Resource Public Key Infrastructure (RPKI) Technical Analysis

ICANN Office of the Chief Technology Officer

Alain Durand
OCTO-014
2 September 2020

TABLE OF CONTENTS

This document is part of the OCTO document series. Please see the OCTO publication page for a list of documents in the series. If you have questions or suggestions on any of these documents, please send them to octo@icann.org.

Executive Summary

Border Gateway Protocol (BGP) is the routing protocol used by Internet service providers (ISPs) over the Internet. It has been around since the early 1990s. BGP routing incidents, like the widely publicized YouTube route leak by Pakistan Telecom in 2008, are known as route leaks and can create Internet-wide traffic diversions. They now occur daily and they take a large toll on ISP operations. These diversions can be the result of configuration errors, software bugs, or active attacks. The root cause of these problems is the lack of built-in security in the BGP protocol.

Retrofitting security has been a long, difficult, and still incomplete endeavor. The most advanced effort available to deploy today is called RPKI origin validation. RPKI origin validation uses the Resource Public Key Infrastructure (Resource PKI, or RPKI), a hierarchical framework of interlocking X.509 public key certificates anchored at the Regional Internet Registries (RIRs). Its objective is to validate that the ISPs originating Internet routes are authorized to do so by the holder of the corresponding Internet Protocol (IP) address blocks. RPKI origin validation has been around since 2011. It is now getting traction as the culmination of several factors, including the RIR-led efforts over many years to promote it and train engineers how to use it; the Internet Society's Mutually Agreed Norms for Routing Security (MANRS) efforts; and the U.S. Department of Homeland Security's funding of RPKI software development. This, combined with the growing impatience with route leaks leading to the sense that "something needs to be done," plus the examples set by some large providers (such as Cloudflare and NTT), has made RPKI origin validation a hot topic in 2020.

Still, the technology is immature. There are serious scaling issues that result in propagation delays which reduce the flexibility ISPs have to deal with emergencies and bring brittleness to the system. The Resource PKI system itself can be attacked. A catastrophic failure scenario could be hard to detect and even harder to recover from. Those risks are compounded by the deployment model which uses five trust anchors, opening up the possibility for data inconsistency and paving the way to an even larger number of trust anchors. Parties that do not use RPKI at all can also become collateral victims of a breach at any one of the trust anchors. The liability risks from those scenarios is considered so high by the American Registry for Internet Numbers (ARIN) that the RIR requires indemnification from any relying parties for their use of its RPKI data. The system has thrown the RIRs into the daily operation of the Internet as active participants, a role they may or may not be best suited for, as some recent incidents have demonstrated.

More critically, by limiting the scope to the origin of route announcements, RPKI origin validation only protects from the most naive attacks on the routing system. A robust routing security system requires full path validation, but that is significantly more complex.

A number of ISPs, Internet exchange points (IXPs), and cloud providers consider stopping route leaks coming from misconfigurations and software bugs with RPKI origin validation enough of an operational improvement that it is worth the cost of deploying this rather complex system. Still, anybody contemplating deploying RPKI origin validation should be aware of the current maturity issues and operational risks associated with it. Securing the routing infrastructure is not (yet) a simple matter of deploying a piece of software. The trade-off between protocol security and operational complexity must be carefully weighted.

1 Introduction

ICANN has the mission to ensure the stable and secure operation of the Internet's unique identifier systems and a commitment to "... Preserve and enhance the administration of the DNS and the operational stability, reliability, security, global interoperability, resilience, and openness of the DNS and the Internet;".[1] Undoubtedly routing and IP address are important elements of a stable operation of the DNS globally. For that reason, ICANN org took a deep dive into RKPI, its operational environment and practice to understand the current state of its deployment and more importantly, identify areas of potential risk, challenges and opportunity of improvement. We hope this document will help open the discussion in appropriate forums in order to find solutions to some of the issues this document highlights.

The widely reported[2] 2008 route leak of YouTube by Pakistan Telecom provided one of the seminal moments that brought security problems in the Border Gateway Protocol (BGP)[3] to a more general awareness. BGP helps Internet Service Providers (ISPs) interconnect and exchange routes in order to maintain a connected Internet. In the 2008 event, Pakistan Telecom was trying to block content that was deemed by the Pakistani government to be objectionable. To block that content, a "local" route to nowhere (a "null route") for YouTube's address blocks was injected by Pakistan Telecom into BGP. This "local" route was supposed to only impact local customers, blocking them from accessing YouTube. But the "route to nowhere" leaked to at least one of the ISPs providing international connectivity service to Pakistan Telecom, and propagated outwards from there, impacting large portions of the Internet and blocking many on the Internet all over the world from accessing YouTube. Another seminal moment was provided later in 2008 when a live demo[4] at the Defcon-16 conference demonstrated how easy it was to perform a BGP hijack.

Internet routing works the following way: An ISP (ISP1) connects a customer (Customer 1) to its network. ISP1 then "advertises" the reachability of Customer 1's public IP addresses via BGP to all its peers, indicating ISP1 has a direct connection to Customer 1. Those ISPs then re-advertise to their own peers that they have a one-hop-away connection to Customer 1 via ISP1. That way, a chain of ISPs, called an autonomous system path (AS path) is propagated by the BGP protocol through the Internet.

The underlying reason the Pakistan Telecom incident was able to occur is that there are no authorization checks built into BGP. Despite early efforts at filtering route announcements based on publicly available routing databases, many ISPs accept - *in good faith* - all routes offered by their routing peers. The effect is that leaked routes, either maliciously or accidentally, propagate quickly to all corners of the Internet.

BGP, the core protocol that underpins all routing on the Internet, traces its origins to the late 1980s and early 1990s. In those times, security issues were not given the prominence they receive today. Interoperability, scalability, flexibility, and stability of the routing system were the paramount concerns. RFC 1163[5] defines the original version of BGP, published in June 1990. It

1 See https://www.icann.org/resources/pages/governance/bylaws-en
2 See https://www.nytimes.com/2008/02/26/technology/26tube.html
3 See https://tools.ietf.org/html/rfc4271
4 See https://www.slideshare.net/nguyenduchaisp21/defcon-16pilosovkapela
5 See https://tools.ietf.org/html/rfc1163

includes a security considerations section that says: “Security issues are not discussed in this memo.” At the time this was not unique to the BGP specification, as many documents published by the Internet Engineering Task Force (IETF) from this era had similar text.

The Pakistan Telecom/YouTube incident is far from unique. Events like these are called “route hijacks” or “route leaks.” Note: The latter term is now preferred to describe this general type of event, covering both intentional and accidental injections of erroneous routing information. A complete taxonomy of route leaks is available in RFC 7908.[6] The Mutually Agreed Norms for Routing Security (MANRS) Observatory[7] supported by the Internet Society (ISOC) reported 328 route leaks for January 2020 alone; MANRS also describes how route leaks are counted.[8] Many high profile route leaks have been reported - among the more notable cases are incidents involving China Telecom circa 2017[9] and, more recently, Rostelecom on 1 April 2020.[10] Intentional or accidental, such route leaks create major disruption of Internet traffic.

Long before the Pakistan Telecom/YouTube incident, the IETF was well aware of BGP’s vulnerabilities. Among many efforts, a paper defining a framework for ensuring protocol correctness, Secure BGP (S-BGP), described in Section 4.4 of this document, was published in 2000. In 2006, following RFC 4272,[11] which was an effort of the Routing Protocol Security Requirements (RPSEC) Working Group, the IETF formed a working group to tackle this problem of routing system security: the Secure Inter-Domain Routing Working Group[12] (SIDR) followed by the SIDR Operations Working Group (SIDROPS).[13]

A wide range of technologies have been proposed to secure BGP.[14] The IETF has published a “BGP Operation and Security” best current practice, RFC 7454.[15] At one end of the spectrum, we find technologies such as the TCP Message Digest 5 (MD5) option to ensure BGP messages have not been tampered with. At the other end, we find technologies that attempt to check the authenticity of the BGP announcements, as described in Section 4.4 of this document. RPKI origin validation (RPKI origin validation) falls in that latter category. It will be the main focus of this document. RPKI origin validation architecture is documented in RFC 7115.[16]

RPKI, or Resource PKI, was built as a general-purpose public key infrastructure to make cryptographically verifiable assertions about Internet number resources, namely IP address blocks and autonomous system (AS) numbers. Some of the technologies described later in Section 4.4 planned to use it to manage all the assertions needed to achieve BGP AS path validation. None of those technologies gathered any significant traction in the operational

6 See https://tools.ietf.org/html/rfc7908
7 See https://observatory.manrs.org
8 See https://observatory.manrs.org/#/about
9 See https://dyn.com/blog/china-telecoms-internet-traffic-misdirection/ and https://arstechnica.com/information-technology/2019/06/bgp-mishap-sends-european-mobile-traffic-through-china-telecom-for-2-hours/
10 See https://www.securityweek.com/russian-telco-hijacked-internet-traffic-major-networks-accident-or-malicious-action?
11 See https://tools.ietf.org/html/rfc4272
12 See https://datatracker.ietf.org/wg/sidr/charter/
13 See https://datatracker.ietf.org/wg/sidrops/charter/
14 See http://www.potaroo.net/papers/BGP_Security_Literature_Review.pdf
15 See https://tools.ietf.org/html/rfc7454
16 See https://tools.ietf.org/html/rfc7115

community. However, a more limited effort in scope, focused only on route origin validation, did. That effort is also known as RPKI, which is a source of confusion between RPKI, the Resource PKI, and RPKI, the use case of the Resource PKI to perform route origin validation. So, sometimes, the latter is called Route Origin Validation, but that terminology also creates confusion when we use the terms Route Origin Authorization (ROA) and Route Origin Validation (ROV).

This document uses the term "RPKI" to mean at the same time the PKI and its use case. This seems to be the choice made in many presentations and articles on this topic. When the distinction between the two will be important, the author will use the terms Resource PKI to talk about the distributed public/private key certification database and RPKI origin validation to talk about the use case.

RPKI origin validation was first specified in its current form in 2011. For nine years, the U.S. Department of Homeland Security's funding of RPKI software development[17] combined with the more recent efforts of the ISOC MANRS operator group and the Regional Internet Registries' (RIRs) workshops have led the industry's efforts to build a Resource PKI infrastructure. Until two years ago, the overall adoption of RPKI origin validation was quite lackluster. However, recently the situation has changed and RPKI origin validation is now a hot topic in network operations circles. RPKI origin validation deployment was one of the main themes at the last North America Network Operation Group (NANOG) in February 2020.[18] It is perceived by its proponents as an important component of the security, stability, and resiliency of the Internet.

This paper explores technical, operational, and legal issues around PRKI. It will answer the following questions:

- What is RPKI and how does it work?
- What protections does RPKI provide?
- Why is RPKI getting traction now?
- Is RPKI likely the last significant step towards BGP security?
- What is the role of ICANN and the RIRs in RPKI deployment?
- What is the impact of RPKI on the IP address market?
- What are RPKI's liability considerations and to whom?
- What could be a catastrophic RPKI failure scenario?
- How could the technology evolve?

1.1 Terminology

It will be useful to define some commonly used terms in this paper.

An *IP address* is a numeric value used to identify end points to the network in the Internet Protocol (IP). There are two protocol families in use in today's Internet: IP version 4 (IPv4), where the numeric values are drawn from the number pool that is representable in 32 bits, and IP version 6 (IPv6), where the number pool is representable in 128 bits.

17 See https://www.pcworld.com/article/157909/feds_.net_security.html

18 See https://www.nanog.org/meetings/nanog-78/agenda/

An *address prefix* is a contiguous sequence of IP addresses that share a common prefix when represented in binary notation. The length of this common prefix in bits is termed the size of the address prefix.

An *Autonomous System* (AS) is a network with a single administrative scope of control, such as an ISP or a large enterprise network. An *Autonomous System Number* (ASN or AS number) is a unique number, generally assigned by one of the RIRs, to an Autonomous System. It is this number that is used in BGP to identify routing domains, typically ISPs.

