Network Working Group
Internet-Draft
Intended status: Standards Track
Expires: March 28, 2016

S. Weiler
Parsons
A. Sonalker
Battelle Memorial Institute
R. Austein
Dragon Research Labs

September 25, 2015

A Publication Protocol for the Resource Public Key Infrastructure (RPKI)

draft-ietf-sidr-publication-07

Abstract

This document defines a protocol for publishing Resource Public Key Infrastructure (RPKI) objects. Even though the RPKI will have many participants issuing certificates and creating other objects, it is operationally useful to consolidate the publication of those objects. This document provides the protocol for doing so.

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 March 28, 2016.

Copyright Notice

Table of Contents

1. Introduction

This document assumes a working knowledge of the Resource Public Key Infrastructure (RPKI), which is intended to support improved routing security on the Internet. [RFC6480]

In order to make participation in the RPKI easier, it is helpful to

have a few consolidated repositories for RPKI objects, thus saving every participant from the cost of maintaining a new service. Similarly, relying parties using the RPKI objects will find it faster and more reliable to retrieve the necessary set from a smaller number of repositories.

These consolidated RPKI object repositories will in many cases be outside the administrative scope of the organization issuing a given RPKI object. In some cases, outsourcing operation of the repository will be an explicit goal: some resource holders who strongly wish to control their own RPKI private keys may lack the resources to operate a 24x7 repository, or may simply not wish to do so.

The operator of an RPKI publication repository may well be an Internet registry which issues certificates to its customers, but it need not be; conceptually, operation of a an RPKI publication repository is separate from operation of RPKI CA.

This document defines an RPKI publication protocol which allows publication either within or across organizational boundaries, and which makes fairly minimal demands on either the CA engine or the publication service.

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].

"Publication engine" and "publication server" are used interchangeably to refer to the server providing the service described in this document.

"Business Public Key Infrastructure" ("Business PKI" or "BPKI") refers to a PKI, separate from the RPKI, used to authenticate clients to the publication engine. We use the term "Business PKI" here because an internet registry might already have a PKI for authenticating its clients and might wish to reuse that PKI for this protocol. There is, however, no requirement to reuse such a PKI.

2. Protocol Specification

The publication protocol uses XML messages wrapped in signed CMS messages, carried over HTTP transport.

The publication protocol uses a simple request/response interaction. The client passes a request to the server, and the server generates a corresponding response.

A message exchange commences with the client initiating an HTTP POST with content type of "application/rpki-publication", with the message object as the body. The server's response will similarly be the body of the response with a content type of "application/rpki-publication".

The content of the POST and the server's response will be a well-formed Cryptographic Message Syntax (CMS) [RFC5652] object with OID = 1.2.840.113549.1.7.2 as described in Section 3.1 of [RFC6492].

2.1. Common XML Message Format

The XML schema for this protocol is below in Section 2.7. The basic XML message format looks like this:

```
<msg
    type="query"
    version="3"
    xmlns="http://www.hactrn.net/uris/rpki/publication-spec/">
  <!-- Zero or more PDUs -->
</msg>

<msg
    type="reply"
    version="3"
    xmlns="http://www.hactrn.net/uris/rpki/publication-spec/">
  <!-- Zero or more PDUs -->
</msg>
```

Common attributes:

version: The value of this attribute is the version of this protocol. This document describes version 3.

type: The possible values of this attribute are "reply" and "query".

A query PDU may be one of three types: <publish/>, <withdraw/>, or <list/>.

A reply PDU may be one of four types: <publish/>, <withdraw/>, <list/>, or <report_error/>.

Each of these PDUs may include an optional tag to facilitate bulk operation. If a tag is set in a query PDU, the corresponding reply(s) or error(s) MUST have the tag attribute set to the same value.

2.2. General Operation

Processing of a query message is handled atomically: either the entire query succeeds or none of it does. When a query message contains multiple PDUs, failure of any PDU may require the server to roll back actions triggered by earlier PDUs.

2.3. Publication and Withdrawal

The publication protocol uses a common message format to request publication of any RPKI object. This format was chosen specifically to allow this protocol to accommodate new types of RPKI objects without needing changes to this protocol.

Both the <publish/> and <withdraw/> PDUs have a payload of an optional tag and a URI. The <publish/> query also contains the DER object to be published, encoded in Base64.

