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

David Clark (MIT), Cecilia Testart (GaTech), Matthew Luckie (CAIDA/UCSD), KC Claffy (CAIDA/UCSD)

14 July, 2023

1 Introduction

In this note, we introduce two related concepts that we believe can lead to improved security of the global Internet routing system (the Border Gateway Protocol or BGP). 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.

A Zone of Trust The first concept we introduce is a *zone of trust*, a connected region of the Internet where providers have taken enhanced steps to improve the security of that region. We focus on embodiments of a zone that also protect the *customers* connected to the providers in the zone. We focus on this sort of zone because in today's Internet the steps necessary to configure and operate systems so that they are secure are sometimes complicated, and smaller ASs may not have the skills to undertake them, nor the resources to give this objective priority. If steps taken by bigger, more skilled and resourced ASs can improve the security of smaller ISPs, the result is a wider scope for the protection.

A "VIPzone" The second concept we introduce is a specific example of a zone of trust, which is intended to protect ASs from two kinds of route hijacks, *invalid origin* hijacks and *invalid path* hijacks. 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 registration of Route Origin Authorizations (ROAs) and filtering BGP announcements based on these ROAs (Route Origin Validation, or ROV) can reduce the impact of the simpler origin hijacks, but is not effective against path hijacks. We describe a set of operational practices for the VIPzone that will also reduce the impact of path hijacks.

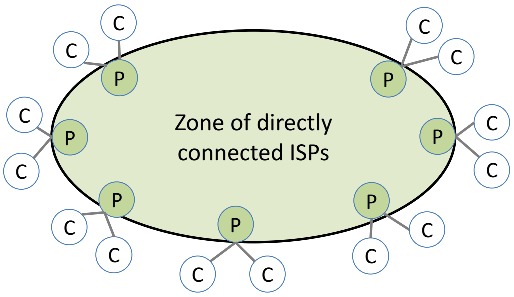

Figure 1: A hypothetical zone of trust, with providers P in the zone providing transit to customers C that are attached to those providers.

Our approach protects customers of the zone (presumably smaller ASs) as well as the members of the zone. It is intended to increase the incentive of ASs to be a part of the zone. It does not depend on the development of a new protocol, but instead depends on existing mechanisms in the routing system, and new operational practices that exploit those mechanisms, which in principle allows quicker deployment than a scheme that requires substantial new software in routers.

2 Zone of Trust

Figure 1 illustrates a simple zone of trust with a number of providers (marked in green) at the edge of the zone providing transit service to a number of customers (white) directly attached to them. The providers (green) are connected by routes that are contained 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 simple zone could protect against origin hijacks as follows. If all providers P in the zone commit to implement ROV and drop invalid announcements, then no invalid announcements will circulate inside the zone, which in turn means that the customers C will never receive a BGP announcement from the zone where the origin is invalid based on a ROA.

This example illustrates four properties of a zone of trust:

- *Collective action* by ASs creates the zone and its attributes.

- Topology matters–the zone depends on the transitive connectivity of the members within the zone

- Customers of the zone obtain protection simply by using a provider that is in the zone. They need take no other action.

- 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 ASs. We call this set of ASs the *local region* of the customer C, and we characterize this residual risk in Section 4.

As ROV is discussed today, the action of each AS is considered in isolation, and security is a statistical measure. We can count the number of ASs that register their ROAs, or the number of ASs that implement ROV, but the consequence for a given AS is a function of what other ASs choose to do. It is thus not clear what specific action an AS should take to reduce its risk profile. Today, invalid announcements propagate to some extent across the Internet, and may or may not reach the AS in question. With this zone, the benefit to a given customer can be stated more clearly: they will receive no announcements from the zone with an invalid origin based on a registered ROA, and the residual risk depends on the size and character of the local region of that AS, which they can know and control according to their own incentives.

A consequence of the coherent zone of trust is that it creates an *incentive* for customers that are concerned about hijacks to 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 ASs 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.

2.1 Two sides of a hijack

Hijacks are a two-sided harm, depending on the relationship of the hijack to the owner of the hijacked address block. From the perspective of the owner of an address block, if that block is hijacked, the *owner* is harmed. If an AS in some other part of the Internet is misled by that hijack, and traffic from that AS ends up at the wrong location, that AS is also harmed. We call the first perspective the *owner harm*, the second the *misdirection harm*.

