Why Is It Taking So Long to Secure Internet Routing?

Routing security incidents can still slip past deployed security defenses

Sharon Goldberg, Boston University

BGP (Border Gateway Protocol) is the glue that sticks the Internet together, enabling data communications between large networks operated by different organizations. BGP makes Internet communications global by setting up routes for traffic between organizations—for example, from Boston University’s network, through larger ISPs (Internet service providers) such as Level3, Pakistan Telecom, and China Telecom, then on to residential networks such as Comcast or enterprise networks such as Bank of America.

While BGP plays a crucial role in Internet communications, it remains surprisingly vulnerable to attack. The past few years have seen a range of routing incidents that highlight the fragility of routing with BGP. They range from a simple misconfiguration at a small Indonesian ISP that took Google offline in parts of Asia,³² to a case of BGP-based censorship that leaked out of Pakistan Telecom and took YouTube offline for most of the Internet,² to a routing error that caused a large fraction of the world’s Internet traffic to be routed through China Telecom,⁶ to highly targeted traffic interception by networks in Iceland and Belarus.³⁴

People have been aware of BGP’s security issues for almost two decades and have proposed a number of solutions, most of which apply simple and well-understood cryptography or whitelisting techniques. Yet, many of these solutions remain undeployed (or incompletely deployed) in the global Internet, and the vulnerabilities persist. Why is it taking so long to secure BGP?

The answer to this question lies in the fact that BGP is a global protocol, running across organizational and national borders. As such, it lacks a single centralized authority that can mandate the deployment of a security solution; instead, every organization can autonomously decide which routing security solutions it will deploy in its own network. Thus, the deployment becomes a coordination game among thousands of independently operated networks. This is further complicated by the fact that many security solutions do not work well unless a large number of networks deploy them.

ROUTING PRIMER

BGP enables networks to route to destination *IP prefixes*. An IP prefix is a set of Internet Protocol addresses with a common prefix that is n bits in length. For example, the set of IP addresses $\{8.0.0.0, 8.0.0.1, \dots, 8.255.255.255\}$ is written as $8.0.0.0/8$, where the notation $/8$ (“slash eight”) implies that the first eight bits (the prefix) are common to all addresses in the set (in this case, those beginning with the numeral 8.). IP prefixes can have variable lengths, and the addresses in one IP prefix may be entirely contained in another IP prefix. For example, the prefix $8.8.8.0/24$, which is allocated to Google, is entirely contained in prefix $8.0.0.0/8$, which is allocated to Level 3; we say that IP prefix $8.0.0.0/8$ *covers* IP prefix $8.8.8.0/24$.

LONGEST-PREFIX-MATCH ROUTING

To decide how to forward an IP packet, an Internet router identifies the *longest* IP prefix that covers the destination IP address in the packet. For example, a packet with destination IP address 8.8.8.8 would be forwarded on the route to the longer 24-bit IP prefix 8.8.8.0/24 rather than to the shorter eight-bit IP prefix 8.0.0.0/8.

AUTONOMOUS SYSTEMS

BGP allows ASes (autonomous systems) to discover routes to destination IP prefixes. ASes are large, autonomous networks operated by different organizations. Each AS is assigned a different AS number (e.g., Google [AS 15169], China Telecom [AS 4134], Comcast [AS 7922], Boston University [AS 111], Verizon Wireless [AS 22394 and AS 6167]) and is allocated a set of IP prefixes. An AS is the *origin* for a prefix that is allocated to it. ASes are interconnected, creating a graph where nodes are ASes and edges are the links between them, as in figure 1. ASes discover routes to IP prefixes through the AS-level graph via BGP *announcements* that they receive from their neighbors. Each BGP announcement contains the AS-level path that the neighbor AS uses to reach the destination IP prefix. In figure 1,^{17,41} IP prefix 66.174.161.0/24 is allocated to Verizon Wireless, whose AS 22394 *originates* the prefix into the routing system by sending the following BGP announcement to AS 6167:

```
22394
66.174.161.0/24
```

AS 6167 selects the route and forwards all traffic for prefix 66.174.161.0/24 to its neighbor AS 22394. AS 6167 then appends its own name to the path and announces the path to its neighbors AS 2828 and AS 3356 as:

```
6167, 22394
66.174.161.0/24
```

Level 3's AS 3356 selects the path and announces it onward to its neighbor AT&T AS 7018 as:

```
3356, 6167, 22394
66.174.161.0/24
```

This process continues, and the AS-level path to prefix 66.174.161.0/24 propagates through the network.

