Survey of Papers on Securing Interdomain Routing

Elaina Wittmer
UIUC
Urbana-Champaign, IL
enw3@illinois.edu

Sam Yuan
UIUC
Urbana-Champaign, IL
jintaoy3@illinois.edu

ABSTRACT

This survey investigates the problems that have existed in interdomain routing security and strategies that exist for mitigating these problems. We also discuss the effectiveness of these strategies and any weaknesses or limitations they have. The survey contains a historical overview of interdomain routing security problems and solutions for BGP, a discussion on the effectiveness of origin authentication through RPKI, and we conclude with a look at SCION, a new internet architecture which offers secure routing without BGP at all.

1 BGP'S SECURITY PROBLEM

The Border Gateway Protocol (BGP) is the de facto protocol used for interdomain routing, allowing autonomous systems to communicate routing paths in a simple, straightforward way. However, while BGP has proven to be highly effective, it has no built-in security guarantees, making interdomain routing announcements susceptible to attacks against confidentiality, message integrity, and validity. Confidentiality attacks can involve "eavesdropping" on announcements, which can be used to discern the flow of traffic and infer customer-provider relationships. Butler et al. [1] note that this type of attack is not unique to BGP and could pertain to any protocol run over TCP, which also does not have guarantees on confidentiality. Man-in-the-middle attacks are threats to message integrity – if an adversary inserts themselves within a line of communication between two ASes and manipulates the messages, therefore manipulating an existing connection either by editing the message to include incorrect information, reestablishing a connection that was supposed to be dropped, or deleting messages such that the connection is dropped.

One type of attack to an announcement's validity is a traffic attraction attack, one of the more common threats to internet security and performance. Attackers can manipulate BGP and maliciously divert traffic to a desired location. One example of this is maliciously flooding a specific network to make it unusable; another is to route traffic through a path that benefits the attacker financially. Addressing such attacks requires a mechanism that both prevents bogus path announcements and police export policies. One version of such traffic attraction attacks is called "prefix hijacking," where an adversary makes an announcement on behalf of an IP prefix that they do not own in an attempt to manipulate the flow of traffic.

According to Butler et al., [1] at the time of publication, most ISPs relied on securing the TCP connection and applying "defensive filtering" of BGP announcements. These strategies protect against attacks on confidentiality, message integrity, and validity, but they are not without their flaws. Defensive filtering in particular is based on a series of heuristics that aim to classify faulty announcements from correct ones, which can over- or under-estimate faulty announcements and are prone to errors. Solutions for more secure BGP announcements include cryptographic solutions, protecting the BGP session from adversaries, defensive filtering strategies, registries with "correct" paths, and securing the physical layer.

Strategies can include elements from multiple solutions; for example, when creating a registry of path information to be able to apply defensive filtering to bad path announcements, it would be a good idea to use public key cryptography to secure this registry and prevent it from being tampered with. New BGP architectures have been proposed to incorporate some of these strategies into the routing architecture. For example S-BGP (Secure BGP) autenticates all information through certificates and public keys, ensuring that messages transmitted are authentic. However, challenges of using this architecture include increased convergence times due to the time cost of validating every message. On the other hand, soBGP (Secure Origin BGP) aims to authenticate the origin of the message, not the message itself. This means that soBGP keeps track of relationships between ISPs and the network topology, reducing the need to validate each message using a public key since only its origin and destination must be checked against the database. In adding new routes using soBGP, the operator verifies the routes, and can choose to do so before accepting them or after accepting them, the latter of which would reduce convergence time. Therefore it's not guaranteed to prevent bogus routes from being accepted into the routing tables, as the operator could choose to accept a bogus route in the interest of decreasing convergence time.

Goldberg et al. [4] evaluates the performance of these alternative secure interdomain routing protocols. They simulate attacks on four protocols: origin authentication, soBGP, S-BGP, and data-plane verification with and without defensive filtering. They found that advanced protocols like S-BGP

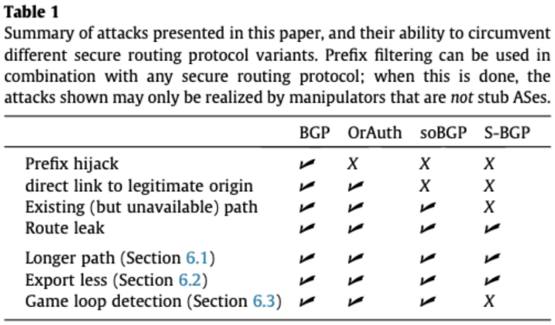

Table 1

Summary of attacks presented in this paper, and their ability to circumvent different secure routing protocol variants. Prefix filtering can be used in combination with any secure routing protocol; when this is done, the attacks shown may only be realized by manipulators that are *not* stub ASes.