Returning to Figure 1, an AS concerned about owner harm, i.e., hijacking of its own addresses, protects itself from this harm *in the zone* by directly connecting to the zone and (as we discuss in Section 4) registering its ROAs. Because the ASs that attach to the zone are protected from owner harm in the zone, other ASs that attach to the zone are thus protected from misdirection harm for the attached ASs. An AS that does not consider owner harm a significant risk need not register its ROAs (although we would encourage universal registrations of ROAs); 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 ASs may have different risk assessments, and a scheme that allows an AS to pick its own options based on its own assessment of risk is more realistic than a scheme that imposes a "one size fits all" solution.

To avoid confusion, we caution against thinking of a zone of trust as a *walled garden*, where users cannot get in or out, or a *gated community* where only privileged users can get in. A better (but not perfect) analogy is a *safe neighborhood*, where neighbors commit to practices that keep an area safer, but anyone can leave the area and go where they want.

Different operational practices, undertaken as a collective action by ASs in a connected region, can produce different sorts of zones that yield different profiles of protection. Section 3 describes one approach, which protects from both invalid origin and invalid path hijacks, and is designed to encourage incremental deployment.

2.2 Is a connected zone practical?

An example of a connected region that exists today is the subset of members of the *Mutually Agreed Norms for Routing Security (MANRS)* initiative that happen to be reachable through each other. The goal of MANRS is to "help reduce the most common routing threats on the Internet'.' [3] MANRS specifies four practices for participating networks, two of which roughly correspond to the RPKI/ROV steps of registering authoritative information about one's prefixes, and verifying BGP announcements against authoritative information. However, to encourage broad uptake, they do not specify the method of verification. 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 directly connected 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. We call the first part of this requirement the Know Your Customer or KYC requirement. A MANRS member must know that the customer at the other end of the connection is allowed to announce the AS that it is announcing. There are no protocols in use today that provide this assurance (BGPsec can provide this assurance if and as it gets deployed and appropriately used). ASs must implement suitable business practices to ensure that this requirement is met.

MANRS does not specify how a member AS should implement the second aspect: the verification of the prefix assertions of its customers. In particular, MANRS does not require the use of RPKI/ROV (ROAs) in this verification. The AS can use ROAs, or can verify against information in the Internet Routing Registry (IRR), or rely on a private arrangement with its customer. (The MANRS requirements do currently specify that network operators must *encourage* their customer network operators to register ROAs.)

Many of the MANRS members make up a conected region today. As of May 2023, there are 830 ISP members of MANRS, with 1011 ASNs. To construct the members that make up the the connected region, we perform a topology exploration using the CAIDA ASrank data from May 2023. 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 499 members with 613 ASNs. Perhaps more interestingly, currently 25,916 customers directly connect to this region. In other words, if MANRS could extend their operational practices to make this region a zone with well-defined security consequences, about one third of the ASs active on the Internet today would receive that protection.[1]

2.3 Enabling flexibility in a zone of trust

The most rigorous approach to a zone of trust is to ensure that *every provider* at the edge fully implements the protective practices, e.g., in the example in Figure 1 that there are no announcements in the zone that are invalid based on a ROA, so that there are no "bad" announcements in the zone. Although cleanest, this approach is not realistic for some operational practices and protection objectives, since in some cases a provider can neither confirm nor reject an announcement. The MANRS initiative has this character. For announcements originated by its directly connected customers, the member must verify the AS and prefix. But MANRS does not assume that a member knows the validity of the full customer cone of that customer. For announcements with more than one AS in the path, a MANRS member may not be able to confirm whether the path announcement is valid. It must forward this announcement onward in this case. This forwarding of potentially invalid announcements is what prevents the current MANRS framework from manifesting a zone of trust.

To accommodate flexibility without sacrificing trust requires a signaling mechanism for an AS to distinquish verified from unverified announcements as it forwards them. In the VIPzone proposall, each AS drops invalid announcements, marks them as VERIFIED if it knows they are correct, and forwards them without marking if the AS is "not sure." There are thus *two* classes of announcement accepted into the zone: VERIFIED and "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 allows for more flexible and incremental deployment of the protections. As members confirm their KYC check for a given customer, they can mark announcements as VERIFIED. But so long as they take the one action of preferring VERIFIED announcements, they can start marking (and thus protecting) the announcements of their customers incrementally. The rule does impose a specific limitation on the routing policies of the zone members, which we characterize in Section 4.