BUSINESS RELATIONSHIPS AND ROUTING POLICIES

If an AS learns multiple routes to a particular IP prefix, then it chooses a single most-preferred route using its local routing policies. BGP provides ASes with considerable flexibility in how they select their routes. Routing decisions are typically independent of the performance of the route at a given instant; instead, they are based on a variety of factors, including route length (i.e., the number of ASes on the AS-level path) and the price of forwarding traffic to the neighbor that announced the route.

The price of forwarding traffic depends on the *business relationships*^{9,19,20} between neighboring ASes. While many business relationships exist, two are particularly relevant here. The first is a *customer-provider* relationship, where the customer AS pays the provider AS to both send and receive traffic; Level 3 and Verizon Wireless have a customer-provider relationship, represented by a directed edge in figure 1 from the customer (Verizon Wireless) to the provider (Level 3)¹⁷. The second relevant business relationship is *settlement-free peering*, where two ASes agree to transit each other's traffic for free; Level 3 and AT&T have a peering relationship, represented by an undirected edge in figure 1.

An AS will almost always avoid forwarding traffic from one neighbor to another if it cannot generate revenue by doing so; for example, China Telecom's AS 4134 in figure 1 will not carry traffic from its peer, Level 3 (AS 3356), to its other peer, AT&T (AS 7018), because neither neighbor pays China Telecom for this service. As such, China Telecom will not send a BGP announcement to AT&T (AS 7018) for the route to the prefix it learned from Level 3 (AS 3356) in figure 1.

This economically motivated behavior^{9,19,20} is often generalized as the following *rule of thumb*: AS *a* will typically announce a route to neighboring AS *n* only if: (1) *n* is a customer of *a*; (2) the route is for a prefix originated by *a*; or (3) the route is through a customer of *a*.

ATTACKS ON BGP

BGP was designed in the early 1990s—a simpler time, when the Internet was less contentious. As a result, BGP lacks basic authentication mechanisms, making it highly vulnerable to attack. We illustrate these vulnerabilities using several real-life routing incidents.

HIJACKS

BGP lacks mechanisms to authenticate the allocation of IP prefixes to autonomous systems; a prefix hijacker exploits this by originating a prefix that was not allocated to its AS. Hijacks can be classified into two types: *prefix* and *subprefix*.

Prefix hijacks. In a prefix hijack, the hijacking AS originates the exact same prefix as the AS(es) that is legitimately allocated the victim IP prefix. The bogus BGP announcement originated by the hijacking AS will be disseminated throughout the routing system, and the other ASes will use their local policies to choose between routes to the legitimate origin AS(es) and bogus routes originated by the hijacking AS.

For 18 minutes on April 8, 2010, China Telecom launched prefix hijacks for 15 percent of the Internet's prefixes.^{6,17} While there is no evidence that this incident resulted from anything other than a misconfiguration, it provides an instructive example of a "classic" prefix hijack. Figure 2¹⁷ shows one of the hijacks¹⁷: China Telecom's AS 22724 hijacks Verizon Wireless's prefix 66.174.161.0/24. The bogus route originated by AS 22724 propagates through the AS-level graph and is eventually selected by AT&T because it is shorter than the legitimate route originating at Verizon Wireless's AS 22394. Meanwhile, Level3 selects the legitimate route, because it is shorter than the bogus route. Thus, network traffic splits between the hijacking AS and the legitimate origin AS, with the nature of the split depending on routing policies used by individual ASes and the topology of the AS-level graph.

Subprefix hijacks. A far nastier attack, the subprefix hijack can potentially allow the hijacker to intercept 100 percent of the network traffic destined for the victim IP prefix. In a subprefix hijack,

the hijacking AS originates a subprefix of the victim's IP prefix—that is, a prefix that is covered by the victim IP prefix.

Perhaps the most famous subprefix hijack occurred on February 24, 2008, when Pakistan Telecom took YouTube offline. The incident² began when Pakistani authorities demanded that YouTube to be censored within Pakistan. To accomplish this, Pakistan Telecom's AS 17557 launched a subprefix hijack by originating the subprefix 208.65.153.0/24 of YouTube's prefix 208.65.153.0/22 to its customer ASes in Pakistan (e.g., Aga Khan University, Lahore Stock Exchange, Allied Bank Pakistan), as in figure 3.^{37,41} This meant that traffic destined for YouTube's servers in AS 36561 would instead be forwarded to the *longer* IP prefix originated by Pakistan Telecom's AS 17557, where traffic could then be dropped.

