

Trust Zones: A Path to a More Secure Internet Infrastructure

David Clark and kc claffy
ddc@csail.mit.edu and kc@caida.org

November 2020

DRAFT: Working version. Please contact authors for latest version.

1 Motivation: persistent insecurity of the Internet infrastructure

We propose a path to measurably improve a particular set of Internet infrastructure security weaknesses. By *Internet infrastructure* we mean the Internet as a packet transport architecture: the transport/network layer protocols (TCP/IP), the Internet routing protocol (BGP), and the naming protocol (DNS). Higher-layer security threats – such as malware, phishing, ransomware, fake news and trolling – get enormous media attention. But the less publicized security concerns with the Internet as a packet transport layer can, and sometimes do, destabilize the foundation on which all higher-level activities occur, and facilitate execution of higher-layer malicious actions. It is the foundational nature of the packet transport layer that motivates our focus.

The insecurity of the Internet infrastructure poses a threat to users, businesses, governments, and society at large. As a further point of concern, many of the known security flaws in these systems have persisted for decades. Insecurity persists for five entangled reasons: lack of agreement on appropriate protective measures; misaligned incentives and negative externalities; inability for relevant actors to coordinate actions—especially across national boundaries; the generality of the Internet as a service platform, which allows malicious actors great fluidity in their attacks; and information asymmetries that leave those who need to act without sufficient knowledge to inform planning and execution. While many of these considerations can apply to security challenges more broadly, the generality of the Internet, the tensions among the different sets of private-sector actors, and the lack of any effective mechanism for high-level direction-setting compound the problem.

We do not imagine that we are going to make the Internet “secure”, if by that we mean free of risk. Risk is a part of living, and the Internet experience will be no exception. Our goal should be to reduce the risk to the level that users are not fearful of using the Internet, while preserving the core benefits of the Internet—the freedom from unnecessary constraint.

The call for better security is aspirational. Any serious attempt to improve security must begin by defining it operationally: breaking the problem into actionable parts; carefully studying the constraints, capabilities and incentives of the relevant actors; analyzing the merits and practicality of different approaches; and developing a strategy to achieve sufficient consensus to motivate progress. These steps are part of any serious system security analysis; our goal is to apply this line of reasoning to the Internet infrastructure layer.

1.1 The core systems of the Internet and their flaws

Figure 1 is a representation of the service layers of the Internet.¹ The hourglass shape reflects the design goal of enabling great diversity in the underlying physical technology over which the Internet operates, and great diversity in the applications that run on top of it. The narrow waist plays an essential role in this model, not as a bottleneck, but as a set of common, well-specified protocols that provide a stable layer of packet transport reliable enough to sustain continual evolution and disruption in layers above and below the narrow waist. The greatest strengths of these protocols – well-specified, non-proprietary, and globally implemented – also make it inherently challenging to improve the security of these layers, because significant changes require global agreement to the increased cost and complexity on the whole ecosystem.

¹Adapted from National Research Council (National Research Council. *Realizing the Information Future: The Internet and Beyond*. Washington, DC: The National Academies Press, 1994. ISBN: 978-0-309-05044-9. DOI: 10.17226/4755. URL: <https://www.nap.edu/catalog/4755/realizing-the-information-future-the-internet-and-beyond>, p. 53)

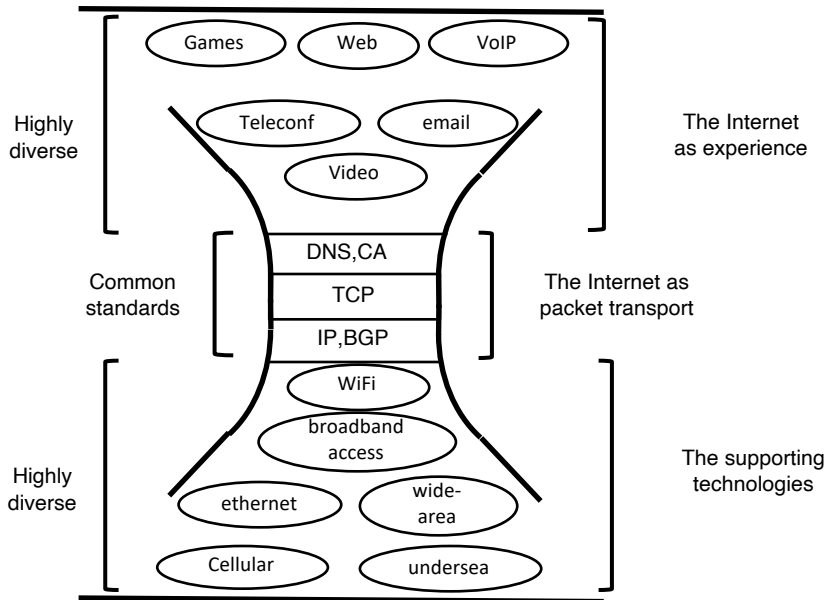


Figure 1: *The hourglass model of the structure of the Internet, capturing the diversity of applications and technology, connected through common agreement on the standards for the core protocols.*

The function of the Internet Protocol (IP) layer is to deliver packets of data.² The Internet Protocol specification states that a router should forward each packet toward its destination address as best it can. This specification says nothing about what else might happen to that packet. The Internet is composed of autonomous systems (ASes) under independent control. An AS might engage in unexpected or unwanted behaviors, such as making a copy of a packet for inspection. End points cannot generally detect such behavior, and the design of the Internet cannot prevent it. Communicating end points protect themselves from unwelcome observation of their traffic by encrypting it.

The IP specification does not include any ability for routers to police or control packets based on their contents. These layers ignore packet content by design. If higher-layer applications facilitate malicious activity such as delivery of malware, expecting the packet layer to identify and stop such packets is comparable to expecting a highway or traffic lights to stop trucks filled with explosives.

The operation of the Internet as a packet carriage layer depends on several critical system elements.

- *Internet (IP) addresses*: every element communicating across the packet layer, i.e., using IP protocols, including end-points and routers, is assigned one or more addresses, so that packets can be delivered to it.
- The *global routing protocol* (the Border Gateway Protocol or BGP)³ propagates topology and routing policy information across 70K+ independent networks called *autonomous systems*. This information enables routers to correctly forward packets to their destinations.
- The *transport protocol* (Transmission Control Protocol or TCP⁴) detects and corrects errors that occur as routers transmit packets across the Internet. Errors might include lost packets, packets with corrupted contents, duplicated or misordered packets, etc. The role of this protocol is only at sending and receiving end-points to detect and remediate these errors, e.g., by retransmitting lost packets. TCP does not operate on packets as they pass through routers. As such, it is less susceptible to abusive manipulation by rogue elements in the network.
- The Domain Name System or DNS translates human-meaningful names (like `www.example.com`) into IP addresses to which routers forward packets. If this system is working in a trustworthy manner, the user will obtain the correct IP address for the intended higher-layer service, e.g., web site, and will not be misled into going to unintended or malicious locations.

²Jon Postel. *Internet Protocol*. RFC 791. Sept. 1981.

³Yakov Rekhter, Tony Lee, and Susan Hares. *A Border Gateway Protocol 4 (BGP-4)*. RFC 4271. Apr. 1996. URL: <https://tools.ietf.org/html/rfc4271>.

⁴Jon Postel, Editor. *Transmission Control Protocol — DARPA Internet Program Protocol Specification*. RFC 793. Information Sciences Institute, University of Southern California, 4676 Admiralty Way, Marina del Rey, CA 90291, Sept. 1981.

- The Certificate Authority system manages and distributes to users encryption keys used for transport connections so that users can confirm the identity of the party with which they are communicating. If this system is working correctly, the user receives a confirmation that the service at the end point receiving the packet is the service the user intended to reach.

To simplify, if these systems are working correctly, the Internet as a packet forwarding system – its “plumbing” – is working correctly. Unfortunately, all of these systems suffer from known vulnerabilities, which attackers regularly exploit, despite decades of attempts to remediate them.

1.2 Internet addressing system (IP)

The network layer of the Internet architecture is most fundamentally defined by Internet Protocol (IP) addresses. IP addresses are an essential part of the Internet. Routers use destination IP addresses in the headers of packets to choose the next hop to forward a packet toward its intended destination. The early designers of the Internet specified the current addressing format (IPv4) in 1981. This format allows for 4.2 billion 32-bit addresses.⁵ In the early 1990s it became clear that the world would require more addresses than fit into a 32-bit field. By then there was a standards organization: the Internet Engineering Task Force (IETF). After much deliberation, in 1998 the IETF standardized on a new addressing format (IPv6) that used 128-bit addresses.⁶ Unfortunately, the IETF decided to make the IPv6 protocol backward-*incompatible* with IPv4, which has greatly slowed if not doomed the transition to the IPv6 protocol. Although parts of the Internet are migrating to IPv6, those parts of the network must support a conversion mechanism in order to communicate with any existing IPv4 network, so long as that network remains IPv4 only.

The framework for allocating IP addresses is hierarchical. The Internet Corporation for Assigned Names and Numbers (ICANN) delegates blocks of addresses to Regional Internet Registries (RIRs), which in turn allocate them to national registries or directly to organizations who operate parts of the Internet called Autonomous Systems (ASes). Because IPv4 addresses are scarce and in demand, an opaque market has emerged for buying and selling IPv4 addresses. The opacity of this market is itself a source of security vulnerabilities related to attribution of IP address ownership.