3 The VIPzone

Four design requirements shape our VIPzone proposal:

- Protect as many ASes as possible against *path hijacks* in addition to *origin hijacks*.
- Avoid the need for new protocols and new mechanism in routers (or route computation servers). Exploit existing BGP mechanisms but incorporate them into new operational practices
- Minimize the effort that small ASs must invest to obtain protection.
- Create an incentive for ASs (both customer and provider) to support the scheme.

Below we discuss the impact of these requirements on routing policy for the VIPzone.

Our **VIPzone** builds directly on MANRS requirements. For announcements originated by its directly connected customers, the MANRS member must check them for validity. In our VIPzone scheme, the member then either drops them or marks them VERIFIED. For announcements that come from customers of the customer, the member forwards them without marking them VERIFIED. We propose the use of a community value to carry the VERIFIED marking, but other BGP mechanisms may be more suitable. The only requirement is that routers propagate this marking as they forward announcements within the zone, and remove this marking if it appears in any announcement entering from outside the zone.

Protections provided by the VIPzone Directly connected customers minimize owner harm, both for origin and path hijacks. Prefixes of attached customers are forwarded 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 MANRS KYC practices. If it launches a path hijack (which must by definition have more than one AS in the path), the member AS may forward it into the zone (a "not sure" situation), but since it is not marked VERIFIED, it will have no impact so long as a corresponding VERIFIED announcement is active.

The VIPzone reduces owner risk, which in turn reduces misdirection risk. No directly connected AS will receive a hijacking BGP announement (neither origin nor path hijack) from inside the zone.

[1] MANRS membership from MANRS, https://www.manrs.org; AS relationship from CAIDA ASRank, https://asrank.caida.org/.

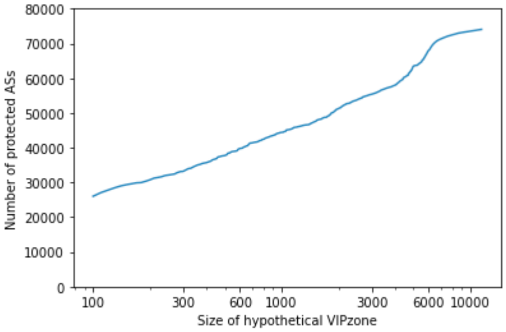

Figure 2: Total number of protected ASs (in the zone or connected directly to it) as a function of zone size, using data from May 2023.

Number of protected ASs We can construct a hypothetical VIPzone, using the current publicly observed Internet AS topology, to explore how the number of protected ASs varies as we increase the zone size. Using CAIDA ASrank data from May 2023, we construct this hypothetical zone by starting with the 100 ASs with the largest local regions, and assume they are in the zone. We then add new members, based on the size of their local region. Figure 2 shows the number of protected ASs (in the zone or connected to it) as we vary the zone size. For completeness, we plot zone size to the maximum value (in this hypothetical case 11,458 ASs), at which point every AS with any customers is in the zone. The only ASs not in the zone are single AS stubs. However, a zone of this size is unrealistic–the large majority of those ASs we added to the zone are small providers with only a few customers, and we would not expect they would have the operational sophistication or the resources to join the zone. If we pick an arbitrary cutoff of 600 members (about the size of the current MANRS zone), a little over half of the ASs in the Internet (in this hypothetical analysis, 39,762) would be protected. (This number is higher than the 29,916 customers of the current MANRS region, because this zone is formed by including all of the largest ASs as measured by their local region.)

In the next section we discuss some details concerning this proposal. In the Appendix, we provide a full specification of the requirements on a member of the VIPzone.

4 Specific considerations

Further reducing the risk of *owner harm*. As we have described the scheme to this point, a remaining harm can penetrate the zone–a hijack based on a sub-prefix (an address block that is a subset of the 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). Trying to distort this rule so that a VERIFIED announcement for a given prefix takes precedence over an unVERIFIED announcement for a longer prefix would risk breaking many operational practices, including local dis-aggregation of prefixes for traffic management purposes.