Events took an unexpected turn when Pakistan Telecom's bogus BGP announcement leaked out of Pakistan. PCCW, a large ISP that provides global network connectivity to Pakistan Telecom, received the bogus routing announcement, selected the bogus route, and announced it to its own neighbors. Because the bogus route was for a longer prefix (/24) than the legitimate route (/22), longest-prefix-match routing meant the bogus route was *always* more preferred by the legitimate route, and within minutes, at least two-thirds of the Internet was sending its YouTube traffic to Pakistan.² The incident was eventually resolved via manual intervention of network operators at YouTube, PCCW, and other ISPs worldwide.

Detecting hijacks. Prefix hijacks might seem to be easy to detect, just by checking that a particular prefix is originated by more than one AS. A single prefix, however, might be originated by multiple ASes for legitimate reasons (e.g., multiple ASes in a disparate part of the AS-level topology might

FIGURE 3

Pakistan Telecom Hijacks YouTube³⁷

originate a single prefix to reduce latency, so other ASes can get “closer” to the prefix). In some situations, only the legitimate holder of a prefix can be absolutely certain that a prefix is being hijacked. The identification of hijacks using anomaly-detection techniques is an active area of research.^{3,21}

ROUTE LEAKS

Route leaks are a separate class of commonly observed routing incidents.²⁸ These leaks are especially interesting because they do not involve the announcement of a bogus route. Instead, the perpetrator announces a legitimate route that it is actually using, but announces it to *too many* of its neighbors. The perpetrator is then overwhelmed by a flood of traffic from neighbors that select the leaked route.

Figure 4^{32,33} illustrates such an incident involving Moratel (AS 23947), a local ISP based in Indonesia. Moratel is not designed to transit large volumes of traffic from an international communications provider such as PCCW (AS 3491) to an important content provider such as Google (AS 15169). Per the rule of thumb in the first section, Moratel therefore should not announce its route to prefix 8.8.8.0/24 to its provider PCCW.

On November 6, 2012, however, a misconfiguration at Moratel did just that, “leaking” the route

23947, 15169

8.8.8.0/24

to PCCW. Understanding why this had an impact requires knowledge of PCCW’s local routing policies. Many routers,^{9,19,20} likely including those in PCCW’s AS, are configured to prefer a route through a neighboring customer over one through a neighboring settlement-free peer. By forwarding traffic through its customers, an AS can generate more revenue. As such, PCCW’s routers preferred the customer route through Moratel over the usual settlement-free peering route directly to Google’s AS 15196. As a result, Moratel received a huge volume of network traffic from PCCW, which quickly

took parts of Moratel's network offline and rendered 8.8.8.0/24 unreachable for PCCW and some of its neighbors, including AS 4436.

IMPACT OF ROUTING INCIDENTS

Incidents of this type can impact routing in different ways, which can be classified as *blackholes* or *interception*.

Blackhole. In a blackhole, network traffic stops at the perpetrator AS and never reaches its legitimate destination; blackholes result in network outages that are visible to end users. Blackholes happen because BGP routing decisions are typically independent of the instantaneous performance of the route. The Moratel incident is a classic example of a route leak leading to a blackhole. Hijacks can also cause blackholes; the Pakistan Telecom/YouTube incident created a blackhole because all of Pakistan Telecom's neighbors had selected its bogus route, leaving Pakistan Telecom without a working route to YouTube and forcing it to drop traffic for YouTube's prefix.

Interception. Traffic interception occurs when the perpetrator AS intercepts traffic for the victim IP prefix and then silently passes it on to the legitimate origin AS. Interception is invisible to end users. Both route leaks and hijacks can lead to traffic interception, as long as the perpetrator has a working route to the legitimate origin AS and enough network capacity to transit the extra traffic it attracts. The 2010 China Telecom hijack is an example. Figure 2 shows how one of China Telecom's routers announced the bogus hijacked routes to its neighbors, while other China Telecom routers maintained a working route to the legitimate origin of the prefix. Traffic then traveled from the hijacking router, through China Telecom's high-capacity network, back out onto the wider Internet, and finally to the legitimate origin AS for the victim IP prefix.^{2,17} Similar incidents were observed last year by Renesys, which reported several short-lived hijacks that caused traffic for targeted IP prefixes to be intercepted by ASes based in Iceland and Belarus.³⁴

DEFENSES

Many of these incidents can be eliminated through security solutions based on simple cryptography or whitelisting techniques. This section looks at these solutions—prefix filtering, RPKI (Resource Public Key Infrastructure), and BGPSEC—and highlights the challenges involved in deploying them on the global Internet.

