A path forward: Improving Internet routing security by enabling trust zones

Abstract

The best currently available practices in routing security (ROV/RPKI) have limited deployment because they do not align the incentives of ISPs and their customers with improved routing security. Even if consistently implemented, the practices target only the simplest form of hijack, i.e,. an *origin hijack*. We propose a new set of operational practices that will block a wider range of hijacks, including *path hijacks*. Our approach relies on existing capabilities and institutions to make measurable progress to prevent both origin and path hijacks: existing RPKI infrastructure and ROV capabilities; the existing MANRS framework, and current techniques for collecting and analyzing interdomain (BGP) topology data. A key insight of our proposal is that *topology matters*, and that there is a coherent core of ISPs that emerges organically in the ecosystem, which we can leverage to create a *zone of trust*, a region that protects not only all networks in the region, but *all directly attached customers* of those networks. The result is a virtuous circle, where customers benefit from choosing ISPs committed to the practices, and ISPs (thus) benefit from committing to the practices. Our approach provides an alternative to consider in the face of increasing regulatory pressure on industry to improve the security of global routing in the Internet.

1 Introduction

The Internet's global routing protocol – Border Gateway Protocol (BGP) – suffers from a well-documented vulnerability: a network (termed an Autonomous System or AS) can falsely announce that it hosts or is on the path to a block of addresses that it does not in fact have the authority to announce. Routers that accept a false route announcement – known as a *route hijack* – will deflect traffic intended for addresses in that block to a rogue AS.

The two clear victims of a route hijack are the owner of the hijacked block and the sender of traffic to the hijacked block. If the attacker hijacks address space in order to impersonate the legitimate holder [6,17,27,57] or to inspect [52] the traffic, then senders of traffic to the hijacked block may fall victim to a scam or surveillance. If the attacker hijacks address space in order to conduct malicious activity [54,74,77], a third victim is the target of the malicious activity. The malicious activity may cause blocklisting of the address block, which impairs the legitimate owner's use of the block.

The simplest form of route hijack is an *origin hijack*, in which a malicious AS falsely announces ('originates an assertion') that it directly hosts (i.e., is the origin for) a prefix that belongs to someone else. In a *path hijack*, an attacker claims to be an AS *in* the path to a prefix, forging the legitimate owner's ASN as the origin of the prefix. Either attack violates the essential assumption of BGP's trust model: mutual trust of the two endpoints exchanging BGP messages. The highly distributed operation of the BGP protocol – ≈75K independent networks around the world – and its role in establishing and maintaining the connectivity we call "the Internet", have contributed to the persistence of this long-standing but increasingly dangerous vulnerability.

All solutions to the route hijack problem require some enhanced operational practices to identify and block propagation of bogus route announcements. The fundamental challenge is that networks that commit to such operational practices face additional costs and operational complexities, but the benefits may not accrue to them or their customers. Collectively, network operators have an interest in a well-performing, reliable public Internet. However, endpoints rather than transit ISPs are the most common target of hijacks and thus primarily benefit from reduced hijacks, while transit ISPs bear the cost of deployment of mitigations. Investment in the mechanism and inevitable errors associated with the learning curve in deployment increases transit ISPs' costs, making them less competitive. There is thus a *last mover advantage*, i.e., the first ISPs to deploy today's BGP security mechanism may see no real benefit, either to themselves or their customers.

The growing prevalence and potential harms of route hijacks have motivated the U.S. Federal Communications Commission (FCC) to issue a Notice of Inquiry into potential regulatory interventions that could reduce the severity of the threat to U.S. networks and traffic [73]. Several U.S. government agencies, including the DHS and a joint filing by the

DOD and DOJ, have urged the FCC to take action [71, 72]. Other commenters emphasized the challenges of regulation in this domain. Tension is clearly increasing on this topic, as multistakeholder activities have continued for over a decade, with the most optimistic predictions for deployment of protocol-based solutions to path hijacks estimating at least another decade. In the meantime, the risk and prevalence of both accidental and malicious BGP hijacks grows, rendering even the largest companies in the world victims of hijacks [17].

There are architectural and economic constraints that prevent a perfect solution to Internet interdomain routing security. The collective-action characteristic of the problem is fundamental: even those who are willing to invest their way to increased routing security cannot do so without commitments from other networks to prevent propagation of bogus routes. We believe the community needs to refocus on a new goal: not to protect the maximum number of ASes, regardless of their interest in security, but to *provide a concrete action that a security-aware AS could take to protect itself from both having its address blocks hijacked, and its traffic to other address blocks hijacked*. In this paper we present such a solution – one that can block both origin and path hijacks by combining well-known best practices, capabilities, and institutions: existing RPKI infrastructure and ROV capabilities; the existing MANRS framework, and current techniques for collecting and analyzing interdomain (BGP) topology data. A key insight of our proposal is that *topology matters*, and that there is a coherent core of ISPs that emerges organically in the ecosystem, which we can leverage to create a *zone of trust*, a region that protects not only all networks in the region, but *all directly attached customers*. These ASes thus benefit from a greatly reduced risk of having their addresses or traffic hijacked. Our proposed practices thus provide a basis for participating networks to market their improved security to potential customers. This critical feature creates an incentive for customers to prefer a participating provider, and for providers to participate in the zone of trust. We call our enhanced program *VIPzone*, for *Verified IP zone*.

The roadmap of this paper is as follows. We first describe barriers to solutions over the last two decades (§2). We describe the threat model in §3. In §4 we introduce the principles and operational practices of our *zone of trust*, and explain how it offers a more incentive-aligned direction for preventing BGP route hijacks. We reason about residual risks of hijacks in various interconnection scenarios (§5) and explore possible deployment trajectories (§6). Our hope is that we can reframe the routing security conversation from purely protocol-based solutions to those that leverage institutionalized self-regulated cooperation and Internet topology analysis capabilities.

2 Background and Related Work

The Internet standards community has long struggled with proposals to tighten the integrity of BGP communications. As with protection of other Internet transport mechanisms (e.g., DNSSEC and TLS certificates), the standards community has grappled with complexities of cryptographic key management, trust anchors, and performance implications that hinders standardization, implementation, and deployment. Over the last 30 years, over 20 proposals to secure BGP have come out of academia, industry and the Internet Engineering Task Force (IETF), some of which Figure 1 highlights. We describe these decades of how the standards and operational communities have tried to tackle this problem, and how it motivates our proposal.

2.1 Interdomain Routing

ASes use BGP to exchange *routes* that describe paths to destinations in the global Internet. Two important components of a route are the *prefix* that specifies the block of addresses of a route, and the *AS path* that reports the sequence of ASes that received the route. In order to prevent packet forwarding loops, a router chooses the *most specific* route to a destination IP address – i.e., for 192.0.31.8, it would prefer a route with a prefix 192.0.31.0/24 over 192.0.30.0/23 because the /24 prefix is more specific than the /23 prefix. Operators typically use this property for *traffic engineering* – steering traffic into their network through specific routers and links. BGP also provides a mechanism to annotate announcements with meta-data – known as *BGP communities* [14] – providing sophisticated signaling within and between ASes, enabling innovations in traffic engineering [18] and automated blocking of denial-of-service attack traffic in transit networks on the path to the victim [15, 39].

There are typically two types of relationship between neighboring ASes: customer-to-provider (c2p), where the customer pays a provider to obtain global reachability, and peer-to-peer (p2p), where the peers exchange routes to their customers without involving an intermediate provider [21]. An AS in the Internet is typically a rational actor; if it has the choice of multiple routes to the same prefix, it prefers routes received from customers (these are a source of revenue), over routes received from peers (these typically do not cost the AS), over routes received from providers (these cost the AS) [21]. Other ASes that an AS X can reach through a customer link are within the *customer cone* of X.

A few ($\approx$15) ASes in the Internet obtain global routing with routes received from their peers and customers. These ASes connect in a full mesh (a peering clique) that enables packet delivery between arbitrary networks in the Internet that each have different transit providers. The ASes in this group that do not pay for peering are known as Tier-1 providers; because any payment between ASes is confidential and speculative [42], we use the term *Tier-1* in this paper to refer to the peering clique.

