



Penn Law
UNIVERSITY *of* PENNSYLVANIA LAW SCHOOL

Public Law and Legal Theory Research Paper Series
Research Paper No. 19-02

Lowering Legal Barriers to RPKI Adoption

Christopher S. Yoo
UNIVERSITY OF PENNSYLVANIA

David A. Wishnick
UNIVERSITY OF PENNSYLVANIA

This paper can be downloaded without charge from the Social Science Research Network
Electronic Paper collection: <https://ssrn.com/abstract=3308619>.



Center for Technology, Innovation and Competition



<http://www.law.upenn.edu/academics/institutes/ctic/>

University of Pennsylvania Law School
3501 Sansom Street
Philadelphia, PA 19104-6204

Lowering Legal Barriers to RPKI Adoption

December 31, 2018

Christopher S. Yoo
David A. Wishnick*

Center for Technology, Innovation and Competition
University of Pennsylvania

EXECUTIVE SUMMARY

Across the Internet, mistaken and malicious routing announcements impose significant costs on users and network operators. To make routing announcements more reliable and secure, Internet coordination bodies have encouraged network operators to adopt the Resource Public Key Infrastructure (“RPKI”) framework. Despite this encouragement, RPKI’s adoption rates are low, especially in North America.

RPKI is a two-sided framework, and on both sides of the framework North American networks lag behind their peers around the globe. For RPKI to provide increased security, networks must first issue route origin authorizations (“ROAs”) to cover their Internet Protocol (“IP”) address space. Other networks must also conduct route origin validation (“ROV”) on the basis of these ROAs to ensure that routing announcements originate from authorized parties.

Recent analysis has suggested that North American networks have issued proportionately fewer ROAs than networks around the globe. Further, when conducting ROV, networks worldwide are less likely to validate routes against North American ROAs than against ROAs issued by networks in other regions of the world. In other words, a material portion of networks conducting ROV do not validate routes originating from North American networks. Taken together, these two analyses paint a problematic picture. North American networks are less likely than their global

* This research is supported by the U.S. National Science Foundation (“NSF”) Award No. 1748362. See *EAGER: Legal Barriers to Securing the Routing Architecture*, NSF (Sept. 1, 2017), https://www.nsf.gov/awardsearch/showAward?AWD_ID=1748362&HistoricalAwards=false (Award Abstract #1748362). The contents of this report, however, reflect the independent views of the authors. They do not in any way represent the views of the NSF.

For valuable discussions and suggestions, the authors wish to thank kc claffy, John Curran, Andrew Gallo, Yossi Gilad, Greg Hankins, Olaf Kolkman, Aris Lambrianidis, Martin Levy, Doug Montgomery, Sandra Murphy, Karl Newell, Anita Nikolich, John O’Brien, Andrei Robachevsky, Steve Ryan, Michael Sinatra, Job Snijders, Rüdiger Volk, Matthias Wählisch, and many others who preferred to remain unnamed. Any mistakes are, of course, the authors’ own.

peers to publish ROAs, and when they do, their ROAs are less likely to be utilized in determining routing tables.

This report presents the results of a year-long investigation into the hypothesis—widespread within the network operator community—that legal issues pose barriers to RPKI adoption and are one cause of the disparities between North America and other regions of the world. On the basis of interviews and analysis of the legal framework governing RPKI, the report evaluates the issues raised by community members and proposes a number of strategies to reduce or circumvent the barriers that are material. The report also describes substantial action taken this year by the American Registry for Internet Numbers (“ARIN”) and other private organizations in light of public dialogue about RPKI.

RPKI presents a classic “chicken-and-egg” problem: While adoption may be net beneficial to the Internet community as a whole, its attractiveness to any individual network operator depends on expectations about other operators’ willingness to adopt. The legal issues analyzed in this report are sources of friction. They deter individual actors from adopting RPKI on the merits, which in turn leads to reduced expectations about the prospects for mass adoption in the future. Reduction of these sources of friction would spur RPKI adoption through two mechanisms. First, reduced legal barriers would make RPKI more attractive to network operators considering whether to implement RPKI. Second, a marginal shift may alter other network participants’ expectations about the future of RPKI. When everyone expects more adoption, more adoption will take place. Any reduction in the legal barriers to adoption thus may contribute to a positive feedback loop that can promote wider deployment of RPKI.

Our recommendations mainly focus on the two sides of the RPKI framework: issuing ROAs for inclusion in RPKI repositories and conducting ROV on the basis of RPKI repository information. Though ROAs are logically prior to ROV in the production sequence, the legal issues surrounding ROV are more significant than those surrounding the issuance of ROAs. Therefore, we address them first. On the ROV side of the equation, the principal legal obstacles stem from the terms and conditions governing access to the RPKI repository offered by ARIN in its Relying Party Agreement (“RPA”) and the manner it employs to ensure the agreement is binding.

Over the course of the past year, ARIN—a private, nonprofit organization that serves as the Regional Internet Registry (“RIR”) for the United States, Canada, and part of the Caribbean—has made a salutary change to its RPA in response to community dialogue. Specifically, ARIN enabled third-party software developers to incorporate acceptance of the RPA into their software workflow. If developers capitalize on this change, network operators will be able to rely on third-party software to access the ARIN RPKI repository more easily. This change built on ARIN’s decision in 2016 to move from a cumbersome email-based method of RPA acceptance to a browser interface-based method. ARIN has also been an active participant in dialogue about further potential changes to support RPKI adoption.

This report is meant to spur further dialogue by clarifying key topics of debate. It recommends the following:

1. The goal of widespread ROV counsels in favor of ARIN reviewing its current approach to repository distribution, embodied in the RPA. We conclude that two paths would be reasonable. First, ARIN should consider dropping the RPA altogether. This would remove the most significant legal barriers to widespread utilization of the ARIN RPKI repository. Second, because the legal risks faced by ARIN in an RPA-free world are ultimately uncertain, it would also be reasonable for ARIN to maintain the RPA for the purposes of contractually allocating risks to the parties best positioned to reduce and mitigate them. If ARIN keeps the RPA, ARIN should consider removing the RPA's indemnification clause, instead of relying solely on the RPA's disclaimers of warranties and limitations of liability, or at least reducing the indemnification clause's scope to eliminate the problem of moral hazard.
2. Developers of RPKI validation software should consider integrating acceptance of ARIN's RPA into their software workflows. ARIN recently enabled this possibility, and developers should deliberate on whether to capitalize on the opportunity.
3. The network operator community and ARIN should more broadly publicize ARIN's policy of revising various RPA clauses for government entities that are prohibited from agreeing to them.
4. In addition to the important step ARIN has already taken to enable third-party software developers to integrate RPA acceptance into their software workflows, ARIN should consider reducing the barriers to third-party service development imposed by the RPA's prohibited conduct clause. Specifically, ARIN should consider methods for allowing approved developers to make use of RPKI information as an input into more sophisticated services.
5. Separately, ARIN should consider revising the prohibited conduct clause to allow broader distribution of information created with RPKI as an input for research and analysis purposes.
6. As a general alternative, the Internet community should consider whether to develop a separate corporate entity that would be responsible for operational aspects of RPKI repository provision. That corporation could conduct such activities for the North American region, or on a worldwide basis.

Regarding the ROA-issuance side of the equation, the principal legal obstacles stem from the terms and conditions found in ARIN's Registration Services Agreement ("RSA"), Legacy Registration Services Agreement ("LRSA"), and RPKI Terms of Service. Regarding these, the report recommends the following:

1. ARIN should consider adopting a pathway to provide RPKI services that would explicitly refrain from altering the existing balance of property and transferability rights associated with legacy IP address allocations.

2. The network operator community and ARIN should broadly publicize ARIN's policy of revising certain RSA/LRSA and RPKI Terms of Service clauses for government entities that are prohibited from agreeing to them, including indemnification, arbitration, and choice of law clauses. ARIN should also begin presenting the RPKI Terms of Service to newly-onboarded members alongside their RSA/LRSA, so that organizations spend less time dealing with legal issues overall.

Separately, the report recommends that the network operator community consider whether to encourage companies and the federal government to include RPKI adoption in procurement best practices or requirements.

In tandem with recommendations designed to encourage adoption, the report also makes two recommendations concerning operational readiness for widespread RPKI deployment. Specifically:

1. To reduce any legal risks associated with RPKI, the network operator community should focus on adopting operational best practices. No system is 100% reliable across all contingencies; as a result, operators should prepare for outages and other headaches. RPKI implementations should be resilient in the face of such contingencies.
2. The five RIRs should work to ensure readiness for widespread RPKI adoption and strive to publicize deeper details on their service-level intentions to the Internet community.

All of these recommendations are meant to be consistent with the goals espoused in the IETF Requests for Comments ("RFCs") that set standards for RPKI¹ and with ARIN's Articles of Incorporation.²

1. INTRODUCTION

The networks constituting the global Internet employ the Border Gateway Protocol ("BGP") to publish routing announcements. These announcements advertise potential pathways across which data can travel from one endpoint to another. Individual networks rely on them to build forwarding tables, which determine data paths. The contents of those tables are important. If a table contains a route derived from an erroneous or fraudulent announcement, then data might be sent in a direction that prevents them from reaching their true destination or sent through a network controlled by a malicious actor.

¹ See, e.g., Matt Lepinski & Stephen Kent, *An Infrastructure to Support Secure Internet Routing*, IETF RFC 6480 (rel. Feb. 2012), <https://tools.ietf.org/html/rfc6480>; Randy Bush, *Origin Validation Operation Based on the Resource Public Key Infrastructure (RPKI)*, IETF RFC 7115 (rel. Jan. 2014), <https://tools.ietf.org/html/rfc7115> [hereinafter "RFC 7115"].

² See *Articles of Incorporation*, ARIN (Aug. 7, 1997), https://www.arin.net/about_us/corp_docs/artic_incorp.html.

Despite these security threats, network operators do not typically authenticate route announcements. BGP does not contain security features to ensure routing accuracy. Rather, BGP operates on a “transitive trust” model, where networks often assume that the routes advertised by their neighboring networks are, in fact, viable. This leaves BGP “surprisingly vulnerable to attack.”³ In April 2018, malicious actors successfully executed a routing hijack to redirect traffic meant for Amazon Route 53, Amazon’s authoritative Domain Name System (“DNS”) service. This attack facilitated the theft of approximately \$150,000 in cryptocurrency.⁴ Months later, hijackers attacked major credit card processors based in the United States.⁵ The public record of similar attacks and routing mistakes is growing,⁶ and it is possible that other incidents have taken place but escaped public notice.

One partial solution to the problem of BGP security is the Resource Public Key Infrastructure (“RPKI”). RPKI complements BGP’s “transitive trust” system with an additional layer of security generated by an “anchored trust” system. RPKI’s system is based on public-key cryptography. Under RPKI, regional Internet registries (“RIRs”)—the organizations responsible for allocating and managing Internet Protocol (“IP”) addresses and Autonomous System (“AS”) numbers to the Internet’s many participating networks—serve as trust anchors for an authentication system. They do so by allocating private cryptographic keys to the holders of IP addresses. These keys allow their holders to publish secure digital objects called “Route Origin Authorizations” (“ROAs”), which establish which networks are authorized to originate routes associated with particular IP addresses. The existence of these ROAs enables other parties to validate the authenticity of route announcements: One can validate a route announcement by comparing its point of origin with the ROAs contained in the RPKI repository maintained by the RIR that issued the address prefix. This process is known as Route Origin Validation (“ROV”).⁷ Networks can then adopt various practices to filter routes appropriately based on ROV information.