	BGP	OrAuth	soBGP	S-BGP
Prefix hijack	✓	*X*	*X*	*X*
direct link to legitimate origin	✓	✓	*X*	*X*
Existing (but unavailable) path	✓	✓	✓	*X*
Route leak	✓	✓	✓	✓
Longer path (Section 6.1)	✓	✓	✓	✓
Export less (Section 6.2)	✓	✓	✓	✓
Game loop detection (Section 6.3)	✓	✓	✓	*X*

Figure 1: Table from Goldberg et al. [4] showing several security protocols and what they can and cannot protect against.

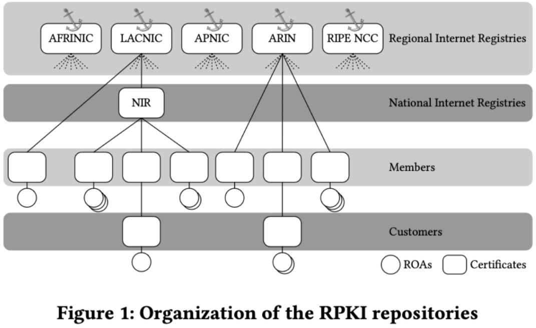

Figure 1: Organization of the RPKI repositories

Figure 2: A visualization from Chung et al. [2] of the RPKI hierarchy showing how anchor entities are linked to ROAs through certificates.

and data-plane verification do not significantly outperform soBGP. Defensive filtering could provide protection that is comparable and even better than the secure protocols. Tier 2 ASes are often easily used to launch effective attacks, even with defensive filtering. While the use of simulation data could raise questions about the validity of these results, the authors claim that their findings *underestimate* real-world attacks and are therefore more conservative estimates, rather than the reverse. The authors conclude that no protocol adequately addresses the security needs of interdomain routing as no one protocol or combination of protocols adequately addresses all of the types of attacks studied, as shown in Figure 1.

All of the proposed solutions to BGP security come with trade-offs in terms of cost and ease of implementation and no one solution can cover all the security weaknesses. Finding an optimal attack is NP-hard, but an attacker can use counterintuitive strategies to attract more traffic. Goldberg et al. [4] found that origin authentication is one of the weaker defense mechanisms in terms of defending against a diverse range of types of attacks, as Figure 1 shows. However, origin authentication is involved in all other strategies, making it a fundamental aspect of any approach [3], and it also creates a defense against prefix-hijacking, one of the more common attacks. The following sections evaluate the strengths and weaknesses of one such origin authentication system: RPKI.

2 RPKI: ORIGIN AUTHENTICATION

2.1 RPKI Background

RPKI is a public key database started in 2008 which assigns a certificate to an AS or router which is bound to a set of IP prefixes through the use of Route Origin Authorization (ROA) objects. The structure of RPKI is hierarchical – Regional Internet Registries hold the certificates and private keys which are used to validate route announcements made by ROAs. Figure 2 shows the structure of said hierarchy. These certificates validate the connection between, for example, a set of IP prefixes and a specific public key. The intention is to validate prefixes as belonging to the origin from which the announcement claims they originate, therefore authenticating the origin of the announcement. This prevents adversaries from hijacking a specific set of prefixes since RPKI will be able to determine whether or not this is true depending on whether the specified origin-IP binding matches what is in the database.

2.2 Recent Studies

RPKI began to gain popularity in 2011 and Chung et al. [2] aim to study the effect this infrastructure has had in securing BGP in the years since. The authors study eight years of BGP data, collected from datasets published by Regional Internet Registries (RIRs) as well as publicly available BGP announcement data. The paper focuses on two different types of traffic routing attacks: prefix-hijacking and sub-prefix hijacking. Sub-prefix hijacking refers to making an announcement for a more specific IP prefix than was previously announced, therefore re-routing traffic meant for IPs with that prefix, as routers will match to the longest matching prefix.