2.2 Routing Security in the 1980s

In 1982, Rosen [59] documented that it is possible to corrupt interdomain routing in RFC 827, in the context of a predeces-

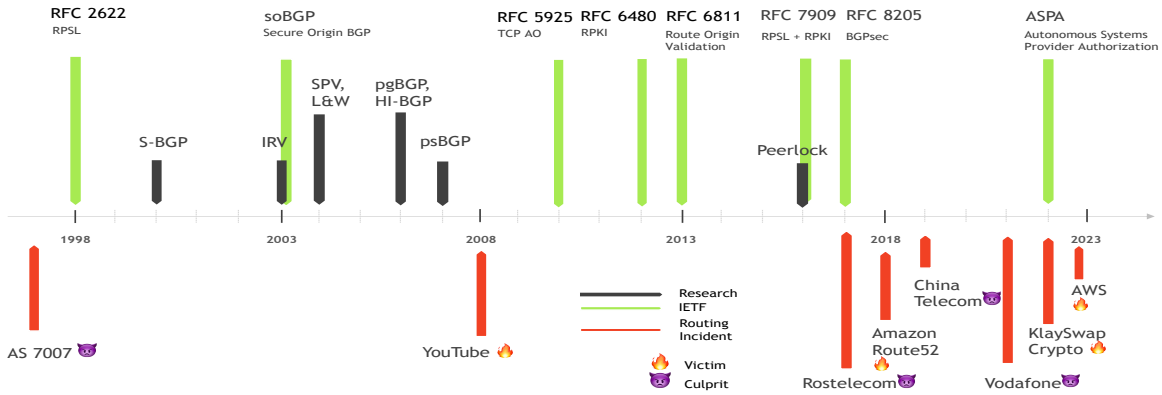

Figure 1: Decades of proposed routing security approaches; sample of high-profile hijacks.

sor of BGP called the Exterior Gateway Protocol (EGP):

> *If any gateway sends an NR [neighbor reachability] message with false information, claiming to be an appropriate first hop to a network which it in fact cannot even reach, traffic destined to that network may never be delivered. Implementers must bear this in mind.*

This warning to implementers suggests the perceived threat in 1982 was accidental misconfiguration, rather than malicious operators. When the IETF started to study BGP security in the late 1990s, they did not initially assume that an AS operator was an important threat actor. Instead, they focused on the threat that a third party could intercept the traffic between two well-behaved ASes and then modify the BGP update to inject a false assertion. To defend against this threat, in 1998 the IETF added an optional extension to TCP to allow end-points to authenticate the contents of a TCP segment [29, 70].

2.3 Routing Security in the 1990s: IRR

The IRR system enables network operators to publish address ownership and routing policy records [1], which other operators can use to build filters that permit or deny routes according to these operator-registered policies. The IRR system was created in the 1990s, when the goal was to prevent route leaks rather than hijacks. More problematic, though, is that some IRR systems do not validate registration data, allowing the IRR to be used by attackers who falsely claim ownership of resources which they use in a hijack [19, 50, 51, 58, 69].

2.4 Routing Security in the 2000s: BGPsec

Aiming for a rigorous approach to routing security, in 2006 the IETF's Secure Inter-Domain Routing (SIDR) Working Group began designing a variant of BGP that would support path validation. During this decade over a dozen competing approaches came out of academic and industry [26, 28, 30, 35, 37, 38, 49, 53, 55, 64, 66, 75, 76, 78]. The protocol that became an IETF standard (RFC 8205) in 2017 is called BGPsec [41]. BGPsec update messages have two important differences from BGP: the router includes in the announcement the AS to which it is sending that announcement; and a cryptographic signature over the message to enable any router along the path to verify that the series of signatures are valid. This mechanism prevents path hijacks: a malicious AS cannot prepend a valid origin announcement to a bogus AS because the AS number of the next AS in the path must match the prepended number. Cryptographic attestation of paths requires propagation of a new layer of cryptographic transaction at each hop, which is computationally expensive but also poses a router-level (rather than AS-level or prefix-level) key distribution challenge, since every router must have its own public key signed by a certificate authority. Furthermore, every AS along the path must implement BGPsec for the path to be protected. Partial deployment (inevitable during a transition) implies inconsistent and unpredictable implementation of the required checking. The complexity, overhead, and misaligned incentives have prevented significant operational deployment of BGPsec, despite a decade-long standardization process that completed 6 years ago.

2.5 Routing Security in the 2010s: Analysis, Compromises, Collective Action

This decade vivified the intractability of the routing security problem. We review three areas of endeavor: rigorous analyses of the misaligned incentives to deploy routing security solutions; technology, standardization, and operational mechanisms to mitigate the simpler problem of origin hijacks; and a collective action (MANRS) to overcome the counter-incentives to deploying these mechanisms.

2.5.1 Analyzing deployment incentives

As early as 2009 researchers began to survey the array of efforts and analyze why they failed to gain traction [13, 47]. Such reviews continued throughout this decade [45, 63, 68]. Researchers also explored approaches to overcome the economic counter-incentives to deployment of protocol-based approaches to routing security, and analyzed the implications of partial deployment [16, 24, 25, 43]. The deepest body of work

on this topic was by Sharon Goldberg and Michael Schapira and collaborators. In 2011, Gill, Schapira, and Goldberg proposed a strategy that would create market pressure to adopt BGP path validation. (They referred to the set of options at the time as *S*BGP*). Their proposal required (e.g., by regulation) a few Tier 1 ISPs to first deploy S*BGP, and required those participating in S*BGP to prefer secure routes over otherwise equivalent (e.g., equal path) routes [24]. This scheme also reduced deployment complexity by allowing transit providers to cryptographically sign routes on behalf of their stub customers. Their simulations on realistic AS topologies showed that under these conditions, the S*BGP ASes would draw traffic away from other ASes, and most of the rest of ASes would then switch to S*BGP to get their traffic (revenue) back. A followup study led by Lychev, Schapira and Goldberg [25,43] two years later acknowledged that having Tier 1 ISPs lead a market-driven deployment would not work because they are typically so well connected that they would have shorter routes to most other networks, and economic incentive would override any longer secure route option. A key insight of this work is that because commercial ISPs must compete in the market, and thus behave economically rationally, there is no way for the market to evolve naturally to pervasive S*BGP deployment. That is, it is not reasonable to expect competitive ISPs to voluntarily pick more secure paths over shorter paths.

2.5.2 Preventing origin hijacks: RPKI and ROV

While BGPsec has been undergoing implementation and evaluation for a decade, operators have focused the more tractable challenge of *Route Origin Validation* (ROV), which is recognized as the best current practice in routing security. The IETF SIDR WG specified ROV in 2013 as a mechanism to mitigate the risk of *origin hijacks* (the simplest form of hijack) [61]. ROV uses a Resource Public Key Infrastructure (RPKI) [31], i.e., an authoritative database maintained outside of BGP to store *Route Origin Authorization* (ROAs), cryptographic signatures (certificates) that authorize designated ASes to originate routes to address blocks. Routers using ROV drop BGP announcements that do not match their corresponding ROAs, if one exists. RFC 6811 [61] specifies the ROV protocol with important caveats: its dependence on the integrity of the database used to validate routes, and its inability to prevent path hijacks. In particular, an attack can impersonate the valid source AS by appending it to a forged BGP announcement (recently observed in the wild [50]). RFC 6811 cautioned: "*..this system should be thought of more as a protection against misconfiguration than as true 'security' in the strong sense.*"

Use of ROAs present other operational challenges. A ROA contains a prefix and a single ASN; if an operator wishes to announce a prefix with different ASNs depending on the interconnection, it must issue multiple ROAs, each with a different ASN. A ROA may also contain a *maxLength* attribute that defines the maximum prefix length allowed for the prefix; for example, a ROA for 192.0.30.0/23 with a maxLength of 24 enables an operator to announce 192.0.31.0/24, which the operator can use for traffic engineering (§2.1). A route is RPKI-valid if any ROA asserts the origin AS in the AS path is valid. In 2017, Gilad *et al.* showed that use of the maxLength attribute could enable an attacker to announce more-specific prefixes that other networks select when reaching addresses within that prefix [23]. Best current practice is to not use the maxLength attribute [22].

Although RIRs have supported RPKI registration of ROAs since 2013, until 2019 there was little evidence of ISPs using ROAs to validate BGP announcements. By late 2022, many large ISPs, including AT&T, KPN, Arelion, and Comcast had started to use ROV to drop invalid announcements [11, 36, 40, 67]. According to NIST's public RPKI monitor based on RouteViews data [48], as of December 2022, 38% of /24s in unique prefix-origin pairs advertised in BGP were covered by RPKI and observed as valid, i.e., the origin AS in the BGP announcement matched the registered ROA. These statistics vary by region: for 31 December 2022, NIST found valid 54% of observed prefix-origin pairs in the RIPE region, 50% in LACNIC, 44% for APNIC, 28% for ARIN, and 19% in the AFRINIC region [46]. One factor limiting the use of ROAs in some regions is the position that a resource holder has sign a registry services agreement (RSA) in order to create ROAs, and many legacy address resource holders in the ARIN region have been reluctant to sign ARIN's agreement due to its controversial position that address holders have no property rights to their address space. In September 2022, ARIN finally removed this clause from their RSA [4].

2.5.3 Collective action attempt: MANRS

In 2014, several network operators established a voluntary initiative to promote operational practices to "help reduce the most common routing threats on the Internet" – which they called Mutually Agreed Norms for Routing Security (MANRS) [32]. MANRS specifies four practices for participating networks, two of which correspond to the RPKI/ROV steps of registering authoritative information about one's prefixes, and verifying BGP announcements against authoritative information. The exact wording of these two practices are: (1) *Prevent propagation of illegitimate routes from customer networks or one's own network.*; and (2) *Document in a public routing registry the prefixes that the AS will originate.*

To conform with the first practice, a MANRS member must verify two aspects of an announcement from a customer: it must confirm that the customer has used an ASN that it is legitimately allowed to use, and for any prefix originated by that customer, that the ASN is allowed to announce that prefix. However, to encourage broad uptake, MANRS does not specify how a member AS should verify the assertions of its customers, and in particular *does not require the use of RPKI/ROV (ROAs) in this verification*. The AS can use ROAs, or can verify against (less authoritative) information in the Internet Routing Registry (IRR), or rely on a private arrangement with its customer.

