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

David Clark, Cecilia Testart,Matthew Luckie, kc claffy

Abstract

Although Internet routing security best practices have recently seen auspicious increases in uptake, ISPs have limited incentives to deploy them. They are operationally complex and expensive to implement and provide little competitive advantage. The practices with significant uptake protect only against origin hijacks, leaving unresolved the more general threat of path hijacks. We propose a new approach to improved routing security that achieves four design goals: improved incentive alignment to implement best practices; protection against path hijacks; expanded scope of such protection to customers of those engaged in the practices; and reliance on existing capabilities rather than needing complex new software in every participating router. Our proposal leverages an existing coherent core of interconnected ISPs to create a *zone of trust*, a topological region that protects not only all networks in the region, but *all directly attached customers* of those networks. Customers benefit from choosing ISPs committed to the practices, and ISPs thus benefit from committing to the practices. We discuss the concept of a zone of trust as a new, more pragmatic approach to security, that improves security in a region of the Internet, as opposed to striving for a global improvement. We argue that the aspiration for global improvement is unrealistic, since the global Internet includes malicious actors. We compare our approach to other schemes, and discuss how a related proposal, ASPA, could be used to increase the scope of protection our scheme achieves. We hope this proposal inspires discussion of how the industry can make practical, measurable progress against the threat of route hijacks in the short term by leveraging institutionalized cooperation rooted in transparency and accountability.

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 (a *prefix*) that it does not in fact have the authority to announce. Routers that accept a forged route announcement – known as a *route hijack* – will route traffic intended for addresses in that block to the rogue AS. 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. 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.

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 [1, 2, 3, 4] or to inspect [5] 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 [6, 7, 8], 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 best currently available practices in routing security require two steps to identify and block propagation of bogus route announcements. First, each ISP must register its own address space in a trusted database (ideally, the Resource Public Key Infrastructure aka RPKI) and routers across the Internet must check announcements against such a database and drop those announcements that are not consistent with the registered information (*route origin validation* aka ROV). An AS who engages in the first step gains no security unless other ASes correctly deploy the second (ROV) step. ROV is sufficiently operationally complex that smaller or lower-resourced ISPs are reluctant to risk misconfigurations that

impair their own service availability. Thus, networks take on additional costs and operational risks, but the benefits may not accrue to them or their customers. Even if consistently implemented, which is a lofty aspiration in a global context, these two practices target only the simplest form of hijack, an *origin hijack*. A proposed approach for protection against a broader range of hijacks is BGPsec, an even more complex and expensive protocol-based solution that is at least a decade away from significant operational deployment [9].

Concerns over slow progress on routing security solutions led the U.S. Federal Communications Commission (FCC) to issue a February 2022 Notice of Inquiry into potential regulatory interventions that could reduce the severity of the threat to U.S. networks and traffic [10]. Several U.S. government agencies, including the DHS and a joint filing by the DOD and DOJ, urged the FCC to take action [11, 12]. Other commenters emphasized the risks of regulation in this domain.

Tension is increasing on this topic, as multistakeholder efforts to advance routing security have continued for over a 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 [3]. The scope of the problem is elusive to measure given lack of disclosure – and sometimes lack of awareness of – incidents.

The collective-action characteristic of the problem is fundamental: even those who are willing to invest in order to increase their own routing security cannot achieve protection without commitments from other networks to prevent propagation of bogus routes. We propose a more practical solution that refocuses on a new goal: to *provide a concrete action that a security-aware AS can take to protect itself from both having its address blocks hijacked, and its traffic to other address blocks hijacked.*

We propose an approach that achieves four related goals. First, it aligns incentives of actors toward improved routing security. Second, it offers protection against not only origin hijacks but the larger looming problem of path hijacks. Third, it allows ASes participating in the approach to protect their customers without additional work on the part of the customer, thus allowing highly-resourced ISPs to protect other parts of the Internet. This feature is compelling because in today's Internet the steps necessary to securely configure systems are sometimes complicated, and smaller ASes may not have the skills or resources to undertake them. Even more compelling is the resulting alignment of incentives: customers will prefer a participating provider since they offer enhanced security, giving providers an incentive to participate so they can market their improved security to potential customers. Finally, our approach requires no new capabilities in routers, relying on existing capabilities and institutions, and current techniques for analyzing interdomain (BGP) topology data.

The roadmap of this paper is as follows. We first describe barriers to routing security over the last two decades (§2). We describe the threat model in §3. In §4 we introduce the principles of a *zone of trust*, a connected region of the Internet where providers take enhanced steps to improve the security of that region, including the security of customers connected to providers in the region. We introduce a specific example of a routing zone of trust which offers a more incentive-aligned direction for protecting ASes from both *origin* hijacks and *path* hijacks. We analyze the residual risks of our scheme and how to minimize them (§5), auditing requirements (§6), and comparison to other proposals (§7).

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, TLS), the standards community has grappled with complexities of cryptographic key management, trust anchors, and performance implications that hinder standardization, implementation, and deployment. Over the last 30 years, over 20 proposals to secure BGP have come from academia, industry and the Internet Engineering Task Force (IETF), some of which Figure 1 highlights. We describe 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. Internet traffic will flow back to the destination prefix following the AS path. To prevent 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. Operators use this property for traffic engineering. BGP also provides a mechanism

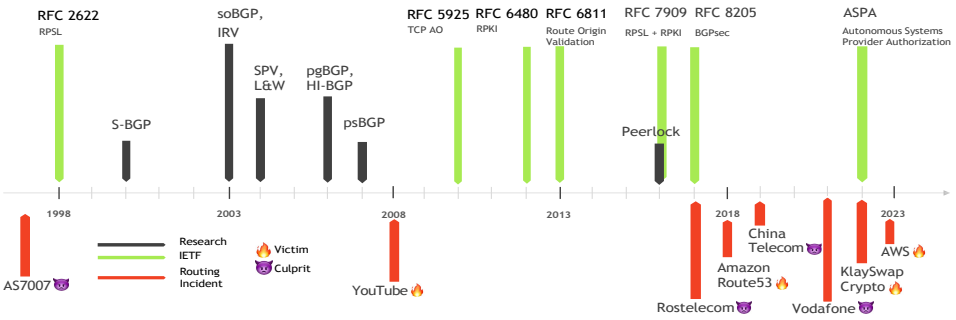

Figure 1: Decades of proposed routing security approaches; sample of high-profile hijacks.

to annotate announcements with attributes – known as *BGP communities* [13] – to enable signaling within and across ASes, facilitating traffic engineering innovations [14] such as automated blocking of denial-of-service attack traffic on the path to the victim [15, 16].

There are two general 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 two peers exchange routes to their customers without involving an intermediate provider [17]. If an AS has multiple routes to the same prefix, the rational choice is to prefer routes received from customers (a source of revenue), over routes received from peers (typically settlement-free, i.e., no cost), over routes received from providers (which cost the AS) [17]. Other ASes that an AS X can reach through a customer link are within the *customer cone* of X.

A few (≈15) large ASes obtain global routing using routes received only from their peers and customers, i.e., they do not pay any transit providers. These ASes connect in a full mesh (a peering clique) that enables packet delivery between arbitrary networks with different transit providers. The ASes in this group that also do not pay for peering are known as Tier-1 providers. However, payments between ASes are confidential; we thus use the term *Tier-1* to refer to the peering clique.