The authors identify four different errors that produce invalid route announcements: Invalid announcements made by entities which own multiple Autonomous System Numbers (ASNs), invalid announcements between customers and providers, mistakes from DDoS protection services, and invalid announcements from a different source altogether. The trends in these announcements are shown in Figure 3. Very few errors have been due to errors in automated announcements made by DDoS services and while in 2011-2012 a large number of errors were made by mismanagement of multiple ASNs, in recent years these types of errors have diminished

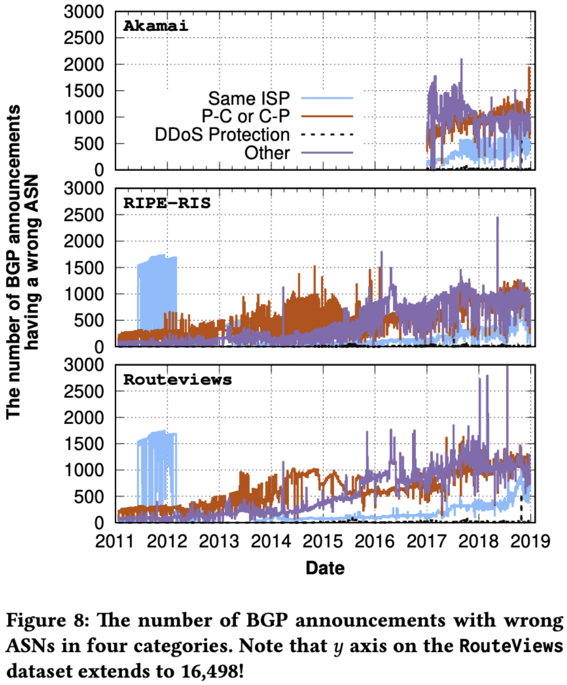

Figure 8: The number of BGP announcements with wrong ASNs in four categories. Note that y axis on the RouteViews dataset extends to 16,498!

Figure 3: Chung et al. [2] shows four different types of invalid path announcements and their prevalence from 2011 to 2019.

significantly. In more recent years, the most common mistakes have been due to providers misallocating an ROA to itself rather than the customer and other uncategorized errors. The highest, most recent spikes in errors have been due to errors labelled "Other," described further below.

ROAs allow entities such as AS's to announce a specific range of prefixes, the longest of which is specified by an ROA attribute called MaxLength, which specifies the length of the most specific prefix this particular entity may announce. Chung et al. [2] find that the majority of misconfigurations occur due to misconfigurations with this attribute specifically, causing authentications of the incorrect origin. The authors also find that adversarial attacks are caught and corrected earlier than misconfigurations, the latter of which can remain uncorrected as long as a year, meaning that misconfigurations are a major weakness of RPKI, even though it is effective in preventing adversarial attacks.

Similar to Chung et al. [2], Cooper et al. [3] also discuss the potential for internal misconfigurations, but in this case, the authors believe the capacity for misconfigurations is an indication that authority within the RPKI system goes unchecked and is therefore a potential threat to routing security. Due to the hierarchical nature of RPKI, there is little

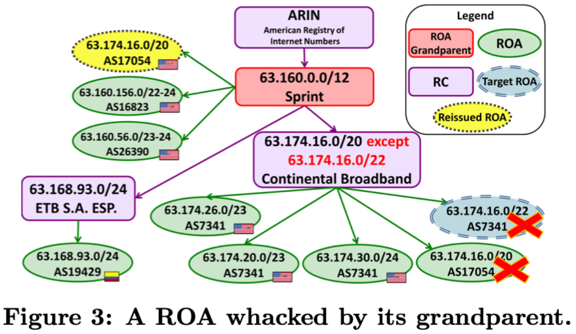

Figure 3: A ROA whacked by its grandparent.

Figure 4: A visualization from Cooper et al. [3] of the RPKI hierarchy showing that when one ROA is the target of an internal attack, multiple IPs can be affected.

defense to protect against rogue entities at the top of the hierarchy who may be able to use their authority to direct traffic in accordance to their own desires or revoke certificates at will. Cooper et al. also point out that these certificates cover prefixes and not specific IPs, so attacks to a particular certificate could potentially compromise large portions of the network, as shown in Figure 4. However, this hierarchical nature cannot be changed for a hypothetical attack – this hierarchy allows for more efficient route announcements and without it, forwarding tables could become prohibitively large.

