

Network Working Group
Internet-Draft
Intended status: Informational
Expires: January 3, 2015

G. Huston
G. Michaelson
APNIC
C. Martinez
LACNIC
T. Bruijnzeels
RIPE NCC
A. Newton
ARIN
A. Aina
AFRINIC
July 2, 2014

RPKI Validation Reconsidered
draft-ietf-sidr-rpki-validation-reconsidered-00.txt

Abstract

This document reviews the certificate validation procedure specified in RFC6487 and highlights aspects of potentially acute operational fragility in the management of certificates in the RPKI in response to the movement of resources across registries, and the associated actions of Certification Authorities to maintain continuity of validation of certification of resources during this movement.

Status of this Memo

This Internet-Draft is submitted in full conformance with the provisions of BCP 78 and BCP 79.

Internet-Drafts are working documents of the Internet Engineering Task Force (IETF). Note that other groups may also distribute working documents as Internet-Drafts. The list of current Internet-Drafts is at <http://datatracker.ietf.org/drafts/current/>.

Internet-Drafts are draft documents valid for a maximum of six months and may be updated, replaced, or obsoleted by other documents at any time. It is inappropriate to use Internet-Drafts as reference material or to cite them other than as "work in progress."

This Internet-Draft will expire on January 3, 2015.

Copyright Notice

Copyright (c) 2014 IETF Trust and the persons identified as the document authors. All rights reserved.

This document is subject to BCP 78 and the IETF Trust's Legal Provisions Relating to IETF Documents (<http://trustee.ietf.org/license-info>) in effect on the date of publication of this document. Please review these documents carefully, as they describe your rights and restrictions with respect to this document. Code Components extracted from this document must include Simplified BSD License text as described in Section 4.e of the Trust Legal Provisions and are provided without warranty as described in the Simplified BSD License.

Table of Contents

- 1. Introduction 3
- 2. Certificate Validation in the RPKI 3
- 3. Operational Considerations 4
- 4. Alternatives Approaches 7
- 5. Security Considerations 7
- 6. IANA Considerations 8
- 7. Acknowledgements 8
- 8. References 8
 - 8.1. Normative References 8
 - 8.2. Informative References 8
- Authors' Addresses 8

1. Introduction

This document reviews the certificate validation procedure specified in RFC6487 and highlights aspects of potentially acute operational fragility in the management of certificates in the RPKI in response to the movement of resources across registries, and the associated actions of Certification Authorities to maintain continuity of validation of certification of resources during this movement.

2. Certificate Validation in the RPKI

As currently defined in section 7.2 of [RFC6487], validation of PKIX certificates that conform to the RPKI profile relies on the use of a path validation process where each certificate in the validation path is required to meet the certificate validation criteria. This can be considered to be a recursive validation process where, in the context of an ordered sequence of certificates, as defined by each pair of certificates in this sequence having a common Issuer and Subject Name respectively, a certificate is defined as valid if it satisfies basic validation criteria relating to the syntactic correctness, currency of validity dates and similar properties of the certificate itself, as described in [RFC5280], and also that it satisfies certain additional criteria with respect to the previous certificate in the sequence (the Issuer part of the pair), and that this previous certificate is itself a valid certificate using the same criteria. This definition applies recursively to all certificates in the sequence apart from the initial sequence element, which is required to be a Trust Anchor.

For RPKI certificates, the additional criteria relating to the previous certificate in this sequence is that the certificate's number resource set, as defined in [RFC3779], is "encompassed" by the number resource set contained in the previous certificate.

Because [RFC6487] validation demands that all resources in a certificate be valid under the parent (and recursively, to the root), a digitally signed attestation, such as a Route Origin Authorization (ROA) object [RFC6482], which refers only to a subset of RFC3779-specified resources from that certificate validation chain can be concluded to be invalid, but not by virtue of the relationship between the RFC3779 extensions of the certificates on the putative certificate validation path and the resources in the ROA, but by other resources described in these certificates where the "encompassing" relationship of the resources does not hold. Any such invalidity along the certificate validation path can cause this outcome, not just at the immediate parent of the end entity certificate that attests to the key used to sign the ROA.

For example, in the certificate sequence:

Certificate 1:

Issuer A, Subject B, Resources 192.0.2.0/24, AS64496-AS64500

Certificate 2:

Issuer B, Subject C, Resources 192.0.2.0/24/24, AS64496-AS64511

Certificate 3:

Issuer C, Subject D, Resources 192.0.2.0/24

Certificate 3 is considered to be an invalid certificate, because the resources in Certificate 2 are not encompassed by the resources in Certificate 1, by virtue of certificate 2 describing the resources of the range AS64501 - AS64511 in this RFC3779 resource extension. Obviously, these Autonomous Systems numbers are not related to the IPv4 resources contained in Certificate 3.

