

The Global RPKI Report

Vol. 1, No. 1

Job Snijders
job@bsd.nl

July - September 2026

Introduction

RPKI, everyone knows it!¹ But how *well* do we know it? Is the RPKI in good shape? Welcome to the inaugural quarterly technical report for RPKI professionals!

This series covers observations related to scalability and reliability of the global RPKI distributed database, comprehensive analysis of incidents in certification authority and trust anchor operations, and executive summaries of developments in standardization forums and the field itself. It is hoped that these reports will contribute toward an interoperable dependable global RPKI.

Methodology

The primary data source for reports is [RPKIViews](#), an open data project aiming to discover and store all the world's RPKI objects in their original form. Data is gathered through a wide array of geographically diverse vantage points using RSYNC and RRDP synchronization transport protocols in parallel. Collection and analysis are performed using [rpki-client](#), [rrdp-tools](#), and proprietary tooling.

RPKIViews is regarded as the most complete global RPKI capture available. This report includes various statistics produced from the fully merged and deduplicated capture dataset. This approach filters out any noise resulting from retransmissions of data caused by protocol idiosyncrasies and intermittent network failures between vantage points and publication points.

The IETF [SIDROPS](#) mailing list is also used to track ecosystem developments.

¹Resource Public Key Infrastructure, a specialized X.509 PKI framework to support improved security for the global Internet routing system.

Toward understanding global RPKI convergence

Because operators may need to urgently issue or revoke ROAs as part of incident response, the RPKI’s usefulness depends on the data’s ability to propagate globally in a timely fashion. When synchronization timeouts occur, any newly signed ROAs or revocations that are queued up for distribution on the server side will only become effective after the client’s next successful run. For some publication points, many more timeout events were observed in the second and third quarter of 2026 than for other publication points, as shown in Figure 1.




Figure 1: Each dot represents a timeout logged for specific publication points.

Causes of synchronization issues

A number of factors influence whether a given {publication point, client} pair can establish a synchronized state:

- The overall size of the repositories hosted at the publication point.
- The total number of objects in a repository.
- The rate at which new objects replace old objects (churn rate).
- The size of the objects.
- Packet loss on the connectivity paths between the publication point and the client.

The overall size of a repository is the product of both object count and object size. Object count mostly depends on whether an issuing CA can pack many payloads into fewer objects. Conversely, the more objects concurrently exist for a given CA, the larger that CA’s manifest will be. Signed object size follows from unavoidable factors such as the size of the cryptographic signature and dynamic factors such as the choice of directory hierarchy and filename naming scheme. The size of a given CA’s CRL relates to the default validity period of the certificates the CA issued: the longer the validity period the higher the chances the certificate will need to be revoked before expiry and the longer the certificate’s serial remains on the CRL.

A CA issuing ROAs containing just a single prefix and using a lengthy naming scheme and overly long validity periods will see a compounding detrimental effect: single prefix ROAs incur the maximum possible overhead, directory and filename strings (long or short) are unavoidably repeated in various fields in the end-entity certificate and transport layer meta data (RSYNC & RRDP), causing bloat, and many objects means larger manifests and CRLs.

From this it follows that the smaller the set of data that needs to be distributed from publication point to clients, the higher the chances of swiftly achieving successful synchronization.

Initial State Statistics

The initial state represents what a fully synchronised client will at a minimum have to store in its internal data store for each of the listed CA systems.

Number of objects in initial state

The following overview compares the world’s largest CA systems in terms of object count. Each object carries one cryptographic signature. If the client uses a relational database it will reflect in table size, if it uses a filesystem as backend these numbers will reflect in inode consumption.

Table 1: Number of objects stored client-side per base location in a fully populated cache. Shorthands are used to denote locations, see [Table 11](#).