The MANRS initiative has a key strength: it illustrates that ISPs can institutionalize their recognition of the need for a collective commitment to operational practices to reduce threats to the routing system. However, as the FCC observed [73], the MANRS program has had limited success. Many of the largest ISPs do not participate, and some participating ISPs are not conforming to the practices [20]. Of those ASes that conformed in December 2022, 95% of them used the IRR rather than the more authoritative (but more complex) RPKI.

The limited success of MANRS is rooted in misaligned incentives that manifest in three ways. First, although if consistently implemented, the MANRS practices will reduce the incidence of invalid origin hijacks, there is no direct relationship between the action of any given MANRS member and the overall security of the Internet, or even the security of any customer of a MANRS member.

Second, the current MANRS practices, even the stronger RPKI/ROV options, only aim to prevent origin hijacks rather than path hijacks, which are arguably not much harder to perpetrate than origin hijacks. Thus, many network operators believe RPKI/ROV does not justify the cost and complexity of participation.

Third, there is currently insufficient auditing of conformance to lend confidence to the assumption of consistent implementation. One independent assessment in 2022 found a significant fraction (16%) of members do not actually conform with the practices [20]. Instead the MANRS Observatory [44] reports trends of all ASes on the Internet, without focusing on or identifying the conformance of individual members. More rigorous auditing would be expensive and further reduce the incentive to participate.

2.5.4 AS Provider Authorization (ASPA)

Recognizing the barriers to BGPsec deployment, and the lack of path validation capability in ROV, in 2019 several engineers proposed AS Path Authorization (ASPA) as a mechanism to protect against route leaks and forged-origin prefix hijacks [7]. As of January 2023, this proposal is still in the IETF development and standardization process [8]. In the ASPA scheme, customers register their set of legitimate transit providers in a globally visible database. That database allows any AS to examine a BGP announcement to detect and reject invalid path announcements, so long as all the ASes along the path have registered their providers in ASPA. ASPA ignores the issue of peering connections, assuming that properly filtering routes from peers can avoid most problems. ASPA aims to solve *route leaks* rather than *path hijacks*; it does not prevent an attacker from spoofing a sequence of ASes in the path that passes ASPA. After explaining details of our approach, we compare it to ASPA in terms of feasibility and protection (§7).

2.5.5 SCION: Leveraging an Already Secure Backbone

In 2022, Birge-Lee *et al.* leveraged the ideas of the SCION network architecture to bootstrap a secure routing system [10]. Their proposal assumes the existence of a Secure Backbone AS (SBAS) that can be a private network or a federated network using BGPsec or SCION or similar protocol. The paper describes techniques to leverage that SBAS to secure routing for external ASes. They simulated their idea with a four-PoP network in their SCIONLab, and used the PEERING testbed peering to emulate hijacks to ASes connected to the SCION core. This research project has similar objectives to our proposed approach, but our motivation was to create a framework that operators could deploy in today's BGP ecosystem.

2.6 Routing Security in the 2020s: Regulatory Interest Grows

In the last ten years, the risk and prevalence of both accidental and malicious BGP hijacks has grown, rendering even the largest companies in the world victims of hijacks [17]. Researchers have discovered hijacks of unannounced address space, and of RPKI-valid address space where the attacker forged the ASN in the ROA as the origin of their path [50]. (The owner had registered a ROA for their address space in November 2014, but abandoned the address space in July 2020 without removing the ROA.)

After earlier hijacks of AWS address space [27] motivated Amazon to register ROAs for most of its address blocks, attackers developed more sophisticated path hijacking techniques. The August 2022 hijack of AWS space [17] succeeded for multiple reasons. Amazon signed multiple ROAs that allowed different ASNs to originate their prefix; these ROAs had maxLength attributes that the attacker exploited to announce an IPv4/24 that hosted the crypto-currency service; and the attacker registered that IPv4/24 in an unauthenticated IRR entry to convince upstream providers to permit the prefix announcement. However, even if Amazon had announced a competing more specific, the attacker's path would have been preferred for networks that were customers of AS1299 who did not have a route to Amazon via a p2p route. The bottom line, and our motivation, is that there is no deployed protocol that defends against path hijacks in the Internet.

The persistent failure of market-driven solutions to routing security has recently triggered government interest and inquiry into potential interventions. In 2022, the OECD [2], ICANN [62], and BITAG [3], and the U.S. FCC [73] all published reports with extensive references related to routing security challenges, and limitations of proposed solutions.

We anticipate governments may feel compelled to intervene in the Internet infrastructure ecosystem to improve routing security, and we seek to provide an alternative that leaves as much control as possible with the participating networks, but provides a demonstrable competitive advantage to networks complying with secure routing practices. We suspect that doing so requires establishing a coherent topological zone of

participants. Per the pessimistic conclusion of [43] regarding market-driven evolution of secure routing: *"We hope that our work will call attention to the challenges that arise during partial deployment, and drive the development of solutions that can help surmount them.. Alternatively, one could find deployment scenarios that create 'islands' of secure ASes that agree to prioritize security 1st for routes between ASes in the island; the challenge is to do this without disrupting existing traffic engineering or business arrangements."* [43] Our proposal pursues this challenge: islands committed to secure practices, architected to prevent path hijacks without requiring the key management complexity of BGPsec. We believe it is a alternative worth serious debate before pursuing more blunt regulatory measures.

3 Threat Model

We next describe the capabilities of defenders (§3.1), to contrast defender capabilities with attacker capabilities (§3.2).

3.1 Defender Capabilities

While there are ≈75K ASes as of January 2023, the vast majority of ASes (≈70K) are stub ASes that rely on transit providers and IXPs for Internet routing. Representatives at these transit providers engage in contractual agreements when they interconnect with their neighbors. Operators at these large networks regularly interact at industry events, such as at network operator group meetings (e.g., NANOG) and thus have established relationships. In our threat model, the defenders are these transit providers. Defenders have the capability to establish parameters with their customers in terms of what prefix announcements the customer is expected (and allowed) to make, and thus to automatically accept or reject routes through configuration capabilities present on routers. Defenders can access external databases, e.g., IRR, RPKI, to support their assessment of their customer routes.

A defender does not usually have the ability to verify the announcements of their customers' customers, due to the temporal dynamism in the interdomain relationships of their customers. Further, some defenders, and their customers, are limited in how they use RPKI. For example, some legacy resource holders are hesitant to obtain ROAs, as doing so would require they enter a contractual agreement with an RIR (§2.5.2). Finally, a defender cannot control the route selection policies of their peers or customers; these ASes might select hijacked routes from their neighbors.

3.2 Attacker Capabilities

We assume that the attacker either controls or has subverted an AS that connects to the Internet using one or more transit providers, which provide routing to the rest of the Internet for that AS and deliver traffic intended for that AS. The attacker has the ability to corrupt unauthenticated databases, such as IRRs, with false claims that they are the legitimate holder of a prefix (§2.3). Finally, an attacker has the ability to commit to security practices that they have no intention to follow.

An attacker does not have the ability to hide their activities, due to the public nature of BGP routing; in order for their attack to be effective, their hijacked route must propagate. Nor does an attacker have the ability to issue ROAs for address space that they do not control, unless they compromise the RIR (an insider) or the prefix holder's RIR account. Multiple RIRs provide the ability for their account holders to automatically update their records via email, provided they include their password in clear text in their email [5, 56]. Therefore, nation states may have obtained RIR account credentials as part of their pervasive surveillance capabilities [9], which they could use to corrupt authenticated RIR databases.

4 Toward a Routing Zone of Trust

Given that some operators are more motivated than others to invest in security, an ideal set of operational practices would bring security benefits to networks that choose to deploy them. To accommodate the reality that some operators will not have the resources to invest in stronger security, an ideal set of operational practices would also enable customer ASes *who are not conformant to them* to improve their own protection from hijacks by obtaining transit service from a network who is conformant. The fundamental premise of our proposal is that the network operator community could develop such a set of operational practices, such that accountable conformance to these practices within a *connected* set of ISPs could overcome the essential barriers of previous routing security frameworks, including achieving protection against path hijacks.

4.1 Benefits of a connected region

Imagine a connected region of the Internet composed of ISPs that commit to perform ROV on all announcements coming into that region. Then, within that region there will be no propagation of origin hijacks. Furthermore, any AS that directly attaches to that region (i.e., the AS is a customer of an ISP in that region, not necessarily in the region itself) cannot be the victim of an origin hijack inside that region (Figure 2). Similarly, if a new set of operational practices (§4.2) prevents the propagation of path hijacks in the region, then customers directly attached to the region will not be the victim of a path hijack inside that region. We call this region a *zone of trust* because the protection arises at the perimeter of the zone. This protection requires that ASes in the zone be able to trust that the routers at the perimeter function correctly, which requires some degree of transparency and accountability.

The idea of a coherent perimeter around a zone is missing from today's interdomain routing system. Global deployment has always been a routing security protocol design assumption. Yet, recognition that ASes themselves can be the threat actors sheds doubt on any aspiration to make BGP *globally* secure.

