

Secure Inter-Domain Routing
Internet-Draft
Intended status: Standards Track
Expires: August 17, 2014

E. Barnes
BBN Technologies
February 13, 2014

Resource Public Key Infrastructure (RPKI) Resource Transfer Protocol and
Transfer Authorization Object (TAO)
draft-barnes-sidr-tao-00

Abstract

This document defines an extension to the rpki-updown protocol to provide support for transferring Internet Number Resources from one INR holder to another. Such transfers take place external to the RPKI, using procedures defined within and between RIRs. This protocol facilitates automation of the maintenance of RPKI data in the context of INR transfers. The protocol supports asynchronous transfers of live or unused INRs within an RIR or between RIRs. The scope of this protocol is limited to the transfer of Internet Number Resources within the Resource Public Key Infrastructure. In support of this protocol, this document also defines a new signed object type for the RPKI repository system, the Transfer Authorization Object (TAO).

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 August 17, 2014.

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
1.1. Terminology	3
2. Scope	4
3. Protocol Specifications	4
3.1. INR Source Path	5
3.2. INR Recipient Path	7
3.3. Swing Point	8
3.4. Transfer Authorization Object	9
3.4.1. TAO Type	9
3.4.2. TAO Validation	9
3.5. ASN.1 Specification of the TAO	10
3.5.1. transferFromSKI	10
3.5.2. transferToSKI	10
3.5.3. ipAddrBlocks	10
3.5.4. asIdentifiers	10
3.5.5. liveXfer	10
3.5.6. overlapPeriod	11
3.6. Common Message Format	11
3.7. End Entity Certificate Constraint	11
3.8. INR Transfer	11
3.8.1. Transfer	11
3.8.2. Request-Not-Performed Response	13
3.8.3. Timeout Response	13
3.8.4. Overlap Failure Response	13
3.9. XML Schema	13
4. Security Considerations	16
5. IANA Considerations	16
6. Acknowledgements	16
7. References	17
7.1. Normative References	17
7.2. Informative References	17

1. Introduction

This document defines an extension to the rpki-updown protocol, defined in [RFC6492], to provide support for transferring Internet Number Resources from one INR holder to another. The protocol supports asynchronous transfers of live or unused INRs. The scope of the protocol is limited to the transfer of Internet Number Resource within the Resource Public Key Infrastructure, defined in [RFC6480]. In support of this protocol, this document also defines a new signed object type, the Transfer Authorization Object (TAO), which makes use of the signed object format defined in [RFC6488].

Many of the messages in this protocol are identical to those in [RFC6488], and the result of the protocol, updated certificates published in the RPKI repository system [RFC6481], is the same for both protocols. To initiate a transfer, an INR holder, or source, creates a TAO and publishes it in its publication point. The TAO is a declaration of the proposed transfer, signed by the transfer source. The source communicates the location of the TAO to the INR recipient. Both entities then pursue the transfer independently, recursively requesting the transfer from their parents until the lowest common ancestor, the swing point is reached. The swing point acts as the ultimate arbiter of the transfer, although any Certification Authority (CA) involved in the transfer is able to deny the transfer. The protocol assumes that the source of the transfer, and the recipient have gained preliminary approval for the transfer, out-of-band (OOB), prior to publishing the TAO and initiating the protocol.

1.1. Terminology

Terms used in this document are:

"Internet Number Resource" (or "resource" or "INR") used in the context of this document to refer to Autonomous System (AS) numbers and IP version 4 or IP version 6 addresses.

"swing point" the lowest common ancestor (Certification Authority) of both the INR source and the INR recipient in the RPKI hierarchy. It is assumed that the swing point is neither the source nor the recipient.

"source" (or "INR source") the INR holder that initiates the transfer

"recipient" (or "INR recipient") the INR holder that is the destination of the transfer

"live" a live INR is a resource that is currently in use

The key words "MUST", "MUST NOT", "REQUIRED", "SHALL", "SHALL NOT", "SHOULD", "SHOULD NOT", "RECOMMENDED", "MAY", and "OPTIONAL" in this document are to be interpreted as described in BCP 14, [RFC2119].

2. Scope

This Resource Public Key Infrastructure (RPKI) INR transfer protocol defines a basic set of interactions that allows:

- o an INR holder to initiate the transfer of Internet Number Resources,
- o the INR source and INR recipient to pursue the transfer asynchronously,
- o and each Certification Authority (CA) along the path between the source and recipient (including the swing point) to validate and approve, or deny, any such transfer.

The resource allocation database and INR transfer policies of each CA along the path are authoritative when determining whether the resources in question may be transferred.

This protocol specification does not encompass:

- o the specification of interactions with the each CA's resource allocation database, nor the specification of a protocol to manage the publication repository.
- o transfers where the source or recipient is also the swing point. Both situations are already handled by rpki-updown as explained in Section 3.3.

3. Protocol Specifications

The INR source MUST initiate the transfer by creating and publishing the Transfer Authorization Object (TAO, see Section 3.4) at its publication point [RFC6481]. The URL of the TAO SHOULD be communicated to the transfer recipient, e.g., via email. Once the TAO is published, and the recipient has received the URL of the TAO, two separate processes begin: the first from the INR source to the swing point, the second from the transfer recipient to the swing point. These two processes proceed independently and recursively. The following steps occur between each parent and child along the specified paths in the hierarchy.

In both cases, when a CA receives an updated certificate from its immediate parent, it MUST promptly update the certificate for the child involved in the transfer. This certificate is published in its publication point and sent to the child using a transfer_response message (Section 3.8.1.2. If this CA is the INR source or INR recipient, no updates are necessary since receipt of the updated certificate indicates that the parent has updated the end point of the transfer. Similarly, when a CA receives an error message from a parent, the CA MUST forward the message code to its immediate child along the path towards the INR source or INR recipient.

Both the INR source and the INR recipient MUST NOT rekey during a transfer; their SKIs are captured in the TAO and the validity of the TAO requires the SKIs not change during the process. A new key would invalidate the TAO and require restarting the transfer process. To avoid this problem, the source SHOULD NOT initiate a transfer that is expected to take longer than the notAfter date in its, or the recipient's, CA certificate. The source should contact (OOB) the CA's along the path to receive an estimate of the time required to complete a transfer, to aid in making this determination.

The process described below is used for transferring either live or unused INRs. The process is identical for both types of transfers except where otherwise specified.

3.1. INR Source Path

The source MUST NOT request a transfer of any INRs that are delegated to one of the source's children (i.e., appear in a CA certificate issued by the source). This requirement avoids one way that a TAO that is valid at the beginning of a transfer could become invalid before the end of the transfer. In particular, in the instance where the source of this transfer is the swing point in another transfer, this prevents the swing point from transferring INRs to a different recipient than specified in the first transfer.

Along the path from the INR source to the swing point, with the INR source as the initial "child", the following messages MUST be transmitted in the specified order.

1. The child sends a transfer_request (Section 3.8.1.1) to its parent.
2. The parent confirms the validity of the transfer_request, responding with an error code 1203 for invalid requests. An invalid request cancels the transfer. A transfer_request is valid if all of the following are true:

- * The attached TAO is valid. See Section 3.4.2 for TAO validation steps.
 - * The TAO in the request is identical to the TAO published in the source's publication point.
 - * The transfer is allowed by the transfer policy of the CA
3. The parent replies with a `transfer_response` (Section 3.8.1.2). For transfers of unused INRs, the `transfer_response` contains an updated certificate, which **MUST** have the same INRs as the certificate it replaces, minus the INRs specified in the TAO. For live transfers, the `transfer_response` contains an error code 1104 response, indicating that the transfer is valid and being pursued asynchronously.
 4. The parent determines if it (the parent) is the swing point. See Section 3.3 for this procedure. If it is not the swing point, the parent repeats this process from step 1, acting as the child. If the parent is not the swing point but is a self-signed CA, an error code 1401 **MUST** be returned.

If, after an excessive wait, a child does not receive a response from its parent, the child **SHOULD** return error 1402 indicating a timeout. This error declares cancellation of the transfer request by the child, and **MUST** be propagated up **AND** down the path. This informs any parents waiting further up the path that the child is no longer waiting for an updated certificate, and indicates that the parent **MUST** time out as well. Ultimately, what constitutes an excessive wait is determined by each CA. However, it is **RECOMMENDED** that each CA not time out a transfer prior to the `notAfter` value in the TAO.

For live transfers, the source waits until the `notAfter` value in the TAO expires. If the recipient has successfully received the INRs at that point, the source **MUST** use the following process to relinquish control of the transferred INRs:

1. The child sends a `transfer_request` (Section 3.8.1.1) to the parent.
2. The parent confirms that the `transfer_request` matches a previous `transfer_request`, with the exception that the `notAfter` **MUST** be in the past. The parent responds with an error code 1203 for invalid requests. A `transfer_request` is valid if all of the following are true:

- * The attached TAO is valid, with the exception of the notAfter value which MUST be in the past. See Section 3.4.2 for TAO validation steps.
 - * The TAO in the request MUST be identical to the TAO published in the source's publication point.
3. The parent replies with a transfer_response (Section 3.8.1.2). This response MUST include an updated certificate which MUST have the same INRs as the certificate it replaces, minus the INRs specified in the TAO.
 4. The parent determines if it (the parent) is the swing point. See Section 3.3 for this procedure. If it is not the swing point, the parent repeats this process from step 1, acting as the child. If the parent is not the swing point but is a self-signed CA, an error code 1401 MUST be returned.

3.2. INR Recipient Path

A recipient will have multiple parents within the RPKI if it has received INR allocations from multiple sources. In such cases, the recipient MUST select the parent via which the resources will be received. The means by which a recipient makes this decision are outside the scope of this protocol. (INR transfers require OOB coordination among the affected organizations. This coordination is expected to provide the recipient with a basis for selecting a parent for the transfer.)

Along the path from the transfer recipient to the swing point, with the INR recipient as the initial "child", the following messages MUST be transmitted in the order specified below.

1. The child sends a transfer_request, Section 3.8.1.1, to the parent.
2. The parent confirms the validity of the transfer_request, responding with an error code 1203 for invalid requests. A transfer_request is valid only if all of the following are true:
 - * The attached TAO is valid. See Section 3.4.2 for TAO validation steps.
 - * The TAO in the request is identical to the TAO published in the source's publication point.
 - * The transfer is allowed by the transfer policy of the CA

3. The parent determines if it (the parent) is the swing point. See Section 3.3 for this procedure.
4. If it is not the swing point, the parent replies with a `transfer_response` containing an error code. If the parent is a self-signed CA and it is not the swing point, an error code 1401 MUST be returned. If the parent is not a self-signed CA, an error code 1104 response MUST be returned, indicating that the transfer is valid and being pursued asynchronously. The parent then repeats this process from step 1, acting as the child.

If, after an excessive wait, a child does not receive a response from its parent, the child SHOULD return error 1402 indicating a timeout. This error declares cancellation of the transfer request by the child, and MUST be propagated up AND down the path by each parent. See the previous section for a discussion of what constitutes "excessive".

During live transfers, CAs in the recipient path have an additional responsibility after receiving an updated certificate. The `overlapPeriod` field of the TAO MUST be less than that number of seconds from the current time to the `notAfter` value of the TAO. If this test fails, this CA MUST forward an error code 1403 up and down the path, ending the transfer. This minimizes the likelihood that the source and recipient do not have an adequate overlap in ownership of the INRs in question during a live transfer.

3.3. Swing Point

A CA determines that it is the swing point by verifying that both the INR source and the INR recipient SKIs, as defined in the TAO, are below the CA in the hierarchy. Because this determination is performed for both paths, starting at the source and the recipient, this will uniquely determine the swing point. This document does not cover the case where the swing point is the source or the recipient. If the swing point is the recipient, the INRs are being relinquished and returned to that organization. If the swing point is the source, the INRs are being assigned. This procedure is already accommodated by use of the up/down protocol. Because the RPKI hierarchy is intended to have a unique root, there should always exist a swing point.

The swing point MUST behave as follows:

1. Confirm that it is the swing point.
2. Confirm the validity and uniqueness of the Subject Key Identifiers (SKI) of the CAs (source and recipient) in the TAO.

3. Confirm that it controls the INRs to be transferred.
4. Wait to receive both transfer_requests, one from the path to the source and one from the path to the recipient.
5. Create an updated certificate for the CA on the path from the swing point to the transfer recipient. Publish this certificate in the swing point's publication point and send the updated certificate to the child CA using a transfer_response message. This updated certificate MUST have the same INRs as the certificate it replaces, plus the INRs specified in the TAO. (The swing point MUST still control the INRs being transferred, but this is a side effect of its normal certificate issuance process.)

Should a swing point receive an error code 1403 message from the CA in the recipient path, the swing point must forward the error code to the CA on the source path, indicating a cancellation of the transfer.

3.4. Transfer Authorization Object

The TAO is encapsulated in a CMS object as defined in [RFC6492] Section 3.1.

3.4.1. TAO Type

TAO OID TBD

3.4.2. TAO Validation

The TAO must be validated by each participant in the process. The creator of the TAO MUST validate the TAO after creation. All CAs that receive a Transfer Request MUST perform the following actions:

1. Determine that the TAO is valid as defined by the steps in [RFC6488] Section 3.
2. Verify that either the transferFromSKI or the transferToSKI (or both) correspond to CAs that are descendants of this CA

Note: This requires that the transfer recipient hold some address space and thus hold a valid CA Certificate before this process is initiated.

3. Verify that the transferFromSKI and the transferToSKI SKIs are valid, corresponding to the SKI extension of a CA within the RPKI, and unique, such that only one CA has an SKI extension that matches each of these values. (This check SHOULD be performed

using the RPKI data acquired by the participant in its role as a relying party [RFC6480].)

4. The parent of the source checks that the source holds the INRs in question. Each CA above that checks that the INRs are held by the CA that made the request.

3.5. ASN.1 Specification of the TAO

```
TransferAuthorization ::= SEQUENCE {  
    transferFromSKI OCTET STRING,  
    transferToSKI OCTET STRING,  
    ipAddrBlocks [0] IPAddrBlocks OPTIONAL,  
    asIdentifiers [1] ASIdentifiers OPTIONAL,  
    liveXfer BOOLEAN DEFAULT FALSE,  
    overlapPeriod INTEGER OPTIONAL  
}
```

Either ipAddrBlocks or asIdentifiers, or both, MUST be included.

3.5.1. transferFromSKI

The transferFromSKI MUST be equal to the SKI of the CA that holds the resources.

3.5.2. transferToSKI

The transferToSKI MUST be equal to the SKI in a valid CA within the RPKI.

3.5.3. ipAddrBlocks

IPAddrBlocks is specified in [RFC3779] Section 2. If the ipAddrBlocks attribute is included, it MUST NOT be empty and it MUST NOT have any resources marked as inherit.

3.5.4. asIdentifiers

ASIdentifiers is specified in [RFC3779] Section 3. If the asIdentifiers attribute is included, it MUST NOT be empty and the inherit flag MUST NOT be TRUE.

3.5.5. liveXfer

This flag is set TRUE only for a transfer of live resources.

3.5.6. overlapPeriod

overlapPeriod is the minimum number of seconds which the source and recipient MUST both hold the INRs. This field MUST hold a non-zero number for live transfers. The value MUST be omitted for transfers of unused space. Thus this field is present only if liveXfer is TRUE.

3.6. Common Message Format

This document defines version 2 of the Common Message Format for the up/down protocol. Version 1 is defined in [RFC6492]. The format in version 2 is identical to version 1, but with several added attributes, defined in Section 3.8, and one additional constraint defined in Section 3.7. The checks specified in [RFC6492] Section 3.2 still apply and MUST be applied.

3.7. End Entity Certificate Constraint

This section corresponds to Section 3.1.1.4 in [RFC6492]. The End Entity (EE) certificate that is required here MUST have its resources marked as inherit. This convention is imposed to ensure that this certificate remains valid during the life of the TAO before, during, and after the transfer takes place.

3.8. INR Transfer

3.8.1. Transfer

This query is used for all requests and responses made during a transfer. This includes messages between the initial sender and its parent, the receiver and its parent, and between each intermediate CA and its parent.

3.8.1.1. Request

The value of the message "type" element for this request is:

```
type="transfer_request"
```

Payload:

```
<request
tao_url="url">
[tao]
</request>
```

tao_url: value is the pointer to the location where the INR source has published the TAO.

[tao] value is the Base64 encoding of the DER-encoded TAO. After decoding, this object MUST be identical to the object published by the source in its publication point.

3.8.1.2. Response

The value of the message "type" element for this response is:

```
type="transfer_response"
```

Payload:

```
<class
tao_url="url"
cert_url="url"
resource_set_as="as resource set"
resource_set_ipv4="ipv4 resource set"
resource_set_ipv6="ipv6 resource set"
resource_set_notafter="datetime"
suggested_sia_head="[directory uri]">
<certificate cert_url="url"
req_resource_set_as="as resource set"
req_resource_set_ipv4="ipv4 resource set"
req_resource_set_ipv6="ipv6 resource set" >
[certificate]
</certificate>
<issuer>[issuer's certificate]</issuer>
</class>
```

In the case where the transfer is for live resources, not all responses will contain a certificate. For the CAs in the path with the INR source, an updated certificate, with the transferred INR removed, will be available once the transfer is complete and the INR source is prepared to relinquish control of the INRs. In contrast, the CAs along the path to the transfer recipient each receive a new certificate after the swing point receives and approves the messages from both the source and the recipient.

tao_url is identical to the tao_url in the request. The definition of all other attributes can be found in [RFC6492] Section 3.3.2.

3.8.2. Request-Not-Performed Response

This response is an extension of [RFC6492] Section 3.6. In addition to the error codes defined there, Error Code 1401 is used when a self-signed CA determines that it is not an ancestor of both the source and the recipient. This indicates a failure of the automated transfer and a manual transfer must take place.

3.8.3. Timeout Response

This response is an extension of [RFC6492] Section 3.6. In addition to the error codes defined there, Error Code 1402 is used when a CA determines that it has waited an excessive duration for a response from its parent. This indicates a failure of the transfer.

3.8.4. Overlap Failure Response

This response is an extension of [RFC6492] Section 3.6. In addition to the error codes defined there, Error Code 1403 is used when a CA in the recipient path determines that the overlapPeriod value is less than the number of seconds between the current time and the notAfter value in the TAO. This indicates a failure of the transfer.

3.9. XML Schema

The following is a RELAX NG compact form schema [ISO.19757-2.2003] describing version 2 of this protocol.