Both the <publish/> and <withdraw/> PDUs also have a "hash" attribute, which carries a hash of an existing object at the specified repository URI. For <withdraw/> PDUs, the hash is mandatory, as this operation makes no sense if there is no existing object to withdraw. For <publish/> PDUs, the hash MUST be present if the publication operation is overwriting an existing object, and MUST be omitted if this publication operation is writing to a new URI where no prior object exists. Presence of an object when no hash attribute is specified is an error, as is absence of the hash attribute or an incorrect hash value when an object is present. Any such errors MUST be reported using the <report_error/> PDU.

The hash algorithm is SHA-256 [SHS], to simplify comparison of publication protocol hashes with RPKI manifest hashes.

The intent behind the hash attribute is to allow the client and server to detect any disagreements about the effect that a <publish/> or <withdraw/> PDU will have on the repository.

Note that every publish and withdraw action requires a new manifest, thus every publish or withdraw action will involve at least two objects.

2.4. Listing the repository

The <list/> operation allows the client to ask the server for a complete listing of objects which the server believes the client has published. This is intended primarily to allow the client to recover upon detecting (probably via use of the "hash" attribute, see Section 2.3) that they have somehow lost synchronization.

The <list/> query consists of a single PDU.

The <list/> reply consists of zero or more PDUs, one per object published in this repository by this client, each PDU conveying the URI and hash of one published object.

2.5. Error handling

Errors are handled at two levels.

Errors that make it impossible to decode a query or encode a response are handled at the HTTP layer. 4xx and 5xx HTTP response codes indicate that something bad happened.

In all other cases, errors result in an XML <report_error/> PDU which takes the place of the expected protocol response PDU. Like the rest of this protocol, <report_error/> PDUs are CMS-signed XML messages and thus can be archived to provide an audit trail.

<report_error/> PDUs only appear in replies, never in queries.

Like all other reply PDUs, if a "tag" attribute was set on the query that generated the error, the <report_error/> PDU MUST have its tag attribute set to the same value.

The error itself is conveyed in the error_code attribute. The value of this attribute is a token indicating the specific error that occurred.

The body of the <report_error/> element contains two sub-elements:

1. An optional text element <error_text/>, which if present, contains a text string with debugging information intended for human consumption.

2. An optional element <failed_pdu/>, which, if present, contains a verbatim copy of the query PDU whose failure triggered the <report_error/> PDU. The quoted element must be syntactically valid.

The position of a <report_error/> element in a reply corresponds to the point in processing the query message where the error occurred. In the simple case of a query message containing only a single element, the <report_error/> element will be the only element in the reply. If, however, the query message contains more than one element, the <report_error/> element may be preceeded by normal responses indicating operations that would have succeeded.

There are several ways that a client can match up elements in a response message with the corresponding elements in the query message:

- For a one-element query, this is trivial.

- For multi-element queries, the simplest way of matching resposes uses the optional tag attribute. The protocol requires tags from query elements to be copied into reply elements, so simply giving each query element a unique tag will suffice.

- If for some reason the client implementation is not able or willing to use unique tags within a multi-element query message, the client can still match queries to responses by counting

elements in the reply message. This approach is not recommended.

See Section 3.8 for examples of a multi-element query and responses.

2.6. Error Codes

These are the defined error codes as well as some discussion of each. Text similar to these descriptions may be sent in an <error_text/> element to help explain the error encountered.

permission_failure: Client does not have permission to update this URI.

bad_cms_signature: Bad CMS signature.

object_already_present: An object is already present at this URI, yet a hash attribute was not specified. A hash attribute must be specified when overwriting or deleting an object. Perhaps client and server are out of sync?

no_object_present: There is no object present at this URI, yet a hash attribute was specified. Perhaps client and server are out of sync?

no_object_matching_hash The hash attribute supplied does not match the hash attribute of the object at this URI. Perhaps client and server are out of sync?

consistency_problem: Server detected an update that looks like it will cause a consistency problem (e.g. an object was deleted, but the manifest was not updated). Note that a server is not required to make such checks. Indeed, it may be unwise for a server to do so. This error code just provides a way for the server to explain its (in-)action.

other_error: A meteor fell on the server.

