Creating a Secure Underlay for the Internet

Henry Birge-Lee, *Princeton University;* Joel Wanner, *ETH Zürich;*
Grace H. Cimaszewski, *Princeton University;* Jonghoon Kwon, *ETH Zürich;*
Liang Wang, *Princeton University;* François Wirz, *ETH Zürich;* Prateek Mittal,
Princeton University; Adrian Perrig, *ETH Zürich;* Yixin Sun, *University of Virginia*

https://www.usenix.org/conference/usenixsecurity22/presentation/birge-lee

This paper is included in the Proceedings of the
31st USENIX Security Symposium.

August 10–12, 2022 • Boston, MA, USA

978-1-939133-31-1

Open access to the Proceedings of the
31st USENIX Security Symposium is
sponsored by USENIX.

Creating a Secure Underlay for the Internet

Henry Birge-Lee, *Princeton University*
Joel Wanner, *ETH Zürich*
Grace Cimaszewski, *Princeton University*
Jonghoon Kwon, *ETH Zürich*
Liang Wang, *Princeton University*
François Wirz, *ETH Zürich*
Prateek Mittal, *Princeton University*
Adrian Perrig, *ETH Zürich*
Yixin Sun, *University of Virginia*

Abstract

Adversaries can exploit inter-domain routing vulnerabilities to intercept communication and compromise the security of critical Internet applications. Meanwhile the deployment of secure routing solutions such as Border Gateway Protocol Security (BGPsec) and Scalability, Control and Isolation On Next-generation networks (SCION) are still limited. How can we leverage emerging secure routing backbones and extend their security properties to the broader Internet?

We design and deploy an architecture to bootstrap secure routing. Our key insight is to abstract the secure routing backbone as a virtual Autonomous System (AS), called Secure Backbone AS (SBAS). While SBAS appears as one AS to the Internet, it is a federated network where routes are exchanged between participants using a secure backbone. SBAS makes BGP announcements for its customers' IP prefixes at multiple locations (referred to as Points of Presence or PoPs) allowing traffic from non-participating hosts to be routed to a nearby SBAS PoP (where it is then routed over the secure backbone to the true prefix owner). In this manner, we are the first to integrate a federated secure non-BGP routing backbone with the BGP-speaking Internet.

We present a real-world deployment of our architecture that uses SCIONLab to emulate the secure backbone and the PEERING framework to make BGP announcements to the Internet. A combination of real-world attacks and Internet-scale simulations shows that SBAS substantially reduces the threat of routing attacks. Finally, we survey network operators to better understand optimal governance and incentive models.

1 Introduction

The de facto inter-domain routing protocol, the Border Gateway Protocol (BGP), is infamously insecure. Adversaries can exploit vulnerabilities in BGP to advertise bogus routes and hijack or intercept communications towards a victim [20, 34, 36, 37, 58]. The initial secure routing efforts have focused on achieving *origin validation*, i.e., validating the owner—the origin Autonomous Systems (AS)—of an IP prefix, in order to prevent prefix hijacking attacks. A standardized mechanism is the Resource Public Key Infrastructure (RPKI), which generates records that bind an IP prefix to the origin AS [24]. However, origin validation is insufficient for preventing more sophisticated interception attacks that manipulate routing paths [52]. Recent works demonstrate the severity of interception attacks, including surveillance and compromising critical internet applications [14, 18, 54, 73]. Other proposals achieve path security in the routing backbone by authenticating the entire path information. Border Gateway Protocol Security (BGPsec) augments BGP by cryptographically signing and validating BGP paths [15]. However, BGPSec requires significant changes to the existing routing infrastructure and has yet to see production deployment. Private backbones and clean-slate Internet architectures such as Scalability, Control and Isolation On Next-generation networks (SCION) [62] have also been proposed. While they are deployed in production networks, they cannot yet be used pervasively.

In this context, can we design a system usable in today's Internet for improving routing security by leveraging an emerging secure routing backbone, such as SCION? Toward this goal, we propose the Secure Backbone AS (SBAS), a novel federated backbone infrastructure. SBAS abstracts away the secure backbone as a virtual AS that interacts with traditional ASes through conventional BGP. Within SBAS, the secure backbone allows for routing between participating customers and is immune to BGP attacks. SBAS ensures that announcements received from participating customers are indeed authorized using existing techniques such as RPKI, and SBAS customers will prioritize routes received from SBAS. Although the customers' connection to SBAS may cross a short tunnel, we demonstrate that even in this scenario SBAS offers significant security benefits.

More specifically, a participating customer (which can be either an AS or a single end host) connects to SBAS via a secure tunnel—a Virtual Private Network (VPN)—if it is not a direct neighbor. Participating customers can connect to SBAS via one or multiple distributed Points of Presence (PoPs). A

customer can bring its own prefix to announce through SBAS (e.g., an AS with its own address space), or use IP prefixes assigned by SBAS (e.g., individual clients and servers without control over their address space). Within SBAS, customer address space is distributed via an SBAS-internal iBGP mesh between the PoPs allowing PoPs to announce customer prefixes to traditional BGP neighbors (providing connectivity and improved security for non-participating hosts). We emphasize that SBAS is compatible with conventional BGP. This enables SBAS to route traffic between SBAS customers and non-participating ASes,providing security benefits even when one communication endpoint does not participate in SBAS.

We have implemented and deployed SBAS on real networks using SCIONLab [47] to emulate a secure SCION backbone, the PEERING framework [69] to send/receive BGP announcements from non-participating ASes, and the WireGuard VPN to establish secure tunnels with SBAS customers. Our implementation minimizes the need for new software and composes existing networking components to implement SBAS routing. Our key evaluation results are as follows:

- In our proof-of-concept deployment using SCIONLab and PEERING, we perform BGP attacks on IP prefixes of SBAS customers and of the VPN endpoints (in an ethical manner). SBAS successfully protects all customer-to-customer communication from our attacks and significantly improves the resilience of communication between SBAS customers and non-participating hosts.
- Our Internet topology simulations further confirm that SBAS improves resilience to routing attacks on communication with *non-participating* Internet hosts. Using a SBAS deployment with just six PoPs improves resilience by 61.8%. Furthermore, SBAS integrates well with the existing effort on Route Origin Validation (ROV): if the broader Internet enforces ROV, 98.5% of adversaries are topologically incapable of hijacking SBAS-announced routes.
- Our proof-of-concept deployment only incurs an 11% latency overhead on average (compared to the Internet), which decreases as more SBAS PoPs are deployed (as participants are closer to their nearest PoP).

An important benefit of the SBAS architecture is its compatibility with diverse secure (possibly non-BGP) backbone approaches [15, 62]. Moreover, given its strong security benefits with only a small latency overhead, SBAS represents a promising new abstraction for securing inter-domain routing that can provide much needed momentum and accelerate real-world adoption of secure backbones. To better understand the path towards a production deployment of SBAS, we surveyed network operators on appropriate incentives and governance structures, and found a potential community of early adopters as well as viable governance models. Given the promising experimental and simulation results of our proof-of-concept SBAS implementation, we dare the community of network operators to realize SBAS in a production environment.

2 Overview of Interdomain Routing Security

BGP and BGP attacks. BGP is the inter-domain routing protocol today. However, BGP lacks authentication of routing information, which allows for BGP attacks where an adversary maliciously sends BGP updates to hijack or intercept traffic to a victim AS [72]. Research shows that BGP attacks can have devastating consequences on critical Internet applications, including those that use cryptographic security mechanisms [18, 54, 73]. Also, BGP attacks are routinely seen in the wild, impacting availability of Internet services and generating millions in revenue for miscreants [23, 37, 75].

In *equally-specific* BGP attacks, the adversary makes a malicious BGP announcement for a victim's prefix that has the same prefix length as the victim's prefix. Consequently, traffic may reach either the adversary or the victim depending on the routing polices. In *more-specific* BGP attacks, the adversary announces a longer prefix than the victim's prefix. Because forwarding is based on longest prefix match, traffic destined for the more specific prefix will be routed to the adversary. This enables an adversary from almost any location in the Internet to attract a significant amount of network traffic destined to the victim. While more-specific BGP attacks are highly effective, they are not always viable given that most routers filter BGP announcements for prefixes longer than /24 [5, 40] (thus protecting /24 prefixes from more-specific attacks), and RPKI can also be used to filter malicious more-specific prefix announcements [24].

The adversary may drop/respond to the traffic (*hijack* attack) or forward the traffic back to the victim (*interception* attack) via tunneling or existing BGP paths by deliberately shaping the malicious BGP announcements [20, 64]. Interception attacks are more sophisticated but also stealthier because the victim may see little to no difference (other than potentially increased latency) in its data plane traffic.

Current Secure Routing: RPKI and route filtering. RPKI mitigates BGP attacks by providing a cryptographically secure database of IP address ownership that can be used to filter out bogus announcements [49]. In RPKI, each AS has a public–private key pair that is used to sign IP Route Origin Authorizations (ROAs) that associate IP prefixes with the ASes of their authorized origin ASes. The ASes compile the databases into a set of route filters which block announcements that do not contain a valid ⟨IP address, origin⟩ pair. However, RPKI is vulnerable to forged origin attacks. In this type of attack, the adversary claims a non-existent link to the victim in a malicious BGP update. Since RPKI only validates the origin of the IP prefix in a BGP update, the malicious update will propagate even in the presence of ROV.

An AS can implement route filters on neighbors' BGP announcements to allow announcements only from approved IP prefixes or AS paths [53, 71]. However, strict prefix-based route filtering is difficult to scale to peer-to-peer links and larger networks with a substantial number of IP prefixes; AS-

path filtering cannot prevent an adversary from announcing a malicious prefix with a legitimate-looking AS path.

Secure Internet backbone candidates. A wide variety of secure Internet routing technologies can be used as a secure backbone, ranging from BGP extensions to entirely new Internet architectures [26, 42, 43, 46, 50, 56, 57, 62, 77, 78, 80, 81]. Due to space limitations, only a few approaches are discussed below, but a more comprehensive overview of secure routing architectures is also available [25, 55, 66].

Federated backbones. BGPsec offers the properties required for a secure backbone by augmenting BGP to provide cryptographic verification of routes in the control plane [15]. BGPsec requires each AS to sign outbound BGP announcements, thus allowing ASes along the path to verify the authenticity of the announcement. BGPsec not only prevents ASes from falsely originating prefixes that were not allocated to them, but also prevents ASes from claiming fake adjacencies. Several shortcomings hampered wide-spread BGPsec deployment so far, e.g., scalability issues, slow convergence, high overhead for update verification, and vulnerabilities that remain unaddressed.

New Internet architectures can also be used for constructing a federated backbone, such as NEBULA [8], NIRA [82], and SCION [62, 84]. Specifically, SCION has been suggested as a clean-slate Internet architecture to provide secure inter-domain routing. SCION provides strong security properties: in-network per-packet source authentication, sovereignty and transparency for trust roots, and attack resilience for inter-domain routing. Of these architectures, SCION is available today as a production network from several ISPs.

Private backbones. Several corporations have developed proprietary private backbones that allow for secure data delivery, e.g., AWS and Cloudflare Argo [17, 27]. While not federated, some of these backbones allow for participants to connect via VPN tunnels and even announce their own address space. While these commercial offerings are promising, they are challenging to scale as the competing providers do not seem to move towards a federated offering.