Note: As discussed in [W3C.REC-xml-names-20091208], "the namespace name, to serve its intended purpose, SHOULD have the characteristics of uniqueness and persistence. It is not a goal that it be directly usable for retrieval of a schema (if any exists)".

default namespace = "http://www.apnic.net/specs/rescerts/up-down/"

```
grammar {
  resource_set_as = xsd:string { maxLength="512000"
                                pattern="[\-,0-9]*" }
  resource_set_ip4 = xsd:string { maxLength="512000"
                                pattern="[\-,/.0-9]*" }
  resource_set_ip6 = xsd:string { maxLength="512000"
                                pattern="[\-,/:0-9a-fA-F]*" }

  class_name = xsd:token { minLength="1" maxLength="1024" }
  ski = xsd:token { minLength="27" maxLength="1024" }
  label = xsd:token { minLength="1" maxLength="1024" }
```

```

cert_url = xsd:string { minLength="10" maxLength="4096" }
base64_binary = xsd:base64Binary { minLength="4"
                                   maxLength="512000" }
tao_url = xsd:string { minLength="10" maxLength="4096" }

start = element message {
  attribute version { xsd:positiveInteger {
                                   maxInclusive="1" } },
  attribute sender { label },
  attribute recipient { label },
  payload
}

payload |= attribute type { "list" }, list_request
payload |= attribute type { "list_response" }, list_response
payload |= attribute type { "issue" }, issue_request
payload |= attribute type { "issue_response" }, issue_response
payload |= attribute type { "revoke" }, revoke_request
payload |= attribute type { "revoke_response" }, revoke_response
payload |= attribute type { "error_response" }, error_response
payload |= attribute type { "transfer_response" },
                                   transfer_response

list_request = empty
list_response = class*

class = element class {
  attribute class_name { class_name },
  attribute cert_url { cert_url },
  attribute resource_set_as { resource_set_as },
  attribute resource_set_ipv4 { resource_set_ipv4 },
  attribute resource_set_ipv6 { resource_set_ipv6 },
  attribute resource_set_notafter { xsd:dateTime },
  attribute suggested_sia_head { xsd:anyURI { maxLength="1024"
                                   pattern="rsync://.+"} }?,
  element certificate {
    attribute cert_url { cert_url },
    attribute req_resource_set_as { resource_set_as }?,
    attribute req_resource_set_ipv4 { resource_set_ipv4 }?,
    attribute req_resource_set_ipv6 { resource_set_ipv6 }?,
    base64_binary
  }*,
  element issuer { base64_binary }
}

issue_request = element request {
  attribute class_name { class_name },
  attribute req_resource_set_as { resource_set_as }?,

```

```
    attribute req_resource_set_ipv4 { resource_set_ip4 }?,
    attribute req_resource_set_ipv6 { resource_set_ip6 }?,
    base64_binary
  }
  issue_response = class

  revoke_request = revocation
  revoke_response = revocation

  revocation = element key {
    attribute class_name { class_name },
    attribute ski { ski }
  }

  error_response =
    element status { xsd:positiveInteger { maxInclusive="9999" } },
    element description { attribute xml:lang { xsd:language },
                          xsd:string { maxLength="1024" } }*
  }

  transfer_request = element request {
    attribute tao_url { tao_url },
    element tao { base64_binary }
  }

  transfer_response = element response {
    attribute tao_url { tao_url },
    attribute cert_url { cert_url },
    attribute resource_set_as { resource_set_as },
    attribute resource_set_ipv4 { resource_set_ip4 },
    attribute resource_set_ipv6 { resource_set_ip6 },
    attribute resource_set_notafter { xsd:dateTime },
    attribute suggested_sia_head { xsd:anyURI { maxLength="1024"
                                              pattern="rsync://.+"} }?,
    element certificate {
      attribute cert_url { cert_url },
      attribute req_resource_set_as { resource_set_as }?,
      attribute req_resource_set_ipv4 { resource_set_ip4 }?,
      attribute req_resource_set_ipv6 { resource_set_ip6 }?,
      base64_binary
    }*,
    element issuer { base64_binary }
  }
```

4. Security Considerations

The checks described at each stage are designed to ensure that these four security goals are met:

- o the TAO was generated by the indicated INR source, that source holds the INRs being transferred, and the TAO has not been modified by another party
- o the transfer recipient is the intended recipient of the resources as per the INR source
- o each CA that processes a transfer request either holds the resources being transferred, or it is on the path between the swing point and the transfer recipient
- o each CA along the path approved the transfer (or has rejected it)

Up/down protocol messages contain a time-based anti-reply feature, so replays of these signed messages can be detected. If a request message is redirected, a CA receiving it will detect and reject this because the request will not be from one of its children. A redirected response message also will be detected because the response will not be from the child's immediate parent. Because all messages (both requests and responses) are contained within a CMS object, the sender of a message is validated through signature verification.

For live transfers, the source initiates the relinquishment of the INRs that were transferred. If they fail to initiate the relinquishment in a timely manner, the recipient may choose to contact any or all of the source's ancestors (up to the swing point) to pursue a forced relinquishment of resources. Any legal or contractual processes used are outside the scope of this document.

5. IANA Considerations

An OID is requested for the TAO object defined above.

6. Acknowledgements

The author would like to acknowledge the valued contribution of Steve Kent for providing a top level description of the TAO protocol, David Mandelberg for his contributions to the security of the protocol, and the authors of the rpki-updown protocol ([RFC6492]) Geoff Huston, Robert Loomans, Byron Ellacott, and Rob Austein.

7. References

7.1. Normative References

- [RFC6492] Huston, G., Loomans, R., Ellacott, B., and R. Austein, "A Protocol for Provisioning Resource Certificates", RFC 6492, February 2012.
- [RFC2119] Bradner, S., "Key words for use in RFCs to Indicate Requirement Levels", BCP 14, RFC 2119, March 1997.
- [RFC3779] Lynn, C., Kent, S., and K. Seo, "X.509 Extensions for IP Addresses and AS Identifiers", RFC 3779, June 2004.
- [RFC6481] Huston, G., Loomans, R., and G. Michaelson, "A Profile for Resource Certificate Repository Structure", RFC 6481, February 2012.
- [RFC6488] Lepinski, M., Chi, A., and S. Kent, "Signed Object Template for the Resource Public Key Infrastructure (RPKI)", RFC 6488, February 2012.

7.2. Informative References

- [W3C.REC-xml-names-20091208]
Bray, T., Hollander, D., Layman, A., Tobin, R., and H. Thompson, "Namespaces in XML 1.0 (Third Edition)", World Wide Web Consortium Recommendation REC-xml-names-20091208, December 2009,
<<http://www.w3.org/TR/2009/REC-xml-names-20091208>>.
- [RFC6480] Lepinski, M. and S. Kent, "An Infrastructure to Support Secure Internet Routing", RFC 6480, February 2012.
- [ISO.19757-2.2003]
International Organization for Standardization,
"Information technology -- Document Schema Definition Language (DSDL) -- Part 2: Regular-grammar-based validation -- RELAX NG", ISO International Standard 19757-2, December 2003.

Author's Address

Internet-Draft

Transfer Authorization Object

February 2014

Edric Barnes
BBN Technologies
10 Moulton St
Cambridge, MA
US

EMail: ebarnes@bbn.com

Secure Inter-Domain Routing
Internet-Draft
Intended status: Best Current Practice
Expires: August 14, 2014

D. Mandelberg
BBN Technologies
February 10, 2014

Simplified Local internet nUmber Resource Management with the RPKI
draft-dseomn-sidr-slurm-00

Abstract

The Resource Public Key Infrastructure (RPKI) is a global authorization infrastructure that allows the holder of Internet Number Resources (INRs) to make verifiable statements about those resources. Internet Service Providers (ISPs) can use the RPKI to validate BGP route origination assertions. Some ISPs locally use BGP with private address space or private AS numbers (see RFC6890). These local BGP routes cannot be verified by the global RPKI, and SHOULD be considered invalid based on the global RPKI (see RFC6491). The mechanisms described below provide ISPs with a way to make local assertions about private (reserved) INRs while using the RPKI's assertions about all other INRs.

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 August 14, 2014.

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	2
1.1. Terminology	3
2. Validation Output Filtering	3
3. Locally Adding Assertions	3
4. Configuring SLURM	4
5. Combining Mechanisms	4
6. IANA Considerations	5
7. Security Considerations	5
8. Acknowledgements	5
9. References	5
9.1. Informative References	5
9.2. Normative References	6
Appendix A. Example SLURM File	6
Author's Address	7

1. Introduction

The Resource Public Key Infrastructure (RPKI) is a global authorization infrastructure that allows the holder of Internet Number Resources (INRs) to make verifiable statements about those resources. For example, the holder of a block of IP(v4 or v6) addresses can issue a Route Origination Authorization (ROA) [RFC6482] to authorize an Autonomous System (AS) to originate routes for that block.

Internet Service Providers (ISPs) can then use the RPKI to validate BGP routes. However, some ISPs locally use BGP with private address space ([RFC1918], [RFC4193], [RFC6598]) or private AS numbers ([RFC1930], [RFC6996]). These local BGP routes cannot be verified by the global RPKI, and SHOULD be considered invalid when using the RPKI. For example, [RFC6491] recommends the creation of ROAs that would invalidate routes for reserved and unallocated address space.

This document specifies two new mechanisms to enable ISPs to make local assertions about private INRs while using the RPKI's assertions about all other INRs. Both mechanisms are specified in terms of abstract sets of assertions. For Origin Validation [RFC6483], an assertion is a tuple of {IP prefix, prefix length, maximum length, AS number} as used by rpki-rtr [RFC6810]. Output Filtering, described

in Section 2, filters out any assertions by the RPKI about locally reserved INRs. Locally Adding Assertions, described in Section 3, adds local assertions about locally reserved INRs. Note that both of these mechanisms can later be extended to cover any assertions made by the RPKI for use in BGPSEC [I-D.ietf-sidr-bgpsec-protocol].

In general, the primary output of an RPKI relying party is the data it sends to routers over the rpki-rtr protocol. The rpki-rtr protocol enables routers to query a relying party for all Origin Validation assertions it knows about (Reset Query) or for an update of only the changes in Origin Validation assertions (Serial Query). The mechanisms specified in this document are to be applied to the result set for a Reset Query, and to both the old and new sets that are compared for a Serial Query. Relying party software MAY modify other forms of output in comparable ways, but that is outside the scope of this document.

This document is intended to supersede [I-D.ietf-sidr-ltamgmt] while focusing only on local management of private INRs. Another draft [I-D.kent-sidr-suspenders] focuses on the other aspects of local management.

1.1. Terminology

The key words "MUST", "MUST NOT", "REQUIRED", "SHALL", "SHALL NOT", "SHOULD", "SHOULD NOT", "RECOMMENDED", "MAY", and "OPTIONAL" in this document are to be interpreted as described in [RFC2119].

2. Validation Output Filtering

To prevent the global RPKI from affecting routes with locally reserved INRs, a relying party is locally configured with a list of IP prefixes and/or AS numbers that are used locally, and taken from the reserved INR spaces. Any Origin Validation assertions where the IP prefix is equal to or subsumed by a locally reserved IP prefix, are removed from the relying party's output. Any Origin Validation assertions where the IP prefix contains a locally reserved IP prefix are removed and the relying party software SHOULD issue a warning.

3. Locally Adding Assertions

Each relying party is locally configured with a (possibly empty) list of Origin Validation assertions. This list is added to the relying party's output.

4. Configuring SLURM

Relying party software SHOULD support the following configuration format for Validation Output Filtering and Locally Adding Assertions. The format is defined using the Augmented Backus-Naur Form (ABNF) notation and core rules from [RFC5234] and the rules <IPv4address> and <IPv6address> from Appendix A of [RFC3986]. Each command specifies an INR to use for Validation Output Filtering. Each <add> command specifies an assertion to use for Locally Adding Assertions. See Appendix A for an example SLURM file.

```
SLURMFile = header *line

header = %x53.4c.55.52.4d SP "1.0" CRLF ; "SLURM 1.0"

line = *WSP [comment] CRLF
      / *WSP command [ 1*WSP [comment] ] CRLF

comment = "#" *(VCHAR / WSP)

command = add / del

add = %x61.64.64 1*WSP IPprefixMaxLen 1*WSP ASnum

del = %x64.65.6c 1*WSP inr

inr = IPprefix / ASnum

IPprefix = IPv4prefix / IPv6prefix

IPprefixMaxLen = IPv4prefixMaxLen / IPv6prefixMaxLen

IPv4prefix = IPv4address "/" 1*2DIGIT

IPv6prefix = IPv6address "/" 1*3DIGIT

; In the following two rules, if the maximum length component is
; missing, it is treated as equal to the prefix length.
IPv4prefixMaxLen = IPv4prefix ["-" 1*2DIGIT]
IPv6prefixMaxLen = IPv6prefix ["-" 1*3DIGIT]

ASnum = 1*DIGIT
```

5. Combining Mechanisms

In the typical use case, a relying party uses both output filtering and locally added assertions. In this case, the resulting assertions MUST be the same as if output filtering were performed before locally

adding assertions. I.e., locally added assertions MUST NOT be removed by output filtering.

If a relying party chooses to use both SLURM and Suspenders [I-D.kent-sidr-suspenders], the SLURM mechanisms MUST be performed on the output of Suspenders.

6. IANA Considerations

TBD

7. Security Considerations

The mechanisms described in this document provide an ISP additional control over its own network. Care should be taken in how that control is used.

8. Acknowledgements

The author would like to thank Stephen Kent for his guidance and detailed reviews of this document.

9. References

9.1. Informative References

- [I-D.ietf-sidr-bgpsec-protocol]
Lepinski, M., "BGPSEC Protocol Specification", draft-ietf-sidr-bgpsec-protocol-08 (work in progress), November 2013.
- [I-D.ietf-sidr-ltamgmt]
Reynolds, M., Kent, S., and M. Lepinski, "Local Trust Anchor Management for the Resource Public Key Infrastructure", draft-ietf-sidr-ltamgmt-08 (work in progress), April 2013.
- [I-D.kent-sidr-suspenders]
Kent, S. and D. Mandelberg, "Suspenders: A Fail-safe Mechanism for the RPKI", draft-kent-sidr-suspenders-00 (work in progress), September 2013.
- [RFC1918] Rekhter, Y., Moskowitz, R., Karrenberg, D., Groot, G., and E. Lear, "Address Allocation for Private Internets", BCP 5, RFC 1918, February 1996.
- [RFC1930] Hawkinson, J. and T. Bates, "Guidelines for creation, selection, and registration of an Autonomous System (AS)", BCP 6, RFC 1930, March 1996.

- [RFC4193] Hinden, R. and B. Haberman, "Unique Local IPv6 Unicast Addresses", RFC 4193, October 2005.
- [RFC6482] Lepinski, M., Kent, S., and D. Kong, "A Profile for Route Origin Authorizations (ROAs)", RFC 6482, February 2012.
- [RFC6483] Huston, G. and G. Michaelson, "Validation of Route Origination Using the Resource Certificate Public Key Infrastructure (PKI) and Route Origin Authorizations (ROAs)", RFC 6483, February 2012.
- [RFC6491] Manderson, T., Vegoda, L., and S. Kent, "Resource Public Key Infrastructure (RPKI) Objects Issued by IANA", RFC 6491, February 2012.
- [RFC6598] Weil, J., Kuarsingh, V., Donley, C., Liljenstolpe, C., and M. Azinger, "IANA-Reserved IPv4 Prefix for Shared Address Space", BCP 153, RFC 6598, April 2012.
- [RFC6810] Bush, R. and R. Austein, "The Resource Public Key Infrastructure (RPKI) to Router Protocol", RFC 6810, January 2013.
- [RFC6890] Cotton, M., Vegoda, L., Bonica, R., and B. Haberman, "Special-Purpose IP Address Registries", BCP 153, RFC 6890, April 2013.
- [RFC6996] Mitchell, J., "Autonomous System (AS) Reservation for Private Use", BCP 6, RFC 6996, July 2013.

9.2. Normative References

- [RFC2119] Bradner, S., "Key words for use in RFCs to Indicate Requirement Levels", BCP 14, RFC 2119, March 1997.
- [RFC3986] Berners-Lee, T., Fielding, R., and L. Masinter, "Uniform Resource Identifier (URI): Generic Syntax", STD 66, RFC 3986, January 2005.
- [RFC5234] Crocker, D. and P. Overell, "Augmented BNF for Syntax Specifications: ABNF", STD 68, RFC 5234, January 2008.

Appendix A. Example SLURM File

SLURM 1.0

```
# Reserve 192.0.2.0/24 and 2001:DB8::/32 for local use.
del 192.0.2.0/24
del 2001:DB8::/32

# Allow either 65536 or 65537 to originate routes to 192.0.2.0/24.
add 192.0.2.0/24 65536
add 192.0.2.0/24 65537

add 2001:DB8::/48-52 65536 # 65536 originates 2001:DB8::/48 and
                        # sub-prefixes down to length 52.
add 2001:DB8:0:42::/64 65537 # However, 65537 originates
                        # 2001:DB8:0:42::/64.
add 2001:DB8:1::/48 65537 # 65537 also originates 2001:DB8:1::/48
```

Author's Address

David Mandelberg
BBN Technologies
10 Moulton St.
Cambridge, MA 02138
US

Email: david@mandelberg.org

Secure Inter-Domain Routing
Internet-Draft
Intended status: Best Current Practice
Expires: November 14, 2015

D. Mandelberg
BBN Technologies
May 13, 2015

Simplified Local internet nUmber Resource Management with the RPKI
draft-dseomn-sidr-slurm-02

Abstract

The Resource Public Key Infrastructure (RPKI) is a global authorization infrastructure that allows the holder of Internet Number Resources (INRs) to make verifiable statements about those resources. Network operators, e.g., Internet Service Providers (ISPs), can use the RPKI to validate BGP route origination assertions. In the future, ISPs also will be able to use the RPKI to validate the path of a BGP route. Some ISPs locally use BGP with private address space or private AS numbers (see RFC6890). These local BGP routes cannot be verified by the global RPKI, and SHOULD be considered invalid based on the global RPKI (see RFC6491). The mechanisms described below provide ISPs with a way to make local assertions about private (reserved) INRs while using the RPKI's assertions about all other INRs.

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 November 14, 2015.

Copyright Notice

Copyright (c) 2015 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	2
1.1. Terminology	4
2. Validation Output Filtering	4
3. Locally Adding Assertions	4
4. Configuring SLURM	4
5. Combining Mechanisms	7
6. IANA Considerations	7
7. Security Considerations	8
8. Acknowledgements	8
9. References	8
9.1. Informative References	8
9.2. Normative References	9
Appendix A. Example SLURM File	10
Author's Address	11

1. Introduction

The Resource Public Key Infrastructure (RPKI) is a global authorization infrastructure that allows the holder of Internet Number Resources (INRs) to make verifiable statements about those resources. For example, the holder of a block of IP(v4 or v6) addresses can issue a Route Origination Authorization (ROA) [RFC6482] to authorize an Autonomous System (AS) to originate routes for that block.

Internet Service Providers (ISPs) can then use the RPKI to validate BGP routes. (Validation of the origin of a route is described in [RFC6483], and validation of the path of a route is described in [I-D.ietf-sidr-bgpsec-overview].) However, some ISPs locally use BGP with private address space ([RFC1918], [RFC4193], [RFC6598]) or private AS numbers ([RFC1930], [RFC6996]). These local BGP routes cannot be verified by the global RPKI, and SHOULD be considered invalid when using the RPKI. For example, [RFC6491] recommends the creation of ROAs that would invalidate routes for reserved and unallocated address space.