Number Resources refers to IP addresses and ASNs.

An *X.509 public key certificate* is a digital attestation made by the certificate's issuer that a given public key belongs to the certificate's subject.

A *resource certificate* is a variant of the X.509 public key certificate where the certificate's issuer is certifying that the holder of the public key is also the holder of a listed set of Number Resources. The system of resource certificates is termed the Resource Public Key Infrastructure (RPKI or Resource PKI).

A *trust anchor location* (TAL) is a description of the location where a RPKI Trust Anchor can be retrieved. A TAL contains both a *uniform resource identifier* (URI) that points to a trust anchor self-signed root certificate and the base64 encoded public portion of the key that is used to sign that repository.[19] The TAL is used as a starting point to walk the RPKI tree of CA repositories. The root certificate has a special attribute which points to the CA repository. Each trust anchor, or root certificate, is the logical equivalent of the *key signing key* (KSK) for DNSSEC.

A *route leak* is "the propagation of routing announcement(s) beyond their intended scope" (RFC 7908[20]). The causes of route leaks include route mis-origination, route hijacks, route policy violation, and others.

2 RPKI: Background

IP address allocation and IP network routing are two different, although related, topics.

IP addresses are allocated to Regional Internet Registries (RIRs) from ICANN as a component of the IANA Numbering Services performed by ICANN's affiliate PTI as described in the "Service Level Agreement for the IANA Numbering Services" between ICANN and the five RIRs. RIRs allocate or assign them to network operators such as ISPs. There are five RIRs, one per "continental" service region: the African Network Information Center (AFRINIC)[21] for Africa, the Asia Pacific Network Information Center (APNIC)[22] for Asia and Oceania, the American Registry for Internet Numbers (ARIN)[23] for North America and parts of the Caribbean, the Latin American and Caribbean Internet Addresses Registry (LACNIC)[24] for Latin America and other parts of the

19 See https://tools.ietf.org/html/rfc8630
20 See https://tools.ietf.org/html/rfc7908
21 See https://afrinic.net
22 See https://www.apnic.net
23 See https://www.arin.net
24 See https://www.lacnic.net

Caribbean, and the Réseaux IP Européens Network Coordination Centre (RIPE-NCC)[25] for Europe, the Middle East, and parts of Central Asia. (In the Europe service region, RIPE designates the community and the RIPE Network Coordination Center (RIPE NCC) designates the RIR providing services. This is a similar distinction as ICANN, the community versus ICANN org, the organization supporting the community. The other regions do not make this distinction.) The policies around the allocation of IP addresses, which may differ depending on the interests of the Internet community within each region, are developed within those communities and implemented by each of those five RIRs.

When a policy has the agreement of all five RIRs according to their policy development process and requires a specific action or outcome in order to be implemented, it can be subject to the Global Policy Development Process. The ICANN Address Supporting Organization (ASO)[26] Address Council,[27] also known as the Number Resource Organization (NRO)[28] Number Council,[29] coordinates the Global Policy Development Process.

Obtaining an IP address block from one of the RIRs, or via the IP address market now that the RIR free pool of IPv4 addresses is exhausted, does not provide any guarantee that those addresses will be reachable over the Internet. To obtain routing service for an address block, one needs to turn that block into a routing prefix and get one or more ISPs to announce that prefix to the rest of the Internet. ISPs follow their own policies, i.e., independent rules, to accept or reject such announcements from their peers and/or their customers. While ISPs do share experiences (and loosely coordinate) in Network Operator Groups (NOGs) in various regions of the world, there are no enforceable global routing policies that set what is acceptable behavior or not. The Mutually Agreed Norms for Routing Security (MANRS)[30] effort supported by the Internet Society is a step in that direction.

When two ISPs interconnect, they exchange routes through a BGP session. Due to the nature of the BGP protocol, there is no agreed-upon single “roadmap” shared by the networks – each router has their own representation of the Internet’s interconnections influenced by the network interconnection topology, as well as its own policies and the policies of its peers. Specifically, BGP is a path-vector protocol, a variation of distance vector protocols.

There was a discussion in the early 1990s at the IETF IPIDRP to migrate from BGP4 to an alternative inter-domain routing protocol known as IDRP developed as part of the Open Systems Interconnect suite of protocols, which is a link state protocol that would provide, among other things, such an agreed upon full map of the Internet.[31] This effort did not succeed, interest in IDRP faded, and development on BGP4 continued.

An ISP only has vision into its network and what its peers are telling it. The way routing among ISPs works is the following: ISP1 tells ISP2 which prefixes it can reach and vice versa. There is a similar process when an ISP connects with an end customer. One question immediately

25 See https://www.ripe.net
26 See https://aso.icann.org
27 See https://aso.icann.org/advisory-council/
28 See https://www.nro.net
29 See https://www.nro.net/about/address-supporting-organization/
30 See https://www.manrs.org/about/
31 See https://ftp.unpad.ac.id/ietf/ietf/ipidrp/ipidrp-minutes-92nov.txt

arises: what filter, if any, should ISPs put in place to accept or reject announcements coming through such peering or customer relationships?

In the case of ISP-to-customer peering (that is, where a customer can be either another ISP or a leaf network), the answer to this question is generally straightforward. The customer's address space comes from either the ISP itself, a different ISP, or has been obtained directly from an RIR and the prefix for that address space has been communicated to the ISP, potentially during the service contract negotiation phase. The ISP can set up a filter that says: I allow my customers to advertise routing prefixes for addresses registered to them, but nothing else. In a typical provider/customer relationship, the provider is expected to advertise either a default route (a route that can be used when no other routes are available) pointing to a router with more information or a full enumeration of the reachable routes to its customer.

In the case of ISP-to-ISP peering, what gets filtered gets more complicated. The expected behavior for an ISP to announce routes is the following:

- The ISP will advertise to its customers all routes learned from all other customers, all providers, and all peers.
- The ISP will advertise to its upstream providers all customer-learned routes, but no peer-learned routers and no provider-learned routes, and all ISP-local routes.
- The ISP will advertise to peers all customer-learned routes and all ISP local routes, but no peer-learned routes and no provider-learned routes.

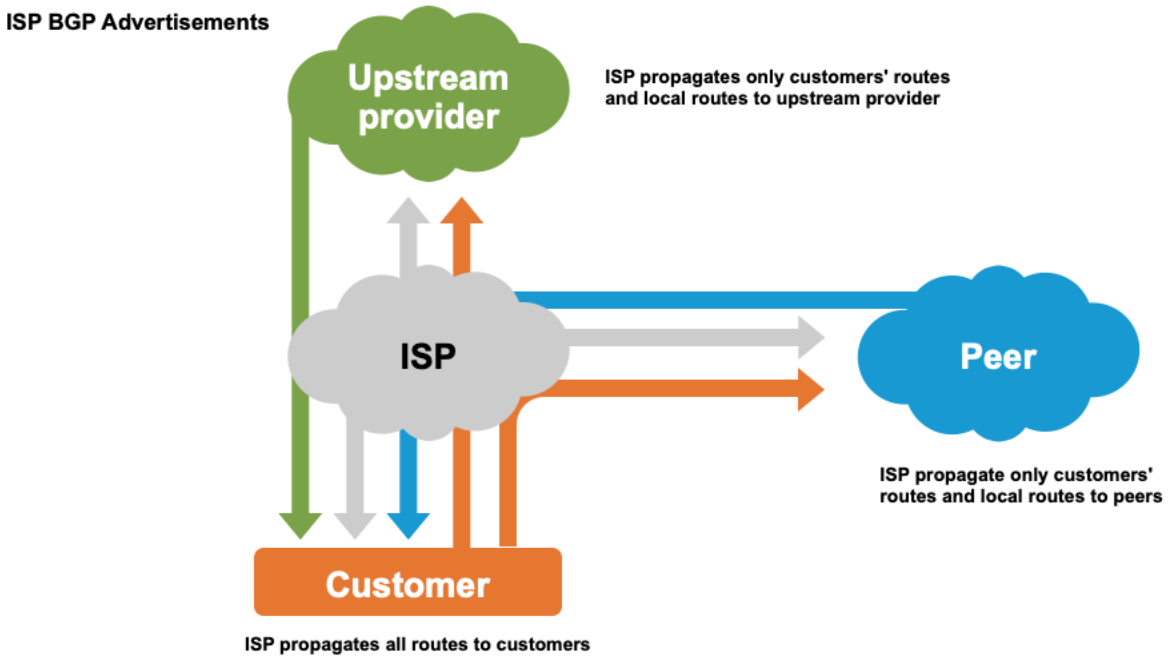

The question becomes, what should ISP2 accept from ISP1? Ideally, ISP1 would have a list of ISP2's customers' prefixes and would only accept announcements for those. However, such customer relationships are often considered private information and are not always disclosed. The situation becomes more complicated when there is a chain of ISPs. The further away the ISP is from the original customer announcing a prefix, the more complicated it gets to assess

the validity of the route announcement. The Internet Routing Registry (IRR) and the Routing Policy Specification Language (RPSL) RFC 2622[32] were created to help fix this problem.

2.1 Routing Registries: From IRR to RPKI

Conceptually, the Internet Routing Registry (IRR) system comprises a set of independent registries containing routing information, with different policies and overlapping data. It is one of the tools developed by a number of organizations independently to share information among ISPs about their intent with regard to announcing IP address blocks. The IRR contains "route objects", i.e., information used by ISPs to configure routers and peering associations. The IRR contains other objects. One such object is "Aut-num," which is considered a starting point to document the routing policy of a network, which was the main objective of the RPSL. The origin of the IRR can be traced back to the mid-1990s, as part of a 1994 National Science Foundation (NSF) award[33] to Merit Network[34] to set up a "routing arbiter." The routing arbiter was intended to address the increasing complexity within the emerging Internet, which had discrete private and publicly funded service platforms, and the policies that applied to the points of interconnection were increasing in complexity.

The global IRR is actually made of more than 20 independent (and sometimes inconsistent) databases. Some of these databases are public, such as the RIPE IRR[35] and the Routing Arbiter Database (RADb),[36] some are not. ISPs publish their routing policies and announcements in one or more of these databases and then use the IRR to build route announcement filters. The premise is that end customers would get their prefix registered by their ISP in the IRR. This would allow their ISP to filter out any erroneous prefix coming from the customer.

One major issue facing the IRR system is that in many cases route registries use an "open write access" policy: there is little to no validation of what goes into some of these databases, and efforts to create an amalgam of all such IRRs inherit the issue of dubious authenticity of individual entries. It is unknown if the objects within these databases are genuine, or even if they have been subsequently tampered with. An additional concern is that objects in the database tend to go stale quickly due to lack of maintenance. These issues render the quality of the IRR data questionable and make its use to automatically build route filters problematic.

The RPKI origin validation can be understood as an effort to go beyond the current IRR in some respects and to actually do less than the IRR in other respects. It relies on a new public distributed database, the Resource PKI, that will provide assurance as to which party can speak authoritatively for an IP address block. A digitally signed attestation that provides permission for a network with a specified AS number to originate an IP address block combines both authentication and non-repudiation. However, Resource PKI will not guarantee that the data is up-to-date: nothing forces an IP address holder to maintain their information, but all certificates in the RPKI have expiry dates and at some point the certificate times out, which means that the digital signature can no longer be verified and the signed attestation should be disregarded.

32 See https://tools.ietf.org/html/rfc2622
33 See https://www.nsf.gov/awardsearch/showAward?AWD_ID=9321060
34 See https://www.merit.edu/research/projects/
35 See https://www.ripe.net/manage-ips-and-asns/db/support/managing-route-objects-in-the-irr
36 See https://www.radb.net

Root CAs are usually valid for 10+ years. ISP CAs are typically good for a year, depending on the contract with the RIR. Resource RPKI also will not guarantee that the associated routing action will actually happen. A prefix holder does not need the permission of the originating autonomous system to make those assertions, and the AS holder is under no obligation to make a particular announcement (see example in Section 2.5.1). Finally, RPKI origin validation is limited in scope to the origination point of a prefix and will say nothing about the actual routing path taken in propagating the route object through the network.

