGP router to verify if its neighbors use correct routing and forwarding engines by attesting if the BGP implementation in the neighbors' routing engines and the software implementation in their packet forwarding engines are correct and what it expects. Hence, the protocol manipulation attack can be detected and prevented.

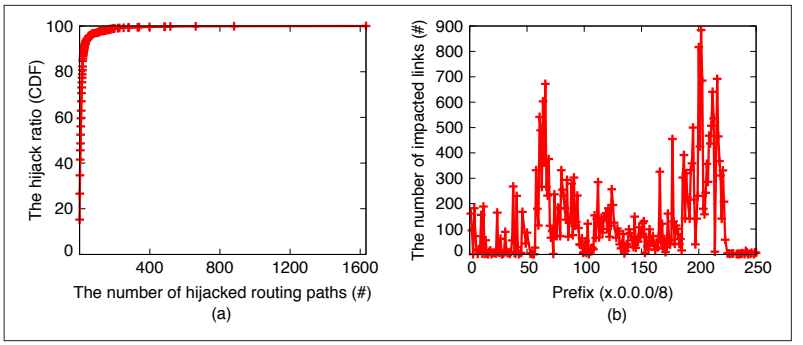

FIGURE 6. The impacts of the vulnerabilities on inter-domain routing: a) impacts of wormhole attacks; b) impacts of mole attacks.

Conclusion

In this article, we review the vulnerabilities in BGPsec. We observe that BGP armed with BGPsec cannot achieve the desired security properties of inter-domain routing due to its fundamental design flaws of BGP. We discuss several enhancements to secure inter-domain routing designs that can ensure strong security properties of packet forwarding.

Acknowledgments

This work was supported in part by the National Key R&D Program of China under Grants 2017YFB0803202 and 2016YFC0901605; the National Natural Science Foundation of China under Grants 61572278, 61625203, U1736209, 61771374, 61771373, and 61601357; NSF under Grants 09-53600 and 17-17313; the Science & Technology Program of Beijing under Grant Z171100005217001; the China 111 Project under Grant B16037; and the Fundamental Research Fund for the Central Universities under Grant JB171501, JB181506, JB181507, and JB181508.

References

[1] K. Butler *et al.*, "A Survey of BGP Security Issues and Solutions," *Proc. IEEE*, vol. 98, no. 1, 2010, pp. 100–22.
[2] J. Durand, I. Pepelnjak, and G. Doering, "BGP Operations and Security," RFC 7454, 2015.
[3] Y.-C. Hu, A. Perrig, and M. Sirbu, "SPV: Secure Path Vector Routing for Securing BGP," *Proc. SIGCOMM*, 2004, pp. 179–92.
[4] M. Lepinski and K. Sriram, "BGPsec Protocol Specification," RFC 8205, Sept. 2017.
[5] Q. Li, Y.-C. Hu, and X. Zhang, "Even Rockets Cannot Make Pigs Fly Sustainably: Can BGP be Secured with BGPsec," *SENT*, 2014.
[6] Y. Song, A. Venkataramani, and L. Gao, "Identifying and Addressing Reachability and Policy Attacks in 'Secure' BGP," *IEEE/ACM Trans. Networking*, vol. 24, no. 5, 2016, pp. 2969–82.
[7] S. Goldberg *et al.*, "How Secure are Secure Interdomain Routing Protocols," *Proc. ACM SIGCOMM*, 2010, pp. 87–98.
[8] U. Hengartner *et al.*, "Detection and Analysis of Routing Loops in Packet Traces," *Proc. ACM SIGCOMM IMW*, 2002, pp. 107–12.
[9] G. Huston and R. Bush, "Securing BGP with BGPsec," *The ISP Column*, June 2011.
[10] Y. Rekhter, T. Li, and S. Hares, "A Border Gateway Protocol 4 (BGP-4)," RFC 4271, 2006.
[11] A. Ramachandran and N. Feamster, "Understanding the Network-Level Behavior of Spammers," *Proc. SIGCOMM*, 2006, pp. 291–302.
[12] P. Mérindol *et al.*, "Quantifying ASes Multiconnectivity using Multicast Information," *Proc. IMC*, 2009, pp. 370–76.
[13] X. Cai and J. Heidemann, "Understanding Block-Level Address Usage in the Visible Internet," *Proc. SIGCOMM*, 2010, pp. 99–110.

[14] F. Le, G. G. Xie, and H. Zhang, "On Route Aggregation," *CoNEXT*, 2011, p. 6.
[15] E. Shi, A. Perrig, and L. V. Doorn, "Bind: A Fine-Grained Attestation Service for Secure Distributed Systems," *Proc. 2005 IEEE Symposium on Security and Privacy*, 2005, pp. 154–68.

Biographies

Qi Li received his Ph.D. degree from Tsinghua University. He is now an associate professor at the Graduate School at Shenzhen, Tsinghua University. His research interests are in network and system security, particularly in Internet security, mobile security, and security of large scale distributed systems. He is currently an editorial board member of *IEEE Transactions on Dependable and Secure Computing*, and has served on the organization or program committees of numerous conferences.

Jiajia Liu received his B.S. and M.S. degrees, both in computer science, from Harbin Institute of Technology in 2004 and from Xidian University in 2009, respectively, and received his Ph.D. degree in information sciences from Tohoku University in 2012. He has been a full professor at the School of Cyber Engineering, Xidian University, since 2013, and was selected into the prestigious "Huashan Scholars" program by Xidian University in 2015. His research interests cover a wide range of areas including load balancing, wireless and mobile ad hoc networks, fiber-wireless networks, Internet of Things, network security, LTE-A and 5G, SDN and NFV.

Yih-Chun Hu received his Ph.D. in 2003 from the Computer Science Department at Carnegie Mellon University. He is an associate professor in the Department of Electrical and Computer Engineering at the University of Illinois at Urbana-Champaign, an affiliate faculty of Computer Science, and faculty of the Advanced Digital Sciences Center. His current research interests are in network security and wireless networks. His general research interests are in security and systems, with an emphasis on the areas of secure systems and mobile communications.

Mingwei Xu received the Ph.D. degree and the B.Sc. degree from Tsinghua University. He is currently a full professor in the Department of Computer Science, Tsinghua University. His research interests include computer network architecture, high-speed router architecture, network protocol design, and network security. He has served on the organization or program committees of various networking conferences.

Jianping Wu received his Ph.D. degree from Tsinghua University. He is currently a professor in the Department of Computer Science, Tsinghua University. He has authored over 200 technical papers in academic journals and proceedings of international conferences, in the research areas of network architecture, high-performance routing and switching, protocol testing, and network security. He is a member of the Chinese Academy of Engineering (CAE).