This document specifies two new mechanisms to enable ISPs to make local assertions about some INRs while using the RPKI's assertions about all other INRs. These mechanisms support the second and third use cases in [I-D.ietf-sidr-lta-use-cases]. The second use case describes use of [RFC1918] addresses or use of public address space not allocated to the ISP that is using it. The third use case describes a situation in which an ISP publishes a variant of the RPKI hierarchy (for its customers). In this variant some prefixes and/or AS numbers are different from what the RPKI repository system presents to the general ISP population. The result is that routes for consumers of this variant hierarchy will be re-directed (via routing).

Both mechanisms are specified in terms of abstract sets of assertions. For Origin Validation [RFC6483], an assertion is a tuple of {IP prefix, prefix length, maximum length, AS number} as used by rpki-rtr version 0 [RFC6810] and version 1 [I-D.ietf-sidr-rpki-rtr-rfc6810-bis]. For BGPsec [I-D.ietf-sidr-bgpsec-overview], an assertion is a tuple of {AS number, subject key identifier, router public key} as used by rpki-rtr version 1. Output Filtering, described in Section 2, filters out any assertions by the RPKI about locally reserved INRs. Locally Adding Assertions, described in Section 3, adds local assertions about locally reserved INRs. The combination of both mechanisms is described in Section 5.

To ensure local consistency, the effect of SLURM MUST be atomic. That is, the output of the relying party must be either the same as if SLURM were not used, or it must reflect the entire SLURM configuration. For an example of why this is required, consider the case of two local routes for the same prefix but different origin AS numbers. Both routes are configured with Locally Adding Assertions. If neither addition occurs, then both routes could be in the unknown state [RFC6483]. If both additions occur then both routes would be in the valid state. However, if one addition occurs and the other does not, then one could be invalid while the other is valid.

In general, the primary output of an RPKI relying party is the data it sends to routers over the rpki-rtr protocol. The rpki-rtr protocol enables routers to query a relying party for all assertions it knows about (Reset Query) or for an update of only the changes in assertions (Serial Query). The mechanisms specified in this document are to be applied to the result set for a Reset Query, and to both the old and new sets that are compared for a Serial Query. Relying party software MAY modify other forms of output in comparable ways, but that is outside the scope of this document.

This document is intended to supersede [I-D.ietf-sidr-ltamgmt] while focusing only on local management of private INRs. Another draft [I-D.kent-sidr-suspenders] focuses on the other aspects of local management.

1.1. Terminology

The key words "MUST", "MUST NOT", "REQUIRED", "SHALL", "SHALL NOT", "SHOULD", "SHOULD NOT", "RECOMMENDED", "MAY", and "OPTIONAL" in this document are to be interpreted as described in [RFC2119].

2. Validation Output Filtering

To prevent the global RPKI from affecting routes with locally reserved INRs, a relying party may be locally configured with a list of IP prefixes and/or AS numbers that are used locally, and taken from reserved INR spaces. Any Origin Validation assertions where the IP prefix is equal to or subsumed by a locally reserved IP prefix, are removed from the relying party's output. Any Origin Validation assertions where the IP prefix contains a locally reserved IP prefix are removed; the relying party software SHOULD issue a warning when this action is taken. (Note that an Origin Validation assertion is not removed due to its AS number matching a locally reserved AS number.) Any BGPsec assertion where the AS number is equal to a locally reserved AS number is removed from the relying party's output.

3. Locally Adding Assertions

Each relying party is locally configured with a (possibly empty) list of assertions. This list is added to the relying party's output.

4. Configuring SLURM

Relying party software SHOULD support the following configuration format for Validation Output Filtering and Locally Adding Assertions. The format is defined using the Augmented Backus-Naur Form (ABNF) notation and core rules from [RFC5234] and the rules <IPv4address> and <IPv6address> from Appendix A of [RFC3986]. See Appendix A for an example SLURM file.

A SLURM configuration file, <SLURMFile>, consists of a head and a body. The head identifies the file as a SLURM configuration file, specifies the version of SLURM for which the file was written, and optionally contains other information described below. The body contains the configuration for Validation Output Filtering and Locally Adding Assertions.

```

SLURMFile = head body

head = firstLine *(commentLine / headLine)

body = *(commentLine / bodyLine)

firstLine = %x53.4c.55.52.4d SP "1.0" EOL ; "SLURM 1.0"

commentLine = *WSP [comment] EOL

headLine = *WSP headCommand [ 1*WSP [comment] ] EOL

bodyLine = *WSP bodyCommand [ 1*WSP [comment] ] EOL

comment = "#" *(VCHAR / WSP)

EOL = CRLF / LF

```

The head may specify a target. If present, the target string identifies the environment in which the SLURM file is intended to be used. The meaning of the target string, if any, is determined by the user. If a target is present, a relying party SHOULD verify that that the target is an acceptable value, and reject the SLURM file if the target is not acceptable. For example, the relying party could be configured to accept SLURM files only if they do not specify a target, have a target value of "hostname=rpki.example.com", or have a target value of "as=65536". If more than one target line is present, all targets must be acceptable to the RP.

```

headCommand = target

target =
    %x74.61.72.67.65.74 1*WSP ; "target"
    1*VCHAR

```

The body contains zero or more configuration lines for Validation Output Filtering and Locally Adding Assertions. Each command specifies an INR to use for Validation Output Filtering. Each <add> command specifies an assertion to use for Locally Adding Assertions.

```

bodyCommand = add / del

add =
    %x61.64.64 1*WSP ; "add"
    addItem

del =
    %x64.65.6c 1*WSP ; "del"

```

```
delItem

addItem = addItemPrefixAS / addItemASKey

; Add a mapping from a prefix and max length to an AS number.
addItemPrefixAS =
    %x6f.72.69.67.69.6e.61.74.69.6f.6e 1*WSP ; "origination"
    IPprefixMaxLen 1*WSP
    ASnum

; Add a mapping from an AS number to a router public key.
addItemASKey =
    %x62.67.70.73.65.63 1*WSP ; "bgpsec"
    ASnum 1*WSP
    RouterSKI 1*WSP
    RouterPubKey

delItem = delItemPrefix / delItemAS

; Filter prefix-AS mappings, using the given prefix
delItemPrefix =
    %x6f.72.69.67.69.6e.61.74.69.6f.6e 1*WSP ; "origination"
    IPprefix

; Filter AS-key mappings for the given AS
delItemAS =
    %x62.67.70.73.65.63 1*WSP ; "bgpsec"
    ASnum

IPprefix = IPv4prefix / IPv6prefix

IPprefixMaxLen = IPv4prefixMaxLen / IPv6prefixMaxLen

IPv4prefix = IPv4address "/" 1*2DIGIT
IPv6prefix = IPv6address "/" 1*3DIGIT

; In the following two rules, if the maximum length component is
; missing, it is treated as equal to the prefix length.
IPv4prefixMaxLen = IPv4prefix ["-" 1*2DIGIT]
IPv6prefixMaxLen = IPv6prefix ["-" 1*3DIGIT]

ASnum = 1*DIGIT

; This is the Base64 [RFC4648] encoding of a router certificate's
; Subject Key Identifier, as described in
; [I-D.ietf-sidr-bgpsec-pki-profiles] and [RFC6487]. This is the
; value of the ASN.1 OCTET STRING without the ASN.1 tag or length
; fields.
```

RouterSKI = Base64

; This is the Base64 [RFC4648] encoding of a router public key's
; subjectPublicKeyInfo value, as described in
; [I-D.ietf-sidr-bgpsec-algs]. This is the full ASN.1 DER encoding
; of the subjectPublicKeyInfo, including the ASN.1 tag and length
; values of the subjectPublicKeyInfo SEQUENCE.
RouterPubKey = Base64

Base64 = 1*(ALPHA / DIGIT / "+" / "/") 0*2"=

An implementation MAY support the concurrent use of multiple SLURM files. In this case, the resulting inputs to Validation Output Filtering and Locally Adding Assertions are the respective unions of the inputs from each file. The typical use case for multiple files is when the files have distinct scopes. For example, an organization may belong to two separate networks that use different private-use IP prefixes and AS numbers. To detect conflict between multiple SLURM files, a relying party SHOULD issue a warning in the following cases:

1. There may be conflicting changes to Origin Validation assertions if there exists an IP address X and distinct SLURM files Y,Z such that X is contained by any prefix in any <addItemPrefixAS> or <delItemPrefix> in file Y and X is contained by any prefix in any <addItemPrefixAS> or <delItemPrefix> in file Z.
2. There may be conflicting changes to BGPsec assertions if there exists an AS number X and distinct SLURM files Y,Z such that X is used in any <addItemASKey> or <delItemAS> in file Y and X is used in any <addItemASKey> or <delItemAS> in file Z.

5. Combining Mechanisms

In the typical use case, a relying party uses both output filtering and locally added assertions. In this case, the resulting assertions MUST be the same as if output filtering were performed before locally adding assertions. I.e., locally added assertions MUST NOT be removed by output filtering.

If a relying party chooses to use both SLURM and Suspenders [I-D.kent-sidr-suspenders], the SLURM mechanisms MUST be performed on the output of Suspenders.

6. IANA Considerations

TBD

7. Security Considerations

The mechanisms described in this document provide a network operator with additional ways to control its own network while making use of RPKI data. These mechanisms are applied only locally; they do not influence how other network operators interpret RPKI data. Nonetheless, care should be taken in how these mechanisms are employed.

8. Acknowledgements

The author would like to thank Stephen Kent for his guidance and detailed reviews of this document. Thanks go to Wesley Wang for the idea behind the target command, to Declan Ma for the idea behind use of multiple SLURM files, and to Richard Hansen for his careful reviews.

9. References

9.1. Informative References

- [I-D.ietf-sidr-bgpsec-overview]
Lepinski, M. and S. Turner, "An Overview of BGPsec", draft-ietf-sidr-bgpsec-overview-06 (work in progress), January 2015.
- [I-D.ietf-sidr-lta-use-cases]
Bush, R., "RPKI Local Trust Anchor Use Cases", draft-ietf-sidr-lta-use-cases-02 (work in progress), December 2014.
- [I-D.ietf-sidr-ltamgmt]
Reynolds, M., Kent, S., and M. Lepinski, "Local Trust Anchor Management for the Resource Public Key Infrastructure", draft-ietf-sidr-ltamgmt-08 (work in progress), April 2013.
- [I-D.ietf-sidr-rpki-rtr-rfc6810-bis]
Bush, R. and R. Austein, "The Resource Public Key Infrastructure (RPKI) to Router Protocol", draft-ietf-sidr-rpki-rtr-rfc6810-bis-03 (work in progress), March 2015.
- [I-D.kent-sidr-suspenders]
Kent, S. and D. Mandelberg, "Suspenders: A Fail-safe Mechanism for the RPKI", draft-kent-sidr-suspenders-03 (work in progress), April 2015.

- [RFC1918] Rekhter, Y., Moskowitz, R., Karrenberg, D., Groot, G., and E. Lear, "Address Allocation for Private Internets", BCP 5, RFC 1918, February 1996.
- [RFC1930] Hawkinson, J. and T. Bates, "Guidelines for creation, selection, and registration of an Autonomous System (AS)", BCP 6, RFC 1930, March 1996.
- [RFC4193] Hinden, R. and B. Haberman, "Unique Local IPv6 Unicast Addresses", RFC 4193, October 2005.
- [RFC6482] Lepinski, M., Kent, S., and D. Kong, "A Profile for Route Origin Authorizations (ROAs)", RFC 6482, February 2012.
- [RFC6483] Huston, G. and G. Michaelson, "Validation of Route Origination Using the Resource Certificate Public Key Infrastructure (PKI) and Route Origin Authorizations (ROAs)", RFC 6483, February 2012.
- [RFC6491] Manderson, T., Vegoda, L., and S. Kent, "Resource Public Key Infrastructure (RPKI) Objects Issued by IANA", RFC 6491, February 2012.
- [RFC6598] Weil, J., Kuarsingh, V., Donley, C., Liljenstolpe, C., and M. Azinger, "IANA-Reserved IPv4 Prefix for Shared Address Space", BCP 153, RFC 6598, April 2012.
- [RFC6810] Bush, R. and R. Austein, "The Resource Public Key Infrastructure (RPKI) to Router Protocol", RFC 6810, January 2013.
- [RFC6890] Cotton, M., Vegoda, L., Bonica, R., and B. Haberman, "Special-Purpose IP Address Registries", BCP 153, RFC 6890, April 2013.
- [RFC6996] Mitchell, J., "Autonomous System (AS) Reservation for Private Use", BCP 6, RFC 6996, July 2013.

9.2. Normative References

- [I-D.ietf-sidr-bgpsec-algs]
Turner, S., "BGP Algorithms, Key Formats, & Signature Formats", draft-ietf-sidr-bgpsec-algs-09 (work in progress), January 2015.

- [I-D.ietf-sidr-bgpsec-pki-profiles]
Reynolds, M., Turner, S., and S. Kent, "A Profile for BGPSEC Router Certificates, Certificate Revocation Lists, and Certification Requests", draft-ietf-sidr-bgpsec-pki-profiles-10 (work in progress), January 2015.
- [RFC2119] Bradner, S., "Key words for use in RFCs to Indicate Requirement Levels", BCP 14, RFC 2119, March 1997.
- [RFC3986] Berners-Lee, T., Fielding, R., and L. Masinter, "Uniform Resource Identifier (URI): Generic Syntax", STD 66, RFC 3986, January 2005.
- [RFC4648] Josefsson, S., "The Base16, Base32, and Base64 Data Encodings", RFC 4648, October 2006.
- [RFC5234] Crocker, D. and P. Overell, "Augmented BNF for Syntax Specifications: ABNF", STD 68, RFC 5234, January 2008.
- [RFC6487] Huston, G., Michaelson, G., and R. Loomans, "A Profile for X.509 PKIX Resource Certificates", RFC 6487, February 2012.

Appendix A. Example SLURM File

SLURM 1.0

```
# This file is only intended to be used on a relying party running
# on rpki.example.com.
target hostname=rpki.example.com # this is a comment

# Reserve IP prefixes for local use.
del origination 10.0.0.0/24
del origination fd0b:ddld:2dcc::/48

# Reserve AS numbers for local use.
del bgpsec 64512
del bgpsec 64513

# Allow either 64512 or 64513 to originate routes to 10.0.0.0/24.
add origination 10.0.0.0/24 64512
add origination 10.0.0.0/24 64513

# 64512 originates fd0b:ddld:2dcc::/52 and sub-prefixes up to length
# 56.
add origination fd0b:ddld:2dcc::/52-56 64512

# However, 64513 originates fd0b:ddld:2dcc:42::/64.
add origination fd0b:ddld:2dcc:42::/64 64513

# 64513 also originates fd0b:ddld:2dcc:100::/52
add origination fd0b:ddld:2dcc:100::/52 64513

# Authorize router keys to sign BGPsec paths on behalf of the
# specified ASes. Note that the Base64 strings used in this
# example are not valid SKIs or router public keys, due to line
# length restrictions in RFCs.
add bgpsec 64512 Zm9v VGHpcyBpcyBub3QgYSByb3V0ZXIgcHVibGljIGtleQ==
add bgpsec 64512 YmFy b3IgcYSBmbG9jayBvZiBkdWNNrcw==
add bgpsec 64513 YWJj bWF5YmUgYSBkaWZmZXJlbnQgYXZpYW4gY2Fycmllcj8=
```

Author's Address

David Mandelberg
BBN Technologies
10 Moulton St.
Cambridge, MA 02138
US

Email: david@mandelberg.org

Network Working Group
Internet-Draft
Intended status: Informational
Expires: August 18, 2014

F. Mejia, Ed.
AEPROVI
R. Gagliano
A. Retana
Cisco Systems
C. Martinez
G. Rada
LACNIC
February 14, 2014

Implementing RPKI-based origin validation one country at a time. The
Ecuadorian case study.

draft-fmejia-opsec-origin-a-country-00

Abstract

One possible deployment strategy for BGP origin validation based on the Resource Public Key Infrastructure (RPKI) is the construction of islands of trust. This document describes the authors' experience deploying and maintaining a BGP origin validation island of trust in Ecuador.

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 August 18, 2014.

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	2
1.1. Policer Network	3
1.2. The resource holders	4
1.3. RPKI certificate authorities and repository	5
1.4. The technical support	5
2. Objective	5
3. Planning	6
3.1. RPKI-based origin validation support	6
3.2. Deploying a RPKI cache into the network	7
3.3. Populating the RPKI database	7
3.4. Action to take with NotFound and Invalid prefixes	8
4. Deployment	8
4.1. RPKI Validation servers	9
4.2. Origin validation setting	10
5. Training and RPKI signing event	11
6. Outcome and post-event activities	11
7. Lessons learned and best practices	12
8. IANA Considerations	13
9. Security Considerations	13
10. Acknowledgements	13
11. Informative References	13
Appendix A. Router configuration templates	14
Authors' Addresses	17

1. Introduction

BGP origin validation based on RPKI [RFC6811] is in early stages of deployment. As with other new technologies, there are impediments to its global adoption as its full value is not yet perceived. Particularly, RPKI based origin validation involves on one side the creation of a large enough set of signed objects and on the other side the application of network policies based on these signed objects by network operators. An operator that does not see a large enough set of signed objects in the RPKI repository system is not encouraged to implement these set of policies. Conversely, IP address space resource holders that are not required by network operators (i.e. transit providers, peers or operators community in general) to create and maintain their RPKI objects have little incentive to do so.

To overcome this bootstrap problem, it is necessary to create a success story that brings enough value to both: network operators and resource holders. Moreover, one possible strategy for the adoption of a security technology is the creation of islands of trust where the technology is fully deployed in a reduced environment. In this direction, some organizations carried forward a full implementation of an island of trust in Ecuador. This was a multi stakeholder project where each party (resource holders, an Internet Exchange Point manager, a Regional Internet Registry and an equipment manufacturer) contributed to its success.

This document describes the experience of implementing RPKI-based origin validation in Ecuador and it is expected to be an useful guide to start other similar projects.

Below, it is described the different roles in the project and the involved parties.

1.1. Policer Network

In this document, the "Policer Network" is the networking infrastructure where the origin validation based on RPKI will be deployed to apply polices on BGP announcements. NAP.EC (www.nap.ec) was selected for this role.

NAP.EC is the Internet Exchange Point (IXP) in Ecuador with two Points of Presence (POPs): Quito (UIO) and Guayaquil (GYE). It has a BGP route-server in each location with a mandatory multilateral routing policy (i.e. all participants have a BGP session to the route-server). Each location uses a different IP address block and Autonomous System Number (ASN). NAP.EC is a meeting point where many organizations (Internet Service Providers, content providers, root servers, etc.) exchange routing information.

The participants connected to NAP.EC announce almost 100% of the total address space used in Ecuador (be believe it is 100% but we cannot be certain though). In some cases they announce their own address space and in some cases they are transit providers for their customers' resources.

AEPROVI (www.aeprovi.org.ec) manages the NAP.EC infrastructure. It is a non-profit organization, based on membership and brings together around 30 Ecuadorian ICT-related companies. AEPROVI also has an excellent reputation as an innovator in the local networking community thanks to projects such as IPv6 adoption and CDN cache servers hosting. These projects have given the local community concrete value and have build the trust on the team that manages the local IXP.