2.7. XML Schema

The following is a RelaxNG compact form schema describing the Publication Protocol.

```
# $Id: rpki-publication.rnc 3407 2015-09-25 21:05:28Z sra $
# RelaxNG schema for RPKI publication protocol.

default namespace =
    "http://www.hactrn.net/uris/rpki/publication-spec/"

# This is version 3 of the protocol.

version = "3"

# Top level PDU is either a query or a reply.

start |= element msg {
  attribute version { version },
  attribute type    { "query" },
  query_elt*
}

start |= element msg {
  attribute version { version },
  attribute type    { "reply" },
  reply_elt*
}

# PDUs allowed in  queries and replies.

query_elt = publish_query | withdraw_query | list_query
reply_elt = publish_reply | withdraw_reply | list_reply | error_reply

# Tag attributes for bulk operations.

tag = attribute tag { xsd:token { maxLength="1024" } }

# Base64 encoded DER stuff.

base64 = xsd:base64Binary

# Publication URIs.

uri = attribute uri { xsd:anyURI { maxLength="4096" } }

# Digest of an existing object (hexadecimal).
```

```
   hash = attribute hash { xsd:string { pattern = "[0-9a-fA-F]+" } }

   # Error codes.

   error |= "permission_failure"
   error |= "bad_cms_signature"
   error |= "object_already_present"
   error |= "no_object_present"
   error |= "no_object_matching_hash"
   error |= "consistency_problem"
   error |= "other_error"

   # <publish/> element

   publish_query = element publish { tag?, uri, hash?, base64 }
   publish_reply = element publish { tag?, uri }

   # <withdraw/> element

   withdraw_query = element withdraw { tag?, uri, hash }
   withdraw_reply = element withdraw { tag?, uri }

   # <list/> element

   list_query = element list { tag? }
   list_reply = element list { tag?, uri, hash }

   # <report_error/> element

   error_reply = element report_error {
     tag?,
     attribute error_code { error },
     element   error_text { xsd:string { maxLength="512000" }}?,
     element   failed_pdu { query_elt }?
   }
```

3. Examples

Following are examples of various queries and the corresponding replies for the RPKI publication protocol.

Note the authors have taken liberties with the Base64, hash, and URI text in these examples in the interest of making the examples fit nicely into RFC text format.

3.1. <publish/> Query, No Existing Object

```
<msg
    type="query"
    version="3"
    xmlns="http://www.hactrn.net/uris/rpki/publication-spec/">
  <publish
      uri="rsync://wombat.example/Alice/60d730635fce156f.cer">
      WW91IGNhbiBoYWNrIGFueXRoaW5nIHlvdSB3YW50Li4u
    </publish>
</msg>
```

3.2. <publish/> Query, Overwriting Existing Object

```
<msg
    type="query"
    version="3"
    xmlns="http://www.hactrn.net/uris/rpki/publication-spec/">
  <publish
      hash="60d730635fce156f"
      uri="rsync://wombat.example/Alice/60d730635fce156f.cer">
      WW91IGNhbiBoYWNrIGFueXRoaW5nIHlvdSB3YW50Li4u
    </publish>
</msg>
```

3.3. <publish/> Reply

```
<msg
    type="reply"
    version="3"
    xmlns="http://www.hactrn.net/uris/rpki/publication-spec/">
  <publish
      uri="rsync://wombat.example/Alice/60d730635fce156f.cer"/>
</msg>
```

3.4. <withdraw/> Query

```
<msg
    type="query"
```

```
      version="3"
      xmlns="http://www.hactrn.net/uris/rpki/publication-spec/">
    <withdraw
        hash="60d730635fce156f"
        uri="rsync://wombat.example/Alice/60d730635fce156f.cer"/>
  </msg>
```


3.5. <withdraw/> Reply

```
  <msg
      type="reply"
      version="3"
      xmlns="http://www.hactrn.net/uris/rpki/publication-spec/">
    <withdraw
        uri="rsync://wombat.example/Alice/60d730635fce156f.cer"/>
  </msg>
```

3.6. <report_error/> With Optional Elements

```
  <msg
      type="reply"
      version="3"
      xmlns="http://www.hactrn.net/uris/rpki/publication-spec/">
    <report_error
        error_code="no_object_matching_hash">
      <error_text>
        Can't delete an object I don't have
      </error_text>
      <failed_pdu>
        <publish
            hash="60d730635fce156f"
            uri="rsync://wombat.example/Alice/60d730635fce156f.cer">
        WW91IGNhbiBoYWNrIGFueXRoaW5nIHlvdSB3YW50Li4u
      </publish>
      </failed_pdu>
    </report_error>
  </msg>
```