Just like DNSSEC,[37] RPKI origin validation has two operational components: route origin attestation and signing, and route origin validation.

Generating RPKI origin validation objects entails an IP address block holder making a verifiable authorization for an AS to originate a routing advertisement for the corresponding prefix in the form of a digital signature attached to the IP address holder authority. Signing provides the AS with the authority to announce the IP address block holder's prefix into the routing system, which any other party, e.g., peers or providers, can validate.

Validating the origination of route objects is a process used by network operators. By maintaining a local copy of all current attestations that have been published in the distributed RPKI framework, a network operator can validate the authenticity of all such origination statements. The aggregate set of all such attestations (route objects and the associated origin ASes) can be used to construct a filter set, and this filter can be used on a BGP speaking router to automatically accept, reject, or prioritize received BGP route announcements and to have some confidence that those actions are being performed based on authentic data.

2.2 RPKI origin validation signing: Route Origin Authorization

The Resource Public Key Infrastructure (RPKI) is a hierarchical framework of interlocking X.509 public key certificates. The RPKI currently has five trust anchors, each operated by one of the RIRs. The resource sets in each of these trust anchors cover the entirety of the Internet number space (both IPv4 and IPv6), allowing each of these RIRs, acting as a Certification Authority (CA), to issue a certificate for any subset of the entire Internet's number space.

Each RIR maintains a Certificate Practice Statement that describes the practice in which the RIR issues a certificate to a subject that matches the current resource allocation state for the subject entity. These certificates are Certificate Authority certificates (CA certificates) that permit the subject entity to issue further RPKI certificates. These RPKI certificates may be CA certificates when the subject entity is itself a local Internet registry (LIR) that performs number allocations, or they may be End Entity certificates (EE certificates) that are used to certify a key used to generate a digital signature.

As certificates are themselves signed objects, the sequence of issued certificates from issuer to subject forms an interlocking chain, where each certificate is signed by its issuer. The root of this chain is the trust anchor, which is a self-signed certificate. Validation of a certificate entails forming an interlocking sequence of issuer/subject certificates that start with a trust anchor and terminate in the certificate being validated.

[37] See https://www.icann.org/resources/pages/dnssec-what-is-it-why-important-2019-03-05-en

One class of signed objects in the RPKI framework is that of Route Origin Authorization (ROA) objects. A ROA is a cryptographically verifiable statement that says: "Prefix P (and its sub-prefixes, up to the specified maximum prefix length) can be announced by the Autonomous System Number A, signed: the holder of the IP address block corresponding to prefix P." In RPKI terms, prefix P is now "covered" by a ROA.

There can be multiple ROAs covering the same prefix. The simplest example is multihoming, where a prefix is originated by two (or more) ISPs. The first ROA will say that the prefix can be originated by the network that uses AS 64496, and the second ROA will say the same prefix can also be originated by the network that uses AS 64497. Another example is the case of an ISP transitioning some address space from one ASN to another. For a period of time, that address space will be covered by two ROAs, one for the old ASN and one for the new.

As mentioned earlier, the permission of the ASN holder is not required, nor even is it strictly necessary for the ASN to be consulted for approval before issuing the ROA. Although it happens (see Section 2.5.1), it is at best a questionable practice for an IP address holder to create a ROA allowing its prefixes to be originated by an ASN it has no relationship with.

The process to sign a ROA follows a conventional X.509 signing process. The IP address block holder creates a self-generated private key/public key pair and an EE certificate signing request. That EE certificate is issued by the holder's CA and is published in the CA's repository of signed subordinate products. The IP address block holder signs the ROA with the private key. It then assembles a digital object that contains the EE certificate, the digital signature, and the route origin assertion (IP address block, maximum prefix length, and origin AS). That digital object is the actual ROA ready to be published. The ROA is then placed into the publication point listed in the EE certificate.

Most ISPs and large networks have received multiple allocations of IP addresses. Not all of these blocks of addresses derive from the same delegation point. ISPs receiving addresses from multiple sources will need to operate multiple CA certificates, one for each delegation path.

All CA and EE certificates expire. In general, the expiration times are aligned to the current contractual arrangements relating to the resource allocation or assignment. Where there are no formal agreements in place, the RIRs have used the community policy processes to develop an agreed certification policy. For example, in the RIPE region, those certificates are valid for 18 months.

2.3 Resource PKI: Delegated Repository vs. Hosted Model

The original Resource PKI model described in RFC 7115[38] envisaged that the RIRs would issue CA certificates to LIRs and ISPs, and these entities would run their own delegated repository of certificates and signed products (such as ROAs). The repository structure is described in RFC 6481.[39] A repository contains all certificates issued by the CA, ROAs, a Certificate Revocation List (CRL), and a manifest of all objects. This model requires a considerable investment on the

[38] See https://tools.ietf.org/html/rfc7115
[39] See https://tools.ietf.org/html/rfc6481

part of LIRs and ISPs, both in expertise and operational capabilities. The server hosting the repository must be maintained 24/7. A recent experiment performed by ICANN OCTO with RPKI validator software showed a number of RPKI repositories being offline or having stale data. Also, until recently, another key consideration was the lack of quality software tools that could perform these functions.

In response to these considerable barriers to adoption, the RIRs developed a hosted model. This model was used initially to bootstrap the system and is used now to facilitate the Resource PKI adoption. In this model, each RIR hosts a large repository of certificates and signed products on behalf of its LIRs and/or members. The process to create ROAs is thus greatly simplified: access your account on your RIR's web server, describe the prefix and AS that will be the route origin, and click "sign." The strict RPKI structure of issuer and subjects is preserved in this model, so the certificates are unaltered, but rather than many discrete parts to the certificate infrastructure with many discrete points of publication, the host structure is considerably simpler.

It is fair to say that the delegated model is more targeted at organizations that have experience running cryptography and managing certificates and have an information and operational security requirement to operate such functions in-house. The hosted model is better suited to organizations wanting to use the Resource PKI but not willing or necessarily able to build up sufficient cryptography expertise in-house.

2.4 RPKI origin validation: Route Origin Validation

RPKI origin validation "Route Origin Validation" (ROV) is a process used by routers at ISP BGP peering points to automatically accept or reject BGP route announcements based on a pre-computed filter list derived from the Resource PKI and ROAs. Participating ISPs are known as "relying parties" of the Resource PKI, a term of art in the X.509 PKI community.

One of the key strengths of the RPKI origin validation model that made its adoption relatively easy compared other approaches to secure BGP (described in Section 4.4 of this document) is that routers performing ROV used in peering do not need to perform the CPU intensive cryptographic functions necessary to validate the ROAs' digital signatures, the precursor to forming and maintaining route announcement filters. Most high-end routers today use dedicated hardware to accelerate packet processing and forwarding. The only intensive role that the CPUs on these routers accomplish is participating in the routing protocols and calculating the routing tables to be downloaded in the forwarding tables on dedicated line cards. As such, routers tend to have fairly simple general-purpose processors.

Instead, BGP routers rely on external computers running validator software to perform the cryptographic validation and build prefix filters based on validated ROAs, augmented potentially by other local sources of routing policies. These prefix filters are then synchronized at regular intervals with the routers using the RPKI-Router protocol defined in RFC 8210.[40] Most router vendors have implemented the corresponding RPKI logic, although a recent "deployathon" (large scale testing event) organized during the APRICOT 2020[41] summit found issues with at least one major vendor.[42]

40 See https://tools.ietf.org/html/rfc8210
41 See https://2020.apricot.net
42 See https://nsrc.org/blog/rpki_deployathon

The validation software constructs a set of all prefixes/authorized ASN combinations described by any valid ROA that validator has obtained; this is a very different approach than DNSSEC validation that happens on-demand. To perform this task, the software needs to be constantly synchronized with all ROA repositories across the Internet.

Building this list starts with downloading five TALs, one from each of the RIRs. TALs are either pre-configured in the validation software or configured manually by the network operator. Once the TALs are obtained, the validator software can follow the publication pointers in the chain of certificates and recursively download certificates and signed product ROAs from the various Resource PKI repositories.

Local knowledge can be added to the list of synchronized ROAs. An example of how to do this is described in RFC 8416:[43] "Simplified Local Internet Resource Management with the RPKI (SLURM)." A simple use case for this is when an ISP uses private IP addresses (RFC 1918)[44] or private ASNs (RFC 6996)[45] within its network and wants ROA validation to apply to those networks in addition to public networks. As those resources are not, by definition, globally unique, they cannot be covered by a ROA covering public space. Section 5 of RFC 6491 (describing operation with an initial model of a single root) has a provision for that.

Once the list of all valid ROAs is obtained, the next step is the synchronization of the peering routers with validator software. As mentioned above, this is done using the RPKI-Router protocol defined in RFC 8210.[46] Now those peering routers are ready to apply the prefix filter to all received BGP announcements.

A BGP announcement that contains an announced prefix and an origin ASN can fall into one of the following three categories:

- Valid: there is a matching valid ROA (or example, a ROA can be invalidated by expiring/revoking the EE certificate) covering that prefix and origin ASN and the announced prefix length is compatible with the maximum prefix length specified in the ROA.
- Invalid: there are one or more ROAs covering that prefix, but either or different origin ASNs, or the announced prefix length is longer than the maximum prefix length specified in the ROA .
- Unknown: there are no covering ROAs.

The logic of ROV known as "reject invalids" used by some network operators today is the following:

- If the announcement falls in a valid state, it is accepted.
- If the announcement is invalid, it is rejected.
- If it is unknown ("NotFound" state), it is accepted.

43 See https://tools.ietf.org/html/rfc8416
44 See https://tools.ietf.org/html/rfc1918
45 See https://tools.ietf.org/html/rfc6996
46 See https://tools.ietf.org/html/rfc8210

The last part of the logic is necessary for the initial graceful transition into an RPKI origin validation world. However, the more difficult question of what should be done about "unknown" or "uncovered" prefixes remains unanswered. The end game is probably to reject them, but how to get there is still unclear. The current recommendation in Section 5 of RFC 7115[47] is that "As origin validation will be rolled out incrementally, coverage will be incomplete for a long time. Therefore routing on NotFound validity state SHOULD be done for a long time." A network operator rejecting "NotFound" today will probably cut itself off from a large portion of the Internet.

This out-of-band ROA distribution mechanism is an important element greatly simplifying the adoption of RPKI origin validation: there are no upgrades needed to the BGP protocol itself. However, as will be discussed in subsequent sections, the out-of-band propagation mechanism is not without problems.

Note there is sometimes some confusion around the term "RPKI invalid." An "invalid ROA" is a ROA whose digital signature cannot be verified in the Resource PKI, and must be discarded by validation software. A "BGP announcement marked invalid" is a BGP announcement for which valid ROAs do exist, but none of those ROAs match the origin ASN (or maximum prefix length) and can be discarded by the BGP speaker.

2.5 Data Quality

ICANN org set up two RPKI validators, RIPE validator[48,49] and NLnet Labs Routinator,[50,51] and compared the output of both validators. They matched 100%.

While running those validators, a few surprises appeared. They are reported here not to point fingers but to illustrate some of the current issues facing RPKI.

It should be remembered that ROAs are cryptographically verifiable attestations. Cryptographically verifying an attestation guarantees the attestation has been signed by the private key of the party claiming it made it. It does not necessarily guarantee that the attestation is objectively true.

2.5.1 Frivolous Origin ASN

As seen previously, a ROA is signed by the owner of an IP prefix, not by the owner of the AS number. As such it is entirely possible to insert frivolous ROAs into the system, signed by the rightful owner of an address block, but asserting some arbitrary AS number as the authorized origin AS.