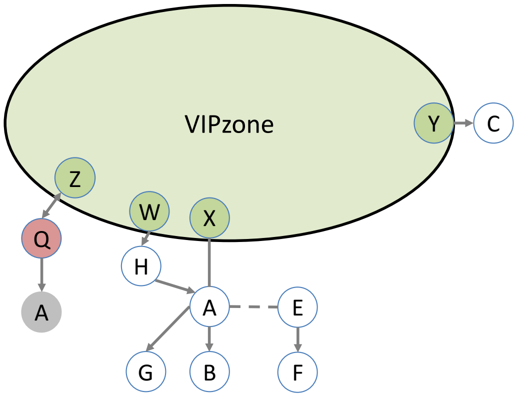

Figure 3: Various customers of a VIPzone, including A with a small local region, C with no local region, and a malicious AS Q pretending that A is a customer.

An AS concerned about owner harm resulting from a sub-prefix attack protects itself by registering ROAs for the prefix. However, the degree of risk mitigation depends on how it configures ROAs. The ROA option *max length* allows a prefix owner to register a ROA that allows a range of valid prefix lengths in announcements. An AS using this option is trading off increased owner risk of a sub-prefix hijack in the zone in exchange for flexibility in how it announces its addresses. The benefit of this flexibility depends on how easy it is for an AS to register new ROAs, and how rapidly they propagate through the Internet. But the choice is up to the AS.

Local regions. An AS directly connected to a zone may have other paths from which it gets BGP announcements. These include the ASs in the customer cone of that AS, the peers of that AS and their customer cones, and any providers (and their customer cones) of that AS that are not in the zone. Consider AS A in Figure 3. A has a provider X in the zone. In addition, its local region includes the provider H, the customers B and G, the peer E and its customer F. Any of these could potentially launch a hijack that represents a misdirection harm. (The mitigation of the owner harm that A achieves by connecting to the zone is not affected by its local region. That mitigation results from the direct attachment to X, and is further improved (from sub-prefix hijacks) if A registers ROAs for its prefixes.)

We make three observations about local regions. First, the risk of hijack by (for example) your own customer is low. Further, ASs outside the zone may or may not implement a robust KYC practice, but if they do, they can detect if their customers are attempting misdirection using a forged-origin attack.

Second, the misdirection from a hijack in the region is restricted to the region. If malicious AS Q launches a path

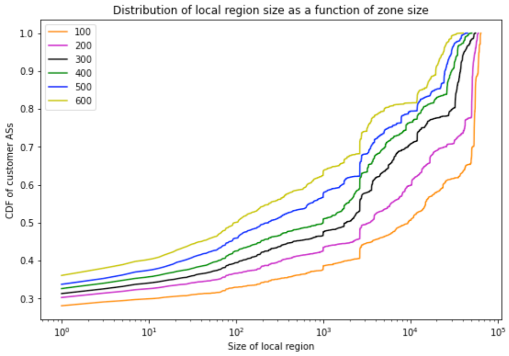

Figure 4: Sizes of local regions for customers of a hypothetical VIPzone, for various zone sizes. For a zone of 600 members, well over half of the attached ASs have a small local region.

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. Such an unverified announcement cannot reach the local region of A.

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 local regions) and compute the size of the local regions for all attached customers. We add to the zone 100 ASs at a time, and at each step compute the size of the local region for the customers. (We start with a sample of 10,000 customers of the initial 100 ASs in this preliminary analysis, to reduce the running time of the analysis).

Figure 4 plots the resulting distribution. Once the zone grows to 600 member ASes, over half of the attached ASs have small local regions. However, a significant number of directly attached ASs have large local regions. In general, these ASs have engaged in open peering with many other ASs to reduce their use of (and thus payment to) transit providers. A realistic consequence of this approach to traffic engineering, i.e., when an AS accepts routes from many peers that do not take known steps to verify their announcements, is an increased risk of hijack.

Some peers may take steps to verify their own customers, and the practical risk of using routes from such a region would be minimal.[2] But again, each AS gets to make its own risk assessment, and act accordingly.

