RPKI Retrieval Delta Protocol - RRDP

Tim Bruijnzeels
Oleg Muravskiy
Bryan Weber
Rob Austein
David Mandelberg

IETF91 | Honolulu | 10 November 2014

Rsync discussions

- WG has spent a lot time discussing
- This is a proposal for something that we think would be a good replacement
- Based on now outdated draft
- Doing this as proof of concept first has been a great help in fixing some of the protocol details

The Concept - Notification File

- Publication Servers publish a notification file with
  - session id
  - serial number
  - a reference to the latest snapshot
  - X references to the latest deltas

  This file is small and cheap to fetch

- Relying Parties fetch the notification file and can decide to:
  - Do nothing (in sync)
  - Get the snapshot
  - Get one or more deltas

The Concept - Snapshot File

- A snapshot file contains **all** objects in the repository at some point in time

  - It is unique and immutable for any given session id and serial And the reference url to it is unique as well..
  - So it can be cached aggressively by Content Delivery Networks (CDNs)
  - Format based on the publication protocol draft
  - So it is easy for publication servers to include the publish messages in these files

The Concept - Delta File

- A delta file contains **all** objects added or removed at some unique point in time

  - It is unique and immutable for any given session id and serial And the reference url to it is unique as well..
  - So it can be cached aggressively by Content Delivery Networks (CDNs)
  - Format based on the publication protocol draft
  - So it is easy for publication servers to include the publish messages in these files

How does the RP learn about this?

- SIA pointer with new OID
  - Different semantics: this is not a unique publication point for this certificate. It includes other products

- Points to notification.xml using http

- RP can learn about new notification locations when validating, and can pro-actively recheck these locations for changes from then on

