Even Rockets Cannot Make Pigs Fly Sustainably: Can BGP be Secured with BGPsec?

Qi Li
ETH Zurich
qi.li@inf.ethz.ch

Yih-Chun Hu
UIUC
yihchun@illinois.edu

Xinwen Zhang
Huawei Research
xinwenzhang@gmail.com

***Abstract*—The Border Gateway Protocol (BGP) suffers from numerous security vulnerabilities, which the BGPsec protocol is supposed to fix. In this paper, we argue that fundamental properties of BGP have inherent security vulnerabilities, and that a complete re-design of BGP is needed to achieve strong security guarantees.**

I. Introduction

The Border Gateway Protocol (BGP) is the de-facto protocol to ensure the inter-AS connectivity of the Internet. However, since BGP does not have built-in mechanisms to verify if a route is genuine, it suffers from severe security vulnerabilities. Any AS (or BGP router) can announce any arbitrary route. For example, on Feb. 24th, 2008, Pakistan Telecom (AS17557) started an unauthorized announcement of prefix 208.65.153.0/24 [4]. One of Pakistan Telecom's upstream providers, PCCW Global (AS3491), forwarded this announcement to the rest of the Internet, resulting in the hijacking of YouTube traffic on a global scale for more than two hours. Many similar traffic blackholes and interceptions with active routing attacks and misconfigurations have been reported [1], [2].

To prevent false routing updates, a wide array of secure BGP schemes has been proposed [6], [12], [16], [19], [20], [30], [32]. Among these, BGPsec [20] has recently been proposed by the IETF.

As we show in this paper, despite almost two decades of attempting to fix BGP security vulnerabilities, new vulnerabilities have been identified and we also present additional weaknesses in this paper. Rather than pointing out new vulnerabilities, the goal of this paper is to argue that the design of BGP has fundamental security weaknesses, and that we need to change to a different protocol to achieve strong interdomain routing security.

Permission to freely reproduce all or part of this paper for noncommercial purposes is granted provided that copies bear this notice and the full citation on the first page. Reproduction for commercial purposes is strictly prohibited without the prior written consent of the Internet Society, the first-named author (for reproduction of an entire paper only), and the author's employer if the paper was prepared within the scope of employment.
SENT '14, 23 February 2014, San Diego, CA, USA
Copyright 2014 Internet Society, ISBN 1-891562-36-3
http://dx.doi.org/10.14722/sent.2014.23001

II. Background: BGPsec

A. Desirable Properties for BGP Security

In this paper, we examine the following four necessary properties to secure BGP.

(1) *Routing availability:* BGP should ensure convergence in the presence of different network events. Even under routing attacks, the routing protocol should quickly converge on correct paths.

(2) *Path Predictability:* A sender should know the exact path that traffic will traverse, so as to prefer routes that avoid known-adversarial networks [2], [3], [5], [11].

(3) *Blackhole-Resistant Routing:* No malicious AS can hijack network traffic. Typically, a blackhole is used to attract traffic to an AS that would otherwise not traverse that AS. This security property, will, for example, prevent prefix hijacking.

(4) *Loop-Free Routing:* No traffic will enter a forwarding loop. Because forwarding loops serve as an attack amplification mechanism, the existence of routing loops can prevent connectivity, overload links, or even disrupt the network.

B. Securing BGP by BGPsec

Prior schemes for securing BGP, such as Secure-BGP (S-BGP) [16], Secure Origin BGP (SoBGP) [32], Pretty Secure BGP (psBGP) [30], and IRR [12], focus on the authenticity and authorization of BGP updates. In particular, S-BGP provides both prefix origin and routing path validation. However, S-BGP introduces prohibitive computation and communication overhead. Recently, the IETF has been working towards standardizing a new approach called BGPsec [20], a protocol based on S-BGP that aims to provide similar security guarantees, specifically, authenticating prefix origin and routing paths. BGPsec is on track to be deployed in the Internet [20].

