

Network Working Group
Internet-Draft
Intended status: Standards Track
Expires: 3 September 2026

J. Snijders
BSD
F. Vompe
DT
2 March 2026

The RPKISPOOL Format for Materializing Resource Public Key
Infrastructure (RPKI) Data
draft-snijders-rpkispool-format-00

Abstract

This document describes a format and data storage approach for materialization of RPKI data in order to support a range of use cases such as auditing Certification Authorities and analytical research. The rpkispool format can be used for high-latency replication of raw RPKI data and associated validation outcomes as efficiently compressed durable objects. The method uses widely available standardized tooling and is designed to support long-term preservation of RPKI data in a cost-effective way.

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 <https://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 3 September 2026.

Copyright Notice

Copyright (c) 2026 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 (<https://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 Revised BSD License text as described in Section 4.e of the Trust Legal Provisions and are provided without warranty as described in the Revised BSD License.

Table of Contents

1. Introduction	2
2. Storage Concept and File Formats	3
3. Filename Scheme	3
3.1. Initstate filenaming scheme	3
3.2. Rpkispool filenaming scheme	4
4. Rpkispool Construction	5
5. RPKIViews Mirrors	5
6. Security Considerations	6
7. IANA Considerations	6
8. References	6
8.1. Normative References	6
8.2. Informative References	7
Authors' Addresses	7

1. Introduction

The ability to economically archive multiple years worth of RPKI data produced by Certification Authorities (CAs) worldwide is essential for ongoing protocol maintenance, development of best practises, and incident research. This document describes a format and data storage approach for efficiently materializing RPKI data for long-term preservation in compact archives using standardized tooling.

The [RPKIViews] project adheres to the `_Sushi Principle_ ([Sushi])`: "raw" data is better than "cooked" data. The guiding principle is that for maximum flexibility, to allow for future and unforeseen use cases, the data should be accessible in its original form, rather than some aggregated or processed form. In order to collect RPKI material the RPKIViews project employs multiple topologically and geographically diverse vantage points and synchronizes using both Rsync and RRDP.

In February 2026, using the method described in this document, the [RPKIViews] project discovered and stored 4,961,325 RPKI objects ([RFC6487], [RFC6488]) and produced 53,826 CCRs ([I-D.ietf-sidrops-rpki-ccr]). Together this data would consume 1.2 TB in uncompressed form, however after compression only 16.3 GB remained, a 98.6% reduction. The daily checkpoints together consumed 307 GB in uncompressed form and 14 GB in compressed form, a 95.4% reduction in size. In other words, a full month's worth of RPKI data only mandated 30 GB of disk space. Storing all the world's RPKI data at a rate of roughly 1 GB per day makes research fairly accessible and affordable.

2. Storage Concept and File Formats

To capture the global RPKI's endless stream of data, batch processors divide the data stream into chunks of fixed duration, processing a day's worth of data at the end of every day.

Each day starts with a set of initial and internally consistent snapshots, which together form the initstate, and throughout the day all change data, i.e., all newly discovered RPKI objects and associated validation outcome states (CCRs) are appended to a log: the rpkispool.

Bundling the RPKI objects together with CCRs and sorting the archive members in a specific way optimally allows compression algorithms to find redundant data and significantly improves compression ratios.

Archive files are formatted following the [ustar] specification and compressed in Zstandard [RFC8878] form with windowLog 27 (--long) at compression level 19.

3. Filename Scheme

This section describes the filenames used for the archive members. The filenaming scheme was designed to allow researchers to extract multiple rpkispool archives in a single directory without naming conflicts.

3.1. Initstate filenaming scheme

The filenames of the members of an initstate archive are constructed as follows: \${RPKIVIEWS_NODE_ID}/\${PUBLICATION_POINT_FQDN}/path/to/object.\${EXTENSION}.