Any non-encompassed resource set can cause invalidation, be it an ASN, IPv4 or IPv6 resource, if it is not encompassed by the resource set in the parent (Issuer) certificate.

The underlying observation here is that this definition of certificate validation treats a collection of resources as inseparable, so that a single certificate containing a bundle of number resources is semantically distinct from an equivalent set of certificates where each certificate contains a single number resource. This semantic distinction between the whole and the sum of its parts is an artifice introduced by the particular choice of a certificate validation procedure, as distinct from meeting any particular operational requirement, and the result is the introduction of operational fragility into the handling of RPKI certificates, particularly in the case where number resources are moved between the corresponding registries, as described here.

3. Operational Considerations

There are two areas of operational concern with the current RPKI validation definition.

The first is that of the robustness of the operational management procedures in the issuance of certificates. If a subordinate Certification Authority (CA) issues a certificate that contains an Internet Number Resource (INR) collection that is not either exactly equal to, or a strict subset of, its parent CA, then this issued certificate, and all subordinate certificates of this issued certificate are invalid. These certificates are not only defined as

invalid when being considered to validate an INR that is not in the parent CA certificate, but are defined as invalid for all INRs in the certificate.

This constraint creates a degree of operational fragility in the issuance of certificates, as all CA's are now required to exercise extreme care in the issuance and reissuance of certificates to ensure that at no time do they overclaim on the resources described in the parent CA, as the consequences of an operational lapse or oversight implies that all the subordinate certificates from the point of INR mismatch are invalid. It would be preferred if the consequences of such an operational lapse were limited in scope to the specific INRs that formed the mismatch, rather than including the entire set of INRs within the scope of damage from this point of mismatch downward across the entire sub-tree of descendant certificates in the RPKI certificate hierarchy.

The second operational consideration described here relates to the situation where a registry withdraws a resource from the current holder, and the resource is transferred to another registry, to be registered to a new holder in that registry. The reason why this is a consideration in operational deployments of the RPKI lies in the movement of the "home" registry of number resources during cases of mergers, acquisitions, business re-alignments, and resource transfers and the desire to ensure that during this movement all other resources can continue to be validated.

If the original registry's certification actions are simply to issue a new certificate for the current holder with a reduced resource set, and to revoke the original certificate, then there is a distinct possibility of encountering the situation illustrated by the example in the previous section. This is a result of an operational process for certificate issuance by the parent CA being de-coupled from the certificate operations of child CA.

This de-coupled operation of CAs introduces a risk of unintended third party damage: since a CA certificate can refer to holdings which relate to two or more unrelated subordinate certificates, if this CA certificate becomes invalid due to the reduction in the resources allocated to this CA relating to one subordinate resource set, all other subordinate certificates are invalid until the CA certificate is reissued with a reduced resource set.

In the example provided in the previous section, all subordinate certificates issued by CA B are invalid, including all certificates issued by CA C, until CA A issues a new certificate for CA B with a reduced resource set.

At the lower levels of the RPKI hierarchy the resource sets affected by such movements of resources may not encompass significantly large pools of resources. However, as one ascends through this certification hierarchy towards the apex, the larger the resource set that is going to be affected by a period of invalidity by virtue of such uncoordinated certificate management actions. In the case of a Regional Internet Registry (RIR) or National Internet Registry (NIR), the potential risk arising from uncoordinated certification actions relating to a transfer of resources is that the entire set of subordinate certificates that refer to resources administered by the RIR or the NIR cannot be validated during this period.

Avoiding such situations requires that CA's adhere to a very specific ordering of certificate issuance. In this framework, the common registry CA that describes (directly or indirectly) the resources being shifted from one registry to the other, and also contains in subordinate certificates (direct or indirect) the certificates for both registries who are parties to the resource transfer has to coordinate a specific sequence of actions.

This common registry CA has to first issue a new certificate towards the "receiving" registry that adds to the RFC3779 extension resource set the specific resource being transferred into this receiving registry. The common registry CA then has to wait until all registries in the subordinate certificate chain to the receiving registry have also performed a similar issuance of new certificates, and in each case a registry must await the issuance of the immediate superior certificate with the augmented resource set before it, in turn, can issue its own augmented certificate to its subordinate CA. This is a "top down" issuance sequence."

It is possible for the common registry to issue a certificate to the "sending" registry with the reduced resource set at any time, but it should not revoke the previously issued certificate, nor overwrite this previously issued certificate in its repository publication point without specific coordination. Only when the common registry is assured that the top down certificate issuance process to the receiving registry CA chain has been completed can the common registry commence the revocation of the original certificate for the sending registry. However, it should not do so until it is assured that the immediate subordinate registry CA in the path to the sending registry has issued a certificate with a reduced resource set, and so on. This implies that on the sending side the certificate issuance and revocation is a "bottom up" process.