BGPsec leverages Resource Public Key Infrastructure (RPKI) to authenticate prefix origins [20]. The roots of trust for RPKI consists of the Regional Internet Registries (RIRs), including RIPE, APNIC, and ARIN, each of which signs certificates for the resources it allocates [15]. RPKI provides a certificate, called a Route Origination Authorization (ROA), to an entity authorized to advertise a given prefix; the ROA specifies the prefix address and size, and the AS authorized to originate the prefix. Each ISP receiving a routing update verifies the ROA, and rejects unauthorized prefix announcements. Figure 1 illustrates an ROA that specifies that AS_x is allowed to originate and announce prefix 10.0.0.0/16. With this ROA,

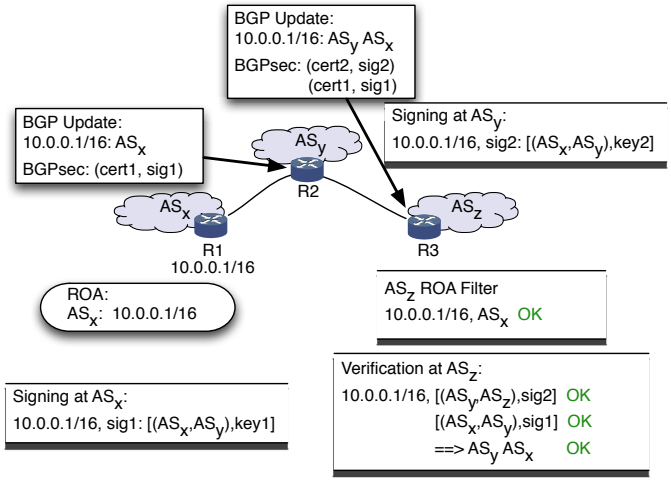

Fig. 1. Securing BGP with BGPsec

AS$_z$ can successfully validate that AS$_x$ is indeed the origin of the prefix.

Like S-BGP, BGPsec attempts to ensure that a BGP router inserts the correct AS number into routing paths such that the routing paths in BGP updates correctly represent inter-AS topologies. Routing path validation in BGPsec relies on the RPKI. In addition, it requires other certificates for signing and verifying BGP updates in each AS. For routing path validation, each BGP router signs the routing path before sending it to the next hop. Different from S-BGP, BGPsec only signs an AS key pair that specifies the local AS number and the AS number to which the update is sent.

Figure 1 shows an example of routing path validation in BGPsec. AS$_x$ signs an AS number pair (AS$_x$, AS$_y$), where AS$_y$ is the next hop AS for the update, and embeds the signature in the routing update sent to AS$_y$. AS$_y$ first verifies the signature and validates the routing path. If successful, it signs AS pair (AS$_y$, AS$_z$) and embeds the signature and the corresponding certification in the routing update sent to AS$_z$. Note that in practice, each AS has a Relying Party, e.g., a RPKI cache server, to verify the received certificates from RPKI and route updates, and then distribute the trusted route records to all BGPsec routers within the AS. A trusted route record specifies prefixes, the maximum lengths of the prefixes, and the origin ASes. In this setting, the BGPsec routers can directly check that the received routing updates are valid by comparing them to the stored trusted records [20].

It has been claimed that BGPsec is secure, and provides authenticated prefix origins and routing paths announced in routing updates. Unfortunately, BGPsec cannot provide the security properties we list above. Moreover, BGPsec introduces new security vulnerabilities. In the following sections, we will elaborate on these vulnerabilities.

III. Routing Availability

BGP is known to suffer from slow convergence [17], and ASes experience severe availability problems during route convergence. Normally, after a failure occurs, e.g., a BGP session timeout, BGP will experience a long period to explore an available route, in particular in cases with flapping routes. Route Flap Damping (RFD) [31] was proposed to damp flapping routes and expedite routing convergence. However, RFD can be attacked so that good routes can be falsely damped. For example, Sriram *et al.* [27] exploit RFD by resetting BGP sessions so as to disrupt the networks. Song *et al.* [26] recently proposed manipulation attacks where malicious ASes can disable a good route by permanently damping the route. The situation becomes worse if routing policies of BGP are conflicting. Under routing policy conflicts, the routes will never converge and the ASes cannot obtain any valid routes [13].

Moreover, BGP is vulnerable to several data plane attacks that use data plane traffic to attack routing control plane and force ASes to change their routes [33], [24], which will also raise a serious availability problem. For example, Zhang *et al.* [33] present BGP session attacks by leveraging low-rate TCP DDoS to force routers to withdraw all previously learned routing paths. Schuchard *et al.* [24] further use this attack technique to generate massive routing updates and overwhelm the computational capacity of routers.

BGPsec does not aim to address the routing availability problem. In addition, BGPsec is unable to throttle the RFD attacks and the data plane attacks, which are aimed at disabling good routes. Instead, the security mechanisms in BGPsec makes the availability problem worse, e.g., it introduces additional delays in verifying route updates, which further prolongs the route convergence time.

