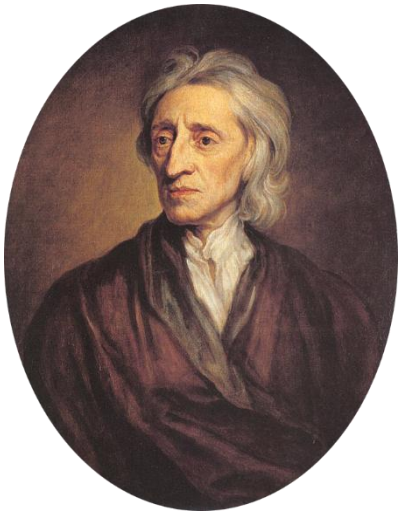

From the Consent of the Routed:

Improving the Transparency of the RPKI.

Ethan Heilman

Danny Cooper Leonid Reyzin Sharon Goldberg

Boston University

Aug 2014

BOSTON UNIVERSITY

Overview

Motivation: The RPKI* (2011 to present) secures interdomain routing,
… but creates a new danger of misbehaving authorities.

Drop RPKI invalid routes?	Route is reachable during …	
	BGP attack	RPKI misbehavior
Yes	✓	X RPKI
No	X	✓

We propose changes to the RPKI to detect misbehavior.

- We have a window of opportunity to influence RPKI design.
- Changes being still being made to RPKI specification.
- Concurrent to our work, IETF is drafting misbehavior defenses

* RPKI = Resource Public Key Infrastructure **[RFC 6480]**

Outline

1. **Background.**
   1. Interdomain routing is not secure: BGP Prefix hijacks.
   2. How the RPKI is designed to prevent these attacks.
   3. Misbehaving RPKI authorities and takedowns.
2. **Our proposed changes.**

The RPKI is designed to prevent prefix hijacks.

The Telegraph

Home | News | World | Sport | Finance | Comment | Culture | Travel | Li

Politics | Obits | Education | Earth | Science | Defence | Health | Scotland

HOME » NEWS » UK NEWS

Pakistan b

The site was banned

By Bonnie Malki

12:48PM GMT 25 F

TV: CNNUS | CNNI | CNN

Home | TV & Video

Report: Ch

From Dugald McConnell, C

November 18, 2010 3:19 a.m. E

renesys

Products | Solutions | Company | Contact Us | Blog

Con-Ed Steals the 'Net

22 JAN, 2006 | 11:06 PM | BY TODD UNDERWOOD

Indonesia Hijacks the World

03 APR, 2014 | 3:09 PM | BY EARL ZMIJEWSKI

Yesterday, Indosat, one of Indonesia's largest telecommunications providers, leaked large portions of the global routing table multiple times over a two-hour period. This means that, in effect, Indosat claimed that it "owned" many of the world's networks. Once someone makes such an assertion, typically via an honest mistake in their routing policy, the only question remaining is how much of the world ends up believing them and hence, what will be the scale of the damage they inflict? Events of this nature, while relatively rare, are certainly not unheard of and can have geopolitical implications, such as when China was involved in a similar incident in 2010.