Thanks to this trust, and to the fact that all local BGP announcements are performed through the route servers, NAP.EC is uniquely positioned to become the "policer network" for this project. At the same time, it can be said that implementing origin validation at the NAP.EC route servers is equivalent to implementing it for all inter-domain routing in the country.

1.2. The resource holders

In this document, an organization which operates their own IP prefixes is called resource holder or simply the holder. They may have resources allocated/assigned from a Regional Internet Registry (RIR) and/or legacy resources (if the allocation was done before RIR formation). Resource holders are responsible for creating the RPKI signed objects for this project.

NAP.EC routing tables involve a number of holders, including organizations like Internet Service Providers, content providers, universities, .ec domain administrator and root servers. Most of them are Ecuadorian companies and have received IP resources only from LACNIC, but some have both RIR and legacy resources. Moreover, a few holders are foreign companies and their resources are legacy or from other RIRs (e.g. root servers and content providers).

Not all resource holders are directly connected to the NAP.EC fabric; some have IP address resources but not an ASN and some others are small networks that receive traffic from other bigger networks. In this case their IP address prefixes are announced by their transit providers. One of the main challenges for this project was to identify all the resource holders that needed to be contacted and to encourage network administrators from these organizations to participate.

In addition, some resource holders are part of a larger (and sometimes international) organization, with strong change management processes. This means that any change on their configurations needed to be planned ahead of time and consulted outside of the country.

In NAP.EC - UIO, the routing table includes prefixes used in Ecuador and other countries.

In NAP.EC - GYE, the routing table includes prefixes from companies operating only in Ecuador.

For the project, the target was limited to prefixes used in Ecuador by Ecuadorian holders that had received resources from LACNIC until mid-2013.

1.3. RPKI certificate authorities and repository

The five Regional Internet Registries (RIRs) have a critical role in the RPKI trust model since they manage the trust anchors of the RPKI hierarchic design. Additionally, due to some reasons (e.g. economics, skills) the scenario where the Certification Authority (CA) certificate is hosted by a RIR will be the most popular for a long time, in which case, RIR's online software tools to manage RPKI objects are imperative.

The RIR-hosted RPKI CA model was used for this project. Local RPKI validation servers (validation and cache) were locally deployed. This means that all resource holders had to create and manage their RPKI signed objects using the online tools implemented by LACNIC and that the local validation servers retrieve these objects from the RIR's public global repositories. No local RPKI CA nor repository were configured.

LACNIC also runs a RPKI testbed (test CA with correspondent GUI and Trust Anchor material). This infrastructure was used during the training activity.

1.4. The technical support

RPKI and origin validation are in the early stage of deployment. Few people have full knowledge about its RFCs, the implementation support in different routers and the maintenance of RPKI signed objects. To involve trained people and train new ones is very important.

People from an equipment manufacturer (Cisco) contributed with support in the startup stage and to train the holders' staff. LACNIC's staff contributed developing new online RPKI tools and training about how to use them.

2. Objective

Considering all the definitions given during the introduction and after several discussions through face and online meetings among the involved parties, the following objective was agreed on:

"Deploy RPKI-based BGP origin validation in NAP.EC's route servers. For the success of the project, 80% of the Ecuadorian prefixes (both IPv4 and IPv6) received by those routers should have a valid origin."

In order to monitor the progress, NAP.EC - GYE was taken as reference because NAP.EC - UIO had non-Ecuadorian prefixes announced.

3. Planning

The project started with an initial idea from a very reduced number of enthusiasts that identified a suitable network (the island of trust), involved the appropriate organizations and set milestones in order to carry forward a full implementation of the technology. Into the process, all parties identified the gaps and proposed solutions to overcome them.

One point that it was wanted to guarantee is that we would be able to create the appropriate "buzz" around the project. So, a communication strategy should not be overlooked. In this case, LACNIC and AEPROVI signed a MoU in April 2013 and all parties (LACNIC, AEPROVI and Cisco) announced the project and issued a press release at the LACNIC event in May 2013.

Some points that required specific discussion by the core team included:

1. RPKI-based origin validation support in the route-servers equipments
2. How to deploy a RPKI cache into the Network
3. How to populate the RPKI database with the correct and necessary information
4. Action to take with NotFound and Invalid prefixes

3.1. RPKI-based origin validation support

NAP.EC uses Cisco equipment. The project started with the initial idea of to implement origin validation into existing routers used as route servers, simply after a software update or upgrade. However, the vendor had no plans to support it in the existing platform. AEPROVI had future plans to carry forward a routers renewal, then this issue was overcome but it stopped the project for some time. Describing the equipment renewal process is beyond the scope of this document.

For Cisco equipment, the vendor has made available some online software tools to check the support. About origin validation, the routers must support: RTR protocol [RFC6810] and RPKI-based origin validation [RFC6811]. Moreover, among other things, four octects ASN support [RFC6793] and IPv6 routing support ([RFC2545] and some others) are mandatory in NAP.EC.

The selected routers were two Cisco ASR-1000 series routers (one for Quito and other for Guayaquil).

3.2. Deploying a RPKI cache into the network

Based on available resources and existing skills, it was decided to use Virtual Machines (VM) as RPKI caches, which would run GNU Linux.

The validating software is in the early stages of its development and there could be bugs or reliability problems, so it was decided using two different packages (processes) in each VM.

To ensure high availability, it was decided to deploy two VMs, each one in different host server.

There are no servers in Guayaquil, therefore both VMs would be in Quito within the NAP.EC management network and connect via the RTR protocol to the route-servers located in Quito and Guayaquil.

Additionally, the firewall rules allow RTR connections from the NAP.EC LANs to the RPKI validator servers in order to facilitate participants to perform origin validation in their edge equipments (if they wish to in the future).

3.3. Populating the RPKI database

The IP resource holders must create all needed RPKI data for the project, at least certificates and Route Origin Authorizations (ROAs). Moreover, the technical staff needed training about RPKI and origin validation because it is a new technology. Accordingly, a reasonable method to achieve it should be contrived.

It was decided to organize an event with two objectives: training and RPKI object signing. One key planning activity was to create the list of participants and to make sure that at least one participant per network had the authentication credentials to the LACNIC system to create its RPKI signed objects.

The target community was limited to Ecuadorian organizations that had received IP resources from LACNIC until mid-2013. That meant around fifty (50) organizations including Internet Services Providers, universities, banks, etc., or expressing it in prefixes: around 8600 IPv4 and 60 IPv6 blocks.

Some weeks before the deployment, there was informal dissemination meetings between the NAP.EC administrator and the participants. The project milestones were reported and the attendees received information about RPKI and origin validation for the first time. A

complete training was offered as a project milestone in the next few weeks.

All organizations were contacted and received an invitation to the event. More information about this can be found in Section 5.

3.4. Action to take with NotFound and Invalid prefixes

Despite the efforts, the RPKI database information may be incomplete, therefore the routing tables often will have NotFound prefixes. Moreover, it is needed some time after the first contact with RPKI-based origin validation technology to fix possible errors (e.g. invalid prefixes) and to assess the impact. A strict policy of dropping prefixes did not seem convenient as a starting point for the project.

It was decided that NAP.EC proceeds as follows:


- At the beginning, the NAP.EC's routers only would monitor the RPKI origin state of prefixes without action.
- In the near future, NAP.EC administration might change the action based on results of the signing-party event and community consensus.

As part of the second stage, some days after the signing event, each prefix is being marked with a BGP community to identify its RPKI origin state, sending that information to the participants.

Finally, some months later, a date was set to begin applying a strict policy. The policy was defined as follows: dropping Invalid prefixes and setting a lower local preference for NotFound prefixes.

4. Deployment

Following is the NAP.EC topology during the deployment:



4.1. RPKI Validation servers

A virtual machine (named VM1) on VMware ESXi was deployed, running GNU Linux, Centos distribution. The other one (named VM2) was cloned from this one.

Each virtual machine has access to the Internet through 1 (one) ethernet interface with a public IPv4 address within the same subnet like this: 192.0.2.2/27 for VM1, 192.0.2.3/27 for VM2 and 192.0.2.1/27 for network gateway.

The 192.0.2.0/27 network is within AS 27814. AS 27814 contains NAP.EC monitoring equipment (among other things) and it is connected to NAP.EC - Quito and NAP.EC - Guayaquil.

Each VM has the following packages (installation and configuration guides of these packages are beyond the scope of this document):

- ntpd (as NTP client)
- iptables (as firewall)
- apache (as web server)
- validating software from RIPE (<http://www.ripe.net/lir-services/resource-management/certification/tools-and-resources>)
- validating software from the rpki.net project (<http://rpki.net/wiki/doc/RPKI/Installation>)

Each validating software was setup on a different port:

- validating software from RIPE, port 65001,

- validating software from the rpki.net project, port 65002.

Each validating software has a monitoring web page, each one was configured on different port:

- RIPE software, port 65081,
- rpki.net software, port 65082.

4.2. Origin validation setting

Since the routers renewal, NAP.EC - Quito and NAP.EC - Guayaquil have each one a Cisco ASR-1000 series router as route server. These routers have IOS-XE version 3 which runs a process with IOS version 15.

First, it is required to configure communication via RTR protocol between the routers and the RPKI validation servers (caches). In this step, the necessary data are: IP addresses and service ports of the caches and the time which the router will re-query the cache (refresh time). Currently, RPKI information does not change quickly, therefore 600 seconds (10 minutes) may be considered enough for refresh time.

Cisco IOS 15 drops Invalid prefixes by default, but there is a command to avoid this behavior (ibgp bestpath prefix-validate allow-invalid). This must be applied while the policy is no action.

Later, a route-map is required to configure any action as applying different BGP local preference or marking each prefix with a BGP community based on its RPKI origin state (also send-community option must be enabled within the BGP session configuration):

The following community assignment policy was applied:

- <IXP-ASN>:21 --> Valid origin
- <IXP-ASN>:22 --> NotFound origin
- <IXP-ASN>:23 --> Invalid origin

Where <IXP-ASN> equals:

- 52482 for NAP.EC - Quito, or
- 52483 for NAP.EC - Guayaquil.

The template used in NAP.EC is in Appendix A.

5. Training and RPKI signing event

The event was called "Seminario sobre seguridad en el encaminamiento de Internet: BGP RPKI - Validacion de origen" and was scheduled for September 4-5, 2013. The agenda included theoretical and practical training, plus two time slots to sign RPKI objects: one at the end of the first day and other one during the second day.

Lack of training materials was a issue to overcome during preparatory work of the event. Some necessary activities were:

- The instructors (four people) prepared materials to cover topics such as BGP, RPKI, origin validation and the new NAP.EC platform.
- LACNIC's staff developed two new on-line tools: RPKI ROA wizard (<http://tools.labs.lacnic.net/roa-wizard/>) and RPKI announcement (<http://tools.labs.lacnic.net/announcement/>), further improved the demo environment of the RPKI system (<http://rpkidemo.labs.lacnic.net/>).
- Cisco's staff implemented a temporary virtualized network with many routers supporting RPKI and origin validation.

The event took place in a hotel and had Internet to access the training tools and the real LACNIC's hosted RPKI system.

Not all organizations sent a representative. The attendance represented around 80% of the target prefixes.

6. Outcome and post-event activities

Before the event, less 1% of the Ecuadorian prefixes were signed. At the start of the second day, less than 20% of the Ecuadorian prefixes were covered by a ROA. At the end of the event, around 80% of the Ecuadorian prefixes had a RPKI origin state as Valid.

MRTG graphs were implemented to monitor the amount of Valid, NotFound and Invalid prefixes after the event.

Feedback was received from attendees before closing the event. Some people recommended applying an acceptable policy in order do not waste the successful effort.

A few days after the event, some non-attending organizations were contacted by the NAP.EC administrator and meetings were coordinated for ROA creation. After these activities, almost 100% of Ecuadorian prefixes are covered for a ROA.

Communication activities performed after the event included:

- This document and presentation at relevant IETF Working Groups.
- Presentation at IEPG, LACNIC and other NOG events
- Publication at tech sites
- Note at local regulator newsletter
- Document and presentation at CITEC (Organization of American States)
- Blogging and social media in relevant platforms

As subsequent operational tasks, an update of validating software was performed. Overall, management has been simple and without major problems.

7. Lessons learned and best practices

- Implementation support needs to be verified in all target platforms.
- The IP resource holders community need RPKI-based origin validation training. Operators are less conservative than original though by organizers and once RPKI local space was full, support for removing invalid was unanimous.
- One day for a RPKI signing party is insufficient. the participants may not be confident about their skills or may need further authorization. Two days is a better practice (people need to sleep over what they learned the first day).
- From now on, when a new ISP wants to join NAP.EC, it receives information about RPKI-based origin validation and it is invited to create its ROAs.
- The event was a great opportunity to assemble the local community, particularly resource holders that had no previous participation at the local IXP.
- Initial work to have the "right people" in the room is a key to success. Particularly, operators need to have access to their RIR account.
- Post event communication needs to be discussed ahead of time.

8. IANA Considerations

No IANA requirements

9. Security Considerations

This document describes the experience of implementing a BGP origin validation island of trust in Ecuador. The actions taken are explicitly to be able to validate the origin in a BGP advertisement. There were no security-related issues identified during the deployment.

10. Acknowledgements

The authors wish to thank:

- all attendees at the training and RPKI signing event, without them this would not have happened.
- AEPROVI, LACNIC and Cisco for supporting the project.
- Arturo Servin for supporting the project from the start.
- Francisco Balarezo, Andres Piazza, Nicolas Fiumarelli and Chip Sharp as well as ISOC and Andean-Trade.

11. Informative References

- [RFC2545] Marques, P. and F. Dupont, "Use of BGP-4 Multiprotocol Extensions for IPv6 Inter-Domain Routing", RFC 2545, March 1999.
- [RFC6793] Vohra, Q. and E. Chen, "BGP Support for Four-Octet Autonomous System (AS) Number Space", RFC 6793, December 2012.
- [RFC6810] Bush, R. and R. Austein, "The Resource Public Key Infrastructure (RPKI) to Router Protocol", RFC 6810, January 2013.
- [RFC6811] Mohapatra, P., Scudder, J., Ward, D., Bush, R., and R. Austein, "BGP Prefix Origin Validation", RFC 6811, January 2013.

Appendix A. Router configuration templates

TEMPLATE 1

Policy: only marking prefixes based on RPKI origin state.

```
router bgp <IXP-ASN>

  bgp rpki server tcp 192.0.2.2 port 65001 refresh 600
  bgp rpki server tcp 192.0.2.2 port 65002 refresh 600
  bgp rpki server tcp 192.0.2.3 port 65001 refresh 600
  bgp rpki server tcp 192.0.2.3 port 65002 refresh 600
  !
  neighbor <neighbor-IPv4> remote-as <neighbor-IPv4-ASN>
  neighbor <neighbor-IPv4> version 4
  !
  neighbor <neighbor-IPv6> remote-as <neighbor-IPv6-ASN>
  neighbor <neighbor-IPv6> version 4
  !
  address-family ipv4
    bgp bestpath prefix-validate allow-invalid
    neighbor <neighbor-IPv4> send-community
    neighbor <neighbor-IPv4> route-map <route-map-name> out
  exit-address-family
  !
  address-family ipv6
    bgp bestpath prefix-validate allow-invalid
    neighbor <neighbor-IPv6> send-community
```

```
        neighbor <neighbor-IPv6> route-map <route-map-name> out
    exit-address-family
!
!
ip bgp-community new-format
!
!
route-map <route-map-name> permit 10
    match rpki valid
    set community <IXP-ASN>:21 no-export
!
route-map <route-map-name> permit 20
    match rpki not-found
    set community <IXP-ASN>:22 no-export
!
route-map <route-map-name> permit 30
    match rpki invalid
    set community <IXP-ASN>:23 no-export
!
```

TEMPLATE 2

Policy: Dropping Invalid prefixes and setting lower local preference for NotFound prefixes.

```
router bgp <IXP-ASN>
    bgp rpki server tcp 192.0.2.2 port 65001 refresh 600
    bgp rpki server tcp 192.0.2.2 port 65002 refresh 600
```

```
bgp rpki server tcp 192.0.2.3 port 65001 refresh 600
bgp rpki server tcp 192.0.2.3 port 65002 refresh 600
!
neighbor <neighbor-IPv4> remote-as <neighbor-IPv4-ASN>
neighbor <neighbor-IPv4> version 4
!
neighbor <neighbor-IPv6> remote-as <neighbor-IPv6-ASN>
neighbor <neighbor-IPv6> version 4
!
address-family ipv4
    neighbor <neighbor-IPv4> send-community
    neighbor <neighbor-IPv4> route-map <route-map-name> out
exit-address-family
!
address-family ipv6
    neighbor <neighbor-IPv6> send-community
    neighbor <neighbor-IPv6> route-map <route-map-name> out
exit-address-family
!
!
ip bgp-community new-format
!
!
route-map <route-map-name> permit 10
```

```
match rpki valid

set community <IXP-ASN>:21 no-export

!

route-map <route-map-name> permit 20

match rpki not-found

set local-preference 50

set community <IXP-ASN>:22 no-export

!

!
```

Authors' Addresses

Fabian Mejia (editor)
AEPROVI
Av. Republica de El Salvador N34-211
Quito
EC

Email: fabian@aeprovi.org.ec

Roque Gagliano
Cisco Systems
Avenue des Uttins 5
Rolle 1180
Switzerland

Email: rogaglia@cisco.com

Alvaro Retana
Cisco Systems
7025 Kit Creek Rd.
Research Triangle Park, NC 27617
US

Email: aretana@cisco.com

Carlos Martinez
LACNIC

Email: carlos@lacnic.net

Gerardo Rada
LACNIC

Network Working Group
Internet-Draft
Intended status: Informational
Expires: September 4, 2015

F. Mejia, Ed.
AEPROVI
R. Gagliano
A. Retana
Cisco Systems
C. Martinez
G. Rada
LACNIC
March 3, 2015

Implementing RPKI-based origin validation one country at a time. The
Ecuadorian case study.

draft-fmejia-opsec-origin-a-country-02

Abstract

One possible deployment strategy for BGP origin validation based on the Resource Public Key Infrastructure (RPKI) is the construction of islands of trust. This document describes the authors' experience deploying and maintaining a BGP origin validation island of trust in Ecuador.

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 September 4, 2015.

Copyright Notice

Copyright (c) 2015 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	2
1.1. Policer Network	3
1.2. The resource holders	4
1.3. RPKI certificate authorities and repository	5
1.4. The technical support	5
2. Objective	5
3. Planning	6
3.1. RPKI-based origin validation support	6
3.2. Deploying a RPKI cache into the network	7
3.3. Populating the RPKI database	7
3.4. Action to take with NotFound and Invalid prefixes	8
4. Deployment	8
4.1. RPKI Validation servers	9
4.2. Origin validation setting	10
5. Training and RPKI signing event	11
6. Outcome and post-event activities	11
7. Lessons learned and best practices	12
8. IANA Considerations	13
9. Security Considerations	13
10. Acknowledgements	13
11. Informative References	13
Appendix A. Router configuration templates for Cisco IOS	14
Authors' Addresses	17

1. Introduction

BGP origin validation based on RPKI [RFC6811] is in early stages of deployment. As with other new technologies, there are impediments to its global adoption as its full value is not yet perceived. Particularly, RPKI based origin validation involves on one side the creation of a large enough set of signed objects and on the other side the application of network policies based on these signed objects by network operators. An operator that does not see a large enough set of signed objects in the RPKI repository system is not encouraged to implement these set of policies. Conversely, IP address space resource holders that are not required by network operators (i.e. transit providers, peers or operators community in general) to create and maintain their RPKI objects have little incentive to do so.