The author found on 26 February 2020 the following ROAs while running validator software:

47 See https://tools.ietf.org/html/rfc7115
48 See https://www.ripe.net/manage-ips-and-asns/resource-management/certification/tools-and-resources
49 See https://github.com/RIPE-NCC/rpki-validator
50 See https://nlnetlabs.nl/projects/rpki/routinator/
51 See https://github.com/NLnetLabs/routinator

ASN	Prefix	Max Prefix Length	Originating RIR
137519	103.118.18.0/24	24	APNIC
5	103.118.18.0/24	24	APNIC
11	103.118.18.0/24	24	APNIC
10	103.118.18.0/24	24	APNIC

A WHOIS check reveals that the prefix 103.118.18.0/24 has been allocated by APNIC to Healthcare Pharmaceuticals, Ltd. in Bangladesh. AS137519 is also assigned to Healthcare Pharmaceuticals, Ltd. in Bangladesh, so the first ROA is expected. However, AS5 is assigned to Symbolics Inc., a defunct computer manufacturer based in Massachusetts, U.S. AS10 is assigned to the CSNET Coordination and Information Center hosted by BBN Systems and Technologies Inc., based in Massachusetts, U.S. AS11 is assigned to Harvard University in Massachusetts, U.S. The author contacted the 103.118.18.0/24 network manager in an effort to understand the latter three attestations. It turned out it was a configuration mistake that was corrected by the time this document was written.

2.5.2 Different Maximum Prefix Length

Another example of a surprising set of ROAs on 26 February 2020 is:

ASN	Prefix	Max Prefix Length	Originating RIR
8987	100.20.0.0/14	24	ARIN
14618	100.20.0.0/14	24	ARIN
16509	100.20.0.0/14	24	ARIN
16509	100.20.0.0/14	14	ARIN

The prefix 100.20.0.0/14 and the ASNs 8987, 14618, 16509 are all allocated to Amazon Inc.

With the first three ROAs, Amazon is asserting prefix 100.20.0.0/14 can originate from either AS8987, AS14618, or AS16509. Those AS numbers all belong to Amazon. Because the maximum prefix length is set to 24, Amazon is also asserting that any sub-prefix of 100.20.0.0/14 of length greater or equal to 24 can also be originated by the same ASN.

However, the last of those four ROAs has a maximum prefix length set to 14. It says that only the exact prefix 100.20.0.0/14 can be accepted from AS16509. It also says that any sub-prefix should be rejected. The no sub-prefix part of this fourth ROA contradicts the third ROA that said sub-prefixes were acceptable.

Route validators today are configured to accept any prefix that matches any ROA, effectively doing a logical OR between all ROAs. The last ROA with maximum prefix length 14 would be simply ignored.

It appears there was an operational reason for this configuration. The author contacted Amazon and the explanation given was that this situation was the result of operational difficulties in deleting ROAs in the ARIN region using the ARIN-hosted RPKI tool set. The larger ROA will stay in the system until it expires.

This example is not unique, and it has become a fairly common practice to leave in place the ROA associated with the exact RIR-allocated prefix length after more specific ROAs have been introduced. The current recommendation from the IETF SIDROPS Working Group is that "whenever possible, operators SHOULD use "minimal ROAs" that include only those IP prefixes that are actually originated in BGP, and no other prefixes."[52]

See Section 4.2 for a more in-depth analysis of maximum prefix length related issues.

2.5.3 Stale Data

A third example of operational problems is found in the output of the Routinator 3000 RPKI validator:

```
rsync://rpki-repo.registro.br/repo/HqatGkF4QDP6Set7UcbXnGGj2TkehDBZ24LGiaLAbd
zu/0/ABE2B87282296266F69CE04223629CFC8BFD6354.mft: stale manifest
rsync://rpki-repo.registro.br/repo/HqatGkF4QDP6Set7UcbXnGGj2TkehDBZ24LGiaLAbd
zu/0/ABE2B87282296266F69CE04223629CFC8BFD6354.crl: stale CRL.
rsync://rpki-repo.registro.br/repo/DmyLDjMgaeUsYnfjUCfi8BYTx4tsZvsFPDws5wDs4x
Fa/0/AB06C1515EDB643DB9DF8E7E1361A8BB0683DE7A.mft: stale manifest
rsync://rpki-repo.registro.br/repo/DmyLDjMgaeUsYnfjUCfi8BYTx4tsZvsFPDws5wDs4x
Fa/0/AB06C1515EDB643DB9DF8E7E1361A8BB0683DE7A.crl: stale CRL.
```

The author contacted staff at registro.br. The explanation given was that registro.br members are using the RPKI delegated repository model. At least two members did set up their own repositories at some point in the past, but stopped maintaining them. These repositories are now offline, resulting in the data associated with those repositories becoming stale.

2.5.4 Impact of Bad Quality Data in Resource PKI

The operational impact of false-but-verifiable ROAs is minimal. ROAs are not route announcements, thus bogus ROAs as discussed above will not direct undesired traffic to a victim AS or create any similar harm. However, they will open the address block holder who signed the bogus ROAs to route leaks as seen in Section 2.5.1. As such, this issue could be best described as a self-inflicted wound, and likely, a self-correcting problem with little to no impact on the global Internet. There is also the scenario where a would-be attacker gains control of the intended victim's RPKI credentials – which can cause these bogus ROAs to be published - then follow up with a routing attack that is aligned to the bogus ROAs.

Bad quality data, however, may create a reputation risk for RPKI.

[52] See https://tools.ietf.org/html/draft-ietf-sidrops-rpkimaxlen

2.6 RPKI and Legacy Addresses in the ARIN Service Region

ARIN provides registry services both to standard customers via a registry services agreement and to legacy resource holders (that is, entities that were assigned resources prior to the existence of the RIRs) who receive basic registry services from ARIN without any fee or contract (but only receive the services that they were receiving at the time of ARIN's formation). Some legacy IP address holders in the ARIN region have signed a Legacy Registration Services Agreement (LRSA),[53] effectively giving them membership to ARIN. However, a number of those legacy IP address holders have so far refused to enter into an LRSA, claiming the agreement would take away some of their perceived rights. ARIN refuses to offer Resource PKI services to legacy holders that have not signed the LRSA. As most legacy addresses are within the ARIN region, many of those legacy blocks cannot be covered by ROAs.

The RIPE community has directed the RIPE-NCC (the organization that provides RIR and other services to the RIPE community) to offer a non-member service contract[54] to enable legacy holders to use Resource PKI. The terms and conditions of this contract are perceived friendlier by the RIPE community than the LRSA is by the ARIN community. Therefore, it is less of an hindrance to RPKI deployment.

While RPKI origin validation adoption is still low, this legacy address issue at ARIN has not been a major problem.

2.7 Other Usage of Resource PKI

RFC 7909[55] suggests using the Resource PKI to sign Routing Policy Specification Language (RPSL) objects. Similarly, some policy proposals in several RIRs have been adopted, or are under discussion, to use the Resource PKI to validate the data in the IRR and remove bogus entries. For example, see RIPE's object clean-up proposal.[56]

As seen in Section 4.4, new efforts toward path validation propose to leverage the Resource PKI.

IPv6 Secure Neighbor Discovery (SEND) also makes use of the Resource PKI in RFC 6494.[57]

53 See https://www.arin.net/resources/guide/legacy/
54 See https://www.ripe.net/manage-ips-and-asns/legacy-resources/ripe-ncc-services-to-legacy-internet-resource-holders
55 See https://tools.ietf.org/html/rfc7909
56 See https://www.ripe.net/publications/docs/ripe-731
57 See https://tools.ietf.org/html/rfc6494

3 From One Root at IANA to Five (or More) Roots

The trust anchor model of the Resource PKI has evolved over time. This evolution had an impact on the overall system.

3.1 Original Model: Single Root

The 2010 Internet Architecture Board (IAB) recommendation on RPKI called for a single authoritative trust anchor, "aligned with the root of the address allocation hierarchy (now part of the IANA function)."[58]

For a number of technical and political reasons, this model was not implemented.

From a political perspective related to the autonomous control of national infrastructure, there was lively discussion about the necessity (or not) of a single root of trust for the Resource PKI with a key or keys that would likely be hosted in the U.S. (for example at ARIN[59]).

From a technical perspective, having a single trust anchor adds some complexity when transferring an IP address block covered by a ROA from one RIR to another.[60] Because certificates cannot be routed to multiple parents, this process might have required the single root trust anchor to take an active role during the transfer process, or necessitated a more complex setup with a level of indirection where RIRs could certify which resources they transferred out to another RIR. A recent blog post from APNIC contains a more complete analysis.[61]

In 2018, the IAB issued a new statement to recognize the de facto situation.[62]

3.2 Five RIR TALs for RPKI Origin Validation

Having five separate Resource PKI trust anchors addressed the political problem. The system could have essentially operated the same way as initially designed, where each RIR would claim coverage only for the set of resources they control. This would still generate complexities in the certificate infrastructure when number resources are moved from one RIR to another. The RIRs adopted a slightly different approach intended to reduce these complexities and associated brittleness of the certificate infrastructure.

To avoid giving the single root an active role during inter-RIR IP address block transfers, a technical fix was deployed by the RIRs: each of them, by virtue of this new technical capability,

58 See https://www.iab.org/documents/correspondence-reports-documents/docs2010/iab-statement-on-the-rpki/
59 See https://www.arin.net/vault/participate/meetings/reports/ARIN_XXVIII/ppm1_transcript.html
60 See https://tools.ietf.org/html/draft-rir-rpki-allres-ta-app-statement
61 See https://blog.apnic.net/2020/04/21/rpki-and-trust-anchors/
62 See https://www.iab.org/documents/correspondence-reports-documents/2018-2/iab-statement-on-the-rpki/

claims coverage for the entire address space and AS number space. That way, the transferred block could still be covered (for a time) both by the old certificate issued by the outgoing RIR and by the new certificate from the incoming RIR.

This technique simplified the certificate issuance and resource tracking in the Resource PKI for the RIRs. Instead of listing the complete set of IP address blocks ever delegated, they only list the prefixes covering the entire IPv4 and IPv6 address space, and all AS numbers.

As a result, there is not a single trust anchor for the Resource PKI system, but five such trust anchors. This is a new situation on the Internet, halfway between the DNSSEC system with a single root and the Web PKI with hundreds of trust anchors.

With more than one root, it is now possible to have unintentionally contradicting ROAs (that is, not part of an inter-RIR transfer) coming from different but valid Resource PKI validation paths. The holder and signer can now potentially be different entities. In most cases, this should not be a major issue; as long as at least one ROA categorizes a route object as "valid," then the other ROAs will be ignored. Harm can still happen, though, as will be seen in Section 5.5. This leads to three observations:

- Coordination is key. How is it achieved and how can it be maintained?
- Conflicts can arise. How will they be detected and fixed?
- Once you have more than one TAL, why stop at five? There can be any number, as we have already witnessed with the proliferation of Internet Routing Registries.

3.2.1 RIR Coordination

Coordination between RIRs is part and parcel of their operations. For example, IP address blocks and ASNs are regularly transferred from one RIR to another. Examples of transfers can be traced back at least to 2002, at the start of the Early Registration Transfer (ERX) Project.[63] The exhaustion of the IPv4 address free pool has made such inter-RIR transfers more frequent. Updating and coordinating RPKI is part of the normal transfer process. However, the RIRs are evolving. If the RIRs become global organizations, the current model defined in principle 1 the "Criteria for Establishment of New Regional Internet Registries"[64] would be altered and it is unclear if the current level of technical collaboration would continue.

3.2.2 Resource PKI Conflicts

Despite best practices, conflicts could still arise. How will they be detected and how will they be fixed is still a work in progress. The Resource PKI system would certainly benefit from global monitoring of the consistency of the five databases.

3.2.3 More Than Five TALs

RPKI origin validation validators handle ROAs irrespectively of their origin. They essentially perform a logical OR on all the verified ROAs. So, adding another source of ROAs outside of

63 See https://www.arin.net/vault/participate/meetings/reports/ARIN_X/PDF/erx.pdf
64 See https://www.icann.org/resources/pages/new-rirs-criteria-2012-02-25-en

the five RIR TALs is simply a matter of configuration. It is possible that an entity outside of the control of the RIRs could issue Resource PKI certificates based not solely on RIR registration data but on reputation data (or something else entirely, like legacy status for example). A prime example would be very large IP address holders that could directly publish a TAL pointing to their own repository.