3 Design Principles

3.1 Goals and Challenges

We seek to design a secure routing architecture that provides to the Internet **high resilience against BGP hijacking attacks**. However, our intention is not to introduce another secure routing protocol. Although secure routing protocols provide clear advantages over the currently used BGP, they have so far only achieved partial adoption. The main obstacle to large-scale adoption is that a participating entity requires a sizable financial investment while gaining limited benefits at the early stage of deployment. To overcome this, the architecture has to be: (1) readily deployable without modifications to existing Internet infrastructure and protocols, (2) readily available for customers who want to use the system, requiring minimum changes in setup, and more importantly, (3) readily beneficial to the customer even with a partial deployment of the architecture. Considering incremental deployability, we aim to leverage an already-deployed secure routing infrastructure (as a secure backbone) to mediate communication between traditional IP endpoints. This extends the benefits of the secure backbone to the broader Internet, and kick-starts routing security for Internet communication. We call this system SBAS, the Secure Backbone AS. It is important to note that SBAS does not compete with any other secure routing methods; SBAS can benefit from them since it is a complementary system, improving security in a synergistic manner. From this approach, the following research challenges arise:

Architectural continuity. Coupling of a secure routing infrastructure and the rest of the Internet requires architectural continuity. That is, the secure backbone understands BGP's control plane and seamlessly bridges remote BGP peers while leaving the leveraged secure routing infrastructure and its security guarantees intact. To this end, the secure backbone must achieve an architectural abstraction of the underlying infrastructure and provide a transparent interface to customers.

End-to-end security. In the context of mediating customer's IP endpoints via a secure backbone, the end-to-end communication path can be segmented into: an external (insecure) segment, which is comprised of the Internet links between an IP endpoint and the SBAS ingress/egress point, and an internal segment between an arbitrary ingress and egress pair of the secure routing infrastructure. To ensure end-to-end secure routing, (1) a customer must be able to select trusted ingress/egress points and securely exchange packets with hijack resilience, and (2) the secure backbone must deliver the security properties it promised to any pair of ingress/egress points even in the presence of internal adversaries.

Routing priority. To enable customers to route traffic from/to the Internet through a secure backbone, SBAS needs to disseminate the customers' prefix announcements to all other customers and external entities. Prefixes will then be announced via SBAS and the Internet, resulting in competing announcements. To maximize the ability to route securely, SBAS must be able to convince the entities receiving the announcements to prioritize routing paths through the secure backbone over the insecure Internet paths.

3.2 Threat Model

Adversary types. SBAS considers two distinct types of adversaries. The first type is an external adversary, who controls an AS on the Internet and is able to make arbitrary BGP announcements. The adversary performs BGP attacks to hijack or intercept the traffic originated from or destined to customers, which enables more sophisticated attacks such as domain validation attacks [18] and traffic analysis [73]. The second type is an internal adversary, who may compromise

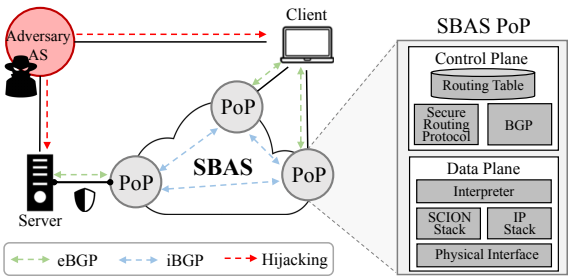

Figure 1: SBAS Overview

entities with various roles in SBAS, attempting to disrupt connectivity to legitimate customers. We also allow the two types of adversaries to collude. Attacks that do not target routing (e.g., exploiting implementation vulnerabilities, or DDoS attacks) are considered out of scope.

BGP attack types. The primary threat that SBAS aims to defend against is equally-specific prefix attacks. This is justified since all prefixes controlled by SBAS are announced as 24-bit prefixes (or 48-bit IPv6 prefixes), which can only be attacked via equally-specific prefix attacks (recall that prefix announcements longer than 24 bits in IPv4 and 48 bits in IPv6 are typically filtered). Even though the primary threat we considered is equally-specific prefix attacks, we show that communication between two SBAS customers benefits from increased resilience even in the presence of more-specific prefix attacks. We demonstrate this property in Section 7.1.

4 Design of SBAS

This section describes the control plane, data plane, and operational aspects in the design of SBAS.

4.1 SBAS Overview

SBAS is an abstraction that enables a federated backbone network to act as a single AS toward the outside Internet. As shown in Figure 1, customers of the system can connect via secure connections (e.g., VPN tunnels) to one or more PoPs, which are located at the edge of SBAS. The system supports both (1) customers that control their own address prefixes to be routed via SBAS, and (2) customers that operate smaller network domains. The latter can simply obtain addresses from an SBAS-owned address range. Internally, the PoPs form a full-mesh BGP topology over the internal routing protocol of the backbone, which is used to distribute customer announcements to the globally distributed PoPs and to achieve maximum security for traffic to secured prefixes. SBAS is fully compatible with conventional BGP security practices and internally performs validation checks to ensure that only legitimate announcements are redistributed by the system.

Moreover, the secure routing protocol used internally, along with additional security mechanisms, ensure that the redistribution scheme tolerates misbehaving SBAS members. This enables the system to extend the benefits of the secure federated backbone to the broader internet, while addressing the challenge of partial deployment incentives that are limiting the practical use of such approaches.

SBAS distinguishes between the following roles:

Customer. A customer is an entity that resides outside the backbone and obtains service from SBAS through a contract, which enables it to route traffic securely through the system. SBAS supports both (1) customers that only control single hosts (e.g., server operators or end users), and (2) entities that own entire address ranges and AS numbers.

Point of Presence (PoP). A PoP is a member of SBAS that is located at the edge, i.e., provides connectivity to SBAS customers and interfaces with the regular Internet.

Backbone operator. Such entities participate in the backbone network, but are not located at the edge; they simply participate in the internal routing and forwarding. This type of member does not need to be aware of the SBAS infrastructure running on top of the backbone.

External entity. This term refers to entities on the Internet that are unaware of SBAS.

SBAS distinguishes among three address categories:

Secure. This includes prefixes announced by SBAS customers and SBAS-owned address ranges, which are assigned to customers. Secure address ranges are announced publicly via BGP.

Internal. To provide an internal addressing scheme among PoPs, e.g., to set up iBGP sessions between PoP routers, the PoPs reserve address space for SBAS internal operation. This address space is not visible outside the SBAS infrastructure.

Global. We use this term to refer to all globally routable addresses to which the categories above do not apply.

4.2 The SBAS Abstraction

Toward the Internet, SBAS is abstracted as a single AS making BGP announcements. The defining characteristic of the system is that it employs a federated structure internally: various entities may participate by connecting to the backbone network, which runs a secure inter-domain routing protocol. Compared to offering secure routing through a tier-1 ISP or IXP that allows any customers to connect via a secure channel (e.g., VPN or direct physical link), SBAS's federated structure and abstraction provide the following benefits: (1) federation lowers the potential for centralization of the Internet and surveillance of traffic at the hyper-connected single node, (2) incremental deployability by allowing other ASes to participate, and (3) an expanding SBAS network results in a reduction of the hop distance to customers which increases resilience to routing attacks.

Virtualized full-mesh iBGP. The internal structure of SBAS can be abstracted to a full-mesh topology between the PoPs, independent of the routing protocol of the backbone. Over these connections, the PoPs redistribute announcements from

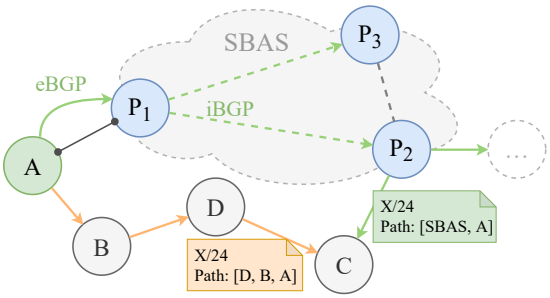

Figure 2: The route redistribution process for a prefix X owned by customer AS A. The prefix is being announced in parallel through SBAS and to neighbors of A.

SBAS customers as well as the Internet, akin to the operation of iBGP in a regular AS. In order to prevent tampering by non-PoP members, the iBGP sessions run over an encrypted and authenticated connection (such as a VPN tunnel).

4.3 Secure Route Redistribution

SBAS offers a high degree of flexibility to its customers through support for dynamic route redistribution. Contrary to a traditional AS, which is controlled by a single entity, the redistribution scheme to be used in SBAS must support its federated structure and remain secure in the presence of malicious members. In the following, we describe the design and security aspects of the route redistribution mechanism.

Federated bring-your-own IP prefix. Customers that already control one or multiple IP prefixes can use them directly with SBAS. For this purpose, SBAS implements a route redistribution mechanism that enables a customer to route incoming traffic from the Internet through the secure backbone. The process is depicted in Figure 2: the customer (AS A) initiates a BGP session with the PoP (AS P_1) over a VPN connection. Using this session, A makes an announcement for its prefix X, which is then redistributed to all other PoPs over the full-mesh iBGP topology. A remote PoP such as P_2, upon receiving such an announcement over the iBGP session with P_1, sends it out to its eBGP neighbors, i.e., Internet peers as well as SBAS customers.

Enhanced RPKI-based security. The RPKI system provides strong security properties for the first hop of BGP advertisements, but does not protect subsequent hops. The design of SBAS complements this property, as it eliminates attack surfaces on the path through its secure backbone. SBAS leverages RPKI to defend against two distinct threats: (1) customers advertising prefixes that they do not own, and (2) PoPs falsely claiming authorization for a prefix from a customer.

The first threat is prevented in SBAS using route validation at the ingress. Each announcement from a customer must carry a valid ROA that authorizes the AS to originate the prefix, which is verified both at the ingress PoP and by the other PoPs that receive the redistributed announcement. To prevent sophisticated routing attacks, SBAS additionally verifies that the AS path of these announcements does not contain any ASNs other than that of the origin (but still allowing for customer traffic engineering using path pre-pending).

An example of the second type of threat would be P_3 (in Figure 2) forwarding the announcement received from P_1 in an attempt to attract traffic to A. To prevent such malicious behavior, a customer can use RPKI to authenticate a single or multiple PoPs that are authorized to re-distribute a given prefix. This approach is similar to *path-end validation* [28], but in this case, it can be used purely by SBAS members and customers without requiring any external deployment.

SBAS-only prefix. Using an SBAS-defined BGP community tag, the customer can instruct the PoPs to only redistribute the announcement internally, i.e., to connected customers. This enables full protection of an address range against hijacking attacks, since secure prefixes are always prioritized by SBAS members and customers (as described in Section 4.5).

4.4 Customer Perspective

Service management. Customers sign up for SBAS via an interface (e.g., SBAS portal) by setting up a contract at their local PoP (details on governance aspects are given in Section 8.2). The connection to SBAS is managed by a client software that receives information about existing PoPs, including publicly reachable IP addresses and a VPN public key for each PoP. SBAS can suggest a default SBAS PoP based on a network proximity metric.