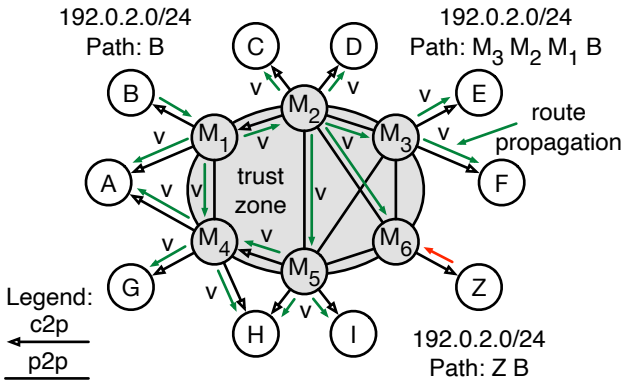

Figure 2: A Routing Zone of Trust can defend members and their customers from path hijacks in the zone if members (M) mark routes from their customers as verified (v) as they enter the zone, and other zone members select verified routes over unverified routes. In this example, M_1 expects their direct customer B to announce 192.0.2.0/24, so it marks that route as verified, and propagates the route to other members. Lines with hollow arrows show c2p links, lines without arrows show p2p links, and lines with solid arrows show route propagation. The hijacked route via Z does not propagate in the zone, because Z is not a member, and the zone has a verified route.

Our premise is that creating a zone of trust through perimeter protection (a trust-but-verify regime) offers a more pragmatic approach for today's routing system. Most important is the alignment of incentives. A zone of trust approach would be able to clearly articulate the benefit that the practices of that zone are bringing to their customers.

Could such a coherent topological region exist? Fortunately, it already does, in the context of the MANRS initiative. The current MANRS program is already centered around operational practices, and already manifests a significant zone of trust. In §6.2 we describe how one could leverage the existing MANRS initiative to bootstrap a zone of trust.

4.2 Proposed Solution: VIPzone

We now describe the details of how our proposed operational practices in a coherent zone of trust, which we call *VIPzone*, will limit path hijacks. For an AS to be in the VIPzone, it must commit to the practices we specify next, and either be a Tier-1 provider, or have a member of the VIPzone as a transit provider.

VIPzone members must use the following operational practices. First, VIPzone members that can participate in these enhanced practices must be part of a connected zone, as illustrated in Figure 2. Second, if a VIPzone member receives a BGP announcement from a neighbor that is not in the zone, and the announcement is for a prefix that the neighbor *originates* and the member can verify as legitimate, then the member will tag the route with a new BGP community value [14], which we call *VERIFIED*. Third, VIPzone members must propagate this community value as they forward announcements to other ASes. This allows neighbors to establish the authenticity of the route, regardless of the distance they are from the origin. Fourth, inside the zone, any AS receiving multiple announcements for the same prefix must prefer one marked VERIFIED. By this rule, no member will prefer a path hijack announcement over a legitimate announcement from customers directly attached to the zone, since those will be marked VERIFIED.

The operational practices that a VIPzone member must configure their routers to follow are:

1. **Prevent false VERIFIED routes:** If the member receives an announcement from a non-member AS, then it MUST remove the VERIFIED community if present. This is to prevent an attacker from injecting a hijacked route that other VIPzone members prefer.

2. **Drop RPKI-invalid routes:** If the member receives an announcement where the origin is RPKI-invalid, the member MUST drop the announcement. This is to prevent origin hijacks.

3. **Prevent propagation of forged routes:** If the member receives an announcement where the AS used by the neighbor is not consistent with the AS numbers legitimate for the neighbor, the member MUST drop the announcement. This is consistent with a know-your-customer requirement, to prevent malicious routes from entering the VIPzone.

4. **Forward VERIFIED routes:** If the member receives an announcement from another member with a VERIFIED community tag set, it MUST retain that tag when forwarding the route to other members. Further, the member MUST retain the VERIFIED tag when it provides the route to non-member neighbors. This is to enable VIPzone members and other neighbors to know which routes have been VERIFIED on entry to the zone, and thus are not path hijacks.

5. **Verify routes with one AS in the path from non-member customers:** If the member receives an announcement with one AS in the path from a non-member customer, it MUST drop the announcement if the route contains a prefix that the customer has no authority to announce (it is not RPKI-valid, or is not from a list of allowed prefixes that the member has previously established their customer is able to legitimately announce) to prevent possible hijacks from propagating. If the prefix is RPKI-valid, is registered by the owner in an authenticated IRR, or from a list of allowed prefixes, it MUST add a VERIFIED community to the route so that other members know that the route is valid.

6. **Forward unverified routes without the VERIFIED tag.** If the member has not established that the announcement is valid (because it has not yet obtained the list of

allowed prefixes, or because the AS path in the route contains more than one unique ASN and so cannot be verified) the member can announce the route to its neighbors but MUST NOT add a VERIFIED community to the route, so that other members do not trust the validity of the route. This ensures routes are still available for networks outside of the VIPzone.

7. **Export routes to a route collector for auditing.** Finally, to allow for auditing behavior of trust zone members, members must export routes to a route collector.

Members that receive a route from a non-member can only perform verification steps if the non-member originates the prefix. The member can use RPKI validation, an authenticated IRR database, or a manually-configured prefix list (ACL) to verify the non-member's announcement is correct. An member receiving any route from a non-member that does not originate the route cannot tag that route as VERIFIED because the AS path could be forged.

4.3 Auditing Requirements

A trust zone, by definition, relies on trust between competing organizations to prevent path hijacks. The VIPzone takes a trust-but-verify approach: checking conformance of members with its requirements, and suspension or ejection of non-compliant members. In support of this auditing, every VIPzone member is required to provide a BGP view to a route collector. The audit process does not use the member's view to audit the member's behavior, as the member could lie; rather, the process uses the views provided by the member's neighbors that are also members and thus provide views of their own. Using these neighbor views, we can establish that the member correctly propagates verified routes with the VERIFIED tag, and does not use the VERIFIED tag on routes that other members have not tagged as VERIFIED.

This practice will allow for automated verification of routing behavior. We argue that an industry-led body, analogous to the CA/Browser forum, should decide on necessary actions if a VIPzone member propagates routes that it should not have selected over a VERIFIED route. The CA/Browser forum uses a similar approach to enforce trustworthy behavior of root certificate authorities. We are not suggesting that in detail the CA/Browser forum is a role model for what is needed here, but rather that it uses similar practices. Its goal is not to detect and block every issuance of a false certificate in real time, but rather to identify CAs that are shown to be untrustworthy and remove them from the list of trusted root CAs included with distributions of browsers. The idea is to enforce proper behavior by making the consequence of misbehavior a substantial penalty. In that context, the CA/Browser community has shown a willingness to take action against providers that do not conform. For the VIPzone to provide protection in practice, the routing community must have the same will. But note that here the penalty is not being disconnected from the Internet, but just losing the right to initiate VERIFIED announcements.

The tests we propose for VIPzone conformance are:

- Rule 1: If an announcement (observed anywhere in the VIPzone) has more than one AS number in the path before it enters the VIPzone, and is marked VERIFIED, the member that introduced the announcement into the core is non-conformant. Our trust model assumes that verification and checking of announcements occurs at specific locations: the ASes at the edge of the VIPzone that have with customers not in the VIPzone. This requirement makes it possible to identify members that do not implement the required practices.

- Rule 2: If an announcement has an invalid origin, as determined by a ROA, independent of path length, the VIPzone member that introduced the announcement is non-conformant.

- Rule 3: ASes in the VIPzone must forward the VERIFIED community value from other VIPzone members.

Cost of checking for conformance The conformance checking requirements imply non-trivial costs. Whether operated by an independent private-sector group such as RouteViews, or some more formally chartered institution or agency, the data collection and curation infrastructure would require staffing to maintain. Then one or more technically capable organizations must perform the auditing and provide the information necessary to judge untrustworthy behavior.

5 Evaluation of Protection and Residual Risk in Various Interconnection Scenarios

By design, our proposed approach does not seek to protect every part of the Internet. Our proposal provides protection from invalid origin and invalid path hijacks to ASes that are in the VIPzone, and to customers that are directly connected to transit providers in the VIPzone. The VIPzone leaves ASes that are not in those classes no worse off than today.

To understand the residual risk in various interconnection scenarios, we introduce the concept of what we call a *local region* of an AS. For any customer C not in the VIPzone but attached to a zone member, the customer C's local region includes the customer cone of C, any transit providers of C that are not in the VIPzone, the peers of C, and the customers in the cones of those peers. Formally, for a customer C, any AS that can originate a BGP announcement that is received by C without passing through the VIPzone is in the local region of C. In general, all of those ASs will be outside the VIPzone. (it might happen that a customer in the cone is in the VIPzone, but this would be unusual.) The customer C is not protected from a hijack launched against it from within this local region. We illustrate scenarios of local-region hijacks in the remainder of this section.

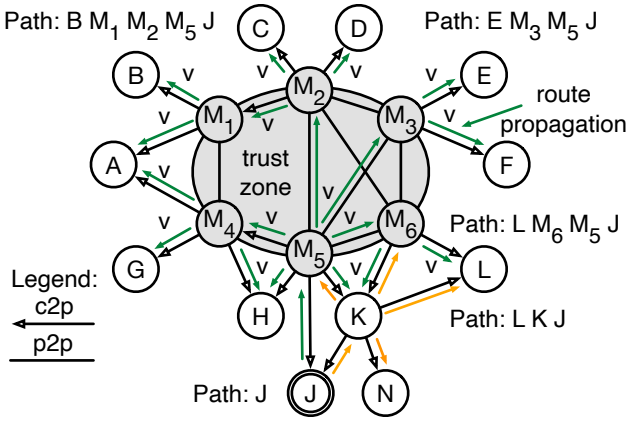

Figure 3: The protection that a Routing Trust Zone provides for customer routes depends on local regions. If customer J has two transit providers (M_5 in zone, and K out of zone) then other ASes that also have out-of-zone providers (e.g. L) may select the unverified route.