A better-known vulnerability embedded in the network layer is the ability to *spoof source IP addresses*. To reach its destination, a packet must have the destination’s IP address in its header. Similarly, if the destination is to return a packet to the original source in order to initiate two-way communication, the source address listed in the first packet must correspond to the actual source of the packet. But if a malicious source sends a packet to a destination using a fake source address, e.g., one belonging to a third endpoint, the receiver of the packet will reply to that third endpoint’s address rather than to the original sender. In fact the receiver cannot respond to the original sender since it does not know the actual source; it trusts the authenticity of the source address field in the header. Malicious actors have exploited this vulnerability to mount a variety of attacks, e.g., volumetric denial-of-service⁷ (DoS), resource exhaustion,⁸ policy evasion,⁹ cache poisoning,¹⁰ and impersonation.¹¹ A volumetric denial-of-service attack arises when an attacker can marshal enough traffic to overwhelm a destination or region of the network. An impersonation attack arises when an attacker uses a victim’s address space to launch scanning or other activity likely to induce blocking of the victim’s address.¹²

Note that if an attacker can marshal enough distinct sources of traffic for a *distributed* denial-of-source attack, such as with a botnet, the attack may not need to use spoofed source addresses, although spoofing still offers the attacker the advantage of making attribution difficult if not impossible. Non-spoofed DDoS attacks arise not from design limitations of the network layer, but from persistent vulnerabilities in end-points and applications that allow malicious actors to take over machines without the owners of those machines being

⁵Postel, *Internet Protocol*.

⁶Steve Deering and Robert Hinden. *Internet Protocol, Version 6 (IPv6) Specification*. RFC 1883. Dec. 1995.

⁷Sam Kottler. *GitHub Engineering DDoS Incident Report*. <https://githubengineering.com/ddos-incident-report/>. Mar. 2018.

⁸W. Eddy. *TCP SYN Flooding Attacks and Common Mitigations*. RFC 4987. Aug. 2007.

⁹Christopher Morrow. *BLS FastAccess internal tech needed*. <http://www.merit.edu/mail.archives/nanog/2006-01/msg00220.html>. 2006.

¹⁰US-CERT. *Multiple DNS implementations vulnerable to cache poisoning VU#800113*. 2008.

¹¹Sean Lyngaas. *Someone is spoofing big bank IP addresses – possibly to embarrass security vendors*. <https://www.cybercoop.com/spoofed-bank-ip-address-greynoise-andrew-morris-bank-of-america/>. Apr. 2019.

¹²Matthew Luckie et al. “Network Hygiene, Incentives, and Regulation: Deployment of Source Address Validation in the Internet”. In: CCS ’19. London, United Kingdom: Association for Computing Machinery, 2019, pp. 465–480. ISBN: 9781450367479. DOI: 10.1145/3319535.3354232. URL: <https://doi.org/10.1145/3319535.3354232>.

aware of it. In the early days of the Internet, the designers appreciated this risk in principle but the idea of an attacker subverting perhaps hundreds of thousands of end points for malicious purposes seemed remote. Today, attacks that involve hundreds of thousands of machines, with tens of gigabits of malicious traffic, are regular events on the Internet. Because these attacks are rooted in a higher-layer vulnerability, we do not focus on them in this paper.

1.3 Internet routing system (BGP)

There are about 70K Autonomous Systems that make up the Internet today. Each AS may own a set of IP addresses, and every AS in the Internet must know how to forward packets to these addresses. The Border Gateway Protocol, or BGP, is the mechanism that ASes use to propagate this knowledge across the network topology. Addresses are organized into *address blocks* of various sizes, identified by the *prefix* (the first part) of the addresses in the block. Each AS uses BGP to announce to its directly connected neighbor ASes the prefixes that it hosts. The receiving ASes pass this announcement on to their neighbors, and so on, until (in principle) it reaches all parts of the Internet. As each AS passes an announcement along, it adds its own AS number to the announcement, so the form of the announcement is a series of AS numbers that describe the path (at the AS level) back to the AS owning the associated address block.

The critical security flaw with BGP is well-known: a rogue Autonomous System can announce a falsehood into the global routing system, i.e., a false announcement that it hosts or is the path to a block of addresses that it does not have the authority to announce. Traffic addressed to that block may travel to the rogue AS, which can drop, inspect, or manipulate that traffic. The simplest form of the resulting harm is that traffic goes to the wrong part of the Internet, and is then (in the best case) discarded. This outcome leads to a loss of availability between the parties intending to communicate. A more pernicious kind of harm is that a rogue endpoint can mimic the behavior of the intended endpoint, and carry out an exchange that seems to the victim to be with a legitimate party. This attack can lead to theft of information such as user credentials, which the malicious actor can then exploit. It can also lead to the download of malicious software, or malware, onto the victim's computer. Another possible harm is that the malicious actor may launch some abusive traffic from addresses in that block, damaging the reputation of the owner of the block.¹³

News of some damaging route hijack episodes has appeared in the press or on mailing lists, but the overall level of hijacking is not clear, since victims have a disincentive to publicize that they have fallen victim to such attacks. To understand the current level of abuse, and the importance of seeking ways to mitigate it, a team at MIT and CAIDA¹⁴ developed a scheme to identify malicious routing announcements based on their intrinsic characteristics. Working with five years of data curated by CAIDA, they demonstrated that there are Autonomous Systems that persist as malicious players in the Internet for years, issuing malicious routing announcements and deflecting (“hijacking”) traffic away from its intended destination. Applying some machine learning methods to routing data, that team identified about 400 of the 70K active ASes as highly likely serial hijackers, and another 400 that are probable hijackers.

This BGP vulnerability has been known for decades: it was first documented in a predecessor of BGP in 1982.¹⁵ The fact that the vulnerability has persisted for so long is an indication of the difficulty of reaching resolution on a preferred path forward. The Internet standards community has debated and developed approaches to improve BGP security for at least two decades, but only recently has made what appears to be substantial progress on a small piece of the problem: *origin validation* (Section 5.2).¹⁶

1.4 Internet Naming system (DNS)

The DNS translates a name of the form `www.example.com` into an Internet destination address to use in the packet header to forward the packet. Structurally, the DNS is a hierarchical, distributed database with built-in redundancy. Assignment of responsibility for domains occurs through a process of *delegation*, in

¹³This attack may seem to be an abuse of the addressing system, but it is the routing system that allows one user to appropriate another user's addresses. Spammers will hijack a small block of addresses, send a large volume of spam, and withdraw the hijack. This makes it seem as if the spam came from a legitimate sender.

¹⁴C. Testart et al. “Profiling BGP Serial Hijackers: Capturing Persistent Misbehavior in the Global Routing Table”. In: *ACM Internet Measurement Conference (IMC)*. Oct. 2019.

¹⁵Eric Rosen. *Exterior Gateway Protocol (EGP)*. 1982. URL: <https://tools.ietf.org/html/rfc827>, pg 31, section 9.

¹⁶For a survey of the history of proposed schemes to secure BGP, see (Cecilia Testart. “Reviewing a Historical Internet Vulnerability: Why Isn't BGP More Secure and What Can We Do About it?” In: *TPRC 46: The 46th Research Conference on Communication, Information and Internet Policy*. Sept. 2018. URL: https://papers.ssrn.com/sol3/papers.cfm?abstract_id=3141666).

which an entity at a higher level in the hierarchy assigns responsibility for a subset of names to another party. The hierarchy starts at the *root* of the DNS, which delegates top-level domains (TLDs) such as *.com*, *.net*, *.nl*, etc. to *TLD registries*. These TLD registries in turn delegate to second-level domains, which may further delegate parts of the name space. Administration of these delegations can be a complex task involving many stakeholders, most obviously *registries*, *registrars* and *registrants*. A DNS *registry* administers a TLD. The *registrar* provides an interface between registrant and registry, managing purchase, renewal, changes, expiration, and billing. The *registrant* is the customer that registers a domain. Other players, e.g. Cloudflare, may buy and host domain resources on behalf of registrants. The organization with overall responsibility for the stewardship of the DNS namespace is the Internet Corporation for Assigned Names and Numbers (ICANN) (§??).

Today, harms that leverage the DNS protocols and supply chain represent some of the most pernicious security threats on the Internet. Malicious actors can subvert existing names or register their own names by penetrating databases operated by either registries or registrars, and then use those names for malicious purposes. By penetrating a registry or registrar database, one can add invalid registrant information, or change the binding from a name to an address. Lack of oversight of the competitive for-profit DNS supply chain contributes to these security risks. But the complexity of the DNS also leads to misconfiguration of the name resolution mechanisms by owners of domain names, which can allow malicious actors to take control of them. Finally, and most challenging, is the registration of domain names intended for malicious use such as phishing or malware delivery. Every month the Internet Corporation for Assigned Names and Numbers (ICANN) reports the number of active domain names associated with abusive practices.¹⁷ Since the beginning of the year, the numbers range from a low of 572K in July to a high of 926K in October. Some registrars support operational practices that seem tailored to the needs of malicious actors, such as automatic registration of bulk, meaningless domain names, the creation on demand of “look-alike” or “impersonation” names, or lax attention to capturing the identity of the registrant. Often the DNS is only one component of malicious activities, and stakeholders disagree on whether the DNS is a suitable or effective system through which to combat them.

1.5 Internet Certificate Authority system

The Certificate Authority (CA) system plays a critical role in Internet security. When operating correctly, it provides a means for a user (typically via a web browser) to verify that a connection is to the intended destination—e.g., the correct banking site—rather than a rogue copy of the site. However, the CA system itself is vulnerable to attack and manipulation. Some certificate authorities may issue misleading certificates providing the wrong public key (the verification credential) to a user. The assumption behind the design of the current CA system was that all CA authorities would be trustworthy, even in a competitive for-profit environment with no regulatory oversight. Not surprisingly, this has proven false in practice.