SBAS connection setup. A customer can connect to SBAS by connecting to one or multiple PoPs. A customer looking to maximize resilience to BGP attacks should generally prioritize the PoP that is the fewest BGP hops away. The connection to a PoP is set up over a VPN tunnel with the PoP's keypair, or where possible, the customer can connect directly at the PoP. The latter case has the additional benefit of eliminating the possibility of any BGP attack on the connection, as a local layer-2 connection can be used. If the customer uses multiple PoPs, one connection is designated as the primary ingress point, with the others serving as backups for improved failure resilience. In order to prevent routing loops, the VPN endpoint that is used to connect to the SBAS PoP must be assigned a non-secure address. A customer may wish to designate a part of its address range to be routed via SBAS (advertised via SBAS), and separate this from their remaining address space (advertised normally on the Internet).

Secure address assignment. For customers that do not control an address space, SBAS can offer a (paid) feature to assign single addresses from a secure SBAS-owned prefix. This option is configured via SBAS client software. Upon assigning such an address to a customer, the PoP announces it to the other PoPs over the existing iBGP sessions. This allows them to route traffic to the appropriate location and keep track of addresses that have already been assigned.

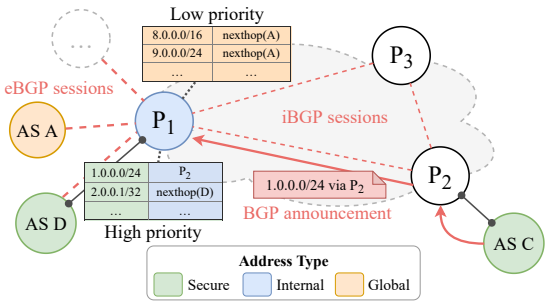

Figure 3: Control-plane configuration of P_1. nexthop(X) refers to the address used as the next hop to AS X.

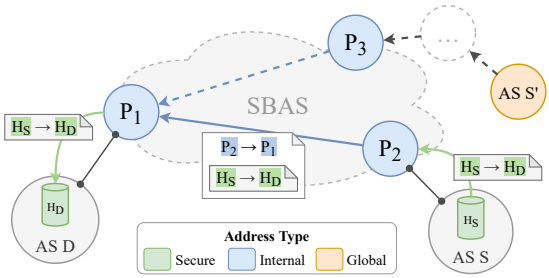

Figure 4: Routing logic for incoming traffic to a customer D who owns a secure address range. The packet shown is sent from a secure address H_S and encapsulated across SBAS using the internal addresses for the PoPs P_1 and P_2, before it is delivered to the secure address H_D. Packets may also originate from global addresses such as from AS S'.

4.5 Routing Logic at PoPs

The following sections describe how an SBAS PoP maintains routing information for the different types of addresses, and how packets are forwarded accordingly.

Control plane: Virtualized iBGP. As shown in Figure 3, each PoP maintains several iBGP sessions to other PoPs over the backbone network, as well as eBGP sessions to customers and Internet peers. From the information received over these sessions, two routing tables are constructed: The first table, which is given the highest priority, maps secure addresses to internal addresses. Each entry may be either a remote customer prefix that is mapped to an internal address representing another PoP (in the example, 1.0.0.0/24 to P_2), or a local customer prefix that can be delivered to the customer's VPN endpoint (2.0.0.1/32 to nexthop(D)). The advertisements for such routes are received over the iBGP session from other PoPs (in the former case) or over the eBGP session from customers (in the latter case).

As a lower-priority table, P_1 maintains an Internet routing table for routes obtained from its Internet peers. These routes will likely also be received via iBGP from other PoPs distributing prefixes they received from their respective neighbors. In this case, route selection can follow standard Internet policies or custom logic implemented by P_1.

Data plane: Secure address prioritization. Next, we describe data-plane forwarding decisions for different scenarios of source and destination locations, as illustrated in Figure 4.

By keeping a strict priority hierarchy between secure routes and external routes, SBAS provides resilience to BGP hijacking attacks by design. The detailed security properties achieved by this design are explored in Section 7.1.

Customer-to-customer ($H_S \rightarrow H_D$) In the simplest case, a packet originates from a secure address H_S in a customer AS S and is destined to another secure address H_D. The packet from H_S is routed through the VPN tunnel to the ingress PoP P_2. There, P_2 looks up the secure address H_D and finds the internal address for the egress PoP P_1 associated with it. The original packet is encapsulated over the backbone's internal protocol to P_2, which delivers it across the VPN tunnel to the destination in AS D.

External origin ($S' \rightarrow H_D$) We consider a packet that is destined to the secure address H_D, but originates from a source in AS S' that is unaware of SBAS. In this case, the data plane operations follow the same sequence: Having received a BGP announcement from P_3 for the secure prefix that contains H_D, AS S' will forward the packet to the nearest SBAS PoP. From this point, the same logic is applied as in the previous case.

External destination ($H_D \rightarrow S'$) For traffic with global destination addresses, the routing decision offers more options through the choice of the egress PoP. Whereas in the previous cases, the traffic was directed to the destination's preferred PoP, the decision to select an egress location is up to the ingress PoP. This makes it possible to optimize for different metrics such as hijack resilience, which can be achieved by routing based on shortest AS path length, combined with hijack detection. For instance, if an egress point notices a recent change in the AS origin of an IP prefix, this egress point may be avoided for traffic to this IP prefix until the origin change can be validated. This approach further improves the system's resilience to external BGP hijacking attacks.

Through a number of key design principles and by leveraging the secure backbone for internal routing, SBAS is able to disseminate routes securely to customers and out to the Internet. Using a strict priority hierarchy on the control plane, traffic to/from customers benefits from strong hijack resilience.

5 Implementation and Deployment

We have implemented the SBAS design and deployed it on a globally distributed infrastructure. A key feature of our implementation is that it minimizes the need for new software and composes existing networking components in a synergistic manner. Driven by this, we implement the prototype of the SBAS system on top of the globally distributed future Internet research network SCIONLab. The SBAS system is comprised of four PoPs, two PEERING announcement nodes, and three customer locations (Appendix B presents the deployment

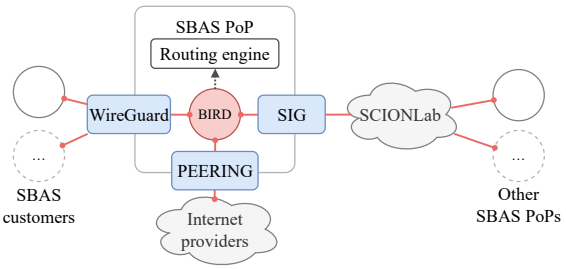

Figure 5: The interfaces (shown in filled rectangles) and BGP sessions (solid lines) maintained by an SBAS PoP.

map). Our SBAS implementation running on the PoP consists of approximately 1000 lines of code. The software automatically configures and runs the various PoP components based on configuration files that describe the setup of the SBAS instance. The full source code is publicly available.

Instantiating SBAS with SCION. We employ SCION as the secure internal routing architecture of SBAS for several reasons: (1) SCION already provides a strong PKI system by design, which is essential for the core SBAS properties such as secure route redistribution, heterogeneous trust for the federated participants, and cryptographic protection in routing, (2) SCION provides network programmability with a high degree of freedom, helping us to virtualize the internal network structure and build a full-mesh intra-SBAS topology, and (3) SCION possesses sufficient system maturity with real-world deployment and operation.

To integrate SBAS with the existing SCION architecture, we leverage the SCION-IP Gateway (SIG), which translates between IP and SCION through en- and de-capsulation of packets. The operation of a SIG at each PoP provides transparent IP connectivity without requiring any changes to customers' networking stacks. We construct the SBAS protoype deployment on SCIONLab [47], the global SCION research network spanning over 50 infrastructure ASes across the world. In our deployment, four ASes instantiated at AWS datacenters in Oregon, Frankfurt, Singapore, and Tokyo are directly connected to the SCIONLab core infrastructure, operating as SBAS PoPs.

Data-plane interfaces. Each PoP has three interfaces for different types of destinations, as shown in Figure 5: (1) A WireGuard instance to send/receive packets to/from SBAS customers that are connected to that PoP, (2) A traditional Internet interface with IP transit/peering and a BGP routing table, and (3) A SIG that encapsulates IP packets in SCION packets and sends them over the SCION backbone. This modular decomposition of interfaces enables a high degree of flexibility for SBAS. For instance, a different backbone architecture can be configured to replace SCION as a drop-in replacement without requiring changes to the other parts of the PoP software.

Control-plane management. In addition to these data plane interfaces, each SBAS PoP maintains BGP sessions with customers, IP transit providers/peers, and other SBAS PoPs. These BGP sessions are handled by the BIRD Internet routing daemon [3]. However, BIRD does not make the final routing decision; it simply exports the routes learned from its various BGP sessions into routing tables, which are then processed with different priorities by the SBAS routing engine. More details are presented in Appendix D.

Routing engine. The SBAS routing engine compiles the routes from these BGP sessions and produces the final forwarding table that enforces the security/route preference requirements of SBAS. In addition to enforcing that secure SBAS customer routes are used over standard Internet routes, the SBAS routing engine can be extended to consider which SCION paths are used to reach specific SBAS PoPs, enabling advanced route selection models such as carbon-emission-based routing [32]. See Appendix D for more details.

Packet handling. Communications inside SCION (e.g., from the SIG interface to remote PoPs) do not use IP for addressing. In SBAS, we encapsulate a SCION packet into an IP packet, using Generic Routing Encapsulation [39], and maintain a *static* table at the SIG that contains a single entry per PoP, mapping its internal IP address to its SCION address. This enables us to have a unified IP routing table for both IP and non-IP packets. The actual routing of each packet can be performed efficiently by the Linux kernel.

BGP connectivity. To provide PoPs with BGP connectivity and an Internet routing table, we use the PEERING framework [70], which allows researchers to make BGP announcements and forward packets through peers and upstream transit providers. Using this component, our announcements are propagated to the worldwide BGP ecosystem.

6 Latency Evaluation

Using the global deployment illustrated in Figure 12, we conduct a series of experiments measuring the latency achieved by SBAS in realistic scenarios. SBAS is competitive with Internet latency and even improves upon it in some cases, despite running in a research testbed using mostly overlay links. Since the PoP components only introduce sub-millisecond overhead (Appendix A), we focus on SCIONLab overhead and end-to-end latency in this section.

The final end-to-end latency $\ell_{S\to D}$ between a source customer S (connected to ingress PoP I) and a destination customer D (connected to egress PoP E) is composed as follows:

$$\underbrace{\ell_{S\to D} = \ell_{S\to I} + \ell_{E\to D} + \overbrace{2\cdot delay_{PoP}}^{\text{Appendix A}} + \overbrace{\ell_{I\to E}}^{\text{Section 6.1}}}_{\text{Section 6.2}}$$

6.1 Latency Optimization Between PoPs

We demonstrate that the backbone network can be leveraged to optimize and even reduce latency between PoPs. By of-

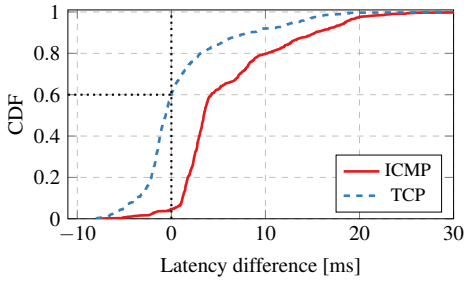

Figure 6: Difference between SCIONLab and Internet latency between all pairs of PoPs. A negative difference indicates that SCIONLab achieves better latency than the Internet. The individual latency measurements (before computing the difference) have an overall mean of 149 ms.