Impact on routing policy The rule that an AS in the zone must prefer a VERIFIED route implies that in some cases, an AS would have to prefer a route inside the zone over a route via its own customer cone, or to a peer not in the zone. Figure 5 shows that for most members of this hypothetical zone, there would be few of these routing exceptions in practice. There are a few outliers that would have many routing exceptions. The largest number of exceptions in the plot is associated with Hurricane Electric. In the AS Rank data for May 2023, HE can reach 20,876 ASs through peering and customer links to ASs outside the zone. Of these, 714 would trigger an exception–about 3.5%. The large number of ASs that HE can reach through its peering gives rise to this large number of exceptions.

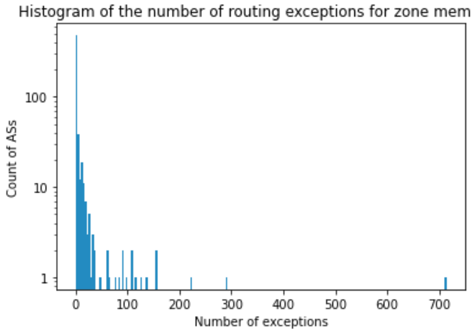

Figure 5: For each member of our hypothetical zone (600 members), the number of exceptions for each AS in the zone, where we define an exception as a destination AS for which the member of the zone, when it prefers a VERIFIED route, must use a provider rather than a peer or a customer to reach that AS.

Protection for ASs not directly connected to the zone . In Figure 3, AS B shares the local region of A, but is not directly connected to the zone. What sort of protection does it receive from hijacks? With respect to owner risk, it 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 exactly the same situation as A: no hijacks will come into the region from the zone, but a hijack in the local region can still cause misdirection harm.

Many smaller ASs 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.

Note that implementation of this approach would likely induce changes in Internet interconnection as ASes shift to exploit the benefits. We use the existing topology for our hypothetical analysis, but there is no reason to imagine that the topology would remain unchanged. Indeed, a key benefit of our scheme is that it creates an incentive for ASs to attach to a member of the zone.

Hardening the KYC requirement 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 that AS) requests that the provider provision a new link to serve that AS. If the provider does not recog-

[2] 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 risk-free at a practical level.

nize 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.

Transit providers will have to harden their business practices to detect these sorts of attacks. This requirement applies equally to the existing MANRS, VIPzone, and the ASPA approaches (Section 6.1).

5 Auditing for conformance

To provide credible protection, a zone must include a capability for checking member conformance with the requirements, and a willingness to suspend or eject members that do not conform. Without a commitment to conformance auditing, the zone may end up nothing but a marketing badge.

Our proposal for a VIPzone does not use real-time detection of suspicious announcement. Real time prevention requires adding code to the BGP processing path in routers or route computation servers. This approach would increase the need for new mechanism, and as well potentially lead to a more brittle scheme, where a harmless error by an ISP originating an announcement would lead to its being dropped by the real-time checking, thus causing loss of legitimate connectivity.

Instead we propose to detect and document failures and hold members accountable. Independent third parties can check conformance off-path, by looking at public BGP announcements. This approach is similar in spirit to how the CA/Browser forum evaluates the correct behavior of certificate authorities. But there must be the will (and the institution) to undertake conformance auditing.

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 two major 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 zone is not conformant.
- Rule 2: If an announcement has an invalid origin, independent of path length, the VIPzone member that introduced the announcement is not conformant. Documented exceptions may be acceptable.

To facilitate conformance checking, the VIPzone could require that members agree to peer with one of the major route collection projects.

6 Relationship to other approaches

6.1 AS Provider Authorization (ASPA)

ASPA is a mechanism that lets a customer register a list of the 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 that is administered by the five RIRs. The ASPA data is globally visible, so any AS receiving a BGP announcement can look at the sequence of ASs in the path, and check to see if there is an ASPA that covers any adjacent pair of ASs in the path. If there is, and the announcement is inconsistent with the ASPA, the AS receiving the announcement can drop it.

ASPA can be used to limit both route leaks and hijacks.We see ASPA as a potentially useful complement to the VIPzone to increase the range of ASs that are protected, if ASPA is used to control hijacks (as opposed to route leaks, which might lead to different deployment practices).

In terms of its design approach, ASPA differs in a number of ways from our proposal.