If an attacker can cause the issuance of an invalid certificate, whether by penetrating a CA and subverting it, paying an untrustworthy CA to issue such a certificate, or simply (and in particular for state-level attackers) working with a CA that acts as an agent of the state in issuing false certificates, an attacker can pretend to be an end point that it is not, even if the victim end point uses encryption and authentication to attempt to verify the identity of the other end. This attack complements DNS or BGP hijacks that bring traffic to that rogue end point, which then emulates the expected end point, even to the point of cryptographically identifying as valid.

2 Historical roots of insecurity

Few people other than Internet historians know that the first Internet backbones were created to connect scientific researchers to high performance computing facilities, and that the first general purpose Internet backbone was funded by the U.S. National Science Foundation (NSF) in the 1980s and 1990s. The NSFNET backbone fostered the intermediate evolution of the TCP/IP protocols, as it allowed an operational network to scale to millions of users. In 1994, the U.S. government decommissioned this backbone, and launched ambitious industrial policies to promote competition, and thus innovation, in the emerging Internet transport and domain name industries. The policy goal was to transition Internet communication services to the private sector, make it a commercial undertaking, and have competition be a substitute for regulation.

¹⁷This data is reported in the monthly DAAR reports, which can be found at <https://www.icann.org/octo-ssr/daar>.

But this transition left the world with an Internet architecture not prepared for all the malicious actors that would try to exploit its weaknesses. The original designers of the Internet understood that there would be malicious users on the Internet, and that those users might attack other end-points. However, they concluded that it was not the job of the Internet to police the traffic sent across it. End-points needed to take on the responsibility of protecting themselves. Otherwise, end-points would be trusting the network to protect them, and it did not seem realistic to place that level of trust in the network.

However, the designers did not assume adversaries would be operating parts of the infrastructure, and thus the protocols did not require authentication of addresses, routes, and names. Once it was clear how universal the Internet infrastructure would become, and that malicious actors would compromise parts of the infrastructure layers, the Internet engineering standards community spent years debating and proposing technical solutions to retrofit layers of authentication into these protocols.¹⁸ However, those various solutions have mostly not overcome the misaligned incentives that hinder deployment. In hindsight, it is easy to understand why profit-seeking firms may not be able to justify investment to enhance security. But that realization does not yield a clear path forward.

Securing these central elements of the Internet requires some level of global governance to guarantee consistent interpretation of addresses and names. As part of the commercial transition, and “lessening the burdens of government”,¹⁹ the U.S. government led the private sector in establishing ICANN as the private, multistakeholder organization responsible for global coordination of the Internet identifier systems for the infrastructure industry, including preserving their security and stability.²⁰ Similar competitive market pressures that inhibit investment in security have challenged this multistakeholder model of governance of Internet identifiers.²¹

The history of failed security solutions teaches us that market forces and existing institutions alone will not remedy the harms that these vulnerabilities pose to the Internet, and to commerce that relies on it. Improving the security of these layers is not only technical, but a multidisciplinary challenge with many tensions among divergent stakeholder interests. This complexity applies to the development and deployment of risk-mitigation strategies, but also to understanding their effectiveness, or even to what extent defenses have been deployed.

Given the fundamental architectural weaknesses of the Internet protocol suite, and the Internet’s increasing status as critical infrastructure around the globe, we predict that society, and the governments that represent it, will not tolerate the continued circumstances that put so many unaware Internet users at risk. However, the lack of any significant governmental focus on the Internet for the last 25 years has left a daunting knowledge gap. Although relevant data sources exist in various forms, knowledge is elusive, and where it emerges, often proprietary. Even if governments decide that intervention is indicated, they do not necessarily have enough knowledge to inform strategy. We believe the policy goal of governments should be to enable reliance on transparency, in this case regarding operational practices associated with trustworthy infrastructure, to minimize the need for stronger government interventions. If interventions are necessary, a similar level of transparency is necessary to inform them.

3 Proposed approach: zones of trust

Past attempts to remediate these vulnerabilities have considered technical remedies, such as protocol enhancements. A purely technical approach has often proved unsuccessful. First, the global and multistakeholder nature of protocol development makes consensus difficult or impossible. More problematic, proposing or even standardizing a new technology does not mean that actors will deploy it. Deployment is costly, can have undesirable side-effects, or bring benefit only to others. The Internet ecosystem includes over 70k ASes, more than 1500 DNS registries, all the sovereign countries of the world, billions of users and uncounted application

¹⁸For example, (Testart, “Reviewing a Historical Internet Vulnerability: Why Isn’t BGP More Secure and What Can We Do About it?”) reviews the history of attempts to improve Internet routing security.

¹⁹Internet Corporation for Assigned Names and Numbers. *ICANN Articles of Incorporation*. <https://www.icann.org/resources/pages/governance/articles-en>.

²⁰Internet Corporation for Assigned Names and Numbers. *ICANN Bylaws*. <https://www.icann.org/resources/pages/governance/bylaws-en/>.

²¹Brian Cute. *A conversation about evolving the effectiveness of our multistakeholder model*. <https://www.icann.org/en/system/files/files/draft-evolving-multistakeholder-model-issues-list-25apr19-en.pdf>. Mar. 2019; Internet Corporation for Assigned Names and Numbers. *Board Action on Competition, Consumer Trust, and Consumer Choice Review*. <https://www.icann.org/news/blog/board-action-on-competition-consumer-trust-and-consumer-choice-review>. Mar. 2019.

developers. Not all of them are equally trustworthy. Some may be actively malicious; some just have mutually adverse interests. We must accept this situation and adapt to it. But this reality implies that we must scope our solutions carefully so that they depend only on the actors that are motivated to implement them. Lack of care in shaping the design process can actually allow actors with adverse interests to participate, which will doom it. The Internet is global, but that does not mean that solutions to security problems need to be global.

The premise of our approach is that improving the security and trustworthiness of the Internet will require moving from approaches that require global agreement to approaches that can be incrementally deployed within regions of the Internet. More specifically, our experience of the last thirty years has convinced us that the path to better security does not lie in proposals for global changes to the Internet protocols, but in finding *operational practices* that *regions* of the Internet can implement to improve the security profile of those regions. This approach allows groups that choose to trust each other to define and circumscribe the systems they trust. It is more consistent with trust models in the physical world, where we accept that there are malicious actors, and we attempt to arrange circumstances to minimize our interaction with them, and to interact with potentially untrustworthy actors only in constrained ways.

In our proposed approach, we call regions that embody a common sense of commitment and a decision to distance themselves from the global pool of bad actors a *zone of trust*. The basis for security inside the region is not technical constraints that prohibit bad actors and actions, but a collective decision by actors in the region to behave in more trustworthy ways. Critically, actors that make up a zone of trust must agree on steps that allow monitoring that zone of trust to detect misbehavior. The operational practices must be based on a trust-but-verify framework.

The rules that define a zone of trust are not likely to be defined “top-down”. Zones of trust are likely to be transnational, and not amenable to creation by domestic regulation within one nation. While a set of like-minded nations might come together to draft regulations and practices, the current private sector dominance of the Internet ecosystem suggests that the rules will emerge “bottom-up”, as has happened in some cases—see our discussion of the Certificate Authority system in Section 5.5. The success of a set of rules that define a zone of trust will depend on a set of checks and balances that respect the interests of the various legitimate actors. The leadership of a dominant actor may be an effective starting point for the creation of rules, so long as that powerful actor takes care that it not create rules that benefit itself.

This idea is not new, even on the Internet. The premise of shared blocklists or threat intelligence is to exclude actors known to be untrustworthy from an otherwise trusted environment. Response Policy Zone (RPZ)²² is a technology to implement a customized DNS policy for recursive DNS resolvers to modify responses to DNS queries in order to block user access to malicious hosts. But scaling this aspiration beyond the scope of a few networks, including to broad regions of the world, requires a more rigorous, general, and measurement-based approach.

We believe the current trends toward a flatter topology, accompanied by regionalization of connectivity to improve performance (§6), provide a basis that facilitates our proposed approach. As users more commonly depend on only a region of the Internet infrastructure for what they do, operators can construct a more secure and trustworthy experience inside that region, by preventing, or at least hindering, actors outside that region from disrupting it. This approach requires identifying operational practices for which incremental deployment brings collective benefit to those groups who collectively deploy them. Groups who choose to explicitly trust each other can then define rules that protect the systems on which they depend. Importantly, these rules must include detection and management of violations.

We emphasize that such regions may be *topological* rather than, or in addition to *geographic*. Also, different threats may imply/require different region shapes. For one threat, the region might be jurisdictional, for another connected set of ASes. So long as an activity operates within a zone of trust relative to the corresponding threat, the activity will benefit from enhanced security.

With respect to the Internet addressing and routing systems, a zone of trust might be the set of interconnected regions (Autonomous Systems) that agree to verify address ownership of their customers, flag unverified announcements as coming from outside the zone, and reject announcements from outside the zone if they conflict with announcements from within the zone.

With respect to the naming system (DNS), a trust zone might be defined by a commitment to block access to domains or URLs based on a determination that they host abusers, and only use registries and

²²Paul Vixie and Vernon Schryver. *DNS Response Policy Zones (RPZ)*. Dec. 2016. URL: <https://tools.ietf.org/html/draft-vixie-dns-rpz-04>; Paul Vixie. *Taking Back the DNS*. <http://www.isc.org/community/blog/201007/taking-back-dns-0/>. July 2010.

registrars that comply with operational practices to minimize and combat abuse.