The success or failure of such an endeavor would rest in the hands of the relying parties, that is, the network operators running RPKI validators. They could choose to add (or not) the TALs for such alternate sources of trust in their validation process. If enough of the Tier 1 networks and IXPs start to use these data sources, it is possible to imagine a future where additional TALs could find their way into the default configuration of validation software, as has already happened with IRRs.

3.3 AS0: Covering Unallocated Space

Beside route leaks, another routing security related phenomenon has been observed: IP address space squatting. Essentially, this is when an ISP, wittingly or unwittingly, is convinced to announce unallocated IP address space. This usually gets detected and fixed fairly quickly, but for a period of time, hours or days, that address block is connected. IP address squatting is often used to mount various types of attacks and abuses, for example to send spam.

As the entire address IPv4 space is now almost fully allocated, one could think IP address squatting should be less and less of a problem with IPv4. However, as many large legacy blocks are not advertised, in whole or in part, this may still be an issue worth investigating.

In IPv6, as only a small fraction of the address space is allocated, IP address squatting is clearly a security threat because it is unknown who is using that space.

This topic was originally covered in RFC 6491,[65] which described the operation of a single root for the Resource PKI. The recommendation was that IANA should issue an AS 0 ROA for all reserved IPv4 and IPv6 resources not intended to be routed. One (or multiple) set(s) of specific ROAs covering all unallocated address space would need to be issued by either a single entity in a model reminiscent of the original single root, or five times, one by each RIR.

A new policy proposal has been adopted at APNIC to create ROAs with AS0 for all APNIC unallocated prefixes.[66] This policy will be implemented by APNIC staff by creating a new TAL[67] for that purpose, bringing the total number of TALs to six. A similar proposal (with similar implementation) is under discussion at LACNIC, potentially bringing the number of TALs to seven. RIPE is considering a similar policy, but implemented differently[68]: the list of unallocated space would be contained in a RIPE-NCC managed SLURM file, available out of band and unencrypted. At the time of this writing, neither ARIN nor AFRINIC have considered such policies.

Besides the multiplication of TALs, the different (and uncoordinated) approaches taken by the different RIRs will create additional complexities in properly configuring RPKI validator software.

65 See https://tools.ietf.org/html/rfc6491
66 See https://www.apnic.net/community/policy/proposals/prop-132
67 See https://www.apnic.net/community/security/resource-certification/#st-anchor
68 See https://www.ripe.net/participate/policies/proposals/2019-08

Creating ROAs with AS 0 for unallocated space has other consequences that are analyzed in Section 5.6.

4 RPKI Technical Risks

As with any technology, particularly security-related technology, RPKI origin validation has been built on a series of technical choices that have consequences. First of all, X.509 is a complex system. Even though the RIR hosted model has made the signing part relatively easy, the delegated model retains the full complexity of X.509. The validation side has also been greatly simplified by recent software, but the actual operation still requires a level of expertise that may or may not be available in all networks trying to deploy it. As seen with the deployment of DNSSEC, even if such cryptography management expertise is developed in the initial phase of deployment by enthusiastic engineers, as the staff moves on or rotates to other tasks, the level of expertise often fades away and maintaining the processes and systems becomes problematic.

4.1 Resource PKI Scaling

A consequence of the RPKI origin validation design choice to not modify BGP is that BGP itself cannot be used to flood ROAs to validators. Instead, an out-of-band ROA distribution mechanism must be put in place to enable RPKI origin validation validator software to access the data it needs. The chosen model is for validators to synchronize directly with all the ROA repositories. This is a key difference with DNSSEC, where a validating DNS resolver performs validation on demand: it only needs to download (and cache) the necessary information when needed.

There are roughly 65,000 ASNs. If all of them were to participate in RPKI origin validation in the delegated repository model, each RPKI origin validation validator would have to constantly maintain synchronization with 65,000 repositories, knowing that some of them could be offline or unreachable. That means 65,000 x 65,000 connections, about 4 billion connections to maintain globally, multiple times per day. Such a full synchronization might be out of reach – a 2012 study indicates the synchronization time might be over 30 days.[69]

The current protocol to perform this synchronization is *rsync*.[70] *Rsync* was designed to minimize the amount of data transferred - at the cost of significant resources in terms of memory and CPU. This was not a significant problem in the early days of RPKI origin validation, but now that adoption is ramping up, limitations are becoming apparent: RIRs have indicated that synchronization time could be up to 24 hours.[71]

There are different factors that make up this propagation delay. First, the relevant ROA repository has to be updated, and a new version of the repository published. This publication of the repository is a process that happens in batches. The publication interval varies, between a few minutes to a few hours. Then, the RPKI validators need to update their caches. Again, this is a batch process that takes place at specific intervals specified in the configuration of each

69 See https://pdfs.semanticscholar.org/f536/a85c49e8c9754b201b18610fd5b35bd70252.pdf
70 See https://rsync.samba.org
71 See https://www.ripe.net/publications/docs/ripe-549

RPKI validator. Next, the BGP routers have to synchronize with the RPKI validator they are associated with, and this also can take time.

The upper bound of 24 hours is a far cry from what is usually called “Internet speed”. Such a long delay can create critical operational problems if an organization has to make emergency changes to its routing structure. It is fair to say that reducing flexibility in the operation of the routing system risks introducing brittleness into the system.

A new protocol, RPKI Repository Delta Protocol (RRDP), RFC 8182,[72] based on web objects, is being deployed to progressively replace *rsync*. A combination of RRDP with judicious usage of web caches may mitigate (to some extent) the above scaling issue. (Because we are in the early days of RRDP deployment, this promise remains to be verified.) It is too early to fully evaluate if this will hold true, as the scaling issues of every ISP performing this synchronization every hour (or faster) are considerable. As the number of route objects increases and as the number of ASes increases, the scale requirements of the RPKI system also rises.

A current mitigation factor in the scaling issue is that, so far, most networks signing their ROAs have decided to not go for the delegated ROA repository model, but instead rely on the RIR hosted model. For example, in the RIPE region, only two parties have chosen the delegated repository model: NLnet Labs and Randy Bush. The situation is very different in the APNIC region that has extensively deployed the delegated repository model. In security terms this outsourcing of critical security functions presents its own risks, but as considered in Section 2.3, the cost-benefit analysis leans in favor of hosting models for all but the largest and most cryptography savvy ISPs and LIRs.

There are security implications related to those scalability issues. A first example: an RIR is compromised and generates a complete set of allocated address prefixes with AS 0 ROAs, and then revokes all other ROAs. The second: any RPKI publication point is compromised and appears to be publishing individual ROAs for each /128 IPv6 address inside an IPv6 prefix. In both cases, the intent of the attacker is to flood the efforts of relying parties to synchronize with a repository that is expanding to extremely large proportions, causing the entire relying party software to come to a halt.

4.2 Maximum Prefix Length

When a ROA is created, there is an additional parameter - prefix X of length Y can be originated by ASN Z, signed by the owner of X. This added parameter, Y, is the maximum prefix length. Let us look at a ROA, as seen while running the Routinator 3000 RPKI validator:

```
AS16509,100.20.0.0/14,24,arin
```

The third parameter, 24, indicates that any sub-prefix of 100.20.0.0/14 can be originated by AS16509 as long as its prefix length is between 14 and 24 (inclusively). It also says that any sub-prefix of 100.20.0.0/14 longer than 24 should be rejected.

So, for example, 100.20.0.0/16, 100.21.4.0/22, and 100.22.7.0/24 originating from AS16509 should all be accepted, but 100.22.7.4/28 with the same origin AS number should be rejected.

[72] See https://tools.ietf.org/html/rfc8182

The initial recommendation[73] was to set maximum prefix length to match the allocated address block. If you get a /14, set maximum prefix length to /14. This would avoid the potential for more specific route leak, as any sub-prefix would be automatically rejected.

This created problems – for example, when network operators were doing traffic engineering to respond to traffic incidents or balance traffic across multiple links. In particular, during a Distributed Denial of Service (DDOS) attack, a common mitigation is to inject more specific routes to divert part of the traffic. Because the propagation time of ROAs can be up to 24 hours as seen in the previous section, setting the maximum prefix length in a ROA to something too long would make it impossible to respond quickly to incidents.

Because it is impossible to predict which part of the address space would be the subject of a DDOS attack, the current operational practice is to simply set the maximum prefix length to /24, voiding the initial defense it provided. (/24 is the longest prefix length currently routable across the Internet.) However, this exposes the address holder to the original set of route leak scenarios. Because there is no protection of the AS path, the attacker can synthesize more specific prefixes with the ROA-defined origin AS and simply inject these more specific prefixes into the routing system. While the AS path may be longer than the genuine advertisement, the use of more specifics will take precedence and the ROA validation systems will mark these advertisements as valid, so the attack will succeed.

4.3 Protection Provided (or not) by RPKI

RPKI origin validation is often presented as a tool to protect from route incident mis-origination, caused by configuration mistakes or deliberate attacks. This section will analyze the effectiveness of RPKI origin validation to address both.

4.3.1 Involuntary Errors

Incidents involving route mis-origination fall into several categories:

- Typos, also known as “fat fingering”. The syntax of BGP commands in most, if not all, router user interfaces, is fairly error-prone and mistakes are easily (and frequently) made. These mistakes are not always caught by internal checks.
- Configuration errors. Configuring BGP requires a high level of skill and expertise, and again, mistakes are easily made.
- Bugs in software. Some of the most well-known root causes of route leaks come from faulty software, e.g., bandwidth optimizer software leaking routes that should have stayed local out to the Internet.
- Actual malicious attacks.

A general assumption is that the first three types of leaks make up the bulk of the total BGP route mis-origination seen daily. The actual proportion is hard to quantify, as it involves deciding the intent (or lack thereof) of the leak, which is very difficult to establish with certainty. The author could not find any data to quantify the relative proportions of route leaks.

[73] See https://tools.ietf.org/html/rfc7115, Section 3

Depending on the parameters of the published ROAs, accidental route mis-origination may be stopped by ROAs. For example, in the case of leakage of more specifics, the careful use of the ROA maximum length parameter can result in the more specifics being marked invalid and rejected by ISPs that have enabled ROV "reject invalids." However, the current practice seems to not use that maximum prefix length parameter as it would prevent emergency changes to network announcements. So it remains unclear how many route leaks would actually be stopped by RPKI origin validation. A recent study looks into this question.[74]

As exemplified in Section 2.5.1, it should also be observed that RPKI origin validation would not protect against making configuration errors in the ROAs themselves.

4.3.2 Attacks

It is often claimed that ROV can protect against some attacks which rely on BGP leaking to succeed.

One such well publicized attack happened in 2018. Thieves used a BGP hijack to divert traffic being sent to the authoritative DNS name servers of MyEtherWallet.com to a set of name servers under the attacker's control. Unwitting customers of MyEtherWallet were, for a few hours, redirected to a fake website where the thieves stole their credentials. An estimated $17 million was stolen. This attack would have been constrained by ISP deploying RPKI ROV if the prefix for the name servers had been protected by a ROA. It is worth noting that this attack would have also been constrained by ISP deploying DNSSEC validation if MyEtherWallet.com had signed their zone.

However, ROV does not protect from slightly more sophisticated attacks. Route Origin Validation only validates the origin of a BGP announcement. It does not validate the full BGP path of route announcements. An attacker just needs to prepend an AS path to a valid origin AS in order to divert traffic. If the attackers had modified their attack in this way, a BGP attack on MyEtherWallet.com would not have been prevented by RPKI origin validation.

