Building a new governance hierarchy:

RPKI and the future of Internet routing and addressing

7 September, 2010

Drafters:

Milton Mueller
Professor, School of Information Studies, Syracuse University
XS4All Professor, Delft University of Technology

Brenden Kuerbis
School of Information Studies, Syracuse University

Concurring:

Michel van Eeten
Professor, Delft University of Technology

Abstract: *This paper attempts to broaden awareness of routing as a site of Internet governance. It sheds light on the historical processes and policy issues raised by the resource public key infrastructure (RPKI). The paper will explain why the RPKI has important governance implications for, 1) the autonomy of ISPs, 2) the centralization of institutional power and global compatibility, 3) the business models of the RIRs and their relationship to ICANN, and 4) the role of governments in Internet governance.*

1. Introduction

Routing and addressing are at the core of how the internet works. Every second, routing arrangements must be able to successfully move trillions of individual data packets from any originating network in the world to any one of millions of destinations anywhere in the world. Routing and addressing are more than just technical-operational processes. They are important forms of Internet governance.

Major changes are underway in routing and addressing policy. They are caused by a combination of factors: the perceived need for greater security in routing, the depletion of the IPv4 address space, and the attempt to migrate to a new Internet standard (IPv6). One of the most significant changes is an attempt to create a Public Key Infrastructure for Internet protocol addresses and routes. Resource Public Key Infrastructure (RPKI) is a security technology that would create a hierarchy of digital certificates which would be used to authenticate the allocation of address blocks and route announcements using those blocks.

RPKI could have a profound effect on the future of the internet. It will, at the very least, challenge the independence of Internet service providers, alter the governance role of the regional Internet address registries (RIRs) and change the relationship between IANA and the RIRs. At best, it can prevent deliberate and unintentional routing problems and enable a more efficient and flexible use of address resources. At worst, it could greatly centralize power over the internet, creating political struggles to seize or capture the leverage over Internet users and suppliers it produces. Yet public awareness and discussion of RPKI is practically nonexistent. Whatever dialogue exists is confined to a few technical circles, many of them obscure from public view. Even within the technical community, discussion of their implications for Internet governance has barely begun.[1]

[1] We do not wish to imply that RPKI issues have been deliberately kept secret, or that there is a menacing conspiracy underway. Routing/RPKI policy is under-appreciated for the same reason that many other internet governance issues are: analysis of it requires a rare combination of technical, operational, institutional, economics and policy knowledge about the internet. The people who understand the technology standards and operations rarely have expertise in public policy, institutional design, or

This paper attempts to broaden awareness of routing as a site of Internet governance. It focuses in particular on the policy issues raised by RPKI. The paper will explain why RPKI has important implications for the accountability and freedom of internet users and the efficiency, openness and security of Internet service.

The paper's primary purpose is informational; it does not advocate a particular solution or policy. It will describe RPKI and the current state of its development and implementation. It will also describe the fascinating debates held during the standardization and implementation processes. These debates make it clear that important policy and governance issues are raised by RPKI implementation strategies. The purpose of this discussion is to give industry actors, government policy makers and civil society public interest advocates a better sense of what the issues are, where things stand, and thus where their intervention might be fruitful.

2. Routing and Security

Routing is the automated process that directs Internet protocol packets from their origin to their destination. IP addresses can be described as part of the language that routers speak to each other. Internet routing protocols consider the IP address to be composed of two parts: the address of the network (the *prefix*) and the address of the connected computer (the *host*). Routing through the Internet is based on the network portion of the address. For each prefix, a router stores information telling it how to find a path to it and uses this information to construct a forwarding table (the routing table) that controls the movement of each incoming packet to the next hop in its journey. Routers also transmit announcements to other routers about the address prefixes to which it is able to deliver packets, and this information is incorporated into the tables of other routers.

Thus, routers are engaged in constant, automated conversations with each other that exchange network prefixes and other routing policy information to keep every router informed about how to reach tens of thousands of other networks on the Internet. Currently, interactions among routers are based on an Internet standard known as Border Gateway Protocol (BGP). As originally described in RFC 1771 (1995), and as later updated by RFC 4271 (2006), BGP is the dominant inter-domain routing protocol of the Internet. (Rekhter and Li, 1995; Rekhter et al, 2006)

In the original BGP protocol, all routers in all Autonomous Systems (ASes) were assumed to be trustworthy. As the Internet grew, the assumption of ubiquitous trust made less and less sense. (Hu, McGrew et al, 2006) The existence of malicious actors on today's internet is a given. But the Internet's routing infrastructure is also vulnerable to unintentional misconfigurations that can cause harmful results. (ENISA, 2010; Barbir, Murphy, & Yang, 2006) Extensive work has been done in the technical community exploring the issue of routing security and proposing various solutions to improve it.

economics. The people with expertise in economics and public policy rarely have a deep understanding of the technology and operations. Even for the experts, RPKI and its cost-benefit trade-offs are complicated. Moreover, the stakeholders who are already involved in negotiating solutions and approaches – such as the address registries, US military research contractors, the ISPs, equipment vendors and so on – have little incentive to foster a critical and wide-ranging dialogue about the governance implications of their actions. Their incentive is to avoid politics as much as possible and come to an agreement that rocks the boat as little as possible.