PREFIX FILTERING

Prefix filtering is a whitelisting technique used to filter out bogus BGP announcements. It is based on the rule of thumb of the first section, which implies that an AS (e.g., Pakistan Telecom in figure 3) will announce BGP routes to its provider (PCCW) only if those routes are: (1) for its own allocated prefixes; or (2) through its own customers (Aga Khan University, Lahore Stock Exchange, etc.). As such, the provider can usually enumerate the small set of IP prefixes that are announced by its customer: that is, the set of IP prefixes allocated to Pakistan Telecom and its customer Pakistani ASes. The provider can therefore keep a *prefix list* of these IP prefixes for each customer and discard BGP announcements from a customer when they are not for prefixes on the list.

Benefit: Prefix filtering is simple and effective. Because a prefix filter is a simple whitelist, it does not usually present a large computational burden to routers. Prefix filtering has been used by various ISPs since the late 1990s and is a highly effective defense against hijacks and leaks perpetrated by

customer ASes. Our research shows that if every Internet provider with at least 25 customer ASes were to deploy prefix filters properly, this would prevent a large fraction of the Internet's ASes from launching routing leaks or hijacks.¹² Moreover, if PCCW had had properly configured prefix filters in April 2008, the Pakistan Telecom's hijack of YouTube might never have happened. The same is true of the Moratel route leak of November 2012.

Challenge: Prefix filtering works only on customer links. Prefix filters, however, typically filter BGP announcements only from *customer* ASes; this is because prefix filters are built on the assumption that the filtered AS will announce only a *small* number of IP prefixes to the filtering AS. Prefix filtering is not typically used to filter BGP announcements from providers or settlement-free peers. For example, the 2010 China Telecom incident in figure 2 could not have been prevented by prefix filtering, since China Telecom announced the bogus route along a settlement-free peering edge between China Telecom (AS 4134) and AT&T (AS 7018).

Challenge: Lopsided incentives. The incentives for deploying prefix filters are somewhat lopsided. For example, the "victims" of the 2008 Pakistan Telecom/YouTube incident were YouTube and all the impacted ASes that could not reach YouTube's hijacked prefix. However, the only AS that could have prevented the incident by using prefix filtering is PCCW itself; deploying prefix filters on the other victim ASes would do nothing to prevent the hijacked route from propagating through the Internet. Thus, the AS deploying the prefix filter (e.g., PCCW) does not have particularly strong incentives to do so, other than protecting *the rest of the Internet* from attacks by its *own customers*.

RPKI: CRYPTOGRAPHIC ORIGIN VALIDATION

The issues with prefix filtering have led to the development of many alternative security solutions. An approach that currently has significant traction is the RPKI.²⁶ Deployed since the start of this decade, the RPKI provides a trusted mapping from allocated IP prefixes to ASes authorized to originate them in BGP. To do this, the RPKI establishes a cryptographic hierarchy of *authorities* that allocate and suballocate IP address space, as well as authorize its use in BGP.

The RPKI hierarchy is rooted at the RIRs (regional Internet registries); there is roughly one RIR per continent: RIPE in Europe; ARIN in North America; LACNIC in Latin America; AfriNIC in Africa; and APNIC in Asia Pacific. Figure 5⁵ shows how ARIN allocates the prefix 8.0.0.0/8 to Level 3, which suballocates prefix 8.8.0.0/24 to Google; these allocations are accomplished using cryptographic certificates. The holder of a cryptographic certificate for a prefix can then sign an ROA (route origin authorization) authorizing a prefix (or its subprefix) to be originated in BGP; in figure 5, for example, Google issues an ROA authorizing its AS 15169 to originate 8.8.0.0/24.

Benefit: Offline cryptography. RPKI does not require any modifications to BGP message formats; nor does it require any cryptography to be performed online during routing. Instead, each day an AS syncs its local cache to the public repositories that store RPKI objects, cryptographically verifies the RPKI objects in its local cache, and pushes the resulting whitelist (mapping IP prefixes and their authorized origin AS(es)) to border routers in its AS.²⁶

Benefit: Protection from hijacks. Routers use this whitelist to filter hijacked BGP routes (i.e., those with an unauthorized origin AS). For example, in figure 2 AT&T can use the RPKI to determine that the route:

```
3356, 6167, 22394, 22394
66.174.161.0/24
```


is legitimate; AS 22394 is the origin of the route, and there is an ROA in the RPKI of figure 5 authorizing AS 22394 to originate 66.174.161.0/24. Meanwhile, the route originating at China Telecom's AS 22724 in figure 2 is bogus, since there is no ROA authorizing AS 22724 to originate 66.174.161.0/24.