2.2 Routing Security in the 1980s

In 1982, Rosen [18] documented that it is possible to corrupt interdomain routing in RFC 827, in the context of a predecessor 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.

2.3 Routing Security in the 1990s

To mitigate the prevalent risk of accidental misconfigurations, in the 1990s network operators developed the *Internet Routing Registry (IRR)* system of distributed databases. The IRR system enabled network operators to publish address ownership and routing policy information [19], which other operators could use to build filters that permit or deny routes according to these operator-registered policies. Unfortunately, some IRR databases do not fully authenticate registration data, allowing attackers to compromise the IRR by falsely registering ownership of resources which they then use in a hijack [20, 21, 22, 23, 24].

When the IETF started to study malicious BGP security threats, 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[25, 26].

2.4 Routing Security in the 2000s: proposals for a secure BGP

Aiming for a more complete 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* – ensuring that each AS appearing in a received AS path was legitimately in the path. During this decade over a dozen competing approaches came out of academia and industry [27, 28, 29, 30, 31, 32, 33, 34, 35, 36, 37, 38, 39, 40]. The protocol that became an IETF standard (RFC 8205) in 2017 is called BGPsec [41]. BGPsec update messages include two important new fields: the AS to which the router is sending that announcement; and a cryptographic signature over the message that enables any router along the path to verify that the series of signatures are valid. This mechanism prevents path hijacks: a malicious AS cannot forge the AS path because the malicious AS cannot sign records for the forged ASes. Cryptographic attestation of paths requires propagation of a new layer of cryptographic transaction at each hop, which is computationally expensive and 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, full protection of the path requires every AS along the path to implement BGPsec. Partial deployment, inevitable during a transition, implies unpredictable protection. The complexity, overhead, and misaligned incentives have prevented significant operational deployment of BGPsec, despite a decade-long standardization process that completed in 2017 [41].

2.5 Routing Security in the 2010s

The 2010s brought three areas of endeavor: rigorous analyses of the incentives to deploy routing security solutions; technology, standardization, and operational mechanisms to mitigate the simpler problem of origin hijacks; and a collective action effort (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 had failed to gain traction [42, 43]. Such reviews continued throughout the subsequent decade [44, 45, 46]. 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 [47, 48, 49, 50]. The deepest body of work on this topic was by Sharon Goldberg and Michael Schapira and their 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 other routes to the same prefix [47]. 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. Followup work two years later [48, 49] acknowledged that having Tier 1 ISPs lead a market-driven deployment would not work because economic incentive would override any secure route received from a peer when an insecure route via a customer is available.

In 2011, researchers proposed a new Internet architecture, SCION [51], that separated ASes into independent trust domains which provide isolation of routing failures and human misconfiguration. Researchers recently used SCION to bootstrap a secure routing system [52]. SCION is now being used commercially to secure the routing among sets of locations on the Internet. SCION assumes a hierarchical architecture, where one or more highly trusted ASes connect the domains to each other.

In 2016, Cohen *et al.* [50] proposed an approach similar to the recently proposed ASPA protocol (see §7). Their simulations focused on the length of paths that an attacker must construct if the AS announcing the prefix has registered what we today call an ASPA. The authors discussed deployments in select geographic regions, perhaps driven by government pressure. They did not propose a connected region, so partial deployment of this approach yields only a probabilistic assessment of protection, as with ROV (and ASPA).

2.5.2 Preventing origin hijacks: RPKI and ROV

While BGPsec has been undergoing implementation and evaluation for a decade, operators have focused on 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) [53]. ROV uses a Resource Public Key Infrastructure (RPKI) [54], an authoritative database maintained outside of BGP that closely matches Internet resource delegation by the Regional Internet Registries (RIRs). Each of the 5 RIRs is the root of trust enabling the holders of the IP addresses blocks delegated by them to issue*Route Origin Authorization* (ROAs). ROAs are cryptographic signatures that authorize designated ASes to originate routes to address blocks. Routers using ROV drop BGP announcements that are not consistent with a registered ROA for the prefix. RFC 6811 [53] 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. This residual risk includes the forged-origin path hijack mentioned above, where the malicious AS impersonates the valid source AS by appending it to a forged BGP announcement (recently observed in the wild [24]). RFC 6811 thus cautioned: "*..this system should be thought of more as a protection against misconfiguration than as true 'security' in the strong sense.*"

Use of ROAs presents other operational challenges. A ROA contains a set of prefixes and a set of origin ASNs. For a BGP route to be valid according to RPKI, there needs to be a ROA with a prefix that is equal to or covers the route prefix, the ASN originating the BGP route has to be in the allowed set, and the length of prefix needs to be allowed. Indeed, ROAs may also contain a *maxLength* attributes that defines the maximum prefix length allowed for each prefix; for example, a ROA for 192.0.30.0/23 with a maxLength of 24 enables the AS to originate 192.0.31.0/24. Operators use this feature for traffic engineering (see the discussion of Interdomain Routing in §2). In 2017, Gilad *et al.* showed that use of the maxLength attribute could enable an attacker to hijack more-specific prefixes that victim networks then unwittingly communicate with. Best current practice is to not use the maxLength attribute [55]. Furthermore, if a BGP route has a prefix covered by a ROA but the route is not valid either because the origin AS in not on the allowed set or the prefix length is not allowed, the route is invalid according to RPKI. As a consequence, if a network provider registers a ROA for a large prefix (*e.g.,* a /16), any sub-delegation to another (smaller) network will be covered by the ROA and routes to those smaller prefixes may be considered invalid. RPKI requires coordination between network operators to prevent making routes unreachable.

Although RIRs have supported RPKI registration of ROAs since 2013, until 2019 there was little evidence of Internet Service Providers (ISPs) using ROAs to validate BGP announcements. But by late 2022, many large ISPs, including AT&T, KPN, Arelion, and Comcast had started to use ROV to drop invalid announcements [56, 57, 58, 59]. According to NIST's public RPKI monitor based on RouteViews data [60], as of May 2024, 51% of IPv4 /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 May 24, 2024, NIST reported 70% of observed prefix-origin pairs in the RIPE region were valid, 58% in LACNIC, 51% for APNIC, 38% for ARIN, and 30% in the AFRINIC region [61]. Recent work examining the state of ROV deployment in the Internet between December 2021 and September 2023 reported that 12.3% of tested ASes had behavior suggesting that they or all of their transit providers had consistently implemented ROV [62]. They reported that larger ASes (i.e., those networks with technical capacity) were more likely to have implemented ROV. APNIC has a live ROV measurement based on the availability of a path to a prefix that switches RPKI status every 2 to 3 days [63], and reports results at a per-country level in a world map [64]. In May 2024, the world map shows the disparities of ROV adoption: many countries including the US, France, Spain, Sweden, Finland and Australia have over 50% adoption; and others such as Russia, China, Brazil and Mexico still have less than 10% ROV adoption. Using BGP and RPKI data to infer when networks drop invalid announcements following the methodology from [65], we measure that about 60% of the 360 networks sharing their data with BGP collectors have adopted ROV. Tier-1 and large networks in the US and Europe are more likely to share BGP data with collectors. 14 out of 17 Tier-1 networks measured with that methodology have deployed ROV.