Typical dialogue - RP has no data

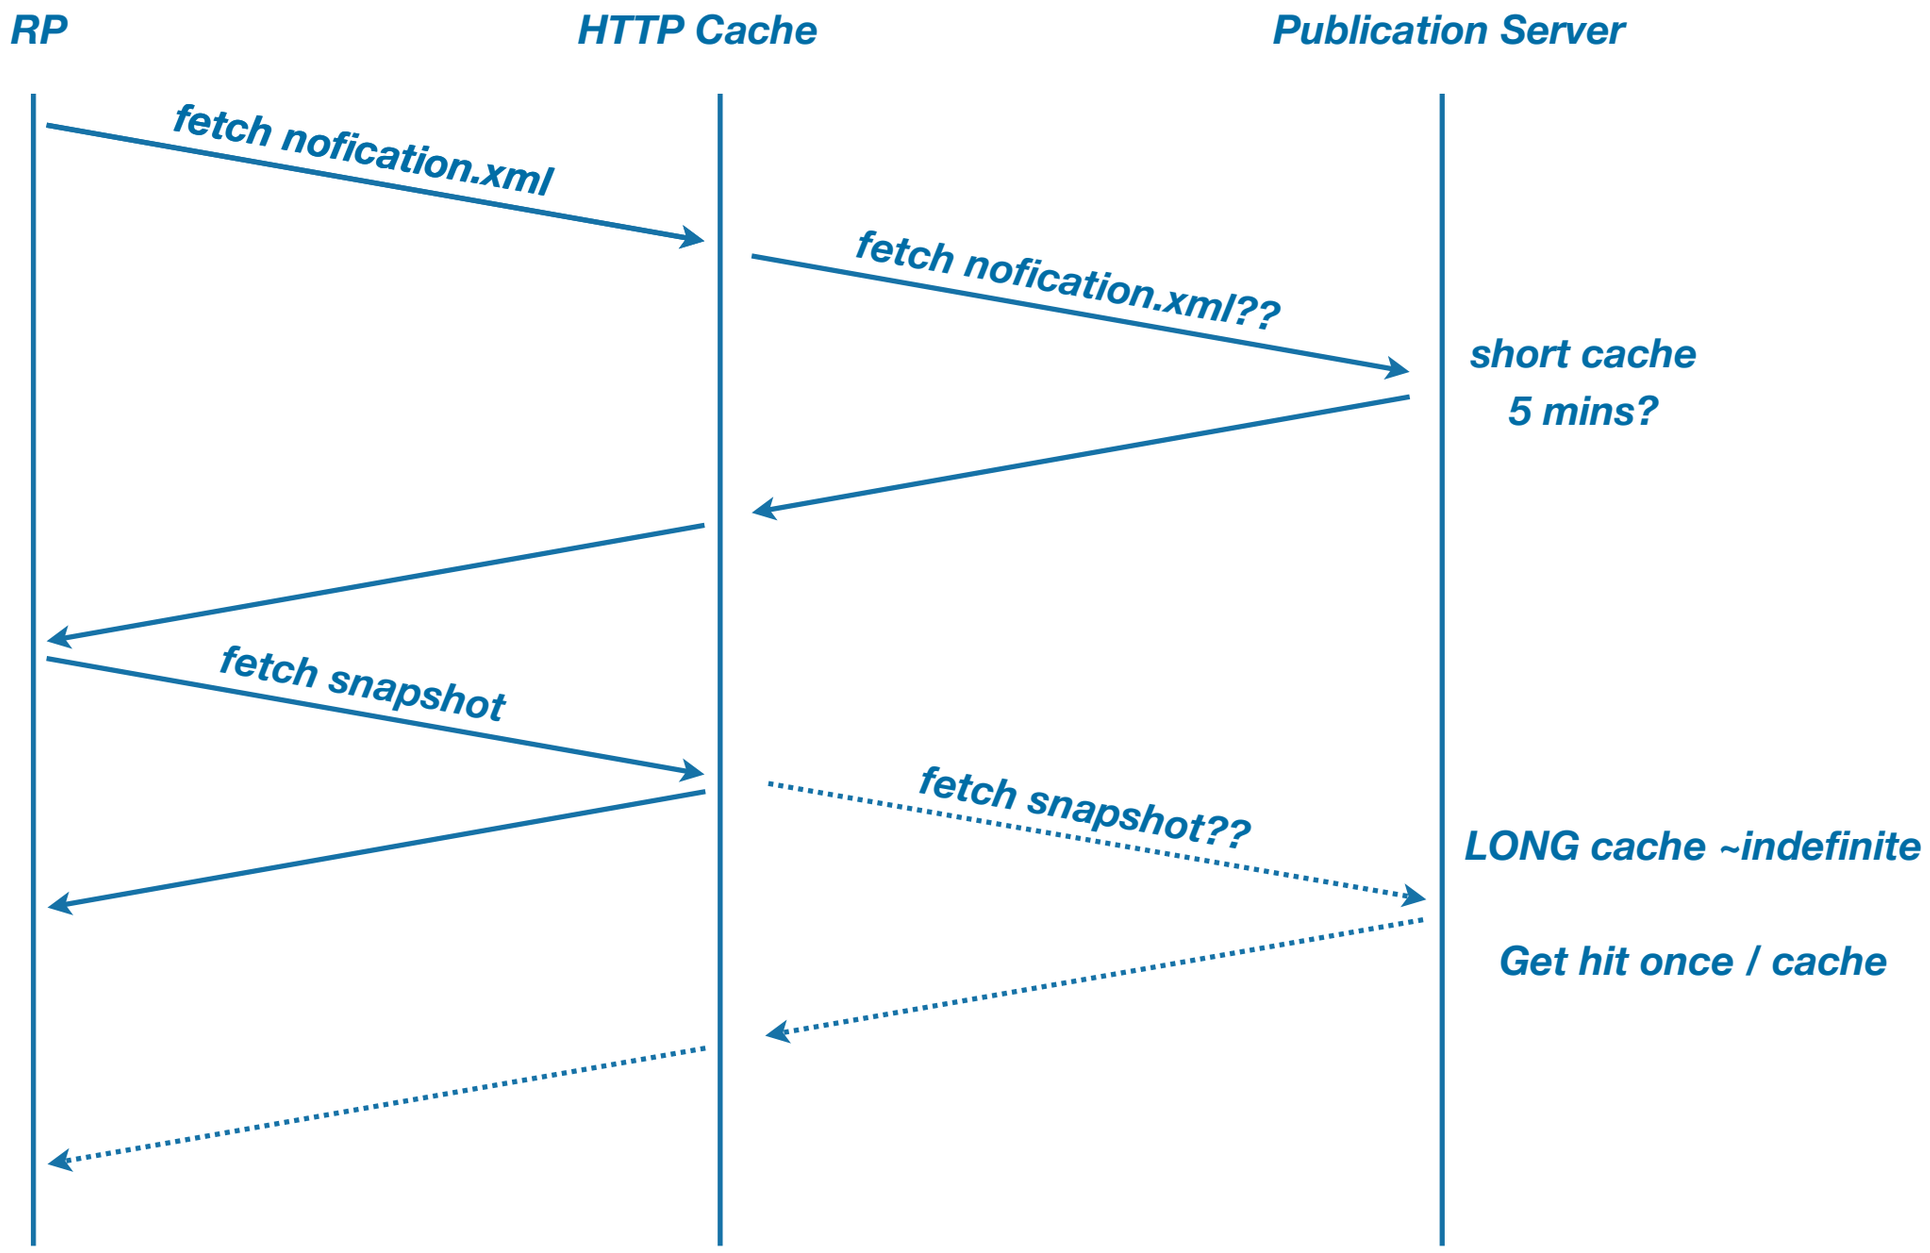