Benefit: Effective incentives. The RPKI also avoids the two problems that plague prefix filtering: it can be used to filter BGP announcements made by *any* neighbor (not just neighboring customers), and it avoids lopsided deployment incentives. During the first phase of RPKI deployment, an AS that wants to protect the routes it *originates* can populate RPKI repositories with ROAs for its originated routes. (Today, RPKI contains ROAs for about 4 percent of the routes announced in BGP.³¹) During the second phase of RPKI deployment, an AS can use RPKI to discard bogus routes, thus protecting the routes it selects. (Currently we are in the very early steps of this phase, with a few ASes worldwide are experimenting with RPKI-based filtering.)

Challenge: RPKI takedowns and misconfigurations. A key challenge to RPKI deployment stems from abuse of RPKI itself.^{5,7,8,30} RPKI is designed as a threat model where BGP is under attack but RPKI is trusted. Can RPKI itself be attacked, misconfigured, or legally compelled to misclassify a legitimate BGP route as bogus? (DNS is subject to legal orders to take down domains;^{14,35} could RPKI be used to take down IP prefixes? Similar issues have already come up in several court cases.^{13,22,29}) Since routers

use RPKI to filter bogus BGP routes, the routers would then lose access to the misclassified route. This means that RPKI creates a new attack vector that can be used to blackhole routes. These issues are known to the RPKI standards community, and there are ongoing efforts to harden RPKI against this type of abuse through the development of configuration tools^{24,31,36} and fail-safe mechanisms.^{16, 23} It is too early to tell what the outcome of these efforts will be.

Challenge: RPKI can be circumvented. Unfortunately, the RPKI cannot prevent some classes of attacks.

The first is a route leak. The RPKI is designed to detect routes with an unauthorized origin AS, but in a route leak, the perpetrator leaks a legitimate route with an authorized origin AS. For example, even if nLayer (AS 4436) in figure 4 had been filtering routes based on the RPKI, it would still have selected the “leaked” Moratel route, since Google is a legitimate origin for prefix 8.8.8.0/24.

The second is a *path-shortening attack* in which an attacker announces a short bogus path to a prefix that terminates at the authorized origin AS. For example, even if RPKI were fully deployed, China Telecom (AS 4134) could still intercept traffic if it announced the route

```
4134, 22394
66.174.161.0/24
```

to AT&T in figure 2. To see why, notice that the route has a legitimate *origin* AS (AS 22394), but the route is actually bogus: there is no edge between AS 4134 and AS 22394. Thus, even if AT&T used RPKI to filter routes, it would still select the bogus route to China Telecom because it has a legitimate origin AS and is shorter than the legitimate route via Level 3.

Fortunately, however, research^{1,12,27} suggests that fewer ASes are likely to select a leaked or shortened route than one that is subprefix hijacked. During a subprefix hijack, the hijacker exploits longest-prefix-match routing to (potentially) convince *all* of the ASes on the Internet to select the bogus route. Meanwhile, both route leaks and path-shortening attacks do not exploit longest-prefix-match routing. Instead, they cause traffic to split between legitimate routes and the leaked/shortened route, with a majority of the traffic taking legitimate routes.^{12,27} The nature of the split is determined by routing policies and the AS-level topology (since ASes closer to the attacker are more likely to select the attacker’s route).

BGPSEC: CRYPTOGRAPHIC PATH VALIDATION

The community has considered a number of solutions that can eliminate the attacks that can be launched against the RPKI. Excellent surveys of these solutions (ranging from cryptographic protocol modifications to anomaly-detection techniques) are available.^{3,21} Here we focus on BGPSEC, the protocol that is currently being standardized by IETF (Internet Engineering Task Force).²⁵ Building on the RPKI’s guarantees that a BGP route has an authorized origin AS, BGPSEC also provides *path validation*.