RPKI is only a partial solution to the problem of BGP security because it does not account for the entire routing path. But its value should not be discounted merely because it is not a panacea. As Internet topology increasingly shifts toward a world

³ Sharon Goldberg, *Why Is It Taking So Long to Secure Internet Routing?*, ACM QUEUE (Sept. 11, 2014), <http://queue.acm.org/detail.cfm?id=2668966>.

⁴ See Doug Madory, *BGP Hijack of Amazon DNS to Steal Cryptocurrency*, ORACLE DYN VANTAGEPOINT (Apr. 25, 2018), <https://dyn.com/blog/bgp-hijack-of-amazon-dns-to-steal-crypto-currency/>. This incident illustrates that routing attacks can be used to affect critical Internet services like authoritative DNS resolution, which, given migration to the cloud, can affect what used to be internal systems.

⁵ See Doug Madory, *BGP/DNS Hijacks Target Payment Systems*, ORACLE DYN VANTAGEPOINT (Aug. 3, 2018), <https://dyn.com/blog/bgp-dns-hijacks-target-payment-systems/>.

⁶ See SANDRA MURPHY, ROUTING SECURITY AND RPKI at 3-4 (Nov. 17, 2015), *available at* <https://www.nanog.org/sites/default/files/04-Murphy-StLouis.pdf>.

⁷ Validators typically cache ROAs instead of looking them up in the RPKI repository every time.

where there are fewer hops between origin and endpoint,⁸ a routing announcement's origin represents a comparatively significant portion of its full path. In the limit, where there is only one hop between origin and endpoint, origin validation *is* path validation. Thus, the value of RPKI is even higher among parties that utilize short paths to reach each other. Among near neighbors in network topology, RPKI's origin validation framework is a particularly important contribution to routing security.

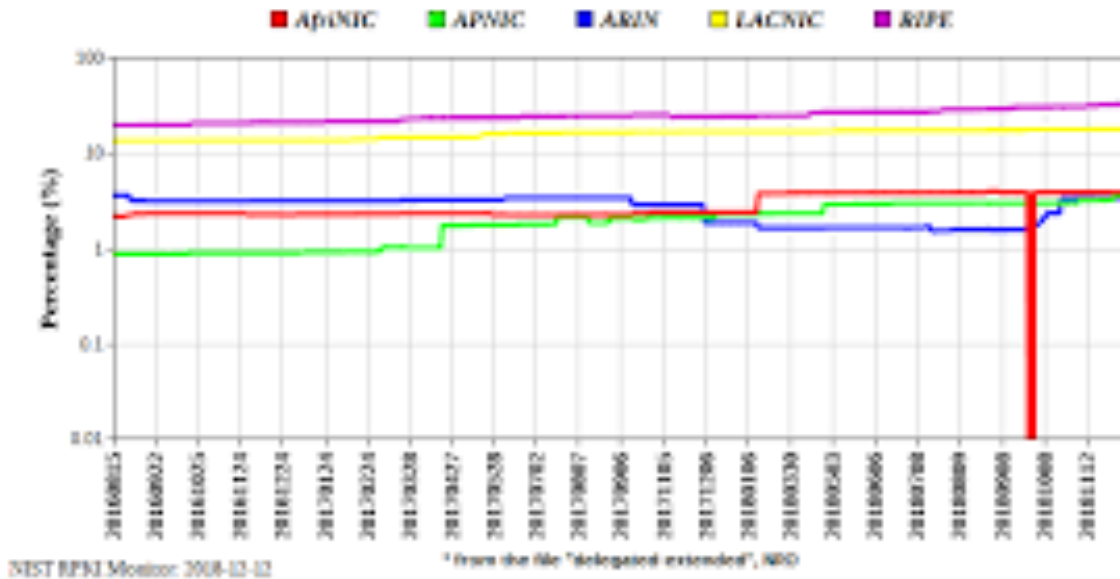
RPKI is a two-sided framework. For it to be successful, network operators must use private keys to sign and publish certificates authenticating the origins of routes for IP addresses under their control. Then they must also filter routes in their routing tables based on the certificates provided by other networks, while following best practices to ensure such filtering is done safely and wisely. Roughly speaking, the value of adopting RPKI increases with the number of other networks that have also adopted RPKI—a classic network effect. Moreover, the value on each side of the framework increases with the number of participants on the other side, which makes the network effect two-sided. The more ROAs there are (assuming accuracy), the higher the value of engaging in ROV. In turn, the more participants that engage in ROV, the higher the value of issuing ROAs.

Despite the fact that leading coordination bodies such as the Internet Engineering Task Force (“IETF”) Working Group on Secure Inter-Domain Routing (“SIDR”), RIRs, and the U.S. National Institute of Standards and Technology (“NIST”) have long promoted RPKI adoption,⁹ adoption rates remain low globally. North American adoption levels (the subject of this paper) remain below the adoption rates seen in Europe and Latin America (as depicted in Figure 1). On the signing side, data collected by U.S. National Institute of Standards and Technology (“NIST”) indicate that the percentage of IPv4 address space covered by ROAs in the ARIN region has lagged behind the levels achieved in the areas governed by Réseaux IP Européens (“RIPE”) (covering Europe and the portions of Asia within the Middle East and Russia) and the Latin American and Caribbean Information Centre (“LACNIC”) (covering Latin America and some of the Caribbean). Somewhat curiously, ROA coverage in the ARIN region declined from September 2016 to September 2018. The recovery in the last quarter of 2018 only succeeded in bringing ARIN ROA coverage back to near the 2016 levels.

⁸ See Christopher S. Yoo, *Paul Baran, Network Theory, and the Past, Present, and Future of the Internet*, 17 COLO. TECH. L.J. (forthcoming 2019).

⁹ See WILLIAM HAAG, DOUG MONTGOMERY, WILLIAM C. BARKER, & ALLEN TAN, NIST SPECIAL PUBLICATION 1800-14: PROTECTING THE INTEGRITY OF INTERNET ROUTING: BORDER GATEWAY PROTOCOL (BGP) ROUTE ORIGIN VALIDATION (2018), *available at* <https://www.nccoe.nist.gov/sites/default/files/library/sp1800/sidr-piir-nist-sp1800-14-draft.pdf> (draft); *Secure Interdomain Routing (sidr)*, IETF, <https://datatracker.ietf.org/wg/sidr/about/> (last visited Dec. 27, 2018); *Resource Certification (RPKI)*, RIPE NCC, <https://www.ripe.net/manage-ips-and-asns/resource-management/certification> (last visited Dec. 27, 2018); KOTIKALAPUDI SRIRAM & DOUGLAS MONTGOMERY, DRAFT NIST SPECIAL PUBLICATION 800-189: SECURE INTERDOMAIN TRAFFIC EXCHANGE: BGP ROBUSTNESS AND DDOS MITIGATION (Dec. 2018) (on file with authors and likely available upon resumption of U.S. government funding to NIST at <https://csrc.nist.gov/publications/detail/sp/800-189/draft>).

Global: RPKI ROA Deployment Status Over Time
 % of *Delegated IPv4 Address Space Covered by ROAs



Source: Nat'l Inst. of Standards and Tech., *RPKI Deployment Monitor*, <https://www.nist.gov/services-resources/software/nist-rpki-deployment-monitor-and-test-system> (last visited Dec. 21, 2018).

On the filtering side, a recent study indicated that the number of networks engaging in ROV filtering in the ARIN region falls far below the number of networks engaging in ROV filtering in the RIPE region.¹⁰ Another study concluded that 20% of the networks engaging in ROV filtering are not including the RPKI repository serving the ARIN region.¹¹

There are some indications that networks are beginning to show greater interest in RPKI adoption. Major network participants including Google and Cloudflare have indicated that they plan to begin implementing RPKI in the near future.¹² Discussions on the mailing lists and at the meetings of the North American

¹⁰ See [Measuring RPKI Route Origin Validation Deployment. ROV DEPLOYMENT MONITOR](https://rov.rpki.net/), <https://rov.rpki.net/> (last visited Dec. 27, 2018).

¹¹ See Ben Cartwright-Cox, *The State of RPKI: Q4 2018*, BEN'S BLOG (Dec. 20, 2018), https://blog.benjojo.co.uk/post/state-of-rpki-in-2018_

¹² See, e.g., Chris Morrow, Network Security Engineer, Google, So I Need to Start Route Filtering Peers, Remarks at the Route Security Track: BGP Route Security, 74th Meeting of the North American Network Operators Group (NANOG 74) (Oct. 2, 2018), available at https://pc.nanog.org/static/published/meetings/NANOG74/1760/20181003_Tzvetanov_Security_Track_Bgp_v1.pdf; Jérôme Fleury & Louis Poinsignon, *RPKI and BGP: our path to securing Internet Routing*, CLOUDFLARE BLOG (Sept. 19, 2018), <https://blog.cloudflare.com/rpki-details/>; Martin J. Levy, *RPKI – The required cryptographic upgrade go BGP routing*, CLOUDFLARE BLOG (Sept. 19, 2018), <https://blog.cloudflare.com/rpki/>.

Network Operators Group (“NANOG”) and ARIN suggest that others are thinking of joining them soon.

Over the past few years, a number of network operators in North America and around the world have suggested that legal issues might be playing some role in holding back adoption. These claims raise concerns with the background law governing liability for using systems like RPKI and specific issues with the agreements that the RIRs have been using to govern the treatment of RPKI resources. But so far, the networking community’s understanding of these legal issues has remained somewhat informal. On the basis of stakeholder interviews and independent analysis, this report aims to clarify and assess those issues, and then to recommend potential solutions.

The rest of the report is organized as follows. Section 2 provides background on where legal issues fit into the “stack” of considerations facing individuals within organizations contemplating adopting RPKI. Section 3 focuses on the route-filtering side of the RPKI equation, assessing the perceived legal barriers and proposing potential solutions designed to help remove or mitigate the effect of those barriers. Section 4 conducts a similar analysis on the ROA-signing side of the equation. Section 5 broadens the lens beyond potential legal barriers to discuss how affirmative strategies, including procurement contracts and increased activity by industry organizations, might be used to promote RPKI adoption. Section 6 discusses the importance of operational best practices among network operators and RIRs. Section 7 concludes.

2. RPKI IN ORGANIZATIONAL CONTEXT

Decisions about whether to adopt RPKI are made by the networks that comprise the global Internet, including large end users, “last mile” Internet service providers, transit providers, backbone providers, and Internet exchange points, among others. Within those organizations, judgments about RPKI’s costs and benefits are most likely to be made by network engineers. As the actors primarily responsible for dealing with the problems RPKI is meant to address, they are the natural advocates for it. Furthermore, these engineers pay attention to the standards-development work that has promoted RPKI. They are often members of professional communities like the North American Network Operators Group (“NANOG”), which offer technical support to would-be adopters.