Given how trivial it is to prepend an AS path, one can only conclude that RPKI origin validation can only protect from the most naive attacks but will not provide any significant defense against a deliberate attack. It is worth noting that prepending might reduce the "competitiveness" of the announcement, thus limiting the announcement propagation scope. However, many attacks intentionally do not have a global scale and operate in a confined (e.g., IXP) environment. Such attacks would not be impacted much by this reduced scope.

4.4 The Road to Path Validation

Securing BGP is a complicated endeavor. Doing it fully requires validation of the entire path in route announcements, which has proven to be a very difficult task.

[74] See https://ripe80.ripe.net/presentations/14-3dleak_viz_madory_ripe.pdf

4.4.1 Previous Attempts

A number of attempts at performing path validation have been made (and abandoned) in the past. Among those, we can mention:

- Secure BGP (sBGP).[75] sBGP was originally described in 2000. It offers a comprehensive approach at securing BGP by placing digital signatures on both the prefixes and AS paths included in BGP messages. This approach is CPU intensive on border routers and requires modifications to the BGP protocol. It is also vulnerable to downgrade attacks that can nullify its benefits.
- BGPsec. BGPsec is an evolution of sBGP, standardized by the IETF in RFC 8205[76]. It uses the Resource PKI to store data about AS relationships. It suffers from the same drawbacks as sBGP, as it is CPU intensive on routers, requires modifications to BGP, and is hard to deploy incrementally.
- Secure Origin BGP (soBGP).[77] The soBGP Internet-Draft was last updated in 2006. SoBGP focuses on AS origin validations, just like RPKI, but adds a test of plausibility of the AS Path through certified AS adjacencies. The work on soBGP predated the Resource PKI by many years, and the proposal's designers chose to use a web of trust as its foundation, but this is not a fundamental aspect of soBGP and the Resource PKI could be readily used in place of the web of trust. The IETF worked for many years to resolve the distinction between a rigorous test applied to route object propagation (AS path validation) and the concept of AS path plausibility as used in soBGP. Note: soBGP would also have required modifications to the BGP protocol.
- Inter-Domain Route Validation (IRV).[78] IRV was originally published in 2003. Similar to RPKI origin validation and contrary to the two approaches mentioned above, IRV did not require a modification to the BGP protocol, but uses an out of band validation method. IRV is based on the idea that ASes should be responsible for maintaining a database (similar to the IRR) of all prefixes they originated or provide transit to. A validating router could use the IRV protocol to query the IRV database of the originating AS to perform origin validation. IRV was never fully explored nor analyzed at IETF.

4.4.2 Mitigation Approaches

Path validation remains an elusive goal today. There are partial mitigation measures that can be put in place:

- Increasing the number of direct peering relationships. With direct peering, the path is a single hop, so path validation is logically equivalent to route origin validation.
- Similarly, the shorter the AS path, the less likely an AS path attack is to succeed.
- Bringing the content closer to end users with Content Delivery Networks (CDN) or caches will have an application level effect to shorten the IP level AS path.

[75] See Seo, K., S. Kent, and C. Lynn, "Secure Border Gateway Protocol (Secure-BGP)", IEEE Journal on Selected Areas in Communications Vol. 18, No. 4, April 2000, pp. 582-592.
[76] See https://tools.ietf.org/html/rfc8205
[77] See https://tools.ietf.org/html/draft-white-sobgp-architecture-02
[78] See http://patrickmcdaniel.org/pubs/ndss03.pdf

- Peerlock[79] is a set of empirical practices to AS path filtering that can prevent a number of leaks. For example, it observes that no more than two transit-free (tier 1) networks should appear in an AS path. As such, announcements which contain more than two should be filtered out.

Reducing the AS path will only work if the injection of more specific routes is prevented by using the shortest possible maximum prefix length in the ROAs. However, this creates other problems as seen in Section 4.2, and is in conflict with the current recommended practice to set it to 24.

None of those techniques deployed individually or even collectively can solve the problem of path validation. They might provide some mitigation against some subset of configuration errors and software bugs, however none of these measures are capable of providing any realistic form of security defense.

It then is reasonable to ask: is RPKI origin validation the end of the road? Is path validation out of reach for the foreseeable future?

4.4.3 New Proposal: ASPA

New ideas keep coming in the IETF SIDROPS Working Group on how to tackle the path validation issue. One such idea worth mentioning is the Autonomous System Provider Authorization[80] (ASPA) extension to RPKI.

ASPA leverages some of the key ideas of soBGP, namely AS path plausibility. The critical point here is that there is no need for total and complete route protection – the aim is to constrain the set of undetectable lies in the AS path. The proposed path validation procedure does not require modifications to BGP or the processing of BGP updates, but relies on a new Resource PKI object - Autonomous System Provider Authorization (ASPA) – to create and distribute a database of ISP-level customer-to-provider relationships. An ASPA is the equivalent of a ROA for ISPs. It attests that “a Customer AS holder (CAS) has authorized a particular Provider AS (PAS) to propagate the Customer's IPv4/IPv6 announcements.” Just like soBGP, ASPA does not provide a complete solution to the path validation issue, since it is AS-granularity authorization. Inadvertent route leak with correct AS relationships cannot be protected by ASPA. However, the relative simplicity of ASPA compared to BGPsec may make it more palatable to router vendors and network operators.

Being based on the Resource PKI, ASPA shares some of properties of RPKI origin validation. Its low CPU impact on border routers and its leverage of an existing PKI and management tools can effectively reduce the barrier to adoption. Conversely, ASPA also inherits some of the risks associated with the Resource PKI model, such as the sheer complexity of the X.509 model, the scaling issues of the distributed repository model, and potentially catastrophic scenarios due to a breach of one of the five roots.

79 See https://archive.nanog.org/sites/default/files/Snijders_Everyday_Practical_Bgp.pdf and https://youtu.be/CSLpWBrHy10

80 See https://tools.ietf.org/html/draft-ietf-sidrops-aspa-verification

5 RPKI Operational Risks

The RPKI origin validation system brings a number of operational risks that need to be considered. A taxonomy of these “adverse effects” can be found in RFC 8211.[81]

5.1 Accidentally Rejecting Valid Routes

Route origin validation brings inherent risks of false positives - that is to say, rejecting real routes – and, as a consequence, dropping customer traffic. The ROV logic focuses on finding any ROAs that could cover a prefix. In most cases, corrupted ROAs should not bring false positives. However, there are cases where the absence of a ROA for a more specific announcement of a ROA-covered aggregate, or the presence of a corrupted ROA for a previously uncovered prefix, could lead to false positives.

A number of ISPs and IXPs have tested the “reject invalid” ROV logic and observed that, although the number of BGP announcements deemed invalid by RPKI is still important (about 1%),[82] the actual traffic that flows toward those prefixes is very small.[83] This observation has convinced those operators to turn on automatic announcement filtering. ISPs who have announced their ROV plans include: NTT,[84] AT&T,[85] Telia Carrier,[86] Orange IC,[87] and Seacom.[88] A number of other carriers are in the process of deploying ROV “reject invalid” but have not publicly announced their plans yet. On the IXP side, Amsterdam IX (AMX-IX),[89] DE-CIX,[90] Seattle IX,[91] and Calgary IX (YYCIX)[92] are among the ones deploying ROV “reject invalid”.

It is important to remember that those observations are a reflection of the current state of the deployment of RPKI. It is still in its early stages, and not all prefixes are covered by ROAs. As of 28 February 2020, there were:

- 112,848 IPv4 ROAs covering 104,369 unique IPv4 prefixes. This data is extracted from Routinator. Multiple ROAs may cover the same prefix, either with a different originating ASN (e.g., multihoming) or with a different maximum prefix length.
- 210,754 IPv4 assignments made by the RIRs. This number is collected from NRO delegated-extended statistics.[93]

81 See https://tools.ietf.org/html/rfc8211
82 See https://observatory.manrs.org/#/overview
83 See https://blog.benjojo.co.uk/post/the-year-of-rpki-on-the-control-plane
84 See https://www.us.ntt.net/support/policy/rr.cfm#RPKI
85 See https://mailman.nanog.org/pipermail/nanog/2019-February/099501.html
86 See https://www.teliacarrier.com/Our-Network/BGP-Routing/Routing-Security-.html
87 See https://twitter.com/OrangeIC/status/1233013893771005952
88 See https://www.ripe.net/participate/mail/forum/routing-wg/PDZlMzAzMzhhLWVhOTAtNzIxOC1lMzI0LTBjZjMyOGI1Y2NkM0BzZWFjb20ubXU+
89 See https://www.ripe.net/support/training/ripe-ncc-educa/presentations/use-cases-stavros-konstantaras.pdf
90 See https://www.de-cix.net/en/resources/route-server-guides/rpki
91 See https://www.seattleix.net/route-servers
92 See https://yycix.ca/communities.html
93 See https://www.nro.net/about/rirs/statistics/

- 822,085 IPv4 prefixes announced in BGP.[94] The number of prefixes in BGP has historically always been several times the number of blocks allocated by the RIRs. This is mostly due to the ISP practice of traffic engineering, where prefixes are often de-aggregated and announced as more specifics.
- 18,956 IPv6 ROAs covering 17,322 unique IPv6 prefixes.
- 49,476 IPv6 assignments made by the RIRs.
- 47,074 IPv6 prefixes announced in BGP.

The above statistics show that the global penetration rate of RPKI to the number of BGP announcements is roughly 14% in IPv4 and 40% in IPv6.

Another point to keep in mind is those observations about the small amount of lost traffic due to false positives are made in a best-case scenario. If a major incident were to happen in the RPKI system (see Section 5.5), it is quite possible that the amount of lost traffic would jump. How much traffic would be lost, what damage could result, how this would be detected, and how quickly this would be remediated remain open questions.

Relatedly, another open question to explore: can an ISP rely on data managed by a third party to make critical business decisions, even if the data is cryptographically signed? So far, a number of ISPs (such as ATT, NTT and others) and enterprises (such as Cloudflare and others) have decided that the tradeoff between the potential risk of RPKI failure and the current state of affairs of route leaks, software bugs, and configuration errors is worth the risk of deploying ROV. Other network operators are still waiting for a more compelling case to activate this form of route filtering.

5.2 Self Disconnecting

There is a special risk that an inexperienced engineer using its RIR-managed RPKI portal may face when configuring ROAs for the network. This risk is to only associate the prefix to ASes they have no relationship with. For example, the network engineer could create a single ROA authorizing AS 0 to announce its network prefix. As this ROA will propagate to the various networks deploying RPKI origin validation “reject invalid,” the network prefix will stop being propagated on the Internet. At that point, it might become difficult for the engineer to fix this problem: they might not be able to connect anymore to the RIR servers to correct the configuration.

5.3 Availability Risk: Downtime

Relying parties (networks that perform ROV on their BGP feeds) need to maintain a local cache of all valid published material in the Resource PKI (all certificates, all current CRLs, and all published signed products). This process depends on periodic synchronization with all the Resource PKI publication points. Thus, their constant availability is a requirement. The closer those repositories are to the roots of the Resource PKI roots, the more critical their availability is.

Despite best engineering efforts, downtimes have happened, notable of which would include:

[94] See https://www.cidr-report.org/as2.0/#General_Status

- RIPE-NCC had a prolonged outage on 3 February 2013;[95]
- ARIN had a prolonged outage on 24 October 2018;[96]
- APNIC had an outage on 13 December 2019;[97]
- AFRINIC had an outage on 30 March 2020;[98] and
- RIPE-NCC had four outages in early 2020 (the most publicized one was on 6 April 2020).[99]

Two of the recent RIPE-NCC outages are worth exploring for the lessons that can be derived. The root cause of the 22 February 2020[100] outage was when an unmonitored disk partition exceeded its quota. It resulted in one part of the RIPE repository, the Certificate Revocation List (CRL), to expire. This expiration triggered a cascading effect for several relying parties. For some, the RPKI validator software invalidated the whole repository. For others, their validator software either did nothing or simply gave a warning. That incident triggered a discussion in the SIDROPS IETF Working Group on what the correct behavior should be for validator software. The root cause of the 6 April 2020 outage was a system integration issue where one of the components made a query to another system that was under a maintenance window. This resulted in a loss of data that took several hours to recover from.