BGPSEC builds on RPKI by adding cryptographic signatures to BGP messages. It requires each AS to digitally sign each of its BGP messages. The signature on a BGPSEC message covers (1) the prefix and AS-level path; (2) the AS number of the AS *receiving* the BGPSEC message; and includes (3) all the signed messages received from the previous ASes on the path. For example, in figure 2, AT&T’s AS 7018 would receive the following BGPSEC message from Level 3’s AS 3356:

$[66.174.161.0/24: 7018; 3356; 6167; 22394]_{3356}$

$[66.174.161.0/24: 3356; 6167; 22394]_{6167}$

$[66.174.161.0/24: 6167; 22394]_{22394}$

where the notation $[m]_A$ means message m signed by AS A . Upon receipt of a BGPSEC announcement, an AS validates the signatures and filters the route if the signatures are invalid.

Benefit: No path-shortening attacks. BGPSEC eliminates path-shortening attacks. In figure 2, China Telecom (AS 4134) announced the path

4134, 22394
66.174.161.0/24

to AT&T. With BGPSEC, this attack would fail. China Telecom (AS 4134) would not receive the BGPSEC announcement

$[66.174.161.0/24: 4134; 22394]_{22394}$

from Verizon Wireless (AS 22394), since AS 22394 and AS 4134 are not neighbors, and thus could not form a “shortened” bogus path that passes the digital signature checks required by BGPSEC.

Challenge: Online cryptography. Unlike the solutions discussed thus far, BGPSEC is an *online* cryptographic protocol; routers must cryptographically sign and verify every BGP message they send. This high computational overhead, which could require routers to be upgraded with crypto hardware accelerators, could slow down BGPSEC deployment.

Challenge: The transition to BGPSEC. All the security solutions considered here face the challenge that each AS will decide whether or not to deploy them based on its own business objectives. This challenge is particularly acute with BGPSEC, because an AS cannot validate the correctness of an AS-level path (and therefore filter bogus routes) unless all the ASes on the path have applied their signatures to the message. This means the security benefits of BGPSEC apply only after *every* AS on the path has deployed BGPSEC. This is in stark contrast to the other two solutions discussed here—prefix filtering and RPKI—where only the AS doing the filtering needs to deploy the security solution. This creates a chicken-and-egg problem; the security benefits of BGPSEC apply only after a large number of ASes have deployed BGPSEC, but there is little security incentive for anyone to be the first to deploy BGPSEC.

There are a number of ways around this chicken-and-egg problem. One idea is that a set of early-adopter ASes would deploy BGPSEC (e.g., for regulatory compliance, because of subsidies, or for public-relations purposes) and then trigger a cascade of BGPSEC deployment.^{4,10} One argument in favor of deploying BGPSEC is that because BGPSEC necessarily influences route selection, an AS that has deployed BGPSEC could attract more revenue-generating traffic from its customers that prefer to select BGPSEC-secured routes. Our research suggest that these economic incentives, along with

several other conditions, could create a cascade that led to BGPSEC adoption at a majority of ASes on the Internet.¹⁰

Beyond economic incentives, there is also the question of what security benefits are provided during the transition to BGPSEC, when some ASes have adopted it but others have not. The answer is, unfortunately, less positive. Given the routing policies that are likely to be most popular¹¹ during the transition to BGPSEC, our recent work argues that BGPSEC can provide only meager improvements to security over what is already possible with RPKI.²⁷ This is because ASes may prioritize economic considerations over security concerns. For example, given a choice between an *expensive*, BGPSEC-secured route through a provider and a *cheap*, insecure BGP route through a customer, an AS might choose the cheap, insecure path. Thus, even ASes that have deployed BGPSEC can suffer from *protocol downgrade attacks*, where an attacker convinces them to select a bogus path instead of a legitimate BGPSEC-secured path.

CONCLUSION

Today we live in an imperfect world where routing-security incidents can still slip past deployed security defenses, and no single routing-security solution is a panacea against routing attacks. Research suggests, however, that the combination of RPKI with prefix filtering could significantly improve routing security; both solutions are based on whitelisting techniques and can reduce the number of ASes that are impacted by prefix hijacks, route leaks, and path-shortening attacks. There are still several deployment challenges to overcome, since prefix filtering is limited by lopsided deployment incentives, while RPKI introduces a new dependence on centralized authorities.

This article has concentrated on protocol-based attacks on BGP. Recent research^{38,39} and media revelations^{15,18,40} indicate that routers themselves could be compromised in a manner that circumvents *protocol-based* defenses such as prefix filtering, RPKI, and BGPSEC. Thus, while we continue to make progress toward protocol-based defenses for routing security, the next frontier in routing security could very well be hardening the software and hardware used in Internet routers.

