

# RPKI Repository Analysis and Delta Protocol

---

Oleg Muravskiy  
SIDR, IETF86

# New Internet-Drafts

---

- RPKI Repository Analysis and Requirements  
[draft-tbruijnzeels-sidr-repo-analysis-00](https://datatracker.ietf.org/doc/draft-tbruijnzeels-sidr-repo-analysis-00)
- RPKI Repository Delta Protocol  
[draft-tbruijnzeels-sidr-delta-protocol-01](https://datatracker.ietf.org/doc/draft-tbruijnzeels-sidr-delta-protocol-01)

Tim Bruijnzeels  
[tim@ripe.net](mailto:tim@ripe.net)

Oleg Muravskiy  
[oleg@ripe.net](mailto:oleg@ripe.net)

Bryan Weber  
[bryan@cobenian.com](mailto:bryan@cobenian.com)

# RPKI Repository Analysis and Requirements

---

- Scalability of rsync server
- Lack of rsync standard and implementations
- Lack of consistent view of repository
- New requirements

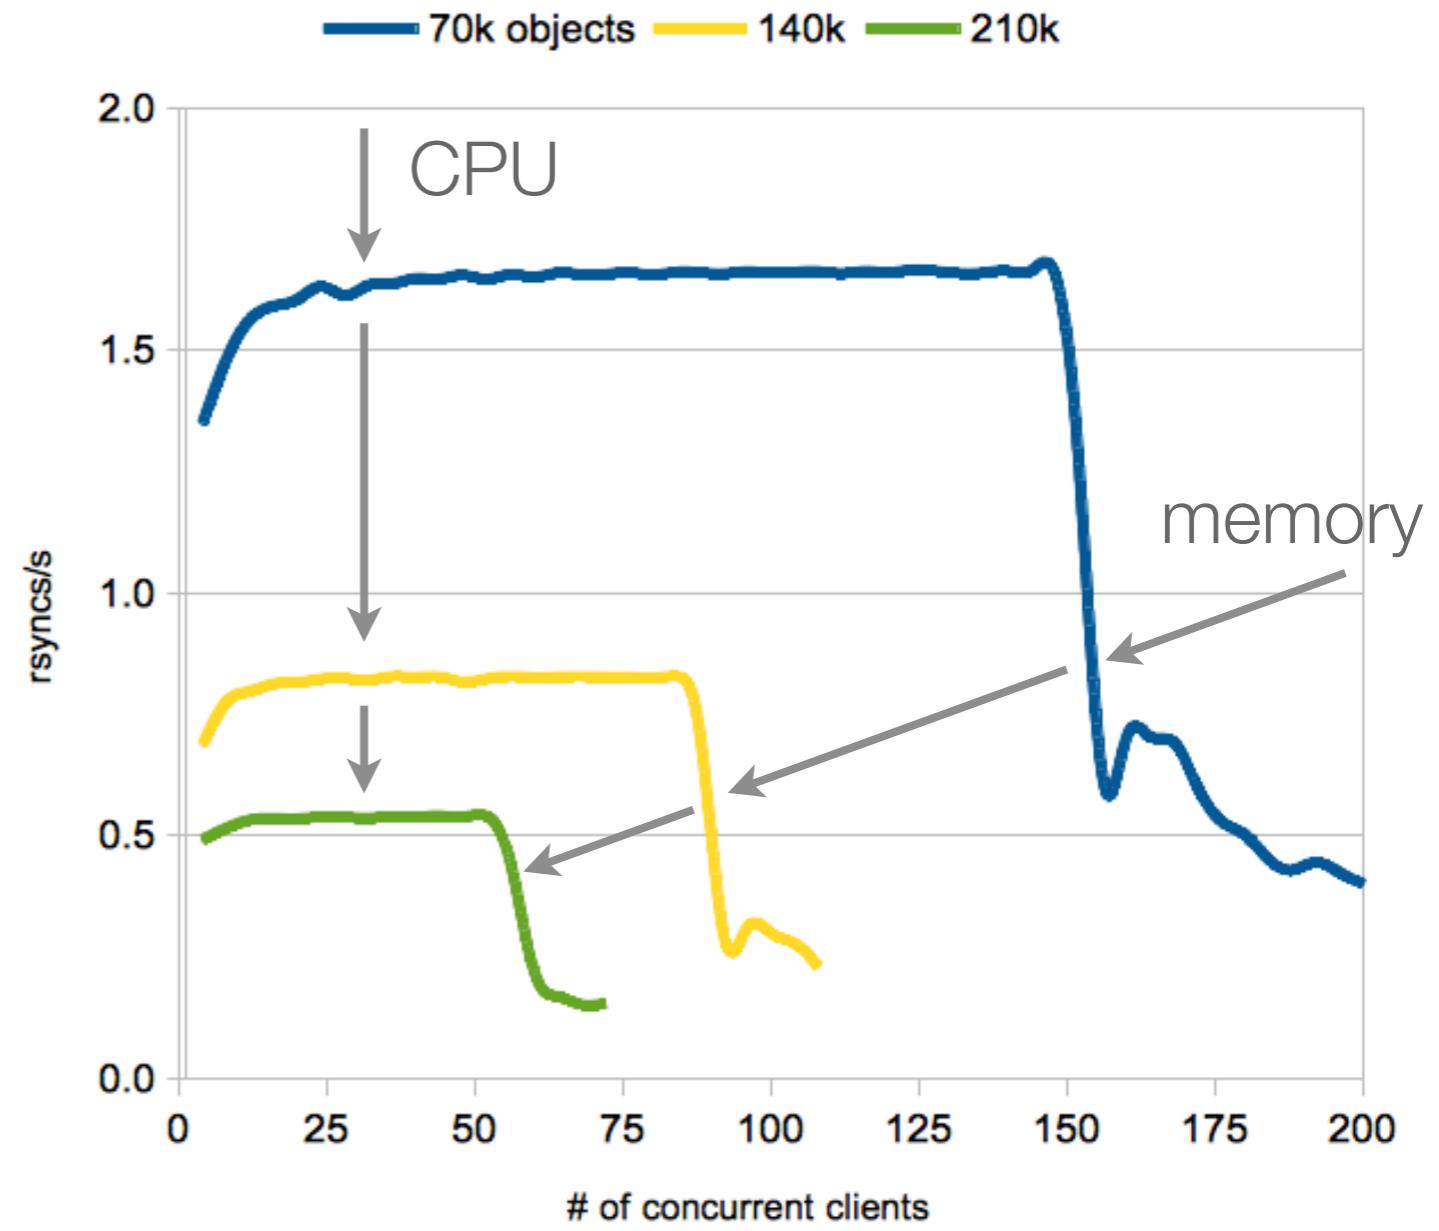
# Scalability of rsync server

- 5 Mac Minis
- 2 GB Memory
- Core2 CPU @2GHz
- 2 cores, no hyperthreading
- 1 repository server
- 4 clients
- Ubuntu 12.04, 3.2.0 kernel



# Scalability of rsync server (2)

- recursive fetch
- client is up to date
- no latency (local network)
- so just the rsync overhead...
  
- rsync is CPU and memory bound
  - CPU: 70% system, 30% user
  - needed cache size rel. to repo size
  - rsync mem depends on number of clients (forks) and repo size