Base location	Apr '26	May'26	Jun '26	Jul '26	Aug '26	Sep '26
arin.net	217k	222k	227k	230k	234k	236k
ripe.net	114k	115k	116k	121k	119k	121k
apnic.net	41k	41k	41k	41k	41k	46k
registro.br	27k	27k	27k	28k	29k	29k
lacnic.net	26k	26k	27k	27k	27k	27k
amazonaws.com	21k	21k	21k	21k	21k	21k
afrinic.net	15k	15k	16k	16k	16k	16k
idnic.net	16k	13k	16k	16k	16k	16k
paas.ripe.net	8k	9k	9k	9k	9k	9k
rps.cnnic.cn	2k	6k	8k	9k	9k	9k
nic.ad.jp	8k	8k	8k	8k	8k	8k
fiti.net.cn	0	0	4k	4k	4k	4k
rps.arin.net	2k	2k	3k	3k	2k	3k
twNIC.tw	6k	6k	3k	3k	3k	3k

Overall size of all objects in initial state

The overview in Table 2 shows the overall size calculated by adding the length of all DER encoded objects together. Clients will always incur some additional overhead to actually store this data (index offsets, B-trees). When a client bootstraps a cold cache this is the minimum amount of data that would need to be transferred (modulo transport overhead).

Table 2: Client-side data stored per base location in a fully populated cache (bytes).

Base location	Apr '26	May'26	Jun '26	Jul '26	Aug '26	Sep '26
arin.net	464M	475M	485M	493M	503M	506M
ripe.net	186M	188M	190M	198M	191M	194M
apnic.net	72M	73M	74M	75M	76M	77M
registro.br	48M	49M	49M	51M	53M	54M
lacnic.net	49M	49M	50M	50M	50M	51M
amazonaws.com	45M	43M	45M	45M	46M	46M
afrinic.net	28M	29M	30M	30M	31M	31M
idnic.net	26M	23M	27M	27M	27M	27M
paas.ripe.net	17M	18M	18M	19M	19M	20M
rps.cnnic.cn	4M	11M	15M	18M	17M	17M
nic.ad.jp	13M	13M	13M	13M	14M	14M
fiti.net.cn	0	0	7M	8M	8M	8M
rps.arin.net	5M	5M	5M	5M	5M	5M
twmic.tw	10M	10M	5M	5M	5M	5M

RRDP Statistics

RRDP statistics should be considered separately from the aforementioned *initial state* statistics. For example, an object transmitted through RRDP might not be a valid object and as such not count toward any particular CA system in the initial state. It is also technically possible to produce internally inconsistent RRDP deltas in which a coherent state is *smeared* across multiple deltas, in such cases the number of deltas will be higher than the number of validation states.

Each RRDP delta is its own XML document containing binary payloads. The volume represents the total size of the dataset newly presented to clients for consideration. Actual network traffic will be much higher than the numbers below might suggest because these numbers entirely exclude retries and any waste caused by partial downloads. As such, these numbers represent a *lower bound* for network traffic.

While the total number of RRDP deltas did not increase significantly month-to-month, the total size grew by almost 50%: from 101.8 GiB in the second quarter to 146.4 GiB in the third quarter of 2026.

Table 3: Number of new RRDP deltas and volume per month.

Period	Apr '26	May'26	Jun '26	Jul'26	Aug'26	Sep'26
# of Deltas	215,035	241,329	230,450	240,356	238,359	220,073
Volume (bytes)	33.3G	33.3G	34.5G	40.8G	49.3G	54.6G

A per-location breakdown shows that one location in particular exhibited significant growth, far exceeding its *initial state* in terms of churn.

Table 4: Newly discovered RRDP deltas, publish elements, and volume in bytes per RRDP location per quarter.

Location	Q2 2026			Q3 2026		
	Deltas	Pub	Volume	Deltas	Pub	Volume
idnic.net	168k	455k	1.8G	202k	569k	2.0G
registro.br	133k	955k	37.2G	129k	990k	78.6G
rps.cnnic.cn	57k	184k	1.2G	38k	121k	0.8G
lacnic.net	50k	315k	1.0G	48k	315k	1.0G
nic.ad.jp	44k	756k	1.8G	46k	763k	1.8G
twNIC.tw	39k	102k	0.4G	36k	98k	0.3G
paas.ripe.net	35k	124k	2.5G	37k	182k	7.8G
apnic.net	34k	1,051k	5.2G	34k	1,042k	5.1G
ripe.net	22k	11,814k	32.5G	23k	11,630k	28.4G
amazonaws.com	18k	67k	2.9G	20k	71k	3.5G
arin.net	17k	2,032k	11.9G	18k	2,231k	14.0G
rps.arin.net	13k	40k	0.3G	13k	42k	0.3G
afrinic.net	11k	290k	2.5G	11k	293k	2.4G