ECURITY | DESIGN | OPIN

Share 90

Tweet 452

g+1 172

Share

Pin it

The Indosat prefix hijack incident from 03/04/2014

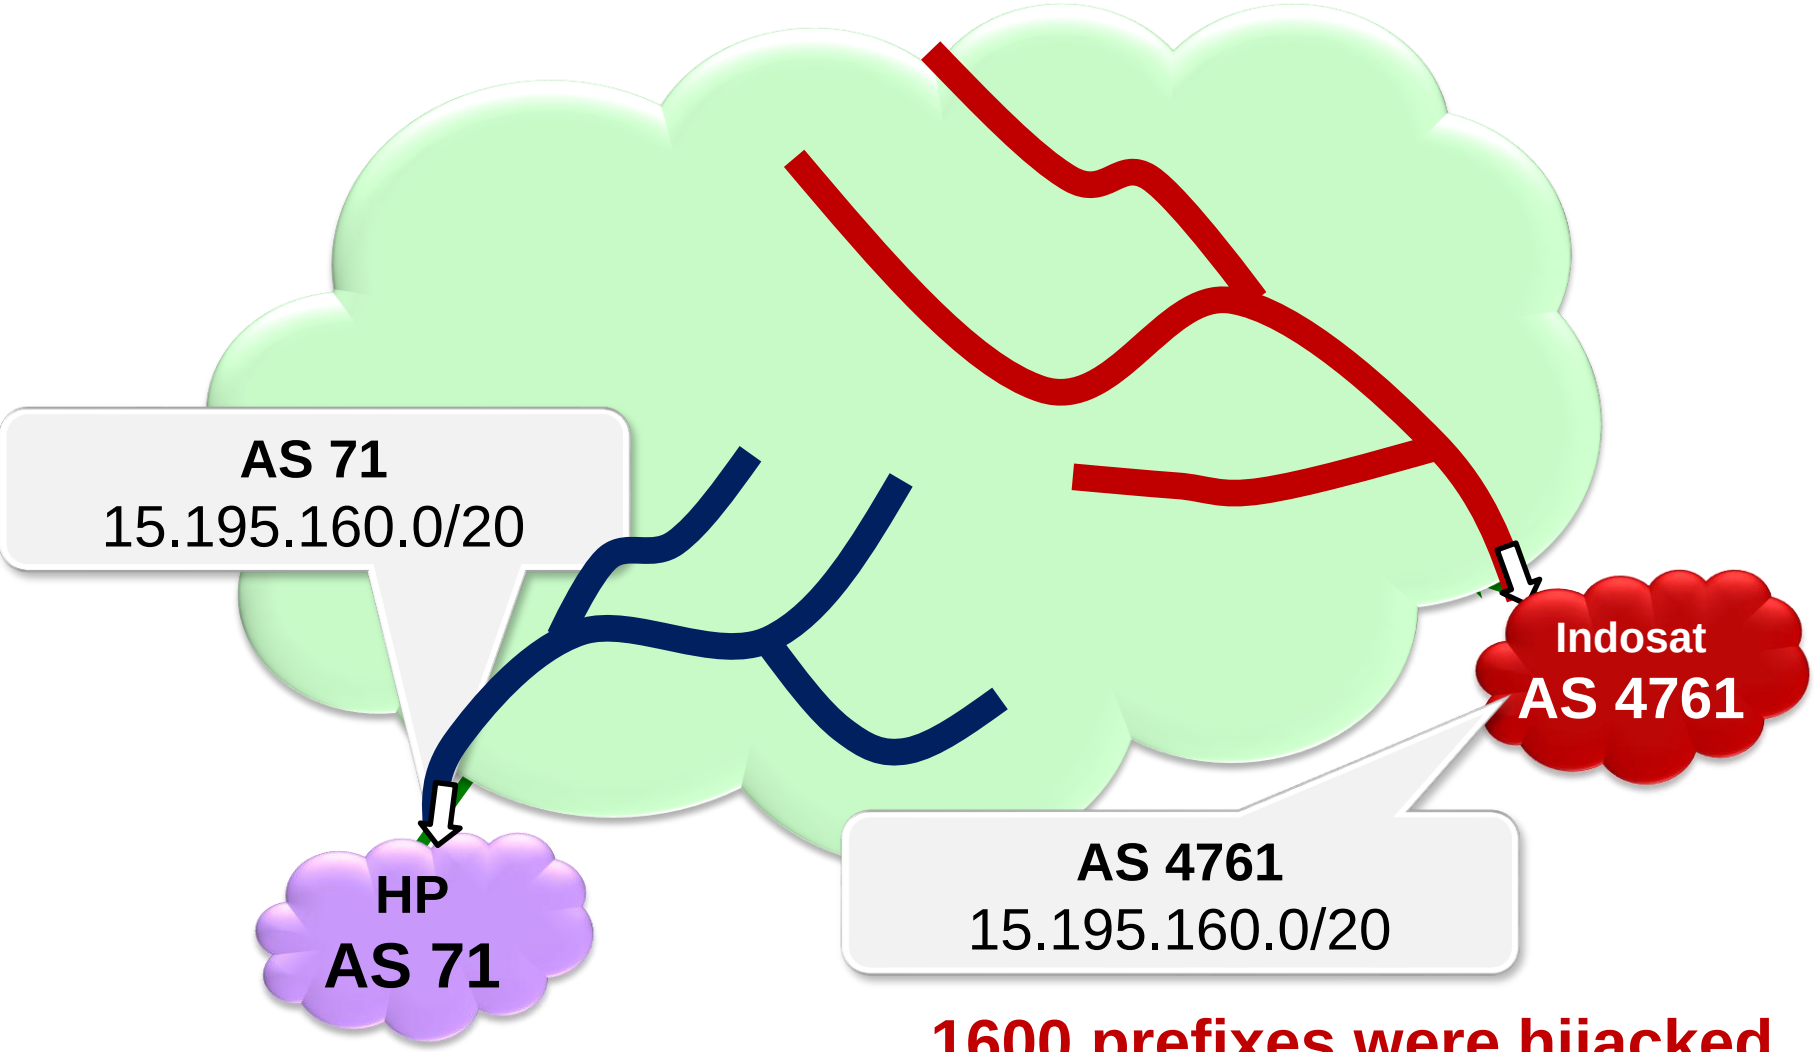

1600 prefixes were hijacked.

Source: http://portal.bgpmon.net/data/indosat-us.txt

What is the fundamental vulnerability?

Problem: Route origin announcements are not authenticated.

Solution: The RPKI authenticates route origins.