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) [66]. MANRS specified 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: (1) it must confirm that the customer has used an ASN that it is legitimately allowed to use; and (2) 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 [10], the MANRS program has had limited success. In May 2024, MANRS had 938 ISP and 30 CDN organizational members that covered 1268 and 30 ASNs, respectively [67]. This constitutes 1.7% of the ≈75K routed ASes. Many of the largest ISPs do not participate, and some participating ISPs are not conforming to the practices. Du *et al.* reported that 5% of MANRS ISPs did not conform with the requirement to register their prefixes in either RPKI or IRR as of May 2022 and 16% did not conform with the filtering requirement [68].

The limited success of MANRS (and its underlying practices) 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. Some network operators believe this benefit does not justify the cost and complexity of RPKI/ROV.

Third, there is insufficient auditing of conformance to lend confidence in assuming consistent implementation [69]. Independent auditing has detected significant non-conformance [68]. 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 [70]. As of June 2024, ASPA is still in IETF development. ASPA builds on presumed use of RPKI and ROV but enables customer ASes to go further by registering a list of their transit providers in the globally visible RPKI database. That database allows any AS to examine a BGP announcement to detect and reject many types of invalid path announcements, so long as the ASes along the path have registered their providers in ASPA. The authors describe ASPA as preventing route leaks as well as some forms of path hijacks; it does not prevent an attacker from spoofing a sequence of ASes in the path if those ASes do not implement ASPA. Our proposed scheme provides a more predictable level of protection and improves incentives for deployment. We compare our scheme to ASPA, and describe a way in which the use of ASPA could expand the protection provided by both approaches (§7).

2.6 Routing Security in the 2020s

This decade, routing security caught the attention of regulators. Researchers discovered hijacks of unannounced address space [71], and forged-origin hijacks of RPKI-valid address space [24]. After earlier hijacks of AWS address space [2] motivated Amazon to register ROAs for most of its address blocks, attackers developed more sophisticated path hijacking techniques. The high-profile hijack of AWS space in August 2022 [3] motivated by the opportunity to steal cryptocurrency, 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 hijack 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 prefix, the attacker's path would have been preferred for networks that were customers of AS1299 who did not have a more-preferred route to Amazon.

The persistent failure of market-driven solutions to routing security has recently triggered government interest and inquiry into potential interventions. In 2022, the OECD [72], ICANN [73], BITAG [74], and the U.S. FCC [10] all published reports with extensive references related to routing security challenges, and limitations of proposed solutions.

We expect governments to 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. Our approach draws inspiration from Lychev *et al.*'s conclusion a decade ago [48] 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."* [48]

We pursue this challenge with an approach that leverages a coherent topological region to achieve our design goals: incentive alignment, competitive advantage to participating networks; proportional responsibility, in that larger players can invest to protect their customers, providing this competitive advantage; and protection against origin as well as path hijacks without the operational complexity of BGPsec. We believe our proposed alternative is worth open debate before pursuing more blunt regulatory measures.

3 Threat Model

We next describe the capabilities of defenders, to contrast defender capabilities with attacker capabilities.

3.1 Defender Capabilities

As of May 2024, ≈85% of the ≈75K ASes on the Internet have no customers. They are in many cases small ASs with limited operational resources to defend themselves. While they may use peering connections to handle some of their traffic (see §5), they connect to transit providers to reach most parts of the Internet. These transit networks engage in contractual agreements when they interconnect with their neighbors. These transit network operators regularly interact at peering forums and other industry events (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 in general 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.[1] Finally, a defender cannot control the route selection policies of their peers or customers; these ASes might select hijacked routes from other neighbors they have.

3.2 Attacker Capabilities

We assume that the attacker 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. 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 completely hide their activities; in order for their attack to be effective, their hijacked route must propagate, and multiple route collector projects today publish the set of AS paths they collect. Nor does an attacker have the ability to issue ROAs for address space that they do not control. An attacker could compromise the RIR (an insider) or the prefix holder's RIR account, but that is out of scope for the proposed approach as it is a generic and well-understood security problem of systems connected to the Internet.

[1] In particular ARIN's agreement embedded a controversial position that in order to register a ROA, holders of legacy space, (i.e., those allocated before ARIN existed) must contractually agree that they have no legal property rights to their address space [75]. In September 2022, ARIN removed this clause from their Registry Services Agreement [76].

4 A Routing Zone of Trust

We first introduce the concept of a zone of trust in a routing context before specifying one in more detail below in our discussion of the Verified IP Zone. Figure 4 depicts a zone with member providers (in green) at the edge of the zone providing transit service to directly attached customers (white). The providers connect within the zone, and must know when they are exchanging traffic with another member of the zone, and when they are communicating with an AS outside the zone.

A zone could protect against origin hijacks as follows. If all providers P in the zone commit to implement ROV and drop invalid announcements from customers outside the zone, then no invalid announcements will circulate inside the zone, which means that customers C will never receive a BGP announcement from the zone where the origin is invalid based on a ROA. These practices turn this zone into a *zone of trust*.

This example illustrates three properties of a zone of trust:

- *Collective action* by ASes creates the zone and its trust attributes.
- ASes in the zone must have paths that connect them with all hops within the zone.
- Customers of the zone obtain protection by using a provider in the zone. They need take no other action.

We call this region a *zone of trust* because the protection in the zone arises from actions of ASes 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. We introduce this design assumption in exchange for one that routing security protocols have always included: global deployment of a protocol. If ASes themselves are threat actors, we are skeptical of an aspiration to make BGP *globally* secure. Creating a zone of trust through perimeter protection (a trust-but-verify regime) offers a more pragmatic approach for today's routing system.

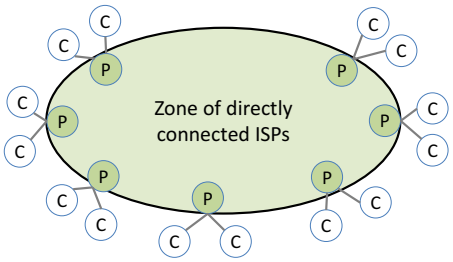

Figure 2: One requirement for our conceptual zone of trust is a coherent topological region with providers in the zone providing transit to customers attached to those providers.

Given the history of routing in the Internet, where each Autonomous System can choose to interconnect based on its own needs, the idea of a coherent perimeter around a zone is missing from today's interdomain routing system. ROV deployment discussions today consider each AS in isolation, leaving security a statistical measure. We can count the number of ASes that register their ROAs, or the number of ASes that implement ROV, but the consequence for a given AS is a function of what other ASes choose to do. It is thus not clear what specific action an AS should take to reduce its own risk profile. Today, invalid announcements may propagate across the Internet, and may or may not reach any given AS. In contrast, a connected zone of trust allows clear articulation of the benefit to a given AS to joining the zone: ASes in the zone will receive no announcements from the zone with an invalid origin based on a registered ROA.

The incentive alignment extends beyond the zone: customers concerned about hijacks can seek out providers that are in the zone, which in turn creates an incentive for providers to commit to the required practices that define the zone and join it. Today, there is little direct benefit to an AS that chooses to implement ROV. Many of the larger ASes do so, as part of a collective action to improve security, but recognizing that these actions can create a coherent zone with direct benefit to their customers will increase their incentive.