2.3 RPKI Conclusions

Whether RPKI can be protected against authorities who abuse their access privileges is an open research question [3]. Cooper et al. point out the same issue as Chung et al. – misconfigurations are easy to implement and difficult to catch. This same issue makes it difficult to flag potential attacks from the top of the hierarchy and it's difficult if not impossible to determine whether issues arise intentionally or through error. As such, RPKI is effective against one particular type of attack, but the hierarchical control of information means that should internal issues arise, these issues could persist for extended periods of time, even up to a year [2]. Chung et al. [2] take a more optimistic viewpoint of RPKI, stating that issues are rare and are typically non-adversarial when they do arise. Both papers point out the same issue: RPKI is vulnerable to internal issues but whether these issues are malicious or not is difficult to determine. While the end result of a misconfiguration or an internal attack is the same, strategies to mitigate these issues may differ. In the case of misconfigurations, it could be made clearer what exactly the role of MaxLength is and how to appropriately assign this attribute, as this is a common source of error. Defense against internal attacks is a more complicated issue

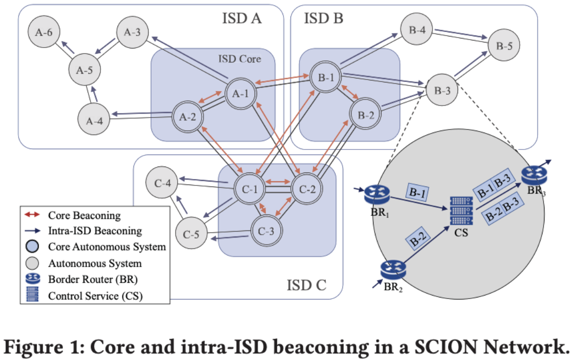

Figure 5: From Krähenbühl et al. [5], showing the architecture of SCION, consisting of interconnected core ASes, control systems, and links to more peripheral ASes.

and as mentioned above, is an open question. Perhaps the best strategy may be to abandon the hierarchical nature of RPKI and explore an entirely new system, as Krähenbühl et al. [5] propose.

3 SCION: A BGP ALTERNATIVE

Krähenbühl et al. [5] propose the next generation Internet architecture SCION, which has been deployed natively in production networks without relying on BGP. SCION utilizes Path Aware Networking (PAN) approach which enables efficient point-to-point packet delivery, even in the presence of malicious network operators. The PAN Internet architecture SCION aims to achieve high security and native interdomain multipath routing. SCION introduces Isolation Domains (ISDs) where groups of ASes agree on a set of trust roots, and an AS can belong to multiple ISDs, as shown in Figure 5. Each ISD has a Trust Root Configuration (TRC) that defines the cryptographic keys and policies. A set of core ASes govern ISDs, provide connectivity to other ISDs, and mange the trust roots. ISDs allows SCION to achieve high scalability and sovereignty in interdomain routing.

SCION provides secure interdomain routing through several mechanisms. For example, SCION offers a global framework for authentication and key establishment for secure network operations. Control plane PKI is used to enable AS authentication by verifying the authenticity of routing information exchanged between ASes with cryptographic signatures and certificates. DRKey is used to defend against DDoS attacks by allowing routers and end hosts to derive cryptographic symmetric keys and verifying the integrity of packets without relying on any state. SCION employs a hierarchical control plane that constructs AS-level path segments based on local policies and disseminates them through a global path server infrastructure. The path segments are cryptographically protected to prevent unautorhized path combinations or alteration. Additionally, the Packet-Carried Forwarding State (PCFS) simplifies the data plane by allowing routers to perform only packet forwarding without requiring any local state, thereby reducing the attack surface and simplifying the router design. SCION also provides mechanisms for fast failover, path revocation, and path exploration, enabling rapid recovery from link failures or attacks. Endpoints can also monitor the performance and quality of paths and switch to alternative paths as necessary.

