

# The *Erik Synchronization Protocol* for use with the RPKI

[draft-spaghetti-sidrops-rpki-erik-protocol](#)

Job Snijders, Tim Bruijnzeels, Tom Harrison, Wataru 'Alt' Ohgai  
[job@sobornost.net](mailto:job@sobornost.net), [tim@ripe.net](mailto:tim@ripe.net), [tomh@apnic.net](mailto:tomh@apnic.net), [alt@nic.ad.jp](mailto:alt@nic.ad.jp)

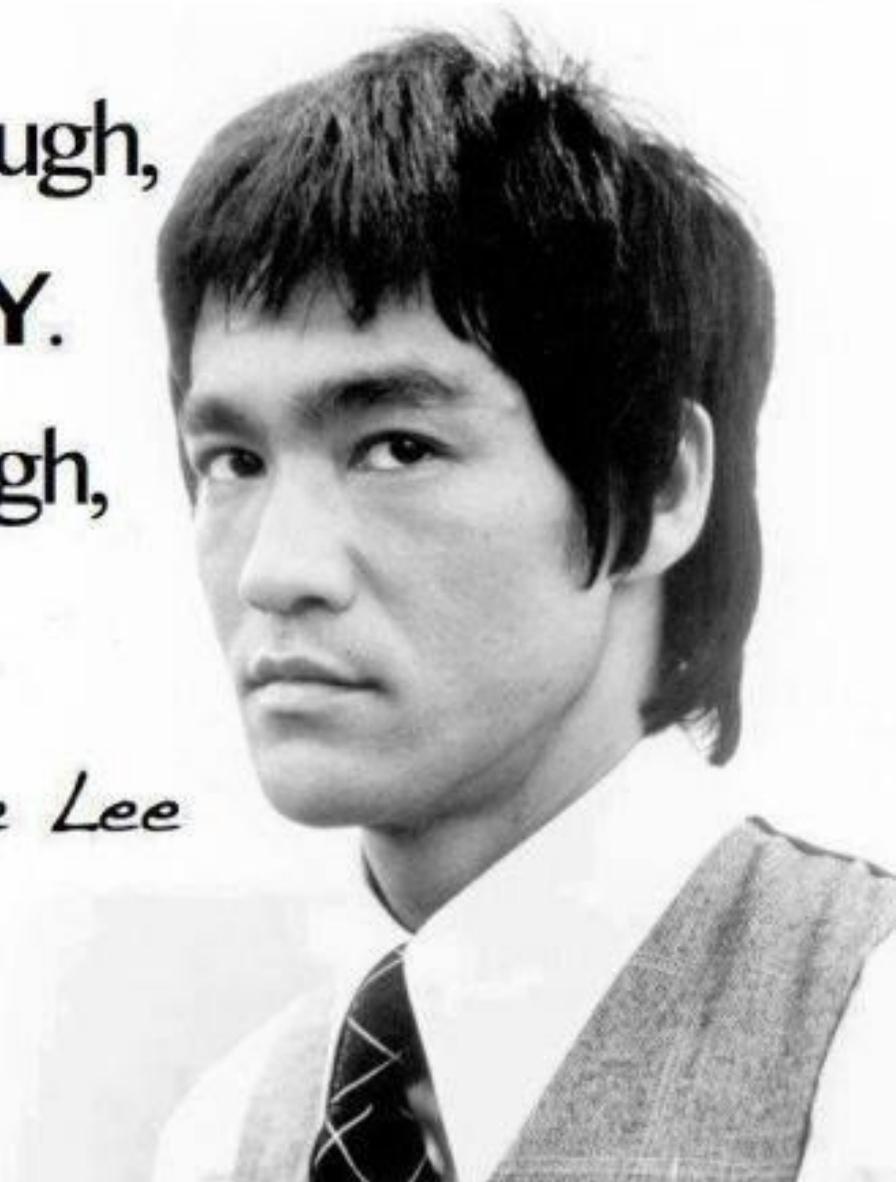
Knowing is not enough,

**We must APPLY.**

Willing is not enough,

**We must DO.**

*- Bruce Lee*



# Proportions of the RPKI data distribution problem

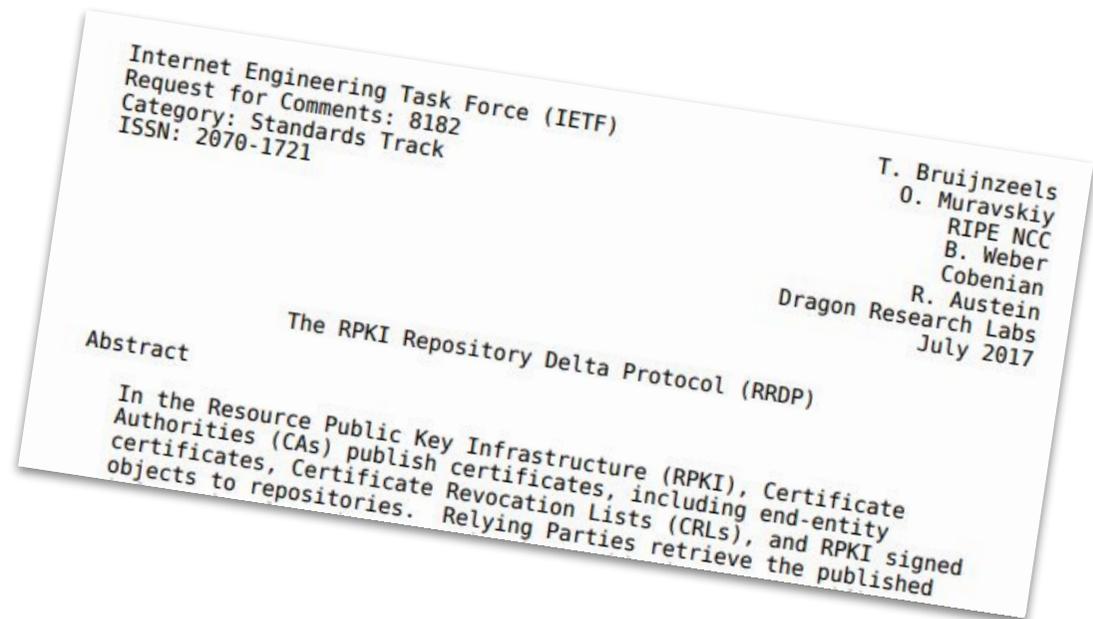
<b>Count:</b>	today, the RPKI is ~ <b>456,000</b> tiny objects (files)
<b>Churn:</b>	~ half those files change at least once a day
<b>Churn rate:</b>	~ 2 new objects issued every second
<b>Total size:</b>	~ 0.8 gibibytes
<b>Consumers:</b>	~ 5000 validators
<b>Trend:</b>	<i>up and to the right</i>

[https://labs.ripe.net/author/job\\_snijders/rpkis-2024-year-in-review/](https://labs.ripe.net/author/job_snijders/rpkis-2024-year-in-review/)

So, what's used to distribute RPKI data today?