Note that the zone does not provide absolute protection from origin hijacks. If a customer C has its own customers, peers, or other providers not in the zone, it could still receive a hijack from those nearby ASes. We call this set of ASes the *local region* of the customer C, and we characterize this residual risk in §5. Importantly, the residual risk depends on the size and character of the local region of each AS, which they can know and control according to their own risk profile.

Figure 3 allows us to consider how the trust zone gives each AS control over the two types of hijack harm: having one's addresses hijacked or having one's traffic hijacked. An AS (e.g., B) can protect

against hijacking of its own addresses (which we call *owner harm*) in the zone by directly connecting to the zone, and registering its addresses in the RPKI. Other ASes that attach to the zone are thus protected from hijacking of their traffic to B's addresses (which we call *misdirection harm*). An AS that does not consider owner harm a significant risk need not register its addresses in a database (although we encourage universal use of the RPKI). The AS may care more about misdirection harm and might thus minimize its local region and get as many route announcements as possible from providers in the zone. Different ASes may have different risk assessments, and unlike today's routing ecosystem, this trust zone is structured to allow an AS to pick its own options based on its own assessment of route hijack risk.

4.1 Does a coherent zone exist?

Could such a coherent topological region exist? In fact, it already does, in the context of the MANRS initiative. Many of the MANRS members make up a connected region today. In May 2024, MANRS had 938 ISP members, with 1268 ASNs [67]. To derive the connected region, we perform a topology exploration using the CAIDA ASrank data [77] for May 2024. We start with members with no providers (Tier 1 providers), and recursively add directly-connected customers that are also MANRS members. The resulting region has 581 members with 766 ASNs. Currently 28,592 customers directly connect to this region. If MANRS could extend their operational practices to make this region a zone of trust, more than one-third of the ASes active on the Internet today would receive that protection.

4.2 Verified IP Zone (VIPzone)

We now describe how a set of proposed operational practices in a coherent zone of trust, which we call *VIPzone* (for Verified IP zone), will limit path hijacks. For an AS to be in the VIPzone, it must commit to these practices, and must be part of a connected zone. To be part of the connected zone, it must either be a Tier-1 provider, or have a member of the VIPzone as a transit provider.

Figure 3 illustrates the basic VIPzone operation. We describe the VIPzone practices below, and provide a finer-grained specification of these practices in the Appendix.

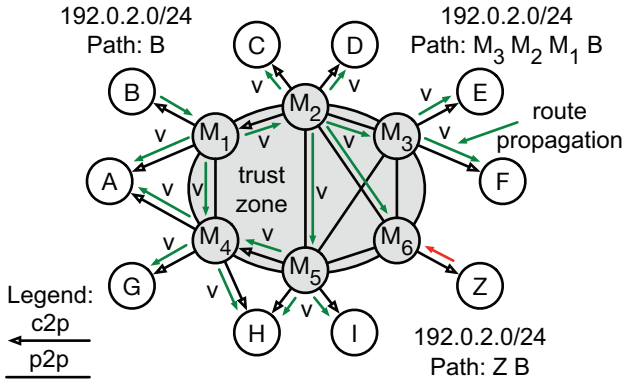

Figure 3: A Routing Zone of Trust can defend members and their customers from path hijacks in the zone if members (M) mark routes from their customers as VERIFIED (v) as they enter the zone, and other zone members select VERIFIED routes over unverified routes. Above, M_1 expects its direct customer B to announce 192.0.2.0/24, so M_1 marks that route as VERIFIED, and propagates it 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 an alternative VERIFIED route.

However, for announcements that VIPzone members have verified, they must propagate the VERIFIED marking as they forward announcements within the zone. A member must remove this marking if it appears in any announcement entering from outside the zone. This allows VIPzone members to establish the authenticity of VERIFIED announcement, regardless of their distance from the origin. Finally, inside the zone, any AS receiving multiple announced routes for the same prefix must prefer one marked VERIFIED. By this rule, no member will prefer a path hijack route over a legitimate route from customers directly attached to the zone, since legitimate routes will be marked VERIFIED.

Customers directly connected to the zone minimize owner harm, both for origin and path hijacks. Zone members verify prefixes received from attached (non-zone) customers and then forward them into the zone marked VERIFIED. If a malicious AS directly connected to the zone tries to launch an invalid origin hijack, zone members will discard it based on the KYC practices. If the AS launches a path hijack (which must by definition have more than one AS in the path), the member AS may forward it

unverified into the zone (a "not sure" situation), but it will have no impact so long as a corresponding VERIFIED announcement is active.

We emphasize the essential role of the VERIFIED tag. When a MANRS member cannot verify whether the path announcement is valid (e.g., multiple ASes in path) the member can forward this announcement onward. Forwarding potentially invalid announcements without *any* signal of risk prevents the current MANRS framework from manifesting a zone of trust. A key requirement of the VIPzone approach is that members propagate two sorts of announcements in the zone: VERIFIED and "not sure". The feature allows for more flexible and incremental deployment of the protections. In our VIPzone proposal, each AS drops invalid announcements, marks announcements as VERIFIED if it knows they are correct, and forwards announcements without the VERIFIED marking if the AS is "not sure." The rule that makes the zone trustworthy in this case is that if there is a VERIFIED announcement for a particular prefix, and one that is not VERIFIED (e.g., "not sure") for the *same prefix*, the zone members must prefer the VERIFIED announcement. This rule does constrain the routing policies of zone members to some degree, which depends on how their local region is configured.

4.2.1 Protection against route leaks

A route leak is an event in which an AS inappropriately (i.e., violating routing policy) forwards a route it legitimately received. The consequence is often that large flows of traffic reach this AS, which is not provisioned to carry them. A classic route leak occurs when a multi-homed AS that takes the routes it receives from one of its transit providers and inadvertently propagates these routes to its other transit provider.

In addition to preventing path hijacks of ASes directly attached to the zone, the VIPzone prevents leaks of announcements of prefixes by ASes *not* in the VIPzone. If the leak occurs within the zone, the announcement would be VERIFIED and thus propagated within the zone. This potential harm from accidental misconfiguration suggests an important insight: most ASes should not be in the VIPzone, but should get the protections by being a *customer* of a VIPzone member. We consider it preferable that only operators with sufficient technical abilities and resources join the VIPzone.

4.2.2 Protection against sub-prefix hijacks

One hijack that can penetrate the zone is based on a sub-prefix (an address block that is a subset of a VERIFIED prefix). Normal routing rules require that an AS, when selecting among routes for an arriving packet, must prefer the announcement with the longer prefix (i. e., smaller address block). Note that requiring that a VERIFIED announcement for a given prefix take precedence over an unVERIFIED announcement for a longer prefix risks breaking traffic management practices that disaggregate prefixes. Such a requirement could introduce loops. An AS concerned about owner harm resulting from a sub-prefix attack protects itself by registering ROAs for the prefix.

4.3 Evaluating VIPzone protections

We explore how many ASes would receive protection from hypothetical zones based on today's Internet topology. Using CAIDA's AS Rank data from May 2024, we initialize a zone with the 100 ASes with the largest customer cones. We then add new members, again ordered by the size of their customer cone. Figure 4 shows the number of protected ASes expanding rapidly with zone size, up to 11,781 ASes (in this data set), at which point every AS with any customers is in the zone. The only ASes not in the zone are single AS stubs.