- The VIPzone design tried to minimize the need for new uses of databases. ASPA depends on (new) records stored in the RPKI.
- The VIPzone design tried to minimize the effort required of a small AS to get protection. It requires that the small AS connect to a transit provider in the zone and (ideally) register its ROAs. ASPA would require that the small AS register an ASPA describing its providers.
- The VIPzone design tries to minimize the need for new mechanism in the routers (or route computation server). The only requirements relate to the VERIFIED flag. ASPA checking requires a new processing check, which includes downloading the relevant ASPA data and inspecting the announcement for validity.
- The VIPzone design assigns clear responsibility for various actions to different ASs. ASs with non-member customers must perform a set of steps. The clear assignment of responsibility allows for conformance checking. The creation of a zone allows a clear description of protection as well as the residual harm. ASPA, as described in the current proposal [1], is not a specification of how it might be used. It is a mechanism, not a proposal for operational practices based on this proposal. Without that companion work, it is not clear which ASs should do ASPA checking, which would have the motivation to register their ASPA, and (as a result) what protection it will achieve.
- The VIPzone requires member ASs to implement KYC procedures. It does not discuss the rigor of these procedures, which may change over time, but these procedures play an essential role. ASPA has a similar requirement

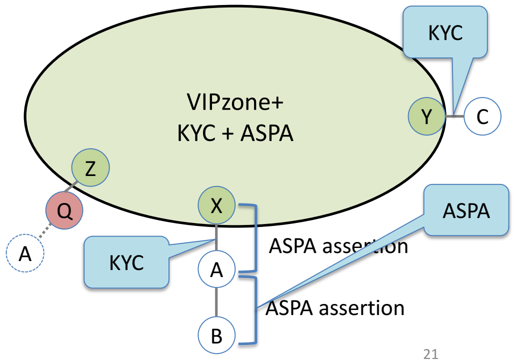

Figure 6: The member X can use the ASPA registered by B to confirm that B is a customer of A. Assuming that A actually has B as a customer (A must have adequate KYC practices to confirm that), X can assume that the announcement originated by B is correct, and mark it VERIFIED, thus protecting B, even though B is not directly connected to the zone.

> for a KYC procedure. Consider the case of an announcement from a customer to a provider. If the customer has registered an ASPA, the provider can verify that the customer AS has listed this provider as acceptable. However, the provider must still confirm that the AS used by the customer is an AS that this customer is allowed to use.

Figure 6 shows how ASPA might augment the creation of a zone of trust to provide protection to more ASs that are outside the zone. In this illustration, Y uses our normal VIPzone practices to verify the announcements originated by C, as does X to verify A. But X now has the option of using ASPA from B to confirm that A is a valid provider of B. Assuming A performs an adequate KYC test to confirm that B is really B, X can now safely mark the announcement of B as VERIFIED, thus extending the zone protection to B.

Section 2.3 described two approaches to creating a zone. ASPA could be used to create a zone of trust in either of those ways. One way is to use ASPA to make sure that no bad announcements enter the zone. In this approach, *every* AS at the edge of the zone must implement ASPA checking. In this illustration, Z must correctly detect that since A has registered its ASPA, Q is not a valid provider of A. If ASPA is deployed in this way, with a strict requirement that every access to the zone do ASPA checking, the VERIFIED mechanism may serve little additional purpose. But if the VERIFIED mechanism is used, then Z has no need to do any ASPA checking, or indeed to even know what ASPA is. Whether the attacker Q pretends that it has A or B as a customer, the announcement will not be marked VERIFIED, and will not be preferred.

While ASPA can be used to verify a valid subscriber-provider announcement, as we illustrate for AS A in Figure 6, it is not clear that using ASPA in this way is useful. In the case of Figure 6, AS A must implement a sufficiently robust KYC process to confirm that the attached customer is actually entitled to be AS B, rather than an imposter. Once A has taken this step, is an ASPA registered by A of any additional value to B, in order that it make a correct decision?

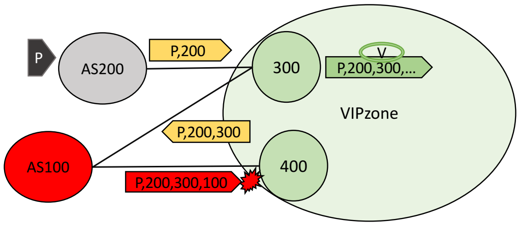

Figure 7: AS 100 legitimately receives from its provider AS300 in the zone a route to AS200. If AS400 leaks this route to AS400, the route (which includes multiple AS hops, will not be marked VERIFIED, and since there is a verified announcement for the prefix, it will not be preferred. The leak has no effect.