In general, relying parties are expected to use their local cache in the event of unreachability of an RPKI publication point. In detail, the situation is more complicated. Unavailability that extends past the next update time of a CA's CRL poses a dilemma for the relying party. Should they continue to apply the outdated locally cached CRL revocation information with the knowledge that the information is out of date, or declare the entire local cache of this publication point invalid? (There is a parallel in the DNS world discussed in RFC 8767.[101]) The CRL and manifest staleness should not be confused with data expiration. Staleness means that the publisher of the data has told you to expect another update by now, but it has not arrived. It does not automatically make any data invalid. The situation is different for locally cached certificates that have expired. They should not be used in the RPKI validation process.

Short outages on the order of hours should not present operational issues as long as CAs use a certificate refresh practice that refreshes a certificate well before its expiration time. Protracted outages that extend over one or more days may present more visible operational issues. In most cases, the consequence would simply be that some prefixes would no longer be covered by a ROA. In the validation process, they would be marked as state "unknown." ISPs filtering on "invalid" would see no consequences. However, there are cases where it could be a problem. An example is the case where a prefix is covered by a ROA with a matching maximum prefix length option, and ROAs that exist to cover longer sub-prefixes. If the more specific prefix ROAs disappear but the aggregate prefix ROA remains, the routes for the more specific will be marked "invalid" and thus rejected. Another circumstance is where a relying party (inadvertently or otherwise) ends up filtering prefixes "NotFound" in the RPKI system. See Section 2.4 for a discussion on rejecting "NotFound".

95 See https://www.ripe.net/support/service-announcements/service-announcements/ripe-ncc-rpki-repository-outage
96 See https://www.arin.net/vault/announcements/2018/20181024.html
97 See https://www.apnic.net/about-apnic/service-updates/service-announcement-13-december-2019/
98 See https://lists.afrinic.net/pipermail/rpki-discuss/2020-March/000108.html
99 See https://www.ripe.net/support/service-announcements/rsync-rpki-repository-downtime
100 See https://www.ripe.net/ripe/mail/archives/routing-wg/2020-February/004015.html
101 See https://tools.ietf.org/html/rfc8767

Given the above, the process of operational management of RPKI publication points should be regarded as an operationally critical activity, as should the proper utilization of RPKI data by relying parties.

5.4 Consistency Risk: Five or More Trust Anchors

The “overlapping” trust anchor set of five self-signed certificates, each of which lists the entire IP resource set, does lead to a potential situation where a misconfiguration could result in inconsistent information being published by two (or more) RIRs. It also implies that if a successful attack was mounted against a single RIR’s Resource PKI publication system, the attacker could publish Resource PKI material for any number resource, not just the number resources under the purview of the RIR.

Given the operational reliability of the RIRs, this appears to be an exceedingly unlikely situation, however, one special case for concern is what would happen if an RIR was presented with a court order or other government mandate that required the RIR to unilaterally update its database - for example, with the goal to take down a network in a country.

The potential consequences of this consistency risk are analyzed in the next section.

5.5 Integrity Issue: Breach

A more worrisome scenario is a security breach at one of the RIRs.

5.5.1 A Potential Failure Scenario

This scenario could happen as a result of a software bug, a compromised employee, or a direct attack on the RIR infrastructure. Here is an example catastrophic failure scenario to contemplate:

- Addresses used by a provider of critical service/infrastructure (party A) are covered by ROA R.
- Major Tier 1 transit ISPs deploy automatic ROV.
- The RIR holding ROA R is a victim of a security breach and the attacker takes control of the RPKI RIR repository.
- Attacker wants to inflict harm to party A.
- Attacker revokes the valid certificate that was used to sign ROA R and issues a new valid certificate and signs a new ROA, R2, for a different (potentially nonexistent) origin ASN.
- Change propagates to relying parties and may not be detected for a long period of time, due to the lack of systematic monitoring
- RPKI ROV keeps functioning without raising alarms (all other certificates look OK).
- All ISPs using automatic ROV drop route announcement from party A.
- Party A becomes inaccessible from most parts of the Internet, potentially impacting systems that depend on party A.

This situation is not easy to detect or deal with. Besides the hacked prefix, the rest of the ROV/RPKI infrastructure keeps working as usual. It is unlikely that the problem would be noticed before damage is done.

5.5.2 Recovery

Once the issue is discovered and brought to the compromised RIR's attention, that RIR will fix the security breach and issue new certificates and new ROAs. This process will certainly take time. The RPKI propagation time of 5 minutes to 24 hours will certainly be an aggravating factor. It is probable that all ISPs would turn off ROV during that period.

In a variation of this scenario, because of the five RIR RPKI roots, any breach at any RIR could have consequences for a covered ROA. An attacker taking control of the Resource PKI system of an RIR other than the one issuing an IP address block could still do harm to the corresponding prefix. Validator software usually performs a logical OR between all ROAs, so a duplicate but contradicting ROA should not directly take down any network. However, it might allow a route leak to happen. In particular, it could allow an orchestrated route "leak to nowhere" to effectively take a network down.

5.5.3 Variations

Another scenario can also impact an organization that is not using RPKI and has no ROA. A breach at any RIR could allow an attacker to take such an organization off the Internet, simply by the attacker issuing a bogus but valid ROA, associated with ASN 0, for the target prefix. In this case no injection of false information into BGP is required. The addition of the bogus ROA labels the existing authentic BGP route object as invalid.

Any breach at any RIR can cause harm, with some variation, to any prefix, covered by a ROA or not. The situation is summarized in the following table:

	Breach at RIR issuing the prefix	Breach at another RIR
Prefix covered by a ROA	Takedown of prefix	Open to route leaks via more specifics that could result in a partial or a complete takedown of the prefix
Prefix not covered by a ROA	Takedown of prefix	Takedown of prefix

This situation can be compared to what would have happened with a single root to the system. A breach at the root could take down any prefix, but a breach at any RIR could only impact the prefixes it assigned/allocated.

	Breach at the root	Breach at RIR issuing the prefix	Breach at another RIR

Prefix covered by a ROA	Takedown of prefix	Takedown of prefix	Nothing
Prefix not covered by a ROA	Takedown of prefix	Nothing	Nothing

There is, however, another point to observe: if an RIR went rogue, was compromised, or was subject to a court order, the five trust anchors system might offer a path to recovery. The other four RIRs could start issuing CAs to affected parties so they could sign their EEs and thus sign their ROAs. At that point, the relying parties could remove the TAL of the affected RIR. This would partially mitigate the problem, but would certainly take time to implement.

To summarize, the decision to have each of the five roots claiming coverage for the entire IP address space has increased both the attack surface and the consequences of breach at any of the RIRs. Conversely, it offers a potential, although very slow, path to recovery in case of a breach.

5.6 RPKI-Induced Loss of Reachability of RPKI Repositories

If the ROA distribution mechanism were to be attacked, or the ROAs received by a relying party were to be corrupted in transit, it is possible that the routes to one of the RIRs managed ROA repositories (or a delegated repository under them) would become invalidated, and as a consequence, the repository would become unreachable.

If an RIR ROA repository were to become unreachable, the problem could probably be detected quickly. However, it might take more time to detect (and fix) a problem with a delegated repository further down the chain. (Section 2.5.3 gives an example of stale data from a delegated repository.) In the absence of a monitoring system that sweeps the publication points for “liveness,” none would know for sure if this scenario had already occurred or not. It is also unclear how to recover operationally from such a situation. A manual intervention might be required to restore the lost route. Such an intervention is clearly within the scope of what a savvy operator (well versed in managing cryptographic systems in general and RPKI in particular) can do, but a less experienced one might have more difficulties dealing with the situation.

5.7 Routing Police

A key question is what happens when an RIR member becomes delinquent and stops paying its registration bills. Prior to RPKI, the RIR would reclaim the IP address blocks and invalidate the associated reverse DNS entries. This may or may not have an immediate impact on the operation of the delinquent member network. However, with RPKI, the RIRs have a much bigger stick. They can now invalidate the CAs for that member, invalidating all its ROAs; for example, RIPE already has.[102] This is still of limited effect as the current expected behavior is to accept “RPKI NotFound” routed.

[102] See https://www.ripe.net/publications/docs/ripe-716

However, if policies like the AS 0 one described in Section 3.3 are implemented by the delinquent member RIR, new ROAs tied with AS 0 will be created for the corresponding prefixes. The effect would be to take the delinquent member network off the Internet within the 24 hour Resource PKI propagation delay window. To correct that situation, the network operator will have to use an out-of-band mechanism to contact its RIR, as its network connectivity will be suspended.

One of the famous claims about the Internet is that there are no routing police. RPKI origin validation is, to some extent, thrusting the RIRs partially into that role. At minimum, it is giving the RIRs a new, critical role in the day-to-day active operation of the routing system. RIRs may or may not be fully prepared for that role.

5.8 RIR Support Model

This new role imposed by RPKI carries high-risk exposure to the RIRs, thus the level of operational support provided by each RIR is of critical importance. The RIRs do offer 24/7 contact support for the systems hosting the five Resource PKI roots, however, how to contact an RIR for emergency RPKI issues is not listed on the RIR's "contact us" page. In all cases, a more typical Monday to Friday, 7 to 7, or 9 to 5 (in the time zone local to the RIR) support is advertised. Contact information for the RIRs:

- AFRINIC: https://afrinic.net/contact
- APNIC: https://www.apnic.net/about-apnic/organization/contact-apnic/
- ARIIN: https://www.arin.net/contact/
- LACNIC: https://www.lacnic.net/630/2/lacnic/contact-us
- RIPE-NCC: https://www.ripe.net/support/contact/technical-emergency-hotline

In a recent development, ARIN is now including a link to "Report Service Issue" at the bottom of each page on its web site.

Although the RIRs are in a natural position to assert information relative to address block holders, it may be that the RIR communities would need to evaluate whether they are willing to ensure their RIR has the resources and remit necessary to ensure their RIR is the appropriate place to operate the roots of a routing control system.

5.9 SLURM

Beyond the simple use case to handle local private networks introduced in Section 2.4, Simplified Local Internet Resource Management with the RPKI (SLURM), RFC 8416[103] was developed with a larger goal: create a local override capability to protect routes from adverse actions that are described in RFC 8211.[104]

An administrator at a relying party can create SLURM files expressed in json format to list a number of local exceptions to override data coming from the Resource PKI. This file is to be processed by the validation software when creating BGP announcement filters.

[103] See https://tools.ietf.org/html/rfc8416
[104] See https://tools.ietf.org/html/rfc8211

The first question to ask is how would the administrator know what exactly to put in the SLURM file beyond the local knowledge information mentioned in Section 2.4. The operational model in case of a major Resource PKI incident is unclear. Would a SLURM file be distributed? By whom? How would it be authenticated? How would it be revoked? Those questions call for further deliberations on the SLURM deployment guide and operation model if the community wants to make use of SLURM on a larger scale.

Several SLURM files from different sources might have to be combined. RFC 8416 Section 4.1 explains that the application of a SLURM file must be atomic, i.e., all assertions included in the file must validate for it to be taken into consideration. Section 4.2 explains that there should be no conflicts between different SLURM files, and if there are, the entire set of SLURM files must be discarded. Those two considerations will ensure the integrity of the local database of assertions, but can create operational challenges, making the use of SLURM brittle.

6 RPKI Liability Issues

RPKI relying parties may or may not be members of one RIR or more, but few relying parties are members of all five RIRs. However, all relying parties need Resource PKI data from all five RIRs to perform a complete RPKI origin validation ROV. Thus, the RIRs find themselves in a position where they need to offer services to non-members with whom they do not have prior formal relationships.