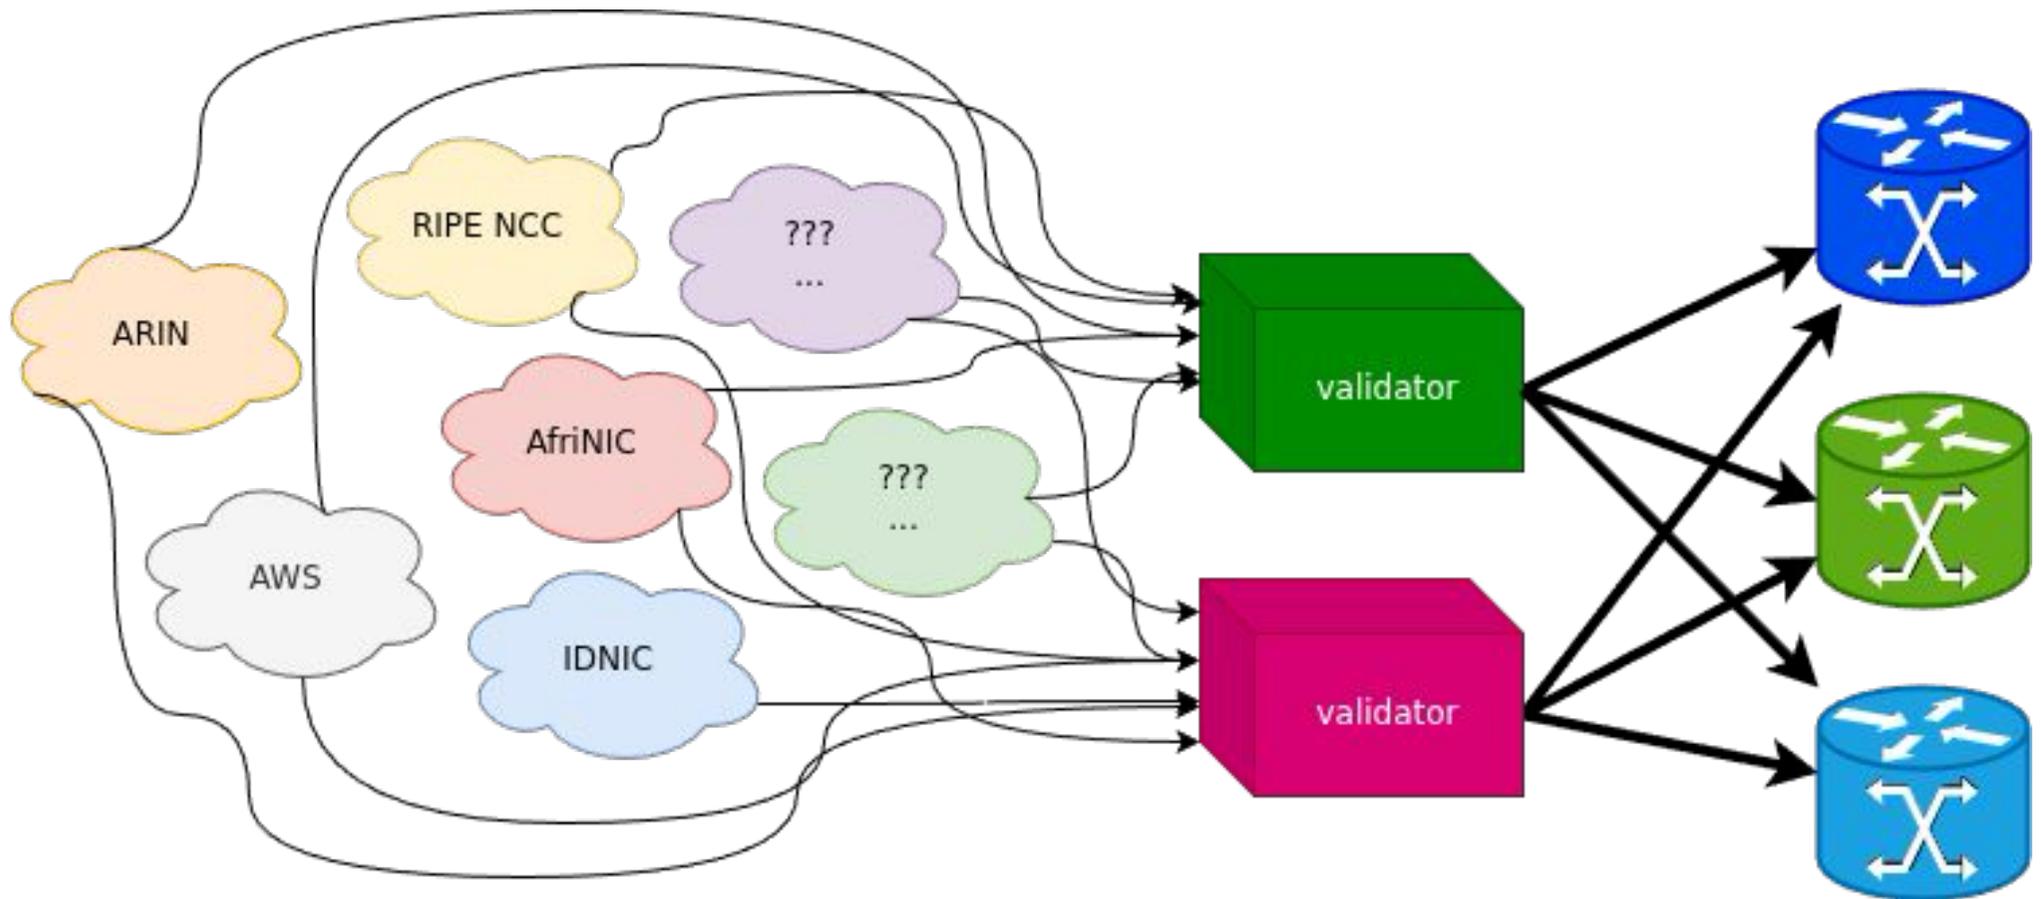


+



# High level overview of RPKI data supply chain

RPKI Repositories → *Rsync / RRDP* → Validators → BGP Routers



But... rsync is efficient, right?!