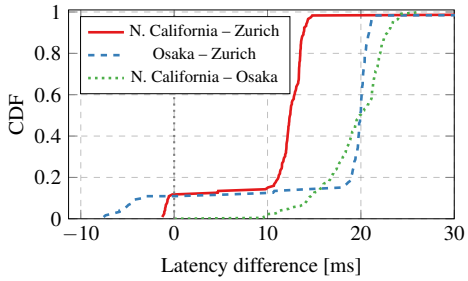

Figure 7: Difference of SBAS versus Internet latency between customers. The individual latency measurements (before computing the difference) have an overall mean of 154 ms.

fering a choice between multiple paths, SCION enables applications to choose optimal paths based on various metrics, whereas BGP by design always selects a single path between each ⟨source, destination⟩ pair (which in many cases could result in sub-optimal latency [47]).

We measure latency between all pairs of PoPs across both SCIONLab (used in our SBAS deployment) and the Internet. These measurements are compared in Figure 6. For the Internet latency baseline, we used different measurement methods to simulate real-world traffic, eliminating protocol-specific factors: echo requests from the ICMP protocol, and TCP handshakes. The SBAS latency is measured over SCION. Note that, although the same packet generator is used for the Internet latency and SBAS latency measurements, in SBAS it appears to the data center network as generic UDP traffic to locally deliver SCION packets. We observed that the latency is consistently lower when measured using ICMP as opposed to the TCP-based measurements. Following the methodology of Kwon et al. [47], we use TCP as the point of comparison for SCION latency.

The results in Figure 6 show that **SCIONLab achieves lower latency across PoP nodes than the Internet for approximately 60% of the measurements**, despite consisting largely of overlay links. The improvement stems from the sophisticated path control that SCION provides; SCION can steer packets through latency-optimized paths (e.g., Tokyo-Singapore-Frankfurt) while BGP selects a detour path (e.g., Tokyo-Seattle-Frankfurt). This indicates the potential for improved latency using an inter-domain routing architecture like SCION, which is able to leverage its advantage even in this setting with relatively few path choices, and can compensate for the overhead of tunneling to create overlay links.

6.2 End-to-End Latency

The final latency experiment evaluates the performance achieved for end-to-end connections between two customer hosts that communicate across SBAS. To evaluate this performance, we ran one round of the latency measurement between all pairs of customer machines in our testbed every hour over the span of two weeks (336 rounds in total). In each round, we ran ping 30 times with a one-second interval and computed the averaged latency of ping packets. We picked three customer locations across multiple continents and chose the closest PoP as the ingress for each customer. Therefore, these measurements are expected to compare less favorably to the Internet baseline than the latency over SBAS, as traffic may need to take a detour from the source over the nearest ingress PoP and from the egress PoP to the destination. The results shown in Figure 7 confirm this hypothesis and align with the measured latency from customers in the testbed to the nearest PoP, which is displayed in Table 1. **On average, end-to-end latency over SBAS is approximately 17 ms higher than over the Internet.** The end-to-end differences between SBAS and the Internet are relatively minor when the large intercontinental latencies are taken into account: for instance, latency between Zurich and Osaka has a mean of 241 ms over the Internet and 259 ms over SBAS, which amounts to a relative increase of less than 7.5%.

In some scenarios, the latency improvements offered by SCIONLab even enable end-to-end connections to achieve better latency than the Internet. The large variance in this difference over time can be attributed to the instability of BGP routes, which change frequently over time.

By using an expanded network of PoPs (reducing the latency from Table 1) as well as dedicated SCION links in the backbone network (as mentioned in the previous section), SBAS's end-to-end latency can be improved further.

6.3 Discussion: Scalability

Several aspects of SBAS's design enable it to scale to a large, real-world deployment. The bandwidth and compute-memory expense of SBAS is roughly proportional to the number of its clients. The PoP module (which includes the WireGuard tunnel, SIG, and routing logic) maintains no per-flow state. Adding a client to SBAS involves only adding a few internal routing table entries and the client's IP/VPN key to the PoP WireGuard configuration. SBAS capacity can be scaled to

Customer	PoP	Latency [ms]
Zurich	Frankfurt	12.36
N. California	Oregon	25.38
Osaka	Tokyo	12.75

Table 1: Latency from customers to the respective closest PoP.

serve additional customers by increasing the computational capacity of the PoPs and increasing the bandwidth of the secure internal inter-PoP network. Key management and exchange between PoPs is handled by SCION, which improves upon RPKI scalability by delegating trust roots and key management to a limited number of isolation domains [62].

We recognize that the growing size of Internet routing tables is a present concern [11]. Although SBAS infrastructure announces /24 prefixes (or /48 for IPv6 infrastructure, the longest publicly routable prefix in IPv6 [9]) for sub-prefix hijack protection, the number of such announcements is overall few and proportional to the number of SBAS POPs combined with the number of connected customer ASes (as each PoP or customer needs one prefix for their WireGuard endpoint). While SBAS customers are encouraged to use /24 length-prefixes (for select security-critical services), SBAS does not disaggregate customer announcements.

7 Security Analysis

We used two primary methods to evaluate the security of SBAS: (1) real-world attacks using the PEERING BGP research framework [69], and (2) simulated attacks using Internet topology simulations [33]. These two methods are intended to complement each other. The PEERING framework allows us to launch ethical BGP attacks against the real SBAS prototype deployment and accurately captures the dynamics of Internet routing, but does not let us experiment with many different adversary locations (as we are limited by the nodes of the PEERING framework). Topology simulations allow us to experiment with a large number of adversary locations but yield less accurate results [41]. However, when considered side-by-side, the results of these two types of evaluations complement each other, allowing for a more accurate understanding of the security of SBAS.

Recall from the threat model (Section 3.2) that we focus our evaluation on equally-specific prefix attacks because more-specific prefix attacks can be avoided by providing services through disaggregated IP prefixes, which we use for all SBAS-related addresses. We do additionally consider more-specific prefix attacks in Section 7.1.

7.1 Evaluating SBAS Deployment Against Ethical Real-World BGP Attacks

We launched real-world BGP attacks against our SBAS deployment using the PEERING framework. First, we set up two PEERING locations (known as *muxes*) to act as SBAS customers. One customer location was the mux utah01 located at the University of Utah. The other one was the mux grnet01 located in the Greek education and research network GRNET. Next, we used the mux neu01 at North Eastern University to serve as an adversary and launch BGP attacks.

Attack target	No SBAS	Using SBAS
utah01 prefix	Failed	Failed
grnet01 prefix	**Succeeded**	Failed
utah01 tunnel	N/A	Failed
grnet01 tunnel	N/A	Failed

Table 2: Summary of ethically conducted real-world attack experiments (via PEERING) from an adversary at neu01.

Ethical considerations. An important ethical principle underlying our experiment setup is that we only made BGP announcements for prefixes that we are authorized to use. Even though we used the neu01 mux as an adversary to model our attacks, it had proper authorization from the PEERING framework for all of its BGP announcements. Further, we used IP prefixes explicitly delegated to our infrastructure (from PEERING and participating educational institutions) that hosted no production services and served no real users. Finally, we also followed the PEERING framework's acceptable use policy as to not overwhelm or crash internet routers.

Our experiments validate that SBAS can mitigate both non-adaptive attacks that target customer prefixes as well as adaptive attacks that target customer tunnels to the SBAS PoPs. Moreover, SBAS can even enhance communication security between a SBAS customer and external hosts on the Internet. Table 2 depicts a summary of our results.

Control case: Successful attack without SBAS. We started with a control case where SBAS was not used and the customers in utah01 and grnet01 announced their prefixes via their traditional BGP providers/peers. The adversary node at neu01 then attempted to hijack communication between the two customers by announcing utah01's prefix and grnet01's prefix. We found that while traffic from grnet01 to utah01 was routed successfully to utah01, **traffic from utah01 to grnet01 was routed to the adversary**. This let the adversary observe (and potentially modify) packets sent from utah01 to grnet01. Thus, the connection between utah01 and grnet01 was successfully attacked by the adversary in the absence of SBAS because utah01 had a superior BGP route to the adversary to the one it had to the victim.

Attack mitigation using SBAS. Next, we connected the utah01 and grnet01 customers to the prototype SBAS implementation using the Oregon and Frankfurt SBAS PoPs respectively. We then had utah01 and grnet01 make the announcements for their IP prefixes through SBAS. Recall that at each customer node, SBAS-learned routes are given higher priority (Section 4.5) than standard Internet routes. We consider two types of adversaries:

Non-adaptive adversary. The non-adaptive adversary is not aware of SBAS and launches BGP attacks against customers' IP prefixes as usual, as in the control case. Because both utah01 and grnet01 were communicating through SBAS, **the adversary was incapable of hijacking any of the traffic between grnet01 and utah01** in either direction of communication. Note that this result is independent of customer or adversary location. Two customers will always successfully resist BGP attacks where an adversary targets a customer's IP prefix announced through SBAS. This is due to route prioritization and holds even in the case of more-specific BGP attacks. SBAS PoPs load secure routes into a separate routing table that is given higher priority than the Internet routing table. SBAS customers' outbound traffic will always go through the connected SBAS PoPs. More-specific routes in the global routing table do not affect routes between SBAS customers.

Adaptive adversary. An adaptive adversary who is aware of SBAS may instead chose to attack the tunnels that each customer uses to communicate with SBAS.While this adaptive attack is inherently less devastating because tunnels are end-to-end encrypted, there are still powerful attacks that can be launched against encrypted traffic [59, 73].

Using our PEERING setup, we had the adversary at neu01 attack the IP prefixes used by utah01 and grnet01 to establish their WireGuard sessions with SBAS.[1] Even when the adversary maliciously announced the IP prefix of the WireGuard endpoint of both victims with an equally-specific BGP attack, **communication between utah01 and grnet01 was uninterrupted and was never routed to the adversary.** Note that SBAS infrastructure prefixes, like the one used for the VPN endpoint, are required to be /24s; more-specific attacks against SBAS VPN endpoints are not viable.

We note that the success of this type of adaptive adversary against SBAS depends highly on the customer's choice of ingress points. As a contrived example, had utah01 chosen the Frankfurt SBAS PoP as its ingress point and grnet01 chosen Oregon as its ingress point, communication along both of the tunnels would have been routed to the adversary. It is because of the proximity of the SBAS ingress point to the SBAS customer (relative to the adversary's location) that SBAS offers improved security even against this type of adaptive adversary. In the optimal case, a customer may even be able to obtain a direct layer-2 connection with a PoP, thwarting this attack entirely.

Characterizing communication security with external hosts. In addition to running experiments to measure the security of communication between SBAS customers in utah01 and grnet01, we evaluated the security benefit that SBAS offers when a SBAS customer is communicating with an non-SBAS-protected (i.e., external) host or server on the Internet. We build upon the methodology presented by Birge-Lee et al. [20]. We constructed a sample of 1k IP addresses from the Censys Internet-wide IPv4 scans [31] to serve as external hosts. The sample was chosen at random and filtered to only include hosts that responded to ICMP echo (ping) requests.