Newly Signed Data Statistics

Newly signed data (by object count)

After initial synchronization, a client will continuously update its local cache replacing or adding new objects based on new information from the CAs. It appears between five and six million objects pass through the ecosystem every month.

Table 5: Total number of newly discovered RPKI objects per month.

	Apr '26	May'26	Jun '26	Jul '26	Aug '26	Sep '26
Objects	5,312,637	5,731,081	5,587,355	5,927,525	5,936,302	5,817,566

Newly signed data (by volume)

The five to six million new objects together consume between 19 and 35 gigabytes per month.

Table 6: Newly discovered objects as data volume in bytes per month.

	Apr '26	May'26	Jun '26	Jul '26	Aug '26	Sep '26
Volume	19,598M	21,836M	22,791M	26,526M	31,687M	34,497M

Breakdowns per object type

What follows is a break down per signed object type.

Table 7: Newly discovered signed objects per type per quarter.

	Q2 2026	Q3 2026
Manifests	8,122,073	8,493,374
CRLs	8,116,283	8,487,012
ROAs	357,013	532,103
Certificates	37,050	49,000
ASPAs	2,385	4,137

The above can be broken down into the originating location of the new data. For readability the data is presented in two separate tables (below): one for the second quarter of 2026 followed by one covering the third quarter of 2026.

Note that CAs related to locations where the new manifest count is not exactly the same as the new CRLs count might not be following established practice. [RFC 9286 section 7](#) specifies: “Upon creation of the publication point associated with a CA, the CA *MUST* create and publish a manifest as well as a CRL.”

Table 8: Newly discovered signed objects per type per location in Q2 2026.

Location	Certs	Manifests	CRLs	ROAs	ASPAs
ripe.net	3,674	5,350,417	5,350,417	24,994	1,296
arin.net	2,677	851,660	851,660	220,341	639
registro.br	19,593	425,295	425,295	8,072	84
apnic.net	3,384	413,529	413,529	11,971	31
nic.ad.jp	2,012	357,311	357,190	325	0
idnic.net	2,481	214,542	214,542	10,863	0
lacnic.net	208	148,501	148,501	2,630	0
afrinic.net	95	128,849	128,849	6,296	0
paas.ripe.net	1,049	48,957	48,957	12,508	237
twNIC.tw	416	48,342	48,342	1,989	0
rps.cnnic.cn	875	68,151	68,151	16,555	0
amazonaws.com	0	14,313	8,644	31,325	2
rps.arin.net	242	17,070	17,070	2,423	32

Table 9: Newly discovered signed objects per type per location in Q3 2026.

Location	Certs	Manifests	CRLs	ROAs	ASPAs
ripe.net	2,559	5,502,213	5,502,213	33,740	1,336
arin.net	3,160	935,774	935,774	413,095	1,441
registro.br	29,682	449,861	449,861	10,436	101
apnic.net	3,208	433,859	433,859	13,745	787
nic.ad.jp	2,149	379,418	379,280	250	0
idnic.net	26	280,838	280,838	612	0
lacnic.net	291	154,802	154,802	2,987	24
afrinic.net	113	134,444	134,444	6,758	0
paas.ripe.net	6,810	61,324	61,324	15,196	193
twNIC.tw	64	48,276	48,276	177	0
rps.cnnic.cn	67	43,003	43,003	1,728	0
amazonaws.com	0	15,519	9,295	28,771	2
rps.arin.net	331	18,103	18,103	2,319	14

Object size comparison

As mentioned in [causes of synchronization issues](#), a multitude of intertwining factors determine the object size. An uptick in the *Mean* size suggests that the associated repository now contains a small but substantial number of significantly larger objects.