Yes and no!

Rsync is efficient because it transfers “only the difference”...

**But, calculating the difference on-the-fly, also consumes CPU & network resources!**

# Measuring Rsync

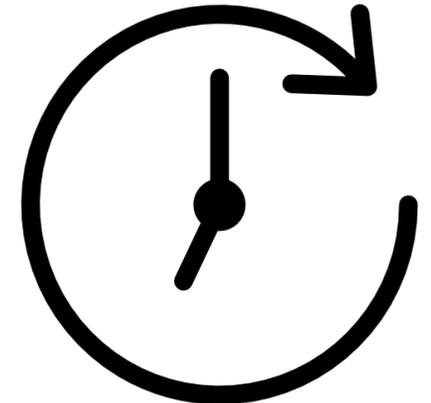
<b>Interval</b>	<b>Bandwidth consumption</b>
Minimum	~ 4 megabytes
Rsync every 15 minutes	~ 40 megabytes
Rsync every 60 minutes	~ 52 megabytes

# RRDP in a nutshell

- HTTP-based protocol, static pre-calculated content
- All add/update/delete operations are written into a “journal”
- RPs download the “journal” and replay it

***Consequently:***

*Fetches of data **already overtaken by events***



# Measuring RRDP

<b>Interval</b>	<b>Bandwidth consumption</b>
Minimum	~ 0.5 megabytes
RRDP every 15 minutes	~ 5 megabytes
RRDP every 60 minutes	~ 100 megabytes

# Comparing Rsync and RRDP

Interval	Rsync	RRDP
Minimum	~ 4 megabytes	~ 0.5 MB
Every 15 minutes	~ 40 megabytes	~ 5 MB
Every 60 minutes	~ 50 megabytes	~ 100 MB

# Rsync / RRDP intrinsic design issues

- Rsync servers vulnerable to easy denial of service (CPU hogging)
- Rsync expensive traversal of resources (no “if-modified-since”)
- RRDP servers can easily DoS their clients (disk space hogging)
- RRDP “loss of state” → reinitialize through snapshot download
- In RRDP FSM, many error paths lead to “snapshot download”

*Neither Rsync nor RRDP are really optimal*

# What *is* the Erik\* Synchronization Protocol?

A data replication system using the following components:

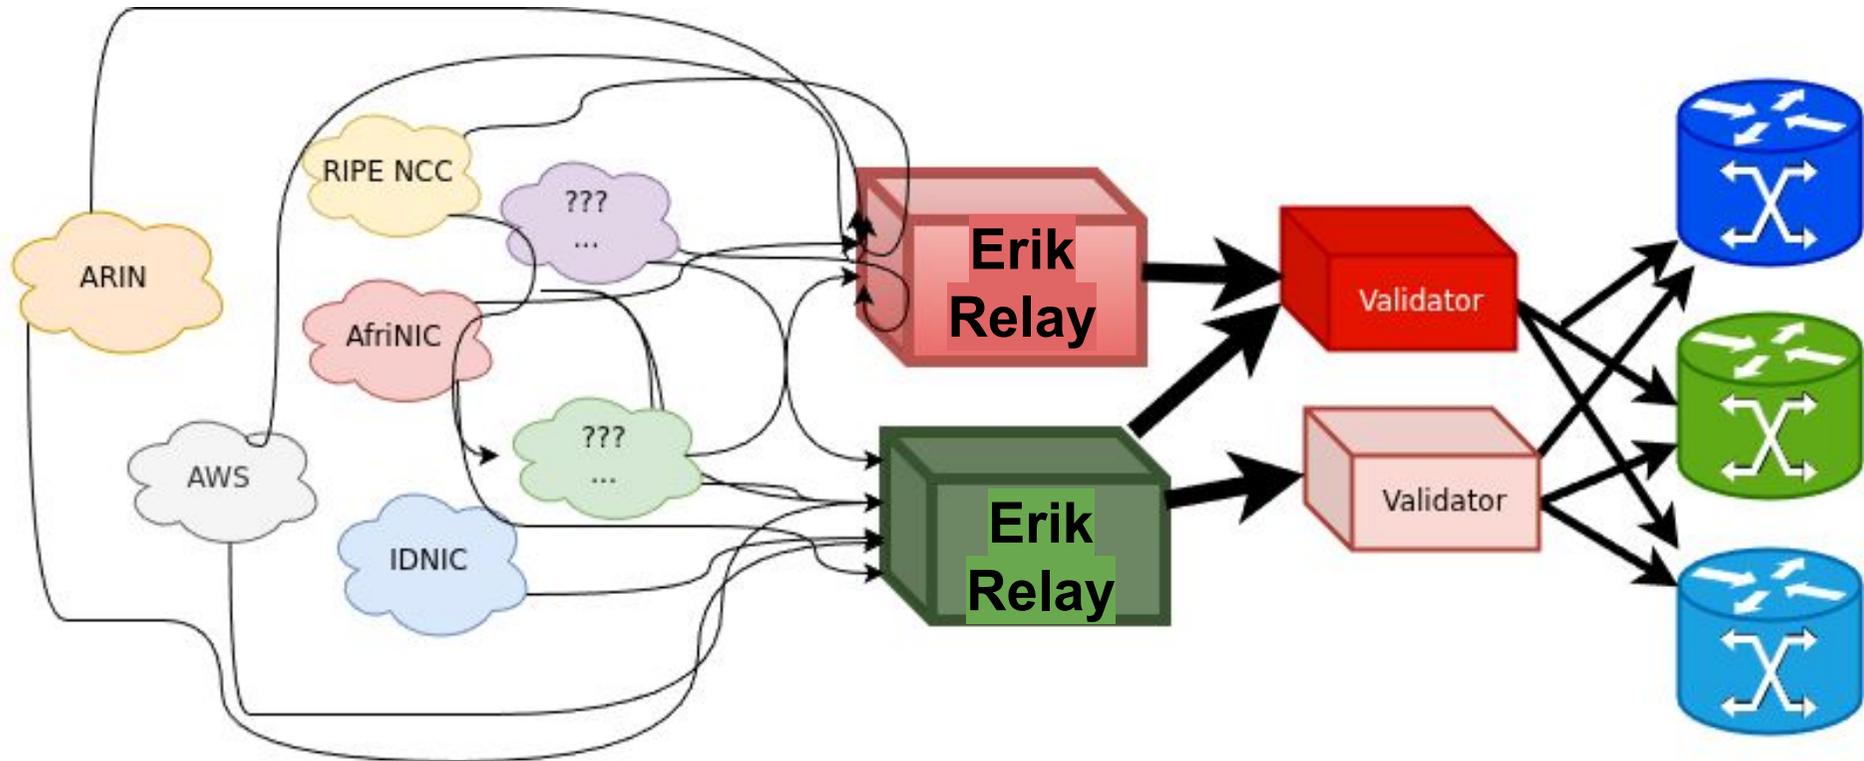
- Merkle Trees (1979)
- Content-addressable naming scheme (1950s?)
- Timestamp-based concurrency control (1975)
- HTTP transport (1989)

\*Named in honor of [Erik Bais](#) who passed away in 2024.



# Possible Erik Protocol deployment strategy

Repositories → *Rsync / RRDP* → **Erik** relay → Validators → BGP Routers



# Advantages of the Erik Synchronization Protocol

- Enables clients to “jump straight to latest” (like Rsync)
- Mirroring / replication / offloading is a supported use-case
- Fetch “only what changed” (like Rsync)
- Static pre-calculated content (like RRDP)
- HTTP-based (like RRDP)
- Light on state, no “session” (like Rsync, unlike RRDP)
- Fallback from Erik to Rsync and back is easy

FAST – EFFICIENT – CHEAP

***Erik relays are accelerators!***

# Comparing Rsync and RRDP and Erik

Interval	Rsync	RRDP	Erik
Minimum	~ 4 megabytes	~ 0.5 MB	~ 0.5 MB
Every 15 minutes	~ 40 megabytes	~ 5 MB	~ 4 MB
Every 60 minutes	~ 50 megabytes	~ 100 MB	~ 20 MB

# The plan

- Write IETF draft
- Write software
- Measure
- Experiment
- Optimize
- Iterate
- Deploy