However, note that such a large zone is unrealistic. Most ASes in that zone are small providers with few customers, likely without sufficient operational sophistication or resources to join the zone. If we pick an arbitrary cutoff of 600 members (about the size of the current MANRS zone), that would protect a little over two thirds of the ASes in the Internet (in this hypothetical analysis, 53,563). This number is higher than the 28,592 customers of the current MANRS region we discussed above, because this VIPzone is formed by including all of the largest ASes (even non-MANRS members) as measured by their customer cone.

4.4 Social engineering attacks

As protection against traditional hijacks improves, attackers devise new ways to disrupt routing. One

is a social engineering attack in which an attacker contacts a provider of a target AS, and (pretending to be an agent of the target AS) requests that the provider provision a new link to serve that target AS. If the provider does not recognize that the request is not legitimate, the attacker now has a BGP connection to the provider that the provider thinks is associated with the target AS. At this point, the attacker can announce routes (e.g., hijack them) associated with the target AS, and the provider will accept these announcements.

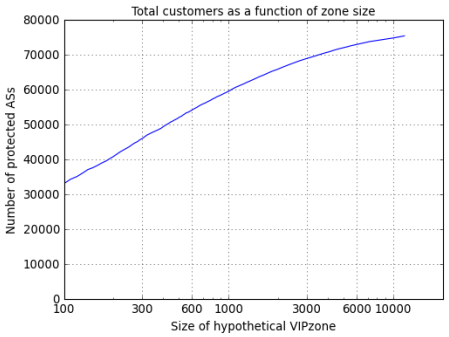

Figure 4: Protected ASes (in the zone or connected directly to it) as a function of zone size (ASRank data, May 2024)

Transit providers will have to harden their implementation of the MANRS Know Your Customer requirement to detect these sorts of attacks. If the attacker can bypass the KYC test, ROAs and ROV are of no use in prevention, since in that context, it appears a correct AS number is being used. This requirement applies equally to the existing MANRS, our VIPzone, and ASPA.

5 Evaluating Residual Risk (Local Regions)

We review the residual risks that an AS faces even if it is directly connected to the VIPzone, and what that AS can do to further reduce these risks. We have already described how an AS mitigates the risk of *owner harm* simply by connecting to the zone. The residual risk of *misdirection harm* if they connect to the zone is a function of the size of the *local region* and the probability that a malicious AS operates in that region.

A local region of a VIPzone customer arises due to its interconnection arrangements outside the zone, from which it receives BGP announcements. These include the ASes in the customer cone of that AS, the peers of that AS and their customer cones, and any providers (and their neighbors, recursively) of that AS that are not in the zone. In Figure 5, A has provider X in the zone. Its local region includes customers B and G, peer E and E's customer F, provider H (which is not in the zone), and its customer J, and peer S of provider H and its customer T. If provider H itself had a provider that was not in the zone, that provider, its customers, and any peers and customers of those peers would also be in the local region of A. Any of these could launch a hijack that triggers a *misdirection harm* to A.

We make three observations about local regions. First, the risk of hijack by one's own customer (A's customer B in this example) is a function of the risk of malicious behavior in the local region. But A (or any AS outside the zone) can mitigate this risk by implementing a robust KYC practice, which can generally detect forged-origin attacks by customers.

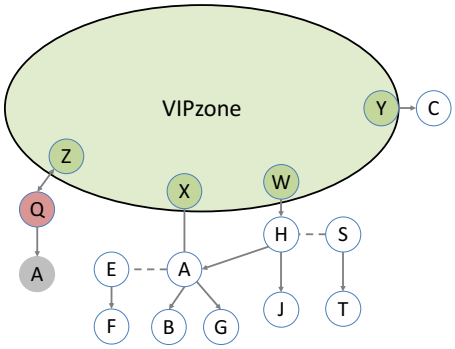

Figure 5: Various customers of a VIPzone, including A with a local region, C with no local region, and a malicious AS Q pretending that A is a customer.

Second, misdirection from a hijack in the region is restricted to the region. In Figure 5, if malicious AS Q launches a path hijack asserting that it has A as a customer, that announcement may penetrate the zone but without a VERIFIED mark, so zone members will prefer the VERIFIED announcement from X.

Third, for many attached customers the local region is small. To examine the size distribution of local regions, we return to our hypothetical VIPzone (i.e., seeded with 100 ASes with the largest customer cone) and compute the size

of the local regions for all attached customers. We add to the zone 100 ASes at a time, and at each step compute the size of the local region for the attached customers.

Figure 6 plots the resulting distribution. To compensate for the limited observability of peering relationships, we use two methods to compute the size of the local region. Figure 6a plots the local region size using the customer-provider and peering relationships from CAIDA's ASRank data [79]. Figure 6b relies on the method in [48], which augments observable peering relationships by assuming that any two ASes that attach to the same IX have a peering relationship. We use data from from PeeringDB and PCH [80] to augment the set of peering relationships inferred by AS Rank.

Figure 6a underestimates the sizes of local regions, since CAIDA's ASRank data is derived from BGP announcements collected by RouteViews and RIPE RIS [81, 82], and those vantage points do not have sufficient density to capture all peering relationshops. Figure 6b probably overestimates the sizes of the local regions, since many ASes that connect to IXs have selective peering policies. So the actual distribution probably lies between these two set of curves.

Note that the distribution of local region sizes is bimodal. Depending on the zone size, between 30% and 60% of the customer ASes have a very small local region–close to 1 AS. These are stub ASes that obtain access to the Internet using a transit provider, and do not peer to obtain connectivity. The right side of the plots shows large local regions, which represent ASes that peer widely to reduce their dependency on transit providers, or else use multiple providers, one of which is not in the zone, and which itself uses massive peering. A realistic consequence of extensive peering with ASes that do not take known steps to verify their announcements is an increased risk of hijack. That expanded attack surface in the ecosystem is a motivation for the approach we propose.

5.1 Why we cannot assess realistic risk using current topology data.

This analysis provides a hypothetical indication of the level of protection and residual risk that a VIPzone would yield under current interconnection patterns. But it is a problematic approach to assessing residual risk, since the architecture of the VIPzone will affect peering incentives, by design. That is, the goal of VIPzone proposal is to devise a set of practices that allow an AS concerned about security risk (in particular the risk of hijack) to take action that minimizes this risk. In other words, they will shift interconnection patterns to exploit the benefit of the zone.

For many small to mid-size networks, connecting to an Internet exchange (IX) is an efficient way to establish many peering connections. For our hypothetical zone with 900 members, of the 57,288 customers of the zone, we identify 15,337 that are attached to at least one IX. Some IXs may choose to take steps to reduce the risk of hijacks among their members, such as requiring that their members document the ASs that they will legitimately announce. Some peers may take steps to verify their own customers, and the practical risk of using routes from such a peer would be minimal.[3] It may not be practical for every AS that connects to an exchange to assess the pragmatic level of risk at that exchange, but if one actor can make that assessment on behalf of the exchange, all the members can take that information into account in deciding what action to take: e.g., whether to peer with all of the other exchange members, or peer selectively with those peers that offer significant volume levels.