IV. Path Predictability

One important routing security property is *path predictability*. Specifically, in the context of routing, path predictability means that upon receiving a routing announcement, a sender can know which route the packets will follow if she sends packets to the announcing router. The incremental and distributed path computation of BGP makes it fundamentally impossible to exactly predict the path a packet will follow, and moreover, the forwarding tables may also be inconsistent with respect to the routing updates.

Prior research has already shown that on the Internet, the data plane forwarding behavior differs nearly 8% of the time from paths advertised on the control plane [21]. Though researchers have developed mechanisms for detecting such inconsistencies (e.g., [28]), such techniques are less effective against colluding adversaries.

The central problem that BGPsec faces in path predictability is that though BGPsec can ensure that the advertised path exists administratively, it cannot ensure that the advertised path is the one along which packets are being sent. In fact, without a mechanism for sharing cryptographic keys between the sender and intermediate ASes, and for efficiently using such keys for authenticating packets along a path [23], the problem of path predictability seems ill-suited to cryptographic solutions.

V. Blackhole-Resistant Routing

BGPsec aims to secure the routing control plane to prevent blackhole attacks caused by route hijacking and propagation of

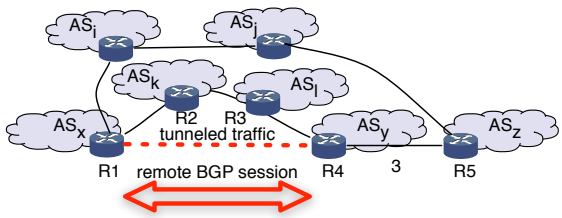

Fig. 2. Basic configurations to launch a wormhole attack. AS_x and AS_y are colluding ASes, and AS_z is victim AS.

forged routes. However, in this section, we will show that route hijacking attacks are still possible on the Internet even with full deployment of BGPsec, by employing wormhole attacks [14].

A. Wormhole Attack

In a nutshell, wormhole attacks can be launched by any AS by tweaking router configurations, which do not require any modification to the BGP protocol nor its implementation. A basic wormhole attack can be launched by the following configuration changes in two colluding ASes. Let us assume that AS_x and AS_y want to attract traffic sent by AS_z (cf. Figure 2). To achieve this, these two ASes collaborate to conceal the intermediate ASes between them, i.e., ASk and AS_l, in the routing path announced to AS_z so that the fake routing path is shorter from AS_z's point of view.

Step 1 Routers in the colluding ASes, i.e., R1 in AS_x and R4 in AS_y, build tunnels, e.g., IP-in-IP tunnels [25], Layer two Tunnel protocol (L2TP) tunnels [29], or Generic Routing Encapsulation (GRE) tunnels [9] between themselves. With the tunneled traffic, R1 and R4 create a virtual link between them.

Step 2 R1 and R4 build a BGP session (called a wormhole session) with each other with the tunnel link AS_x-ASy. That is, the network operators in AS_x sets R4 as R1's BGP peer in BGP session configuration, and the operator in AS_y sets R1 as R4's BGP peer in BGP session configuration. In this way, R1 and R4 can directly exchange their routing updates via the tunnel link.

After these configurations, AS_x and AS_y can successfully generate fake route updates even under the deployment of BGPsec. As shown in Figure 3, AS_x signs the AS number pair (AS_x, AS_y), embeds the signature and its ROA certificate in a route update, and sends it to AS_y through the built BGP session between AS_x and AS_y. In this setting, AS_y directly obtains all required "authentic" signatures from AS_x, though the session is built remotely to announce the existence of the fake link AS_x-AS_y. Next, AS_y only needs to sign the AS number pair (AS_y, AS_z) if it wants to attract the traffic from AS_z. It is clear that AS_z can successfully verify the prefix origin by ROA filters and verify the forged routing path $\{AS_z, AS_y, AS_x\}$ by verifying the AS number pairs (AS_x, AS_y) and (AS_y, AS_z). Therefore, AS_z will select the forged path if it has the shortest path length among all learned routing paths.