An example is as follows:

```

$ zstdcat 20260301-initstate.tar.zst | tar -tf- | head -n 20
ams1/akane.maru.co.jp/repo/1073c6/1/3134302e3233352e3139392e302f32342d3234203d3e203231
34363735.roa
blr1/akane.maru.co.jp/repo/1073c6/1/3134302e3233352e3139392e302f32342d3234203d3e203231
34363735.roa
blr2/akane.maru.co.jp/repo/1073c6/1/3134302e3233352e3139392e302f32342d3234203d3e203231
34363735.roa
dus1/akane.maru.co.jp/repo/1073c6/1/3134302e3233352e3139392e302f32342d3234203d3e203231
34363735.roa
miso/akane.maru.co.jp/repo/1073c6/1/3134302e3233352e3139392e302f32342d3234203d3e203231
34363735.roa
nyc1/akane.maru.co.jp/repo/1073c6/1/3134302e3233352e3139392e302f32342d3234203d3e203231
34363735.roa
sng1/akane.maru.co.jp/repo/1073c6/1/3134302e3233352e3139392e302f32342d3234203d3e203231
34363735.roa
syd1/akane.maru.co.jp/repo/1073c6/1/3134302e3233352e3139392e302f32342d3234203d3e203231
34363735.roa
syd2/akane.maru.co.jp/repo/1073c6/1/3134302e3233352e3139392e302f32342d3234203d3e203231
34363735.roa
yyz1/akane.maru.co.jp/repo/1073c6/1/3134302e3233352e3139392e302f32342d3234203d3e203231
34363735.roa
zur1/akane.maru.co.jp/repo/1073c6/1/3134302e3233352e3139392e302f32342d3234203d3e203231
34363735.roa
zur2/akane.maru.co.jp/repo/1073c6/1/3134302e3233352e3139392e302f32342d3234203d3e203231
34363735.roa
ams1/akane.maru.co.jp/repo/1073c6/1/F03205B3993400CC3FC657EFE68E0696DA332458.crl
blr1/akane.maru.co.jp/repo/1073c6/1/F03205B3993400CC3FC657EFE68E0696DA332458.crl
blr2/akane.maru.co.jp/repo/1073c6/1/F03205B3993400CC3FC657EFE68E0696DA332458.crl
dus1/akane.maru.co.jp/repo/1073c6/1/F03205B3993400CC3FC657EFE68E0696DA332458.crl
miso/akane.maru.co.jp/repo/1073c6/1/F03205B3993400CC3FC657EFE68E0696DA332458.crl
nyc1/akane.maru.co.jp/repo/1073c6/1/F03205B3993400CC3FC657EFE68E0696DA332458.crl
sng1/akane.maru.co.jp/repo/1073c6/1/F03205B3993400CC3FC657EFE68E0696DA332458.crl
syd1/akane.maru.co.jp/repo/1073c6/1/F03205B3993400CC3FC657EFE68E0696DA332458.crl

```

3.2. Rpkispool filenames scheme

The filenames of the members of an rpkispool archive are constructed in one of two ways. Members whose filename starts with static/ are DER-encoded RPKI objects where the filename is the SHA-256 ([SHS]) of the object encoded using Base64 with the filename safe alphabet (Section 5 of [RFC4648]). For performance reasons, the directory hierarchy in static/is constructed using the last few bytes of the SHA-256.

Any other members are auxiliary data and grouped according to production date, using the form
 \${YEAR}/\${MONTH}/\${DAY}/\${ISO8601}-\${NODEID}.\${EXTENSION}. The filename extension signifies the type of file.

An example is as follows:


```
$ zstdcat 20260301-rpkispool.tar.zst | tar -tf- | grep -m 1 -B 10 -A 10 2026/
static/xs/Ec/-fNEEKsF1NLwAtLAxQtqolj7UXWj9I9nJjBWGwlxsEc
static/yI/iA/Ix5u73xRGkXAwjg2-93ULtv-yHXqV9_ucDIXZ-5yIiA
static/ya/Z8/vfnWEX976G9VluRL_-i6G-G4DICffwpq7drYOnJyaZ8
static/zM/Q4/rleR5X9N8slvj_zuwV7JyleKmfcAfglrdlCCuHEzMQ4
static/-j/lc/lJ7XMmaXnUt0OstDsW4rKBVi8XKE5lr6iC4sxsq8-jlc
static/0K/LU/HvnnCbGHIE42NpG6Fpx8NK_94LLNHHbZsuh1Q710KLU
static/0s/8Y/eQmMs0N8T2FObu-k7HorubQqUrVQd3lkM7Mm_kZ0s8Y
static/1A/HE/QWFu2t5XTsuMGhaVowKVMrKyRmLlHJlmlqL7uf0M1AHE
static/1D/XQ/sKeIyf9Fj71hW0R0SPcFViBchkZvcNhi6VL45V-1DXQ
static/1Z/jU/2SMhqOUgK5NYtq3L06ZqFHvZdKbhdI9HYAbCPlclZjU
2026/03/01/20260301T000035Z-miso.ccr
2026/03/01/20260301T000035Z-miso.log
2026/03/01/20260301T000035Z-miso.metrics
static/2C/M8/gRAUD18Gla4Sr5A0jrh7kTpX4Yj1Zfrywjt5yeD2CM8
static/2N/j4/TVkptGmN3prJANdeWxQVS1Bt-UpnjtDX6RzIG-B2Nj4
static/2k/7U/vKNQKfG--QClvWrie6lq_LZSPiXhkcxHK_Thfkf2k7U
static/2p/YM/jRRbQXFhjhlk-Xd6jAgShu3vleyFxfIekY_BUSc2pYM
static/3I/rM/P-9fPiA2KjBGj9WTOzv2ESP9JDpbUzjPeIaSjfb3IrM
static/3K/0s/j_TEWaLDXKqgHPxUC-im0MxTfwwJCpI5tkMHoG93K0s
static/3Q/Lk/gM_ST0n22aL3Di66t90ntLACau2tQHv12g8rb7J3QLk
static/4Q/0I/3LP1SRW2KXKgA8SpsyF-F4LHyBAoCh9_tJxkNir4Q0I
```