To measure the impact of SBAS on communication between a SBAS customer and external hosts, we performed ethical equally-specific prefix hijacks from the adversary neu01 against IP prefixes originated by utah01 and grnet01 and announced by SBAS (using nodes at amsterdam01 and seattle01). Similarly, we also performed ethical hijacks where no SBAS was used as the control case. Then, for each hijack, we launched a ping scan of the 1k external hosts in our sample from an IP address in the prefix that was being hijacked. When each host in the sample received the ping request, it generated a ping response with a destination IP address that was under attack by the adversary. Computing the fraction of hosts whose ping responses were routed to the adversary allowed us to measure the impact of each hijack we launched.

SBAS significantly enhances communication security with external hosts. When the utah01 and grnet01 nodes were not using SBAS, the adversary at neu01 was able to hijack traffic from 72% and 76% of the 1k external hosts respectively. When we connected utah01 and grnet01 to SBAS and launched an attack on the SBAS announced-prefix, **the adversary's hijacking capability reduced threefold to affecting only 25% of hosts.** We emphasize that the presented security improvements are conservative, as this experiment was performed against the SBAS prototype deployment that uses only the seattle01 and amsterdam01 PEERING muxes. We are actively working to expand this deployment, and present recommended expansion steps and the associated security improvement in Section 7.2.1. As more nodes are added, hosts will have a shorter route to the nearest SBAS PoP which will further reduce the spread of the adversary's attacks.

7.2 Quantifying Hijack Resilience via BGP Attack Simulations

To evaluate SBAS security beyond PEERING mux adversaries and client locations, we employ Internet-scale attack simulations. The Internet topology was constructed using the CAIDA AS topology dataset [1], augmented with peering information inferred from the bdrmap tool [51] and BGP Routing Information Base (RIB) data collected from Route Views [68] and RIPE NCC RIS [2] to correctly model route selection at the AWS datacenters and PEERING nodes that are part of the SBAS deployment. We build upon the methodology developed by Birge-Lee et al. [19] to perform prefix-level (as opposed to AS-level) simulations. We evaluate how likely traffic from external sources will still be routed to SBAS (via the PEERING framework) in the event of an equally-specific BGP attack. We also explore how security improves with more BGP-announcing SBAS nodes and with full deployment of RPKI in the broader internet.

[1] Another attack can be launched by hijacking the IP prefix of the SBAS PoPs for their WireGuard endpoint. We could not conduct this in the wild, as we were not authorized to hijack the SBAS PoPs' AWS-controlled prefix.

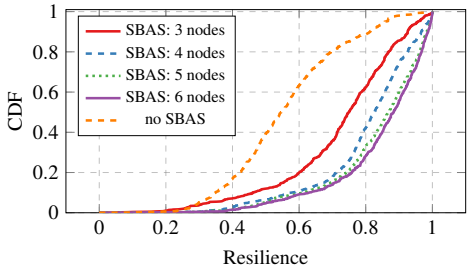

Figure 8: Cumulative distribution of SBAS resilience against (randomized) AS-level adversary.

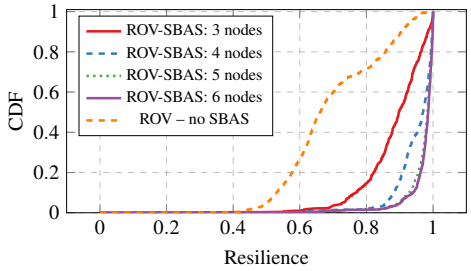

Figure 9: Cumulative distribution of SBAS resilience against (randomized) AS-level adversary, assuming adoption of ROV.

Computing attack resilience metric. We use the notion of *resilience* [79] to quantify the fraction of total potential adversaries that are topologically unable to launch an equally-specific prefix hijacking attack for an arbitrary prefix announced by a victim AS. The detailed definition of resilience is in Appendix B. Resilience is affected by the relative locations of the victim and adversary ASes, their peering and provider relationships, and the application of further security measures, including RPKI.

Internet topology simulations. We run BGP simulations against an AS-level adversary considering a random sample of 1k adversary ASes as the attacker set $\mathcal{A}$, corresponding to approximately 1.39% of the $N = 71669$ ASes profiled in the CAIDA AS topology. Against each attacker AS, we consider all other $N-1$ ASes as the set of external hosts (traffic sources) $\mathcal{B}$. Given these fixed sets of adversary ASes and external hosts, we run BGP simulations for two scenarios: (1) a victim prefix is announced via SBAS BGP announcement nodes; (2) a victim prefix is announced in a conventional manner without SBAS. We consider a random sample of 1k victim prefixes, which are selected based on the methodology in Section 7.1. We evaluate SBAS configurations with varying number and location of BGP announcement nodes. When varying the number of BGP-speaking SBAS nodes, we present results for configurations (node locations) that are globally optimal (more details in Appendix B).

7.2.1 SBAS Significantly Enhances Resilience Across Adversary ASes and Customer Locations

We analyze the distribution of prefix-level resilience of routes announced by SBAS against sampled AS-level adversaries performing BGP hijack attacks and compare them to the scenario where SBAS is not used (Figure 8). Our results show that SBAS deployment significantly improves routing security across adversary ASes and customer locations even with a small number of BGP announcement nodes. Increasing the number of announcement nodes in the SBAS backbone further enhances security.

Figure 8 shows that **a conservative deployment of SBAS with just 5–6 BGP announcement nodes can lead to more than a 60% improvement in median resilience.** A three-node SBAS with announcements at Amsterdam, Seattle, and ISI yielded a resilience of 0.750, a 37.2% improvement compared to the baseline median resilience of 0.545 (without SBAS). Including another announcement node at GRNet further increased median resilience to 0.825, a 50.9% improvement. Additional announcement nodes provide diminishing security returns: a fifth node at UWashington, to 0.870 (59.3% improvement), and finally a sixth node at GATech, to 0.884 (61.8% improvement).

Resilience improvements against known serial hijacking attack ASes. As an illustrative case study, we examine the threat of BGP attacks launched by a sample of 11 ASes identified as serial hijacking offenders in prior work [74]. This smaller sample of attackers allows us to plot resilience for SBAS customers communicating with external hosts for each attacker AS. Notably, **SBAS improved resilience against all 11 of the known serial hijacker ASes, with a median resilience improvement of 64.9%** using a configuration of 6 BGP announcement nodes. For the most devastating adversary, AS 9009, the 6-node SBAS produced a median resilience gain of 602.6%. (Plotted results and further discussion are included in Appendix B.)

7.2.2 ROV Enforcement in the Broader Internet Can Further Boost Resilience Offered by SBAS

We further examine the extent of SBAS's potential security improvements by considering the adoption of ROV (Route Origin Validation using RPKI) in the broader (non-SBAS) Internet. Recall that ROV makes equally-specific hijacks less likely to succeed by verifying the route's origin AS: an attacker thus needs to prepend itself to a route originating from the valid RPKI-signed AS, which increases the attacker's path length by one. We repeat the prior simulation setup with the addition of ROV in the broader internet to evaluate the extent of resilience improvement provided by a combination of SBAS and ROV (see Figure 9).

We use *ROV-aided* to refer to the scenario in which ROV is deployed by the broader internet. ROV substantially improves

resilience offered by SBAS: on average, **a ROV-aided SBAS deployment increased median resilience over its non-ROV counterpart by 45.0%.** For example, a ROV-aided SBAS deployment with 3 announcement nodes improves median resilience up to 0.898, a 35.7% improvement compared to the baseline resilience of 0.661 without SBAS. Similarly, **a ROV-aided SBAS deployment with 6 announcement nodes improves resilience to 0.985 (a 49.0% improvement), meaning 98.5% of adversaries were topologically incapable of hijacking SBAS-announced routes.**

8 Incentives and Governance

Beyond the setup of the technical SBAS components, adoption of SBAS would require coordination between participating ISPs, formation of organizations to handle governance, and presenting/marketing SBAS to customers. To better understand network operator's SBAS deployment incentives and preference for governance models, we have conducted a survey. In Sections 8.1 and 8.2, we discuss incentives and possible models of governance for an SBAS deployment, followed by a discussion of survey results in Section 8.3.

8.1 Deployment Incentives

Current market trends demonstrate demand for reliable and secure Internet connectivity. SD-WAN, leased lines, and Network-as-a-service products specifically designed to mitigate routing outages have seen widespread adoption for businesses in various sectors [4]. The willingness to purchase these services despite substantial costs indicates that customers are willing to pay a premium to protect against routing-induced network outages. A candidate first customer may have incentives for SBAS's security properties that outweigh the difficulties inevitable in early-stage technology deployment, similar to the initial customer of the SCION network [45]. The added cost of deploying SBAS is marginal if the infrastructure already supports SCION connectivity (currently natively supported by 10 ISPs), but the additional customer base that can be reached can provide major financial benefits. Our evaluation in Section 7.2 has shown that 5 SBAS PoPs can already provide immediate security benefits to the first customer. The current SCION-supporting ISPs would thus suffice for bootstrapping SBAS.

We believe SBAS's lightweight implementation will also help it gain early adoption. Several survey responses (Section 8.3) emphasize the necessity for interoperability with current routing hardware and protocols, with minimum effect on operational robustness as requisite for industry adoption of any new routing security solution. Compared to BGPsec and other proposed clean-slate routing protocols, SBAS uses commodity network hardware and does not suffer from all-or-nothing deployment security improvements.[2] SBAS's use of reliable, off-the-shelf networking components (BGP, iBGP, and secure tunnels) reduces most of the effort required to implement and maintain custom routing modules. This relative ease of implementation translates to lower transition costs for customers and more rapid experimentation under real-world traffic flow conditions.

8.2 Governance Models

Due to the federated nature of the SBAS PoP operation, a governance structure is needed to coordinate global operation (e.g., AS management, RPKI ROA distribution, and coordination of secure and internal address ranges). We present four different governance models that all received support in our survey, presented in the order of their degree of centralization and reliance on existing structures.

Scenario 1: ICANN and regional Internet registries. The regional Internet registries (RIRs) already play a major role in coordinating the control plane of the Internet, e.g., by allocating IP ranges, AS numbers, and providing hosted RPKI services. They would therefore be natural entities to govern a shared AS number for SBAS. Such a governance model would also benefit from the strong ties between the RIRs and the network providers. However, albeit they provide coordination activities and services to their members [6, 10, 13, 48, 67], the RIRs do not cover operation of network infrastructure. Operating SBAS would be orthogonal to other efforts by RIRs to improve routing security.

Scenario 2: Multi-stakeholder organization. Under this governance model, a foundation involving interested parties such as ISPs and companies would run SBAS. This would provide the benefit of creating an entity with a clear scope of duties with regards to SBAS, entirely dedicated to guarantee the smooth operation of SBAS, and which could also receive dedicated contributions towards that effort. On the other hand, this would require new structures to be set up.

Scenario 3: Federation of network providers. A governance model relying on the initiative of ISPs working together to join (some of) their resources in SBAS, through a loose coordination at the technical level between the involved parties, building on letters of intent and bilateral agreements. In this governance model, the network effect is less noticeable and the early participants would bear the bulk of the burden of driving the adoption of the initiative.