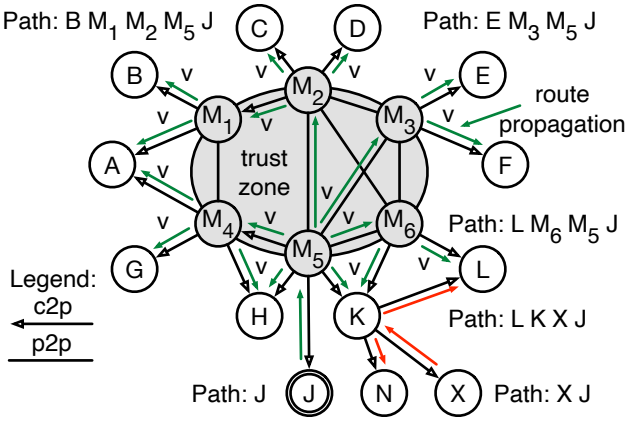

Figure 4: Another illustration of how protection depends on how customers connect to the Routing Trust Zone. Customer L receives two routes for J's prefix, a VERIFIED route via M_6 and an (unverified) hijacked route via K. If L does not prefer the VERIFIED route via M_6 it may select the hijacked route because it has the same AS path length.

In general, an AS in the VIPzone also may have a local region outside the zone, i.e., customers and peers not in the zone that originate and announce routes to that AS. If one AS in that local region tries to hijack the address of another AS in that local region, this attempt might succeed, but that attack is only possible within this local region.

An AS with a large local region may find it advantageous to commit to the practices we specify here and become a member of the VIPzone, in order to improve its protection—and the protection of its directly attached customers—from hijacks. An AS with a small local region may be able to assess the actual level of risk from a locally-generated hijack and conclude that the risk is not material. Such pragmatic risk assessment is part of any realistic approach to security.

5.1 Multihoming Transit Scenarios

VIPzone members, and non-members exclusively connected to VIPzone transit providers, will receive an authentic route from the VIPzone if one is available. The risk for a VIPzone customer increases if it advertises to or accepts routes from a non-member. In the case of transit if a multihomed AS has at least one provider in-zone and at least one provider out-of-zone, then the multihomed AS increases its risk that other ASes with transit providers outside of the zone might use a hijacked route. To illustrate, Figure 3 shows J originating a prefix to transit providers M_5 and K. L has the choice of routes from transit providers M_6 and K, and selects the route from K because it has the shortest AS path. If J were to only announce its routes via in-zone providers, then other customers of the VIPzone are more likely to select a route from another VIPzone member. For example, if J in Figure 3 did not announce its route to K, then L would reach J via VIPzone members M_6 and M_5.

Notice that L, which is also multihomed, could have chosen a VERIFIED route via M_6 if it preferred routes received from VIPzone members that are VERIFIED over unverified routes. In §4.2, rule #4 requires VIPzone members to retain VERIFIED tags so that non-members could select these routes. While the VIPzone practices are not compulsory for non-members, a non-member may choose to configure their routers to remove VERIFIED tags from non-member neighbors, and then prefer routes received from their neighbors who are VIPzone members that are tagged as VERIFIED, to defend themselves from using a malicious hijacked path towards a destination.

Figure 4 illustrates that L is incentivized to so, as it now receives a path hijack by X of equal length to the authentic route via its VIPzone provider, so it may select the hijacked route and suffer associated harms. Note that the risk of L accepting a hijacked route increases if L's relationship with K is p2p or p2c, as these would ordinarily have a higher preference than a provider route regardless of path length.

Figure 5 illustrates a similar scenario of the residual risk in a local region. Here AS J connects to two transit providers (M_5 and X) of which only M_5 is a zone member. AS X and AS Z are in the local region of J, since they can originate BGP announcements that arrive at J without passing through the zone. If J did not use X as a transit provider, or preferred the VERIFIED route from M_5, it would prevent this hijack. If X or Z sends a bogus announcement for a prefix to J, J might decide to prefer it over a valid (VERIFIED) route from M_5. This could happen only if X or Z are malicious–given the local region of J there are no other ASs in a position to launch a hijack.

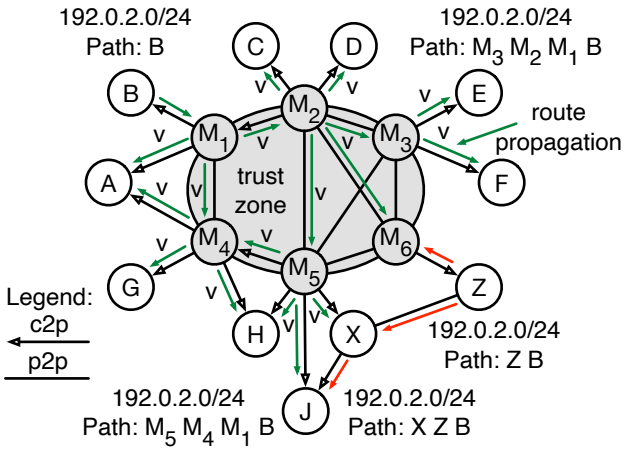

Figure 5: The risk of a hijack of J's traffic, in this case destined to prefix B, is limited to BGP announcements coming from X or Z, but that assumes it chooses that route rather than a VERIFIED route to B from its zone transit provider M_5.

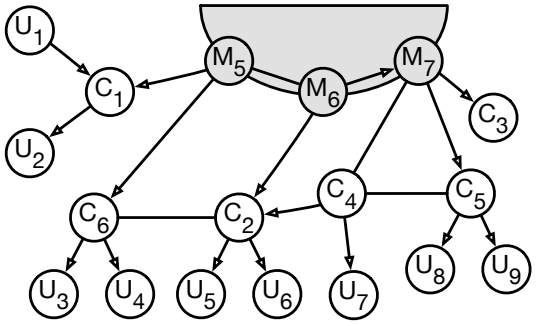

Figure 6: **Peering across the VIPzone perimeter.** *C_2 has two transit providers, only one of which (M_6) is in the VIPzone. M_6 will announce into the VIPzone a verified path to C_2. C_4 peers with M_7 which is in the VIPzone. C_4 will announce to M_7 a route to C_2. M_7 will prefer the VERIFIED announcement, and will send traffic to C_4 through its provider M_6 in the zone, not over the peering link to C_4.*

5.2 Peering Interconnection Scenarios

In most cases, the analysis for a peering connection is similar to transit connections, and straightforward.

5.2.1 Peering with IXP route servers

An IXP-operated route server centralizes peering routes from IXP members and makes these routes available to other IXP members. If the IXP is a member of the VIPzone and has configured the route server to verify routes received from IXP members, then the route server can mark routes as VERIFIED, and VIPzone members can propagate the VERIFIED route. Otherwise, routes received from a route server are unverified.

5.2.2 Peering of zone customers outside zone

If two ASes not in the VIPzone but directly connected to VIPzone providers peer with each other, they may receive announcements of routes to each other via the VIPzone that are marked VERIFIED, and announcements over the peering connection that are not VERIFIED. Because ASes not in the VIPzone are not expected to use that community value to assign a preference to an announcement, their routing policy would normally be the same as today. (Note ASes outside the zone may choose to use this VERIFIED value to prefer routes.)

5.2.3 Peering across the VIPzone perimeter

Peering across the VIPzone perimeter has a straightforward scenario and a complicated scenario. Imagine that VIPzone M_7 in Figure 6 peers with non-VIPzone C_4. In the straightforward case, M_7 will apply the same VIPzone rules to peer C_4 as it does for customers C_3 and C_5, i.e., forward or drop announcements and mark as VERIFIED announcements that the peer legitimately originates. The customers of that peer C_4 would not have their routes VERIFIED. Typical routing policy is that the AS in the zone would only use these announcements from peer C_4 for itself and its customers—it would not forward them on to other peers or providers.

The complicated peering scenario arises when a customer of that non-zone member also obtains transit service from an AS in the VIPzone. Figure 6 shows C_2 with two transit providers, only one of which is in the zone. The transit provider not in the zone (C_4) also peers with an AS in the zone (M_7). In this case, M_7 will receive a VERIFIED announcement to C_2 via M_6, which per the VIPzone rules it must prefer over the route via the peering link from C_4, so M_7 will not benefit from the peering link for traffic to C_2, even if it would normally prefer that peering link.

In §6.3, we analyze the observable connectivity of a major ISP, Hurricane Electric (HE), which has a liberal peering policy, and as a result has more peers than any other ISP in the Internet. We analyze the consequences if HE were to join the VIPzone—in particular what residual risks might arise from the large local region created by their peers.

6 Deployment Trajectories

We consider three deployment trajectories, the first of which demonstrates the value of a VIPzone approach, but has no clear path to achieving it. The second trajectory leverages the MANRS initiative. The third trajectory, or something similar, could derive from regulatory incentives.