Offering services to nonmembers is not something new for the RIRs. For example, WHOIS data has been published for many years by all RIRs. What is new is that with RPKI ROV, RIR Resource PKI data is now directly used by ISPs to make filtering and routing decisions and the RIRs are now embedded in the critical path of the routing system. As a result and due to this real-time dependency, the stakes are much higher than with the publication of WHOIS data. If a RIR makes a mistake, gets hacked, or experiences a prolonged outage, there could be significant damages to third parties. Similarly, a relying party using Resource PKI data outside of the current best practices (for example deploying a policy to reject RPKI “unknown” route announcements) and suffering a direct or indirect outage could sue the RIR originating the data. In the North America region, ARIN has recognized this situation as a major liability risk and an existential threat to the ARIN organization that must be mitigated. The chosen mitigation approach is to require all relying parties using ARIN RPKI data to indemnify ARIN for their use of ARIN’s RPKI data.

6.1 ARIN Specific Legal & Technical Approach

ARIN has put in place a series of mechanisms to protect itself from liabilities due to an RPKI failure.

6.1.1 ARIN Relying Party Agreement

ARIN has created a Relying Party Agreement (RPA)[105] that must be accepted by anyone using the ARIN TAL to perform RPKI validation. To force users to agree to the RPA, ARIN TAL is not available freely, rather a user has to navigate to a web page in a process called "browseware."

According to a study conducted by Christopher S. Yoo from the University of Pennsylvania Law School,[106] the "browseware" process is an effective legal tool to protect ARIN from their perceived liability and may be the only practical way to do so. However, the same study concludes that ARIN's RPA creates a significant barrier to entry to RPKI. Relying parties may or may not be willing to indemnify ARIN for their use of RPKI data. Even if they were willing, they may face internal challenges making it difficult and sometimes impossible to agree to the RPA. (For example, the above paper mentions that government entities might be prohibited from agreeing to indemnification, arbitration, and choice of law clauses. Non-government entity network engineers trying to deploy RPKI would have to get their legal department to review, understand and approve the RPA.) This last step may not be a show stopper, but a hurdle that weighs on the balance. The study observes the net result is that the ARIN region has the least amount of RPKI deployment among the five RIRs.

Recent data confirms this is true on the Route Origin Validation side. On 2 February 2020, one RPKI monitor[107] only listed a few entities in the North America region using RPKI ROV. However, ROA signing does not require agreeing to ARIN's RPA. As seen on a RPKI validator on 28 February 2020, there were:

	ARIN	RIPE	APNIC	LACNIC	AFRINIC	Total
IPv4 ROAs	7,494	68,674	29,440	6,399	841	112,848
Unique IPv4 prefixes covered by ROAs	6,493	63,436	27,679	5,937	824	104,369
IPv4 address blocks allocated	64,235	81,368	44,097	17,321	3,733	210,754
% of IPv4 allocated prefixes covered by ROAs	10%	78%	63%	34%	22%	50%

105 See https://www.arin.net/resources/manage/rpki/rpa.pdf
106 See https://scholarship.law.upenn.edu/cgi/viewcontent.cgi?article=3037&context=faculty_scholarship
107 See https://rov.rpki.net

	ARIN	RIPE	APNIC	LACNIC	AFRINIC	Total
IPv6 ROAs	1,348	9,824	6,150	1,487	147	18,956
Unique IPv6 prefixes covered by ROAs	1,258	9,412	5,166	1,340	146	17,322
IPv6 address blocks allocated	7,117	21,190	10,207	10,032	930	49,476
% of IPv6 allocated prefixes covered by ROAs	18%	44%	51%	13%	16%	35%

Based on this data, it can be argued that there is simply a general lack of interest for RPKI in the ARIN region compared to the RIPE and APNIC regions, regardless of ARIN's RPA liability concerns.

The other observable result of ARIN RPA is that most RPKI validation software does not come with the ARIN's TAL preconfigured. Relying parties have to make an additional manual step to agree to ARIN's RPA, download ARIN's TAL, and configure it on the validation software. As a result, a number of relying parties only use data from the other four RIRs.

ARIN does allow validator software vendors to have installation tools that obtain and install the ARIN TAL once the user has agreed to the RPA. However, some software vendors see this as putting them in a situation to indemnify ARIN on behalf of their users and they are reluctant to take on this liability. Both NLnetLabs Routinator 3000 and RIPE Validator still require an additional manual step to install ARIN's TAL.

6.1.2 ARIN's Non-Repudiation Approach

With the exception of ARIN, RIR members adopting the hosted RPKI model can simply log in to their account and generate a ROA. The RIR holds the certificates to sign those ROAs on their member's behalf. ARIN designed their hosted RPKI system differently. ARIN does not keep the keys to generate ROAs on a member's behalf. ARIN members need to log in to ARIN's portal with SSL using their own private keys and then generate their ROAs using their own key. The absence of member keys on ARIN's server provides non-repudiation to the ARIN organization: even if several staff members at ARIN were compromised, they could not sign ROAs on a member's behalf. Similarly, this prevents ARIN from creating ROAs even if ordered by a court or national government. The flip side is that users need to manage their own public/private key pair to log in to the system (as opposed to simply using a password like in the other RIRs). This

means creating ROAs is more complex in the ARIN region than in the other regions. More importantly, while ARIN's approach does reduce the risk that an insider at ARIN could create unauthorized (by the resource holder) ROAs, this diversion in operational practices creates increased cost and complexity for users. It is not clear if this approach will be retained in ARIN's forthcoming IRR and RPKI refresh coming in late 2020. ARIN indicated in a private communication: "At present, we have a fairly challenging user interface for our hosted RPKI, and this predominantly stems from not ever generating/requesting/storing the RPKI party's private key used for signing requests. Our community has repeatedly requested a very easy ROA generation interface integrated with ARIN online. It might not be possible to meet this ease-of-use requirement without storing the request signing key, and thus, maintaining non-repudiation."

6.2 Other RIRs Perspectives on Liability

The other four RIRs do not appear to have expressed the same level of concern, or rely on an implied agreement for usage. In particular, RIPE-NCC terms and conditions for accessing RPKI data says that "The RIPE NCC is in no way liable for any direct or indirect damages"[108] but does not require explicit indemnification. Some of experts who were interviewed expressed the viewpoint that risk mitigation via transparency and applying current engineering best practices is a better approach than indemnifications.

It should be observed that discussions about liability are happening within each RIR community. As mentioned earlier, relying parties may or may not be members of any particular RIR. As such, there is currently no forum for the relying parties to discuss these liability issues in a global context.

7 Covering Routes to Critical Infrastructure with ROAs

ARIN's Relying Party Agreement[109] states:

> *"You acknowledge and agree that neither the [Online Resource Certification PKI] Services (or any part thereof) nor the Certificate is designed, intended, or authorized for use in connection with equipment in hazardous circumstances or for uses requiring fail-safe performance, including uses in connection with the operation of nuclear facilities, aircraft navigation or communication systems, air traffic control systems, or weapons control systems, where failure could lead to death, personal injury, or severe environmental damage."*

Besides those explicit restrictions, one could wonder if RPKI should be used to cover routes to critical Internet infrastructures. In the ICANN world, the question is: should DNS root servers' prefixes be covered by ROAs? Similarly, should DNS top-level domains (TLD) authoritative servers' prefixes be covered by ROAs?

[108] See https://www.ripe.net/manage-ips-and-asns/resource-management/certification/legal/ripe-ncc-certification-repository-terms-and-conditions

[109] See https://www.arin.net/resources/manage/rpki/rpa.pdf

The first observation is that there are thirteen root servers. If one were taken down by a BGP route leak, the reasonable assumption is that the other twelve would take up the additional load. Root server incidents might happen, and managing BGP announcements of the prefixes covering the IP addresses of those servers when an incident happens is an important issue, as seen in a recent incident.[110] However, it is unclear this is an area where RPKI origin validation could help much. This observation also applies to a certain degree to TLD authoritative servers.

The second observation is about the origin of the addresses used by the thirteen DNS root servers. Nine of those prefixes were allocated by ARIN, two from RIPE NCC, and one from APNIC. A catastrophic failure (such as the one described in Section 5.5) at ARIN would have disproportionate consequences. A possible mitigation would be for some of the root server operators to transfer their prefixes to other RIRs to rebalance the risks.

The third observation, also specific to DNS, is that there is already a mechanism put in place to secure the system: DNSSEC. DNSSEC is based on data integrity, not transport integrity. Which root server instance is used does not matter; what matters is that the root DNS data validates. The trade-off then becomes extra security by combining different security mechanisms at different levels versus extra complexity. Given that RPKI origin validation can only protect against naive routing system attacks, the cost-benefit tradeoff is not clear.

In light of just these three observations, it can be difficult to make a strong case for securing the routes to DNS root servers with ROAs.

Generalizing this discussion to sign (or not) ROAs for critical infrastructure, not just DNS root servers, one must keep in mind the different RPKI risks described in Section 4. As some Tier 1 ISPs are already doing ROV, critical infrastructure operators will have to be prepared to deal with the consequences of security breaches at any of the roots of the Resource PKI, regardless of their decision to sign ROAs or not.

Because it appears from Section 5 analysis that, in case of a breach at a different RIR than the one originating the address block, a prefix is marginally better off with a ROA than without, the actual recommendation to sign ROAs for root servers moves slightly into positive territory.

8 Conclusion

There is significant interest in RPKI, driven by the RIRs and network operators, large and small. Many parties believe there is enough low-hanging fruit with RPKI that the return on investment is positive. Signing ROAs has now been made simple enough that pretty much any IP address holder can do it and RPKI origin validation offers protection against fat fingering, configuration mistakes, and software bugs. Although RPKI origin validation does not protect from non-naive attacks on the routing system, from an operator perspective, both attacks on the routing system and route leaks caused by fat fingers generate tickets that must be dealt with. Any help that RPKI origin validation will provide on that front will certainly be welcome by many ISPs.

However, the overall system, which is based on X.509 certificates, is complex. This complexity introduces the risk that new mistakes, typos, and fat fingers will find their way into the Resource PKI itself. Strong organizational expertise in cryptographic system management will likely

[110] See https://www.isc.org/docs/f-root/incident-2020-01.pdf

remain a prerequisite to turning on ROV. RPKI itself does not come without issues. The propagation delay that can be up to 24 hours, compounded by the lack of widespread systematic monitoring, can be a major operational issue. Also worth noting is that, on top of not addressing all aspects of the routing security problem, RPKI origin validation can introduce new threats to the routing system, such as in case of attacks to the Resource PKI repositories, the various certificates, or the ROAs distribution systems. To date, RPKI origin validation ROV has only been deployed in a limited scale. There are still unanswered questions relating to the scalability of the overall system.

Ultimately, whether the cost of the rather detailed infrastructure and operational complexity of RPKI origin validation is worth the value of the benefit in terms of routing integrity is a determination that will need to be made by network operators. Some network operators worried about the impact on their operations of misconfiguration-induced route leaks clearly believe that this is the case, while others concerned about routing security have not yet been convinced. Perhaps more importantly, RPKI does imply certain changes to critical operational structures of the Internet as a whole. It remains unclear whether the communities involved and impacted by those changes are fully aware of those implications. Further work in communicating the implications of RPKI is clearly warranted.

9 Acknowledgements

While all opinions in the report are those of the author, we would like to acknowledge the following persons who provided input, feedback or reviews during the development of this report:

- Alain Aina, WACREN
- Rob Austein, Hacntr
- John Curran, ARIN
- Kim Davies, ICANN (IANA)
- Geoff Huston, APNIC
- Fredrik Korsback, Amazon
- Nathalie Künnake-Trenaman, RIPE NCC
- Martin Levy, Cloudflare
- Di Ma, ZDNS
- Terry Manderson, ICANN (DNS and Network Engineering)
- Carlos Martinez, LACNIC
- Christopher Morrow, Google
- Ricardo Patara, NIC Brazil
- Amreesh Phokeer, AFRINIC
- Andrei Robachevsky, ISOC
- Job Snijders, NTT
- Bill Woodcock, PCH

Special thanks go to David Huberman, ICANN, for his constant support and willingness to act as a sounding board while writing this document.