Wormhole attacks allow colluding ASes to generate fake links with valid signatures with BGPsec, thus produced forged routing paths also have valid signatures from the point view

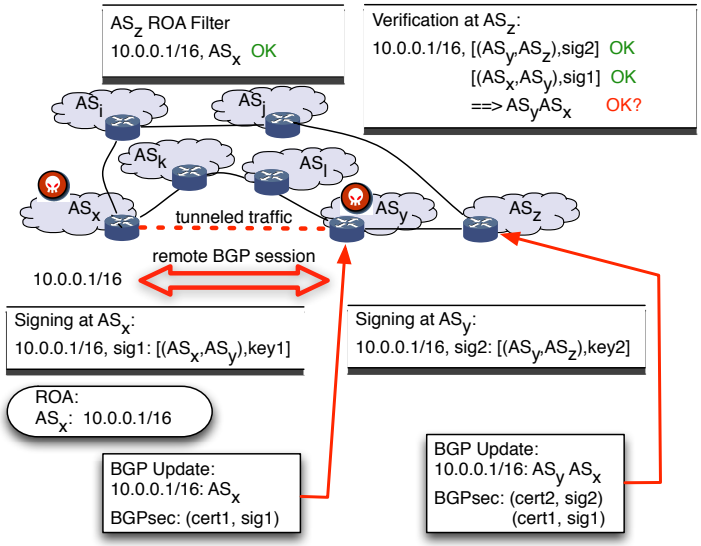

Fig. 3. Launching wormhole attacks to AS_z.

of victim ASes. Any victim ASes receiving the update cannot validate if the announced paths are delivered via a tunneled link. These fake routing paths can be successfully verified and adopted by the ASes deployed with BGPsec. Therefore, the wormhole attacks can easily raise routing blackholes, which cannot be prevented by BGPsec.

B. Impact of Wormhole Attacks

We use two different measured Internet AS topologies in our experiments. We use the measured 830-node AS topology from the SSFNet project[1] that is obtained from a BGP routing table, which is referred to as the 830-set topology, and the measured real Internet AS topology from a CAIDA dataset[2] to generate the graph of ASes. In the CAIDA topology, we focus on all 34 ASes that contribute to the Router-Views repository and their neighbor ASes, which is referred to as the rv-set topology. These two topologies include tier-1 ASes, tier-2 ASes and other ASes, and the relationships between these ASes are set according to the CAIDA AS relationship report. Table I shows the number of links in these two subgraphs. We simulate BGP routing polices on the Internet topology according to Gao-Rexford conditions [10]. To evaluate the impact of wormhole attacks, we select 10 AS pairs with different outdegree in the 830-set and rv-set subgraphs as colluding ASes to launch the attacks. We use three different strategies in the simulations: *high-attack*, *low-attack* and *random-attack* denote that ASes are selected with high, low, and random outdegrees, respectively. These ASes comply with the constraint that they have more than three neighbors. We investigate the number of ASes in the graph that are impacted by the attacks and measure the number of routing paths in each node that are hijacked by the wormhole attacks.

Figure 4 illustrates the number of hijacked routing paths at

[1] http://www.ssfnet.org/Exchange/gallery/asgraph/index.html
[2] http://as-rank.caida.org/data/

# of ASes	# of links	# of routing paths
830	791	~7312
1425	1405	~5510

TABLE I. THE NUMBER OF LINKS AND VALID ROUTING PATHS IN THE 830-SET AND RV-SET SUBGRAPH.

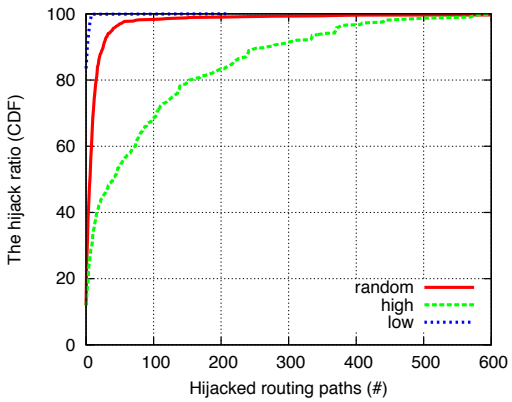

Fig. 4. CDF of hijacked routing paths in the 830-set AS graph by wormhole attacks.

each AS in the 830-set graph. In low-attack, random-attack, and high-attack scenarios, 701, 14,519, and 76,548 routing paths are hijacked by the wormhole attacks, respectively. The wormhole attack only hijacks routing paths at 17% nodes in the low-attack scenario. Since the colluding ASes in the scenario are lower-tier ASes with fewer customer ASes, they hijack fewer routing paths. We observe that in the random-attack and high-attack scenarios, wormholes hijack most nodes' routing paths in the topology. Only about 11% nodes' routing paths are not impacted by the wormhole attacks. In these two scenarios, the colluding ASes have more customer ASes, and 2,197 and 172 routing paths at one ASes are hijacked by the wormhole attacks, respectively. In particular, one AS has more than 2500 routing paths hijacked in the random-attack scenario. Here, a hijacked routing path at one AS indicates that the AS to one destination prefix is hijacked. We assume that each AS only has one prefix.