If this process is not carefully followed, then the risk is that some or all of the subordinate certificates of this common registry CA will be unable to be validated until the entire process of

certificate issuance and revocation has been completed. While this sequenced process is intended to preserve validity of certificates in the RPKI, it is a complex, fragile and operationally cumbersome process.

The underlying consideration here is that the operational coordination of these certificate issuance and revocation actions to effect a smooth resource transfer across registries is mandated by the nature of the particular choice of certificate validation process described in [RFC6487].

4. Alternatives Approaches

If the current definition of the RPKI certificate validation procedure is considered to introduce unacceptable levels of fragility and risk into the operational environment, what alternatives exist?

One approach is to remove the semantic requirement to consider the collection of resources in the extension field of the RPKI certificate as an indivisible bundle. This would allow for a certificate to be considered as valid for some subset of the resources listed in this extension, without necessarily being considered as valid for all such described resources. The implications of this approach is that any mismatch between parent and subordinate over resources where the subordinate certificate lists resources that are not contained in the parent certificate would affect validity questions relating to only those particular resources, rather than invalidating the subordinate certificate for all resources, and all of its subordinate products. This would appear to offer a relatively precise match to the defined problem space, and limits the scope of consequent third party damage in the event of a INR mismatch in the RPKI certification hierarchy.

Another approach may involve the alteration of the RPKI provisioning protocol [RFC6492] to include a specific signal from child to parent ("bottom up") relating to readiness for certificate revocation. At this stage it is entirely unclear how this signalling mechanism would operate, nor is it clear that it would alter the elements of operational fragility nor mitigate to any meaningful extent the risks of failure to ensure strict INR consistency at all times. This is a topic for further study.

5. Security Considerations

The Security Considerations of [RFC6487] and [RFC6492] do not address the topic described here. Obviously, within the current RPKI

validation procedure, any inconsistency in certificates located towards the apex of the RPKI hierarchy would invalidate the entirety of the sub-tree located below the point of this inconsistency. If the RPKI was used to control inter-domain routing in the context of a secure routing protocol, then the implications of this large scale invalidation of certificates would have a corresponding massive impact on the stability of routing. This appears to be a serious situation.

6. IANA Considerations

No updates to the registries are suggested by this document.

7. Acknowledgements

TBA.

8. References

8.1. Normative References

- [RFC3779] Lynn, C., Kent, S., and K. Seo, "X.509 Extensions for IP Addresses and AS Identifiers", RFC 3779, June 2004.
- [RFC5280] Cooper, D., Santesson, S., Farrell, S., Boeyen, S., Housley, R., and W. Polk, "Internet X.509 Public Key Infrastructure Certificate and Certificate Revocation List (CRL) Profile", RFC 5280, May 2008.
- [RFC6487] Huston, G., Michaelson, G., and R. Loomans, "A Profile for X.509 PKIX Resource Certificates", RFC 6487, February 2012.

8.2. Informative References

- [RFC6482] Lepinski, M., Kent, S., and D. Kong, "A Profile for Route Origin Authorizations (ROAs)", RFC 6482, February 2012.
- [RFC6492] Huston, G., Loomans, R., Ellacott, B., and R. Austein, "A Protocol for Provisioning Resource Certificates", RFC 6492, February 2012.

Authors' Addresses

Geoff Huston
Asia Pacific Network Information Centre
6 Cordelia St
South Brisbane, QLD 4101
Australia

Phone: +61 7 3858 3100
Email: gih@apnic.net

George Michaelson
Asia Pacific Network Information Centre
6 Cordelia St
South Brisbane, QLD 4101
Australia

Phone: +61 7 3858 3100
Email: ggm@apnic.net

Carlos M. Martinez
Latin American and Caribbean IP Address Regional Registry
Rambla Mexico 6125
Montevideo 11400
Uruguay

Phone: +598 2604 2222
Email: carlos@lacnic.net

Tim Bruijnzeels
RIPE Network Coordination Centre
Singel 258
Amsterdam 1016 AB
The Netherlands

Email: tim@ripe.net

Andrew Lee Newton
American Registry for Internet Numbers
3635 Concorde Parkway
Chantilly, VA 20151
USA

Email: andy@arin.net

Alain Aina
African Network Information Centre (AFRINIC)
11th Floor, Raffles Tower
Cybercity, Ebene
Mauritius

Phone: +230 403 51 00
Email: aalain@afnic.net