The RPKI adoption decision involves considerations beyond the technology’s merits. Network engineers face myriad demands on their scarce time and financial resources. As a result, they must weigh the value of adopting RPKI against its costs—including the opportunity costs of foregone effort on other projects. They must also consider whether other departments and functions within their organizations have stakes in the decision to adopt RPKI, which varies by organization based on their particular risk and legal review practices. This is one place where legal questions come into play. Because RPKI raises potential operational risks if not implemented properly, network engineers interested in adopting it must consider whether and how

to engage their colleagues in legal and procurement departments. They must weigh the cost of such engagement both in terms of time and institutional capital. Budgets are not infinite, and RPKI is a technically complex framework to explain. As a result, any issue—even a seemingly small one—can put a weight on the scale against adoption, especially early on in the accumulation of network effects, where the value of adoption to first movers can be low. This is not unique to RPKI, of course, but it can be significant.

3. ROUTE VALIDATION AND FILTERING

Just as RPKI can be viewed as two interrelated processes—ROA issuance and ROV-based filtering—the legal framework can also be divided into two distinct portions. The first we will address is ROV-based route filtering.

To improve the security of routing, the design of the RPKI framework contemplates that networks will filter routes based on RPKI information. That is, networks are encouraged to adopt best practices regarding dropping routes that are not authenticated, while also maintaining reliable fallback configurations to account for the risk of faults or unavailability of the RPKI service itself.

Networks deploying RPKI can follow different adoption paths. Some networks may filter based on their own ROV analysis. Or, as is often the case with special-purpose additions to Internet security efforts, networks may seek to rely on information provided by third parties that offer ROV either as a commercial service or as a free service. Due to the benefits of specialization and scale economies, the latter might enable growth in the value of RPKI information—for instance, if a private company or open-source provider offered a set of route filters based on RPKI information in tandem with other information, such as information obtained from Internet Routing Registries (“IRRs”).

In any case, parties conducting ROV need access to the RPKI repositories of the RIRs. From a legal perspective, this means that key issues include (a) access to the RPKI repositories and (b) redistribution of those repositories and information developed based on them. This Section discusses those issues. It is motivated by the objective of fostering broad distribution and use of RPKI-based route analysis, consistent with reasonable allocation of risks and duties among stakeholders.

3.1 Background

Access to RPKI informational repositories is necessary for anyone who wants to conduct ROV and is thus essential to RPKI’s success. The RPKI repositories represent the “public” portion of the public key infrastructure. Repositories for each region of the globe are provided by the RIRs for those regions. The technical designers of the system envisioned they would be widely distributed around the world and that any party engaging in ROV worldwide would do so for all routes worldwide. The nature of repository distribution means that the providers and users are in a relationship with each other with potential legal implications in cases where the use

of RPKI leads to harm. Providers and users may both want to structure that relationship to allocate rights and responsibilities over the production, distribution, and use of repository information in a sensible manner.

Each RIR has the authority to structure access to its own repository of RPKI information as it sees fit. For North America, the source of this information is the American Registry for Internet Numbers (“ARIN”) RPKI repository, which includes ARIN’s repository of RPKI certificates, certificate revocation lists, signed objects, and ARIN’s public key.¹³ Parties wishing to conduct ROV for route announcements for locations originating in North America need access to the authentic ARIN Repository. Such access is provided via a file called a Trust Anchor Locator (“TAL”).

The key legal document governing access to ARIN’s repository is called the Relying Party Agreement (“RPA”). To access ARIN’s repository, parties must download the TAL file from an ARIN website. This webpage contains a statement that “[t]he ARIN Repository is available to anyone under the terms and conditions in the Relying Party Agreement,” immediately followed by a link to the RPA.¹⁴ The agreement is what lawyers call “browsewrap.”¹⁵ This means that the webpage visitor does not need to affirmatively click on an acceptance box in order to access resources, but the webpage states that use of the resources constitutes an agreement to be bound by the terms contained in the document accessible through a link on the webpage. That statement and the link are prominently displayed in a visitor’s visual field. Though courts are wary of enforcing browsewrap against unwitting parties, they are willing to do so where parties have actual or constructive knowledge of the agreement’s existence. This is especially true if the party is sophisticated. As a result, ARIN’s RPA would likely be held to bind network operators utilizing ARIN’s RPKI repository.

Interviewees suggest that the RPA is a major source of stakeholder concern over the legal framework governing the filtering side of RPKI. These concerns are described and evaluated in the Sections that follow. In Section 3.2, we evaluate the legal impediments to *direct* repository access—the kind of access that a party would seek if they were planning on doing their own ROV. We also report on changes made during the past year that alleviate some of these impediments. In Section 3.3, we propose reforms to overcome remaining impediments. In Section 3.4, we address issues specific to governmental entities seeking direct repository access. In Section 3.5, we address legal impediments to *indirect* repository access—the kind of access a party would receive from an intermediary offering ROV as a commercial service or as

¹³ See *Trust Anchor Locator (TAL)*, ARIN, available at <https://www.arin.net/resources/rpki/tal.html> (last visited Dec. 27, 2018) [hereinafter “ARIN TAL”].

¹⁴ See ARIN, Resource Certification Relying Party Agreement, available at <https://www.arin.net/resources/rpki/rpa.pdf> (last visited Dec. 27, 2018) [hereinafter “ARIN RPA”].

¹⁵ The words “Relying Party Agreement” serve as a hyperlink to a PDF of the RPA itself, and the link is further denoted by an Adobe Acrobat logo appearing next to the words. See ARIN TAL, *supra* note 13.

a free service. We describe changes made during the past year to reduce some of these barriers and evaluate and propose reforms to the remaining ones.

3.2 Legal Barriers to Direct Repository Access

Each RIR should ensure that it distributes its RPKI repository in a manner consistent with the goals of widespread RPKI adoption and a proper allocation of the rights and responsibilities for safe RPKI usage. A number of interviewees with whom we spoke as part of our research suggested that some tension exists between those two goals, especially in a relatively litigious region such as North America. In that region, the organization tasked with navigating the tension is ARIN. This Section introduces ARIN and its Relying Party Agreement.

Background on ARIN. ARIN is a private, member-driven non-profit organization founded in 1997 that serves as the RIR for the United States, Canada, and part of the Caribbean and until 2002 also served as the RIR for Latin America and Africa. It has a budget of \$22 million annually, funded entirely by member registration fees and dues without any governmental support. It began participating in RPKI along with its peer RIRs in 2008. Since initiating its RPKI efforts, ARIN has not engaged in attempts at direct cost recovery for RPKI provision. Instead, it funds its RPKI efforts out of its general, member-funded budget. ARIN estimates that over the last ten years, it has spent approximately \$6 million dollars on RPKI service development. These efforts have included software and web interface development; the procurement, configuration, and operation of a hardware security module; system maintenance; legal analysis and evaluation; and promotional activities.

The role of the RPA. Early in our course of interviews, a number of interviewees noted that, unlike the other four RIRs, ARIN required would-be route origin validators to download its RPKI repository from the ARIN website. In contrast, the other four RIRs allow the TALs for their repositories to be included in software downloads without having to affirmatively accept any specific terms of service. For instance, RIPE Network Coordination Centre (“RIPE NCC”) offers the most popular validator software. When a user downloads the RIPE NCC software, the package comes with four of the five RPKI repositories preloaded, but the page states that “[t]o access ARIN’s [resources], you will have to agree to ARIN’s Relying Party Agreement. Please visit [ARIN’s] web page for more information.”¹⁶ Similarly, a validator provided by Dragon Research Labs similarly includes the TALs for four of the RIRs and omits the ARIN TAL. This difference is driven by heightened concerns about legal risk in the litigious North American region and by the requirement of contract doctrine in the United States that agreements be prominently placed in order to be binding.

¹⁶ See RIPE NCC, *rpki-validator*, GITHUB, at lines 140-55 (Mar. 2, 2017), <https://github.com/RIPE-NCC/rpki-validator/blob/rpki-validator-2.24/rpki-validator-app/README.txt>.

The requirement of agreeing to the RPA to gain access to ARIN's repository raised technical and institutional concerns for interviewees. We discuss each in turn.

Technical concerns. As a technical matter, interviewees reported that the placement of the ARIN TAL separately from the others creates friction that makes ROV setup more onerous. In particular, it inhibits automated distribution of the validator software. It also raises the risk that engineers will forget or refuse to download ARIN's repository resources or simply reject ROV altogether.¹⁷ While, given time and focus, the process of downloading and incorporating an extra TAL into validation software is well within the capacity of the average network engineer, time and focus are inevitably in scarce supply. Especially because mistakes in the configuration of RPKI validation software can have negative consequences, network engineers are loath to implement ROV into production unless they are confident it can be managed effectively over the long term.

Over the course of the year of this study, network engineers had the opportunity to present their concerns over the ARIN RPA's placement in a number of fora. As a result of that dialogue, ARIN took the step of explicitly enabling third parties to develop software installation tools that handle the repository-collection process for their users. Specifically, ARIN now allows third-party software providers to collect acceptances to the RPA as a part of a user interface rather than requiring a separate visit to the ARIN website.¹⁸ This has the potential to be a useful step in reducing the barriers to RPKI adoption because it enables the development of turnkey ROV solutions. To capitalize on it, developers of ROV software should consider designing their software to prompt users to accept the ARIN RPA and conduct ROV for all five global regions. This may require deviating from the practice of enabling automated installation. Interviewees were enthusiastic about third-party software to help solve the chicken-and-egg problem facing RPKI, so exploiting this path may prove important to RPKI's ultimate success. Further discussion on this point can be found in Section 3.5, "Legal Barriers to Indirect Use," below.

One other interface design for ensuring RPA acceptance was proposed by interviewees. The proposal was to embed the RPA into the source code of a software package (like OpenBSD) or directly into the TAL file, itself. Indeed, a recent change to the Internet Engineering Task Force ("IETF") TAL definition seems to contemplate the latter approach.¹⁹ This approach to creating a binding contract is common in open source software distribution. It is attractive because it facilitates automation in the installation process. Unfortunately, we do not view it as a sufficiently reliable method

¹⁷ This risk is borne out in data, as discussed *supra* note **Error! Bookmark not defined.** and accompanying text.

¹⁸ See John Curran, President and CEO of ARIN, *Software installation tools retrieving ARIN TAL (was: Re: ARIN RPKI TAL deployment issues)*, NANOG MAILING LIST (Oct. 13, 2018), <https://mailman.nanog.org/pipermail/nanog/2018-October/097528.html>.