Table 10: Mean and median object sizes and total volume per location (bytes).

Location	Mean		Median		Total Volume	
	Q2'26	Q3'26	Q2'26	Q3'26	Q2 '26	Q3 '26
registro.br	25,216	48,147	2,081	1,997	22.1G	45.5G
ripe.net	2,014	1,766	1,924	1,923	21.6G	19.7G
arin.net	4,406	4,693	2,126	2,126	8.5G	10.5G
apnic.net	3,902	3,823	1,979	1,979	3.2G	3.4G
afrinic.net	6,904	6,548	2,049	2,049	1.8G	1.8G
paas.ripe.net	15,628	34,987	2,031	2,083	1.7G	6.0G
amazonaws.com	29,341	36,336	2,056	2,056	1.6G	2.1G
nic.ad.jp	1,581	1,560	1,859	1,857	1.1G	1.2G
idnic.net	1,797	1,641	1,845	1,967	795M	931M
lacnic.net	1,773	1,855	1,998	1,998	532M	584M
rps.cnnic.cn	3,392	3,621	1,936	2,029	521M	429M
rps.arin.net	4,385	4,166	2,106	2,156	162M	162M
twmic.tw	1,556	1,581	1,769	1,885	154M	155M

Table 11: Location Shorthands

Shorthand	Base FQDN	RRDP FQDN
afrinic.net	rpki.afrinic.net	rrdp.afrinic.net
amazonaws.com	rpki-rsync.us-east-2 .amazonaws.com	rpki-rrdp.us-east-2 .amazonaws.com
apnic.net	rpki.apnic.net	rrdp.apnic.net
arin.net	rpki.arin.net	rrdp.arin.net
fiti.net.cn	rsync.fiti.net.cn	rpki.fiti.net.cn
idnic.net	rpki-rsync.idnic.net	rpki-rrdp.idnic.net
lacnic.net	repository.lacnic.net	rrdp.lacnic.net
nic.ad.jp	rpki-repository.nic.ad.jp	rpki-repository.nic.ad.jp
paas.ripe.net	rsync.paas.rpki.ripe.net	rrdp.paas.rpki.ripe.net
registro.br	rpki-repo.registro.br	rpki-repo.registro.br
ripe.net	rpki.ripe.net	rrdp.ripe.net
rps.arin.net	rpki-rps.arin.net	rrdp-rps.arin.net
rps.cnnic.cn	rpki-rps.cnnic.cn	rrdp-rps.cnnic.cn
twmic.tw	rpkica.twmic.tw	rpkica.twmic.tw

RPKI Newswire

This section contains curated news from the field.

ARIN improved standards compliance

ARIN resolved two long-standing standards compliance issues in August 2026.

Background: Leftovers from an older signer implementation (which ARIN phased out in August 2020) remained in the ecosystem in the form of a few thousand incorrectly encoded objects.

Specifically, a number of objects were encoded using BER coding (indefinite length encoding) while the RPKI standards only permit DER coding. The other issue was that a number of objects incorrectly omitted parameters in `sha256WithRSAEncryption AlgorithmIdentifier` instances.

In recent years the RPKI community has taken to reducing variability in encoding, this cleanup is an important step forward. ARIN was the last to use BER coding.

AfriNIC improved standards compliance

AfriNIC completed a multi-year project to reissue all non-conformant certificates in August 2026. See AfriNIC's final update [here](#).

Background: RFC 6487 section 4 unambiguously specifies: “*The CommonName attribute MUST be encoded using the ASN.1 type PrintableString*”, For many years, however, large CAs incorrectly issued certificates with `UTF8String` encoding owing to a documentation bug in [OpenSSL](#).

While the specific choice in RPKI standards for either `PrintableString` or `UTF8String` hardly matters in its context and would almost appear as a minor cosmetic detail, it is very important for both signers and validators to be on the same page. For optimal interoperability (now and in the future), the deployed software and published standards should be fully aligned in order to achieve byte-for-byte identical outputs for identical inputs in a diverse set of systems.