# Lack of rsync standard and implementations

---

- no standard
- single implementation
- no libraries

# Lack of consistent view of repository

---

- rsync serves file system, which is updated at the same time
- manifest vs. content of directory
- RFC6486 leaves it to local policies

# New requirements

---

- Transport agnostic
- Consistent view on repository
- No specific repository layout
- Expected size, number of clients, update rate

# Expected size, number of clients, update rate

- Repository size
    - ~3 million objects
    - ~5 GB
  - Churn per day (baseline)
    - ~1.2 million objects
    - ~1.6 GB
  - $\times 8k$  clients = 12.5 TB traffic
- Relying Parties
- ~40k ASNs  $\times 2(?)$  instances
  - Fetch every 10(?) minutes
  - ...
- 8k conn/min \* 0.5 min/conn  
4k concurrent clients

# RPKI Repository Delta Protocol (RRDP)

---

- Transport-agnostic (HTTP in mind)
  - could use CDN
  - no specific repository layout
- Update notification
  - clients do the math
- Consistent view
  - snapshot and version
  - deltas
  - all immutable

# Update Notification File

---

- session-id
- current version number
- snapshot
  - version
  - list of snapshot segments
    - URI
    - hash
- list of deltas
  - from version
  - to version
  - list of delta segments
    - URI
    - hash