Typical dialogue - Server has updates

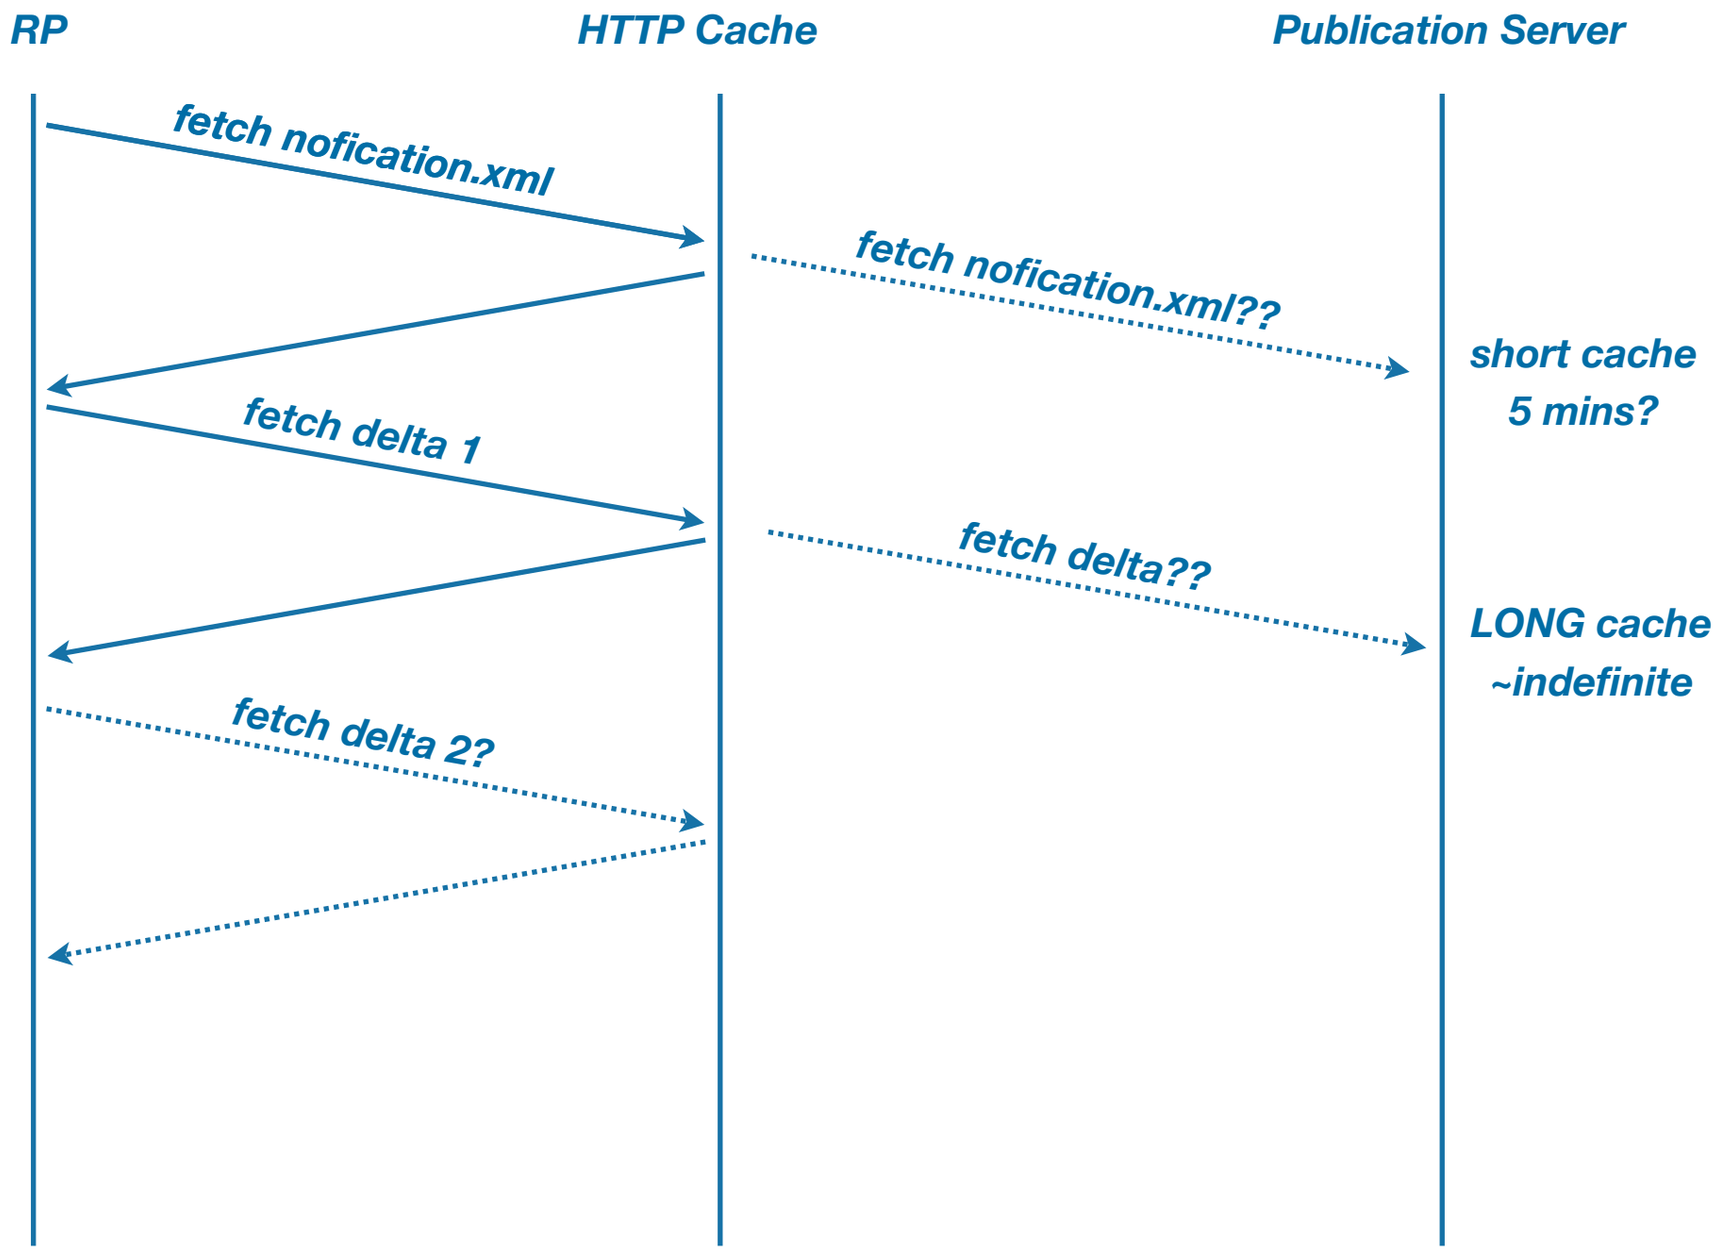

Typical dialogue - Server has no updates

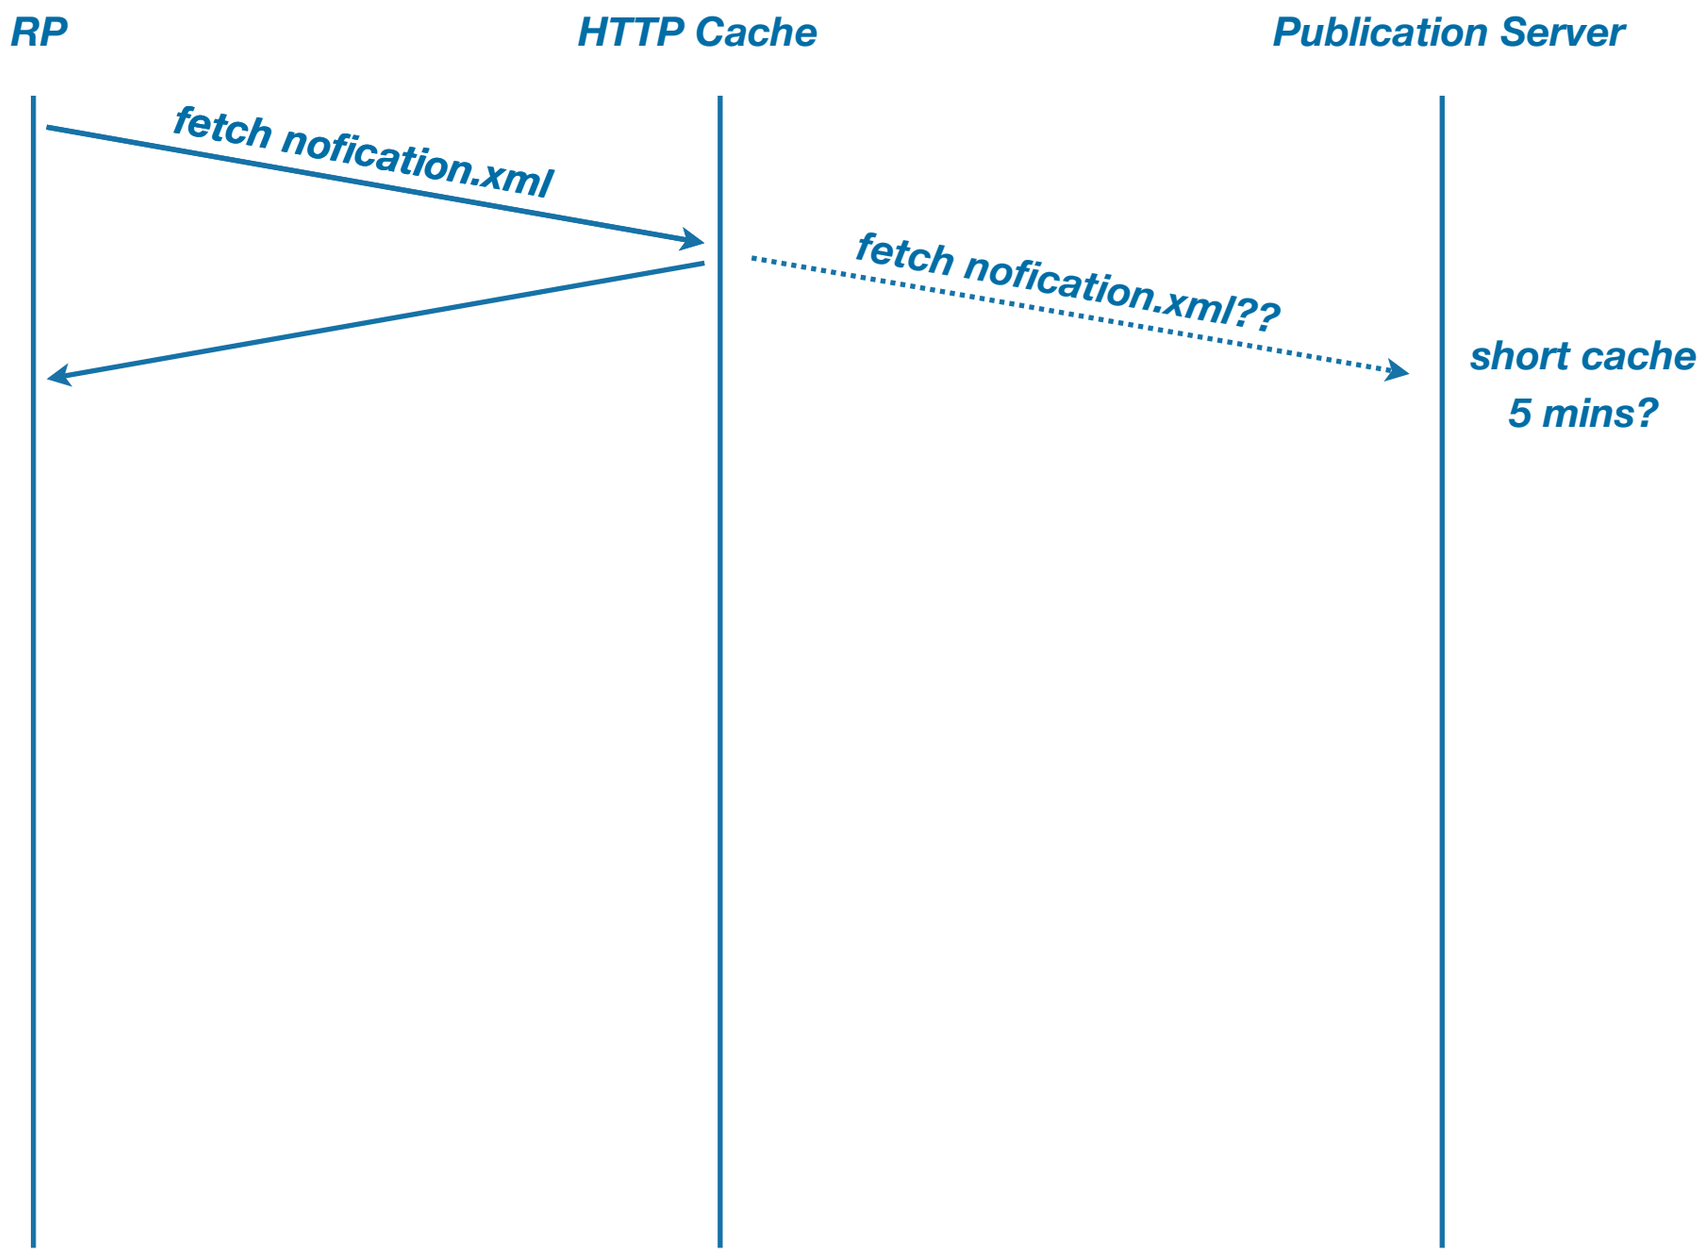

What about path discovery?

- publish / withdraw messages include rsync URI
- but one of the current proof of concept code bases relies on key identifiers, proof of possession for MFT, and hashes..