RIPE NCC RRDp performance issue in Bangalore

In the period August 6 to August 14, and September 7 to September 11, 2026, RIPE NCC's RRDp service in Bangalore, India, appeared to often serve HTTP 503 errors in response to requests for RRDp deltas. An RRDp webservice emitting HTTP errors will cause RPKI validators to retry by downloading a full snapshot, and if that also fails try RSYNC service.

Table 2 shows that RIPE NCC's full snapshot is substantial in size. The likely tangible impact reflected delays in ROA propagation and revocation from RIPE NCC members toward Indian ISPs.

TWNIC improves RRDP reliability by increasing retention

On August 14, 2026, a number of RPKIViews collectors observed synchronisation issues due to missing RRDP delta files on TWNIC's RRDP publication point. Collectors received HTTP 404 error responses. Typically this either means new notification files were published before the associated snapshot/delta files were also made publicly accessible or older snapshot and delta files were removed too soon.

After having been made aware of the issue, TWNIC staff promptly increased the retention time for older RRDP material in order to help clients avoid running into a time-of-check-time-of-use race condition.

IDNIC brief outage

On August 29, 2026, the IDNIC publication point was malfunctioning. A post hoc analysis of the CMS signing-time attribute contained in the newly discovered objects retrieved after the publication point became reachable suggests there was a seven hour delay for newly signed objects.

AfriNIC brief outage

On September 20, 2026 AfriNIC experienced an outage: a critical CRL was not resigned and distributed before the older version expired, causing all subordinate products to become invalid. [This CRL](#) covers the scope of the primary online CA for AfriNIC.

RIPE NCC implements policy to revoke non-functional CAs

On September 24, 2026, the RIPE NCC [shared](#) that they had successfully implemented the [RIPE-847](#) policy which provides RIPE NCC with mandate to revoke so-called “non-functional CAs”. A CA is considered non-functional when no valid and current manifest can somehow be retrieved and validated for an extended period of time.

Global RPKI One Million Payloads Milestone

On September 3, 2026, at 00:53:40 UTC, the global RPKI reached the formidable milestone of 1,000,000 ROA payloads.




Figure 2: Graph depicting the 1M moment

LACNIC introduces support for ASPA Objects

On September 30, 2026, LACNIC introduced support for managing ASPA objects in the *MiLACNIC* members portal (announcement [here](#)).

Four RIRs (APNIC, ARIN, LACNIC, and RIPE NCC) now support management of ASPA objects through their respective portals.

Standardization News

Stricter validation

Following from BER encoding no longer appearing in the ecosystem and the presence of `UTF8String` instances also being a thing of the past, a number of RPKI validator implementers have expressed an interest to release versions of their software packages with stricter object validation and standards conformance checks.

- [BER coding rejection](#)
- [UTF8String rejection](#)

RPKI Trust Anchor Constraints

The SIDROPS chairs [declared consensus](#) to progress a configuration format specification used to constrain RPKI Trust Anchors toward RFC publication. This is an important step forward for the initiative to better isolate Regional Internet Registries from each other and provide improved detection capabilities for “out-of-scope” certificate issuances.

RPKI testing fixture availability

A new vendor-agnostic “RPKI test fixture”, based on snapshots of the global RPKI, is now [available](#).

The fixture tarballs contain everything an RP needs to perform full chain validation: the TALs, TA certs, CA certs, signed objects, and a precomputed CCR to compare against. The fixture is primarily aimed at assisting in the development and maintenance of cryptographic libraries. Given the scope of the fixture, validation of the data set ensures a broad set of library functions and that code paths are exercised. The hope is that this will help detect code changes that would accidentally negatively impact RPKI applications (announcement [here](#)).

Letters to the editor

If you have any questions regarding this report or topics would like to be investigated for the next report, please let me know at job@bsd.nl!

Notice

This publication is distributed on an “AS-IS” basis, without warranty of any kind either express or implied, including but not limited to the implied warranties of merchantability, fitness for a particular purpose, or non-infringement.

This publication could contain technical inaccuracies or typographical errors. Later reports may modify or update information provided in this report. The author shall not have any liability for any loss or damage caused directly or indirectly by the information contained herein.