3.7. <report_error/> Without Optional Elements

```
<msg
    type="reply"
    version="3"
    xmlns="http://www.hactrn.net/uris/rpki/publication-spec/">
  <report_error
      error_code="object_already_present"/>
</msg>
```

3.8. Error Handling With Multi-Element Queries

3.8.1. Multi-Element Query


```
<msg
    type="query"
    version="3"
    xmlns="http://www.hactrn.net/uris/rpki/publication-spec/">
  <publish
      tag="Alice"
      uri="rsync://wombat.example/Alice/3bc51062973c458d.cer">
      QWxpY2U=
  </publish>
  <withdraw
      hash="cd9fb1e148ccd844"
      tag="Bob"
      uri="rsync://wombat.example/Bob/cd9fb1e148ccd844.cer"/>
  <publish
      tag="Carol"
      uri="rsync://wombat.example/Carol/b2dd7d8a70567a0e.cer">
      Q2Fyb2w=
  </publish>
  <list/>
  <withdraw
      hash="809a721743350c0c"
      tag="Dave"
      uri="rsync://wombat.example/Dave/809a721743350c0c.cer"/>
  <publish
      tag="Eve"
```

```
        uri="rsync://wombat.example/Eve/b9bae658d9657985.cer">
        RXZl
      </publish>
  </msg>
```

3.8.2. Successful Multi-Element Response


```
  <msg
      type="reply"
      version="3"
      xmlns="http://www.hactrn.net/uris/rpki/publication-spec/">
    <publish
        tag="Alice"
        uri="rsync://wombat.example/Alice/3bc51062973c458d.cer"/>
    <withdraw
        tag="Bob"
        uri="rsync://wombat.example/Bob/cd9fb1e148ccd844.cer"/>
    <publish
        tag="Carol"
        uri="rsync://wombat.example/Carol/b2dd7d8a70567a0e.cer"/>
    <list
        hash="f842c3e1858df8c8"
        uri="rsync://wombat.example/Fee/f842c3e1858df8c8.cer"/>
    <list
        hash="b139ca23414476bb"
```

```
        uri="rsync://wombat.example/Fie/b139ca23414476bb.cer"/>
    <list
        hash="1995e9544ba80191"
        uri="rsync://wombat.example/Foe/1995e9544ba80191.cer"/>
    <list
        hash="9c00b310c10a022c"
        uri="rsync://wombat.example/Fum/9c00b310c10a022c.cer"/>
    <withdraw
        tag="Dave"
        uri="rsync://wombat.example/Dave/809a721743350c0c.cer"/>
    <publish
        tag="Eve"
        uri="rsync://wombat.example/Eve/b9bae658d9657985.cer"/>
  </msg>
```

3.8.3. Failure Multi-Element Response


```
  <msg
      type="reply"
      version="3"
      xmlns="http://www.hactrn.net/uris/rpki/publication-spec/">
    <publish
        tag="Alice"
        uri="rsync://wombat.example/Alice/3bc51062973c458d.cer"/>
    <withdraw
        tag="Bob"
        uri="rsync://wombat.example/Bob/cd9fb1e148ccd844.cer"/>
    <publish
```

```
        tag="Carol"
        uri="rsync://wombat.example/Carol/b2dd7d8a70567a0e.cer"/>
    <list
        hash="f842c3e1858df8c8"
        uri="rsync://wombat.example/Fee/f842c3e1858df8c8.cer"/>
    <list
        hash="b139ca23414476bb"
        uri="rsync://wombat.example/Fie/b139ca23414476bb.cer"/>
    <list
        hash="1995e9544ba80191"
        uri="rsync://wombat.example/Foe/1995e9544ba80191.cer"/>
    <list
        hash="9c00b310c10a022c"
        uri="rsync://wombat.example/Fum/9c00b310c10a022c.cer"/>
    <report_error
        error_code="no_object_matching_hash"
        tag="Dave">
      <failed_pdu>
        <withdraw
            hash="809a721743350c0c"
            tag="Dave"
            uri="rsync://wombat.example/Dave/809a721743350c0c.cer"/>
      </failed_pdu>
    </report_error>
  </msg>
```

4. Operational Considerations

There are two basic options open to the repository operator as to how the publication tree is laid out. The first option is simple: each publication client is given its own directory one level below the top of the rsync module, and there is no overlap between the publication spaces used by different clients. For example:

```
rsync://example.org/rpki/Alice/
rsync://example.org/rpki/Bob/
rsync://example.org/rpki/Carol/
```


This has the advantage of being very easy for the publication operator to manage, but has the drawback of making it difficult for relying parties to fetch published objects both safely and as efficiently as possible.