6.1 Greedy Construction of Global VIPzone

Figure 7 shows the growth in the number of customer ASes whose routes are protected as the VIPzone size increases. For

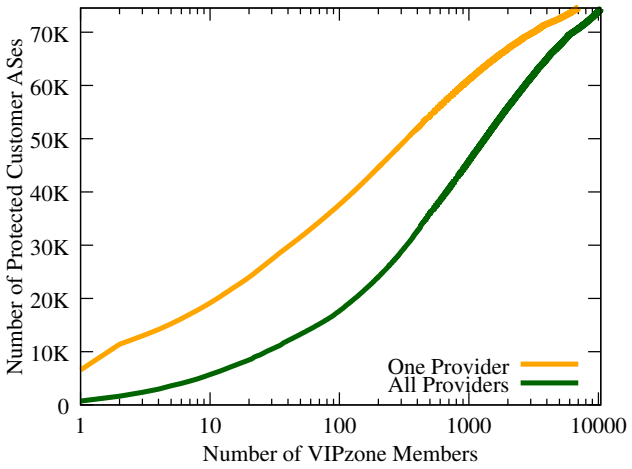

Figure 7: With the 100 most strategically selected ASes as VIPzone members, half of the ASes in the Internet are protected from hijacks within the zone.

this analysis, we added ASes to the VIPzone to maximize the gain in protected customers for each iteration using CAIDA's AS relationship data (1 Dec 2022). Customers are protected from hijacks within the zone when they have one provider as a VIPzone member, and gain maximum benefit when all of their providers are VIPzone members as this reduces the scope for hijacks out of zone. With 100 ASes, as VIPzone members, half of the ASes in the Internet are protected from hijacks within the zone.

6.2 Leveraging MANRS to Launch VIPzone

Surprisingly, the MANRS consortium already represents a substantial coherent topological region, which could bootstrap a VIPzone. On 1 December 2022, MANRS had 747 ISP members (as well as 22 MANRS CDN member organizations) participating with 906 and 25 ASNs, respectively [33]. (Some ISPs participate with more than one AS.) To see how many of these are part of a coherent region, we started with the 8 Tier-1 providers (§2.1) that are MANRS members, and recursively examined their customers' ASes to identify which were also members of MANRS. Using CAIDA's AS2Org data set (the most recent version at time of writing, 1 Oct 2022) and AS Relationship data (1 Dec 2022) [12], we found a connected region with 452 members and 563 ASNs. This is fewer than the number of MANRS members, as not all members obtain transit from at least one other MANRS member. There are 25,330 customer ASes (owned by 22,854 organizations) directly connected to this region.

If all the MANRS members in the currently connected region implemented the enhanced VIPzone practices, about one-third of the ASes in the Internet, those that are *directly connected* to a member of the resulting zone today, would have their prefixes protected from hijacks in the zone. Further, there would be incentive for the VIPzone to expand. ASes that want their announcements protected but whose transit provider is not a VIPzone member may be able to connect to another transit provider that is a member, or the transit provider itself might commit to the enhanced practices and become a member in their own right.

The Internet Society has recently launched the MANRS+ initiative to consider how to enhance the existing MANRS framework [34] to improve protection against routing security threats. We consider the VIPzone proposal to be a feasible candidate enhancement.

6.3 U.S. Case Study

Motivated by the FCC's recent Notice of Inquiry [73], we use the U.S. as a case study to explore a hypothetical deployment scenario and its residual risks of hijacks. We try to estimate what fraction of ASes are protected, and the character (size, routing footprint) of those that are not protected. For this analysis, we consider the actual Internet AS-level topology, including provider-customer and peering relationships. Relevant aspects of this topology, such as the AS path length from the periphery to core (Tier 1) ISPs, varies by region and over time. Deployment of our scheme would likely induce changes in how AS operators configure their (transit and peering) interconnections.

Assume that the MANRS program expanded such that any U.S. network operator (AS) with more than 100 customer ASes (aggregating across all the ASes that belong to that organization) became part of the VIPzone. What would protection of the U.S. region of the Internet look like in this case? We use data from CAIDA's AS Rank service (1 December 2022) to estimate the number and types of ASes receiving protection from this scenario. Note that these numbers rely on a current publicly observable state of the Internet's AS topology, which changes as ISPs change their interconnection patterns. They are only an approximation to frame the discussion.

6.3.1 U.S. ASes with only transit interconnections

Stipulating that all ASes with greater than 100 customer ASes participate in VIPzone practices, Figure 8 depicts the distribution of ASes with different levels of protection. An estimated 542 U.S. ASes would be in the VIPzone, and another 14,219 customer ASes (in the teal section) would also benefit from protection due to their direct transit connection to a provider in the VIPzone. These 14,219 ASes include those with the interconnection arrangements described in Table 1. The first row (12,707 ASes) represents customers who only have transit providers in the zone. The second row (1,512 ASes) is customers that get transit service both from provider(s) in the VIPzone and from provider(s) directly connected to the zone (C_3 in Figure 8). These ASes are protected, but ASes in the VIPzone will always prefer the prefixes of that customer that are directly announced to the VIPzone.

The remaining set of ASes do not get protection, although the stub ASes that are two or more hops away from the zone

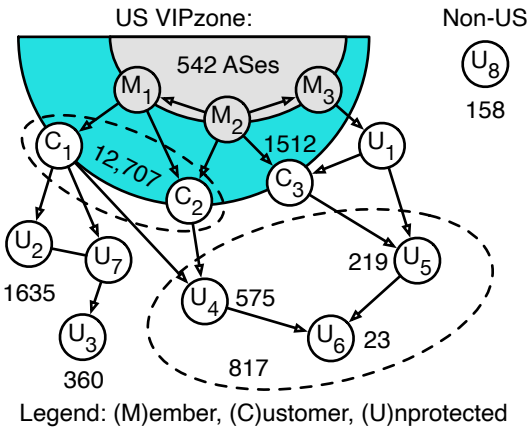

Figure 8: The shape of provider-customer connections in the U.S. region of the Internet if all organizations with more than 100 AS customers were in the VIPzone.

(1,635 + 360) could acquire protection by choosing a transit provider in the zone, which provides incentive for their current transit provider to join the zone. For the 817 multi-homed ASes that are more than one hop from the zone (fifth row), the hijack risk also remains. In our observation period, these ASes were small institutions, e.g., a business or school. A policy question is whether the risk of hijacks of these small targets is a material concern. The 158 ASes that received transit from a non-U.S. provider were mostly foreign subsidiaries of U.S. firms. To understand what protection the VIPzone could secure for these ASes would require a hypothetical scope for the VIPzone outside the U.S., which we do not attempt. Finally, 99 U.S. ASes had no inferred transit provider, so we did not consider them in this analysis.

Connection type	Number
Customers with only transit provider(s) in VIPzone (C_1 and C_2)	12,707
Customers with mixed transit provider types (C_3)	1,512
Total Protected (82.7%):	**14,219**
Single homed stub that are two hops from VIPzone (U_2)	1,635
Single homed ASes that are more than two hops from VIPzone (U_3)	360
Multi-homed ASes more than one hop from the VIPzone (U_4, U_5, U_6)	817
ASes with transit from non-U.S. provider (U_8)	158
Total Unprotected (17.3%):	**2,970**
Total U.S. ASes	**17,189**

Table 1: Protection based on interconnection arrangements of ≈17K U.S. ASes if all U.S. ASes with more than 100 customer ASes were required to join the VIPzone. Note that a metric of the VIPzone's success would be that ASes changed their interconnection arrangements to increase protection coverage offered by the zone. Thus we intend this analysis only to illustrate the coverage benefits of VIPzone deployment.

6.3.2 U.S. ASes that peer outside zone

Figure 8 depicts only transit links, i.e., between a customer and transit provider; peering connections are so dense that we cannot usefully visualize them. But as mentioned in §5.2, peering connections have implications for protection against hijacks. As a hypothetical case study, we examined the U.S. ISP with over 100 AS-level customers and the largest number of peering connections: Hurricane Electric (HE). In our data set, HE had 2079 AS-level customers, and 7677 peers globally, of which 699 were in the U.S. 76 Of these 699 were in the VIPzone, 536 were directly connected customers of ASes in the VIPzone, and 30 were ASes more than one hop from the zone. If HE were in the zone, when HE receives announcements from its peers not in the zone, HE should mark them VERIFIED or not based on VIPzone practices §4.2.

The complication for an AS in the VIPzone with peers that are not in the zone is that some ASes in the customer cone of one of their peers may also have a connection to another provider that is in the VIPzone (see §5.2.3 and Figure 6). HE's rich peering practices yield 2629 ASes it can reach both through a peer not in the VIPzone (and thus using announcements that are not VERIFIED) and through a peer or customer in the VIPzone, using a path that is VERIFIED. We assume that HE would prefer to reach these ≈2K ASes over a peering connection, but for most of them (all but 352) HE has multiple routes via peers to those ASes, one of which is a peering path inside the VIPzone. This outcome arises from the rich peering relationships used by HE. Only for the remaining 352 ASes would HE have to decide to forego the peering path and prefer the VERIFIED option, or undertake to get accurate knowledge about the customers of that peer, so that it could implement an exception to the preference rule and prefer the announcement from the peer that is not VERIFIED.

Our proposal generally assumes it is not practical for an AS to have accurate knowledge of the customers of its customers, or in this case its peers. However, it may be practical in specific cases. If HE were to prefer this route without having information about the customers of the peer, it would put its own customers at risk of a hijack originating in that peer.