4. Rpkispool Construction

The vantage points are relying party instances which periodically perform synchronization and validation and produce a CCR for each iteration. After each iteration the previous and current CCR are compared to deduce which objects are newly discovered. Newly discovered objects are appended to an hourly tar archive together with the CCR and any other auxiliary data files.

These hourly intermediate states are collected and materialized to a durable distributed filesystem serve as input to the daily compactization process. The daily compactization process deduplicates objects, normalizes timestamps using the method specified in [RFC9589], and performs Zstandard compression.

5. RPKIViews Mirrors

This section is to be removed before publishing as an RFC.

The [RPKIViews] project produces a daily rpkispool combined from 10+ vantage points worldwide. Starting January 1st, 2026, every day a new rpkispool file is made available containing all yesterday's RPKI data. These rpkispool files are mirrored at the following publicly accessible locations:

Netherlands rsync://josephine.sobornost.net/rpki/rpkispools/

Japan rsync://dango.attn.jp/rpki/rpkispools/

United States rsync://rpkiviews.kerfuffle.net/rpki/rpkispools/

6. Security Considerations

The storage format provides no authenticity and may appear to be zip bombs.

7. IANA Considerations

This document has no IANA actions.

8. References

8.1. Normative References

- [I-D.ietf-sidrops-rpki-ccr] Snijders, J., Bakker, B., Bruijnzeels, T., and T. Buehler, "A Profile for Resource Public Key Infrastructure (RPKI) Canonical Cache Representation (CCR)", Work in Progress, Internet-Draft, draft-ietf-sidrops-rpki-ccr-02, 4 December 2025, <<https://datatracker.ietf.org/doc/html/draft-ietf-sidrops-rpki-ccr-02>>.
- [RFC4648] Josefsson, S., "The Base16, Base32, and Base64 Data Encodings", RFC 4648, DOI 10.17487/RFC4648, October 2006, <<https://www.rfc-editor.org/info/rfc4648>>.
- [RFC6487] Huston, G., Michaelson, G., and R. Loomans, "A Profile for X.509 PKIX Resource Certificates", RFC 6487, DOI 10.17487/RFC6487, February 2012, <<https://www.rfc-editor.org/info/rfc6487>>.
- [RFC6488] Lepinski, M., Chi, A., and S. Kent, "Signed Object Template for the Resource Public Key Infrastructure (RPKI)", RFC 6488, DOI 10.17487/RFC6488, February 2012, <<https://www.rfc-editor.org/info/rfc6488>>.
- [RFC8878] Collet, Y. and M. Kucherawy, Ed., "Zstandard Compression and the 'application/zstd' Media Type", RFC 8878, DOI 10.17487/RFC8878, February 2021, <<https://www.rfc-editor.org/info/rfc8878>>.

- [RFC9589] Snijders, J. and T. Harrison, "On the Use of the Cryptographic Message Syntax (CMS) Signing-Time Attribute in Resource Public Key Infrastructure (RPKI) Signed Objects", RFC 9589, DOI 10.17487/RFC9589, May 2024, <<https://www.rfc-editor.org/info/rfc9589>>.
- [SHS] National Institute of Standards and Technology, "Secure Hash Standard", FIPS PUB 180-4, August 2015.
- [ustar] IEEE/Open Group, "ustar Interchange Format", IEEE Std 1003.1-2024, DOI 10.1109/IEEESTD.2018.8423794, June 2024, <https://pubs.opengroup.org/onlinepubs/9799919799/utilities/pax.html#tag_20_94_13_06>.

8.2. Informative References

- [RPKIViews] Snijders, J., "The RPKIViews Project", 2026, <<https://www.rpkiviews.org/>>.
- [Sushi] Johnson, B. and J. Adler, "The Sushi Principle: Raw Data Is Better", February 2015, <<https://web.archive.org/web/20161126104941/https://conferences.oreilly.com/strata/big-data-conference-ca-2015/public/schedule/detail/38737>>.

Authors' Addresses

Job Snijders
BSD Software Development
Amsterdam
Netherlands
Email: job@bsd.nl
URI: <https://www.bsd.nl/>

Fedor Vompe
Deutsche Telekom
Münster
Germany
Email: Theodor-Fedor.Vompe@telekom.de