The rv-set graph has a similar distribution of hijacked routing paths to the 830-set graph. Figure 5 shows the CDF of hijacked routing paths per AS by our wormhole attacks in the rv-set AS graph. There are 72% and 44% nodes in the rv-set graph with at least one routing path hijacked in the random-attack and high-attack scenarios. In the random-attack and high-attack scenarios, 23,932 are 108,422 routing paths are hijacked by the wormhole attacks. Similar to the results in the 830-set topology, an AS in the random-attack scenario has more routing paths hijacked than the ASes in the high-attack scenario. In the AS, 1,633 routing paths are hijacked.

Note that, the number of hijacked routing paths by the wormhole attacks may be restricted by the limitations of the inferred real AS topology. Many real eBGP links are missed in these AS topology [8]. These links may be highly preferred victim ASes to deliver packets. Therefore, colluding ASes may hijack more routing paths if the wormhole attacks are launched across the Internet. BGPsec is unable to ensure that BGP can achieve blackhole-resistant routing.

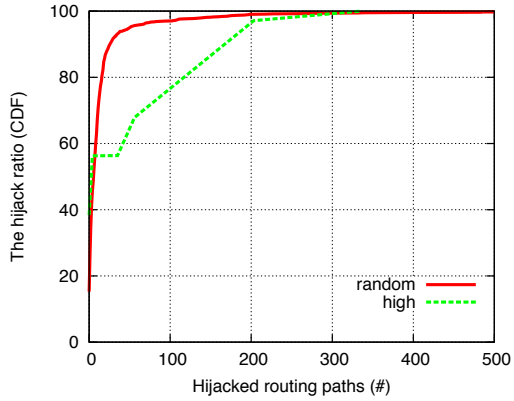

Fig. 5. CDF of hijacked routing paths in the rv-set AS graph by wormhole attacks.

VI. LOOP FREE ROUTING

Loop free routing is an important property for any routing protocol. In this section, we show that attackers can generate forwarding loops and easily overload network links by launching a *mole attacks* in the Internet. Mole attacks violate the loop-free routing property.

A. Mole Attack

In general, a mole attack can be launched if a prefix is allocated to an AS and the AS does not fully consume it, i.e., the AS does not set a specific route for the prefix. In the current Internet, larger ASes will apply for prefix blocks from Regional Internet Registries (RIRs) on behalf of their customer ASes.[3] They announce the prefix blocks to the Internet with the correct signatures, but they cannot know the usage of the prefix allocated to their customers. The customer ASes may always set a static default route to one of their providers [22]. In this setting, if any customer AS does not fully consume the prefix, an attacker can easily launch the mole attacks by generating traffic to the unused prefixes to overload the victim AS link between the provider and customer AS, and exacerbate the packet forwarding performance and even disrupt the network connectivity.

Figure 6 shows an example of the mole attack. AS_y is authorized to announce the prefix 10.0.0.0/24. The routing announcement is legitimate and can be verified by BGPsec. AS_y is multihomed to two provider ASes, i.e., AS_x and AS_z. Meanwhile, AS_y sets a default route to AS_z. We assume that AS_y does not fully use the prefix block 10.0.0.0/24, any traffic to the addresses in the prefix will be forwarded among AS_x, AS_y, and AS_z permanently, which allows an attacker to easily increase the link loads between these ASes and flood the link by generating traffic to the prefix. Here, if AS_y is only attached to one AS, e.g., AS_x, only the link connecting AS_y and AS_x will be affected by the mole attack.

To launch a mole attack and flood the target AS link, the attacker needs to locate a target prefix that will traverse the

[3] Actually, a customer network could be without any AS number. For simplicity, in this paper, we do not differentiate between "network" and "AS".