(Butler, Farley, McDaniel, & Rexford, 2010) One of the security flaws in BGP was illustrated vividly in 2008 when a Pakistani ISP's attempt to block YouTube affected ISPs around the world, effectively knocking YouTube off the Internet for most users for an hour. Several other well known misconfigurations leading to temporary routing outages have occurred in the past, although the overall extent and severity of the routing security problems network operators' deal with is not publicly known.

Some assessments of this problem are more alarmist than others. Some observers ridicule the existing state of affairs as "routing by rumor." (Internet Architecture Board [IAB], 2010) One prominent internet veteran has suggested that unless major changes are made, a major breakdown will occur and the Internet community will be "crucified" in the press as Toyota Motors was for its alleged throttle malfunctions.[2] Other voices are less alarmed. They note that a variety of measures are already in place by ISPs to filter out false route announcements. They claim that the same network operators who don't currently apply filtering of BGP announcements won't deploy a PKI. A major breakdown such as the Pakistan case, they claim, applied only to one site and was remedied in the space of an hour. They claim that routing takes place reliably in the vast majority of cases, and some believe that measures as extreme as RPKI are not needed, or impose costs that exceed the benefits.

3. RPKI as a proposed solution

RPKI uses digital resource certificates to authenticate the assignment of IP address blocks and Autonomous System (AS) numbers. (Kent, 2006) These certificates bind a resource holder to its public cryptographic key, as well as information about the IP address block prefixes and Autonomous System (AS) numbers allocated to them. Subsequently, resource holders can create route origin authorization (ROA) statements, or standardized verifiable attestations that the holder of a certain prefix authorizes a particular Autonomous System (AS) to announce that prefix. Using this information, other network operators (e.g., ISPs) can validate that 1) a specific network, as indicated by a unique AS, is the legitimate holder of an IP address block, and 2) the AS that originates a route announcement using a particular prefix is authorized to do so. Like all PKIs, authenticating certificates therein (and subsequently the associated allocation and routing information) would rely on the system having one or more Certificate Authorities (CAs)[3], which would publish a public key(s) or "trust anchor" to be used to authenticate other certificates.

[2] See presentation, *The RPKI & Origin Validation*, by Randy Bush to RIPE NCC on 3 May 2010, available at http://www.ripe.net/ripe/meetings/ripe-60/presentations/Bush-The_RPKI_Origin_Validation.pdf. Ironically, it has since come to light that driver error, not the car, was responsible for many of these problems.

[3] Most users are familiar with digital certificates through their use of Certificate Authorities (CAs) for web sites. CAs are third parties who are trusted by the subject (publisher) of the certificate and the parties interacting with the subject who rely upon the certificate for authenticating it (the relying party). This allows relying parties to rely upon signatures or assertions made by the private key that corresponds to the public key that is certified. Many private sector companies offer CA services commercially. Government agencies may also act as CAs, or organizations can set up their own, internal CA. A 2009 market share report determined that VeriSign and its acquisitions (which include Thawte and Geotrust) have a 47% share of the certificate authority market, followed by GoDaddy (23%), and Comodo (15%). See Wikipedia http://en.wikipedia.org/wiki/Certificate_authority

The critical feature of the proposed RPKI solution is the attempt to link resource certificates to the authoritative sources of internet resources, namely ICANN and the RIRs. This could fundamentally change their governance role. As noted by a member of the IETF's Internet Architecture Board (IAB), RPKI could be used to give the institutions that control IP address resource allocation operational control over what is routed (and, therefore, what information is accessible) over the Internet.[4]

4. How we got here: the historical process

An RPKI aligned with the existing IP address allocation hierarchy is not the only available approach for making routing more secure. Yet in the past 5 years it has emerged from the IETF and the RIRs as the dominant approach. Despite the many uncertainties associated with governance arrangements for RPKI (see section 5), serious and possibly irreversible implementation steps are already being taken. How and why did the technical community converge on such an implementation of RPKI? The basic answer is that the U.S. government, acting through contractors with Defense Department and DHS grants, pushed a solution that was well aligned with the incentives of the extant Internet governance institutions, just as it has actively promoted DNSSEC as the chief mechanism for securing the domain name system.

The idea for a RPKI emerged in the mid-1990s, from work funded by the National Security Agency and DARPA to address routing security problems. (Kent, Lynn, & Seo, 2000) In 1999 individuals affiliated with U.S. government contractor BBN Technologies first authored Internet drafts proposing to use X.509 extensions (digital certificates) to authenticate IP addresses, AS identifiers and BGP announcements.[5] BBN's specification required the initial development of two PKIs “to verify the identities and authorization of BGP speakers and of owners of ASes and of portions of the IP address space.” (Seo, Lynn, & Kent, 2001)

The first phase of IETF work around routing security, starting in 2002, developed basic requirements. The next phase of standards activity began in late 2005, when the IETF turned towards building “an extensible architecture for interdomain routing security.”[6] The Secure Inter-Domain Routing (SIDR) Working Group (within the IETF Routing Area) was initiated in November 2005, and chartered in February 2006. It was co-chaired by Geoff Huston, Chief Scientist of the Asia Pacific regional Internet registry (APNIC) and Sandra Murphy of SPARTA. SPARTA was awarded a Department of Defense contract in 2005, the objective of which was to “ensure the security of the Internet routing infrastructure so that it is reliable in the event of a deliberate malicious nation-state level attack from adversaries.” (U.S. Army Research, Development and Engineering Command [U.S. Army RDECOM] 2005) SPARTA’s proposal said that it would “progress the address and AS allocation PKI specifications through the IETF process,” “design an architecture for the distribution of the address and AS certificates,” and “identify *acceptable and willing* entities for the roles of any data repositories in the distribution

[4] See D. McPherson, “IPv4 Exhaustion: Trading Routing Autonomy for Security.” Arbor Networks blog, http://asert.arbornetworks.com/2008/03/ipv4-exhaustion-trading-routing-autonomy-for-security/

[5] They built on previous work done within the IETF's Public-Key Infrastructure (PKIX) working group (within the IETF Security Area), which had developed the standards for the widely-used X.509 public key infrastructure. The PKIX Working Group was co-chaired by BBN's Chief Scientist for Information Security, Steve Kent.