To overcome this bootstrap problem, it is necessary to create a success story that brings enough value to both: network operators and resource holders. Moreover, one possible strategy for the adoption of a security technology is the creation of islands of trust where the technology is fully deployed in a reduced environment. In this direction, some organizations carried forward a full implementation of an island of trust in Ecuador. This was a multi stakeholder project where each party (resource holders, an Internet Exchange Point manager, a Regional Internet Registry and an equipment manufacturer) contributed to its success.

This document describes the experience of implementing RPKI-based origin validation in Ecuador and it is expected to be an useful guide to start other similar projects.

Below, it is described the different roles in the project and the involved parties.

1.1. Policer Network

In this document, the "Policer Network" is the networking infrastructure where the origin validation based on RPKI will be deployed to apply polices on BGP announcements. NAP.EC (www.nap.ec) was selected for this role.

NAP.EC is the Internet Exchange Point (IXP) in Ecuador with two Points of Presence (POPs): Quito (UIO) and Guayaquil (GYE). It has a BGP route-server in each location with a mandatory multilateral routing policy (i.e. all participants have a BGP session to the route-server). Each location uses a different IP address block and Autonomous System Number (ASN). NAP.EC is a meeting point where many organizations (Internet Service Providers, content providers, root servers, etc.) exchange routing information.

The participants connected to NAP.EC announce almost 100% of the total address space used in Ecuador (be believe it is 100% but we cannot be certain though). In some cases they announce their own address space and in some cases they are transit providers for their customers' resources.

AEPROVI (www.aeprovi.org.ec) manages the NAP.EC infrastructure. It is a non-profit organization, based on membership and brings together around 30 Ecuadorian ICT-related companies. AEPROVI also has an excellent reputation as an innovator in the local networking community thanks to projects such as IPv6 adoption and CDN cache servers hosting. These projects have given the local community concrete value and have build the trust on the team that manages the local IXP.

Thanks to this trust, and to the fact that all local BGP announcements are performed through the route servers, NAP.EC is uniquely positioned to become the "policer network" for this project. At the same time, it can be said that implementing origin validation at the NAP.EC route servers is equivalent to implementing it for all inter-domain routing in the country.

1.2. The resource holders

In this document, an organization which operates their own IP prefixes is called resource holder or simply the holder. They may have resources allocated/assigned from a Regional Internet Registry (RIR) and/or legacy resources (if the allocation was done before RIR formation). Resource holders are responsible for creating the RPKI signed objects for this project.

NAP.EC routing tables involve a number of holders, including organizations like Internet Service Providers, content providers, universities, .ec domain administrator and root servers. Most of them are Ecuadorian companies and have received IP resources only from LACNIC, but some have both RIR and legacy resources. Moreover, a few holders are foreign companies and their resources are legacy or from other RIRs (e.g. root servers and content providers).

Not all resource holders are directly connected to the NAP.EC fabric; some have IP address resources but not an ASN and some others are small networks that receive traffic from other bigger networks. In this case their IP address prefixes are announced by their transit providers. One of the main challenges for this project was to identify all the resource holders that needed to be contacted and to encourage network administrators from these organizations to participate.

In addition, some resource holders are part of a larger (and sometimes international) organization, with strong change management processes. This means that any change on their configurations needed to be planned ahead of time and consulted outside of the country.

In NAP.EC - UIO, the routing table includes prefixes used in Ecuador and other countries.

In NAP.EC - GYE, the routing table includes prefixes from companies operating only in Ecuador.

For the project, the target was limited to prefixes used in Ecuador by Ecuadorian holders that had received resources from LACNIC until mid-2013.

1.3. RPKI certificate authorities and repository

The five Regional Internet Registries (RIRs) have a critical role in the RPKI trust model since they manage the trust anchors of the RPKI hierarchic design. Additionally, due to some reasons (e.g. economics, skills) the scenario where the Certification Authority (CA) certificate is hosted by a RIR will be the most popular for a long time, in which case, RIR's online software tools to manage RPKI objects are imperative.

The RIR-hosted RPKI CA model was used for this project. Local RPKI validation servers (validation and cache) were locally deployed. This means that all resource holders had to create and manage their RPKI signed objects using the online tools implemented by LACNIC and that the local validation servers retrieve these objects from the RIR's public global repositories. No local RPKI CA nor repository were configured.

LACNIC also runs a RPKI testbed (test CA with correspondent GUI and Trust Anchor material). This infrastructure was used during the training activity.

1.4. The technical support

RPKI and origin validation are in the early stage of deployment. Few people have full knowledge about its RFCs, the implementation support in different routers and the maintenance of RPKI signed objects. To involve trained people and train new ones is very important.

People from an equipment manufacturer (Cisco) contributed with support in the startup stage and to train the holders' staff. LACNIC's staff contributed developing new online RPKI tools and training about how to use them.

2. Objective

Considering all the definitions given during the introduction and after several discussions through face and online meetings among the involved parties, the following objective was agreed on:

"Deploy RPKI-based BGP origin validation in NAP.EC's route servers. For the success of the project, 80% of the Ecuadorian prefixes (both IPv4 and IPv6) received by those routers should have a valid origin."

In order to monitor the progress, NAP.EC - GYE was taken as reference because NAP.EC - UIO had non-Ecuadorian prefixes announced.

3. Planning

The project started with an initial idea from a very reduced number of enthusiasts that identified a suitable network (the island of trust), involved the appropriate organizations and set milestones in order to carry forward a full implementation of the technology. Into the process, all parties identified the gaps and proposed solutions to overcome them.

One point that it was wanted to guarantee is that we would be able to create the appropriate "buzz" around the project. So, a communication strategy should not be overlooked. In this case, LACNIC and AEPROVI signed a MoU in April 2013 and all parties (LACNIC, AEPROVI and Cisco) announced the project and issued a press release at the LACNIC event in May 2013.

Some points that required specific discussion by the core team included:

1. RPKI-based origin validation support in the route-servers equipments
2. How to deploy a RPKI cache into the Network
3. How to populate the RPKI database with the correct and necessary information
4. Action to take with NotFound and Invalid prefixes

3.1. RPKI-based origin validation support

NAP.EC uses Cisco equipment. The project started with the initial idea of to implement origin validation into existing routers used as route servers, simply after a software update or upgrade. However, the vendor had no plans to support it in the existing platform. AEPROVI had future plans to carry forward a routers renewal, then this issue was overcome but it stopped the project for some time. Describing the equipment renewal process is beyond the scope of this document.

For Cisco equipment, the vendor has made available some online software tools to check the support. About origin validation, the routers must support: RTR protocol [RFC6810] and RPKI-based origin validation [RFC6811]. Moreover, among other things, four octects ASN support [RFC6793] and IPv6 routing support ([RFC2545] and some others) are mandatory in NAP.EC.

The selected routers were two Cisco ASR-1000 series routers (one for Quito and other for Guayaquil).

3.2. Deploying a RPKI cache into the network

Based on available resources and existing skills, it was decided to use Virtual Machines (VM) as RPKI caches, which would run GNU Linux.

The validating software is in the early stages of its development and there could be bugs or reliability problems, so it was decided using two different packages (processes) in each VM.

To ensure high availability, it was decided to deploy two VMs, each one in different host server.

There are no servers in Guayaquil, therefore both VMs would be in Quito within the NAP.EC management network and connect via the RTR protocol to the route-servers located in Quito and Guayaquil.

Additionally, the firewall rules allow RTR connections from the NAP.EC LANs to the RPKI validator servers in order to facilitate participants to perform origin validation in their edge equipments (if they wish to in the future).

3.3. Populating the RPKI database

The IP resource holders must create all needed RPKI data for the project, at least certificates and Route Origin Authorizations (ROAs). Moreover, the technical staff needed training about RPKI and origin validation because it is a new technology. Accordingly, a reasonable method to achieve it should be contrived.

It was decided to organize an event with two objectives: training and RPKI object signing. One key planning activity was to create the list of participants and to make sure that at least one participant per network had the authentication credentials to the LACNIC system to create its RPKI signed objects.

The target community was limited to Ecuadorian organizations that had received IP resources from LACNIC until mid-2013. That meant around fifty (50) organizations including Internet Services Providers, universities, banks, etc., or expressing it in prefixes: around 8600 IPv4 and 60 IPv6 blocks.

Some weeks before the deployment, there was informal dissemination meetings between the NAP.EC administrator and the participants. The project milestones were reported and the attendees received information about RPKI and origin validation for the first time. A

complete training was offered as a project milestone in the next few weeks.

All organizations were contacted and received an invitation to the event. More information about this can be found in Section 5.

3.4. Action to take with NotFound and Invalid prefixes

Despite the efforts, the RPKI database information may be incomplete, therefore the routing tables often will have NotFound prefixes. Moreover, it is needed some time after the first contact with RPKI-based origin validation technology to fix possible errors (e.g. invalid prefixes) and to assess the impact. A strict policy of dropping prefixes did not seem convenient as a starting point for the project.

It was decided that NAP.EC proceeds as follows:

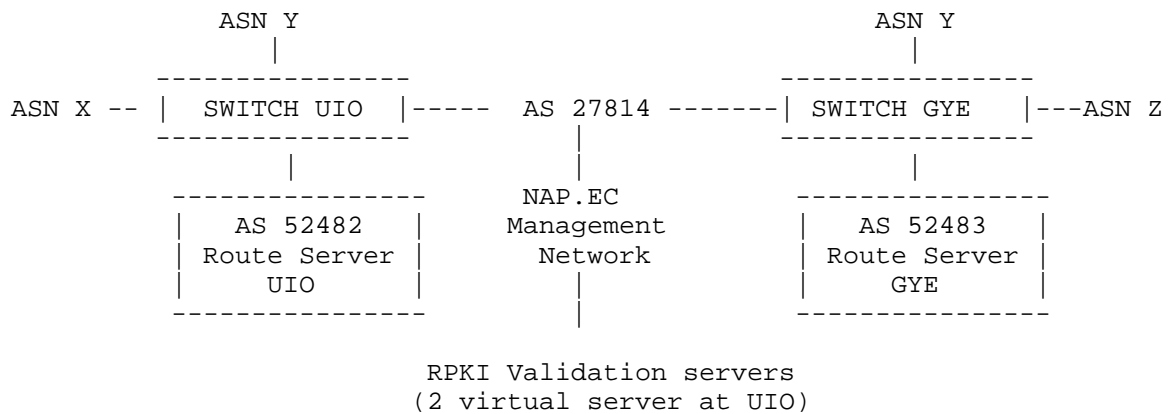
- At the beginning, the NAP.EC's routers only would monitor the RPKI origin state of prefixes without action.
- In the near future, NAP.EC administration might change the action based on results of the signing-party event and community consensus.

As part of the second stage, some days after the signing event, each prefix is being marked with a BGP community to identify its RPKI origin state, sending that information to the participants.

Finally, some months later, a date was set to begin applying a strict policy. The policy was defined as follows: dropping Invalid prefixes and setting a lower local preference for NotFound prefixes.

4. Deployment

Following is the NAP.EC topology during the deployment:



4.1. RPKI Validation servers

A virtual machine (named VM1) on VMware ESXi was deployed, running GNU Linux, Centos distribution. The other one (named VM2) was cloned from this one.

Each virtual machine has access to the Internet through 1 (one) ethernet interface with a public IPv4 address within the same subnet like this: 192.0.2.2/27 for VM1, 192.0.2.3/27 for VM2 and 192.0.2.1/27 for network gateway.

The 192.0.2.0/27 network is within AS 27814. AS 27814 contains NAP.EC monitoring equipment (among other things) and it is connected to NAP.EC - Quito and NAP.EC - Guayaquil.

Each VM has the following packages (installation and configuration guides of these packages are beyond the scope of this document):

- ntpd (as NTP client)
- iptables (as firewall)
- apache (as web server)
- validating software from RIPE (<http://www.ripe.net/lir-services/resource-management/certification/tools-and-resources>)
- validating software from the rpki.net project (<http://rpki.net/wiki/doc/RPKI/Installation>)

Each validating software was setup on a different port:

- validating software from RIPE, port 65001,

- validating software from the rpki.net project, port 65002.

Each validating software has a monitoring web page, each one was configured on different port:

- RIPE software, port 65081,
- rpki.net software, port 65082.

4.2. Origin validation setting

Since the routers renewal, NAP.EC - Quito and NAP.EC - Guayaquil have each one a Cisco ASR-1000 series router as route server. These routers have IOS-XE version 3 which runs a process with IOS version 15.

First, it is required to configure communication via RTR protocol between the routers and the RPKI validation servers (caches). In this step, the necessary data are: IP addresses and service ports of the caches and the time which the router will re-query the cache (refresh time). Currently, RPKI information does not change quickly, therefore 600 seconds (10 minutes) may be considered enough for refresh time.

Cisco IOS 15 drops Invalid prefixes by default, but there is a command to avoid this behavior (ibgp bestpath prefix-validate allow-invalid). This must be applied while the policy is no action.

Later, a route-map is required to configure any action as applying different BGP local preference or marking each prefix with a BGP community based on its RPKI origin state (also send-community option must be enabled within the BGP session configuration):

The following community assignment policy was applied:

- <IXP-ASN>:21 --> Valid origin
- <IXP-ASN>:22 --> NotFound origin
- <IXP-ASN>:23 --> Invalid origin

Where <IXP-ASN> equals:

- 52482 for NAP.EC - Quito, or
- 52483 for NAP.EC - Guayaquil.

The template used in NAP.EC is in Appendix A.

5. Training and RPKI signing event

The event was called "Seminario sobre seguridad en el encaminamiento de Internet: BGP RPKI - Validacion de origen" and was scheduled for September 4-5, 2013. The agenda included theoretical and practical training, plus two time slots to sign RPKI objects: one at the end of the first day and other one during the second day.

Lack of training materials was a issue to overcome during preparatory work of the event. Some necessary activities were:

- The instructors (four people) prepared materials to cover topics such as BGP, RPKI, origin validation and the new NAP.EC platform.
- LACNIC's staff developed two new on-line tools: RPKI ROA wizard (<http://tools.labs.lacnic.net/roa-wizard/>) and RPKI announcement (<http://tools.labs.lacnic.net/announcement/>), further improved the demo environment of the RPKI system (<http://rpkidemo.labs.lacnic.net/>).
- Cisco's staff implemented a temporary virtualized network with many routers supporting RPKI and origin validation.

The event took place in a hotel and had Internet to access the training tools and the real LACNIC's hosted RPKI system.

Not all organizations sent a representative. The attendance represented around 80% of the target prefixes.

6. Outcome and post-event activities

Before the event, less 1% of the Ecuadorian prefixes were signed. At the start of the second day, less than 20% of the Ecuadorian prefixes were covered by a ROA. At the end of the event, around 80% of the Ecuadorian prefixes had a RPKI origin state as Valid.

MRTG graphs were implemented to monitor the amount of Valid, NotFound and Invalid prefixes after the event.

Feedback was received from attendees before closing the event. Some people recommended applying an acceptable policy in order do not waste the successful effort.

A few days after the event, some non-attending organizations were contacted by the NAP.EC administrator and meetings were coordinated for ROA creation. After these activities, almost 100% of Ecuadorian prefixes are covered for a ROA.

Communication activities performed after the event included:

- This document and presentation at relevant IETF Working Groups.
- Presentation at IEPG, LACNIC and other NOG events.
- Publication at tech sites.
- Communication to local telecommunications regulator.
- Document and presentation at CITEC (Organization of American States).
- Blogging and social media in relevant platforms.

As subsequent operational tasks, the following are necessary:

- Updating of validating software.
- Permanent monitoring the origin state of prefixes.
- Alerting about Invalid prefixes.

7. Lessons learned and best practices

- Implementation support needs to be verified in all target platforms.
- The IP resource holders community need RPKI-based origin validation training.
- Initial work to have the "right people" in the room is a key to success for the RPKI signing party. Particularly, operators need to have access to their RIR account.
- One day for a RPKI signing party is insufficient, two days is a better practice. The participants may not be confident about their skills or may need further authorization (people need to sleep over what they learned the first day).
- The event was a great opportunity to assemble the local community, particularly resource holders that had no previous participation at the local IXP.
- Post event communication needs to be discussed ahead of time.

- Operators are less conservative than original though by organizers and once RPKI local space was full, support for removing invalid prefixes was unanimous.
- From now on, when a new ISP wants to join NAP.EC, it receives information about RPKI-based origin validation and it is invited to create its ROAs.

8. IANA Considerations

No IANA requirements

9. Security Considerations

This document describes the experience of implementing a BGP origin validation island of trust in Ecuador. The actions taken are explicitly to be able to validate the origin in a BGP advertisement. There were no security-related issues identified during the deployment.

10. Acknowledgements

The authors wish to thank:

- all attendees at the training and RPKI signing event, without them this would not have happened.
- AEPROVI, LACNIC and Cisco for supporting the project.
- Arturo Servin for supporting the project from the start.
- Francisco Balarezo, Andres Piazza, Nicolas Fiumarelli and Chip Sharp as well as ISOC and Andean-Trade.

11. Informative References

- [RFC2545] Marques, P. and F. Dupont, "Use of BGP-4 Multiprotocol Extensions for IPv6 Inter-Domain Routing", RFC 2545, March 1999.
- [RFC6793] Vohra, Q. and E. Chen, "BGP Support for Four-Octet Autonomous System (AS) Number Space", RFC 6793, December 2012.
- [RFC6810] Bush, R. and R. Austein, "The Resource Public Key Infrastructure (RPKI) to Router Protocol", RFC 6810, January 2013.

[RFC6811] Mohapatra, P., Scudder, J., Ward, D., Bush, R., and R. Austein, "BGP Prefix Origin Validation", RFC 6811, January 2013.