¹⁹ See Geoff Huston, Samuel Weiler, George Michaelson, Stephen Kent, & Tim Bruijnzeels, *Resource Public Key Infrastructure (RPKI) Trust Anchor Locator draft-ietf-sidrps-https-tal-05*, IETF INTERNET DRAFT (Oct. 11, 2018), <https://www.ietf.org/rfcdiff?url2=draft-ietf-sidrps-https-tal-05>.

of establishing a binding contract in the RPKI context. To create a binding browsewrap agreement, a user must have actual or constructive knowledge of the agreement. In the case of open source software, a coder would have such knowledge of the contents of the source code due to the practical necessity of inspecting the source code in the course of software development. But in the RPKI case, there is no reason for a repository user to inspect the TAL file's contents. Indeed, one of the reasons for seeking a way to streamline the RPA-acceptance process is to make it so network engineers do not have to fiddle with TAL files as they begin engaging in ROV. As a result, including the RPA in source code would not reliably ensure that anyone is actually seeing the RPA. Therein lies the problem. The proposed method would not reliably create a binding agreement. While parties that actually know about an agreement buried in source code (for instance, because they saw a presentation at NANOG about it) might be deemed bound by it,²⁰ those ignorant of it would not be bound by it under U.S. common law. Therefore, assuming it is valuable to have an RPA—an assumption we explore in Section 3.3 below—it would need to be presented visually to the parties it is meant to bind.

Institutional concerns. In addition to technical concerns, many interviewees claimed that the RPA was causing institutional friction sufficient to delay or prevent RPKI adoption. Network engineers within some organizations state that they are wary of entering into the RPA out of fear of running afoul of their organizations' procurement rules.

ARIN's efforts to date. In response to both the technical and institutional concerns raised by members, ARIN has devoted significant resources (in terms of employee and board time, along with expenditures on legal counsel) to refine the RPA consent process. As early as 2014, NANOG participants raised concerns over some of the clauses to the agreement and over the fact that ARIN had structured the agreement as "clickwrap"—a term for legal agreements that require affirmative assent via a mouse-click. In response to these concerns, ARIN reviewed its approach to offering the RPKI repository to third parties²¹ and decided in February 2016 to restructure the agreement into the "browsewrap" described above.²² This change obviated the need for end users to make an affirmative mouse-click explicitly accepting the terms and conditions contained in the RPA before accessing ARIN's TAL. According to multiple interviewees, this change resulted in increased willingness among some network engineers to make use of the Repository, as entering

²⁰ See, e.g., *Register.com, Inc. v. Verio, Inc.*, 356 F.3d 393, 402 (2d Cir. 2004) (holding that actual knowledge of a browsewrap agreement sufficed to establish assent).

²¹ See, e.g., John Curran, President and CEO of ARIN, & Steve Ryan, Counsel to ARIN, RPKI Relying Party Agreement (RPA) Change (Jan. 2016), available at https://www.arin.net/vault/about_us/bot/20160111/exhibit_c.pdf.

²² See John Curran, President and CEO of ARIN, *Change re ARIN RPKI Relying Party TAL Access*, NANOG MAILING LIST (Feb. 4, 2016), <https://mailman.nanog.org/pipermail/nanog/2016-February/084042.html>.

into browsewrap agreements falls within their understanding of their authority to act unilaterally within their organizations.

This change did not, however, address the perceived problem for all potential users. Many network engineers are not allowed to enter into the RPA browsewrap agreement unilaterally. That is because their corporate procurement policies prohibit employees from entering into agreements that contain indemnification clauses and other terms not seen in standard licenses without first subjecting those agreements to internal review. These internal review processes require network engineers to invest time in navigating corporate bureaucracy to try out things like RPKI. Internal bureaucracy can be valuable to ensuring that all parts of an organization (for instance, engineering and legal) are on the same page about a new endeavor and any related risks that arise from new dependencies on external services, but they also make new endeavors more time-consuming to undertake.

Some interviewees further stated that the indemnification clause in ARIN's RPA exceeds what their organizations would be willing to accept to participate in ROV. This claim has been hard to vet. On the one hand, some interviewees have stated that their legal departments viewed the indemnification clause as too strict. Further, there may be a large cohort of non-North-American network operators that lack the capacity to assess a legal agreement written in English. On the other hand, the fact that all ARIN members have signed a Registration Services Agreement ("RSA," described in more detail in Section 4.1 below) that contains an indemnification clause that is substantially similar to the RPA's indemnification clause, suggests that indemnification is not a deal-breaker for those organizations, which include many of the most significant networks on the Internet. Rather, it is a weight on the scale in operators' decision-making processes. Our interviews with legal personnel have corroborated the view that indemnification is not typically an automatic deal-breaker, but rather acts as a weight on the scale. In the end, however, every organization will approach the question of indemnification in its own way.

On the basis of all these considerations, we draw three conclusions. First, due to institutional factors, the RPA's indemnification clause poses a nontrivial barrier to participation in ROV. Second, engineers interested in participating in ROV can and should engage their appropriate internal processes to evaluate ARIN's RPA in its current form, and as it may evolve over time. Finally, as we discuss below, ARIN should continue to consider ways to improve RPA acceptance, including by considering changes to the RPA's terms to make them more acceptable to reticent parties.

3.3 Reforming Direct Access

Framing the question of reform. Given the reality that the RPA—and specifically its indemnification clause—alters the institutional calculus for network engineers considering whether to engage in ROV, should ARIN and its members consider removing the RPA, revising the indemnification clause, or trying other potential solutions?

In deciding how best to structure legal limitations on access to the ARIN Repository, it is necessary to frame the proper goal. One sensible goal for the Internet community, as discussed above, is widespread distribution of RPKI repository information. Another goal, of course, is to ensure the ongoing stability and soundness of ARIN—a crucial organization in North America’s Internet governance. The RPA protects ARIN from undue liability, so any proposal to change it or eliminate it should be approached with caution.

Plainly, these two goals are sometimes in tension. Though some RPKI advocates would wish for RPKI information to flow completely freely, the wisdom of that approach cannot be assumed *a priori*. RIRs have important interests—in proper Repository use and appropriate allocation of liability for misuse, for example—that reasonably inform how they offer their TALs to potential users.

By the same token, for an RIR interested in supporting RPKI adoption, complete insulation from potential legal risks is not a feasible goal. Any organization that takes productive action in society cannot completely eliminate the risk of liability or of having to defend against lawsuits—even frivolous ones over entirely legitimate conduct. Instead, the proper objective for an RIR is to balance the risks of incurring legal costs against the benefits of engaging in activities that further the organization’s goals. This means that legal protection is an exercise in optimization and appropriate allocation of risk, not necessarily maximization of legal protections.

The uncertain bounds of potential liability. When it comes to distribution of RPKI repositories, striking the proper balance between potential benefits and risks must take place amid conditions of uncertainty. RPKI is a new service, and we know of no lawsuits dealing with the proper apportionment of the potential sources of liability associated with it. Furthermore, the exact harm scenarios will shift both with increased deployment and as new uses for RPKI information develop. Each RIR and relying party must therefore evaluate its legal risk based on its own best assessment of how RPKI usage might go wrong and where their liability might lie and weigh those risks against the potential benefits of broader RPKI deployment. RIRs and relying parties can gain additional perspective into potential liability from RPKI failure by drawing comparisons to other situations where providers of similar types of trusted information have been subject to legal claims.

Framing the harm scenarios. At its root, an RPKI repository is a body of information. It holds information necessary to conduct ROV, including Resource Certificates, Certificate Revocation Lists, and signed objects (including, most importantly, ROAs). Though RIRs currently publish the leading RPKI repositories, they are not the sole creators of the information contained within them. To the contrary, much of the most important information—specifically, the ROAs pertaining to specific locations—can be produced only by the parties that hold private keys pertaining to specific IP address space.

Providers of information might face claims under a number of different legal theories. In the case of RPKI, scenarios giving rise to a legal claim include incidents that make it impossible or difficult for traffic to reach an Internet endpoint and

incidents that allow traffic to pass into unwanted hands. How might a repository provider like ARIN be implicated in such incidents? Given RPKI's limited track record, it is impossible to be certain, but we conjecture that an aggrieved party might accuse ARIN of failing to issue private keys to IP address holders in a proper manner, facilitating the issuance of faulty ROAs (whether through administrative error or security failure), or improperly revoking private keys or ROAs. Aggrieved parties might include network operators, and they might also include downstream customers whose traffic has been disrupted or misdirected.

One obvious scenario worth considering is how downstream users of RPKI information might react if an RIR's repository were to temporarily become unavailable. This is not fanciful: RIRs' repositories have gone down in the past,²³ and no amount of diligence can completely eliminate the possibility of similar temporary outages in the future. To date, these have had little impact on Internet traffic. That may be because networks that utilize ROV find it easy and sensible to adopt best practices that respond gracefully to outages and similar problems. But it also may be because RPKI is in such early stages of deployment. If some future downstream users are unprepared to handle the occasional outage—a possibility that can never be completely precluded—then misconfigurations could, under certain circumstances, lead to traffic disruptions once RPKI ROV is widely in use.²⁴ In such a situation, an RIR might be accused of contributing to the misfortune, even despite the requirement for relying parties to utilize best practices in their use of RPKI information.

Though the magnitudes and precise harm scenarios vary, every provider of a trusted information service used to direct Internet traffic faces risks of this type of accusation. Their examples can help inform RIR policy. For instance, Internet participants use cryptographic Transport Layer Security ("TLS") certificates to secure application-layer communications. Similarly, parties use the Domain Name System Security Extensions ("DNSSEC") to ensure cryptographic origin validation of data when resolving DNS queries. Finally, many network operators currently use IRR information to support their route filtering decisions. These services thus provide examples that, though imperfectly analogous, are relevant to an RPKI repository provider's risk analysis.

We are not aware of a case in U.S. state or federal court involving allegations of operational harm based on faulty provision of TLS, Secure Socket Layer ("SSL," TLS's predecessor), IRR, or DNSSEC information. That does not mean that providers of key information within these systems are immune from liability (or that they have not settled claims outside of public view), just that no such case has

²³ See, e.g., Mark Kosters, ARIN Chief Technology Officer, *ARIN RPKI Repository (Update)*, ARIN (Oct. 24, 2018), https://www.arin.net/announcements/2018/20181024_update.html (describing an ARIN RPKI service issue); *RIPE NCC RPKI Repository Outage*, RIPE NCC (Feb. 3, 2013), <https://www.ripe.net/support/service-announcements/service-announcements/ripe-ncc-rpki-repository-outage> (describing a RIPE NCC RPKI repository outage).

²⁴ It is worth noting that, at present, even network operators who are already using RPKI today are not perfect at doing so. See, e.g., Nusenu, *Cleaning Up ROAs Inconsistent with the BGP State*, APNIC (Oct. 16, 2018), <https://blog.apnic.net/2018/10/16/cleaning-up-roas-inconsistent-with-the-bgp-state/>.

yet been litigated. Thus, an inquiry into the bounds of a repository provider's liability should also look to analogies to other forms of trusted information. These include maps, financial records and reports, medical information, and the like. Once a (necessarily rough) estimate of liability risk is made, it can then inform an assessment of different paths of action. As discussed below, the inquiry should look at legal techniques to insulate certificate repository providers from liability that do not involve indemnification. These include the use of disclaimers of warranties and liability, express statements that users are assuming various risks, and disclosures that services are being provided on an "as is" basis. The inquiry should also consider whether the net value of dropping the RPA exceeds its potential costs in terms of liability risk.