A further uncertainty in these plots derives from the common use of prefix filters on peering links, precisely to protect themselves from harm due to erroneous or malicious BGP announcements. That practice would reduce the effective size of the local region from which hijacks can come. Many ASes consider such filtering good routing hygiene today. We know no way to measure the extent to which operators have deployed such filters.

But we emphasize that the power of a trust zone approach is that each AS gets to make its own risk assessment, and act accordingly. ASes with small or no local region would not have to take these steps. Larger ASes are more likely to have the operational capacity to protect their local region, e.g., implement prefix filters. If a VIPzone existed, we would expect ASes to take actions to reduce the residual risks from their local regions.

5.1.1 Protection for ASes not attached to the zone

In Figure 5, AS B shares the local region of A, but is not directly connected to the zone. What protection does B receive from hijacks? With respect to owner risk, B can prevent simple hijacks based on an invalid origin by registering ROAs, but it gets no protection from path hijacks. With respect to misdirection risk, it is in the same situation as A: no hijacks will come into B's region from the zone, but a hijack in

[3]Internet2 exemplifies such a region; they track the full customer cone of their members, and use prefix filters to prevent incorrect announcements. Using routes from a region of this sort is practically risk-free.

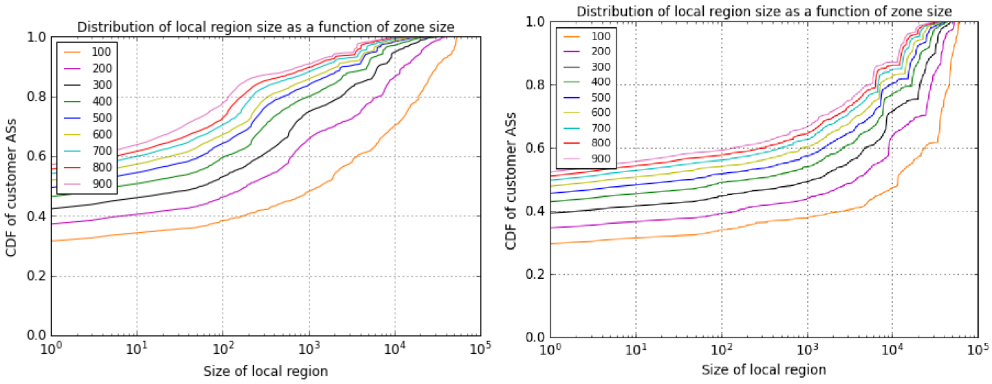

(a) Plot based on data from ASrank

(b) Plot based on data from ASrank plus an assumed peering relationship between any two ASes at an IX

Figure 6: Sizes of local regions for customers of a hypothetical VIPzone, for various zone sizes. Between 30% and >50% of the customer ASes have a region size near 1.

B's local region can still cause misdirection harm. Many smaller ASes offer low-value, limited-interest services, and their owner risk of a hijack is minimal. If the AS does consider the owner risk to be substantial, they can and should obtain transit from a member of the zone.

6 Auditing Requirements

Our proposal for a VIPzone does not use real-time detection of suspicious announcements. Real time prevention requires adding code to the BGP processing path in routers or route computation servers. This approach would potentially lead to a more brittle scheme. Instead the VIPzone uses a trust-but-verify approach: checking conformance of members with its requirements, detection and documenting of failures, and suspension or ejection of non-compliant members. This requirement means that members must have the will, and the institution, to undertake conformance auditing. Independent third parties can check conformance off-path, by looking at public BGP announcements. In support of this auditing, every VIPzone member would be required to provide a BGP view to a route collector. The audit process does not use the member's view to audit their behavior (the member could lie) but rather uses the views provided by the member's neighbors that are also members and thus provide views of their own. Using the neighbor views allows confirmation that the member correctly propagated verified routes with the VERIFIED tag, and did not use the VERIFIED tag on routes that other members had not tagged as VERIFIED.

This approach is similar in spirit to how the CA/Browser forum verifies the correct behavior of certificate authorities. 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 so that providers of browsers can choose to remove them from their list of trusted root CAs. 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. We argue that an industry-led body, analogous to the CA/Browser forum but with a stronger centralized authority, should decide on necessary actions if a VIPzone member does not conform to the required practices. But note that the penalty in this case is not being disconnected from the Internet, but just losing the right to initiate VERIFIED announcements.

Independent of the exact specification of the practices that define a zone, it must be possible to tell by inspection if an announcement is not conformant. The three tests for VIPzone member 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 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.

Another advantage of off-line conformance checking is that it could allow an AS to register its intent to violate Rule 1 in a specific case and announce a route that is non-conformant (e.g., to deal with a specific customer requirement), and accept responsibility for ensuring that it is benign. Allowing benign exceptions, including marking them VERIFIED, enables more nuanced balance of security and availability priorities.

6.1 Cost of checking for conformance

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

7 Comparison to Other Proposed Solutions

We compare our proposal to two leading alternative proposals to advance the collective state of routing security, in particular to prevent path hijacks: BGPsec and ASPA. But we preface this comparison with a comment on the tension between our VIPzone 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 [83]. The collective action aspect of routing security, where it is not feasible to verify implementation by other parties, is at odds with this assumption. The VIPzone approach better aligns incentives, allocates responsibility to specific points in the zone (the perimeter), and ascertains whether zone members are implementing the required operational practices.

7.1 AS Provider Authorization (ASPA)

ASPA [70] is a mechanism that lets a customer AS register a list of providers that the customer uses. This registration (an Autonomous System Provider Authorization or ASPA) is recorded in the same system that is used to store ROAs–the RPKI administered by the five RIRs. The ASPA data is globally visible, so any AS receiving a BGP announcement can look at the sequence of ASes in the path, and check to see if there is an ASPA that covers any adjacent pair of ASes in the path. If there is, and the announcement is inconsistent with the ASPA, the AS receiving the announcement can drop it [70]. ASPA can be used to limit both route leaks and, to some degree, against path hijacks, assuming the appropriate ASes deploy ASPA in the correct places. The ASPA specification describes several deployment scenarios.

ASPA's design differs in several ways from our proposal.

- The VIPzone design tries to minimize the effort required of small ASes to get protection. It requires only that the small AS connect to a transit provider that is in the zone and (ideally) register its ROAs. ASPA requires that the small AS register an ASPA describing its providers. While the mechanics of registration need not be complex, this registration becomes one more data record that the operators of the AS must keep track of, and remember to change if they change providers.
- The VIPzone design does not require new mechanisms in the routers (or route computation servers). The actions required of a VIPzone member (see our discussion in Verified IP Zone) include new operational practices and use of a new community value. ASPA checking requires a new processing check, which includes downloading the relevant ASPA data and inspecting the announcement for validity. This dependency also implies the need for the RPKI to store and manage new (ASPA) records.
- The VIPzone design assigns clear responsibilities: an AS at the edge of the zone has specific requirements to check announcements received from its customers, including a KYC check. This perimeter allows clear description of protection and residual harm. The current ASPA draft [70]

describes use cases without assigning responsibilities to specific ASes. Thus it is not clear which ASes should do ASPA checking, which ASes would have the motivation to register ASPA records, and (thus) what protection ASPA will achieve. For example, if an AS has listed a provider in an ASPA record, and that provider has such poor business/operational practices that it cannot identify an imposter posing as their legitimate customer, an ASPA alone cannot prevent the resulting harm. Assignment of responsibility, as in VIPzone, allows the possibility of conformance checking.