ACKNOWLEDGMENTS

I thank all my collaborators on the research I've drawn upon here: Kyle Brogle, Danny Cooper, Phillipa Gill, Shai Halevi, Ethan Heilman, Pete Hummon, Alison Kendlar, Robert Lychev, Aanchal Malhotra, Leonid Reyzin, Jennifer Rexford, Michael Schapira, and Tony Tauber. I am extremely grateful to the NSF (1017907), Cisco and the Sloan Foundation for funding our research into this topic.

REFERENCES

1. Ballani, H., Francis, P., Zhang, X. 2007. A study of prefix hijacking and interception in the Internet. *Proceedings of the ACM SIGCOMM 2007 Conference*: 265-276.
2. Brown, M. 2008. Pakistan hijacks YouTube. Renesys blog; http://www.renesys.com/blog/2008/02/pakistan_hijacks_youtube_1.shtml.
3. Butler, K., Farley, T. McDaniel, P., Rexford, J. 2010. A survey of BGP security issues and solutions. In *Proceedings of the IEEE 98(1)*: 100-122.
4. Chan, H., Dash, D., Perrig, A., Zhang, H. 2006. Modeling adoptability of secure BGP protocols. *Proceedings of the ACM SIGCOMM 2006 Conference*. 36(4): 279-290.

5. Cooper, D., Heilman, E., Brogle, K., Reyzin, L., Goldberg, S. 2013. On the risk of misbehaving RPKI authorities. In *Proceedings of the 12th ACM Workshop on Hot Topics in Networks (HotNets XII)*.
6. Cowie, J. 2010. China's 18-minute mystery. Renesys blog; <http://www.renesys.com/blog/2010/11/chinas-18-minute-mystery.shtml>.
7. FCC Communications Security, Reliability and Interoperability Council III (CSRIC). 2012. Secure BGP deployment. *Communications and Strategies*; http://transition.fcc.gov/bureaus/pshs/advisory/csric3/CSRICIII_9-12-12_WG6-Final-Report.pdf.
8. FCC Communications Security, Reliability and Interoperability Council, Working Group 6. 2013. Secure BGP deployment, final report. http://transition.fcc.gov/bureaus/pshs/advisory/csric3/CSRIC_III_WG6_Report_March_%202013.pdf
9. Gao, L., Rexford, J. 2001. Stable Internet routing without global coordination. *IEEE/ACM Transactions on Networking* 9(6): 681-692.
10. Gill, P., Schapira, M., Goldberg, S. 2011. Let the market drive deployment: a strategy for transitioning to BGP security. In *Proceedings of the ACM SIGCOMM 2011 Conference*: 14-25.
11. Gill, P., Schapira, M., Goldberg, S. 2013. A survey of interdomain routing policies. *ACM SIGCOMM Computer Communication Review* 44(1):28-34.
12. Goldberg, S., Schapira, M., Hummon, P., Rexford, J. 2010. How secure are secure interdomain routing protocols? In *Proceedings of the ACM SIGCOMM 2010 Conference*: 87-98.
13. Goldman, E. 2006. Sex.com—An update. Technology and Marketing Law blog; http://blog.ericgoldman.org/archives/2006/10/sexcom_an_updat.htm.
14. Government Printing Office. 2011. H.R.3261 - Stop Online Piracy Act.
15. Greenwald, G. 2014. How the NSA tampers with US-made Internet routers. *The Guardian* (May 12).
16. Heilman, E., Cooper, D., Reyzin, L., Goldberg, S. 2014. From the consent of the routed: improving the transparency of the RPKI. *Proceedings of the ACM SIGCOMM 2014*.
17. Hiran, R., Carlsson, N., Gill, P. 2013. Characterizing large-scale routing anomalies: a case study of the China Telecom incident. In *Passive and Active Measurement*: 229-238. Springer Berlin Heidelberg.
18. Horchert, J., Appelbaum, J., Stöcker, C. 2013. Shopping for spy gear: catalog advertises NSA toolbox. *Der Spiegel* (December 29); <http://www.spiegel.de/international/world/catalog-reveals-nsa-has-back-doors-for-numerous-devices-a-940994.html>.
19. Huston, G. 1999. Interconnection, peering and settlements, part I. *Internet Protocol Journal* 2(1).
20. Huston, G. 1999. Interconnection, peering and settlements, Part II. *Internet Protocol Journal* 2(2).
21. Huston, G., Rossi, M., Armitage, G. 2011. Securing BGP: a literature survey. *IEEE Communications Surveys and Tutorials* 13(2): 199-222.
22. Mueller, M. L., Internet Governance Project. 2011. In important case, RIPE-NCC seeks legal clarity on how it responds to foreign court orders; <http://www.internetgovernance.org/2011/11/23/in-important-case-ripe-ncc-seeks-legal-clarity-on-how-it-responds-to-foreign-court-orders/>.
23. Kent, S., Mandelberg, D. 2014. Suspenders: a fail-safe mechanism for the RPKI. Internet Engineering Task Force (IETF); <http://tools.ietf.org/html/draft-kent-sidr-suspenders-01>.
24. LACNIC Labs. RPKI looking glass; www.labs.lacnic.net/rpkitools/looking_glass/.
25. Lepinski, M., ed. 2014. BGPSEC protocol specification. IETF Network Working Group; <http://tools.ietf.org/html/draft-ietf-sidr-bgpsec-protocol-05>.