[6] See http://tools.ietf.org/wg/sidr/charters?item=charter-sidr-2006-07-03.txt

architecture." (SPARTA Inc., 2006, emphasis added) While Murphy would assume an influential leadership role in the SIDR Working Group, numerous Internet-Drafts in the group would also be authored by individuals from BBN Technologies, which had received additional contracts from the US government to develop RPKI software. There were also contributions from individuals affiliated with APNIC, Internet Systems Consortium and the Internet Initiative Japan.

The SIDR working group produced an architectural specification in which the Public Key Infrastructure for validating address holders, AS numbers and route authorizations was linked "to the existing resource allocation structure." This was the critical design decision that made RPKI design and implementation raise so many interesting internet governance issues. The SIDR working group asserted (not entirely accurately) that "existing resource allocation and revocation practices have well-defined correspondents in this architecture."[7] Following this logic, the architectural specification called for the IANA to issue RPKI certificates for the IP address and ASN resources for which it was authoritative.

Within the SIDR working group, there was long debate over the governance implications of linking RPKI to the existing resource allocation structure (see section 5 below for a more detailed discussion of these concerns). As a result the architectural specification made an important concession to these concerns. It allowed organizations to configure their own preferred root certificate authority or trust anchor. RPKI's design was "capable of accommodating a variety of trust anchor arrangements." At the same time, it further codified within the RPKI standard the reliance on the extant IANA and RIR allocation hierarchy. A statement by SPARTA's Murphy summed up the policy in a colorful way – and also revealed how ambiguous the attitudes and specifications were:

> the ability of a relying party to choose a trust anchor is a big get-out-of-jail-free card for those who are allergic to the idea of one root. NOT that I'm recommending using that card.[8]

Simultaneously to the IETF standards work, many efforts to implement secure routing occurred outside of the IETF. These efforts began even before the standardization process was completed. U.S. government contractor BBN corresponded with APNIC and RIPE and worked extensively with the American Registry for Internet Numbers (ARIN) beginning in 2003, on "how to integrate support for the [RPKI] with ARIN's procedures for assignment of IP addresses and AS numbers." (BBN Technologies, 2004, p. 4) SPARTA's Murphy submitted a policy proposal in February 2006 to ARIN that would require it, in all IP address transactions with its members, to collect an optional field recording "a list of the ASNs that the user permits to originate address prefixes within the address block."[9] In arguing for this proposal, Murphy said that the existing routing registries (even ARIN's own) did not have reliable route origin authorization information, and that other IRRs would be unable to collect it accurately because only the RIRs were

[7] See Section 1 of *An Infrastructure to Support Secure Internet Routing*, available at http://tools.ietf.org/html/draft-ietfsidr-arch-09. The statement was partly inaccurate because there are few established practices and few if any precedents for resource revocation in current policies.

[8] Sandra Murphy, Sparta Inc., in post to SIDR WG list, 1 December 2008. http://www.ietf.org/mail-archive/web/sidr/current/msg00733.html

[9] See ARIN *2006-03: Capturing Origins in Templates*. https://www.arin.net/policy/proposals/2006_3.html

able to authenticate the organizations to whom prefixes were allocated. The proposal was adopted by the ARIN Board of Trustees in November 2006.

However, the debate over RPKI among other RIRs' members was far from settled. The RIPE NCC formed a Resource Certification Task Force in 2006 to examine the implications of deploying a RPKI for resource certification in the region. In addition to producing an overview of the RIRs expected certificate authority hierarchy model, it found numerous impacts of resource certification on existing policies. The model had each RIR maintaining its own trust anchor, to be used for signing and validating of IP address blocks and ASNs it allocated. Instead of cross-certifying each others' trust anchor, RIR's would generate additional certificates for any resource that was allocated from outside but used within its region. With regard to policy, an October 2008 proposal for a certification policy met resistance from participating RIPE community members.[10] As of May 2010, the document had stalled in the policy development process.

In 2008, two policy proposals were submitted to APNIC and RIPE NCC members by Randy Bush, affiliated with Internet Initiative Japan.[11] They proposed to develop RIR-based routing registries combining global RPKI authentication of IP address and ASN assignments with route object authorization information. But neither proposal was adopted by the RIRs' respective memberships. A similar document, essentially reflecting the policy proposals made to the APNIC and RIPE communities, was submitted directly to the ARIN Advisory Council by Bush and IAB member Danny McPherson through its "consultation and suggestion" process. But this proposal, as a suggestion rather than as a policy proposal, could not be deliberated by the ARIN membership.