With respect to the Certificate Authority (CA) system, the trust zone is currently defined by the providers of browsers, who determine that they will not trust (e.g., not use) certain certificate authorities.

In summary, a sustainable zone of trust must have clear rules about acceptable behavior, a commitment to measurement to detect rule violation, a commitment to deal with rule violation, constraints that limit the ability of bad actors outside the zone of trust to disrupt its operation, and design of applications so that their dependencies stay within the zone in which they operate. This paper elaborates on this idea, and explains why we believe a measurement-supported *zone of trust* approach is the best trajectory to deal with these security challenges at the Internet infrastructure layer and contribute to a more secure and trustworthy Internet experience.

4 Elements of our approach

Our approach depends not on understanding the details of individual attacks, but rather understanding the degrees of freedom that an attacker has. It depends on analysis, informed by detailed system knowledge, to understand where attackers have the least flexibility or the most vulnerability in the construction of attacks, with the goal of proposing operational practices that exploit these weak points in the attackers' options. This process underpins any defense systems analysis—our goal here is to apply it to the Internet.

4.1 The generality of the Internet

The Internet was designed to be a general purpose platform suited to support a wide range of applications. This generality is part of what has made the Internet so successful. However, malicious actors exploit this generality as they maneuver to avoid detection and disruption. As one example, botmasters that take over vulnerable end points to build a botnet must devise a way to control these so-called zombie computers. Defenders try to disrupt these control systems, and botmasters exploit the generality of the Internet to devise new schemes to control their botnets. A botnet control system is, from the perspective of the Internet infrastructure, just another application, and the Internet was designed to support as wide a range of applications as possible. Just as its generality is a boon to the innovator, it is a boon to the attacker.

In the attempt to make the Internet more secure, this generality has two implications. The first is we must study the overall process by which the malicious activity executes, to find the points in that process where the attacker has the least flexibility. For many criminal activities, that point may have nothing to do with the technical character of the attack, but instead how money flows to the attacker. One must resist the temptation to put in place remedies that just chase the bad guys from place to place, if the result is mild inconvenience to the attackers but large cost to the defenders.

The second implication of this generality is that barriers to malicious activity may risk collateral harm to legitimate activities, because the barrier may have to be broad in design to thwart the ability of the attacker to exploit the intrinsic generality of the Internet. This reality has been a point of great concern to many people responsible for operating Internet infrastructure. The core objective of the Internet is availability. Security by definition degrades availability because it raises protective barriers. The risk of collateral harm from an over-broad remedy is not restricted to security practices online—it can arise as well in the design of law. In fact, the balance of freedom and order is a fundamental and recurring challenge to society. The tension emerges here in particularly stark terms because the very specific goal of the Internet (be available and deliver data) and the goal of security (block suspicious activity) seem in direct contention.

Since drawing a precise line between acceptable and unacceptable behavior is practically impossible, a push for better security must accept inconveniences for legitimate users, in the interest of minimizing room for malicious actors to maneuver. For this reason, many designers are uncomfortable deploying protective mechanisms that may block legitimate activity. Similarly, ISPs, who have the primary responsibility for realizing the availability of the Internet, resist mechanisms that accidentally block legitimate activities, because irate and confused users tend to call customer service, which generates costs for the ISP. Design of mechanisms that may cause collateral harm will work best when the user has the means to circumvent the mechanism by explicit action (e.g., the damage is inconvenience, not total prohibition), and the presence of the mechanism is visible to those legitimate users, so they can understand what happened and why. For example, some top-level domains (TLDs) are relevant world-wide, others may be important only regionally. If a TLD with regional importance is infested with many names used for abusive purposes, requiring explicit acknowledgment of risk for users outside its primary region of utility might be quite acceptable.

4.2 The role of measurement

In tactical defense, gathering security-related data centers on what the bad guys are doing at a given moment. Maintainers of blocklists try to infer which address and naming resources attackers are using, on an ongoing basis. This data is evanescent. Interdiction and forensics may be useful to respond to ongoing attacks, but they do not shift the playing field toward the defenders. For example, defenders that attempt to deal with the registration of domain names for malicious purposes are locked in an endless battle with the malicious actors, who adapt to interdictions as fast as they appear. So long as the defenders only try to find the bad guys and chase them from where they currently are to some new place, the generality of the Internet works against the defenders.

The role of data collection and analysis in our approach is central to the following more strategic objectives. We do not mean to trivialize these tasks by listing them as bullet points—these will be substantial research efforts. However, we believe that this is the most viable path to tilting the playing field toward the defenders.

- Understanding malicious behavior in order to craft operational practices that hinder it. This objective requires modeling the scope of an adversary's options.
- Arguing that a practice will measurably improve security posture, e.g., reduce an attack surface.
- Tracking actual levels of different types of abuse so that we can make a plausible argument that levels of abuse are changing as we deploy new practices. While a given practice may not be easily linked causally to changing levels of abuse, if we cannot get data about levels of abuse, we are shooting in the dark when we claim progress.
- Understanding the baseline behavior, including how application design is evolving, and behavior of users invoking those applications. Establishing a set of operational practices that define the zone of trust requires balancing constraints on a range of acceptable options – to hinder the bad guys– against the risk of inhibiting innovation. This balance is practical only to the extent that applications continue to manifest regionalization behavior (§6), but our approach builds on the forces inducing such regional structure. We discuss this trend in Section 6.
- Verification of compliance with accepted operational practices by actors that have committed to the practice.

Many debates about operational practices occur in a context devoid of data. A core premise of our approach is that long-term gathering, curation and analysis of data is critical to a methodical approach to improving these elements of security.

4.3 Engagement with stakeholders

This approach targets turning knowledge derived from understanding system characteristics and ongoing data analysis into open, actionable knowledge that is relevant and meaningful to various actors in the ecosystem, including those responsible for protecting infrastructure. Different threats will likely call for zones of different shape and dimension. Some zones might be jurisdictional, others topological. Measurement must provide critical knowledge that will serve as the basis for establishing new operational practices, validating commitments, and understanding their effectiveness.

The design of proposed operational practices must also rely on a pragmatic recognition of what various actors are willing and able to undertake. Developing an understanding of incentives, costs, and externalities is as important as developing an accurate model of system operation. This means a large component of this project will be transferring knowledge generated by this effort to policy development and cyber-crime communities. We also believe this approach should include an international advisory team that includes policy makers, operators, and researchers to advise on the role of measurement, analysis, and business incentives in developing operational procedures.

5 Applying measurement-based approach to specific problems

We consider the role of measurement in developing and evaluating enhanced operational practices that networks can deploy incrementally to tackle the specific security challenges we have outlined (§1.2-§1.5).

Dedicated stakeholders have provided a head start in defining operational practices relevant to some of our challenges. For example, a group of network operators, facilitated by the Internet Society, have defined a set of operational practices that can prevent several types of addressing and routing abuse (§1.2 and §1.3). This code of conduct, launched in 2014, is called the *Mutually Agreed Norms for Routing Security (MANRS)* initiative.²³ The MANRS initiative draws on well-established operational practices defined by the Internet standards or operational communities. The practices that MANRS currently requires are modest, but represent an excellent first step, and a natural target for targeted measurement and analysis.

5.1 Abuse of Internet addressing system

One recommended MANRS practice is that ISPs prevent traffic with spoofed source IP address from leaving their network, also known as *Source Address Validation (SAV)*. Incentive misalignment of this practice represents a classic negative externality in the Internet infrastructure market: networks that allow spoofing save on their own operational costs, while imposing costs on others (in the form of attacks and attack risk). The Internet Engineering Task Force defined source address validation as Best Current Practices (BCP) 38 and 84.²⁴ These RFCs specify steps that an ISP should take to ensure that its customers do not abuse, even accidentally, the addressing scheme by sending invalid source addresses. Participation in the MANRS initiative requires that ISPs commit to implement SAV.

Given that BCP 38 is accepted in principle by industry, an obvious role for measurement is to encourage increased uptake of the recommended practice by measuring compliance with the practice by ISPs. The measurement/policing challenge for SAV is that the point of measurement for a given ISP must be internal to that ISP. Thus, one requirement for compliance must be that an ISP that commits to implement BCP 38 must also commit to hosting one or more measurement points. This requirement illustrates the point that commitment to a set of practices must include a commitment to measurement to verify compliance.

The Internet Society, in its role as facilitator of MANRS, has no independent measurement tools to verify compliance with the requirements. They currently depend on data from CAIDA to verify compliance with the MANRS requirement that operators do SAV. CAIDA and collaborators previously spent many years operating this capability in the Spoofer project:²⁵ to prove to an independent third party that a given network has properly deployed SAV.²⁶ In this most recent study, CAIDA showed no evidence that those who asserted a commitment to deploy SAV were any more likely to properly deploy it than others. This discovery illustrates the need for measurement and analysis to support deployment of operational practices, and to assess the effectiveness of initiatives to promote them. Another relevant feature of the Spoofer project is that it transformed technical knowledge into actionable knowledge. The most common use of this tool thus far has been to help operators diagnose their own SAV configurations, a function the private sector has understandably found no incentive to commercially offer.

5.2 Internet routing system (BGP)

We have recently seen growing acceptance of a step toward better BGP security. The Regional Internet Registries, who maintain databases of address block ownership, and the Internet protocol standards community (IETF) have developed a protocol for voluntary use of *Route Origin Authorization (ROA)*, a mechanism to establish definitive authority to originate a specified prefix into the global routing system. Uptake of this technology is low, although growing. Globally, about 19% of the Internet address space is protected by ROAs as of October 2020.²⁷ The MANRS code of conduct (§5.1) includes a requirement that member ISPs will check the BGP origin announcements of their customers to make sure that the AS and prefix are valid. Dropping BGP announcements that fail this test prevents simple forms of BGP hijack, i.e., origin hijacks.