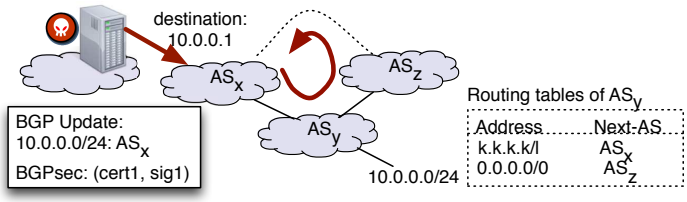

Fig. 6. The mole attack generates a permanent forwarding loop that can be misused to overload the AS's links.

target link and is not fully consumed. The AS that owns the target AS is called the target AS. The attacker can attack the target link by simply generating traffic to the IP addresses belonging to the target prefix. More specifically, the attacker can take the following steps to achieve this:

Step 1 The attacker can firstly check whether the target link can be directly flooded by investigating if the customer AS that the target link is attached to fully consumes its own prefix. If any sub-block of prefix is not used, the customer AS is the target AS and the unused prefix sub-block is the target prefix. The announced prefixes can be obtained from some public services, e.g., by looking into ROA from the RPKI server. The attacker can locate the target prefix by checking if the target link is repeating in the forwarding paths to the prefix using traceroute. If the customer AS of the target link does not have any target prefix, the attacker needs to identify all AS pairs that use the target AS link to deliver their traffic, which can be obtained from the Routeviews data[4]. If any AS pair has a common customer AS that (i) asks one of them to announce its prefixes but set a default route to the other AS, and (ii) does not fully consume its prefix, the unused prefix block is the target prefix. Note that the attacker can have some strategies to select a target prefix in the announced prefixes, e.g., normally a unused prefix appears in some larger prefix block.

Step 2 After locating the target prefix, the attacker can launch the mole attack and start flooding the link by simply generating traffic to the addresses belonging to the target prefix.

Note that although the IP address prefixes are fully allocated, a significant number of IP addresses are not used [7]. Therefore, it is easy to launch mole attacks in the current Internet. The situation in the IPv6 networks will be worse, because it is much easier to identify IPv6 prefixes that are allocated but unused.

If RPKI certificates can be issued in an usage-based approach, i.e., issuing certificates according to the real prefix usage, the authenticity of the used prefix can be verified by BGPsec and then only the traffic to the used prefix will be forwarded. The traffic to the unused prefix will be blackholed, which helps preventing the mole attack. However, it will significantly increase the complexity in operating RPKI when the usage-based approach is adopted, e.g., certificates may be frequently issued and revocated. Actually, the prefix announcement scenario shown in Figure 6 is very common in the current Internet [18]. Although network operators can install the filters

[4]http://www.routeviews.org/

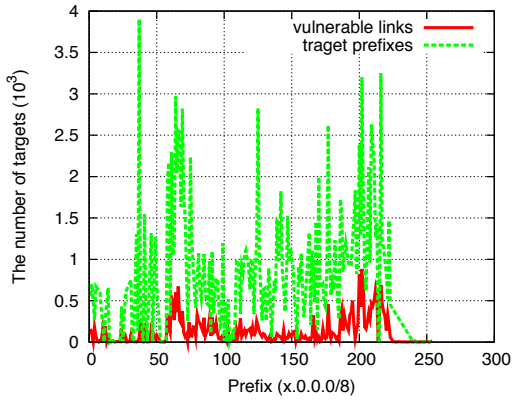

Fig. 7. Distribution of target links and prefixes of mole attacks.

in routers to drop all traffic to unused prefixes, the attack still cannot be prevented because of dynamism of the prefix announcement and usage pattern. In particular, this category of anomalies is very wide, e.g., ASes may aggregate the prefixes from the same AS and announce the aggregated prefixes. Moreover, ASes want to split the prefixes and announce the sub-blocks, e.g., to achieve traffic engineering, and the usage of prefixes are always changing. To detect and prevent the attack, routers should have an automatic mechanism to detect the consistency between the announced prefixes and the used prefixes and block the blocks of unused prefixes. BGPsec, however, does not perform any of these operations.

B. Vulnerability to Mole Attacks

To evaluate the vulnerability of AS links, we use traceroute to measure the routing paths to all /24 prefixes in the IPv4 Internet. We also use real Routeviews data to map different prefixes to the ASes so that we can identify the AS links existing in the routing paths. We measure the number of vulnerable links that can be the target links of mole attacks and the number of target prefixes that can be used to attack the target links. In this experiment, we identify a vulnerable link by identifying the link that is repeating in the routing path.