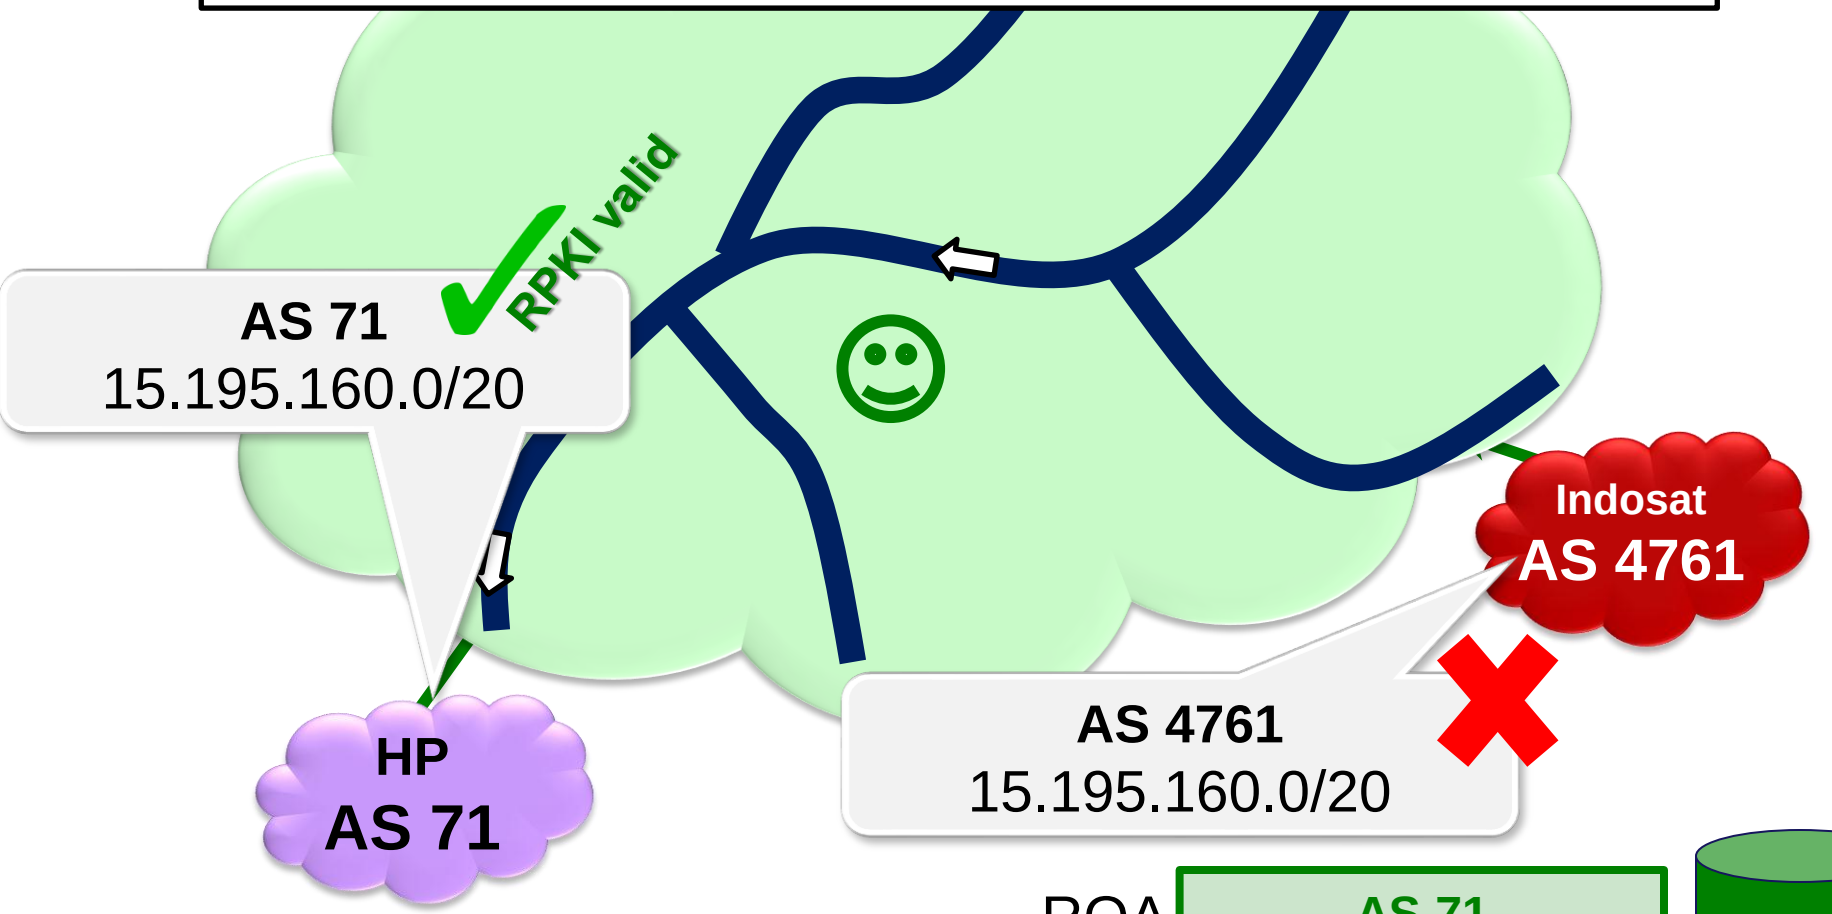

ROA
(Route Origin Authorization)