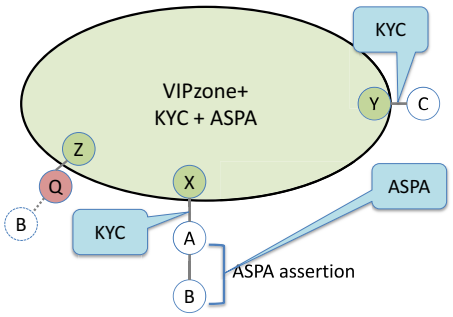

Figure 7: **ASPA-extended VIPzone.** Zone member X can use ASPA to verify an announcement with two ASes in the path outside the zone. If B registers an ASPA recording A as a provider, and X has done KYC on A, then X can mark the route as VERIFIED.

ISPs could use ASPA to further the range of VIPzone protection to customers of zone customers This extension would allow an AS at the edge of a zone to mark as VERIFIED announcements with two or fewer ASes in the path, as opposed only one. In Figure 7, Y uses VIPzone practices to verify announcements originated by C, as does X to verify A. But X has the option of using an ASPA registered by B to confirm that A is a valid provider of B. X can tag as VERIFIED the announcement that includes both A and B in the path only if there is an ASPA registered by B. Otherwise, X must not mark the announcement, but can forward it into the zone unmarked.

The other case in Figure 7 is that Q is malicious, and wants to hijack a prefix belonging to B. If B has registered a ROA for the prefix, then Q cannot validly announce B's prefix. It would have to pretend to be B. If B has registered an ASPA saying that its provider is A, then this ASPA would allow Z to conclude that the announcement is invalid. The attacker Q could add AS A to the path to make a valid path, but then the announcement would have three ASes in the path outside the zone (A, B, Q), and (in the VIPzone we propose) Z must not mark a path VERIFIED if it has more than two ASes in the announcement.

7.2 BGPSEC

Like ASPA, BGPsec is attempting to achieve a zero-trust approach. To the extent that every router that forwards the announcement adds its own cryptographic signature, any router along the path can verify that the series of signatures to that point are valid. This function also means that BGPsec, if pervasively and correctly deployed, provides the technical means to address the KYC requirement that VIPZone and ASPA cannot do in-protocol. That is, BGPsec prevents the social engineering impersonation attack, since the imposter will not have the necessary keys to sign their announcements. However, the requirement for comprehensive deployment dramatically reduces the incentive for ISPs to undertake the cost and complexity of BGPsec deployment, and a lengthy trajectory of partial deployment implies inconsistent and unpredictable implementation of the required checking. We expect that governments will not have the patience to wait for deployment of a global solution to route hijacks.

8 Conclusion

There is currently no consensus as to the next step to secure BGP beyond the simplest type of hijacks. As of 2024, 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 misaligned. We have proposed a path forward that creates incentives for ASes (both customer and provider) to participate, protects ASes against *path hijacks* and *origin hijacks* with no effort or investment needed by small ASes, and avoids the need for new mechanism in routers.

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 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. A coherent core of ISPs has already emerged organically in the ecosystem, which can be leveraged to create a *zone of trust*, a region that protects not only all networks in the region, but all directly attached customers.

A few concerns with VIPzone bear further consideration. First, will it concentrate power in a few trusted networks, those with the authority to verify routes? We believe the VIPzone requirement for transparency, accountability, and independent auditing, provides a counterpoint to potential abuses of power.

Second, will trust zones fragment the Internet? Some Internet fragmentation has already occurred, and trust zones provide a way to bridge some of these fragments using a trust-but-verify framework, like treaties in other global domains. We acknowledge that multiple trust zones may emerge, including on national boundaries. But note that for a VIPzone to be effective, both (1) ASes that produce important services, and (2) ASes that consume those important services, must be attached to that zone. As an extreme example, if each country wants its own trust zone, networks with global customer bases would have to replicate their point of attachments in all trust zones where they serve customers. We imagine trust zones to evolve instead more like global trading zones.

Third, achieving the VIPzone protection requires auditing and enforcing conformance with the practices. The institutional framework required for the necessary data collection already exists in multiple places, e.g., RIPE and RouteViews. But it is still more expensive (and therefore less incentive-compatible) than doing nothing in the current unregulated environment.

Fourth, ISPs have to trade off some autonomy in exchange for routing security. ISPs are required to prefer VERIFIED routes over customer routes, and ISPs would hand some control over to a non-ISP third party (the auditor) similar to the CA/Browser Forum today. But unlike other proposed approaches to routing security, transit ISPs can claim to offer their customers a securely-routed service by participating.

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 debate and analysis in the context of possible regulatory measures. We recognize that ISPs, like most private sector actors, prefer lack of regulation and work to avoid it as long as possible. But the EU has made it clear they will regulate to safeguard their citizens despite private sector objections [84, 85, 86]. We offer this path forward as an approach where the private sector could drive a self-regulatory framework that achieves the accountability regulators are now seeking in digital domains.

9 Acknowledgments

This work is supported in part by funds from the National Science Foundation (NSF: #OAC-2131987).

References

[1] Maria Apostolaki, Aviv Zohar, and Laurent Vanbever. Hijacking Bitcoin: Routing Attacks on Cryptocurrencies. In *2017 IEEE Symposium on Security and Privacy (SP)*, May 2017.

[2] Dan 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/`.

[3] 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/`.

[4] 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.

[5] Pierluigi 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.

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

[7] Pierre-Antoine Vervier, Olivier Thonnard, and Marc Dacier. Mind Your Blocks: On the Stealthiness of Malicious BGP Hijacks. In *Proceedings 2015 Network and Distributed System Security Symposium*, San Diego, CA, 2015.

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

[9] Job Snijders. The bgpsec plan, 2022. `https://www.bgpsec.net/history.html`.

[10] 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`.

[11] 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`.

[12] 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`.

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

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

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

[16] 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.

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

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

[19] Merit. IRR - Internet Routing Registry. `http://www.irr.net`.

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

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

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

[23] Ben Du, Gautam Akiwate, Thomas Krenc, Cecilia Testart, Alexander Marder, Bradley Huffaker, Alex C Snoeren, and KC Claffy. IRR Hygiene in the RPKI Era. In *PAM*, pages 321–337, 2022.

[24] Leo Oliver, Gautam Akiwate, Matthew Luckie, Ben Du, and kc claffy. Stop, DROP, and ROA: Effectiveness of Defenses through the Lens of DROP. In *ACM Internet Measurement Conference*, 2022.

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

[26] Joe Touch, Allison Mankin, and Ronald P. Bonica. RFC 5925: The TCP Authentication Option, Jun 2010.

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

[28] Stephen Kent, Charles Lynn, and Karen Seo. Secure Border Gateway Protocol (S-BGP). *IEEE Journal on Selected areas in Communications*, 18, 2000.

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

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

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

[32] Geoffrey Goodell, William Aiello, Timothy Griffin, John Ioannidis, Patrick D. McDaniel, and Aviel 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.

[33] Yih-Chun Hu, Adrian Perrig, and Marvin Sirbu. SPV: Secure path vector routing for securing BGP. *ACM SIGCOMM Computer Communication Review*, 34(4), 2004.