Appendix A. Router configuration templates for Cisco IOS

TEMPLATE 1

Policy: Only marking prefixes based on RPKI origin state.

```
router bgp <IXP-ASN>

  bgp rpki server tcp 192.0.2.2 port 65001 refresh 600
  bgp rpki server tcp 192.0.2.2 port 65002 refresh 600
  bgp rpki server tcp 192.0.2.3 port 65001 refresh 600
  bgp rpki server tcp 192.0.2.3 port 65002 refresh 600
  !
  neighbor <neighbor-IPv4> remote-as <neighbor-IPv4-ASN>
  neighbor <neighbor-IPv4> version 4
  !
  neighbor <neighbor-IPv6> remote-as <neighbor-IPv6-ASN>
  neighbor <neighbor-IPv6> version 4
  !
  address-family ipv4
    bgp bestpath prefix-validate allow-invalid
    neighbor <neighbor-IPv4> send-community
    neighbor <neighbor-IPv4> route-map <route-map-name> in
  exit-address-family
  !
  address-family ipv6
```

```
    bgp bestpath prefix-validate allow-invalid
    neighbor <neighbor-IPv6> send-community
    neighbor <neighbor-IPv6> route-map <route-map-name> in
    exit-address-family
!
!
ip bgp-community new-format
!
!
route-map <route-map-name> permit 10
    match rpki valid
    set community <IXP-ASN>:21
!
route-map <route-map-name> permit 20
    match rpki not-found
    set community <IXP-ASN>:22
!
route-map <route-map-name> permit 30
    match rpki invalid
    set community <IXP-ASN>:23
!
```

TEMPLATE 2

Policy: Dropping Invalid prefixes and setting lower local preference for NotFound prefixes.

```
router bgp <IXP-ASN>
```

```
bgp rpki server tcp 192.0.2.2 port 65001 refresh 600
bgp rpki server tcp 192.0.2.2 port 65002 refresh 600
bgp rpki server tcp 192.0.2.3 port 65001 refresh 600
bgp rpki server tcp 192.0.2.3 port 65002 refresh 600
!
neighbor <neighbor-IPv4> remote-as <neighbor-IPv4-ASN>
neighbor <neighbor-IPv4> version 4
!
neighbor <neighbor-IPv6> remote-as <neighbor-IPv6-ASN>
neighbor <neighbor-IPv6> version 4
!
address-family ipv4
    neighbor <neighbor-IPv4> send-community
    neighbor <neighbor-IPv4> route-map <route-map-name> in
exit-address-family
!
address-family ipv6
    neighbor <neighbor-IPv6> send-community
    neighbor <neighbor-IPv6> route-map <route-map-name> in
exit-address-family
!
!
ip bgp-community new-format
!
```

```
!  
route-map <route-map-name> permit 10  
    match rpki valid  
    set community <IXP-ASN>:21  
!  
route-map <route-map-name> permit 20  
    match rpki not-found  
    set local-preference 50  
    set community <IXP-ASN>:22  
!  
!
```

Authors' Addresses

Fabian Mejia (editor)
AEPROVI
Av. Republica de El Salvador N34-211
Quito
EC

Email: fabian@aeprovi.org.ec

Roque Gagliano
Cisco Systems
Avenue des Uttins 5
Rolle 1180
Switzerland

Email: rogaglia@cisco.com

Alvaro Retana
Cisco Systems
7025 Kit Creek Rd.
Research Triangle Park, NC 27617
US

Email: aretana@cisco.com

Carlos Martinez
LACNIC

Email: carlos@lacnic.net

Gerardo Rada
LACNIC

Network Working Group
Internet-Draft
Intended status: Informational
Expires: August 13, 2014

G. Huston
G. Michaelson
APNIC
February 9, 2014

RPKI Validation Reconsidered
draft-huston-rpki-validation-01.txt

Abstract

This document reviews the certificate validation procedure specified in RFC6487 and highlights aspects of operational management of certificates in the RPKI in response to the movement of resources across registries, and the associated actions of Certification Authorities to maintain certification of resources during this movement. The document describes an alternative validation procedure that reduces the operational impact of certificate management during resource 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 August 13, 2014.

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.

Table of Contents

1. Introduction	3
2. Operational Considerations	4
3. A Specific Resource RPKI Certificate Validation Process . . .	6
3.1. Resource Transfers and Specific Resource Certificate Validation	8
3.2. A Specification of Specific Resource Validation	8
4. Local Repository Cache Maintenance	10
5. Security Considerations	11
6. IANA Considerations	11
7. Acknowledgements	11
8. References	11
8.1. Normative References	11
8.2. Informative References	12
Authors' Addresses	12

1. Introduction

This document reviews the certificate validation procedure specified in [RFC6487] and highlights aspects of operational management of certificates in the RPKI in response to the movement of resources across registries, and the associated actions of Certification Authorities to maintain certification of resources during this movement. The document describes an alternative validation procedure that reduces the operational impact of certificate management during resource movement. The alternative validation procedure also offers a higher level of robustness in the face of resource inconsistencies in a putative certificate validation path.

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 common Issuer and Subject Name pairs, 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, 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 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 holding 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.

2. 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 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 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 that they do not 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 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 oversight.

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 above example, all subordinate certificates issued by CA C are invalid until CA B issues a new certificate for CA C 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 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 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 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 certificate validation process described in [RFC6487].

3. A Specific Resource RPKI Certificate Validation Process

The question considered here is: Is there an alternate definition of RPKI certificate validity that could remove the requirement for such careful orchestration of certification actions across the RPKI to support resource transfers?

The general definition of certificate validity as defined in [RFC5280] assumes a validation question relating to the relying party's (RP's) level of trust in a subject's signed material, given knowledge of a subject's name, the subject's public key, the RP's

chosen trust anchor(s) and an overall PKI to define the domain of discourse.

The validation question assumed by the [RFC6487] RPKI certificate validation process relates to a RP's level of trust in the combination of some signed material, a certificate that attests to the public key used to sign this material and the set of all number resources that have been assigned or allocated to the subject of the certificate, given knowledge of the certificate, the RP's chosen trust anchor(s), the RPKI, and the application of the same test applied to the superior certificate in the RPKI hierarchy, and so on to a Trust Anchor.

There is a alternative certificate validation procedure that starts with an attestation containing the subject's signed material and an explicit enumeration of a set of number resources. The associated validation question relates to whether a RPKI validation process can attest to the validity of a subject's signed attestation relating to a particular set of number resources, rather than a signed attestation relating to all number resources held by this subject. We will term this alternate certificate validation process "specific resource" validation.

If the certificate validation procedure is specifically restricted to a question of ascertaining the validity of a particular set of number resources in the context of the RPKI, the RPKI validation procedure need not be as strict as a recursive "encompassing" condition for the resources contained in each pair of certificates in the validation path. It would be sufficient in the context of this "specific resource" validation procedure to require only that each certificate in the validation path has a number resource extension that "encompasses" the specific resources described in the original validation question. Rather than a validation test for all possible questions, this is a specific validation question in the context of specific resources.

This validation question can be informally described as: Given a certificate and a given resource set, is there an Issuer-Subject ordered sequence of certificates from a Trust Anchor to the certificate being validated, where each certificate on this sequence is well-formed, not revoked by a valid CRL, where the certificate's lifetimes are valid, and where the RFC3779 resource extension in the certificate encompass the given resource set?

In the example from Section 1, using a this alternate certificate validation process, a validation question of certificate 3 and the resource 10.0.1.0/24, the validation outcome would be positive, in that certificates 1, 2 and 3 all encompass the specific resource

10.0.1.0/24, assuming that the certificates are valid in all other respects.

3.1. Resource Transfers and Specific Resource Certificate Validation

When considering the transfer of resources across registries, and the associated certification actions, then if the validation process was one of "specific resource" validation, then there is no requirement for synchronized orchestration of the process of certificate issuance and revocation by the CAs involved in this transfer in order to preserve the validity of resources described in these certificates.

Along the chain of the "sending" registry CA hierarchy each registry CA can issue a certificate with a reduced resource set that removes the resource being transferred, and revoke the previously issued certificate without regard to the specific timing of similar actions by either it's superior or its subordinate registry CA.

Similarly, in the "receiving" registry hierarchy each CA can issue a certificate with an augmented resource set that includes the resource being transferred without particular regard to the timing of similar actions by the other superior or subordinate registry CAs.

Validation questions relating to the migrating resource made against certificates on the "sending registry" will return an invalid outcome as soon as any registry CA in this chain has performed revocation of the original certificate. Validation questions relating to the migrating resource made against certificates on the "receiving registry" will return an valid outcome only when all the registries in this chain have performed certificate re-issuance and included the resource in the new certificate.

Critically, at all times validation questions relating to any other resource using the "specific resource" validation approach will return the same outcomes throughout this issuance and revocation process. This "specific resource" validation process engenders a more robust outcome in RPKI certificate management. Validation questions relating to resources which are not being transferred from one registry to another cannot be compromised by any failure to adhere to a strict process of issuance and revocation relating to the certification of the resources being transferred.

3.2. A Specification of Specific Resource Validation

The following is a amended specification of certificate validation as described in [RFC6487] that describes the proposed "specific resource" certificate validation process.

Validation of signed resource data using a target resource certificate and a specific set of number resources consists of verifying that the digital signature of the signed resource data is valid, using the public key of the target resource certificate, and also validating the resource certificate in the context of the RPKI, using the path validation process. This path validation process verifies, among other things, that a prospective certification path (a sequence of n certificates) satisfies the following conditions:

1. for all 'x' in $\{1, \dots, n-1\}$, the Subject of certificate 'x' is the Issuer of certificate ('x' + 1);
2. certificate '1' is issued by a trust anchor;
3. certificate 'n' is the certificate to be validated; and
4. for all 'x' in $\{1, \dots, n\}$, certificate 'x' is valid.

Certificate validation entails verifying that all of the following conditions hold, in addition to the Certification Path Validation criteria specified in Section 6 of [RFC5280]:

1. The certificate can be verified using the Issuer's public key and the signature algorithm
2. The current time lies within the certificate's Validity From and To values.
3. The certificate contains all fields that MUST be present, as defined by this specification, and contains values for selected fields that are defined as allowable values by this specification.
4. No field, or field value, that this specification defines as MUST NOT be present is used in the certificate.

5. The Issuer has not revoked the certificate. A revoked certificate is identified by the certificate's serial number being listed on the Issuer's current CRL, as identified by the CRLDP of the certificate, the CRL is itself valid, and the public key used to verify the signature on the CRL is the same public key used to verify the certificate itself.
6. The resource extension data contained in this certificate "encompasses" the entirety of the resources in the specific resource set.
7. The Certification Path originates with a certificate issued by a trust anchor, and there exists a signing chain across the Certification Path where the Subject of Certificate 'x' in the Certification Path matches the Issuer in Certificate 'x + 1' in the Certification Path, and the public key in Certificate 'x' can verify the signature value in Certificate 'x+1'.

A certificate validation algorithm MAY perform these tests in any chosen order.

4. Local Repository Cache Maintenance

This change in the validation process would have some impact on the operation of a local cache of validated RPKI certificates. Given that the validation process requires the specification of a specific resource set, it would appear that it would not be possible to validate an RPKI certificate without also specifying a specific resource set.

However, using a top-down validation process, and an additional local data structure associated with each locally held validated RPKI certificate, it is possible to maintain a local cache of validated certificates, and the set of valid and invalid resources for each certificate.

The additional data structures are the certificate's validated and invalidated resource set. These sets are defined as follows:

- o If the certificate is used as a Trust Anchor, then the local validated resource set is copied from the certificate's RFC3779 resource set. There is no invalid resource set.
- o Otherwise, the certificate's local validated resource set is defined as the intersection of this certificate's RFC3779 resource

set and the parent certificate's local validated resource set. The certificate's invalid resource set is the difference between this set and the certificate's RFC3779 resource set.

If the certificate's validated resource set is empty then the certificate is not valid.

If the invalid resource set is not empty, then any resources that are contained in this invalid resource set cannot be valid by virtue of this certificate.

In all operations on the local repository cache, local applications should use the certificate's local validated resource set in place of the resources described in the certificate's RFC3779 extension.

The invalid resource set can be used as a diagnostic aide in local cache management.

5. Security Considerations

The Security Considerations of [RFC6487] apply to the validation procedure described here.

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.

Authors' Addresses

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

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

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

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

SIDR
Internet-Draft
Obsoletes: 6490 (if approved)
Intended status: Standards Track
Expires: August 15, 2014

G. Huston
APNIC
S. Weiler
SPARTA, Inc.
G. Michaelson
APNIC
S. Kent
BBN

February 11, 2014

Resource Certificate PKI (RPKI) Trust Anchor Locator
draft-huston-sidr-rfc6490-bis-01

Abstract

This document defines a Trust Anchor Locator (TAL) for the Resource Certificate Public Key Infrastructure (RPKI).

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 August 15, 2014.

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
1.1. Terminology	3
2. Trust Anchor Locator	3
2.1. Trust Anchor Locator Format	3
2.2. TAL and Trust Anchor Certificate Considerations	4
2.3. Example	5
3. Relying Party Use	6
4. Security Considerations	6
5. IANA Considerations	7
6. Acknowledgments	7
7. References	7
7.1. Normative References	7
7.2. Informative References	8
Authors' Addresses	8

1. Introduction

This document defines a Trust Anchor Locator (TAL) for the Resource Certificate Public Key Infrastructure (RPKI) [RFC6480]. This format may be used to distribute trust anchor material using a mix of out-of-band and online means. Procedures used by Relying Parties (RPs) to verify RPKI signed objects SHOULD support this format to facilitate interoperability between creators of trust anchor material and RPs.

1.1. Terminology

The key words "MUST", "MUST NOT", "REQUIRED", "SHALL", "SHALL NOT", "SHOULD", "SHOULD NOT", "RECOMMENDED", "MAY", and "OPTIONAL" in this document are to be interpreted as described in [RFC2119].

2. Trust Anchor Locator

2.1. Trust Anchor Locator Format

This document does not propose a new format for trust anchor material. A trust anchor in the RPKI is represented by a self-signed X.509 Certificate Authority (CA), a format commonly used in PKIs and widely supported by RP software. This document specifies a format for data used to retrieve and verify the authenticity of a trust anchor in a very simple fashion. That data is referred to as the TAL.

The motivation for defining the TAL is to enable selected data in the trust anchor to change, without needing to effect re-distribution of the trust anchor per se. In the RPKI, certificates contain extensions that represent Internet Number Resources (INRs) [RFC3779]. The set of INRs associated with an entity acting as a trust anchor is likely to change over time. Thus, if one were to use the common PKI convention of distributing a trust anchor to RPs in a secure fashion, this procedure would need to be repeated whenever the INR set for the entity acting as a trust anchor changed. By distributing the TAL (in a secure fashion), instead of the trust anchor, this problem is avoided, i.e., the TAL is constant so long as the TA's public key and its location does not change.

The TAL is analogous to the TrustAnchorInfo data structure [RFC5914] adopted as a PKIX standard. That standard could be used to represent the TAL, if one defined an rsync URI extension for that data structure. However, the TAL format was adopted by RPKI implementors prior to the PKIX trust anchor work, and the RPKI implementer community has elected to utilize the TAL format, rather than define

the requisite extension. The community also prefers the simplicity of the ASCII encoding of the TAL, vs. the binary (ASN.1) encoding for TrustAnchorInfo.

The TAL is an ordered sequence of:

- 1) a URI section, and
- 2) a subjectPublicKeyInfo [RFC5280] in DER format [X.509], encoded in Base64 (see Section 4 of [RFC4648]).

where the URI section is comprised of one of more of the ordered sequence of:

- 1.1) An rsync URI [RFC5781] ,and
- 1.2) A <CRLF> or <LF> line break.

2.2. TAL and Trust Anchor Certificate Considerations

Each rsync URI in the TAL MUST reference a single object. It MUST NOT reference a directory or any other form of collection of objects.

The referenced object MUST be a self-signed CA certificate that conforms to the RPKI certificate profile [RFC6487]. This certificate is the trust anchor in certification path discovery [RFC4158] and validation [RFC5280][RFC3779].

The validity interval of this trust anchor SHOULD reflect the anticipated period of stability the particular set of Internet Number Resources (INRs) that are associated with the putative TA.

The INR extension(s) of this trust anchor MUST contain a non-empty set of number resources. It MUST NOT use the "inherit" form of the INR extension(s). The INR set described in this certificate is the set of number resources for which the issuing entity is offering itself as a putative trust anchor in the RPKI [RFC6480].

The public key used to verify the trust anchor MUST be the same as the subjectPublicKeyInfo in the CA certificate and in the TAL.

The trust anchor MUST contain a stable key. This key MUST NOT change when the certificate is reissued due to changes in the INR extension(s), when the certificate is renewed prior to expiration or for any reason other than a key change.

Because the public key in the TAL and the trust anchor MUST be stable, this motivates operation of that CA in an off-line mode. Thus the entity that issues the trust anchor SHOULD issue a subordinate CA certificate that contains the same INRs (via the use of the "inherit" option in the INR extensions of the subordinate certificate). This allows the entity that issues the trust anchor to keep the corresponding private key of this certificate off-line, while issuing all relevant child certificates under the immediate subordinate CA. This measure also allows the CRL issued by that entity to be used to revoke the subordinate (CA) certificate in the event of suspected key compromise of this potentially more vulnerable online operational key pair.

The trust anchor MUST be published at a stable URI. When the trust anchor is re-issued for any reason, the replacement CA certificate MUST be accessible using the same URI.

Because the trust anchor is a self-signed certificate, there is no corresponding Certificate Revocation List that can be used to revoke it, nor is there a manifest [RFC6486] that lists this certificate.

If an entity wishes to withdraw a self-signed CA certificate as a putative Trust Anchor, for any reason, including key rollover, the entity MUST remove the object from the location referenced in the TAL.

Where the TAL contains two or more rsync URIs, then the same self-signed CA certificate MUST be found at each referenced location. In order to operational increase resilience, it is RECOMMENDED that the domain name parts of each of these URIs resolve to distinct IP addresses that are used by a diverse set of repository publication points, and these IP addresses be included in distinct Route Origination Authorizations (ROAs) objects signed by different CAs.

2.3. Example

```
rsync://rpki.example.org/rpki/hedgehog/root.cer
MIIBIjANBgkqhkiG9w0BAQEFAAOCAQ8AMIIBCgKCAQEAovWQL2lh6knDx
GUG5hbtCXvvh4AOzjhDkSHlj22gn/loiM9IeDATIwP44vhQ6L/xvuk7W6
Kfa5ygmqQ+xOZOwTWPCrUbqaQyPNxokuivzyvqVZVDecOEqs78q58mSp9
nbtxmLRW7B67SJCBSzfa5XpVyXYEgYAJkk3fpmefU+AcctxvvHB5OVPIa
BfPcs80ICMgHQX+fphvute9XLxjfkJKJWkhZqZ0v7pZm2uhkcPx1PMGcrG
ee0WSDC3fr3erLueagpiLsFjwwpX6F+Ms8vqz45H+DKmYKvPSstZjCCq9
aJ0qANT90tnfSDOS+aLRPjZryCNyvvBHxZXqj5YCGKtwIDAQAB
```

3. Relying Party Use

In order to use the TAL to retrieve and validate a (putative) TA, an RP SHOULD:

1. Retrieve the object referenced by (one of) the URI(s) contained in the TAL.
2. Confirm that the retrieved object is a current, self-signed RPKI CA certificate that conforms to the profile as specified in [RFC6487].
3. Confirm that the public key in the TAL matches the public key in the retrieved object.
4. Perform other checks, as deemed appropriate (locally), to ensure that the RP is willing to accept the entity publishing this self-signed CA certificate to be a trust anchor, relating to the validity of attestations made in the context of the RPKI (relating to all resources described in the INR extension of this certificate).

An RP SHOULD perform these functions for each instance of TAL that it is holding for this purpose every time the RP performs a re-synchronization across the local repository cache. In any case, an RP also SHOULD perform these functions prior to the expiration of the locally cached copy of the retrieved trust anchor referenced by the TAL.