²³Internet Society. *Mutually Agreed Norms for Routing Security*. <https://www.manrs.org/>.

²⁴P. Ferguson and D. Senie. *Network Ingress Filtering: Defeating Denial of Service Attacks which employ IP Source Address Spoofing*. IETF BCP38. May 2000; F. Baker and P. Savola. *Ingress Filtering for Multihomed Networks*. IETF BCP84. Mar. 2004.

²⁵Matthew Luckie, Ken Keys, Ryan Koga, Rob Beverly, kc claffy. *Spoofer Source Address Validation Measurement System*. <http://spoofer.caida.org>. 2016.

²⁶M. Luckie et al. “Network Hygiene, Incentives, and Regulation: Deployment of Source Address Validation in the Internet”. In: *ACM Computer and Communications Security (CCS)*. Nov. 2019.

²⁷NIST. *RPKI Monitor*. URL: <https://rpki-monitor.antd.nist.gov/>.

5.2.1 Toward a stronger set of operational norms for routing security

The growing uptake of ROAs suggests two obvious next steps that we illustrate here: demonstrating how ROAs can improve security by identifying who/how/where networks use them, and evaluating additional operational practices that use ROAs to further enhance routing security.

Once address owners register their addresses using ROAs, ISPs can use that information to detect and discard invalid routing announcements, i.e., those inconsistent with a ROA. This step improves BGP security by preventing the acceptance and propagation of invalid source announcements. MIT and CAIDA have developed a method to track which ISPs are dropping invalid routes,²⁸ and will continue to track this over time. Industry has indicated that this is useful actionable knowledge.²⁹ A growing set of ISPs that drop invalid routes will create pressure on other ISPs to register their ROAs, and correct errors in registration, which can in turn motivate further adoption of dropping invalid routes, creating a virtuous cycle toward improved routing security.

However, hijackers can launch more sophisticated attacks, which involve an *invalid path announcement*. The general form of this attack is that the customer provides a BGP announcement with (perhaps several) AS numbers in the path, where the first is a valid origin (AS/prefix) and the last is the valid AS of the customer. In other words, the hijacker is asserting that it has customers, one of which is this AS, for which there is a valid AS/prefix ROA.

To block this option for route hijacking, MANRS needs to tighten its operational practices over time. We envision an approach, which we call *recursive MANRS*, that requires that every MANRS-compliant ISP know which of its customers is also MANRS-compliant. This information will not change rapidly, so it should not be a burden for ISPs to track it. If the customer of a MANRS-compliant ISP is also MANRS-compliant, then that provider ISP can assume that the customer ISP has checked its own customers, and it can safely accept the path. If the ISP's customer does not participate in MANRS, the ISP should treat any BGP announcement from this customer as suspect. If the ISP receiving this suspect announcement from this customer has another route to the same origin that is not suspect, it should discard the suspect route independent of the AS path length.

There is an alternative approach to preventing the propagation of invalid path announcements, called Autonomous System Provider Authorization, or ASPA.³⁰ It is interesting to contrast the features of the two. ASPA proposes a new, global, cryptographically signed database (perhaps stored in the same locations as the ROA data) in which each AS in the Internet records the list of all valid transit providers of that AS. If all ASes within a zone of trust have recorded an ASPA, then any AS within the zone that receives an invalid (hijack) route announcement can detect it. ASes within the zone are then protected from hijacks based on invalid path announcements.

In recursive MANRS, the knowledge of which providers an AS is using is implicit—the knowledge is not publicly recorded in a global data base but results from business agreements between provider and customer ASes. Since the data is not globally known, only a router at the point where a MANRS-compliant AS receives a route announcement from a non-compliant customer can perform the check. In ASPA, any router can perform the check. Recursive MANRS is thus an enhanced practice that all MANRS-compliant ASes must implement.

The advantage of recursive MANRS is that there is no global, public database. A database of that sort may be a substantial barrier to deployment, as it requires every AS to publicly disclose its potential transit providers, and it may be a target of malicious attack to corrupt the information. Corrupting the database could effectively drop an AS from the Internet. On the other hand, the global database may allow the detection (and blocking) of certain forms of route leaks.

This discussion of preventing hijacks based on invalid path announcements illustrates that a given mechanism, e.g., the use of ROAs, can play a role in a range of operational practices with different security outcomes. Simple dropping of route announcements where the ROA makes the route invalid will prevent hijacks based on an invalid origin. Dropping route announcements according to the recursive MANRS rule additionally prevents hijacks using an invalid path. Different operational practices may trigger different incentives by the various actors to deploy the practices, which the system security analysis must consider.

²⁸C. Testart et al. “To Filter or not to Filter: Measuring the Benefits of Registering in the RPKI Today”. In: *Passive and Active Measurement Conference (PAM)*. Jan. 2020.

²⁹Carl Lagerfeldt and Johan Gustawsson. *Routing Security: RPKI Update Q2/20*. 2020.

³⁰Alexander Aximov et al. *Verification of AS_PATH Using the Resource Certificate Public Key Infrastructure and Autonomous System Provider Authorization*. Nov. 2020.

5.3 Internet Naming system (DNS)

The Domain Name system is more complex than BGP, and the path toward better security is less clear. There are more layers of operation, more players in the ecosystem, more tensions among them, and fundamentally more things to go wrong. In contrast to BGP, where the Internet Society launched the MANRS initiative to promote a set of operational practices that would reduce the attack surface for abusers of routing system vulnerabilities, there are no widely accepted recommendations for practices that would improve security. It is not clear who could play the role of developing and incentivizing operational norms. One candidate for this role is ICANN, with its significant responsibility for stewardship of the DNS supply industry. But there is growing evidence that security and consumer protection concerns have been a casualty of the multistakeholder model, where lowest-cost operational practices take priority over secure ones. That ICANN is funded entirely by this industry creates a set of conflicting incentives that are in direct tension with development of stronger operational practices, much less enforcing them.

Before we investigate the development of DNS operational practices, we must evaluate how to define, discover, quantify and continuously monitor aspects and trends of DNS abuse, in order to find patterns of abusive behavior that can motivate operational practices. An equally important goal is to develop a model that illuminates the incentives of the various actors in the ecosystem, which requires understanding the money flows, contributions of the bad actors to the flows, and degrees of freedom for all parties.

The starting point is to develop data collection infrastructure that enables mapping of trust dependencies, relationships among domains and nameserver infrastructure, and unintended attack surfaces that result from current operational practices. Researchers have already explored the use of TLD zone configuration files (called *zone files*) to study operational practices reflected in those zone files, including anomalous patterns in the DNS that correlate with suspicious or risky registrar or registrant behavior.³¹ Another source of data relies on these TLD zone files to seed active measurements of DNS infrastructure. Three Dutch research institutions (SURFnet, SIDN Labs, and U. Twente) – operate the OpenINTEL project, a system for comprehensive active measurements of the global DNS.³² OpenINTEL uses these zone files to drive DNS queries for all covered domain names once every 24 hours, covering over 232 million domains per day.³³

5.4 Toward a stronger set of operational norms for DNS security

A serious commitment to sustained data infrastructure will be necessary to enable the research that will serve as the basis for proposing and socializing a measurable set of enhanced operational practices such as the Internet Society has stewarded for the routing system. ICANN’s Security and Stability Advisory committee has for many years published documents on how to mitigate the many risks to security of a domain name in its lifecycle, and advised stakeholders accordingly.³⁴ The recommendations in these reports could serve as the basis for discussing a proposed code of conduct.

A two-pronged strategy would prioritize implementation and support for technical capabilities to verify one’s own (or others’) compliance with proposed best practices, and to assess the attack surfaces from (often unintentional) failure to comply. A second prong would be to foster and participate in a cyclic feedback relationship between actionable knowledge that influences policy development and the use of actionable knowledge to refine and inform technical knowledge used in DNS abuse mitigation communities.

³¹**akiwate:usenix21**; Dave Piscitello and Dr. Colin Strutt. *Criminal Abuse of Domain Names: Bulk Registration and Contact Information Access*. <http://interisle.net/sub/CriminalDomainAbuse.pdf>. Oct. 2019; Greg Aaron. *Domain Name Registration Data at the Crossroads: The State of Data Protection, Compliance, and Contactability at ICANN*. <http://interisle.net/sub/DomainRegistrationData.pdf>. Mar. 2020; G. Akiwate et al. “Unresolved Issues: Prevalence, Persistence, and Perils of Lame Delegations”. In: *ACM Internet Measurement Conference (IMC)*. Oct. 2020.

³²Roland van Rijswijk-Deij et al. “A High-Performance, Scalable Infrastructure for Large-Scale Active DNS Measurements”. In: *IEEE Journal on Selected Areas in Communications* 34.7 (2016). ISSN: 0733-8716. DOI: 10.1109/JSAC.2016.2558918. URL: <http://ieeexplore.ieee.org/lpdocs/epic03/wrapper.htm?arnumber=7460220>.

³³<https://openintel.nl>

³⁴ICANN Security and Stability Advisory Committee (SSAC). *SSAC Advisory on Registrant Protection: Best Practices for Preserving Security and Stability in the Credential Management Lifecycle*. <https://www.icann.org/en/system/files/files/sac-074-en.pdf>. Nov. 2015; ICANN Security and Stability Advisory Committee (SSAC). *SSAC Response to the new gTLD Subsequent Procedures Policy Development Process Working Group Initial Report*. <https://www.icann.org/en/system/files/files/sac-074-en.pdf>. Oct. 2018.