The unlikelihood of strict products liability. The most serious fear for any provider of information is that they might be held liable for harm under a doctrine called "strict liability." Strict liability refers to the imposition of legal responsibility for harm, even where the liable party acted with reasonable care in a transaction.

This doctrine is unlikely to be applied to providers of RPKI repositories. The primary reason for this is that strict liability applies to harms involving products, not services. For the most part, information suppliers have been deemed to fall outside the sweep of strict products liability because they are providing services.

Even in the unlikely event that an RPKI repository were deemed to be a product, rather than a service, the application of strict products liability typically requires a finding that a given product has been distributed in a defective condition. To be deemed defective, a product must fail to meet reasonable consumer expectations regarding safety or fail a test pitting the utility of investment in safety against the risks of failing to deploy an available safety measure (also known as the risk-utility calculus). Assuming an RIR adhered to IETF standards, it would be unlikely that a court would find its RPKI materials to be defective under either standard.

While courts have applied strict products liability to information providers in rare cases involving the sellers of defective aeronautical charts, such cases are far afield from likely scenarios involving RPKI. In those cases, publishers had made precise misrepresentations that were causally linked to deadly plane crashes. The reliance of users on the precise details of the maps could not have been more serious or direct. In contrast, occasional problems with specific elements of RPKI information are to be expected and are likely to be resolved through various backup mechanisms. Harms are likely to be less serious than a plane crash, and there are many best practices that network operators should be implementing to prevent bad outcomes in the case of RPKI outages or misconfigurations. As a result, the most plausible factual scenarios involving RPKI failures are unlikely to tempt courts or juries into applying strict products liability to the provision of RPKI repositories.

The possibility of negligence liability. While the application of strict products liability is quite unlikely, the application of doctrines concerning negligence is possible. For instance, the *Restatement (Second) of Torts*—a respected, though

nonauthoritative, source on central matters of legal doctrine—identifies “[i]nformation [n]egligently [s]upplied for the [g]uidance of [o]thers” as a core kind of negligent misrepresentation.²⁵ An allegation of this kind of negligence might, under the right circumstances, be asserted against an RPKI repository provider with some degree of plausibility. Similar claims might allege other forms of negligence, such as negligent hiring and supervision of personnel tasked with operating an RIR’s RPKI facilities. To be victorious, a claimant would need to establish that (i) the provider had a duty to act with a certain level of care with regard to a given aspect of RPKI repository-provision; (ii) the provider breached that duty; (iii) the breach actually and proximately caused a resulting injury; and (iv) that injury resulted in cognizable harm.

The risk of negligence liability is mitigated by a number of factors. First, it is highly likely that RIRs will act with the requisite level of care in providing RPKI information. They are competent institutions, and if they adhere to the practices and procedures developed by the Internet community with regard to RPKI, they will likely avoid liability for negligence. Second, the law often prevents parties that acted negligently themselves from recovering for negligent misrepresentation. Any party utilizing RPKI information—for instance in making route-filtering decisions—would themselves be held to a standard of care based on a reasonable network operator’s approach to such activities and would have to follow established best practices in order to prevail. Third, to the extent claims are brought by parties in privity, courts are likely to uphold limitations on consequential damages like the one found in ARIN’s RPA. Fourth, to the extent that claims could be brought by downstream customers of networks utilizing RPKI information, the liability of an RIR would likely be cabined by the doctrine of proximate cause and by doctrines limiting recovery in negligence for pure economic loss.

Nevertheless, negligence liability is still a reasonable concern. The probability and magnitude of such liability is, in the end, uncertain. Despite the fact that network operators should be using best practices to ensure that RPKI outages do not cause downstream harms, it is of course possible that some operators will not follow best practices. In such cases, a harmed party may attempt to sue ARIN alongside their network provider. The questions generated by negligence claims are often hard to resolve, and this means that the question of legal risk is not cut and dried. Questions about particular breaches of duty and the vagaries of causation often generate knotty questions of fact, particularly when it comes to highly technical matters. Furthermore, even defending against claims that ultimately prove unsuccessful can still be costly—especially when questions of fact are involved. Lawyers’ services are not free. As a result, our analysis proceeds on the premise that the costs of negligence suits—including defense against ultimately unsuccessful claims—are something that RIRs should consider. Though our analysis below explores the possibility of eliminating the RPA altogether, we believe that the ultimate uncertainty about potential negligence claims makes it reasonable for ARIN to retain its RPA as well.

²⁵ RESTATEMENT (SECOND) OF TORTS § 552 (1977).

Further, though we recommend ARIN and its members consider rebalancing the current allocation of risk by removing, or at least revising, the RPA's indemnification clause, the ultimate uncertainty about potential negligence claims again suggests that reasonable minds can disagree about how far to go in this regard.

Dropping the RPA as a strategy to promote RPKI. When distributing an RPKI repository, should an RIR require a relying party to enter into an agreement at all? This was a central question raised by many interviewees. Interviewees noted that three of the five RIRs enable access to their RPKI repositories without placing them behind an explicit relying party agreement like ARIN's RPA.²⁶ Similarly, the Internet Assigned Numbers Authority ("IANA") does not require an RPA for its DNS Root Zone Trust Anchors²⁷; relying parties are unlikely to be bound by anything like an RPA in the TLS context²⁸; and ARIN itself does not require an RPA for parties that utilize its IRR information.²⁹ Interviewees suggested that the value of North American ROAs would vastly improve if ARIN opted for a similar agreement-free path in the RPKI context. This is because parties worldwide would have an easier, less legalistic path to conducting ROV on routes covered by those ROAs. In turn, this would increase the value of route-signing in North America.

We think the RPA-free path is worth consideration. This is because an agreement-free distribution would enable the ARIN Repository to circulate more widely than it does at present. To be sure, ARIN would lose legal protection by doing so. But if one believes that the likelihood of serious harms from RPKI misconfigurations, outages, and accidents is very low, then one should be sanguine

²⁶ These RIRs are the African Network Information Centre ("AfriNIC"), the Asia-Pacific Network Information Centre ("APNIC"), and LACNIC. See, e.g., RIPE NCC, *rpki-validator*, *supra* note 16, at lines 140-155. Of these RIRs, APNIC publishes a document that purports to establish a contractual relationship between APNIC and the recipients of RPKI certificates, but that document is unclear as to its application to relying parties. The ultimate legal status and coverage of that document is beyond the scope of this report but has been a topic of community discussion. See Edward Dore, Freethought Internet, *ARIN RPKI TAL Deployment Issues*, NANOG MAILING LIST (Oct. 15, 2018), <https://mailman.nanog.org/pipermail/nanog/2018-October/097543.html>. We are counting RIPE NCC's Certification Repository Terms and Conditions as an explicit RPA-style agreement. See *RIPE NCC Certification Repository Terms and Conditions*, RIPE NCC (Sept. 30, 2016), <https://www.ripe.net/manage-ips-and-asns/resource-management/certification/legal/ripe-ncc-certification-repository-terms-and-conditions> [hereinafter "*RIPE NCC Certification Repository Terms and Conditions*"]. It is beyond the scope of this report whether that agreement would be legally binding as to all parties who access the RIPE NCC TAL.

²⁷ See *Trust Anchors and Keys*, IANA, available at <https://www.iana.org/dnssec/files> (last visited Dec. 27, 2018).

²⁸ See Steven B. Roosa & Stephen Schultze, *The "Certificate Authority" Trust Model for SSL: A Defective Foundation for Encrypted Web Traffic and a Legal Quagmire*, 22 INTELL. PROP. & TECH. L. J. 3, 6-7 (2010) (discussing documents "purport[ing] to be . . . agreement[s] between the CA and the relying party/end user" and stating that "[t]he end user's assent to these standard documents is generally neither obtained nor sought").

²⁹ See *ARIN's Internet Routing Registry (IRR)*, ARIN, <https://www.arin.net/resources/routing/> (last visited Dec. 27, 2018) [hereinafter "ARIN IRR"].

about losing that protection. The loss of protection would enable wider ROV for North American-originated routes, and it would also enable open-source software developers to automate the installation of validator software that covered all five global regions. Examples of agreement-free distribution in other contexts are suggestive, if imperfect, evidence that such a path would be reasonable for ARIN to consider.³⁰

However, there are risk-reward tradeoffs that are also worth considering. One of these tradeoffs has to do with the allocation of liability and responsibility for potential problems stemming from repository use. We concluded above that it is ultimately uncertain whether RPKI repository providers might face claims of negligence that have some plausibility. At the very least, they might have to expend resources defending against ultimately unwarranted allegations. Judgments about the seriousness of these risks might differ, given the relative novelty of RPKI and the range of possible applications of the law. Certainly, we cannot rule out the risk of legal costs. To mitigate this risk through legal design (in addition to efforts to mitigate risk through investment in service-level quality), there are two broad strategies that an RIR might pursue. The first of these is the one currently pursued by ARIN: the use of an RPA. The second of these is for an RIR to spin off a special-purpose legal entity solely for the provision of RPKI information. The latter approach might enable an RPA-free distribution even in the face of legal risks.

Contract as a strategy for limiting legal risk. The RPA can mitigate much of the legal risk posed by RPKI provision. That is because contract clauses that explicitly limit liability and establish that the agreeing party assumes various risks often suffice to defeat negligence claims (and similar claims of breach of implied warranties) asserted by parties to the agreement. For instance, in Virginia, the

³⁰ Any comparisons of the probabilities and magnitudes of legal risk between RPKI providers, IRR providers, the DNS Root Zone Trust Anchor provider, and TLS providers are inevitably fraught with difficulty. This is due to the near-complete absence of public legal disputes involving their provision. Nevertheless, we may hazard some comparisons. First, IRR is only partially analogous to RPKI. This is due to the fact that there are many publishers of IRR information, whereas RIRs serve (by design) as exclusive publishers of RPKI information for their regions. As a result, any failure or accident involving an RPKI publisher is covered by less redundancy than is present for IRR. In contrast, the risks posed by a failure or accident involving IANA's DNS functions could pose similarly significant downstream ramifications as an RPKI failure or accident. As a result, the fact that IANA continues to publish its Root Zone Trust Anchors without requiring an RPA is instructive in the RPKI context. Finally, the case of TLS certificate authorities ("CAs") is mixed. Failures or accidents on the part of CAs can lead to website unavailability, but such incidents are usually resolved by end users who route around the problem. In some cases, legal risk that, in principle, might be shouldered by CAs is passed onto browser and operating-system vendors via indemnification clauses, though this is not a universal practice. As a result of this diversity, it is difficult to draw direct lessons for RPKI from the TLS context. In any event, as noted above, despite periodic problems with TLS (and, prior to TLS, SSL) certificate-provision, we are not aware of any court cases adjudicating responsibility for end-user harm.

Stepping back, it is worth noting that all of this discussion happens against the backdrop of widely-known best practices for mitigating failures and accidents among upstream service-providers. Harms are most likely to arise when downstream parties have failed to adopt best practices.

jurisdiction whose law the RPA selects to govern disputes, courts have admitted contract clauses as evidence of the express assumption of risk by a party participating in a risky activity. Virginia courts also honor clauses limiting liability in some circumstances. Finally, Virginia law tends to allow parties to disclaim liability for a counterparty's consequential damages in transactions like the RIR-relying party transaction. It is reasonably likely that courts would uphold similar clauses found in the RPA. Such clauses—analogueous to the “as-is” license language that typically accompanies open-source software—enable a service provider to bind direct counterparties to a contractual allocation of risks.