Note that if Hurricane and Cogent agreed to peer with each other, Hurricane would get access via the VIPzone to 220 of the 352 ASes we consider here. Business rather than technical decisions will determine availability of VIPzone protections to customers. Also, deployment of the practices would likely lead to changes in interconnection practices, and thus the Internet topology.

7 Comparison to Other Proposed Solutions

A recent IETF draft proposal uses a community value to allow networks to signal a likely route leak [65] is similar in spirit to our proposal but does not enable a trusted routing zone that incentivizes deployment.

From zero trust to zone of trust. First, we acknowledge a tension between our approach and the philosophy of *zero trust architectures*. Zero trust is usually proposed in a context where each machine or subsystem performs its own verification to protect itself, and the incentives are directly aligned [60]. The collective action aspect of routing security is at odds with this assumption. A zero-trust approach is inauspicious if the incentives of the implementors are not aligned, and if it is not feasible to verify implementation. The VIPzone approach better aligns incentives, and allocates responsibility to specific points in the zone (the perimeter), and has conformance checking to ascertain whether the trusted zone members are properly implementing the required operational practices.

AS Provider Authorization ASPA is more of a zero trust approach: it allows any AS anywhere to inspect any BGP announcement for an invalid route assertion, while VIPzone requires this test only at the boundary between the MANRS zone and non-MANRS customer. However, exactly because any ISP can do this verification, it is not immediately clear which ISPs *should* do it, or that anyone *will* do it. As with ROAs, just because an ISP has registered its prefixes or its providers, there is no guarantee that any ISPs will use this information to filter routes. The VIPzone approach clearly allocates responsibility for checking to specific points in the zone The VIPzone design adds clarity to who should do what when.

ASPA also requires a global database of intended AS neighbors. Some ASes may not want to disclose their potential providers (beyond those revealed in BGP). ASPA would require that all potential providers register, so that route changes would be effective. The VIPzone approach relies on sharing only BGP routing data such as RouteViews collects already. Information that validates the provider-customer relationship can remain a local matter between those two ASes, who can translate it into a more general signal (the VERIFIED community value), which propagates through the VIPzone.

Finally, ASPA specifies customer-provider relationships at the AS granularity [8]. It does not record which customer AS connects to a provider AS at which interconnection point. A large AS (e.g., a global transit provider) likely has many ASes as customers, so an AS posing as one of those customers at one interconnection point would appear to be valid based on the ASPA assertion. As such, ASPA is most useful in a local context, such as for detecting route leaks [8].

BGPSEC Like ASPA, BGPsec is closer to a zero trust approach–every router that forwards the announcement must add its own cryptographic signature, and any router along the path can verify for itself that the series of signatures are valid. Again, one concern is that the lack of incentive for ISPs to undertake this extra effort may lead to inconsistent and unpredictable implementation of the required checking. The complexity and overhead of BGPsec has meant that to date there is no deployment of the scheme in practice.

8 Conclusion

There is currently no consensus as to the next step to secure BGP beyond the simplest type of hijacks. As of 2023, BGPsec has no production deployment, and arouses significant controversy over the operational feasibility of its key management aspects. For all proposed solutions to prevent path hijacks, incentives are badly misaligned: ISPs must bear the cost of action, but users bear the cost of inaction. And even those who are willing to invest their way to increased routing security cannot effectively do so due to the collective-action characteristics of the problem. In particular, the current MANRS framework has nothing to offer directly to ASes that want protection from having their addresses hijacked, or having their traffic to a distant prefix hijacked.

We have proposed an approach to overcome this obstacle, enabling a path forward against both origin and path hijacks. One insight that shapes our proposal is that if there is a coherent topological region of the Internet, and with practices limiting malicious BGP routes entering that region, then the operational practices, if enforced, can provide much stronger protection against abuse for those who join, and thus incentive to participate. The result is a virtuous circle, where customers benefit from choosing ISPs committed to the practices, and ISPs (thus) benefit from committing to the practices.

Our approach relies on existing capabilities and institutions: RPKI and ROV capabilities; the existing MANRS framework, and current techniques for collecting and analyzing interdomain (BGP) topology data. There is already a coherent core of ISPs that emerges organically in the ecosystem, which we can leverage to create a *zone of trust*, a region that protects not only all networks in the region, but *all directly attached customers*. Achieving this protection requires auditing and enforcing conformance with the practices. Conformance checking will bring additional costs, but the institutional framework required for such checking already exists in multiple places, e.g., RIPE and RouteViews.

Our proposal responds to a long-standing need for some medium-term path forward on protection against path hijacks. We believe it is a direction worth serious debate and further analysis in the context of possible regulatory measures.

References

[1] IRR - Internet Routing Registry. http://www.irr.net.

[2] BGP incidents, mitigation techniques and policy actions, 2022. https://www.oecd-ilibrary.org/science-and-technology/routing-security_40be69c8-en.

[3] Security of the Internet's Routing Infrastructure, 2022. https://www.bitag.org/documents/BITAG_Routing_Security.pdf.

[4] American Registry of Internet Numbers. Response to NOTICE OF INQUIRY. PS Docket No. 22-90. In the Matter of Secure Internet Routing, 2023. https://www.fcc.gov/ecfs/document/10120103512712/1.

[5] APNIC. Creating route objects via email updates, Jan 2023. https://www.apnic.net/manage-ip/using-whois/guide/creating-route-objects/.

[6] M. Apostolaki, A. Zohar, and L. Vanbever. Hijacking Bitcoin: Routing Attacks on Cryptocurrencies. In *2017 IEEE Symposium on Security and Privacy (SP)*, May 2017.

[7] A. Azimov, E. Bogomazov, R. Bush, K. Patel, and J. Snijders. Verification of AS_PATH Using the Resource Certificate Public Key Infrastructure and Autonomous System Provider Authorization. https://www.ietf.org/archive/id/draft-ietf-sidrops-aspa-verification-01.txt, July 2019.

[8] A. Azimov, E. Bogomazov, R. Bush, K. Patel, J. Snijders, and K. Sriram. BGP AS_PATH Verification Based on Resource Public Key Infrastructure (RPKI) Autonomous System Provider Authorization (ASPA) Objects. Internet-Draft draft-ietf-sidrops-aspa-verification-11, Internet Engineering Task Force, Oct 2022. https://datatracker.ietf.org/doc/draft-ietf-sidrops-aspa-verification/11/.

[9] R. Barnes, B. Schneier, C. Jennings, T. Hardie, B. Trammell, C. Huitema, and D. Borkmann. Confidentiality in the face of pervasive surveillance: A threat model and problem statement. RFC 7654, Aug 2015.

[10] H. Birge-Lee, J. Wanner, G. H. Cimaszewski, J. Kwon, L. Wang, F. Wirz, P. Mittal, A. Perrig, and Y. Sun. Creating a Secure Underlay for the Internet. In *USENIX Security Symposium*, Aug 2022.

[11] J. Borkenhagen. AT&T/as7018 now drops invalid prefixes from peers. https://mailman.nanog.org/pipermail/nanog/2019-February/099501.html.

[12] Bradley Huffaker, Matthew Luckie, and kc claffy. CAIDA Autonomous System Rankings ASRank. https://asrank.caida.org/.

[13] K. Butler, T. R. Farley, P. McDaniel, and J. Rexford. A Survey of BGP Security Issues and Solutions. *Proceedings of the IEEE*, Jan 2010.

[14] R. Chandra, P. Traina, and T. Li. BGP Communities Attribute. IETF RFC 1997, 1996.

[15] CISCO. Remotely Triggered Black Hole Filtering - Destination Based and Source Based. Cisco White Paper, http://www.cisco.com/c/dam/en_us/about/security/intelligence/blackhole.pdf, 2005.

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

[17] Dan Goodin. How 3 hours of inaction from Amazon cost cryptocurrency holders $235,000, Sep 2022. https://arstechnica.com/information-technology/2022/09/how-3-hours-of-inaction-from-amazon-cost-cryptocurrency-holders-235000/.

[18] B. Donnet and O. Bonaventure. On BGP Communities. *ACM CCR*, 38(2):55–59, Mar 2008.

[19] B. Du, G. Akiwate, T. Krenc, C. Testart, A. Marder, B. Huffaker, A. C. Snoeren, and K. Claffy. IRR Hygiene in the RPKI Era. In *PAM*, pages 321–337, 2022.

[20] B. Du, C. Testart, R. Fontugne, G. Akiwate, A. C. Snoeren, and k. claffy. Mind Your MANRS: Measuring the MANRS Ecosystem. In *ACM Internet Measurement Conference*, 2022.

[21] L. Gao. On Inferring Autonomous System Relationships in the Internet. *IEEE/ACM Trans. Networking*, 9(6), 2001.

[22] Y. Gilad, S. Goldberg, K. Sriram, J. Snijders, and B. Maddison. The use of maxLength in the resource public key infrastructure RPKI. RFC 9319, Oct 2022.

[23] Y. Gilad, O. Sagga, and S. Goldberg. MaxLength Considered Harmful to the RPKI. In *Conference on Emerging Networking EXperiments and Technologies*, CoNEXT '17, 2017.

[24] P. Gill, M. Schapira, and S. Goldberg. Let the market drive deployment: A strategy for transitioning to BGP security. *ACM SIGCOMM Computer Communication Review*, 41(4):14–25, aug 2011.

[25] S. Goldberg. Why is It Taking So Long to Secure Internet Routing? *Comm. of the ACM*, 57(10), 2014.