Despite the lack of enthusiasm expressed by RIR members for resource certification systems, both APNIC and ARIN had already begun to develop the software necessary to support RPKI systems. ARIN and APNIC did this in 2006; RIPE followed suit in 2008. This would coincide with the initiation of the U.S. Department of Homeland Security's Internet Infrastructure Security (IIS) program to facilitate implementation of the National Strategy to Secure Cyberspace. As part of the IIS program, DHS expected to "develop and deploy a Public Key Infrastructure (PKI) with the American Registry for Internet Numbers (ARIN)" by 2008, and to "conclude PKI deployment activities with global registries" by 2010.[12]

When APNIC and ARIN launched their pilot RPKI programs in 2009, they included a public repository of resource certificates and route origin authorizations. By January 2011, all the RIRs intended to offer resource certification as an optional service to their members, with ARIN, APNIC, and LACNIC having determined that no policy was needed.[13] ARIN and the other RIRs acknowledged there would be operational matters

[10] RIPE *2008-08: Initial Certification Policy for Provider Aggregatable Address Space Holders*

[11] RIPE *2008-04:Using the Resource Public Key Infrastructure to Construct Valid IRR Data and* APNIC *prop-0059:Using the Resource Public Key Infrastructure to Construct Valid IRR Data.* See http://archive.apnic.net/policy/proposals/prop-059-v001.html and http://www.ripe.net/ripe/policies/proposals/2008-04.html

[12] See http://www.dhs.gov/xlibrary/assets/SandT5yearplan.pdf, pages 3 and 53

[13] Memo from Andrew de la Haye, RIPE COO, to RIPE Certificate Authority Task Force, 31 March, 2010. http://www.ripe.net/ripe/maillists/archives/ca-tf/2010/msg00014.html

with regard to RPKI, but claimed that certification would be an “opt-in member service,” not a policy, and therefore required no policy decision.[14]
In sum, RPKI implementation has gained considerable momentum even though most of the policy issues raised by it have not been fully aired and resolved. The RIRs have gone ahead with implementation despite the inability of RPKI proposals to gain consensus support in their own policy process.

5. Governance issues raised by RPKI

Implementations of security technologies are never neutral in their impact. They tend to alter power relations and economic dependencies within and across organizations. RPKI is no exception. The design and implementation of a RPKI regime raises important governance issues. As the price of its security benefits, RPKI could reduce the autonomy of Internet service providers and centralize authority over internet resources. Implementing it forces us to reconsider the relationship between ICANN and the Regional Internet Registries. RPKI poses interesting tradeoffs between global compatibility and global security. It is attracting the attention of law enforcement agencies.

Below, we identify and briefly analyze four distinct governance issues raised by RPKI’s design and implementation. By “governance issues” we mean that the choices alter the distribution of power and/or the distribution of costs and benefits. The four points are:

1. RPKI challenges the autonomy of ISPs
2. RPKI requires a fateful tradeoff to be made between simplified global compatibility and centralization of power
3. RPKI affects the policies and business models of the RIRs, and their relationship to IANA
4. RPKI provides important lessons about the role of governments (nation-states) in Internet governance

5.1 The Autonomy of ISPs

The internet has evolved in a way that detaches responsibility for address allocation from operational responsibility for routing. The RIRs register and record address block assignments in order to keep them unique. While ISPs use the RIR’s address allocations database, Internet service providers wholly control and authorize what routes they announce, and they decide for themselves which other ISPs’ routing announcements they trust or filter. Indeed, the RIRs’ authority over address usage itself is almost completely a byproduct of the ISPs’ willingness to use their registries as coordination tools.

RPKI changes all that. It has the potential to give RIRs direct, operational impact on routing. David Conrad, at that time the head of IANA, wrote on the SIDR list,

> Today, RIR influence on routing is essentially advisory in nature -- if an address holder (say) fails to pay their address maintenance fee, RIRs can, at most, remove the address holder's blocks from whois databases. However, as I

[14] Memo from Andrew de la Haye to RIPE Certificate Authority Task Force, 31 March, 2010. http://www.ripe.net/ripe/maillists/archives/ca-tf/2010/msg00014.html
more detail on ARIN's policy development process, specifically Section 2.6

> understand it, this has limited effect on existing [routing arrangements]. The RIR could potentially reallocate the space, but this would likely be a good way of annoying multiple parties (not just the folks the address space was reclaimed from). …[I]f filter lists are built or routers check origin authenticity in real-time by traversing the RPKI tree(s), there would seem to be significantly more control vested in each parent node in the path up to the root of the RPKI hierarchy. My fear is that this will simply be unacceptable in a political or business sense.[15]

Confirming Conrad's point, a university network operator objected to the way RPKI altered "the balance of power" between network operators and the RIRs:

> Today if there is a legal dispute between an allocator [RIR] and an organization with an allocation, it will be solved through existing civil means. This may take some time. In the meantime the status quo continues (from a technical/operational perspective). With RPKI the allocator can revoke the organizations certificate, while the civil process takes its time, causing harm to the organization that is now un-routable. Don't think they won't do the revocation. I have personally seen situations where if one party has "the switch" to enforce their will, they use it.[16]

One of the key governance issues concerns the terms and conditions under which certificates would be revoked. As the quotations above suggest, in an RPKI regime certificate revocation could have direct operational effects. Predictably, revocation of certificates emerged as a critical point of contention in the broader community's debates over RPKI. For example, when RIPE-NCC proposed implementing resource certification, its members refused to support it due to concerns about the length of certificate validity and the linking of certificate revocation to RIPE membership status. Participants commented that "people will be reluctant to [use resource certificates] if they have reasons to fear that routing may be stopped due to unexpected events relating to certificates revocation."[17] Furthermore, there were allegations that some governments viewed the ability of the RIRs to withdraw a certificate (thereby theoretically preventing prefixes from being routable) as an infringement of national sovereignty.[18]

Clearly, RPKI diminishes the autonomy of ISPs. It could be used to replace a looser, networked form of governance based on decentralized associative choices among Internet service providers with a more centralized and hierarchical governance form.