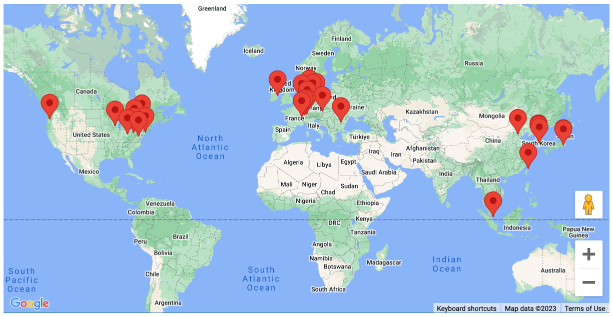

Figure 6: From the SCION website (Courtesy Google Maps), SCION has been deployed successfully in Europe, North America, and Asia.

Compared to BGP, SCION provides a more secure interdomain routing architecture that can defend against route hijacks, eavesdropping, tampering, and other threats. SCION offers high availability, low latency, high bandwidth, application-based path selection, geofencing, and transparency, making a promising approach for evolving the Internet towards a more secure and robust routing system while coexisting with today's Internet infrastructure. According to their website[1], SCION has already been deployed in multiple locations globally, including Europe, North America, and Asia, as Figure 6 shows. The system has been implemented in the Swiss financial network which handles over 3 million transactions a day. Krähenbühl et al. [5] claim that SCION is immune to prefix hijacking by nature of the multi-path design as announcements can traverse the network through various paths in the case that one of the paths fails.

4 CONCLUSION

This survey has shown that BGP has obvious security flaws and each flaw requires a specific patch. RPKI is one such patch which aims to authenticate the origin of route announcements. RPKI is very effective at defending against adversarial attacks from outsiders, but is less effective at defending against internal mistakes or internal abuses of

[1] https://scion-architecture.net/

power. Therefore, although RPKI defends against one threat to interdomain routing, it has holes that can be exploited, intentionally or otherwise.

Some recent work has proposed SCION, a different protocol which does not utilize BGP at all, and may be one solution to addressing multiple security flaws in BGP at once, rather than relying on overlying infrastructures. However, one of the challenges with this approach is convincing network authorities that replacing BGP is worth the cost in both time, money, and effort.

REFERENCES

[1] Kevin Butler, Toni R. Farley, Patrick McDaniel, and Jennifer Rexford. 2010. A Survey of BGP Security Issues and Solutions. *Proc. IEEE* 98, 1 (2010), 100–122. https://doi.org/10.1109/JPROC.2009.2034031

[2] Taejoong Chung, Emile Aben, Tim Bruijnzeels, Balakrishnan Chandrasekaran, David Choffnes, Dave Levin, Bruce M. Maggs, Alan Mislove, Roland van Rijswijk-Deij, John Rula, and Nick Sullivan. 2019. RPKI is Coming of Age: A Longitudinal Study of RPKI Deployment and Invalid Route Origins. In *Proceedings of the Internet Measurement Conference (IMC '19)*. Association for Computing Machinery, New York, NY, USA, 406–419. https://doi.org/10.1145/3355369.3355596

[3] Danny Cooper, Ethan Heilman, Kyle Brogle, Leonid Reyzin, and Sharon Goldberg. 2013. On the Risk of Misbehaving RPKI Authorities. In *Proceedings of the Twelfth ACM Workshop on Hot Topics in Networks (HotNets-XII)*. Association for Computing Machinery, New York, NY, USA, Article 16, 7 pages. https://doi.org/10.1145/2535771.2535787

[4] Sharon Goldberg, Michael Schapira, Peter Hummon, and Jennifer Rexford. 2010. How Secure Are Secure Interdomain Routing Protocols. In *Proceedings of the ACM SIGCOMM 2010 Conference (SIGCOMM '10)*. Association for Computing Machinery, New York, NY, USA, 87–98. https://doi.org/10.1145/1851182.1851195

[5] Cyrill Krähenbühl, Seyedali Tabaeiaghdaei, Christelle Gloor, Jonghoon Kwon, Adrian Perrig, David Hausheer, and Dominik Roos. 2021. Deployment and Scalability of an Inter-Domain Multi-Path Routing Infrastructure. In *Proceedings of the 17th International Conference on Emerging Networking EXperiments and Technologies (CoNEXT '21)*. Association for Computing Machinery, New York, NY, USA, 126–140. https://doi.org/10.1145/3485983.3494862