AS 71
15.195.160.0/20

RPKI

The structure of the RPKI

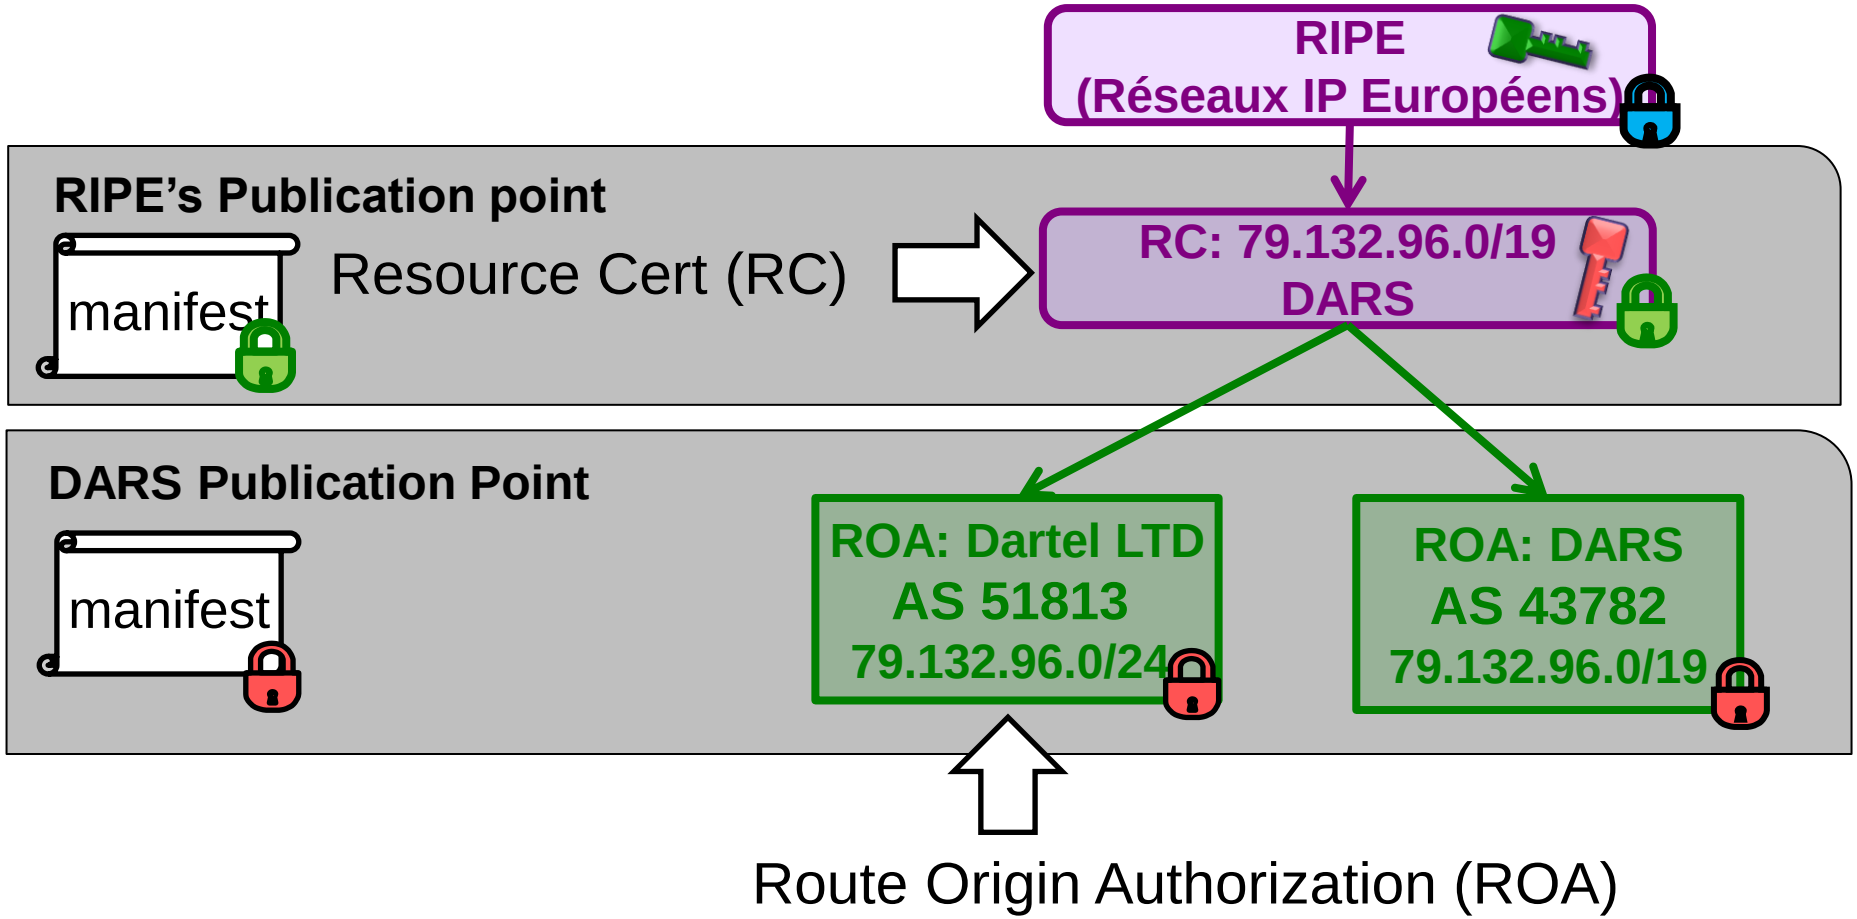

Deployment Status of the RPKI:

- Today: ROAs cover about 4% of interdomain routes.
- Goal: Cover all routes!