5.2 Trust model and global compatibility

If RPKI has the potential to centralize power, the next key issue becomes, "where is this power centralized?" In particular, what organization or institution serves as the root-level trust anchor for the certificate hierarchy?

[15] David Conrad, post to SIDR WG list 17 September 2009. http://www.ietf.org/mail-archive/web/sidr/current/msg01098.html

[16] Jeff Schiller, MIT network operator, post to the SIDR WG email list, September 20, 2009, http://www.ietf.org/mail-archive/web/sidr/current/msg01117.html

[17] See http://www.ripe.net/ripe/maillists/archives/ca-tf/2009/msg00013.html

[18] See http://www.ripe.net/ripe/maillists/archives/ca-tf/2010/msg00008.html

In the classical PKI scenario, everyone trusts a single CA and the sender and recipient of the information rely on the same CA. This kind of centralization is relatively easy to achieve when we are dealing with a single organization.[19] It becomes harder and harder to achieve as the set of organizations using it becomes larger and more diverse. Even the U.S. government could not agree on a single CA for all of its activities. When one is talking about the global internet, which involves ~35,000 autonomous systems and hundreds of thousands more resource holders scattered across hundreds of different language groups and political systems, centralization and unity of a trust anchor become unrealistic – even disruptive and dangerous – goals.

Insofar as one uses a centralized, strictly hierarchical trust model, one is also creating the potential for centralizing political and regulatory authority over the Internet. We have already seen this drama play out in the domain name system. The DNS is a hierarchical name space. From a purely technical standpoint, all the administrator of the root has to do is coordinate top level domain name assignments to ensure that all names are unique and there are no duplications. But the concentration of power at the root creates a big, fat target for political, military and business actors. Almost inevitably, control of the DNS root will be extended beyond coordination of uniqueness to achieve other political and economic objectives. ICANN, for example, regulates the dispensation of top level domains (as well as second-level domains) in ways that respond to political demands for trademark and copyright protection, support economic competition policy objectives, nation-state control of geographic names, etc., etc.

Is it desirable to move toward that kind of centralization over routing? Some of the political implications were noted by IAB member D. McPherson, who wrote “If some country holding the keys (TA) goes to war with another and decides they want to revoke all of their allocations, then ISPs would have zero control over this outside of their own routing domain.”[20] The concern over “bringing down national network infrastructures” and the relationship to a single authoritative trust anchor residing with IANA (which maintains a contractual relationship with a single government) were expressed again in a recent study. (ENISA 2010b)

The SIDR working group – significantly influenced by researchers supported by U.S. military contracts – wanted to map the resource allocation hierarchy directly onto the PKI, to make it as technically simple and unambiguous as possible. However, its deliberations explicitly noted the political and governance issues associated with that. In an attempt to square the circle, it proposed to allow Relying Parties to choose a diverse set of Trust Anchors. But this means that both the software functioning and the security arrangements are more complicated, less predictable, and possibly less secure.

[19] “In the classical PKI scenario, someone receives a document signed with a digital certificate. The recipient must trust the creator of that certificate (the Certification Authority (CA)) to be able to confirm the identity of the sender. This is simple if the sender and recipient are using the same CA. The need for interoperability arises where the document has been signed with a certificate from a CA that the recipient does not know. The obvious approach is to centralise as much trust as possible and avoid this problem entirely. This is reflected in the root CA and hierarchy PKI models discussed below. However, those models require tight central control and unanimous support.” (Galexia 2005 pg. 4)

[20] Danny McPherson, post to SIDR WG list 11 March 2008, http://www.ietf.org/mail-archive/web/sidr/current/msg00346.html

The governance issues associated with control of the root were not clarified when the Number Resource Organization (NRO, an association of the RIRs) and the Internet Architecture Board (IAB) issued statements about the issue in late 2009 and early 2010, respectively. Both statements supported a "single, authoritative trust anchor" for RPKI. But both statements also were willing to issue "get out of jail free" cards allowing relying parties to choose their own trust anchor, at least in the short term. And both statements studiously avoided any serious discussion of the governance issues, in particular the linkage of resource allocation to operational control of routing.

The NRO (2009) statement said:

> The Regional Internet Registries (RIRs) believe that the optimal eventual RPKI configuration involves a single authoritative trust anchor. That configuration may not be achievable in the short-term and the details and timelines for its implementation will depend among other things on discussions within the RIRs' communities and dialogues with others including the Internet Architecture Board (IAB) and the Internet Engineering Task Force (IETF). In the meantime, the RIRs have agreed to undertake pragmatic implementations of RPKI services based on interim trust anchor models...