Scenario 4: Decentralized governance model. Each PoP operator can join SBAS independently, in the same manner as there is no centralized instance governing which TOR nodes can join the network [65]. This model is most flexible for PoP operators, with low barriers to entry for new operators. However the continuity of operation of a sufficient number of PoPs is not guaranteed, and sharing scarce resources such as IP address space and AS numbers would be challenging.

[2] Although the BGPsec signature validation implementation itself can be incrementally deployable, it requires every hop in the routing path to sign its path segments to achieve the desired security properties.

Governance model recommendation. Based on the results of our survey (presented in Section 8.3) and the structure of SBAS, we suggest the federation of network providers. Not only did this structure receive the most votes in our survey, but it has the benefit of placing the governance decisions in the hands of ISPs, who are the ones responsible for running SBAS PoPs and carrying SBAS traffic. Furthermore, it does not require the involvement of any overly powerful or centralized organization, which many operators in our survey expressed concerns over. While issues like IP, domain, and ASN allocation inherently involve distribution of a shared resource, governance of SBAS (which operates out of SCION and IP address already controlled by participants) primarily involves technical and policy coordination that can be done in a more decentralized manner. We feel that in the same way the MANRS project [53] has been successful in bringing together ISPs to standardize routing security practices, a federation of ISPs could also be successful in standardizing SBAS operational practices. Additionally, in a real-world operation of SBAS, the federation of providers may establish an association or foundation to create a more concrete structure to govern, organize, and enforce the operation of all participants (incorporating structures of Scenario 2).

8.3 Survey of Network Operators

Survey participants were recruited through direct contact, and through the RIPE and NANOG mailing lists [60, 61], garnering 31 responses. We summarize important results in this section, and discusses detailed survey results in Appendix C.

Our survey indicates that network security is very important to the operations of the majority of ISPs and there is a community of early adopters that would be interested in deploying SBAS. Specifically, 26 out of 31 (84%) participants said, on a scale of zero to five, that the importance of network security to their ISP's operations was either a four or a five. However, when asked if secure routing was a marketable product, only 15 operators (of 31) responded with a four or five. This gap in responses can be attributed to the usually assumed network effect problem where a large critical mass of participants is needed for substantial security benefits. In these cases, early incentives are insufficient for early stakeholders to undergo the costs of building a new system. This is a major reason why other solutions that require high usages rates to yield security improvements fail to see deployment (as is the case with BGPsec [52]). Even with substantial interest in the community, many network operators do not see the financial benefits of investing in secure routing, creating a self-fulfilling prophecy of low adoption rates.

We proceeded to ask operators to gauge the interest of their ISP offering SBAS to customers and eight operators rated it as a four or five. Furthermore, this group of interested operators seemed strongly convinced of the SBAS design: they reported a mean score of 3.75 points on SBAS's deployment feasibility and a mean score of 4.6 on the effectiveness of SBAS against routing attacks. The chief concerns for deployment cited were mostly related to logistics, such as establishing inter-ISP iBGP sessions and several ISPs sharing an ASN in the routing system. While this is a minority of survey participants, it (1) represents 53% of operators who felt secure routing was a marketable productand (2) shows there is a non-trivial group of convinced, interested early-adopters that would enable substantial security improvements through SBAS. Even with a few early adopters, communication between SBAS participants and the broader Internet achieves substantially higher resilience (e.g., 5 PoPs in Section 7.2).

Among the governance models, the federated model was chosen by 14 participants (≈45%) as the most popular potential governance structure, followed by delegating responsibilities over to the RIRs (35%), a decentralized model (13%), and a multi-stakeholder organization (7%). Several respondents stressed the importance of selecting a structure that would not be dominated by large corporations and with mechanisms to prevent it from growing beyond its needed scope.

By surveying network operators, we see some next steps required for a production deployment of SBAS. We encourage network operators and the research community to work collaboratively to establish SBAS as a production network.

9 Related Work

To improve on the limitations of BGP, various alternatives have been suggested, including studies that use overlay technology to establish new routes [12, 35, 38, 83]. Particularly, Andersen et al. propose RON (Resilient Overlay Network), an architecture that constructs an overlay network using distributed applications, monitors the underlay routing paths in real-time, and constructs new paths [7]. Peter et al. propose the ARROW architecture, which flattens the Internet topology using overlay tunnels between ISPs, and provides a new route if needed [63]. Compared to SBAS, ARROW focuses on availability and only addresses use cases in which customers are fully participating. Network pluralism articulates the need for architectural heterogeneity [22, 44, 76]. Crowcroft et al. introduce Plutarch, which describes each homogeneous network architecture as context and enables communication across a set of contexts by interstitial functions that interpret the encapsulated functionalities of each context [29]. Avramopoulos and Rexford present a security backbone framework connecting various secure routing architectures via a secure mesh of virtual links [16]. Indeed, network pluralism enables the graceful coexistence of diverse network architectures. However, the approaches simply glue network architectures together, supporting them to only *survive*. In contrast, SBAS not only bridges secure routing infrastructures to the Internet in a synergistic manner, but also extends the benefits to the broader Internet, enhancing them to *thrive*.

10 Conclusion

While secure routing enjoyed much attention from the research community over the past two decades, real-world adoption has been lagging, perhaps due to the significant infrastructure changes required. With an ambition to make rapid progress to secure routing, we investigate how to leverage a secure communication backbone to secure communication on the regular Internet. We design and deploy an architecture, SBAS, in which communication between traditional IP endpoints are mediated via a secure backbone that is operated in a federated manner. SBAS substantially reduces the threat of inter-domain routing attacks and only incurs a small latency overhead (and as our results show can even speed up some end-to-end connections compared to the Internet). A core contribution of this work is the incentive-compatible design. SBAS does not compete with other secure routing architectures, but instead demonstrates that an existing secure routing infrastructure with limited deployment can already benefit the rest of the Internet. While several challenges still exist when deploying SBAS in a production setting, our survey shows a potential path forward and our experimental results show promise that sizable security improvements can be achieved with even a small set of early adopters. We hope that SBAS revitalizes the quest for secure inter-domain routing.

Acknowledgements

This work was supported in part by the National Science Foundation under grants CNS-1553437, CNS-1704105, and CNS-193596 and by the United States Air Force and DARPA under Contract No. FA8750-19-C-0079. Additionally, we gratefully acknowledge support from ETH Zurich, from the ETH4D and EPFL EssentialTech Centre Humanitarian Action Challenge Grant, and from the Zurich Information Security and Privacy Center (ZISC). We would like to thank the anonymous reviewers and our shepherd Deepak Kumar for their valuable feedback as well as David Hausheer, Nicola Rustignoli, Kyveli Mavromati, and the anonymous participants of our survey for their contributions to this project. Any opinions, findings and conclusions or recommendations expressed in this material are those of the authors and do not necessarily reflect the views of the United States Air Force, DARPA, or any other sponsoring agency.

References

[1] The CAIDA AS relationships dataset. https://www.caida.org/catalog/datasets/as-relationships/, 2021.

[2] RIS Raw Data – RIPE Network Coordination Centre. https://www.ripe.net/analyse/internet-measurements/routing-information-service-ris/ris-raw-data, 2021.

[3] The BIRD Internet Routing Daemon Project. https://bird.network.cz/, 2021.

[4] P. 81. The Rise of Network as a Service. https://www.perimeter81.com/blog/network/the-rise-of-network-as-a-service, 2021.

[5] E. Aben. Propagation of Longer-than-/24 IPv4 Prefixes. https://labs.ripe.net/author/emileaben/propagation-of-longer-than-24-ipv4-prefixes/, 2014.

[6] AFRINIC. AFRINIC Bylaws (Constitution) 2020. https://afrinic.net/bylaws, 2020.

[7] D. Andersen, H. Balakrishnan, F. Kaashoek, and R. Morris. Resilient Overlay Networks. In *ACM Symposium on Operating Systems Principles (SOSP)*, 2001.

[8] T. Anderson, K. Birman, R. Broberg, M. Caesar, D. Comer, C. Cotton, M. J. Freedman, A. Haeberlen, Z. G. Ives, A. Krishnamurthy, W. Lehr, B. T. Loo, D. Mazières, J. M. Nicolosi, Antonio Smith, I. Stoica, R. v. Renesse, M. Walfish, H. Weatherspoon, and C. S. Yoo. The Nebula Future Internet Architecture. In *The Future Internet Assembly*, 2013.

[9] APNIC. Why is a /48 the recommended minimum prefix size for routing? https://blog.apnic.net/2020/06/01/why-is-a-48-the-recommended-minimum-prefix-size-for-routing/, 2020.

[10] APNIC. Policy SIG. https://www.apnic.net/community/policy/policy-sig/, 2021.

[11] APNIC. What will happen when the routing table hits 1024k? https://blog.apnic.net/2021/03/03/what-will-happen-when-the-routing-table-hits-1024k/, 2021.

[12] M. Apostolaki, G. Marti, J. Müller, and L. Vanbever. SABRE: Protecting Bitcoin against Routing Attacks. In *Network and Distributed System Security Symposium (NDSS)*, 2019.

[13] ARIN. Charter for the ARIN Board Governance Working Group. https://www.arin.net/about/welcome/board/committees/charters/#charter-for-the-arin-board-governance-working-group, 2021.

[14] A. Arnbak and S. Goldberg. Loopholes for Circumventing the Constitution: Unrestricted Bulk Surveillance on Americans by Collecting Network Traffic Abroad. *Michigan Telecommunications & Technology Law Review*, 21:317, 2014.

[15] R. Austein, S. Bellovin, R. Housley, S. Kent, W. Kumari, D. Montgomery, C. Morrow, S. Murphy, K. Patel, J. Scudder, S. Weiler, M. Lepinski, and K. Sriram. BGPsec Protocol Specification. RFC 8205, 2017.

[16] I. Avramopoulos and J. Rexford. A Pluralist Approach to Interdomain Communication Security. In *Workshop on the Economics of Networks, Systems and Computation (NetEcon)*, 2007.

[17] AWS. Bring your own IP addresses (BYOIP) in Amazon EC2. https://docs.aws.amazon.com/AWSEC2/latest/UserGuide/ec2-byoip.html, 2021.

[18] H. Birge-Lee, Y. Sun, A. Edmundson, J. Rexford, and P. Mittal. Bamboozling Certificate Authorities with BGP. In *USENIX Security Symposium*, 2018.

[19] H. Birge-Lee, L. Wang, D. McCarney, R. Shoemaker, J. Rexford, and P. Mittal. Experiences Deploying Multi-Vantage-Point Domain Validation at Let's Encrypt. In *USENIX Security Symposium*, 2021.

[20] H. Birge-Lee, L. Wang, J. Rexford, and P. Mittal. SICO: Surgical Interception Attacks by Manipulating BGP Communities. In *ACM SIGSAC Conference on Computer and Communications Security (CCS)*, 2019.

[21] L. Blunk, M. Karir, and C. Labovitz. Multi-threaded routing toolkit (mrt) routing information export format. RFC 6396, RFC Editor, October 2011.

[22] G. Bouabene, C. Jelger, C. Tschudin, S. Schmid, A. Keller, and M. May. The Autonomic Network Architecture (ANA). *IEEE Journal on Selected Areas in Communications*, 28(1):4–14, 2009.

[23] R. Brandom. Hackers Emptied Ethereum Wallets by Breaking the Basic Infrastructure of the Internet. The Verge, 2018.