In the experiment, we observe that the routing paths to more than 200K /24 prefixes have more than 30 hops. Since traceroute is disabled on some routers, we cannot identify all vulnerable links. Surprisingly, we still identify more than 30K vulnerable links. The result is reasonable because the default routes that exist in 70% of the backbone networks [22] can easily induce permanent traffic forwarding between ASes. Normally, unused prefixes are included in the announced larger prefix blocks. A larger prefix block having one /24 unused prefix may include more /24 unused prefixes. These prefixes can probably be used to attack the same vulnerable links or the vulnerable links belonging to the same ASes. Thus, the distribution of the vulnerable links exhibits locality. Figure 7 shows the distribution of the vulnerable links in different prefix blocks.

The majority of vulnerable links are incurred by the traffic forwarded between two ASes permanently. For example, AS 5541 announces prefix 80.96.192.1/24 for AS 21462 but AS 21462 does not fully consume it, the link connecting AS 21462 and AS 5541 is thus vulnerable. Attackers can easily flood

the link by generating traffic to any address within the prefix. Note that, AS 21462 is also the provider AS of other ASes, e.g., ASes 51654, 34301, and 49591. Therefore, if an attacker floods this vulnerable link, several ASes' Internet connections may be impacted.

Moreover, we identify a significant number of vulnerable links that are incurred by the traffic forwarded among more than two ASes permanently. For example, AS 25914 announces prefix 108.160.80.0/20 but does not fully consume the prefix. Because AS 25914 sets a default route to its provider AS, AS 32881, the traffic to the prefix from AS 25914 will go through AS 32881 but go back to AS 25914 again by the path {AS 32881, AS 11666, AS 11084, AS 25914}. Therefore, the forwarding loop among these four ASes is created, and the links delivering the traffic are vulnerable. Because these links may suffer from the same target prefixes, we only count these links once.

Figure 7 illustrates the distribution of the target prefixes that can be used to flood the vulnerable links. Similar to the vulnerable link distribution, the distribution of the target prefix exhibits strong locality. We find that more than 170K /24 prefixes across the entire IPv4 address space can be used to flood the vulnerable links. On average, a vulnerable link can be flooded by using six /24 prefix blocks. Therefore, the addresses of the traffic flooding the links can be diversified and an attacker can randomly choose the addresses in the prefix blocks to attack the links, which can elude prior DDoS defenses.

VII. Conclusion

This paper investigates whether BGPsec can achieve a set of important security properties. Unfortunately, we find that BGP armed with BGPsec cannot achieve any of the security properties due to their fundamental design principles. For example, we specifically show that in BGPsec, routes can still be hijacked and routing loops still exist today. We hope that this paper will re-launch a dialog to rethink the fundamental tenets of BGP and BGPsec designs.

Acknowledgement

We would like to thank Adrian Perrig for his constructive comments on this paper.

References

[1] China's 18-minute mystery. http://www.renesys.com/blog/2010/11/chinas-18-minute-mystery.shtml.

[2] Defending against BGP man-in-the-middle attacks. http://www.renesys.com/tech/presentations/pdf/blackhat-09.pdf.

[3] Stealing the internet: An Internet-scale Man In The Middle attack. http://www.defcon.org/images/defcon-16/dc16-presentations/defcon-16-pilosov-kapela.pdf.

[4] Youtube hijacking: A RIPE NCC RIS case study. http://www.ripe.net/news/study-youtube-hijacking.html.

[5] H. Ballani, P. Francis, and X. Zhang. A study of prefix hijacking and interception in the Internet. In *SIGCOMM*, pages 265–276, 2007.

[6] K. Butler, T. Farley, P. McDaniel, and J. Rexford. A survey of bgp security issues and solutions. *Proceedings of the IEEE*, 98(1):100–122, 2010.

[7] X. Cai and J. Heidemann. Understanding block-level address usage in the visible internet. In *Proceedings of SIGCOMM*, pages 99–110, 2010.

[8] K. Chen, D. R. Choffnes, R. Potharaju, Y. Chen, F. E. Bustamante, D. Pei, and Y. Zhao. Where the sidewalk ends: extending the Internet as graph using traceroutes from P2P users. In *Proceedings of CoNEXT*, pages 217–228, 2009.

[9] D. Farinacci, T. Li, S. Hanks, D. Meyer, and P. Traina. Generic routing encapsulation (GRE). *RFC 2784*, 2000.