The value of such terms depends on the seriousness of the risks against which they defend compared with the costs they impose on the organizational mission. ARIN's RPA contains terms that seek to limit liability and allocate risks, suggesting ARIN favors a cautious approach to liability risk. RIPE NCC follows a similar approach. Its Certification Repository Terms and Conditions state that users employ the RIPE NCC repository at their “own risk” and that RIPE NCC “is in no way liable for direct or indirect damages” stemming from activities involving the repository.³¹ As noted above, other RIRs have opted not to require an agreement by relying parties.

Ultimately, the decision is a matter of judgment, and both using and foregoing an agreement are reasonable choices for an RIR and its community members. While ARIN would be on sound footing adopting the approach of the three RIRs that offer their TALs without placing it behind an agreement, we view it as reasonable for ARIN to maintain their RPA for the purpose of contractually limiting liability and establishing that relying parties assume various risks of using the Repository information, which is likely to cut off some of ARIN's legal risk. As noted above, in order to make the RPA effective, ARIN would need to continue presenting the agreement visually to users of the RPKI repository. This would require software developers to build manual RPA-acceptance processes into their validator software packages—a deviation from typical methods of automated open-source software installation, but not uncommon in software installation more broadly.

Corporate separation as a strategy for limiting legal risk. There are clear trade-offs to dropping the RPA, and to revising it. For this reason, it may be valuable for the North American Internet community to consider creating a separate corporate entity that would be responsible for the operational aspects of RPKI repository provision.

Why might a new organization be worth considering? At present, the reason why ARIN requires its relying parties to sign an RPA is to ensure that legal liability—even if unlikely—does not threaten the broader mission of the organization. If ARIN were to divest itself of the repository-provision side of RPKI and instead contract with a separate entity to publish a trustworthy repository, it would no longer bear direct legal risk related to outages. The new organization in charge of operating the RPKI repository would then be in a position to make its own decisions regarding how much

³¹ *RIPE NCC Certification Repository Terms and Conditions*, *supra* note 26, at art. 4.

legal risk to bear vis-à-vis relying parties. Because it would not have other functions that it would have to safeguard, it might see fit to bear more legal risk, such as by eliminating the RPA altogether. Assuming ARIN's transfer of responsibilities did not run afoul of legal rules that pass liability through to associated entities, ARIN itself would no longer bear legal risk from repository operation. The two organizations could coordinate via contract regarding RPKI operations. The corporate structures (i.e., boards, procedures) would have to be designed and observed to ensure separateness between the two entities.

Of course, there are many aspects of corporate separation that are beyond the scope of this report, but which would affect the viability of the proposal. At this stage, we recommend that corporate separation be considered alongside other methods of reducing legal barriers to RPKI adoption. It should be discussed by ARIN and the Internet community in the coming months, as dialogue on RPKI progresses.

The benefits and costs of forgoing indemnification. Let us now set aside the idea of corporate separation and turn back to the RPA. If ARIN determines that it is appropriate to maintain the RPA, it may still be valuable to remove (or, in the alternative, revise) the RPA's indemnification clause in favor of an "as is" disclaimer of warranties.

What does the indemnification clause do? At its root, indemnification requires the relying party to bear the burden of various costs associated with a covered set of legal risks. In case of the RPA, the covered set of legal risks is expansive. The RPA's indemnification clause covers "any and all claims" that are "asserted by a third party in connection with" two types of events—(i) the use of RPKI information and services or (ii) the breach of the RPA's terms.³² The clause covers situations where the use or breach was by the relying party or by any "[a]ssociated [p]ersons," such as customers or clients.³³

In such a situation, what would the relying party owe to ARIN? The key terms of ARIN's indemnification agreement—"indemnify, defend, and hold harmless"—impose distinct responsibilities.³⁴ First, the duty to "indemnify" would require the relying party to pay for a covered set of losses suffered by ARIN after they were established through a legal process. Separately, the "duty to defend" would require the relying party to cover the ARIN's "expense of defending suits" alleging harm from covered activities.³⁵ Finally, some courts would treat the obligation to "hold harmless" as a right of ARIN to be released from suit brought by the relying party.

³² Specifically, "any and all claims, demands, disputes, actions, suits, proceedings, judgments, damages, injuries, losses, expenses, costs and fees (including reasonable attorneys' fees and expenses), interest, fines and penalties of whatever nature." ARIN RPA, *supra* note 14, at § 7.

³³ *Id.*

³⁴ *Id.*

³⁵ *Capital Envt'l Servs., Inc. v. N. River Ins. Co.*, 536 F. Supp. 2d 633, 640 (E.D. Va. 2008) (applying Virginia law) (internal quotation marks omitted).

This clause provides benefits to ARIN in terms of reducing legal risk. Indemnity insulates ARIN from the monetary costs of adverse legal outcomes. And even before such an outcome might come to fruition, the “duty to defend” would require a relying party to cover the costs of legal defense of all claims falling within its scope. This right to defense would be available to ARIN early in litigation—before a court reached the merits of an underlying suit. As a result, it would allow ARIN to avoid litigation costs associated with even meritless claims brought against it.

At the same time, some interviewees regarded the RPA’s indemnification clause as a barrier to RPKI adoption. This is due in part to its mere existence and due in part to its particular terms. Recall that many organizations require formalized review of any agreement containing an indemnification clause. This costs time and deters network engineers from proposing RPKI within their organizations. The clause’s terms are also more stringent than some organizations would likely accept for the purpose of participating in RPKI as early adopters. In particular, the clause as currently drafted is quite broad: it requires relying parties to indemnify ARIN even for its own negligence. Interviewees noted that they were wary of binding their organizations to defend and indemnify ARIN for such a broad swath of activity. Further, they stated they were wary of indemnifying ARIN when the value of RPKI is unclear and when they were unsure of ARIN’s investments to ensure that RPKI functions reliably on a day-to-day basis. These factors have deterred a number of potential adoptees from advancing RPKI within their organizations.

Some interviewees suggested that the publication of more materials like ARIN’s Certificate Practices Statement and further outreach to members might allay their concerns about indemnification. Others suggested that the clause be revised or deleted from the RPA.

How should ARIN and its members evaluate the indemnification clause issue? It poses a tradeoff between the legal risk-reduction benefits of the clause and the drag it creates on RPKI adoption. The risk-reduction function of the indemnification clause is only relevant to the extent ARIN is likely to face legal risk for incidents involving RPKI. Though this risk is difficult to estimate, two factors suggest it is not grave. First, as discussed above, the RPA’s disclaimer of warranties and liability would provide substantial (though certainly not total) protection against liability. Second, research has not revealed negligence suits involving comparable security information, such as TLS (and its predecessor SSL), IRR, or DNSSEC. This absence is suggestive, but it is not dispositive. If a widespread RPKI outage were to harm many customers of ARIN’s relying parties, then the indemnification clause would indeed protect it from serious legal risk. On the other side of the scale, the community must weigh exactly how serious a block the indemnification clause is to their organizations’ willingness to agree to the RPA. Our interviews suggest that it is significant. Engineers seeking to adopt RPKI usually are doing so out of a sense that the framework will benefit the entire community, including their organizations. But the gains are not generally perceived to be so significant that they would justify indemnifying ARIN for its own negligence.

Given these equities, the ARIN membership might consider asking ARIN to drop the indemnification clause while maintaining clear disclaimers of warranties and limitations of liability within the RPA. Doing so would eliminate the moral hazard problem associated with having relying parties indemnify ARIN for its own negligence. ARIN is the best-positioned party to reduce its own risk of acting negligently and thus should certainly bear that burden. Further, ARIN is well-positioned to reduce the ultimate risk of harm through investments in the quality of its provision and through clear disclaimers applicable to relying parties. Though the elimination of the indemnification clause would shift some legal costs from relying parties to ARIN at the margin and may drive up ARIN's insurance costs, the shift would simultaneously remove a clear barrier to RPKI adoption. Notably, ARIN *already* makes this tradeoff in some cases: it willingly drops the indemnification clause for certain governmental counterparties (discussed in Section 3.4 below). This suggests that the clause may not be strictly necessary.

Dropping the indemnification clause (while retaining disclaimers and limitations on liability) would also bring ARIN closer to comparable RIRs on this issue. ARIN is the only RIR that clearly imposes an indemnification clause as such on repository access and is the only RIR to impose independent obligations to defend or hold harmless. While RIPE NCC does disclaim liability,³⁶ this falls short of requiring that relying parties provide an affirmative promise to indemnify, defend, and hold RIPE NCC harmless. AfrINIC and LACNIC, for their parts, appear not to have comparable disclaimers of liability at all, let alone indemnification agreements. All in all, ARIN places more burdens on its relying parties than do the other RIRs.

ARIN's choice to impose an indemnification clause is also more burdensome than approaches taken by IANA with regard to its DNS Root Trust Anchors,³⁷ the providers of the OpenSSL Toolkit,³⁸ or ARIN's approach to its Internet Routing Registry.³⁹ To be sure, other security resources are provided against the backdrop of indemnification clauses, as are many Internet services—including residential “last mile” service.⁴⁰ The Terms of Use for ARIN's Whois Terms of Use, which network operators likely encounter through their typical operations, also contain an indemnification clause to which (at least to our knowledge) network operators have

³⁶ See *RIPE NCC Certification Repository Terms and Conditions*, *supra* note 26, at art. 4.

³⁷ See *Trust Anchors and Keys*, IANA, available at <https://www.iana.org/dnssec/files> (last visited Dec. 27, 2018).

³⁸ See *License, OPENSLL CRYPTOGRAPHY AND SSL/TLS TOOLKIT*, available at <https://www.openssl.org/source/license.html> (last visited Dec. 27, 2018).

³⁹ See ARIN IRR, *supra* note 29.

⁴⁰ See, e.g., digicert, Certificate Services Agreement §§ 6.3-6.4 (Apr. 12, 2017), available at <https://www.digicert.com/wp-content/uploads/2017/06/Certificate-Services-Agreement.pdf> (indemnification limited to claims arising out of the actions of the customer or customer's agent); Comcast Cable Commc'ns, LLC, *Comcast Agreement for Residential Services*, XFINITY, <https://www.xfinity.com/Corporate/Customers/Policies/SubscriberAgreement> (last visited Dec. 27, 2018); Charter Commc'ns, Inc., *Charter Residential Internet Service Agreement*, SPECTRUM, <https://www.spectrum.com/policies/residential-internet-tc.html> (last visited Dec. 27, 2018).

not objected.⁴¹ Similarly, providers of DNS services require their users to indemnify them.⁴² But all this is merely suggestive. The real question is one of ARIN's own optimization. Given the costs of the indemnification clause on ROV adoption and ARIN's ability to insulate itself from a significant bulk of liability risk through the use of disclaimers and explicit statements of risks, dropping the clause in favor of an "as is" disclaimer of warranties may be the best path, perhaps in conjunction with an increase in insurance coverage to address residual risk of litigation costs.