5.5 Internet Certificate Authority system

Data collection and analysis are critical to sustaining and improving the security of the CA system. Currently, the CA system is one of the better instrumented of the systems we consider. A consortium led by Google instituted a logging system for the CA system called Certificate Transparency (CT). The idea behind CT is that any authority issuing a certificate must also disclose it in a distributed public log that anyone can examine. CAs can still issue a rogue certificate but they cannot do it secretly. Browsers can search the CT log for a certificate before accepting it. Owners of domain names (or others acting on their behalf) must check the CT logs to detect rogue certificates.

Measurement and analysis to detect misbehavior is essential, since problems continue to arise. A detailed analysis of errors and inappropriate actions by Certificate Authorities³⁵ classified 379 incidents of misbehavior by those authorities between 2008 and 2019, with causes including human error, lack of required auditing, system penetrations, and mis-configured or buggy software. They identified 30 probable or confirmed issuance of rogue certificates. The authors' conclusions echo the conclusions about the DNS: Certificate Authorities are profit-seeking firms whose primary goal is to sell certificates. There is evidence that motivations of profitability leads to actors compromising the expectation that they will serve the public interest by operating at the highest level of integrity and quality.

The CA system also has the most well-developed industry-led effort to discipline this market. The Certification Authority Browser Forum (CA/Browser Forum)³⁶ designs and publishes guidelines regarding the issuance and management of digital certificates, and identifies CAs that do not conform. The various web browsers include a list of CAs that the browser will trust when verifying a query, and the developers of the different browsers have dropped many CAs from their list of trusted authorities, which means that when a user attempts to connect to a web site that uses a certificate from one of these untrusted authorities, they will receive a warning message and have to take explicit steps to bypass the warning and proceed.

There are several lessons we draw from study of the CA system.

- This work illustrates the value of data collection and analysis, both to understand the extent of the problem, and to provide support for proposals as to how to mitigate it.
- The CA/Browser Forum represents what seems to be a functioning industry organization that has taken steps to improve security, led by a heavily capitalized and incentivized player (Google).
- The CA/Browser Forum has accepted the necessity of causing possible collateral harm to improve security. Certificates from untrusted authorities will either be rejected or trigger a warning to a user that they are about to engage in a potentially dangerous action.
- It is not clear what recourse a CA has if it is declared untrustworthy. Designers of systems like this, based on a private sector tribunal, have considerable latitude to determine the checks and balances and the rights of recourse.

6 Regionalization: the evolving character of the Internet

A central element of our approach is to find solutions that do not require global consensus and implementation. Regional approaches (§3) will allow groups that decide they will choose to trust each other to define and control the systems on which they depend. We refer to regions that embody a common sense of commitment and a commitment to distance themselves from the global pool of bad actors as *zones of trust*. To be effective, this approach must include mechanisms to keep typical activities of users inside such a zone. An observation about the Internet's changing character gives us some confidence that we can find a path to success. We provide some empirical evidence for this observation in §6.1.

The design goal of the Internet was and continues to be that any two machines anywhere on the Internet could freely communicate. A packet might cross several ASes to reach its destination, but today most traffic traverses only one, in large part due to the goal of efficient delivery of high-volume content from large providers, e.g., Netflix, Amazon, and YouTube, to access providers. Such content providers strive to stage their content in intermediate servers, and attach them directly to large broadband access providers

³⁵Nicolas Serrano, Hilda Hadan, and L. Jean Camp. "A Complete Study of P.K.I. (PKI's Known Incidents)". en. In: *SSRN Electronic Journal* (2019). ISSN: 1556-5068. DOI: 10.2139/ssrn.3425554. URL: <https://www.ssrn.com/abstract=3425554> (visited on 11/10/2020).

³⁶<https://cabforum.org/>

at geographically distributed points. In this case, not only does the traffic need to cross only one service provider, but the traffic enters the access network at a point close to where it will exit to the consumer.

Many application designers similarly use *cloud* platforms and associated services to host applications and content close to the users, thereby shortening the path data takes across the Internet, optimizing performance for users, infrastructure support cost for themselves, and resilience. In the U.S., the largest cloud providers are Amazon Web Services (AWS), Google Cloud Platform (GCP) and Microsoft Azure. As an example, AWS organizes its cloud platform into regions, within which are availability zones for resilience. As of October 2020, AWS has 24 regions (3 more announced) and 77 availability zones across the globe.³⁷ An access network may connect to multiple of these regions and zones, so depending on the deployment scenario, traffic may originate and terminate at points on the access network that are relatively close to each other.

One outcome of this evolutionary trajectory is that traffic on the public Internet becomes more *localized*; the role of the Internet becomes the consumer-facing mass-market access method to the larger Internet ecosystem, with more of the traffic exiting the public Internet onto these other platform assets as soon as possible..

6.1 Measuring regionalization

Measuring the degree to which the Internet experience has become more localized is challenging, because the degree to which a user's experience is localized will depend on where that user is within the Internet. A user attached to a large U.S. broadband access provider will probably have a more localized experience than a user from the developing world. In addition, measuring user traffic destinations raises issues of data collection and privacy.

As an initial exploration to assess how localized the Internet experience is becoming, we took two different datasets of popular destinations on the Internet, and measured their AS path lengths (number of ASes crossed to reach them) starting from the home location of one of us. This sort of exploration yields only anecdotal insight, and (as mentioned above) is highly colored by the fact that the origin used for the exploration was served by a well-connected U.S. broadband access provider.

For our first experiment, we started with the Cisco Umbrella list,³⁸ which lists the top one million URLs world wide, and extracted the top 1000 second-level domain (SLD) names (names like google.com, or netflix.com). Based on its proprietary sampling methodology, Cisco infers these SLDs to be the most popular world-wide. This global list may not accurately represent a typical U.S. broadband access network user experience, but it provides a starting point. In November 2020, we used our local DNS resolver on our broadband residential broadband connection to map each SLD, or a popular subdomain of the SLD, to an IP address.

Of these 1000 SLDs, the associated address for 630 of them was in an AS directly connected to our access provider. In other words, the path from our home to the destination crossed the access provider and went directly to the AS hosting that address.

87% of those SLDs had addresses that were hosted on large cloud and CDN providers such as Amazon, Google, Microsoft and Akamai. This (again anecdotally) illustrates the migration of applications and services into the cloud.

379 of the SLDs were reached by crossing multiple ASes. (The sum of the two exceed 1000 since we reached a few SLDs directly as well as indirectly.) 76% of the paths to these SLDs went through the traditional Tier 1 providers such as Level 3, Cogent, NTT, or Telia. 49 of the SLDs were in China, reflecting the global nature of this list.

To more faithfully reflect U.S. residential traffic, one could instrument a set of U.S. users, but such data collection raises serious privacy concerns. However, we can experiment on ourselves. The Firefox browser records a history of visited URLs, so one of us looked at the URLs visited from our own browser to see where we had been going. The browser logged 8791 URLs from which we extracted 1452 distinct domains. We could resolve all but 21 of them in late November 2020. We assume the others were no longer active or could not be resolved for some other reason. Of those, we found 747 (52%) of those domains in address space of ASes directly connected to our access network, and 754 by an indirect path. (The sum exceeds 1452 because again some we reached both directly and indirectly.) Again, 81% of the directly connected domains were hosted by a major CDN or cloud provider, illustrating the growing use of cloud services to host many web

³⁷<https://aws.amazon.com/about-aws/global-infrastructure/>

³⁸<https://umbrella.cisco.com/blog/cisco-umbrella-1-million>

sites. Beyond that 81%, most of the destinations were directly connected customers of the access provider. There were a total of 70 directly attached ASes where the domains were hosted.

For the domains that were reached by a path with more than one intermediate AS, 92% of the paths exited the access provider across one of the traditional Tier 1 transit providers.

One further question relates to the character of the AS paths, especially the longer paths. The data here is sometimes ambiguous: the traceroute tool we use sometimes does not properly report all the Autonomous Systems along the path. But of the paths where we have a reasonable confidence that the data is correct, we found 43 domains that were four hops away (located within 20 different ASes) and 9 domains that were 5 hops away (located within 2 different ASes.) A few terminated outside the U.S., but most of them went to a destination that was reached through a Tier 2 provider attached to the Tier 1 that was the exit path from the initial AS. For the set of URLs in this sample, we found no paths that were longer than 5 ASes.

Our high-level conclusion from these preliminary explorations is that many web sites today are provisioned in a distributed and replicated way, which means that the path to them, at least across a major U.S. broadband access provider, is a direct path from that access provider to the location of the web site. On the other hand, connectivity via traditional transit providers seems to still be critical. About half of our observed connections depended on these paths. But these paths were still relatively short: they either crossed a tier-1 provider to a directly connected customer or to a customer attached to a tier-2 provider of that tier-1 provider.

6.2 Incentives to move content and services to the cloud

An additional element of the observed regionalization of the Internet is the behavior of enterprise customers. Enterprise customers, as well as application developers, are moving to the cloud. To improve the performance and security of these enterprise systems, there are networks, distinct from the public Internet but often with global reach, that offer to connect enterprise locations to cloud locations. AWS, for example, partners with such providers, which they call Direct Connect Partners, as alternatives to the public Internet to reach AWS from enterprise sites.