[24] R. Bush and R. Austein. The Resource Public Key Infrastructure (RPKI) to Router Protocol. RFC 6810, 2013.

[25] K. Butler, T. R. Farley, P. McDaniel, and J. Rexford. A Survey of BGP Security Issues and Solutions. *Proceedings of the IEEE*, 98(1):100–122, 2009.

[26] J. M. Camacho, A. García-Martínez, M. Bagnulo, and F. Valera. BGP-XM: BGP Extended Multipath for Transit Autonomous Systems. *Computer Networks*, 57(4):954–975, 2013.

[27] Cloudflare. Argo Smart Routing. https://www.cloudflare.com/products/argo-smart-routing/, 2021.

[28] A. Cohen, Y. Gilad, A. Herzberg, and M. Schapira. Jumpstarting BGP Security with Path-End Validation. In *ACM SIGCOMM Conference*, 2016.

[29] J. Crowcroft, S. Hand, R. Mortier, T. Roscoe, and A. Warfield. Plutarch: An Argument for Network Pluralism. *ACM SIGCOMM Computer Communication Review (CCR)*, 33(4):258–266, 2003.

[30] DPDK Project. Data Plane Development Kit. https://dpdk.org.

[31] Z. Durumeric, D. Adrian, A. Mirian, M. Bailey, and J. A. Halderman. A Search Engine Backed by Internet-Wide Scanning. In *ACM Conference on Computer and Communications Security (CCS)*, 2015.

[32] W. E. Forum. Footprints on the path: how routing data could reduce the Internet's carbon toll. https://www.weforum.org/agenda/2021/03/internet-carbon-emissions-data-path-scion/, 2021.

[33] L. Gao and J. Rexford. Stable Internet Routing without Global Coordination. *IEEE/ACM Transactions on Networking*, 9(6):681–692, 2001.

[34] A. Gavrichenkov. Breaking HTTPS with BGP Hijacking. *Black Hat*, 2015.

[35] P. B. Godfrey, I. Ganichev, S. Shenker, and I. Stoica. Pathlet Routing. *ACM SIGCOMM Computer Communication Review (CCR)*, 39(4):111–122, 2009.

[36] S. Goldberg, M. Schapira, P. Hummon, and J. Rexford. How Secure are Secure Interdomain Routing Protocols? *ACM SIGCOMM Computer Communication Review (CCR)*, 40(4):87–98, 2010.

[37] D. Goodin. How 3ve's BGP Hijackers Eluded the Internet – and Made $29M. Ars Technica, 2018.

[38] P. K. Gummadi, H. V. Madhyastha, S. D. Gribble, H. M. Levy, and D. Wetherall. Improving the Reliability of Internet Paths with One-hop Source Routing. In *USENIX Symposium on Operating Systems Design and Implementation (OSDI)*, 2004.

[39] S. Hanks, T. Li, D. Farinacci, and P. Traina. Generic Routing Encapsulation (GRE). RFC 1701, 1994.

[40] G. Huston. CIDR REPORT for 31 Jan 22. https://www.cidr-report.org/as2.0/, 2022.

[41] J. Juen, A. Johnson, A. Das, N. Borisov, and M. Caesar. Defending Tor from Network Adversaries: A Case Study of Network Path Prediction. *Proceedings on Privacy Enhancing Technologies*, 2015(2):171–187, 2015.

[42] J. Karlin, S. Forrest, and J. Rexford. Pretty Good BGP: Improving BGP by Cautiously Adopting Routes. In *IEEE International Conference on Network Protocols (ICNP)*, 2006.

[43] S. Kent, C. Lynn, and K. Seo. Secure Border Gateway Protocol (S-BGP). *IEEE Journal on Selected areas in Communications*, 18(4):582–592, 2000.

[44] T. Koponen, S. Shenker, H. Balakrishnan, N. Feamster, I. Ganichev, A. Ghodsi, P. B. Godfrey, N. McKeown, G. Parulkar, B. Raghavan, J. Rexford, S. Arianfar, and D. Kuptsov. Architecting for Innovation. *ACM SIGCOMM Computer Communication Review (CCR)*, 41(3):24–36, 2011.

[45] C. Krähenbühl, S. Tabaeiaghdaei, C. Gloor, J. Kwon, A. Perrig, D. Hausheer, and D. Roos. Deployment and scalability of an inter-domain multi-path routing infrastructure. In *Proceedings of the International Conference on Emerging Networking EXperiments and Technologies (CoNEXT)*, 2021.

[46] N. Kushman, S. Kandula, D. Katabi, and B. M. Maggs. R-BGP: Staying Connected in a Connected World. In *Network and Distributed Systems Security Symposium (NDSS)*, 2007.

[47] J. Kwon, J. A. García-Pardo, M. Legner, F. Wirz, M. Frei, D. Hausheer, and A. Perrig. SCIONLAB: A Next-Generation Internet Testbed. In *IEEE International Conference on Network Protocols (ICNP)*, 2020.

[48] LACNIC. Bylaws. https://www.lacnic.net/76/2/lacnic/bylaws, 2018.

[49] M. Lepinski and S. Kent. An Infrastructure to Support Secure Internet Routing. RFC 6480, 2012.

[50] Y. Liao, L. Gao, R. Guerin, and Z.-L. Zhang. Reliable Interdomain Routing through Multiple Complementary Routing Processes. In *ACM Conference on Emerging Networking Experiments and Technologies (CoNEXT)*, 2008.

[51] M. Luckie, A. Dhamdhere, B. Huffaker, D. Clark, and k. claffy. bdrmap: Inference of Borders Between IP Networks. In *ACM Internet Measurement Conference (IMC)*, 2016.

[52] R. Lychev, S. Goldberg, and M. Schapira. BGP Security in Partial Deployment: Is the Juice Worth the Squeeze? In *ACM SIGCOMM Conference*, 2013.

[53] Mutually Agreed Norms for Routing Security. https://www.manrs.org/.

[54] A. Maria, Z. Aviv, and V. Laurent. Hijacking Bitcoin: Routing Attacks on Cryptocurrencies. In *IEEE Symposium on Security and Privacy (S&P)*, 2017.

[55] A. Mitseva, A. Panchenko, and T. Engel. The State of Affairs in BGP Security: A Survey of Attacks and Defenses. *Computer Communications*, 124:45–60, 2018.

[56] M. Motiwala, M. Elmore, N. Feamster, and S. Vempala. Path Splicing. In *ACM SIGCOMM Conference*, 2008.

[57] D. Naylor, M. K. Mukerjee, P. Agyapong, R. Grandl, R. Kang, M. Machado, S. Brown, C. Doucette, H.-C. Hsiao, D. Han, T. H.-J. Kim, H. Lim, C. Ovon, D. Zhou, S. B. Lee, Y.-H. Lin, C. Stuart, D. Barrett, A. Akella, D. Andersen, J. Byers, L. Dabbish, M. Kaminsky, S. Kiesler, J. Peha, A. Perrig, S. Seshan, M. Sirbu, and P. Steenkiste. XIA: Architecting a More Trustworthy and Evolvable Internet. *ACM SIGCOMM Computer Communication Review (CCR)*, 44(3):50–57, 2014.

[58] O. Nordström and C. Dovrolis. Beware of BGP Attacks. *ACM SIGCOMM Computer Communication Review (CCR)*, 34(2):1–8, 2004.

[59] A. Panchenko, L. Niessen, A. Zinnen, and T. Engel. Website Fingerprinting in Onion Routing based Anonymization Networks. In *ACM Workshop on Privacy in the Electronic Society (WPES)*, 2011.

[60] A. Perrig. Incrementally deployable secure Internet routing: operator survey. https://seclists.org/nanog/2021/Dec/202, 2021. NANOG mailing list.

[61] A. Perrig. Incrementally deployable secure Internet routing: operator survey. https://www.ripe.net/ripe/mail/archives/ripe-list/2021-December/002400.html, 2021. RIPE mailing list.

[62] A. Perrig, P. Szalachowski, R. M. Reischuk, and L. Chuat. *SCION: A Secure Internet Architecture*. Springer Verlag, 2017.

[63] S. Peter, U. Javed, Q. Zhang, D. Woos, T. Anderson, and A. Krishnamurthy. One Tunnel is (Often) Enough. *ACM SIGCOMM Computer Communication Review (CCR)*, 44(4):99–110, 2014.

[64] A. Pilosov and T. Kapela. Stealing the Internet: An Internet-scale Man in the Middle Attack. *NANOG 44*, 2008.

[65] T. T. Project. Relay operations: Technical Setup. https://community.torproject.org/relay/setup/, 2021.

[66] J. Rexford and C. Dovrolis. Future Internet Architecture: Clean-slate versus Evolutionary Research. *Communications of the ACM*, 53(9):36–40, 2010.

[67] RIPE. RIPE NCC Charter. https://www.ripe.net/about-us/what-we-do/ripe-ncc-charter, 2016.

[68] University of Oregon Route Views Project. http://www.routeviews.org/routeviews/.

[69] B. Schlinker, T. Arnold, I. Cunha, and E. Katz-Bassett. PEERING: Virtualizing BGP at the Edge for Research. In *ACM Conference on Emerging Networking Experiments and Technologies (CoNEXT)*, 2019.

[70] B. Schlinker, K. Zarifis, I. Cunha, N. Feamster, and E. Katz-Bassett. PEERING: An AS for Us. In *ACM Workshop on Hot Topics in Networks (HotNets)*, 2014.

[71] J. Snijders. Practical everyday BGP filtering with AS_PATH filters: Peer locking. *NANOG-67*, 2016.

[72] Y. Sun, M. Apostolaki, H. Birge-Lee, L. Vanbever, J. Rexford, M. Chiang, and P. Mittal. Securing Internet Applications from Routing Attacks. *Communications of the ACM*, 64(6):86–96, 2021.

[73] Y. Sun, A. Edmundson, L. Vanbever, O. Li, J. Rexford, M. Chiang, and P. Mittal. RAPTOR: Routing Attacks on Privacy in Tor. In *USENIX Security Symposium*, 2015.

[74] C. Testart, P. Richter, A. King, A. Dainotti, and D. Clark. Profiling BGP Serial Hijackers: Capturing Persistent Misbehavior in the Global Routing Table. In *ACM Internet Measurement Conference (IMC)*, 2019.

[75] A. Toonk. The Canadian Bitcoin Hijack. https://www.bgpmon.net/the-canadian-bitcoin-hijack/.

[76] J. S. Turner and D. E. Taylor. Diversifying the Internet. In *IEEE Global Telecommunications Conference (GLOBECOM)*, 2005.

[77] T. Wan, E. Kranakis, and P. C. van Oorschot. Pretty Secure BGP, psBGP. In *Network and Distributed System Security Symposium (NDSS)*, 2005.

[78] F. Wang and L. Gao. Path Diversity Aware Interdomain Routing. In *IEEE International Conference on Computer Communications (INFOCOM)*, 2009.

[79] M. Wübbeling and M. Meier. Improved Calculation of AS Resilience against IP Prefix Hijacking. In *IEEE Conference on Local Computer Networks Workshops (LCN Workshops)*, 2016.