Given that the mandatory-to-implement retrieval protocol for relying parties is rsync, a more efficient repository structure would be one which minimized the number of rsync fetches required. One such structure would be one in which the publication directories for subjects were placed underneath the publication directories of their issuers: since the normal synchronization tree walk is top-down, this can significantly reduce the total number of rsync connections required to synchronize. For example:

rsync://example.org/rpki/Alice/
rsync://example.org/rpki/Alice/Bob/
rsync://example.org/rpki/Alice/Bob/Carol/

Preliminary measurement suggests that, in the case of large numbers of small publication directories, the time needed to set up and tear down individual rsync connections becomes significant, and that a properly optimized tree structure can reduce synchronization time by an order of magnitude.

The more complex tree structure does require careful attention to the base_uri attribute values when setting up clients. In the example above, assuming that Alice issues to Bob who in turn issues to Carol, Alice has ceded control of a portion of her publication space to Bob, who has in turn ceded a portion of that to Carol, and the base_uri attributes in the <client/> setup messages should reflect this.

The details of how the repository operator determines that Alice has given Bob permission to nest Bob's publication directory under Alice's is outside the scope of this protocol.

5. IANA Considerations

IANA is asked to register the application/rpki-publication MIME media type as follows:

MIME media type name: application
MIME subtype name: rpki-publication
Required parameters: None
Optional parameters: None
Encoding considerations: binary
Security considerations: Carries an RPKI Publication Protocol
Message, as defined in this document.
Interoperability considerations: None
Published specification: This document
Applications which use this media type: HTTP
Additional information:
Magic number(s): None
File extension(s):
Macintosh File Type Code(s):
Person & email address to contact for further information:
Rob Austein <sra@hactrn.net>
Intended usage: COMMON
Author/Change controller: Rob Austein <sra@hactrn.net>

6. Security Considerations

The RPKI publication protocol and the data it publishes use entirely separate PKIs for authentication. The published data is authenticated within the RPKI, and this protocol has nothing to do with that authentication, nor does it require that the published objects be valid in the RPKI. The publication protocol uses a separate Business PKI (BPKI) to authenticate its messages.

Each RPKI publication protocol message is CMS-signed. Because of that protection at the application layer, this protocol does not require the use of HTTPS or other transport security mechanisms.

Although the hashes used in the <publish/> and <withdraw/> PDUs are cryptographic strength, the digest algorithm was selected for convenience in comparing these hashes with the hashes that appear in RPKI manifests. The hashes used in the <publish/> and <withdraw/> PDUs are not particularly security-sensitive, because these PDUs are protected by the CMS signatures.

Compromise of a publication server, perhaps through mismanagement of BPKI keys, could lead to a denial-of-service attack on the RPKI. An attacker gaining access to BPKI keys could use this protocol delete (withdraw) RPKI objects, leading to routing changes or failures. Accordingly, as in most PKIs, good key management practices are important.

7. References

7.1. Normative References

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

[RFC5652] Housley, R., "Cryptographic Message Syntax (CMS)", RFC 5652, STD 70, September 2009.

[RFC6492] Huston, G., Loomans, R., Ellacott, B., and R. Austein, "A Protocol for Provisioning Resource Certificates", RFC 6492, February 2012.

[SHS] National Institute of Standards and Technology, "Secure Hash Standard", FIPS PUB 180-4, March 2012, <http://csrc.nist.gov/publications/fips/fips180-4/fips-180-4.pdf>.

7.2. Informative References

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

Authors' Addresses

Samuel Weiler
Parsons

Email: weiler@tislabs.com

Anuja Sonalker
Battelle Memorial Institute

Email: sonalkera@battelle.org

Rob Austein
Dragon Research Labs

Email: sra@hactrn.net