The IAB (2010) statement said:

> The IAB considers a properly designed and deployed RPKI to be an absolute prerequisite to having a secure global routing system, which is in turn a prerequisite to having a reliable worldwide Internet. ... The SIDR architecture and protocols have been designed to support a single trust anchor as well as multiple trust anchors. The IAB however, is in strong agreement with the Number Resource Organization (NRO) regarding the number of trust anchors as well as what and whom they represent:
> 1. the RPKI should have a single authoritative trust anchor
> 2. this trust anchor should be aligned with the registry of the root of the allocation hierarchy

Although the IAB statement made a perfunctory mention of the governance issues, in the end it simply ignored them. Its support for a single trust anchor, it openly admitted, was based entirely on technical reasoning and ignored any concerns about what it called "fairness and equality." Worse, even its technical reasoning was incomplete: it supported a single root in order to avoid "inconsistent and conflicting assertions about to whom a particular address block has been allocated." While these concerns are valid, IAB did not address the linkage of address allocation and assignment to the real-time control of ISP routing, nor did it explain how its support for a single root was consistent with the ability of ISPs to choose their own trust anchor.

Another issue is that the "get out of jail free card" promised by the SIDR working group, NRO and IAB may turn out to be illusory. While the ability of ISPs to choose their own trust anchor might lead to a more heterogeneous yet compatible certification system, it is also possible that governments or the ICANN/RIR system itself could decide to require certification as a condition of receiving address resources. (As a precedent, ICANN decided to require DNSSEC implementation for applicants for new top level domains.) It

is also possible that once the system achieves a critical mass of adopters, powerful network effects will lead to convergence on a single, centralized trust anchor, creating compatibility problems for ISPs that don't use the same trust anchor. As long as it is unclear how RPKI achieves compatibility among multiple roots, it is disingenuous to pretend that RPKI allows ISPs a free choice of trust anchors – just as it is disingenuous to pretend that anyone who wants to create an alternate DNS root can easily do so.
The issue of what type of trust anchor system will be used and how we might avoid some of the problems associated with either a more-centralized or less-centralized system constitutes a global internet governance issue of the highest order. And it is also clear that the entities making the decisions thus far – IETF, IAB and NRO – have not done much to resolve them.

5.3 Future of the RIRs

RPKI puts the RIRs in the center of many internet governance issues by dramatically expanding their authority over Internet resources. It has also raised some tensions regarding the relationship between the RIRs and ICANN/IANA.

At its best, RPKI would not only help to prevent bogus route announcements and address hijacking, it would also facilitate the smooth transfer of ipv4 address resources from one party to another after the free pool is depleted. A fully functional, globally compatible RPKI system would act as an effective property title for IP address blocks, giving the address holder legitimate claim to acquire or transfer address resources, and allow third parties to verify who is the legitimate holder of address blocks. By itself, this would constitute a major change in the role of the RIRs.

But getting comprehensive resource certification into place requires huge changes in the procedures and policies of the RIRs. And depending on the policies adopted, such an implementation could either introduce healthy competitive pressures among the RIRs, or foreclose such competition. The implementation process also could massively increase their workload at a time when their ability to collect fees may be attenuated.

The value of a certificate regime depends to a great extent on global adoption and use of RPKI. This means that the RIRs (or someone) must issue certificates to all current ip address holders, or at least to a critical mass of them. To do this, an RIR would have to first verify the identity of an existing resource holder, then verify the validity of the assignment or allocation, and finally issue a certificate. This process would have to be repeated for every address block holder.

In the ipv4 space, where the RIRs issued most of the allocations years ago, the RIRs would be thrust into the role of auditing each network's address usage with the implicit threat of taking away the resources if the allocation is no longer consistent with policy. As one RIPE-NCC document admitted, "Many resources are now used for other purposes than they were originally assigned for. Certifying such resources would seem to imply that the RIPE NCC has validated this re-assignment."[21] If the RIRs actively re-review each organization's "needs" for and usages of ip addresses, and threaten to refuse certification to address holders on that basis, they would be taking a far more

[21] RIPE document 070206, "Outline new and current services affected by certification." Draft v1.5 https://ripe59.ripe.net/ripe/maillists/archives/ca-tf/2007/doc00000.doc

aggressive intervention into internetworking than they have taken before. This might discourage a large part of the industry from participating in the certification regime, undermining its value. The RIRs do not have any real experience with the revocation of address resources. A RIPE-NCC document noted "Revocation of a resource has never taken place... Policy in this area will need to address whether 'not renewing the certificate' means revoking the resource and returning it to the pool."[22]

These issues become especially thorny when considering so-called "legacy" holders of v4 address blocks. Many organizations and ISPs acquired address blocks before the RIRs existed; therefore they have no contractual relationship with an RIR regarding their address holdings. The RIR is not in any position to "certify" or approve the validity of their address holdings.

The RIRs must also decide whether to make certification contingent upon payment of membership fees. If RPKI became so widely adopted that most ISPs refused to route packets from entities with no certification, such a requirement would make membership in the RIRs compulsory and their fees a kind of monopoly tax rather than a membership payment for a voluntarily selected set of services and organizational rights. One ISP expressed fears about the monopoly power of the RIRs during the SIDR working group:

> Although there is plenty of sense in aligning the RPKI chain of trust with the resource allocation chain, ISPs may have concerns with the RIRs being the trust anchors. The incentive structure for the RIRs is fundamentally different than that of a [private market] certificate provider like Verisign/Thawte/CyberTrust. If these root CAs time and again demonstrate that they are untrustworthy they lose customers, revenue, and potentially their trusted status. What entices an RIR toward vigilance as they validate the supposedly authorized origin of a prefix?[23]