These alternative networks can offer better service commitments than the public Internet. Exactly because the Internet is composed of many interconnected ASes, each operated by a separate firm, cooperation and coordination among these firms is required to ensure a specified level of performance. This level of coordination is hard to achieve among AS providers who are competitors at the same time that they interconnect. The service traditionally provided by the public Internet has been called *best effort* service—the collective set of ASes do their best to deliver traffic, but make no specific performance or reliability commitment. The cloud networks operated by AWS or Google, even though they may have global reach, are under the control of one firm that can engineer and manage its network to make stronger service commitments. Similarly, the third-party networks (like Amazon’s Direct Connect Partners) that provide enterprise interconnect to cloud providers are operated by one entity, which can control network characteristics.

Innovation in the cloud ecosystem provides new options for application designers, and how application designers choose to exploit these assets influences how the ecosystem evolves. This evolution is not a planned process, but an emergent phenomenon. A metrics-based zone-of-trust approach can leverage this evolutionary trend to improve the security of the Internet for most users, and importantly, without threatening the role of the public Internet in enabling permissionless innovation at the edge. As users increasingly depend on only a region of the Internet for what they do, that region can provide them a more secure and trustworthy experience by undertaking operational practices that prevent, or at least hinder, actors outside the region attempting to disrupt operations in the region.

An important assumption underlying our proposal is that neither the process nor outcome will disrupt the globally interconnected character of the Internet. Two end points anywhere on the Internet can still exchange traffic directly. We distinguish this trajectory toward more local connections, which we call the *regionalization* of the Internet, from what has been called the *Balkanization* of the Internet, which implies deliberate *disconnection* of regions.³⁹ Some countries are exploring the extent to which they can isolate their region from the global Internet. Such deliberate isolation is a different phenomenon from what we describe, which is the continuing enrichment of platform assets on which application developers depend, in a

³⁹The term “balkanization” as applied to the Internet may have first been used in a 1997 paper by Van Alstyne and Brynjolfsson (Marshall Van Alstyne and Erik Brynjolfsson. *Electronic Communities: Global Village or Cyberbalkans?* Mar. 1997. URL: <http://web.mit.edu/marshall/www/papers/CyberBalkans.pdf>). The term has been used since by many authors, usually to describe an undesirable outcome where the Internet splinters into disconnected regions.

virtuous cycle with creative exploitation of platform assets by application designers to provide a better user experience.

Of course, this movement toward the cloud could slow if the cloud ecosystem becomes more problematic for the application developer. The research community should continue to track issues in the larger Internet ecosystem that might discourage application designers from locating there, such as resilience, security, or business issues. Issues such as “cloud neutrality”,⁴⁰ could arise and push application designers in different directions.

6.3 Leveraging regionalization to prevent abuse of the address space

The IP address space is administered regionally, with five Regional Internet Registries responsible for address allocation in different parts of the globe. But the address space is used globally, and forcing a geographic structure on its use is an unnecessary and in our view harmful constraint. Global connectivity of the address space is a core value of the Internet’s design.

We think about regionalization in this context in the following way. A commitment by individual ISPs to implement SAV does not create a connected region. The requirement for security through regionalization derives from controlling the *risk of harm* that arises from a lack of SAV, which is increased ability for an attacker to carry out DDoS attacks.

At the packet forwarding layer, the only obvious countermeasures to DDoS attacks are to block or dissipate the traffic. Regionalization may help this approach, although the tradeoffs merit consideration. First, while many Internet services are replicated in many regions, these services typically still have globally reachable addresses, and thus globally attackable addresses. But regionalization removes the need for global reachability of these replicas. An application provider may want some globally reachable service points for resilience, but restrict other instantiations of the service to one region.

The first issue with this approach is that the back end control element of the application needs to reach (to manage) the distributed service points. But cloud-hosted service points could have two interfaces: one connecting to the public Internet but not globally reachable and one in the private cloud network, protected from attack. This pattern is used by some applications today, and might become more common in the future.

Second, many critical Internet services today use *anycast* addressing, a technique that assigns the same address to many different distributed destinations. The Internet routing protocols then automatically take traffic to the closest instance, in terms of the routing path computation, of that address. With anycast-addressed services, DDoS attacks from bots in different parts of the world can only reach the nearest instance of the server, thus dissipating the attack.

Third, if regionalization is empirically true then it implies that links connecting regions will be less important to most activities, thus operators could throttle (but not disable) them during an attack to keep bots outside the region from overwhelming services inside the region. This approach would degrade global connectivity of the region to preserve stable operation internal to the region, although presumably the DDoS attack itself is already degrading connectivity on those links. Operators might be able to throttle/block only those addresses under attack, as with many DDoS scrubbing services today.

6.4 BGP

We proposed an extension to MANRS that we called *recursive MANRS* (§5.2.1): an operational practice that would require a MANRS-compliant ISP to treat BGP announcements from its customers differently depending on whether its customers were themselves MANRS-compliant. For this practice to limit the propagation of malicious BGP announcements with invalid *path* announcements, as opposed to invalid *origin* announcements, it is necessary that the MANRS-compliant ASes be directly connected into a contiguous region. If MANRS ASes are directly connected, recursive application of this rule means that an attacker’s false path announcement will not succeed *within the topological region circumscribed by that set of ASes*—a zone of trust for secure routing. Today, some MANRS members form an interconnected region, but other members are isolated from that region, because they connect to the Internet using transit providers that do not commit to being MANRS-compliant. An important role for data analysis is to map out the topology of the MANRS members and identify prospective participants that would improve the connectivity of individual MANRS members to a directly connected cluster that represents a zone of trust.

⁴⁰Molly Wood. *We Need to Talk About ‘Cloud Neutrality’*. Feb. 2020. URL: <https://www.wired.com/story/we-need-to-talk-about-cloud-neutrality/>

The emergence of a coherent region of directly connected MANRS ASes creates a stronger industry incentive for additional ASes to join MANRS. Customers are better protected from being misled by false BGP announcements if they connect to a MANRS-compliant transit provider, and a customer that is concerned that others will not forward its route announcements needs to connect to a MANRS-compliant transit provider.

6.5 DNS

Improving the security of the DNS through the approach of regionalization is more complex than in the case of BGP, where the actors that commit to a code of conduct (such as MANRS or an enhanced MANRS) have an explicit topological relationship to each other. The DNS by design does not map onto the topology of the Internet. A domain name registered in any TLD can name a service hosted in any part of the Internet, and in principle a user in any part of the Internet might look up a name registered in any TLD. A zone-of-trust approach must find a creative way to regionalize this behavior.

Imagine that a suitable group of experts, assembled so as not to include the bad guys in the group, defines a code of conduct for registries and registrars that reduces the incidence of registrations for malicious purposes, and improves the ability of law enforcement to identify the registrant. How can that first step lead to a zone of trust?

One obvious but possibly over-aggressive answer is that for users that choose to be within a zone of trust with respect to the DNS, resolvers do not resolve names that are registered in registries/registrars that do not conform to the code of conduct. That is, those resolvers issue some error response to a DNS query about those names. This approach assumes the risk that legitimate services may become unavailable, at least if a zone of trust suddenly deployed such a mechanism. However, a well-orchestrated transition would notify providers of legitimate services that they need to register their service names inside a compliant name service. A service could have more than one domain name, and a service that desired a global reach might register several names, each valid inside a given zone of trust.

Such a scheme would not eliminate all DNS-based malicious activity on the Internet. It would motivate malicious actors to try to get within the zone of trust, i.e., to register names with providers that comply with the code of conduct. Thus, the code of conduct must include elements that make it harder for these actors to register names for malicious purposes, and easier to find out who they are.

Users, or their tools, could also install exceptions to local blocking by recursive resolvers of DNS queries.

Also, registries that agree to a code of conduct must be able to refuse registrations from registrars that do not comply with it,⁴¹ or else the zone of trust must require both registry and registrar information to assess whether a name is within the zone of trust.

6.6 The Certificate Authority (CA) system

We noted in §5.5 that browser developers have been willing to remove many CAs from their list of trusted actors. Such removal can cause collateral harm in the form of inconvenience (or outright blockage) to users trying legitimately to get to web sites with certificates issued by those CAs. However, many CAs seem to serve regional markets, and regional decisions about whether to trust a CA may balance the benefits and harms based on observed behavior of users in different regions.

6.7 The role of the application in creating and exploiting a zone of trust

Application behavior, e.g., high volume streaming, has moved the Internet toward its regionalized character, and in so doing provided a way to think about security at a regionalized level. But in a zone of trust, applications must take steps to remain within that zone, or take special action if they must go outside that zone. A study of design practices for modern applications is an important part of this approach.

Email is a quintessential example of a global application that is also a primary vector for malicious behavior. It may be the most challenging application to shape so that it has a regional character. We must consider all the security vulnerabilities that arise inside email and see how we could leverage our *zone-of-trust* concept to mitigate them.

⁴¹Under current ICANN rules, registries may not discriminate in this fashion.

7 Summary thoughts and conclusions

For each set of security risks we identified, we described an approach to establishing and maintaining a zone of trust to mitigate that set of risks within its scope.