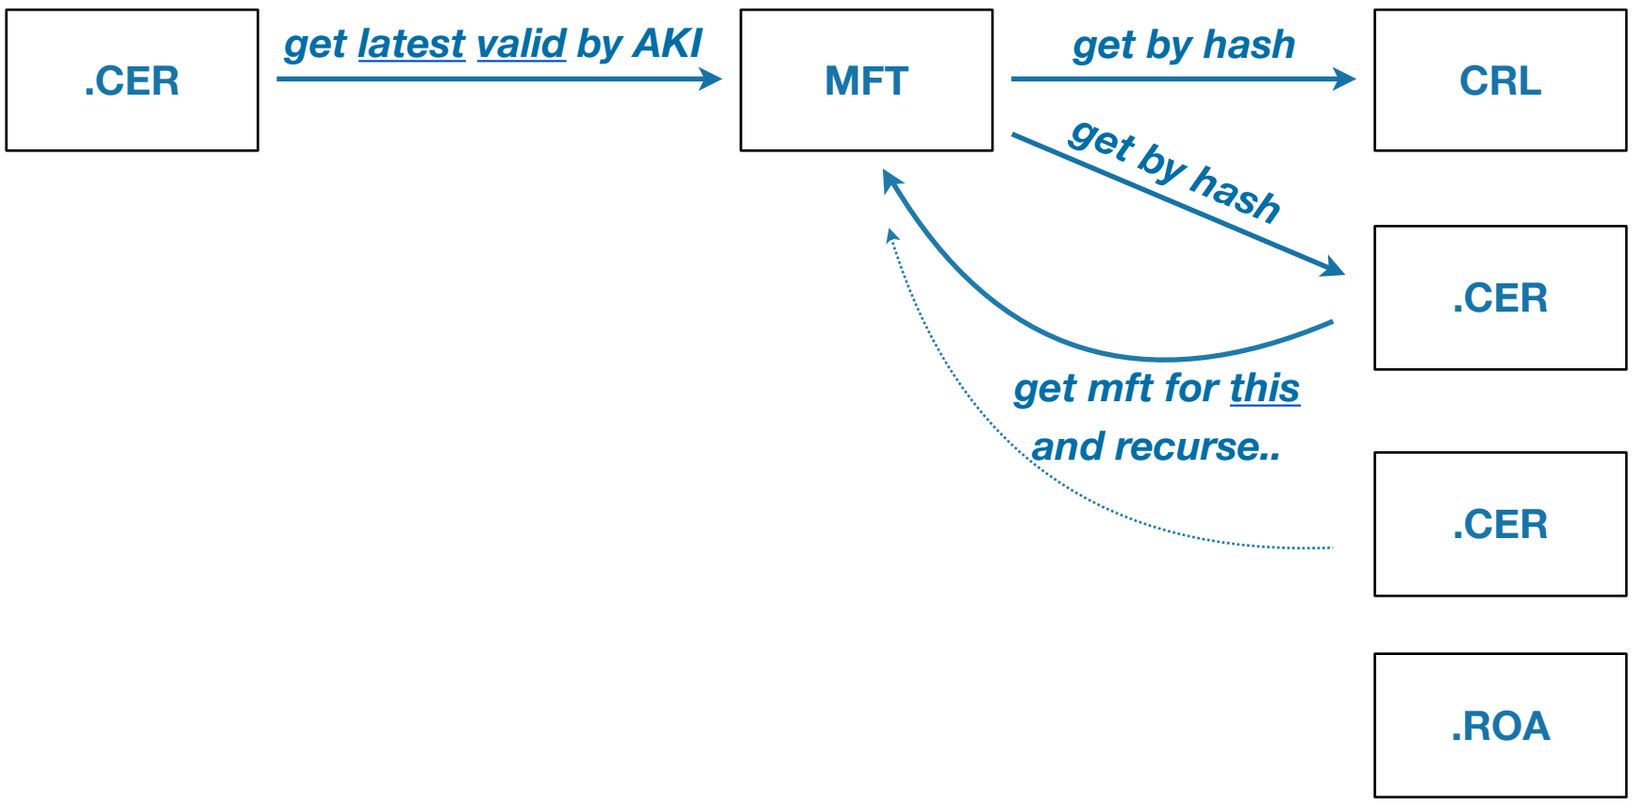

Next steps?

- Proof of concept code works
- Implement production grade code in publication server and RP software
- Add pointers to real production repository so that we can do real world measurements (help desired)
- Update draft to reflect changes and share with working group
- If you want to be involved sooner.. speak up, email, chat..

Questions?