[80] D. Xu, M. Chiang, and J. Rexford. DEFT: Distributed Exponentially-weighted Flow Splitting. In *IEEE International Conference on Computer Communications (INFOCOM)*, 2007.

[81] W. Xu and J. Rexford. MIRO: Multi-path Interdomain Routing. In *ACM SIGCOMM Conference*, 2006.

[82] X. Yang, D. Clark, and A. W. Berger. NIRA: A New Inter-domain Routing Architecture. *IEEE/ACM Transactions on Networking*, 15(4):775–788, 2007.

[83] X. Yang and D. Wetherall. Source Selectable Path Diversity via Routing Deflections. *ACM SIGCOMM Computer Communication Review (CCR)*, 36(4):159–170, 2006.

[84] X. Zhang, H.-C. Hsiao, G. Hasker, H. Chan, A. Perrig, and D. Andersen. SCION: Scalability, Control, and Isolation On Next-Generation Networks. In *IEEE Symposium on Security and Privacy (S&P)*, 2011.

A Overhead of SBAS Components at PoPs

In comparison to regular packet forwarding, an SBAS needs to perform some additional computational steps, such as decryption from the customer VPN tunnel and encapsulation to remote PoPs. This first benchmark measures the overhead incurred by the SBAS components on packets passing through a PoP. To obtain an upper bound on this overhead, we consider communication between a pair of customer hosts. This is the most computationally intensive scenario, as traffic is sent through VPN tunnels on both sides.

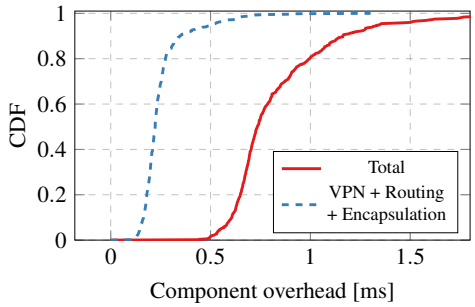

Figure 10: Overhead incurred by SBAS components at a PoP, measured between the packet arrival from the customer to the sending operation of the SCION packet through SBAS. In the dashed data series, the final SCION step is omitted.

We run ping from the source customer to the destination host with a 1 s interval for an hour. The overhead is determined by subtracting timestamps of packets captured at the SIG of the PoP from the corresponding timestamps of the same packets captured at the WireGuard interface (after VPN decapsulation, route selection, encapsulation, and tunneling over SCION). This yields the following results.

The majority of packets have sub-millisecond overhead at PoPs, with a mean of 0.83 ms and a standard deviation of 0.27 ms. The SBAS component delay is also invariant to packet size: repeated experiments with 1KB packets reported a mean overhead of only 0.75 ms. We attribute the tail end of the distribution to operating system factors such as process scheduling and resource contention. Recall that the SBAS components only perform en/decapsulation and routing of packets, exhibiting indistinguishable overhead for different protocols. It is also important to note that, since the current prototype is a software-based implementation, the processing overhead can be further minimized with a production-grade implementation, e.g., using the Data Plane Development Kit (DPDK) [30] or hardware accelerators.

Figure 10 summarizes the latency expense incurred by SBAS PoP component in sending a packet through the backbone. Moving from an overlay-based network like SCIONLab to a native SCION network would reduce the SBAS overhead further, as the outermost layer of encapsulation (SCION in IP) would not be required anymore. Moreover, a production-grade implementation of SCION could be used that performs better than the open-source research prototype. The secondary measurement indicated by the dashed line (which omits the SCION latency) in Figure 10 provides a lower bound estimate on the potential SBAS backbone latency: approximately 70% of the median 0.74 ms latency can be attributed to SCION latency, suggesting that the overhead of the SBAS-specific infrastructure is in fact relatively light.

Since our technology is applied to inter-domain traffic, this additional latency is negligible relative to the propagation delay over larger geographical distances, which is often on the order of 100 ms and higher.

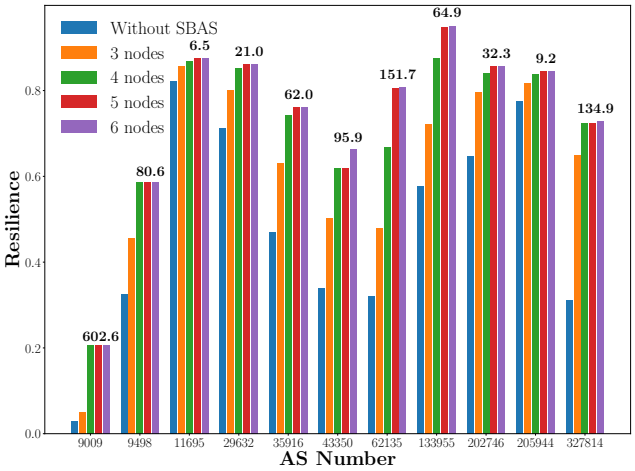

Figure 11: Resilience comparison across known serial hijacker ASes. We mark the maximum resilience gain offered by SBAS for each AS adversary.

B Additional Simulation Results

Resilience definition For a given victim AS $v \in \mathcal{V}$, we consider a set of potential adversary ASes $\mathcal{A}$ along with a set of potential traffic sources $\mathcal{B}$. Let us consider an adversary AS $a \in \mathcal{A}$ which attempts to launch an equally specific-prefix hijack attack against a prefix p originated from v, and a traffic source $b \in \mathcal{B}$ which sends traffic to p.

$$\alpha(v,a,b) = \begin{cases} 1 & \text{if } a \text{ fails to hijack traffic from } b \text{ to } v \\ 0 & \text{otherwise} \end{cases}$$

In our simulations, selection among equally preferred paths is made via a random tiebreak. Aggregating across the adversary and traffic source sets, we compute a normalized resilience measure for the victim:

$$\beta(v,\mathcal{A},\mathcal{B}) = \sum_{a\in\mathcal{A}}\sum_{b\in\mathcal{B}} \frac{\alpha(v,a,b)}{|\mathcal{A}||\mathcal{B}|}$$

Intuitively, a resilience of 1 indicates that a node in set $\mathcal{B}$ attempting to send traffic to a prefix originated by the victim v will always route its traffic to the true origin, even in the presence of equal prefix length attacks by the attackers in $\mathcal{A}$.

Serial Hijacker Simulation The histogram in Figure 11 shows the results of the simulation of SBAS's resilience against a set of ASes with a history of serial BGP hijacking attacks (as identified in Section 7.2.1. As previously mentioned, SBAS routes offer higher resilience than the baseline approach for all of the 11 serial hijacking ASes, with a mean resilience improvement of 114.7%. Although this result focuses on a relatively small adversary set, it demonstrates that

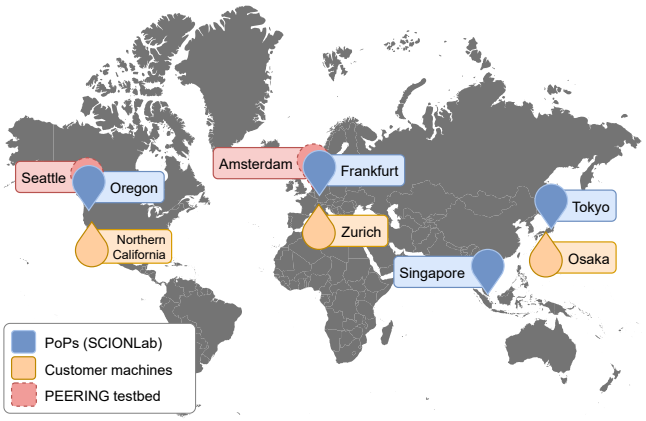

Figure 12: The real-world deployment of SBAS including four PoPs in the United States, Germany, Singapore, and Japan.

SBAS provides sizable security improvements against not only randomly sampled adversaries but also notoriously malicious ones.

Choice of BGP Announcement Node Locations

Locations for the BGP announcement nodes for the three- to six-node SBAS configurations simulated in Section 7.2.1 were computed by simulating all possible N node combinations from the set of 15 total transit-carrying PEERING nodes against another random sample of 1000 AS attackers, and then choosing the best performing node locations within them. In general, we found that the globally optimal node location configuration was equivalent to the iterative greedily optimized one, which can serve as an efficient guiding principle for scaling up the SBAS backbone in the future (starting from the current deployment shown in Figure 12).

While SBAS may leverage any AS-level participant to announce BGP routes, we restricted the simulation to sets of PEERING nodes as a parallel to deployments in the ethical hijacking experiments. We note that choosing SBAS BGP announcement nodes outside of the PEERING network may provide further security gains, which will be a focus for future work. We anticipate adding other announcement nodes in geographical areas where PEERING has no presence, such as Asia, may provide the greatest further resilience gains beyond the simulated deployments.

C Details of Network Operators Survey

Our survey of network operators consisted of a four-minute video describing SBAS (available at https://youtu.be/xsLjcI-qRd0) followed by 19 questions divided into four sections: Background, Incentives, Feasibility and Usefulness, and Governance. We designed the survey to not only find out operators opinion's towards SBAS, but also find out what operators thought of secure routing technologies in general since people with different perspectives on secure routing are likely to have different attitudes towards SBAS. By distributing the survey through direct contact, the RIPE mailing list and the NANOG mailing list, we received 31 responses to the survey. A full description of survey questions and results is contained in our tech report https://arxiv.org/abs/2206.06879.

D System Details of an SBAS PoP

Data Plane In order to handle packets, an SBAS node primarily operates three routing tables for different types of destinations: control, secure, and optimized. When a packet has to egress the kernel of the SBAS PoP (either when coming from the PoP itself or being forwarded), the PoP first checks if the packet contains the source IP of the local router running on the PoP. These are control packets used to enable the iBGP sessions between different SBAS PoPs and are routed using the control table (which contains routes to the IP addresses used by the routers at different SBAS PoPs). This table is loaded with the highest priority and is used exclusively for inter-router communication.

Next, all packets are checked against the secure table. This table contains the RPKI-validated routes to different SBAS customers. If there is a covering prefix in the secure table, a packet is always routed via this prefix (to the appropriate SBAS customer) *regardless of whether an Internet route for that prefix or a more specific prefix exists*. This prevents routing loops (since customer's secure prefixes are also announced to the Internet) and ensures that, even in the event of a sub-prefix BGP hijack, a packet is sent through the secure network to the right customer.

After this, all packets hit the optimized routing table. This table contains routes to Internet destinations. These routes either involve sending the packet to one of the PoPs Internet peers or providers or sending it to another SBAS PoP in the case where the SBAS PoP (this is appropriate in the case where the SBAS PoP does not make BGP announcements or an alternate PoP's Internet route is selected by the routing engine).

Control Plane Each SBAS PoP maintains three types of BGP sessions: iBGP sessions with other PoPs, eBGP sessions with SBAS customers, and eBGP sessions with Internet peers and providers. Prefixes learned from each of these sessions are loaded into different tables and handled in the data plane (see Section D for more details). Furthermore, for Internet routes and routes learned from other SBAS PoPs, the SBAS routing engine performs route selection between different available routes based on a user configurable metric that can vary from security, to greenness, to preference for certain geographic regions. This is done by having the BIRD routing demon (which manages the BGP sessions) output routes in MRT format [21] which is then parsed by the SBAS routing engine. The SBAS routing engine then compares the user-defined metric on the available routes in the MRT file and installed the best one into the optimized routing table.