In the alternative, ARIN should at least consider narrowing the clause significantly, to remove the moral hazard problem arising from indemnification applicable to its own negligence. Specifically, it could limit indemnification to situations where a relying party or downstream customer of a relying party failed to live up to a set of common best practices in relation to RPKI's role in real-time routing practices.

Such practices, of course, are essential to Internet security in a world of widespread RPKI deployment, no matter the legal consequences.⁴³ Collaborative efforts that publicize best practices, such as the Mutually Agreed Norms for Routing Security ("MANRS"),⁴⁴ therefore serve a dual purpose of promoting RPKI adoption directly and reducing the chances of RPKI-related incidents turning into legal issues. These chances are likely low already, but the more investment that network operators put into best practices, the lower they become.

3.4 Direct Access by Governmental Entities

In addition to the general issues surrounding direct access to the ARIN RPKI repository, interviewees also raised issues specifically applicable to government entities. These have to do with terms in the RPA that government agencies regard as problematic. First, federal procurement law prohibits federal actors from authoritatively agreeing to the RPA's indemnification clause.⁴⁵ Second, under some

⁴¹ See, e.g., *Whois Terms of Use*, ARIN, at § B.5 (Apr. 9, 2014), https://www.arin.net/whois_tou.html.

⁴² See, e.g., *Oracle Services Agreement*, ORACLE DYN at § 14 (Apr. 6, 2017), <https://dyn.com/legal/dyn-services-agreement/>.

⁴³ See, e.g., Randy Bush, RFC 7115, *supra* note 1.

⁴⁴ See *Mutually Agreed Norms for Routing Security*, MANRS, <https://www.manrs.org> (last visited Dec. 27, 2018).

⁴⁵ The Anti-Deficiency Act prohibits government employees from authoritatively agreeing on behalf of the government to "unrestricted, open-ended indemnification agreement[s]" like the one in ARIN's RPA. See *The Anti-Deficiency Act Implications of Consent by Government Employees to Online Terms of Service Agreements Containing Open-Ended Indemnification Clauses*, 36 Op. O.L.C. at 1, 2012 WL 5885535 (Mar. 27, 2012), available at <https://www.justice.gov/file/20596/download> ("A government employee with actual authority to contract on behalf of the United States violates the Anti-Deficiency Act by entering into an unrestricted, open-ended indemnification agreement on behalf of the government. A government employee who lacks authority to contract on behalf of the United States does not violate the Anti-Deficiency Act by consenting to an agreement, including an agreement

circumstances, federal agencies are discouraged from agreeing to alternative dispute resolution procedures like arbitration.⁴⁶ Third, similar prohibitions operate at the state and local level and also sometimes forbid accepting agreements that specify the choice of law outside the state in which the governmental entity sits.⁴⁷ A number of interviewees stated that the presence of indemnification and choice-of-law clauses in the RPA were gating issues that prevented them from considering ROV.

This set of barriers is easily resolved. ARIN already has adopted a policy of modifying both clauses for governmental entities to the extent necessary to comply with applicable law or regulations.⁴⁸ This policy eliminates the concerns raised by interviewees about governmental entity access, although it would be valuable for ARIN and its members to publicize this policy more broadly and to inform the community of which government agencies that have previously received waivers or modifications of the problematic clauses.

3.5 Legal Barriers to Indirect Use

Beyond directly downloading an RPKI repository from an RIR or via a third-party software tool to conduct their own ROV, network operators may also benefit from using RPKI repository materials as inputs to support more complete services. For instance, services are emerging that use RPKI information to clean up IRRs. Others combine RPKI information, IRR information, and other inputs to create dynamic route-filtering advice for end users. Many of these services are being offered by third-party providers. End users employing these emerging services do not necessarily need to access the RIR RPKI repositories directly to benefit from RPKI. The same is true with services that translate RPKI ROAs into IRR objects and with public monitoring projects, such as Certificate Transparency reporting. This Section discusses legal barriers to the development of software and services that enable the indirect use of RPKI information.

The prohibited conduct clause. The major barrier to the development of indirect RPKI uses is the “prohibited conduct” clause in ARIN’s RPA. This clause

containing an unrestricted, open-ended indemnification clause, because no binding obligation on the government was incurred.”).

⁴⁶ See 5 U.S.C. § 572(b).

⁴⁷ See, e.g., 1 CAL. STATE CONTRACTING MANUAL § 7.86 (2017), available at <http://www.dgs.ca.gov/Portals/32/Users/141/25/3725/8%20Pages%20from%20SCM%20June%202017-3.pdf> (prohibiting agreement to indemnification clauses); Katherine A. Adams, Contract Law for State Purchasing Officers § III.G (Sept. 23, 2013), available at https://www.naepnet.org/resource/collection/A9EC9928-E0AA-4604-85AE-28941F4BE73C/Contracting_101_Handbook.docx (discussing rules governing choice of law clauses for Kentucky government entities).

⁴⁸ See *Registration Services Agreement (RSA) FAQ*, ARIN, https://www.arin.net/resources/agreements/rsa_faq.html (last visited Dec. 27, 2018) [hereinafter “ARIN RSA FAQ”].

states that information derived from the ARIN Repository may be made available to third parties only “so long as such use and disclosure is solely for informational purposes, namely reporting, educational, summary or statistical purposes, and such use and disclosure of the information *is not in a readily machine-readable format.*”⁴⁹ This imposes limit on parties’ authority to redistribute the repository or to circulate information that uses the information contained in the repository as an input.

The benefits of the prohibited conduct clause. The prohibited conduct clause is intended to protect ARIN against liability for accidents involving certain uses of its repository information. Consider the following hypothetical situation. Imagine that a party relying directly on information obtained from ARIN’s RPKI repository simply redistributed that information free of any agreement to anyone who asked for it. This could potentially open up ARIN to exactly the kinds of tort claims (discussed in Section 3.3 above) against which the RPA is designed to protect. The prohibited conduct clause defends against this hypothetical by placing the onus on relying parties to ensure that all users using information that they download from the ARIN RPKI repository are bound by the RPA. Unfortunately, it also impedes third-party RPKI offerings.

The costs of the prohibited conduct clause. Multiple interviewees stated that the prohibited conduct clause is an impediment to important software- and service-development efforts for RPKI. Recall that RPKI deployment presents a chicken-and-egg problem, where the value of issuing ROAs depends on the extent to which networks filter based on invalid and unsigned routes. Given that reality, the development of “turnkey” validation support services will remove important barriers to the deployment of ROV filtering. As noted above, ARIN took an important step in reducing the legal barriers to such development when it announced that third-party software developers could incorporate acceptance of the RPA into their software packages. This means that open-source validation packages like RIPE NCC’s can now build the process of downloading ARIN’s TAL directly into their workflow if they so choose. However, the prohibited conduct clause still prohibits would-be providers from selling or giving away services that include information derived from the ARIN Repository presented in common formats like JavaScript Object Notation (“JSON”) or Comma-Separated Values (“CSV”) or domain-specific protocols like RPKI-Router. (For instance, a router vendor might operate a ROA cache and enable a turnkey ROV configuration option within its routers.) These prohibitions make it harder than necessary to design useful RPKI-based services. They also block the emergence of service providers that would be most likely to promote best practices for RPKI usage.

Enabling indirect use. Given those costs, ARIN should consider altering the prohibited conduct clause to facilitate third-party RPKI system development. In addition to the beneficial step already taken by ARIN to support third-party validation software, we suggest two further strategies for reform. First, ARIN might

⁴⁹ ARIN RPA § 5, *supra* note 14 (emphasis added).

consider allowing distribution of services making use of RPKI information as an input on the condition that they require users to accept the RPA or an appropriate variant of it for the use case involved. This would be quite similar to the allowance ARIN has made for designers of validation software, but it would extend to more robust service providers. Alternatively, ARIN could require these robust service providers to protect ARIN via an intended third-party beneficiary clause.⁵⁰ This is a common arrangement used in many analogous settings, including free and open-source software,⁵¹ and if drafted well, can be reliably expected to be upheld in court.⁵² ARIN should evaluate such approaches and consider whether to enable them on a permissionless or permissioned basis.

Second, ARIN should also consider revising its prohibition on distributing information in “machine-readable format.” Even if ARIN does not wish to enable wholesale redistribution of its actual RPKI repository, it should still support broader distribution of information created with RPKI as an input for research and analysis purposes. Crucially, such information is far more valuable when in machine-readable format, because it enables sophisticated analysis and trendspotting. To strike a better balance, ARIN may wish to rephrase the clause to prohibit using repository data for purposes *other* than supporting operational RPKI ROV functionality. RIPE NCC, for instance, forbids using its Repository data for “identification purposes, advertising, direct marketing, marketing research or similar purposes.”⁵³ A tailored prohibition along those lines would enable active community members to share their RPKI analyses in computable formats without fearing retribution. Alternatively, ARIN could explicitly allow machine-readable distribution for specific research and analysis purposes.

4. ROUTE ORIGIN AUTHORIZATION

This Section turns from ROV-based filtering to Resource Certification access and ROA-signing. Without active signing of ROAs, there would not be any useful information to support ROV. As a result, it is necessary for RIRs to encourage IP address space holders to obtain the cryptographic keys that enable them to sign ROAs and then use those keys to begin doing precisely that. This Section discusses the legal issues that interviewees flagged as potential hindrances to access to, and active use of, RPKI keys. It is motivated by the objective of ensuring broad access to RPKI keys, consistent with reasonable allocation of risks and duties among stakeholders.

⁵⁰ See, e.g., RESTATEMENT (SECOND) CONTRACTS § 302 (1981).

⁵¹ See, e.g., *NASA Open Source Agreement v1.3 (NASA-1.3)*, OPEN SOURCE INITIATIVE, <https://opensource.org/licenses/NASA-1.3> (last visited Dec. 27, 2018).

⁵² See RESTATEMENT (SECOND) CONTRACTS § 302 (1981) (citing cases).

⁵³ *RIPE NCC Certification Repository Terms and Conditions*, *supra* note 26, art. 3.

4.1 Background

At present, the entire RPKI system is anchored in trusted allocations of RPKI private keys to parties that are authorized to originate routing announcements for particular IP address spaces. The responsibility for authorization is held by each RIR.

Holders of IP space administered by ARIN must sign two agreements to receive access to the private keys that enable the issuance of ROAs. First, entities that wish to sign ROAs must sign either an RSA if they received their IP addresses after ARIN was created or a Legacy Registration Services Agreement (“LRSA”) if they obtained their IP addresses before ARIN came into existence. The RSA and LRSA (which now contain identical terms) cover the entire scope of a relationship between ARIN and a member, including RPKI. Second, IP address holders that wish to sign ROAs must sign an RPKI Terms of Service Agreement. This agreement covers specific aspects of the ARIN-member relationship involving RPKI. Unlike the RPA, none of these agreements is browsewrap. Instead, they are explicitly signed by parties that wish to receive their RPKI keys.

Concerns raised by interviewees over these agreements are described and evaluated in the Sections that follow. In Section 4.2, we address legal impediments created by the LRSA. In Section 4.3, we address issues concerning governmental access to Resource Certification.