26. Lepinski, M., Kent, S. 2012. RFC 6480: an infrastructure to support secure Internet routing. Internet Engineering Task Force (IETF); <http://tools.ietf.org/html/rfc6480>.
27. Lychev, R., Goldberg, S., Schapira, M. 2013. BGP security in partial deployment. Is the juice worth the squeeze? In *Proceedings of the ACM SIGCOMM 2013 Conference*: 171-182.
28. McPherson, D., Amante, S., Osterweil, E., Mitchell, D. eds. 2013. Draft: Route-leaks & MITM attacks against BGPSEC. IETF Network Working Group; <http://tools.ietf.org/html/draft-ietf-grow-simple-leak-attack-bgpsec-no-help-03>.
29. Miller, R. 2014. Court ruling: Israeli and US terrorism victims now “own” Iran’s Internet. Joshuapundit blog (June 25); <http://joshuapundit.blogspot.com/2014/06/court-ruling-israeli-and-us-terrorism.html>.
30. Mueller, M., Kuerbis, B. 2011. Negotiating a new governance hierarchy: an analysis of the conflicting incentives to secure Internet routing. *Communications and Strategies* 81: 125-142.
31. National Institute of Standards and Technology. RPKI deployment monitor; <http://www-x.antd.nist.gov/rpki-monitor/>.
32. Paseka, T. 2012. Why Google went offline today and a bit about how the Internet works. Cloudflare blog (November 6); <http://blog.cloudflare.com/why-google-went-offline-today-and-a-bit-about>.
33. PeeringDB. 2014; <https://www.peeringdb.com/>.
34. Peterson, A. 2013. Researchers say U.S. Internet traffic was re-routed through Belarus. That’s a problem. *Washington Post* (November 20).
35. Piscitello, D. 2012. Guidance for preparing domain name orders, seizures and takedowns. Thought Paper, ICANN (March).
36. RIPE Network Coordination Centre. RPKI validator; <http://localcert.ripe.net:8088/trust-anchors>.
37. RIPE Network Coordination Centre. 2008. YouTube hijacking: A RIPE NCC RIS case study. RIPE NCC Blog; <http://www.ripe.net/internet-coordination/news/industry-developments/youtube-hijacking-a-ripe-ncc-ris-case-study>.
38. Schuchard, M., Thompson, C., Hopper, N., Kim, Y. 2012. Taking routers off their meds: why assumptions of router stability are dangerous. In *Proceedings of the Network and Distributed System Security Symposium (NDSS)*.
39. Schuchard, M., Thompson, C., Hopper, N., Kim, Y. 2013. Peer pressure: exerting malicious influence on routers at a distance. In *IEEE 33rd International Conference on Distributed Computing Systems (ICDCS)*: 571-580.
40. Storm, D. 2014. 17 exploits the NSA uses to hack PCs, routers and servers for surveillance. *ComputerWorld* (January 3); <http://blogs.computerworld.com/cybercrime-and-hacking/23347/17-exploits-nsa-uses-hack-pcs-routers-and-servers-surveillance>.
41. Wang, L., Park, J., Oliveira, R., Zhang, B. Internet AS-level topology archive; <http://irl.cs.ucla.edu/topology/>.

LOVE IT, HATE IT? LET US KNOW

feedback@queue.acm.org

SHARON GOLDBERG is an assistant professor of computer science at Boston University. Her research focuses on network security, especially issues related to the core Internet.