In the case where a TAL contains multiple URIs, RP may use a locally defined preference rule to select the URI from where fetch the Trust Anchor certificate. Some examples are:

- o Using the order provided in the TAL
- o Selecting the URI randomly from the available list
- o Creating a prioritized list of URIs based on RP specific parameters, such as connection establishment delay

If the connection to the preferred URI fails, or the fetched CA certificate public key does not match the TAL public key, the RP SHOULD fetch the CA certificate from the next URI, according to the local preference ranking.

4. Security Considerations

Compromise of a trust anchor private key permits unauthorized parties to masquerade as a trust anchor, with potentially severe consequences. Reliance on an inappropriate or incorrect trust anchor

has similar potentially severe consequences.

This trust anchor locator does not directly provide a list of resources covered by the referenced self-signed CA certificate. Instead, the RP is referred to the trust anchor itself and the INR extension(s) within this certificate. This provides necessary operational flexibility, but it also allows the certificate issuer to claim to be authoritative for any resource. Relying parties should either have great confidence in the issuers of such certificates that they are configuring as trust anchors, or they should issue their own self-signed certificate as a trust anchor and, in doing so, impose constraints on the subordinate certificates.

5. IANA Considerations

[This document specifies no IANA actions.]

6. Acknowledgments

This approach to TA material was originally described by Robert Kisteleki.

The authors acknowledge the contributions of Rob Austein and Randy Bush, who assisted with earlier versions of this document and with helpful review comments.

The authors acknowledge with work of Roque Gagliano, Terry Manderson and Carloa Martinez Cagnazzo in developing the ideas behind the inclusion of multiple URIs in the TAL.

7. References

7.1. Normative References

- [RFC2119] Bradner, S., "Key words for use in RFCs to Indicate Requirement Levels", BCP 14, RFC 2119, March 1997.
- [RFC3779] Lynn, C., Kent, S., and K. Seo, "X.509 Extensions for IP Addresses and AS Identifiers", RFC 3779, June 2004.
- [RFC4648] Josefsson, S., "The Base16, Base32, and Base64 Data Encodings", RFC 4648, October 2006.
- [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.

- [RFC5781] Weiler, S., Ward, D., and R. Housley, "The rsync URI Scheme", RFC 5781, February 2010.
- [RFC6480] Lepinski, M. and S. Kent, "An Infrastructure to Support Secure Internet Routing", RFC 6480, February 2012.
- [RFC6487] Huston, G., Michaelson, G., and R. Loomans, "A Profile for X.509 PKIX Resource Certificates", RFC 6487, February 2012.
- [X.509] ITU-T, "Recommendation X.509: The Directory - Authentication Framework", 2000.

7.2. Informative References

- [RFC4158] Cooper, M., Dzambasow, Y., Hesse, P., Joseph, S., and R. Nicholas, "Internet X.509 Public Key Infrastructure: Certification Path Building", RFC 4158, September 2005.
- [RFC5914] Housley, R., Ashmore, S., and C. Wallace, "Trust Anchor Format", RFC 5914, June 2010.
- [RFC6486] Austein, R., Huston, G., Kent, S., and M. Lepinski, "Manifests for the Resource Public Key Infrastructure (RPKI)", RFC 6486, February 2012.

Authors' Addresses

Geoff Huston
APNIC

Email: gih@apnic.net
URI: <http://www.apnic.net>

Samuel Weiler
SPARTA, Inc.
7110 Samuel Morse Drive
Colombia, Maryland 21046
USA

Email: weiler@sparta.com

George Michaelson
Asia Pacific Network Information Centre

Email: ggm@apnic.net
URI: <http://www.apnic.net>

Stephen Kent
BBN Technologies
10 Moulton St.
Cambridge, MA 02138
USA

Email: kent@bbn.com

Secure Inter-Domain Routing Working Group
Internet-Draft
Updates: 6487 (if approved)
Intended Status: Standards Track
Expires: March 21, 2014

M. Reynolds
IPSw
S. Turner
IECA
S. Kent
BBN
September 17, 2013

A Profile for BGPSEC Router Certificates,
Certificate Revocation Lists, and Certification Requests
draft-ietf-sidr-bgpsec-pki-profiles-06

Abstract

This document defines a standard profile for X.509 certificates for the purposes of supporting validation of Autonomous System (AS) paths in the Border Gateway Protocol (BGP), as part of an extension to that protocol known as BGPSEC. BGP is a critical component for the proper operation of the Internet as a whole. The BGPSEC protocol is under development as a component to address the requirement to provide security for the BGP protocol. The goal of BGPSEC is to design a protocol for full AS path validation based on the use of strong cryptographic primitives. The end-entity (EE) certificates specified by this profile are issued under Resource Public Key Infrastructure (RPKI) Certification Authority (CA) certificates, containing the AS Identifier Delegation extension, to routers within the Autonomous System (AS). The certificate asserts that the router(s) holding the private key are authorized to send out secure route advertisements on behalf of the specified AS. This document also profiles the Certificate Revocation List (CRL), profiles the format of certification requests, and specifies Relying Party certificate path validation procedures. The document extends the RPKI; therefore, this documents updates the RPKI Resource Certificates Profile (RFC 6487).

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."

Copyright Notice

Copyright (c) 2013 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.

1. Introduction

This document defines a profile for X.509 end-entity (EE) certificates [RFC5280] for use in the context of certification of Autonomous System (AS) paths in the Border Gateway Protocol Security (BGPSEC) protocol. Such certificates are termed "BGPSEC Router Certificates". The holder of the private key associated with a BGPSEC Router Certificate is authorized to send secure route advertisements (BGPSEC UPDATES) on behalf of the AS named in the certificate. That is, a router holding the private key may send to its BGP peers, route advertisements that contain the specified AS number as the last item in the AS PATH attribute. A key property that BGPSEC will provide is that every AS along the AS PATH can verify that the other ASes along the path have authorized the advertisement of the given route (to the next AS along the AS PATH).

This document is a profile of [RFC6487], which is a profile of [RFC5280], and it updates [RFC6487]. It establishes requirements imposed on a Resource Certificate that is used as a BGPSEC Router Certificate, i.e., it defines constraints for certificate fields and extensions for the certificate to be valid in this context. This document also profiles the Certificate Revocation List (CRL) and certification requests. Finally, this document specifies the Relying Party (RP) certificate path validation procedures.

1.1. Terminology

It is assumed that the reader is familiar with the terms and concepts described in "A Profile for X.509 PKIX Resource Certificates" [RFC6487], "BGPSEC Protocol Specification" [ID.sidr-bgpsec-protocol], "A Border Gateway Protocol 4 (BGP-4)" [RFC4271], "BGP Security

Vulnerabilities Analysis" [RFC4272], "Considerations in Validating the Path in BGP" [RFC5123], and "Capability Advertisement with BGP-4" [RFC5492].

The key words "MUST", "MUST NOT", "REQUIRED", "SHALL", "SHALL NOT", "SHOULD", "SHOULD NOT", "RECOMMENDED", "NOT RECOMMENDED", "MAY", and "OPTIONAL" in this document are to be interpreted as described in [RFC2119].

2. Describing Resources in Certificates

Figure 1 depicts some of the entities in the RPKI and some of the products generated by RPKI entities. IANA issues a Certification Authority (CA) to a Regional Internet Registries (RIR). The RIR, in turn, issues a CA certificate to an Internet Service Providers (ISP).

The ISP in turn issues End-Entity (EE) Certificates to itself as well as CRLs. These certificates are referred to as "Resource Certificates", and are profiled in [RFC6487]. The [RFC6480] envisioned using Resource Certificates to generate Manifests [RFC6486] and Route Origin Authorizations (ROAs) [RFC6482]. ROAs and Manifests also include the Resource Certificates used to sign them.

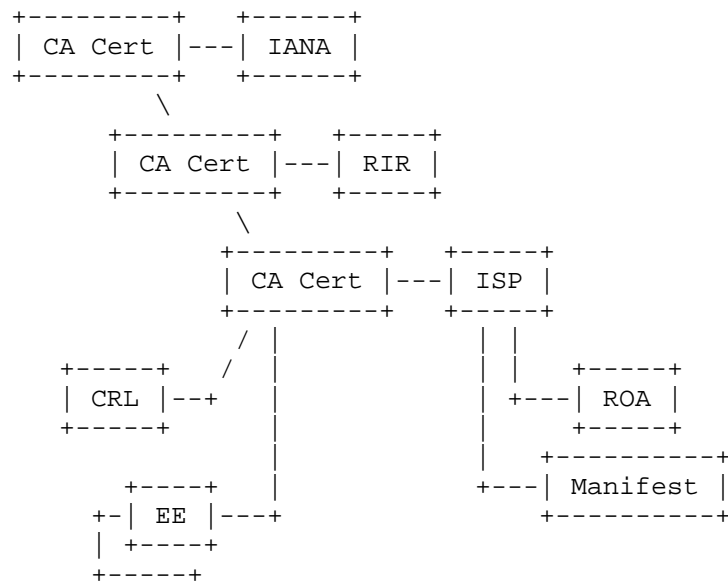


Figure 1

This document defines another type of Resource Certificate, which is referred to as a "BGPSEC Router Certificate". The purpose of this

certificate is explained in Section 1 and falls within the scope of appropriate uses defined within [RFC6484]. The issuance of BGPSEC Router Certificates has minimal impact on RPKI CAs because the RPKI CA certificate and CRL profile remain unchanged (i.e., they are as specified in [RFC6487]). Further, the algorithms used to generate RPKI CA certificates that issue the BGPSEC Router Certificates and the CRLs necessary to check the validity of the BGPSEC Router Certificates remain unchanged (i.e., they are as specified in [RFC6485]). The only impact is that the RPKI CAs will need to be able to process a profiled certificate request (see Section 5) signed with algorithms found in [ID.sidr-bgpsec-algs]. The use of BGPSEC Router Certificates in no way affects RPKI RPs that process Manifests and ROAs because the public key found in the BGPSEC Router Certificate is only ever used to verify the signature on the BGPSEC certificate request (only CAs process these) and the signature on a BGPSEC Update Message [ID.sidr-bgpsec-protocol] (only BGPSEC routers process these).

Only the differences between this profile and the profile in [RFC6487] are listed. Note that BGPSEC Router Certificates are EE certificates and as such there is no impact on process described in [ID.sidr-algorithm-agility].

3. Updates to [RFC6487]

3.1 BGPSEC Router Certificate Fields

A BGPSEC Router Certificate is a valid X.509 public key certificate, consistent with the PKIX profile [RFC5280], containing the fields listed in this section. This profile is also based on [RFC6487] and only the differences between this profile and the profile in [RFC6487] are listed.

3.1.1.1. Subject

This field identifies the router to which the certificate has been issued. Consistent with [RFC6487], only two attributes are allowed in the Subject field: common name and serial number. Moreover, the only common name encoding options that are supported are printableString and UTF8String. For BGPSEC Router Certificates, it is RECOMMENDED that the common name attribute contain the literal string "ROUTER-" followed by the 32-bit AS Number [RFC3779] encoded as eight hexadecimal digits and that the serial number attribute contain the 32-bit BGP Identifier [RFC4271] (i.e., the router ID) encoded as eight hexadecimal digits. If the same certificate is issued to more than one router (hence the private key is shared among these routers), the choice of the router ID used in this name is at the discretion of the Issuer. Note that router IDs are not

guaranteed to be unique across the Internet, and thus the Subject name in a BGPSEC Router Certificate issued using this convention also is not guaranteed to be unique across different issuers. However, each certificate issued by an individual CA MUST contain a Subject name that is unique within that context.

3.1.2. Subject Public Key Info

Refer to section 3.1 of [ID.sidr-bgpsec-algs].

3.1.3. BGPSEC Router Certificate Version 3 Extension Fields

The following X.509 V3 extensions MUST be present (or MUST be absent, if so stated) in a conforming BGPSEC Router Certificate, except where explicitly noted otherwise. No other extensions are allowed in a conforming BGPSEC Router Certificate.

3.1.3.1. Basic Constraints

BGPSEC speakers are EEs; therefore, the Basic Constraints extension must not be present, as per [RFC6487].

3.1.3.2. Extended Key Usage

BGPSEC Router Certificates MUST include the Extended Key Usage (EKU) extension. As specified, in [RFC6487] this extension MUST be marked as non-critical. This document defines one EKU for BGPSEC Router Certificates:

```
id-kp OBJECT IDENTIFIER ::=
  { iso(1) identified-organization(3) dod(6) internet(1)
    security(5) mechanisms(5) pkix(7) kp(3) }
```

```
id-kp-bgpsec-router OBJECT IDENTIFIER ::= { id-kp TBD }
```

Relying Parties MUST require the extended key usage extension to be present in a BGPSEC Router Certificate. If multiple KeyPurposeId values are included, the relying parties need not recognize all of them, as long as the required KeyPurposeId value is present. BGPSEC RPs MUST reject certificates that do not contain the BGPSEC Router EKU even if they include the anyExtendedKeyUsage OID defined in [RFC5280].

3.1.3.3. Subject Information Access

This extension is not used in BGPSEC Router Certificates. It MUST be omitted.

3.1.3.4. IP Resources

This extension is not used in BGPSEC Router Certificates. It MUST be omitted.

3.1.3.5. AS Resources

Each BGPSEC Router Certificate MUST include the AS Resource Identifier Delegation extension, as specified in section 4.8.11 of [RFC6487]. The AS Resource Identifier Delegation extension MUST include exactly one AS number, and the "inherit" element MUST NOT be specified.

3.2. BGPSEC Router Certificate Request Profile

Refer to section 6 of [RFC6487]. The only differences between this profile and the profile in [RFC6487] are:

- o The ExtendedKeyUsage extension request MUST be included and the CA MUST honor the request;
- o The SubjectPublicKeyInfo and PublicKey fields are specified in [ID.sidr-bgpsec-algs]; and,
- o The request is signed with the algorithms specified in [ID.sidr-bgpsec-algs].

3.3. BGPSEC Router Certificate Validation

The validation procedure used for BGPSEC Router Certificates is identical to the validation procedure described in Section 7 of [RFC6487]. The exception is that the constraints applied come from this specification (e.g., in step 3: the certificate contains all the field that must be present - refers to the fields that are required by this specification).

The differences are as follows:

- o BGPSEC Router Certificates MUST include the BGPSEC ECU defined in Section 3.1.3.1.
- o BGPSEC Router Certificates MUST NOT include the SIA extension.
- o BGPSEC Router Certificates MUST NOT include the IP Resource extension.
- o BGPSEC Router Certificates MUST include the AS Resource Identifier Delegation extension.

- o BGPSEC Router Certificate MUST include the "Subject Public Key Info" described in [ID.sidr-bgpsec-algs] as it updates [RFC6485].

NOTE: The cryptographic algorithms used by BGPSEC routers are found in [ID.sidr-bgpsec-algs]. Currently, the algorithms specified in [ID.sidr-bgpsec-algs] and [RFC6485] are different. BGPSEC RPs will need to support algorithms that are needed to validate BGPSEC signatures as well as the algorithms that are needed to validate signatures on BGPSEC certificates, RPKI CA certificates, and RPKI CRLs.

4. Design Notes

The BGPSEC Router Certificate profile is based on the Resource Certificate profile as specified in [RFC6485]. As a result, many of the design choices herein are a reflection of the design choices that were taken in that prior work. The reader is referred to [RFC6484] for a fuller discussion of those choices.

5. Security Considerations

The Security Considerations of [RFC6487] apply.

A bgpsec certificate will fail RPKI validation, as defined in [RFC6487], because the algorithm suite is different. Consequently, a RP needs to identify the EKU before applying the correspondent validation.

A BGPSEC Router Certificate is an extension of the RPKI [RFC6480] to encompass routers. It is a building block of the larger BGPSEC security protocol used to validate signatures on BGPSEC Signature-Segment origination of Signed-Path segments [ID.sidr-bgpsec-protocol]. Thus its essential security function is the secure binding of an AS number to a public key, consistent with the RPKI allocation/assignment hierarchy.

6. IANA Considerations

None.

7. Acknowledgements

We would like to thanks Geoff Huston, George Michaelson, and Robert Loomans for their work on [RFC6487], which this work is based on. In addition, the efforts of Steve Kent and Matt Lepinski were instrumental in preparing this work. Additionally, we'd like to thank Roque Gagliano, Sandra Murphy, and Geoff Huston for their reviews and comments.

8. References

8.1. Normative References

- [RFC2119] Bradner, S., "Key words for use in RFCs to Indicate Requirement Levels", BCP 14, RFC 2119, March 1997.
- [RFC3779] Lynn, C., Kent, S., and K. Seo, "X.509 Extensions for IP Addresses and AS Identifiers", RFC 3779, June 2004.
- [RFC4271] Rekhter, Y., Ed., Li, T., Ed., and S. Hares, Ed., "A Border Gateway Protocol 4 (BGP-4)", RFC 4271, January 2006.
- [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.
- [RFC6485] Huston, G., "The Profile for Algorithms and Key Sizes for Use in the Resource Public Key Infrastructure (RPKI)", RFC 6485, February 2012.
- [RFC6487] Huston, G., Michaelson, G., and R. Loomans, "A Profile for X.509 PKIX Resource Certificates", RFC 6487, February 2012.
- [ID.sidr-bgpsec-algs] Reynolds, M. and S. Turner, "BGP Algorithms, Key Formats, & Signature Formats", draft-ietf-sidr-bgpsec-algs, work-in-progress.

8.2. Informative References

- [RFC4272] Murphy, S., "BGP Security Vulnerabilities Analysis", RFC 4272, January 2006.
- [RFC5123] White, R. and B. Akyol, "Considerations in Validating the Path in BGP", RFC 5123, February 2008.
- [RFC5492] Scudder, J. and R. Chandra, "Capabilities Advertisement with BGP-4", RFC 5492, February 2009.
- [RFC6480] Lepinski, M. and S. Kent, "An Infrastructure to Support Secure Internet Routing", RFC 6480, February 2012.
- [RFC6482] Lepinski, M., Kent, S., and D. Kong, "A Profile for Route Origin Authorizations (ROAs)", RFC 6482, February 2012.
- [RFC6484] Kent, S., Kong, D., Seo, K., and R. Watro, "Certificate Policy (CP) for the Resource Public Key Infrastructure

(RPKI)", BCP 173, RFC 6484, February 2012.

[RFC6486] Austein, R., Huston, G., Kent, S., and M. Lepinski,
"Manifests for the Resource Public Key Infrastructure
(RPKI)", RFC 6486, February 2012.

[ID.sidr-algorithm-agility] Gagliano, R., Kent, S., and S. Turner,
"Algorithm Agility Procedure for RPKI", draft-ietf-sidr-
algorithm-agility, work-in-progress.

[ID.sidr-bgpsec-protocol] Lepinski, M., "BGPSEC Protocol
Specification", draft-ietf-sidr-bgpsec-protocol, work-in-
progress.