[10] L. Gao and J. Rexford. Stable internet routing without global coordination. *IEEE/ACM Trans. Netw.*, 9(6):681–692, 2001.

[11] S. Goldberg, S. Halevi, A. D. Jaggard, V. Ramachandran, and R. N. Wright. Rationality and traffic attraction: Incentives for honest path announcements in BGP. In *Proc. of the ACM SIGCOMM*, pages 267–278, 2008.

[12] G. Goodell, W. Aiello, T. Griffin, J. Ioannidis, P. McDaniel, and A. Rubin. Working around BGP: An incremental approach to improving security and accuracy of interdomain routing. In *Proc. of the ISOC NDSS*, pages 75–85, 2003.

[13] T. G. Griffin, F. B. Shepherd, and G. Wilfong. The stable paths problem and interdomain routing. *IEEE/ACM Transactions on Networking*, 10(2):232–243, 2002.

[14] Y.-C. Hu, A. Perrig, and D. B. Johnson. Packet leashes: A defense against wormhole attacks in wireless networks. In *INFOCOM*, 2003.

[15] G. Huston and R. Bush. Securing BGP with BGPsec. *The ISP Column*, June 2011.

[16] S. Kent, C. Lynn, and K. Seo. Secure border gateway protocol. *IEEE JSAC*, 18(4):582–592, 2000.

[17] C. Labovitz, A. Ahuja, A. Bose, and F. Jahanian. Delayed internet routing convergence. In *Proceedings of SIGCOMM*, pages 175–187, 2000.

[18] F. Le, G. G. Xie, and H. Zhang. On route aggregation. In *CoNEXT*, page 6, 2011.

[19] Q. Li, M. Xu, J. Wu, X. Zhang, P. P. Lee, and K. Xu. Enhancing the trust of internet routing with lightweight route attestation. In *Proc. of the ASIACCS*, pages 92–101, 2011.

[20] E. M. Lepinski. BGPSEC Protocol Specification. Internet-Draft draft-ietf-sidr-bgpsec-protocol-07.txt, IETF Secretariat, 2013.

[21] Z. M. Mao, J. Rexford, J. Wang, and R. H. Katz. Towards an accurate AS-level traceroute tool. In *Proceedings of SIGCOMM*, pages 365–378, 2003.

[22] P. Mérindol, V. Van den Schrieck, B. Donnet, O. Bonaventure, and J.-J. Pansiot. Quantifying ASes multiconnectivity using multicast information. In *Proceedings of IMC*, pages 370–376, 2009.

[23] J. Naous, M. Walfish, A. Nicolosi, D. Mazières, M. Miller, and A. Seehra. Verifying and enforcing network paths with icing. In *Proceedings of CoNEXT*, 2011.

[24] M. Schuchard, A. Mohaisen, D. F. Kune, N. Hopper, Y. Kim, and E. Y. Vasserman. Losing control of the Internet: Using the data plane to attack the control plane. In *NDSS*, 2011.

[25] W. Simpson. IP in IP tunneling. *RFC 1853*, 1995.

[26] Y. Song, A. Venkataramani, and L. Gao. Identifying and addressing protocol manipulation attacks in "secure'' BGP. In *Proc. of ICDCS*, 2013.

[27] K. Sriram, D. Montgomery, O. Borchert, O. Kim, and D. R. Kuhn. Study of BGP peering session attacks and their impacts on routing performance. *IEEE J.Sel. A. Commun.*, 24(10):1901–1915, 2006.

[28] L. Subramanian, V. Roth, I. Stoica, S.Shenker, and R. Katz. Listen and whisper: Security mechanisms for BGP. In *Proc. of NSDI*, 2004.

[29] W. Townsley, A. Valencia, A. Rubens, G. Pall, G. Zorn, and B. Palter. Layer two tunneling protocol 'L2TP'. *RFC 2661*, August 1999.

[30] P. van Oorschot, T. Wan, and E. Kranakis. On inter-domain routing security and pretty secure BGP (psBGP). *ACM TISSEC*, 10(3):1–41, 2007.

[31] C. Villamizar, R. Chandrac, and R. Govindan. BGP Route Flap Damping. *RFC 2439*, 1998.

[32] R. White. Through secure origin BGP. *The Internet Protocol Journal*, 6(3):15–22, 2003.

[33] Y. Zhang, Z. M. Mao, and J. Wang. Low-rate TCP-targeted DoS attack disrupts Internet routing. In *NDSS*, 2007.