- With respect to BGP, the zone of trust might be that set of interconnected Autonomous Systems that commit to the practices defined by MANRS: to verify that their customers are announcing valid blocks of addresses; to flag unverified announcements that come from outside the zone; and to reject announcements from outside the zone if they conflict with announcements from within the zone. Another zone might commit to the enhanced practices we called recursive MANRS. An AS utilizes that zone of trust by registering its ROAs and connecting to the Internet through a transit provider that is MANRS-compliant. An AS connected in this way will have a high level of assurance that the route to any other AS connected in this way will not be hijacked. In turn, applications that host their service points inside those ASes are protected from hijacking.
- With respect to the DNS, the zone of trust is defined by the set of registrars and registries that agree to a code of conduct that makes those domains inhospitable to malicious users. Names in those domains are much less likely to be dangerous, avoiding (or cautiously treating) resolution of a name outside that zone of trust will diminish exposure to risk. Alternatively, a trust zone might be defined by a set of operators of recursive DNS resolvers that commit to block access to domains or URLs based on a determination that they host abusers, and to hold registries and registrars to a high level of operational security performance.
- With respect to the CA system, the zone of trust is the set of CAs that are judged trustworthy.

Our long-term goal is to foster the emergence of zones of trust within the Internet. With proper framing and shaping of incentives, these zones may emerge bottom-up in the existing ecosystem. Alternatively, governments may move to shape the regions of the Internet under their control. Individual governments, or even groups of governments, cannot impose global solutions. The ability to create regions of higher trust across national boundaries is central to any approach to governmental regulation or intervention to improve Internet infrastructure security.

We have some understanding of the requirements for a zone of trust. A sustainable zone of trust requires five elements:

- Clear rules about acceptable behavior,
- A commitment to measurement to detect rule violation,
- A commitment to deal with rule violation,
- Constraints on the ability of bad actors outside the zone to disrupt its operation,
- Applications that limit their dependencies to the extent possible to the zone in which the application operates

The concept of a zone of trust must be general. Different threats will call for zones of different shape and dimension. For one threat, the zone might be jurisdictional, for another the zone might be a connected set of ASes. So long as an activity operates within a zone of trust defined for each threat, the zone will provide enhanced security.

A key component is measurement to provide critical knowledge about topology and connectivity, the basis for validating commitments, and the state of deployment. Pursuit of this approach should include an international advisory team that includes policy makers, operators, and researchers to advise on the role of measurement and analysis in developing these operational procedures.

References

Akiwate, G. et al. “Unresolved Issues: Prevalence, Persistence, and Perils of Lame Delegations”. In: *ACM Internet Measurement Conference (IMC)*. Oct. 2020.

- Alstyne, Marshall Van and Erik Brynjolfsson. *Electronic Communities: Global Village or Cyberbalkans?* Mar. 1997. URL: <http://web.mit.edu/marshall/www/papers/CyberBalkans.pdf>.
- Aximov, Alexander et al. *Verification of AS_PATH Using the Resource Certificate Public Key Infrastructure and Autonomous System Provider Authorization*. Nov. 2020.
- Baker, F. and P. Savola. *Ingress Filtering for Multihomed Networks*. IETF BCP84. Mar. 2004.
- Brian Cute. *A conversation about evolving the effectiveness of our multistakeholder model*. <https://www.icann.org/en/system/files/files/draft-evolving-multistakeholder-model-issues-list-25apr19-en.pdf>. Mar. 2019.
- US-CERT. *Multiple DNS implementations vulnerable to cache poisoning VU#800113*. 2008.
- Dave Piscitello and Dr. Colin Strutt. *Criminal Abuse of Domain Names: Bulk Registration and Contact Information Access*. <http://interisle.net/sub/CriminalDomainAbuse.pdf>. Oct. 2019.
- Deering, Steve and Robert Hinden. *Internet Protocol, Version 6 (IPv6) Specification*. RFC 1883. Dec. 1995.
- Eddy, W. *TCP SYN Flooding Attacks and Common Mitigations*. RFC 4987. Aug. 2007.
- Ferguson, P. and D. Senie. *Network Ingress Filtering: Defeating Denial of Service Attacks which employ IP Source Address Spoofing*. IETF BCP38. May 2000.
- Greg Aaron. *Domain Name Registration Data at the Crossroads: The State of Data Protection, Compliance, and Contactability at ICANN*. <http://interisle.net/sub/DomainRegistrationData.pdf>. Mar. 2020.
- ICANN Security and Stability Advisory Committee (SSAC). *SSAC Advisory on Registrant Protection: Best Practices for Preserving Security and Stability in the Credential Management Lifecycle*. <https://www.icann.org/en/system/files/files/sac-074-en.pdf>. Nov. 2015.
- *SSAC Response to the new gTLD Subsequent Procedures Policy Development Process Working Group Initial Report*. <https://www.icann.org/en/system/files/files/sac-074-en.pdf>. Oct. 2018.
- Internet Corporation for Assigned Names and Numbers. *Board Action on Competition, Consumer Trust, and Consumer Choice Review*. <https://www.icann.org/news/blog/board-action-on-competition-consumer-trust-and-consumer-choice-review>. Mar. 2019.
- *ICANN Articles of Incorporation*. <https://www.icann.org/resources/pages/governance/articles-en>.
- *ICANN Bylaws*. <https://www.icann.org/resources/pages/governance/bylaws-en/>.
- Internet Society. *Mutually Agreed Norms for Routing Security*. <https://www.manrs.org/>.
- Kottler, Sam. *GitHub Engineering DDoS Incident Report*. <https://githubengineering.com/ddos-incident-report/>. Mar. 2018.
- Lagerfeldt, Carl and Johan Gustawsson. *Routing Security: RPKI Update Q2/20*. 2020.
- Luckie, M. et al. “Network Hygiene, Incentives, and Regulation: Deployment of Source Address Validation in the Internet”. In: *ACM Computer and Communications Security (CCS)*. Nov. 2019.
- Luckie, Matthew et al. “Network Hygiene, Incentives, and Regulation: Deployment of Source Address Validation in the Internet”. In: *CCS '19*. London, United Kingdom: Association for Computing Machinery, 2019, pp. 465–480. ISBN: 9781450367479. DOI: 10.1145/3319535.3354232. URL: <https://doi.org/10.1145/3319535.3354232>.
- Lyngaas, Sean. *Someone is spoofing big bank IP addresses – possibly to embarrass security vendors*. <https://www.cyberscoop.com/spoofed-bank-ip-address-greynoise-andrew-morris-bank-of-america/>. Apr. 2019.
- Matthew Luckie, Ken Keys, Ryan Koga, Rob Beverly, kc claffy. *Spoofed Source Address Validation Measurement System*. <http://spoofer.caida.org>. 2016.
- Molly Wood. *We Need to Talk About ‘Cloud Neutrality’*. Feb. 2020. URL: <https://www.wired.com/story/we-need-to-talk-about-cloud-neutrality/>.
- Morrow, Christopher. *BLS FastAccess internal tech needed*. <http://www.merit.edu/mail.archives/nanog/2006-01/msg00220.html>. 2006.
- National Research Council. *Realizing the Information Future: The Internet and Beyond*. Washington, DC: The National Academies Press, 1994. ISBN: 978-0-309-05044-9. DOI: 10.17226/4755. URL: <https://www.nap.edu/catalog/4755/realizing-the-information-future-the-internet-and-beyond>.
- NIST. *RPKI Monitor*. URL: <https://rpki-monitor.antd.nist.gov/>.
- Paul Vixie. *Taking Back the DNS*. <http://www.isc.org/community/blog/201007/taking-back-dns-0/>. July 2010.
- Postel, Jon. *Internet Protocol*. RFC 791. Sept. 1981.

- Postel, Editor, Jon. *Transmission Control Protocol — DARPA Internet Program Protocol Specification*. RFC 793. Information Sciences Institute, University of Southern California, 4676 Admiralty Way, Marina del Rey, CA 90291, Sept. 1981.
- Rekhter, Yakov, Tony Lee, and Susan Hares. *A Border Gateway Protocol 4 (BGP-4)*. RFC 4271. Apr. 1996. URL: <https://tools.ietf.org/html/rfc4271>.
- Rijswijk-Deij, Roland van et al. “A High-Performance, Scalable Infrastructure for Large-Scale Active DNS Measurements”. In: *IEEE Journal on Selected Areas in Communications* 34.7 (2016). ISSN: 0733-8716. DOI: 10.1109/JSAC.2016.2558918. URL: <http://ieeexplore.ieee.org/lpdocs/epic03/wrapper.htm?arnumber=7460220>.
- Rosen, Eric. *Exterior Gateway Protocol (EGP)*. 1982. URL: <https://tools.ietf.org/html/rfc827>.
- Serrano, Nicolas, Hilda Hadan, and L. Jean Camp. “A Complete Study of P.K.I. (PKI’s Known Incidents)”. en. In: *SSRN Electronic Journal* (2019). ISSN: 1556-5068. DOI: 10.2139/ssrn.3425554. URL: <https://www.ssrn.com/abstract=3425554> (visited on 11/10/2020).
- Testart, C. et al. “Profiling BGP Serial Hijackers: Capturing Persistent Misbehavior in the Global Routing Table”. In: *ACM Internet Measurement Conference (IMC)*. Oct. 2019.
- “To Filter or not to Filter: Measuring the Benefits of Registering in the RPKI Today”. In: *Passive and Active Measurement Conference (PAM)*. Jan. 2020.
- Testart, Cecilia. “Reviewing a Historical Internet Vulnerability: Why Isn’t BGP More Secure and What Can We Do About it?” In: *TPRC 46: The 46th Research Conference on Communication, Information and Internet Policy*. Sept. 2018. URL: https://papers.ssrn.com/sol3/papers.cfm?abstract_id=3141666.
- Vixie, Paul and Vernon Schryver. *DNS Response Policy Zones (RPZ)*. Dec. 2016. URL: <https://tools.ietf.org/html/draft-vixie-dns-rpz-04>.