Appendix A. ASN.1 Module

```
BGPSECEKU { iso(1) identified-organization(3) dod(6) internet(1)
  security(5) mechanisms(5) pkix(7) id-mod(0) TBD }

DEFINITIONS EXPLICIT TAGS ::=

BEGIN

-- EXPORTS ALL --

-- IMPORTS NOTHING --

-- OID Arc --

id-kp OBJECT IDENTIFIER ::= {
  iso(1) identified-organization(3) dod(6) internet(1)
  security(5) mechanisms(5) kp(3) }

-- BGPSEC Router Extended Key Usage --

id-kp-bgpsec-router OBJECT IDENTIFIER ::= { id-kp TBD }

END
```

Appendix B. Example BGPSEC Router Certificate

Appendix C. Example BGPSEC Router Certificate Request

Appendix D. Change Log

Please delete this section prior to publication.

D.1. Changes from sidr-bgpsec-pki-profiles-03 to sidr-bgpsec-pki-profiles-04

In s2.1, removed the phrase "another BGPSEC Router Certificate (only BGPSEC routers process these)" because the BGPSEC certificates are only ever EE certificates and they're never used to verify another certificate only the PDUs that are signed.

Added new s3.1.3.1 to explicitly state that EE certificates are only ever EE certs.

D.2. Changes from sidr-bgpsec-pki-profiles-02 to sidr-bgpsec-pki-profiles-03

Updated s3.3 to clarify restrictions on path validation procedures are in this specification (1st para was reworded).

Updated s3.3 to point to s3.1.3.1 for BGPSEC ECU (thanks Tom).

D.3. Changes from sidr-bgpsec-pki-profiles-01 to sidr-bgpsec-pki-profiles-02

Updated references.

D.4. Changes from sidr-bgpsec-pki-profiles-00 to sidr-bgpsec-pki-profiles-01

Added an ASN.1 Module and corrected the id-kp OID in s3.1.3.1.

D.5. Changes from turner-bgpsec-pki-profiles-02 to sidr-bgpsec-pki-profiles-00

Added this change log.

Amplified that a BGPSEC RP will need to support both the algorithms in [ID.sidr-bgpsec-algs] for BGPSEC and the algorithms in [ID.sidr-rpki-algs] for certificates and CRLs.

Changed the name of AS Resource extension to AS Resource Identifier Delegation to match what's in RFC 3779.

D.6. Changes from turner-bgpsec-pki-profiles -01 to -02

Added text in Section 2 to indicate that there's no impact on the procedures defined in [ID.sidr-algorithm-agility].

Added a security consideration to let implementers know the BGPSEC certificates will not pass RPKI validation [RFC6487] and that keying off the EKU will help tremendously.

D.7. Changes from turner-bgpsec-pki-profiles -00 to -01

Corrected Section 2 to indicate that CA certificates are also RPKI certificates.

Removed sections and text that was already in [RFC6487]. This will make it easier for reviewers to figure out what is different.

Modified Section 6 to use 2119-language.

Removed requirement from Section 6 to check that the AS # in the certificate is the last number in the AS path information of each BGP UPDATE message. Moved to [ID.sidr-bgpsec-protocol].

Authors' Addresses

Mark Reynolds
Island Peak Software
328 Virginia Road
Concord, MA 01742

Email: mcr@islandpeaksoftware.com

Sean Turner
IECA, Inc.
3057 Nutley Street, Suite 106
Fairfax, VA 22031
USA

Email: turners@ieca.com

Steve Kent
Raytheon BBN Technologies
10 Moulton St.
Cambridge, MA 02138

Email: kent@bbn.com

Network Working Group
Internet-Draft
Intended status: Informational
Expires: August 10, 2014

R. Bush
Internet Initiative Japan
February 6, 2014

RPKI Local Trust Anchor Use Cases
draft-ietf-sidr-lta-use-cases-00

Abstract

There are a number of critical circumstances where a localized routing domain needs to augment or modify its view of the Global RPKI. This document attempts to outline a few of them.

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 August 10, 2014.

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	2
2. Suggested Reading	2
3. What is 'Local'	2
4. Example Uses	3
5. Notes	3
6. Security Considerations	4
7. IANA Considerations	4
8. Acknowledgments	4
9. References	4
9.1. Normative References	4
9.2. Informative References	4
Author's Address	5

1. Introduction

Today RPKI-based Origin Validation, [RFC6811], relies on widespread deployment of the Global Resource Public Key Infrastructure (RPKI), [RFC6480]. In the future, RPKI-based Path Validation, [I-D.lepinski-bgpsec-overview], will be even more reliant on the Global RPKI.

But there are critical circumstances in which a local, well-scoped, administrative and/or routing domain will need to augment and/or modify their internal view of the Global RPKI.

This document attempts to lay out a few of those use cases. It is not intended to be authoritative, complete, or to become a standard. It merely tries to lay out a few critical examples to help scope the issues.

2. Suggested Reading

It is assumed that the reader understands the RPKI, see [RFC6480], the RPKI Repository Structure, see [RFC6481], Route Origin Authorizations (ROAs), see [RFC6482], and Ghostbusters Records, see [RFC6493].

3. What is 'Local'

The RPKI is a distributed database containing certificates, CRLs, manifests, ROAs, and Ghostbusters Records as described in [RFC6481]. Policies and considerations for RPKI object generation and maintenance are discussed elsewhere.

Like the DNS, the Global RPKI presents a single global view, although only a loosely consistent view, depending on timing, updating,

fetching, etc. There is no 'fix' for this, it is not broken, it is the nature of distributed data with distributed caches.

There are critical uses of the RPKI where a local administrative and/or routing domain, e.g. an end-user site, a particular ISP or content provider, a geo-political region, ... may wish to have a specialized view of the RPKI.

For the purposes of this exploration, we refer to this localized view as a 'Local Trust Anchor', mostly for historical reasons, but also because implementation would likely be the local distribution of one or more specialized trust anchors, [RFC6481].

4. Example Uses

Carol, a RIPE member, is a victim of the "Dutch Court Attack" (someone convinces a Dutch court to force the RIPE/NCC to remove or modify records) and we all want to save the ability to route to Carol's network(s). There is need for some channel through which we can exchange some local trust command and data group necessary to propagate patches local to all our caches.

Bob has a multi-AS network under his administration and some of those ASs use private ([RFC1918]) or 'borrowed' US military space, and he wishes to certify them for use in his internal routing.

Alice runs the root trust for a large organization where upper management has the router geeks pointing their competitors' prefixes to pictures of kittens and unicorns, and Alice is responsible for making the CA hierarchy have validated certificates for those redirected resources as well as the rest of the internet.

5. Notes

In these examples, it is ultimately the ROAs, not the certificates, which one wants to modify. But one can't just hack new ROAs as one does not have the private keys needed to sign them. Hence one has to first hack the 3779 certificates.

But we should not lose sight of the goal that it is the ROAs and Ghostbuster Records which need re-issuing under the new 3779 certificates.

Further, since we're not the NSA, GCHQ, ..., we can not assume that we can reissue down from the root trust anchor at the IANA or from the RIRs' certificates. So we have to create a new trust anchor which, for ease of use, will contain the new/hacked certificates and ROAs as well as the unmodified remainder of the Global RPKI.

And, because Alice, Bob, and Carol want to be able to archive, reproduce, and send to friends the data necessary to recreate their hacks, there will need to be a formally defined set of data which is input to a well-defined process to take an existing Global RPKI tree and produce the desired modified re-anchored tree.

6. Security Considerations

These use cases are all about violating global security, albeit within a constrained local context.

7. IANA Considerations

This document has no IANA Considerations.

8. Acknowledgments

The author wishes to thank Rob Austein.

9. References

9.1. Normative References

- [RFC6481] Huston, G., Loomans, R., and G. Michaelson, "A Profile for Resource Certificate Repository Structure", RFC 6481, February 2012.
- [RFC6482] Lepinski, M., Kent, S., and D. Kong, "A Profile for Route Origin Authorizations (ROAs)", RFC 6482, February 2012.
- [RFC6493] Bush, R., "The Resource Public Key Infrastructure (RPKI) Ghostbusters Record", RFC 6493, February 2012.
- [RFC6811] Mohapatra, P., Scudder, J., Ward, D., Bush, R., and R. Austein, "BGP Prefix Origin Validation", RFC 6811, January 2013.

9.2. Informative References

- [I-D.lepinski-bgpsec-overview] Lepinski, M. and S. Turner, "An Overview of BGPSEC", draft-lepinski-bgpsec-overview-00 (work in progress), March 2011.
- [RFC1918] Rekhter, Y., Moskowitz, B., Karrenberg, D., de Groot, G., and E. Lear, "Address Allocation for Private Internets", BCP 5, RFC 1918, February 1996.

[RFC6480] Lepinski, M. and S. Kent, "An Infrastructure to Support
Secure Internet Routing", RFC 6480, February 2012.

Author's Address

Randy Bush
Internet Initiative Japan
5147 Crystal Springs
Bainbridge Island, Washington 98110
US

Email: randy@psg.com

Network Working Group
Internet-Draft
Intended status: Informational
Expires: January 29, 2017

R. Bush
Internet Initiative Japan
July 28, 2016

Use Cases for Localized Versions of the RPKI
draft-ietf-sidr-lta-use-cases-07

Abstract

There are a number of critical circumstances where a localized routing domain needs to augment or modify its view of the Global RPKI. This document attempts to outline a few of them.

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 29, 2017.

Copyright Notice

Copyright (c) 2016 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	2
2. Suggested Reading	2
3. What is 'Local'	2
4. Example Uses	3
5. Some Approaches	3
6. Security Considerations	4
7. IANA Considerations	4
8. Acknowledgments	4
9. References	4
9.1. Normative References	5
9.2. Informative References	5
Author's Address	5

1. Introduction

Today RPKI-based Origin Validation, [RFC6811], relies on widespread deployment of the Global Resource Public Key Infrastructure (RPKI), [RFC6480]. In the future, RPKI-based Path Validation, [I-D.ietf-sidr-bgpsec-overview], will be even more reliant on the Global RPKI.

But there are critical circumstances in which a local, clearly-scoped, administrative and/or routing domain will want to augment and/or modify their internal view of the Global RPKI.

This document attempts to lay out a few of those use cases. It is not intended to be authoritative, complete, or to become a standard. It merely tries to lay out a few critical examples to help frame the issues.

2. Suggested Reading

It is assumed that the reader understands the RPKI, see [RFC6480], the RPKI Repository Structure, see [RFC6481], Route Origin Authorizations (ROAs), see [RFC6482], and GhostBusters Records, see [RFC6493].

3. What is 'Local'

The RPKI is a distributed database containing certificates, CRLs, manifests, ROAs, and GhostBusters Records as described in [RFC6481]. Policies and considerations for RPKI object generation and maintenance are discussed elsewhere.

Like the DNS, the Global RPKI tries to present a single global view, although only a loosely consistent view, depending on timing,

updating, fetching, etc. There is no 'fix' for this, it is not broken, it is the nature of distributed data with distributed caches.

There are critical uses of the RPKI where a local administrative and/or routing domain, e.g. an end-user site, a particular ISP or content provider, an organization, a geo-political region, ... may wish to have a specialized view of the RPKI.

For the purposes of this exploration, we refer to this localized view as a 'Local Trust Anchor', mostly for historical reasons, but also because implementation would likely require the local distribution of one or more specialized trust anchors, [RFC6481].

4. Example Uses

We explore this space using three examples.

Carol, a resource holder (LIR, PI holder, ...), operates outside of the country in which her RIR is based. Someone convinces the RIR's local court to force the RIR to remove or modify some or all of Carol's certificates, ROAs, etc. or the resources they represent, and the operational community wants to retain the ability to route to Carol's network(s). There is need for some channel through which operators can exchange local trust, command, and data collections necessary to propagate patches local to all their RPKI views.

Bob has a multi-AS network under his administration and some of those ASs use private ([RFC1918]) or 'borrowed' address space which is not announced on the global Internet (not to condone borrowing), and he wishes to certify them for use in his internal routing.

Alice is responsible for the trusted routing for a large organization, commercial or geo-political, in which management requests routing engineering to redirect their competitors' prefixes to socially acceptable data. Alice is responsible for making the CA hierarchy have validated certificates for those redirected resources as well as the rest of the Internet.

5. Some Approaches

In these examples, it is ultimately the ROAs, not the certificates, which one wants to modify or replace. But one probably can not simply create new ROAs as one does not have the private keys needed to sign them. Hence it is likely that one has to also do something about the [RFC6480] certificates.

The goal is to modify, create, and/or replace ROAs and GhostBuster Records which are needed to present the localized view of the RPKI data.

One wants to reproduce only as much of the Global RPKI as needed. Replicating more than is needed would amplify tracking and maintenance.

One can not reissue down from the root trust anchor at the IANA or from the RIRs' certificates because one does not have the private keys required. So one has to create a new trust anchor which, for ease of use, will contain the new/modified certificates and ROAs as well as the unmodified remainder of the Global RPKI.

Because Alice, Bob, and Carol want to be able to archive, reproduce, and send to other operators the data necessary to reproduce their modified view of the global RPKI, there will need to be a formally defined set of data which is input to a well-defined process to take an existing Global RPKI tree and produce the desired modified re-anchored tree.

It is possible that an operator may need to accept and process modification data from more than one source. Hence there is a need to merge modification 'recipes'.

6. Security Considerations

Though the above use cases are all constrained to local contexts, they violate the model of a single global PKI, albeit to meet real operational needs. Hence they MUST be implemented to assure the local constraint.

Authentication of modification 'recipes' will be needed.

7. IANA Considerations

This document has no IANA Considerations.

8. Acknowledgments

The author thanks Chris Morrow, Karen Seo, Rob Austein, and Steve Kent for comments and suggestions.

9. References

9.1. Normative References

- [RFC6480] Lepinski, M. and S. Kent, "An Infrastructure to Support Secure Internet Routing", RFC 6480, February 2012.
- [RFC6481] Huston, G., Loomans, R., and G. Michaelson, "A Profile for Resource Certificate Repository Structure", RFC 6481, February 2012.
- [RFC6482] Lepinski, M., Kent, S., and D. Kong, "A Profile for Route Origin Authorizations (ROAs)", RFC 6482, February 2012.
- [RFC6493] Bush, R., "The Resource Public Key Infrastructure (RPKI) Ghostbusters Record", RFC 6493, February 2012.
- [RFC6811] Mohapatra, P., Scudder, J., Ward, D., Bush, R., and R. Austein, "BGP Prefix Origin Validation", RFC 6811, January 2013.

9.2. Informative References

- [I-D.ietf-sidr-bgpsec-overview] Lepinski, M. and S. Turner, "An Overview of BGPSEC", draft-ietf-sidr-bgpsec-overview-02 (work in progress), May 2012.
- [RFC1918] Rekhter, Y., Moskowitz, B., Karrenberg, D., de Groot, G., and E. Lear, "Address Allocation for Private Internets", BCP 5, RFC 1918, February 1996.

Author's Address

Randy Bush
Internet Initiative Japan
5147 Crystal Springs
Bainbridge Island, Washington 98110
US

Email: randy@psg.com

Internet Engineering Task Force
Internet-Draft
Updates: 6485 (if approved)
Intended status: Standards Track
Expires: July 5, 2014

G. Michaelson, Ed.
G. Huston
APNIC
January 2014

Clarifying RPKI use of CMS SignerInfo"
draft-michaelson-signerinfo-01

Abstract

RFC6485 section 2 mandated a single CMS OID sha256withRSAEncryption from RFC4055 for use in the CMS SignerInfo field. This draft updates RFC6485 and extends it to permit the correct CMS use which includes an option of rsaEncryption for the SignerInfo field.

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 July 5, 2014.

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	2
1.1. Requirements Language	2
2. Revised CMS SignerInfo	2
3. Current Systems Behaviour	3
4. Acknowledgements	3
5. IANA Considerations	3
6. Security Considerations	3
7. References	3
7.1. Normative References	3
7.2. Informative References	4
Authors' Addresses	4

1. Introduction

RFC 6485 [RFC6485] defines The Profile for Algorithms and Key Sizes for Use in the Resource Public Key Infrastructure (RPKI). In that document, Section 2 specifies a single signature algorithm (SHA-256) and a single CMS OID, sha256withRSAEncryption, to be used for the SignerInfo field of the CMS object.

A closer reading of the relevant RFCs RFC 4055 [RFC4055] and RFC 5754 [RFC5754] identified that the CMS SignerInfo field must support use of the rsaEncryption OID for full conformance with the CMS specifications, and the normative references in RFC 6485 inherit the requirement.

To ensure full conformance with the CMS specifications, RFC 6485 is updated by this draft. All of RFC 6485 applies except for a change to the SignerInfo field.

1.1. Requirements Language

The key words "MUST", "MUST NOT", "REQUIRED", "SHALL", "SHALL NOT", "SHOULD", "SHOULD NOT", "RECOMMENDED", "MAY", and "OPTIONAL" in this document are to be interpreted as described in RFC 2119 [RFC2119].

2. Revised CMS SignerInfo

In RFC 6485 Section 2 the following sentence:

The Object Identifier (OID) sha256withRSAEncryption from [RFC4055] MUST be used.

Is replaced by:

One of the Object Identifiers (OID) rsaEncryption or sha256WithRSAEncryption from [RFC4055] MUST be used. RPKI implementations MUST support rsaEncryption for the signatureAlgorithm field and SHOULD support sha256WithRSAEncryption.

3. Current Systems Behaviour

All known RPKI CA implementations already do what this draft recommends.

4. Acknowledgements

Andrew Chi and David Mandelberg discovered this problem.

Russ Housley documented the RFC chain back to 2630.

This draft reflects a discussion between Rob Austein and Matt Lepinski on the SIDR Working group mailing list and a private communication between Rob Austein and Geoff Huston.

5. IANA Considerations

This memo includes no request to IANA.

6. Security Considerations

By conforming more closely to the CMS specifications, RPKI CMS objects are less likely to be rejected as non-conformant with the standards. No change is made to the cryptographic status of the CMS objects produced.

7. References

7.1. Normative References

[RFC2119] Bradner, S., "Key words for use in RFCs to Indicate Requirement Levels", BCP 14, RFC 2119, March 1997.

- [RFC4055] Schaad, J., Kaliski, B., and R. Housley, "Additional Algorithms and Identifiers for RSA Cryptography for use in the Internet X.509 Public Key Infrastructure Certificate and Certificate Revocation List (CRL) Profile", RFC 4055, June 2005.
- [RFC5754] Turner, S., "Using SHA2 Algorithms with Cryptographic Message Syntax", RFC 5754, January 2010.
- [RFC6485] Huston, G., "The Profile for Algorithms and Key Sizes for Use in the Resource Public Key Infrastructure (RPKI)", RFC 6485, February 2012.

7.2. Informative References

- [AUSTEIN] Austein, SR., "RFC 6485 is inconsistent with base CMS specifications", 2012, <<http://www.ietf.org/mail-archive/web/sidr/current/msg04813.html>>.

Authors' Addresses

George Michaelson (editor)
APNIC
6 Cordelia St, South Brisbane
Brisbane, Queensland 4101
AU

Phone: +61 7 3858 3150
Email: ggm@apnic.net

Geoff Huston
APNIC

Email: gih@apnic.net