[34] Lakshminarayanan Subramanian, Volker Roth, Ion Stoica, Scott Shenker, and Randy H Katz. Listen and Whisper: Security Mechanisms for BGP. In *Symposium Networked System Design and Implementation,*, 2004.

[35] Tao Wan, Evangelos Kranakis, and Paul C. van Oorschot. Pretty Secure BGP, psBGP. In *Proceedings of the 2005 ISOC Symposium on Network and Distributed Systems Security*, San Diego, 2005.

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

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

[38] Patrick Reynolds, Oliver Kennedy, Emin Gün Sirer, and Fred B. Schneider. Using External Security Monitors to Secure BGP. *IEEE/ACM Transactions on Networking*, 2006.

[39] Jian Qiu and Lixin Gao. Hi-BGP: A Lightweight Hijack-proof Inter-domain Routing Protocol, 2006.

[40] 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.

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

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

[43] Kevin Butler, Toni R. Farley, Patrick McDaniel, and Jennifer Rexford. A Survey of BGP Security Issues and Solutions. *Proceedings of the IEEE*, Jan 2010.

[44] Muhammad S. Siddiqui, Diego Montero, Rene Serral-Gracia, Xavi Masip-Bruin, and Marcelo Yannuzzi. A survey on the recent efforts of the Internet Standardization Body for securing inter-domain routing. *Computer Networks*, 80:1–26, Apr 2015.

[45] Asya Mitseva, Andriy Panchenko, and Thomas Engel. The state of affairs in BGP security: A survey of attacks and defenses. *Computer Communications*, 124:45–60, Jun 2018.

[46] Cecilia 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.

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

[48] Robert Lychev, Sharon Goldberg, and Michael Schapira. BGP security in partial deployment: is the juice worth the squeeze? In *ACM SIGCOMM*, pages 171–182, Aug 2013.

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

[50] Avichai Cohen, Yossi Gilad, Amir Herzberg, and Michael Schapira. Jumpstarting BGP Security with Path-End Validation. In *ACM SIGCOMM*, 2016.

[51] Xin Zhang, Hsu-Chun Hsiao, Geoffrey Hasker, Haowen Chan, Adrian Perrig, and David G. Andersen. Scion: Scalability, control, and isolation on next-generation networks. In *2011 IEEE Symposium on Security and Privacy*, pages 212–227, 2011.

[52] Henry Birge-Lee, Joel Wanner, Grace H. Cimaszewski, Jonghoon Kwon, Liang Wang, François Wirz, Prateek Mittal, Adrian Perrig, and Yixin Sun. Creating a Secure Underlay for the Internet. In *USENIX Security Symposium*, Aug 2022.

[53] John Scudder, Randy Bush, Pradosh Mohapatra, David Ward, and Rob Austein. RFC 6811: BGP Prefix Origin Validation, Jan 2013.

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

[55] 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.

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

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

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

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

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

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

[62] Weitong Li, Zhexiao Lin, Md. Ishtiaq Ashiq, Emile Aben, Romain Fontugne, Amreesh Phokeer, and Taejoong Chung. RoVista: Measuring and analyzing the route origin validation (ROV) in RPKI. In *ACM IMC*, Oct 2023.

[63] Geoff Houston. ISP Column - March 2021, Mar 2021.

[64] APNIC labs. RPKI I-ROV Filtering World Map.

[65] Cecilia Testart, Philipp Richter, Alistair King, Alberto Dainotti, and David Clark. To Filter or Not to Filter: Measuring the Benefits of Registering in the RPKI Today. In Anna Sperotto, Alberto Dainotti, and Burkhard Stiller, editors, *Passive and Active Measurement*, volume 12048. Springer, 2020.

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

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

[68] Ben Du, Cecilia Testart, Romain Fontugne, Gautam Akiwate, Alex C. Snoeren, and kc claffy. Mind Your MANRS: Measuring the MANRS Ecosystem. In *ACM Internet Measurement Conference*, 2022.

[69] Internet Society. MANRS Network Operator Participants, July 2023. https://www.manrs.org/netops/participants/.

[70] Alexander Azimov, Eugene Bogomazov, Randy Bush, Keyur Patel, Job Snijders, and Kotikalapudi 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-16, Internet Engineering Task Force, Oct 2023. https://datatracker.ietf.org/doc/draft-ietf-sidrops-aspa-verification/16/.

[71] Loqman Salamatian, Todd Arnold, Ítalo Cunha, Jiangchen Zhu, Yunfan Zhang, Ethan Katz-Bassett, and Matt Calder. Who squats ipv4 addresses? *SIGCOMM Comput. Commun. Rev.*, 53(1):48–72, apr 2023.

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

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

[74] Broadband Internet Technical Advisory Group. Security of the Internet's Routing Infrastructure, 2022. https://www.bitag.org/documents/BITAG_Routing_Security.pdf.

[75] Christopher Yoo and David Wishnick. Lowering Legal Barriers to RPKI Adoption. *Faculty Scholarship at Penn Law*, Jan 2019.

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

[77] CAIDA. AS relationships data for May 2023. https://publicdata.caida.org/datasets/as-relationships/serial-1/.

[78] Kotikalapudi Sriram and Alexander Azimov. Methods for Detection and Mitigation of BGP Route Leaks. Internet-Draft draft-ietf-grow-route-leak-detection-mitigation-09, Internet Engineering Task Force, Oct 2022. Work in Progress.

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

[80] CAIDA. CAIDA Internet eXchange Points (IXPs) Dataset, 2023. https://www.caida.org/catalog/datasets/ixps/.

[81] Routeviews Project – University of Oregon. http://www.routeviews.org/.

[82] RIPE Routing Information Service. http://www.ripe.net/ris/.

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

[84] Council of European Union. Council regulation (EU) no 2016/679, general data protection regulation, 2014.

[85] Council of European Union. Council regulation (EU) no 2022/2065 of the European Parliament and of the Council of 19 October 2022 on a Single Market For Digital Services and amending Directive 2000/31/EC (Digital Services Act), 2022. https://eur-lex.europa.eu/eli/reg/2022/2065/oj.

[86] Council of European Union. Proposal for a REGULATION OF THE EUROPEAN PARLIAMENT AND OF THE COUNCIL on horizontal cybersecurity requirements for products with digital elements and amending Regulation (EU) 2019/1020, 2022. https://eur-lex.europa.eu/legal-content/EN/TXT/?uri=celex:52022PC0454.

APPENDIX

A Full specification of required actions for members of the VIPzone

We summarize the required operational practices of VIPZone members in the body of the paper; here we repeat the summary, provide additional details, and diagram specific scenarios.

First, VIPzone members that can participate in these enhanced practices must be part of a connected region.

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 [13], which we call *VERIFIED*. (Some other BGP mechanism with equivalent properties could also be used.)

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 (KYC) 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. Customers of zone members do not need to understand or act on the VERIFIED marking; the zone rules allow them the option to distinguish 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 prefixes that the member has previously established as allowed from their customer). If the prefix is RPKI-valid, is registered by the owner in an authenticated IRR, or from a list of allowed prefixes, the zone member AS 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 zone 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. To preserve Internet connectivity, zome members must forward unverified routes according to normal routing policies.
7. **Export routes to a route collector for auditing.** To allow for auditing behavior of trust zone members, members must export their routes to a route collector.