[26] G. Goodell, W. Aiello, T. Griffin, J. Ioannidis, P. D. McDaniel, and A. D. Rubin. Working around BGP: An Incremental Approach to Improving Security and Accuracy in Interdomain Routing. In *ISOC Symposium on Network and Distributed Systems Security*, 2003.

[27] D. Goodin. Suspicious event hijacks Amazon traffic for 2 hours, steals cryptocurrency, Apr 2018. https://arstechnica.com/information-technology/2018/04/suspicious-event-hijacks-amazon-traffic-for-2-hours-steals-cryptocurrency/.

[28] M. G. Gouda, E. N. Elnozahy, C.-T. Huang, and T. M. McGuire. Hop integrity in computer networks. *IEEE/ACM Transactions on Networking*, 10(3), Jun 2002.

[29] A. Heffernan. RFC 2385: Protection of BGP Sessions via the TCP MD5 Signature Option, Aug 1998.

[30] Y.-C. Hu, A. Perrig, and M. Sirbu. SPV: Secure path vector routing for securing BGP. *ACM SIGCOMM Computer Communication Review*, 34(4), 2004.

[31] G. Huston and G. Michaelson. RFC 6483: Validation of Route Origination Using the Resource Certificate Public Key Infrastructure (PKI) and Route Origin Authorizations (ROAs), Feb 2012.

[32] Internet Society. Mutually Agreed Norms for Routing Security (MANRS). https://www.manrs.org/.

[33] Internet Society. Mutually Agreed Norms for Routing Security (MANRS) network operator participants. https://www.manrs.org/netops/participants/.

[34] Internet Society. MANRS+ Working Group Charter, 2002. https://www.manrs.org/about/manrs-working-group-charter/.

[35] J. Israr, M. Guennoun, and H. T. Mouftah. Credible BGP – Extensions to BGP for Secure Networking. In *Fourth International Conference on Systems and Networks Communications*, Sep 2009.

[36] Jason Livingood, May 2021. https://corporate.comcast.com/stories/improved-bgp-routing-security-adds-another-layer-of-protection-to-network.

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

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

[39] T. King, C. Dietzel, J. Snijders, G. Doering, and G. Hankins. BLACKHOLE community. RFC 7999, Oct 2016.

[40] KPN. AS286 Routing Policy. https://as286.net/AS286-routing-policy.html.

[41] M. Lepinski and K. Sriram. RFC 8205: BGPsec Protocol Specification, Sep 2017.

[42] M. Luckie, B. Huffaker, A. Dhamdhere, V. Giotsas, and k claffy. AS relationships, customer cones, and validation. In *ACM IMC*, pages 243–256, Oct 2013.

[43] R. Lychev, S. Goldberg, and M. Schapira. BGP security in partial deployment: is the juice worth the squeeze? In *ACM SIGCOMM*, pages 171–182, Aug 2013.

[44] MANRS. MANRS Observatory. https://www.manrs.org/manrs-observatory/.

[45] 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, Jun 2018.

[46] National Institute for Stantards and Technology. RPKI Deployment Monitor. https://rpki-monitor.antd.nist.gov/.

[47] M. O. Nicholes and B. Mukherjee. A survey of security techniques for the Border Gateway Protocol (BGP). *IEEE Communications Surveys Tutorials*, 11(1):52–65, 2009.

[48] NIST. NIST RPKI Monitor 2.0: Methodology and User's Guide, 2022. https://rpki-monitor.antd.nist.gov/Methodology#ROV_Donut.

[49] O. Nordstrom and C. Dovrolis. Beware of BGP attacks. *ACM CCR*, 34(2), 2004.

[50] L. Oliver, G. Akiwate, M. Luckie, B. Du, and k. claffy. Stop, DROP, and ROA: Effectiveness of Defenses through the Lens of DROP. In *ACM Internet Measurement Conference*, 2022.

[51] Ostap Efremov. 196.52.0.0/14 revoked, cleanup efforts needed. RIPE NCC Anti-Abuse Working Group, 2021.

[52] P. Paganini. BGP hijacking - Traffic for Google, Apple, Facebook, Microsoft and other tech giants routed through Russia. https://securityaffairs.co/wordpress/66838/hacking/bgp-hijacking-russia.html, Dec 2017.

[53] J. Qiu and L. Gao. Hi-BGP: A Lightweight Hijack-proof Inter-domain Routing Protocol, 2006.

[54] A. Ramachandran and N. Feamster. Understanding the network-level behavior of spammers. In *ACM SIGCOMM Computer Communication Review*, volume 36. ACM, 2006.

[55] P. Reynolds, O. Kennedy, E. G. Sirer, and F. B. Schneider. Using External Security Monitors to Secure BGP. *IEEE/ACM Transactions on Networking*, 2006.

[56] RIPE NCC. Updating the RIPE database, Aug 2022. https://www.ripe.net/manage-ips-and-asns/db/support/updating-the-ripe-database#email-updates.

[57] RIPE Network Coordination Centre. YouTube Hijacking: A RIPE NCC RIS case study. https://www.ripe.net/publications/news/industry-developments/youtube-hijacking-a-ripe-ncc-ris-case-study, Mar 2008.

[58] Ronald F. Guilmette. Cogent & FDCServers: Knowingly aiding and abetting fraud and theft?, 2019.

[59] E. Rosen. Exterior Gateway Protocol (EGP), RFC 827, Oct 1982. DOI 10.17487/RFC0827.

[60] Scott Rose and Oliver Borchert and Stuart Mitchell and Sean Connelly, 2020. https://www.nist.gov/publications/zero-trust-architecture.

[61] J. Scudder, R. Bush, P. Mohapatra, D. Ward, and R. Austein. RFC 6811: BGP Prefix Origin Validation, Jan 2013.

[62] Security and Stability Advisory Committee. SSAC Briefing on Routing Security, 2022. https://www.icann.org/en/system/files/files/sac-121-en.pdf.

[63] M. S. Siddiqui, D. Montero, R. Serral-Gracia, X. Masip-Bruin, and M. Yannuzzi. A survey on the recent efforts of the Internet Standardization Body for securing inter-domain routing. *Computer Networks*, 80:1–26, Apr 2015.

[64] B. R. Smith and J. J. Garcia-Luna-Aceves. Securing the Border Gateway Routing Protocol. In *Global Telecommunications Conference*, Nov 1996.

[65] K. Sriram and A. Azimov. Methods for Detection and Mitigation of BGP Route Leaks. Internet-Draft draft-ietf-grow-route-leak-detection-mitigation-08, Internet Engineering Task Force, Oct 2022. Work in Progress.

[66] L. Subramanian, V. Roth, I. Stoica, S. Shenker, and R. H. Katz. Listen and Whisper: Security Mechanisms for BGP. In *Symposium Networked System Design and Implementation,*, 2004.

[67] Telia. Telia Carrier Takes Major Step to Improve the Integrity of the Internet Core. https://www.businesswire.com/news/home/20190915005013/en/.

[68] C. Testart. Reviewing a Historical Internet Vulnerability: Why Isn't BGP More Secure and What Can We Do About it? In *Telecommunications Policy Research Conference*. SSRN, Aug 2018.

[69] A. Toonk. Using BGP data to find Spammers, Sep 2014. https://bgpmon.net/using-bgp-data-to-find-spammers/.

[70] J. Touch, A. Mankin, and R. P. Bonica. RFC 5925: The TCP Authentication Option, Jun 2010.

[71] U.S. Department of Homeland Security Cybersecurity and Infrastructure Security Agency (CISA). NOTICE OF INQUIRY. PS Docket No. 22-90. In the Matter of Secure Internet Routing, Feb 2022. https://www.fcc.gov/ecfs/document/1022806680214/1.

[72] U.S. Department of Justice National Security Division (Matthew G. Olsen) and U.S. Department of Defense Acquisition and Sustainment (William A. Laplante). Public Comment to NOTICE OF INQUIRY. PS Docket No. 22-90. In the Matter of Secure Internet Routing, Sep 2022. https://www.fcc.gov/ecfs/document/1091496862125/1.

[73] U.S. Federal Communications Commission. NOTICE OF INQUIRY. PS Docket No. 22-90. In the Matter of Secure Internet Routing, Feb 2022. https://www.fcc.gov/ecfs/document/1022806680214/1.

[74] P.-A. Vervier, O. Thonnard, and M. Dacier. Mind Your Blocks: On the Stealthiness of Malicious BGP Hijacks. In *Proceedings 2015 Network and Distributed System Security Symposium*, San Diego, CA, 2015.

[75] T. Wan, E. Kranakis, and P. C. van Oorschot. Pretty Secure BGP, psBGP. In *Proceedings of the 2005 ISOC Symposium on Network and Distributed Systems Security*, San Diego, 2005.

[76] R. White. Securing BGP Through Secure Origin BGP - The Internet Protocol Journal - Volume 6, Number 3. *The Internet Protocol Journal*, 6(3), Sep 2003.

[77] White Ops and Google. The Hunt for 3ve: Taking down a major ad fraud operation through industry collaboration. Technical report, Google, Nov 2018.

[78] X. Zhao, D. Pei, L. Wang, D. Massey, A. Mankin, S. F. Wu, and L. Zhang. Detection of invalid routing announcement in the Internet. In *Proceedings International Conference on Dependable Systems and Networks*, 2002.