# Sample notification file

```
<?xml version="1.0" encoding="us-ascii"?>
<msg type="notification" version="1" xmlns="http://www.ripe.net/rpki/rrdp">

  <notification session_id="d9f6dc91-0394-40b9-9663-66aef4bb623a"
    current_version="203">

    <snapshot version="202">
      <snapshot-segment
        uri="http://host/d9f6dc91-0394-40b9-9663-66aeb623a/snapshot/202/1.xml"
        hash="279b79fd8389e20585f26735ee70e0e4d4b8af23bb2e2e611c70e92d2433e"/>
      . . .
    </snapshot>

    <deltas>
      <delta-segment from="200" to="201"
        uri="http://host/d9f6c91-0394-40b9-9663-66aeb623a/deltas/200/201.xml"
        hash="a2d56ec180f2dde2a46b2e0565932e25829b852a0b43107d5de6e41394c292"/>
      <delta-segment from="201" to="202"
        uri="http://host/d9f6c91-0394-40b9-9663-66aeb623a/deltas/201/202.xml"
        hash="a2d56ec180f2dde2a46b2e0565932e25829b852a0b43107d5de6e41394c292"/>
      <delta-segment from="202" to="203"
        uri="http://host/d9f6c91-0394-40b9-9663-66aeb623a/deltas/202/203.xml"
        hash="a2d56ec180f2dde2a4f92e0565932e25829b852a0b43107d5de6e41394c293"/>
    </deltas>
  </notification>
</msg>
```

# Snapshot segment

- session-id
- version
- segment number
- publish element  
(from draft-ietf-sidr-publication)

```
<?xml version="1.0" encoding="us-ascii"?>
<msg xmlns="http://www.ripe.net/rpki/rrdp"
      type="snapshot" version="1">

  <snapshot session_id="d9f6dc91-0394-40b9-9663-66aef4bb623a"
            repository_version="1" index="2">

    <publish uri="http://host/foo/bar/cer1.cer">
      MIIE+jCCA+KgAwIBAgIBDTANBgkqhkiG9w0BAQsFADAzMTEwLwYDVQQD
      jRBODAxN0U2NkE5RTkxNzJFNDYxMkQ4Q0Y0QzgzRjIzOERFMkEzMB4XE
      .....
      h8zeHoFVu6ghRPy5db0A4akX/KG6b8XIx0iwPYdLiDbdWFbtTdPcXBau
    </publish>

    <publish uri="http://host/foo/bar/cer2.cer">
      MIIE+jCCA+KgAwIBAgIBDTANBgkqhkiG9w0BAQsFADAzMTEwLwYDVQQD
      h8zeHoFVu6ghRPy5db0A4akX/KG6b8XIx0iwPYdLiDbdWFbtTdPcXBau
      .....
      jRBODAxN0U2NkE5RTkxNzJFNDYxMkQ4Q0Y0QzgzRjIzOERFMkEzMB4XD
    </publish>

    <publish uri="http://host/foo/bar/cer3.cer">
      MIIE+jCCA+KgAwIBAgIBDTANBgkqhkiG9w0BAQsFADAzMTEwLwYDVQQD
      h8zeHoFVu6ghRPy5db0A4akX/KG6b8XIx0iwPYdLiDbdWFbtTdPcXBau
      .....
      jRBODAxN0U2NkE5RTkxNzJFNDYxMkQ4Q0Y0QzgzRjIzOERFMkEzMB4XD
    </publish>
  </snapshot>
</msg>
```

# Delta segment

- session-id
  - from version
  - to version
  - segment number
  - delta
    - publish element
    - withdraw element
- (from draft-ietf-sidr-publication)

```
<?xml version="1.0" encoding="us-ascii"?>
<msg type="deltas" version="1"
      xmlns="http://www.ripe.net/rpki/rrdp">

<deltas session_id="d9f6dc91-0394-40b9-9663-66aef4bb623a"
         from="0" to="3" index="2">

  <delta version="2">
    <withdraw uri="http://host/foo/bar/cer1.cer" />
    <publish uri="http://host/foo/bar/cer2.cer">
      MIIE+jCCA+KgAwIBAgIBDTANBgkqhkiG9w0BAQsFADAzMTEw
      .....
      jRBODAxN0U2NkE5RTkxNzJFNDYxMkQ4Q0Y0QzgzRjIzOERFM
    </publish>
    <publish uri="http://host/foo/bar/cer3.cer">
      MIIE+jCCA+KgAwIBAgIBDTANBgkqhkiG9w0BAQsFADAzMTEw
      .....
      jRBODAxN0U2NkE5RTkxNzJFNDYxMkQ4Q0Y0QzgzRjIzOERFM
    </publish>
  </delta>

  <delta version="3">
    <withdraw uri="http://host/foo/bar/cer2.cer" />
  </delta>
</deltas>
</msg>
```

# Questions?



**RIPE**  
**NCC**