4.2 Legacy Access

To receive the private keys that enable the issuance of ROAs for legacy address space, entities must consent to ARIN’s LRSA.⁵⁴ Many interviewees stated that this requirement poses a barrier to RPKI adoption for legacy IP address holders. While some legacy address holders have overcome the barrier and chosen to sign ARIN’s LRSA, others view it as prohibitive.

Perceived problems with the LRSA. Interviewees reported that the LRSA might act as an impediment because of its terms—specifically, the requirement that signatories to “acknowledge[] and agree[] that” they lack property rights in their IP number resources.⁵⁵ Some legacy resource holders view this as an unreasonable concession due to their view that they hold rights that would be given away via such an acknowledgment. They view themselves as the owners or legitimate controllers of their legacy IP resources and do not want to run the risk of turning over any iota of their present control to ARIN.

Assessing the cost of LRSA barrier. The perceived hindrance posed by the LRSA’s clause regarding property rights is real, but it is difficult to measure its impact. At present, the broader issue of legacy resource treatment is negatively

⁵⁴ This includes entities that have already signed RSAs for non-legacy address space.

⁵⁵ ARIN, Registration Services Agreement (LRSA: Version 4.0) § 7 (Aug. 16, 2016), *available at* <https://www.arin.net/resources/agreements/rsa.pdf>.

impacting the comparatively narrow and logically distinct issue of RPKI adoption. Legacy resource holders that are interested in participating in RPKI but are apprehensive about signing ARIN's LRSA must decide which position they value more. Anecdotal evidence indicates that multiple parties faced with that tradeoff have opted to avoid RPKI. That is, the current linkage between the LRSA and RPKI access likely is not driving legacy holders to sign the LRSA. Rather, it is turning them away from RPKI.

It is worth exploring whether decoupling the issues could enable ARIN to better serve its goal of driving RPKI participation while respecting the rights of its full members and without reopening the contentious "property rights" issue. This would be especially valuable in North America, where there is a higher concentration of legacy IP holdings than in other regions. As a result, ARIN's decision to tie the RPKI to the LRSA poses a higher cost on RPKI adoption in North America than it would in other regions.

We do not mean to overstate the importance of the LRSA. It is not clear whether the LRSA is a "but-for" cause of non-adoption. At present, most network operators that have signed LRSAs still have not deployed RPKI. The same is true for the IPv6 address spaces held by IPv4 legacy resource holders that signed RSAs in order to obtain their IPv6 address blocks. Lessening the perceived burden of the LRSA would hardly be a silver bullet. In addition, transfers of legacy IP space continue to reduce the set of legacy holders for whom the LRSA might be barrier. Nevertheless, it would be valuable to remove the LRSA as a roadblock on the path to widespread issuance of ROAs.

A potential path forward. To achieve more widespread ROA-issuance, ARIN may wish to consider altering its approach to the "no property rights" clause. The key role played by the "no property rights" clause in the LRSA is to create a structure that enables ARIN to provide registration services to LRSA signatories under conditions it sees as appropriate for operating its authoritative registry. The LRSA gives a party "[t]he exclusive right to be the registrant" of a given address block, and the "right to transfer the registration" within the ARIN registry under the terms of ARIN's governance.⁵⁶ This set of rights is paired with the "no property rights" clause concept to clearly establish ARIN's control over how transfers and registrations take place within its registry. One can think of the "no property rights" clause as one side of a deal and the rights of registration and transfer as the other side.

ARIN and its members should consider whether to decouple that entire deal from the RPKI Resource Certification process. That is, they should consider offering a transactional pathway to obtaining RPKI private keys that neither requires a "no property rights" admission, nor delivers any rights regarding registration or transfer of IP space. By separating RPKI from the property rights controversy, ARIN would open the RPKI door to LRSA holdouts. ARIN could adopt an at-will, fee-for-service model for this pathway, in which ARIN protects all its other rights as put forth in the

⁵⁶ *Id.* § 2(b).

normal LRSA. Further, this clause could contain a provision allowing termination with explicit reversion to the status quo ante.

This would place ARIN closer to RIPE NCC and the Asia-Pacific Network Information Centre (“APNIC”). Both have constructed multiple pathways to receive RPKI services that do not require the signing of a full member services agreement.⁵⁷ For RIPE NCC, these include the options of (i) signing a “non-member service contract,” (ii) contracting with a sponsoring Local Internet Registry, and (iii) seeking an accommodation for special circumstances.⁵⁸ Each of these pathways separates the question of access to RPKI keys from the broader question of a legacy resource holder’s relationship with the RIR and requires those benefiting from RPKI services to compensate the RIR for the costs of making them possible. Thus, entities wishing to avoid an affirmative consent to the idea that they hold no property rights in registered IP resources (something also contained in RIPE NCC’s Standard Services Agreement)⁵⁹ can opt for one of the alternate pathways.

Under a non-member service contract, legacy resource-holders could be ushered into the RPKI fold without having to overcome their deep-seated opposition to agreeing to the LRSA. In such a contract, ARIN could retain broad rights to deliver or terminate RPKI services and support to parties unwilling to sign the LRSA. In essence, such a structure would give ARIN an ongoing option to deliver RPKI services to non-signatories of the LRSA. This could help bring more participants into the ROA process. The attractiveness of this approach would depend on the interest and willingness of paid-in ARIN members to facilitate greater service-provision to those not yet signed up.

4.3 Governmental Access to Resource Certification

As with access to the ARIN RPKI repository discussed in Section 3.4 above, governmental entities have special concerns when it comes to accessing Resource Certification. Both the RPKI Terms of Service and the RSA or LRSA contain indemnification, arbitration, and choice of law clauses that may be outside the bounds of an agency’s ability to contract. The solution here is identical to what was proposed above: ARIN and the NANOG community should publicize ARIN’s policy of dropping both clauses for governmental entities that are barred by law or regulation from agreeing to them.⁶⁰ ARIN should also present the RPKI Terms of Service to new

⁵⁷ See *APNIC Non-Member Resource Services Agreement*, APNIC (July 1, 2002), <https://www.apnic.net/about-apnic/corporate-documents/documents/membership/non-member-agreement/>; *RIPE NCC Services to Legacy Internet Resource Holders*, RIPE NCC (Mar. 16, 2015), <https://www.ripe.net/publications/docs/ripe-639>; *Policy for Resource Certification for Non-RIPE NCC Members*, RIPE NCC (Oct. 16, 2013), <https://www.ripe.net/publications/docs/ripe-596>.

⁵⁸ See *RIPE NCC Services to Legacy Internet Resource Holders*, RIPE NCC, at § 2 (Mar. 16, 2015), <https://www.ripe.net/publications/docs/ripe-639>.

⁵⁹ See *RIPE NCC Standard Services Agreement*, RIPE NCC, at § 10.2 (Dec. 27, 2016), <https://www.ripe.net/publications/docs/ripe-673>.

⁶⁰ See ARIN RSA FAQ, *supra* note 48.

LRSA (and RSA) signatories during the member onboarding process. This would save repeat visits between lawyers.

Our interviews indicate that governmental agencies also might view themselves as prohibited by internal policy from disavowing property rights in their legacy IP address allocations. The “non-member services” pathway to accessing RPKI keys discussed in Section 4.2 above would help address this concern.

5. USING PROCUREMENT CONTRACTS TO PROMOTE ADOPTION

Beyond remedying explicit barriers to adoption, it is also possible to use other legal mechanisms to affirmatively promote adoption. This Section discusses one such mechanism: procurement requirements.

Procurement requirements are demands that large organizations place on their suppliers. When a major corporation requires its outside law firms to have certain forms of insurance or to engage in certain sustainability practices with their office real estate, it acts as a significant force in the marketplace. The desire to tap into these opportunities encourages more expansive use of insurance and environmentally friendly building design and management.

Procurement requirements could also be a lever to promote RPKI. If a large organization required its network providers (whether last-mile ISPs, backbone providers, or otherwise) and outsourced information technology providers (including providers of email, file-storage, DNS, cloud and similar services) to adopt RPKI signing and filtering as a condition of doing business with it, then those providers would likely adopt the practice. Once they adopted the practice at the behest of one client, they could offer RPKI services to other clients, as well. This positive feedback can help drive RPKI adoption. Recall that RPKI adoption becomes increasingly valuable to every potential participant as more participants come online. This means that if a major participant joins the fray, it adds to the incentives for others to participate. This effect is even more powerful when a major actor uses leverage to convince *others* to participate. Procurement contract requirements function as just that kind of lever.

Who might serve the role of catalyst? One obvious candidate—and a significant potential actor—is the U.S. federal government. There is ample precedent for the federal government promoting new technology through procurement policy. For instance, in 2005, the Office of Management and Budget (“OMB”) sought to promote adoption of IPv6 by requiring agencies to achieve certain transition benchmarks over a series of years. As the network security community makes progress towards RPKI adoption, it should consider whether to encourage OMB to take similar steps with RPKI.

6. ENSURING OPERATIONAL READINESS

The goal of widespread RPKI adoption is at the heart of this report. But adoption is only beneficial to the extent it is implemented properly. Indeed, many interviewees emphasized that widespread RPKI adoption must be accompanied by high levels of operational competence to ensure that RPKI is a source of security, and not a source of newly-introduced problems.

Interviewees who raised this theme tended to coalesce around three main points. First, network operators that begin making routing decisions based on RPKI information must adopt best practices when they do so. They must prepare to handle outages on the part of the RIRs, and they must be ready to failover gracefully. In short, no informational service is reliable and available 100% of the time, and RPKI is no exception. As a result, all network operators must ensure that their networks are resilient in the face of unavailable RPKI publication points and other problems that may arise despite every participant's best efforts. Second, interviewees stated that network operators would benefit from greater clarity regarding how the five RIRs intend to deliver their RPKI services. Interviewees reported that standardized and expanded disclosures of service-level intentions among the RIRs would enable network operators to better prepare themselves for foreseeable contingencies when relying on RPKI. Third, RIRs must prepare to provide real-time support for RPKI services.

Although evaluation of particular best practices and service-level intentions among operators and RIRs is beyond the scope of this report, the general lesson is essential: it is far more valuable to *reduce* risks than to allocate them via well-crafted legal arrangements. At its best, good legal design can incentivize risk reduction, but the lion's share of risk reduction will depend on the initiative and ingenuity of engineers and technical staff at network operators and at the RIRs.

7. CONCLUSION

This report has evaluated barriers to adoption of RPKI driven by issues of law in North America, and it has proposed ways to overcome them. These include alterations to ARIN's agreements applicable on both the filtering and signing sides of the RPKI equation. They also include a proposal for using procurement requirements as a catalyst to adoption. The report is not meant, however, to be a definitive statement of best practices for the North American network security community. Rather, it is meant to open dialogue on salient issues facing the community. Each recommendation would benefit from public discussion and analysis to surface hidden tradeoffs and to take stock of the Internet community's priorities. If the report contributes to more clarity and candor regarding how best to evaluate the inherent tradeoffs in coordinating a cryptographic security framework in the decentralized environment of the Internet, then it will have served its purpose.