Mechanisms such as ASPA or ROV are usually defined without reference to other specified protocols and mechanisms. It is through the crafting of operational practices that different mechanisms are woven together to make an overall approach. The composition of various mechanisms into an overall approach, and the balance of dependencies among them, is an important exercise that defines the overall security of the resulting system.

7 Route leaks

A *route leak* is another kind of BGP announcement that can cause loss of traffic and other operational risks. The distinction between a route hijack and a route leak is that a route leak is not generally a malicious action by the offending AS, but rather the inappropriate forwarding of a BGP announcement that an AS legitimately received. A classic example of a route leak would be 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 belonging to those ASes. Figure 7 illustrates a leak by a multi-homed AS, and how the VIPzone prevents propagation of this leak. AS 100 might incorrectly announce (leak) the path to AS 200 that it receives from one transit provider (AS 300) to its other transit provider (AS 400). Since a VERIFIED path to AS 200 exists in the zone, AS 400 should not propagate its unverified route. If it did, ASes in the VIPzone would never prefer that route, so customers directly attached to the VIPzone would not receive that route, and traffic to AS 200 would never flow from the zone to AS 400.

Blocking route leaks is important for both operational and security reasons, because it is not possible to determine with certainty that a route leak is *not* malicious. A route leak that causes traffic to flow over paths with enough capacity to carry the traffic (so the leak will not impair performance) may persist for a long time without detection. There is no way a network can examine incoming announcements to distinguish an accidental misconfiguration from an attempt by the "leaker" to inspect or perhaps selectively modify traffic as it passes through that AS.

The VIPzone as we have constructed it provides protection from route leaks so long as the AS performing the leak is *not* in the VIPzone. If the leak occurs within the zone, the announcement from AS 300 to AS 100 would be VERIFIED, and when AS100 forwards (leaked) this announcement to AS 400, it would thus be marked VERIFIED, and the routing preference rules will not prevent it from propagating.

Such potential harms from accidental misconfiguration suggest an important insight about VIPzone deployment. A natural but unnecessary — even counterproductive— objective is to maximize the number of ASes in the VIPzone. Smaller ASes (certainly stub ASes) will get the benefit of VIPzone from being a *customer* of a VIPzone member. Actually joining will require that the joining AS correctly implement a range of operational practices, which for smaller ASes with less sophisticated staff may be difficult. Getting these practices wrong may result in malformed announcements in the zone, which will lead to the revocation of their VIPzone status. We consider it preferable that only operators with sufficient technical abilities attempt to join the VIPzone. Other requirements (such as maintaining correct contact information, registering their own prefixes in a public database, implementing anti-spoofing filters) make sense for an AS of any size, and a MANRS-like initiative may want to define two tiers of ISP membership to accommodate different likely capabilities.

8 Conclusion

A zone of trust represents a new way to think about routing security. It is more than "every AS makes its own decisions and defends itself" and less than a global solution. A global solution is not realistic, since malicious actors are inside the system.

A zone of trust, if properly sized so that for some set of users, the zone connects those users to the services they normally use, will provide improved security for those associations, while supporting global connectivity with the level of security available today.

A properly crafted set of rules for a zone will allow a more precise articulation of the level of security, residual harm, and so on, than results from independent action by individual ASs. For this reason, a properly crafted zone can increase the incentive for ASs to join the zone, or to become direct customers of a member of the zone.

As we illustrate with our analysis of the connectivity among the MANRS members, a useful connected region already exists. The challenge for MANRS is to define a set of rules for that region so that it becomes a zone with well-defined security properties, not to bring the connected region into existance.

References

[1] A. Azimov, E. Bogomazov, R. Bush, K. Patel, J. Snijders, and K. Sriram. BGP AS_PATH Verification Based on Autonomous System Provider Authorization (ASPA) Objects, Jul 2023. Work in Progress.

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

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

A Full specification of required actions for members of the VIPzone

VIPzone members must use the following operational practices. 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 [2], 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 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. While our proposal does not require that customers of zone members understand or act on the VERIFIED marking, if they choose to do so they can 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 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. To preserve Internet connectivity, it is critical that unverified routes be forwarded according to normal routing policies.

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.