This ISP was noting that the absence of competition among address allocation authorities makes it difficult for businesses to accept them as root Certificate Authorities. Competition provides a form of accountability that is lacking in this case, unless the RPKI regime allows anyone to certify resources, not just the RIRs. Introducing competition among the RIRs as providers of services related to address holdings (e.g., Whois, routing registries, and authentic titles for address holdings) would alleviate many of these concerns. However, the RIRs were originally organized to be territorially exclusive issuers of ip address blocks. And many of the existing policies and planned address block transfer policies are designed to maintain that exclusivity. The transfer of resources has been described by the aforementioned RIPE-NCC document as possibly "one of the most complex questions of resource certification." If the RIRs act in a coordinated fashion as an "IP address cartel," they could use RPKI to maintain the territorial exclusivity of address allocations, thereby eliminating inter-regional transfers on the servicing of address blocks and thus eliminate a competitive check on the RIRs.

[22] RIPE document 070206, "Outline new and current services affected by certification." Draft v1.5 https://ripe59.ripe.net/ripe/maillists/archives/ca-tf/2007/doc00000.doc

[23] Ryan Shea, Senior Engineer, Network and Info Security, Verizon Business in post to SIDR WG email list 22 September 2009. http://www.ietf.org/mail-archive/web/sidr/current/msg01142.html

There has also been evidence of tension and negotiation between ICANN and the RIRs regarding the implementation of RPKI. It is notable that despite all the support expressed by the NRO and the IAB for a single trust anchor for RPKI, neither explicitly proposes to make ICANN the root. This provides a very interesting clue as to the institutional competition going on between ICANN and the RIRs. Strictly applied, the IAB's principle that "th[e] trust anchor should be aligned with the registry of the root of the allocation hierarchy" means that ICANN, which controls the root of the IP address allocation hierarchy, should be the supreme certificate authority over addressing. If ICANN held this authority, however, it could add yet another revenue stream and its authority over internet governance would grow. It might even be possible to disintermediate the RIRs, and issue certificates and address blocks directly to organizations and end users.

Bear in mind that the Address Supporting Organization has never been formally established as an independent entity and that the RIRs' trade association, the Number Resource Organization, has never signed a formal contract with ICANN binding themselves to its rules. Instead, the RIRs and ICANN are joined through a loose and noncommittal memorandum of understanding. Indeed, whereas ICANN gets over $50 million a year in fees from domain name registries and registrars, it collects less than a million in "voluntary contributions" from the RIRs. From the standpoint of the decentralization of power over Internet governance, these informal relationships are a good thing in certain respects. But RPKI threatens to reconfigure things.

It is clear that ICANN has taken a firm interest in RPKI, and that the RIRs were not, initially, entirely comfortable with this interest. In June 2008, ICANN's Security and Stability Advisory Committee (SSAC) indicated its interest in "management of certificates for the addressing system (RPKI)." Later, the U.S. Department of Homeland Security's IIS program manager joined ICANN's Security and Stability Advisory Committee, and ICANN's 2010-11 fiscal budget included financial support for the activity. Perhaps responding to ICANN's overtures, ARIN's Board determined it did not support "a sole IANA signed key" for the RPKI.[24] ARIN proposed instead a shared key, a position it communicated to the other RIRs.

While there are legitimate concerns about strengthening the power of the RIRs to extract fees and about their status as exclusive suppliers of a critical resource, these concerns apply even more strongly to ICANN. As membership organizations, the accountability arrangements of the RIRs, while not perfect, are superior to those of ICANN. (Mueller, 2009) And the RIRs, unlike ICANN, are not under a contractual obligation to a single government.

5.4 Role of governments in IG

RPKI's potential to centralize authority over internet routing would provide an inviting target for the many regulatory agencies and interest groups who want to assert stronger control over the internet. As noted in section 5.2, the centralization of authority over the domain name space in ICANN has already led to a counterproductive politicization of many issues. A similar fate could befall routing if the institutional arrangements for RPKI are not handled properly.

[24] See ARIN Board minutes, available at https://www.arin.net/about_us/bot/bot2009_0206.html

There is another, more subtle argument to be made regarding the role of governments. Some observers have asserted that the internet is immune to traditional forms of regulation by nation-states and intergovernmental organizations; others have reacted by claiming that the internet can and should be subordinated to traditional governmental authority and jurisdictions. The history of RPKI development provides insight into a newer, more Internet-specific role that governments play in nongovernmental Internet institutions. In this case we see that the U.S. government has succeeded in participating in and shaping the bottom-up standards and policy development process of the IETF and RIRs in a unique way. The influence took place not through the formal participation of governmental representatives, nor by the existence of a special Governmental Committee with formal supervisory powers, but by *contracting* with scientists and researchers to do research on the relevant standards and then take their proposals directly into the process. The US government exerted a strong influence on the outcome by funding BBN and SPARTA, but nevertheless these contractors had to interact more or less as peers with other participants in the standards development and RIR policy development processes.

This model is not ideal, but it seems superior to more bureaucratic and hierarchical forms of interaction between the internet community and governments. Although there are some troubling signs that the bottom up policy development process has been minimized in this case, that problem is not inherent in the mode of participation. In general, it is possible for governments to influence Internet governance institutions without transforming those institutions into intergovernmental organizations, and without subjecting them to hierarchical "oversight" or supervision.

6. Concluding observations

Routing is now at the center of the debate over the security and governance of the internet. The policies ISPs use for routing have major policy and cost implications for themselves, their users and interconnection partners. Routing and the allocation and assignment of IP addresses are deeply interdependent, and thus IP address policy, from ICANN down through the regional address registries, will be deeply affected by RPKI. Routing methods and policies also have regulatory and law enforcement implications – both for those who would try to control the internet and for those who would try to keep it free from arbitrary interference.

There is still dissension about the technical wisdom of RPKI. We need to have a richer discussion of RPKI and its implications for internet governance. The technical expertise embodied in the IAB and IETF is invaluable, but not dispositive; other forms of expertise and participation are required.

The Internet Governance Project (IGP) is an alliance of academics that puts expertise into practical action in the fields of global governance, Internet policy, and information and communication technology.

To download publications or learn more about the IGP, please visit our website at http://internetgovernance.org

References

BBN Technologies. (2004). Transitioning Secure Border Gateway Protocol (S-BGP) into the Internet (AFRL-IF-RS-TR-2004-63 Final Technical Report). Retrieved from http://www.dtic.mil/srch/doc?collection=t3&id=ADA422110.

Barbir, A., Murphy, S., & Yang, Y. (2006). *Generic Threats to Routing Protocols. RFC 4593*. Internet Engineering Task Force. Retrieved from http://tools.ietf.org/html/rfc4593.

Butler, K., Farley, T., McDaniel, P., & Rexford, J. (2010). A Survey of BGP Security Issues and Solutions. *Proceedings of the IEEE*, *98*(1), 100-122. doi: 10.1109/JPROC.2009.2034031.

European Network and Information Security Agency [ENISA]. (2010). Report on secure routing technologies. Retrieved from http://www.enisa.europa.eu/act/res/technologies/tech/routing/report-on-secure-routing-technologies/at_download/fullReport.

European Network and Information Security Agency [ENISA]. (2010b). State-of-the-art Deployment and Impact on Network Resilience. Retrieved from http://www.enisa.europa.eu/act/res/technologies/tech/routing/state-of-the-art-deployment-and-impact-on-network-resilience/at_download/fullReport.

Galexia. (2005). *PKI Interoperability Models* (p. 23). Sydney, Australia. Retrieved from http://www.galexia.com/public/research/assets/pki_interoperability_models_2005/pki_interoperability_models_2005.pdf.

Hu, Y., McGrew, D., Perrig, A., Weis, B., & Wendlandt, D. (2006). (R)Evolutionary Bootstrapping of a Global PKI for Secure BGP. In *Workshop on Hot Topics in Networks (HotNets'06)*. Irvine, CA. Retrieved from http://sparrow.ece.cmu.edu/group/pub/hu_mcgrew_perrig_weis_wendlandt_bgp.pdf.

Internet Architecture Board [IAB]. (2010). *IAB statement on the RPKI*. Retrieved from http://www.ietf.org/mail-archive/web/ietf-announce/current/msg07028.html.

Kent, S. (2006). An Infrastructure Supporting Secure Internet Routing. In A. S. Atzeni & A. Lioy, *Public key infrastructure : Third European PKI Workshop: Theory and Practice*, Lecture Notes in Computer Science (Vol. 4043, pp. 116-129). Berlin, Heidelberg: Springer Berlin Heidelberg. doi: 10.1007/11774716.

Kent, S., Lynn, C., & Seo, K. (2000). Design and analysis of the Secure Border Gateway Protocol (S-BGP). In *Proceedings DARPA Information Survivability Conference and Exposition. DISCEX'00* (pp. 18-33). IEEE Comput. Soc. Retrieved from http://ieeexplore.ieee.org/lpdocs/epic03/wrapper.htm?arnumber=824939.

Mueller, M. (2009). ICANN, Inc.: Accountability and participation in the governance of critical Internet resources. Internet Governance Project. Retrieved from http://internetgovernance.org/pdf/ICANNInc.pdf.

Number Resource Organization [NRO]. (2009). *NRO Statement on RPKI*. Retrieved from http://www.nro.net/news/nro-declaration-rpki.html.

Rehkter, Y., Li, T., & Hares, S. (2006). *A Border Gateway Protocol 4 (BGP-4). RFC 4271*. Internet Engineering Task Force. Retrieved from http://tools.ietf.org/html/rfc4271.

Rekhter, Y., & Li, T. (1995). *A Border Gateway Protocol 4 (BGP-4). RFC 1771*. Internet Engineering Task Force. Retrieved from http://tools.ietf.org/html/rfc1771.

SPARTA Inc. (2006). *Secure Protocols for the Routing Infrastructure (SPRI) Initiative: A Road Map*. Retrieved from https://www.cyber.st.dhs.gov/docs/spriRoadmap.pdf.

Seo, K., Lynn, C., & Kent, S. (2001). Public-key infrastructure for the Secure Border Gateway Protocol (S-BGP). In *Proceedings of DARPA Information Survivability Conference and Exposition II. DISCEX'01* (pp. 239-253). Anaheim, CA: IEEE Comput. Soc. doi: 10.1109/DISCEX.2001.932219.

U.S. Army Research Development and Engineering Command [U.S. Army RDECOM]. (2005). *Securing the Routing Infrastructure (Award Contract: W911NF-05-C-0113)* (p. 30). Retrieved from http://internetgovernance.org/pdf/Sparta FOIA, Award, etc....pdf.