How relying parties sync to the RPKI

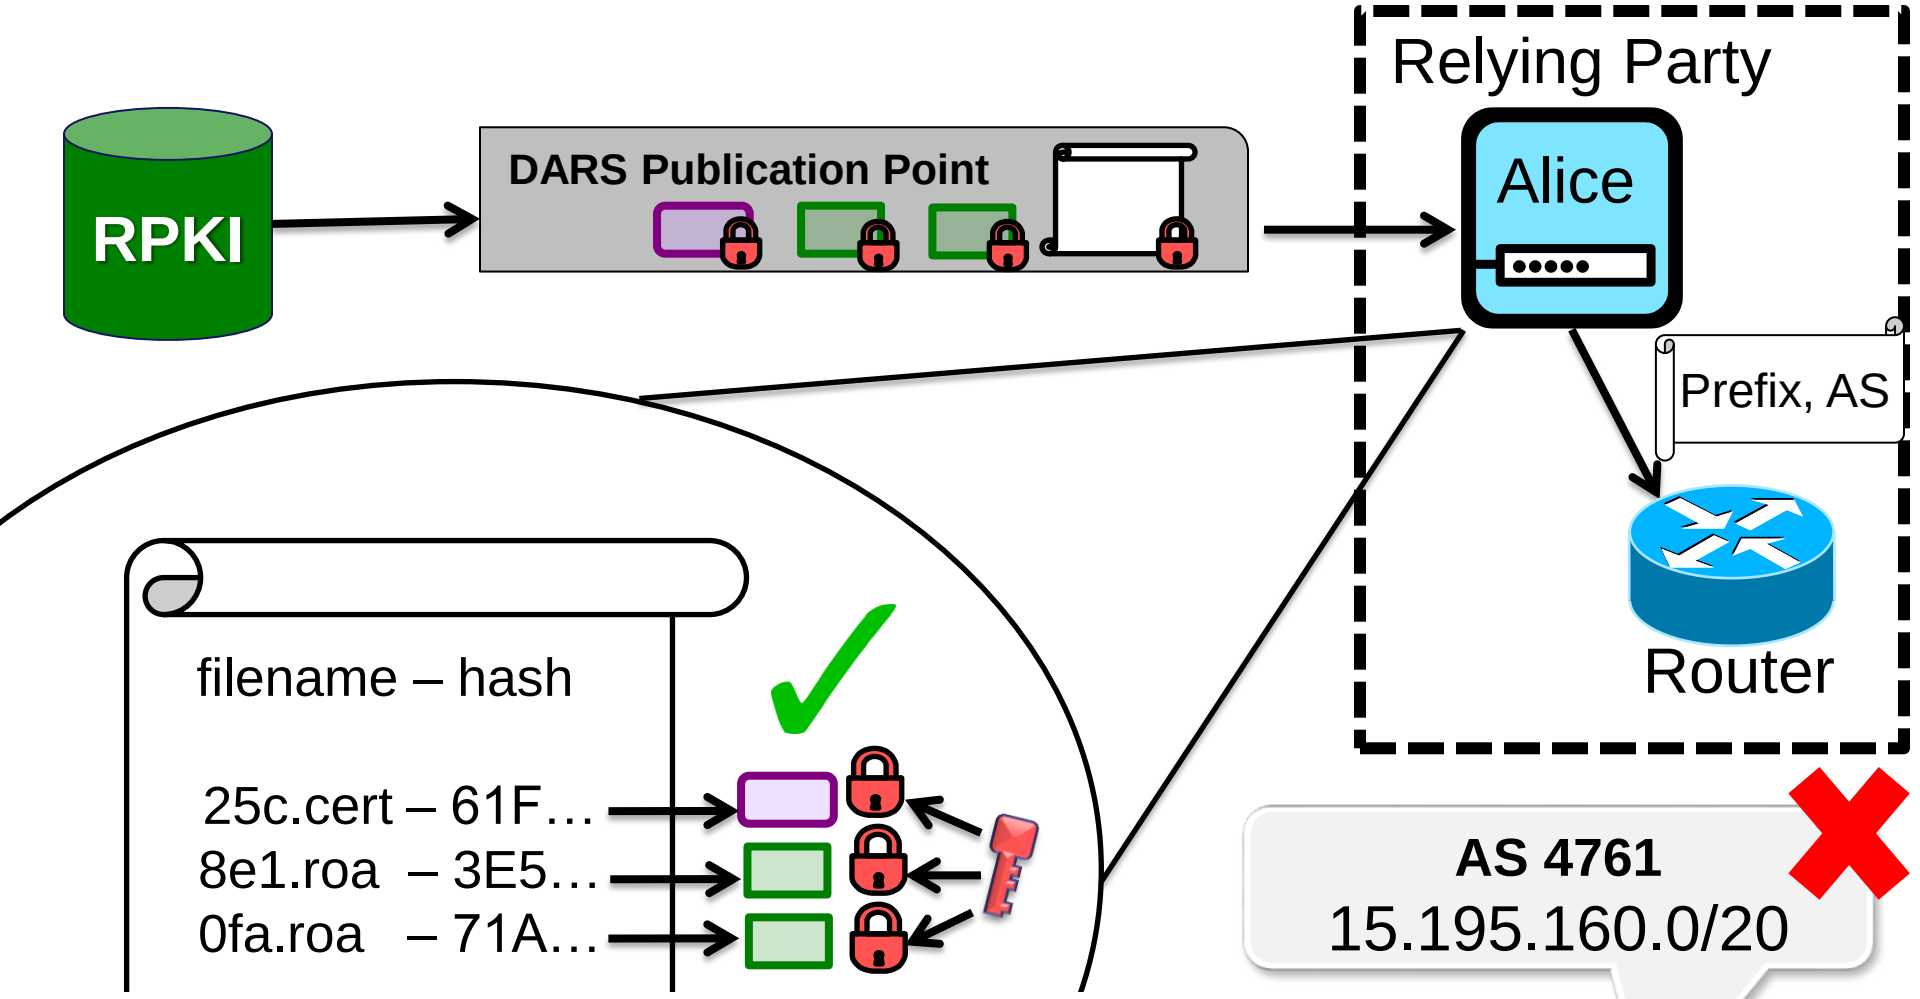

Status of the RPKI today:

- Today, few routers discard “RPKI invalid” routes

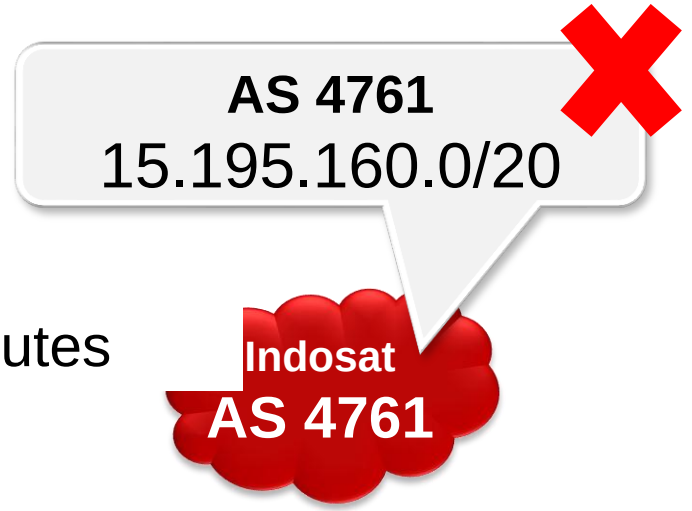

Misbehaving RPKI authorities.

- Prior to the RPKI, authorities could allocate IPs but not revoke them.
- But RPKI authorities **can** revoke allocations!
- Creates a risk that the RPKI can be used for unilateral takedowns.
  - Law enforcement? Business disputes? Extortion?
  - The RPKI designed to secure routing, not enable takedowns.
  - **[Mueller-Kuerbis'11, Mueller-Schmidt-Kuerbis'13, Amante'12, FCC'13,…]**
- States seem to want the ability to takedown IP prefixes…
  - Dutch court ordered RIPE to takedown prefixes (Nov'11)
  - US court issued a writ of attachment on Iran's IP prefixes (June'14)
  - IP allocation does not reflect jurisdiction.

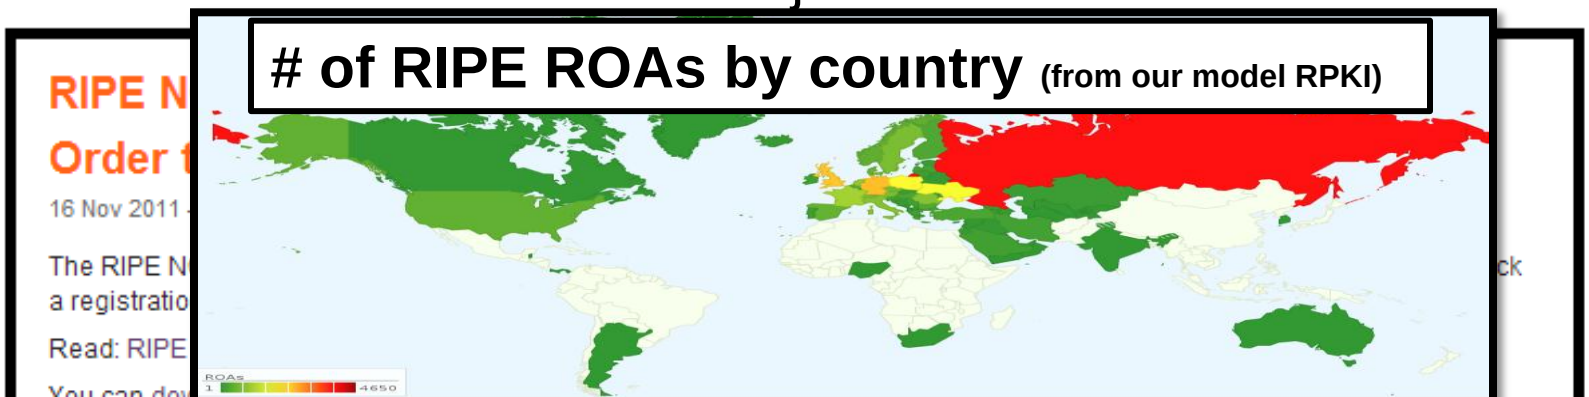

An RPKI takedown?

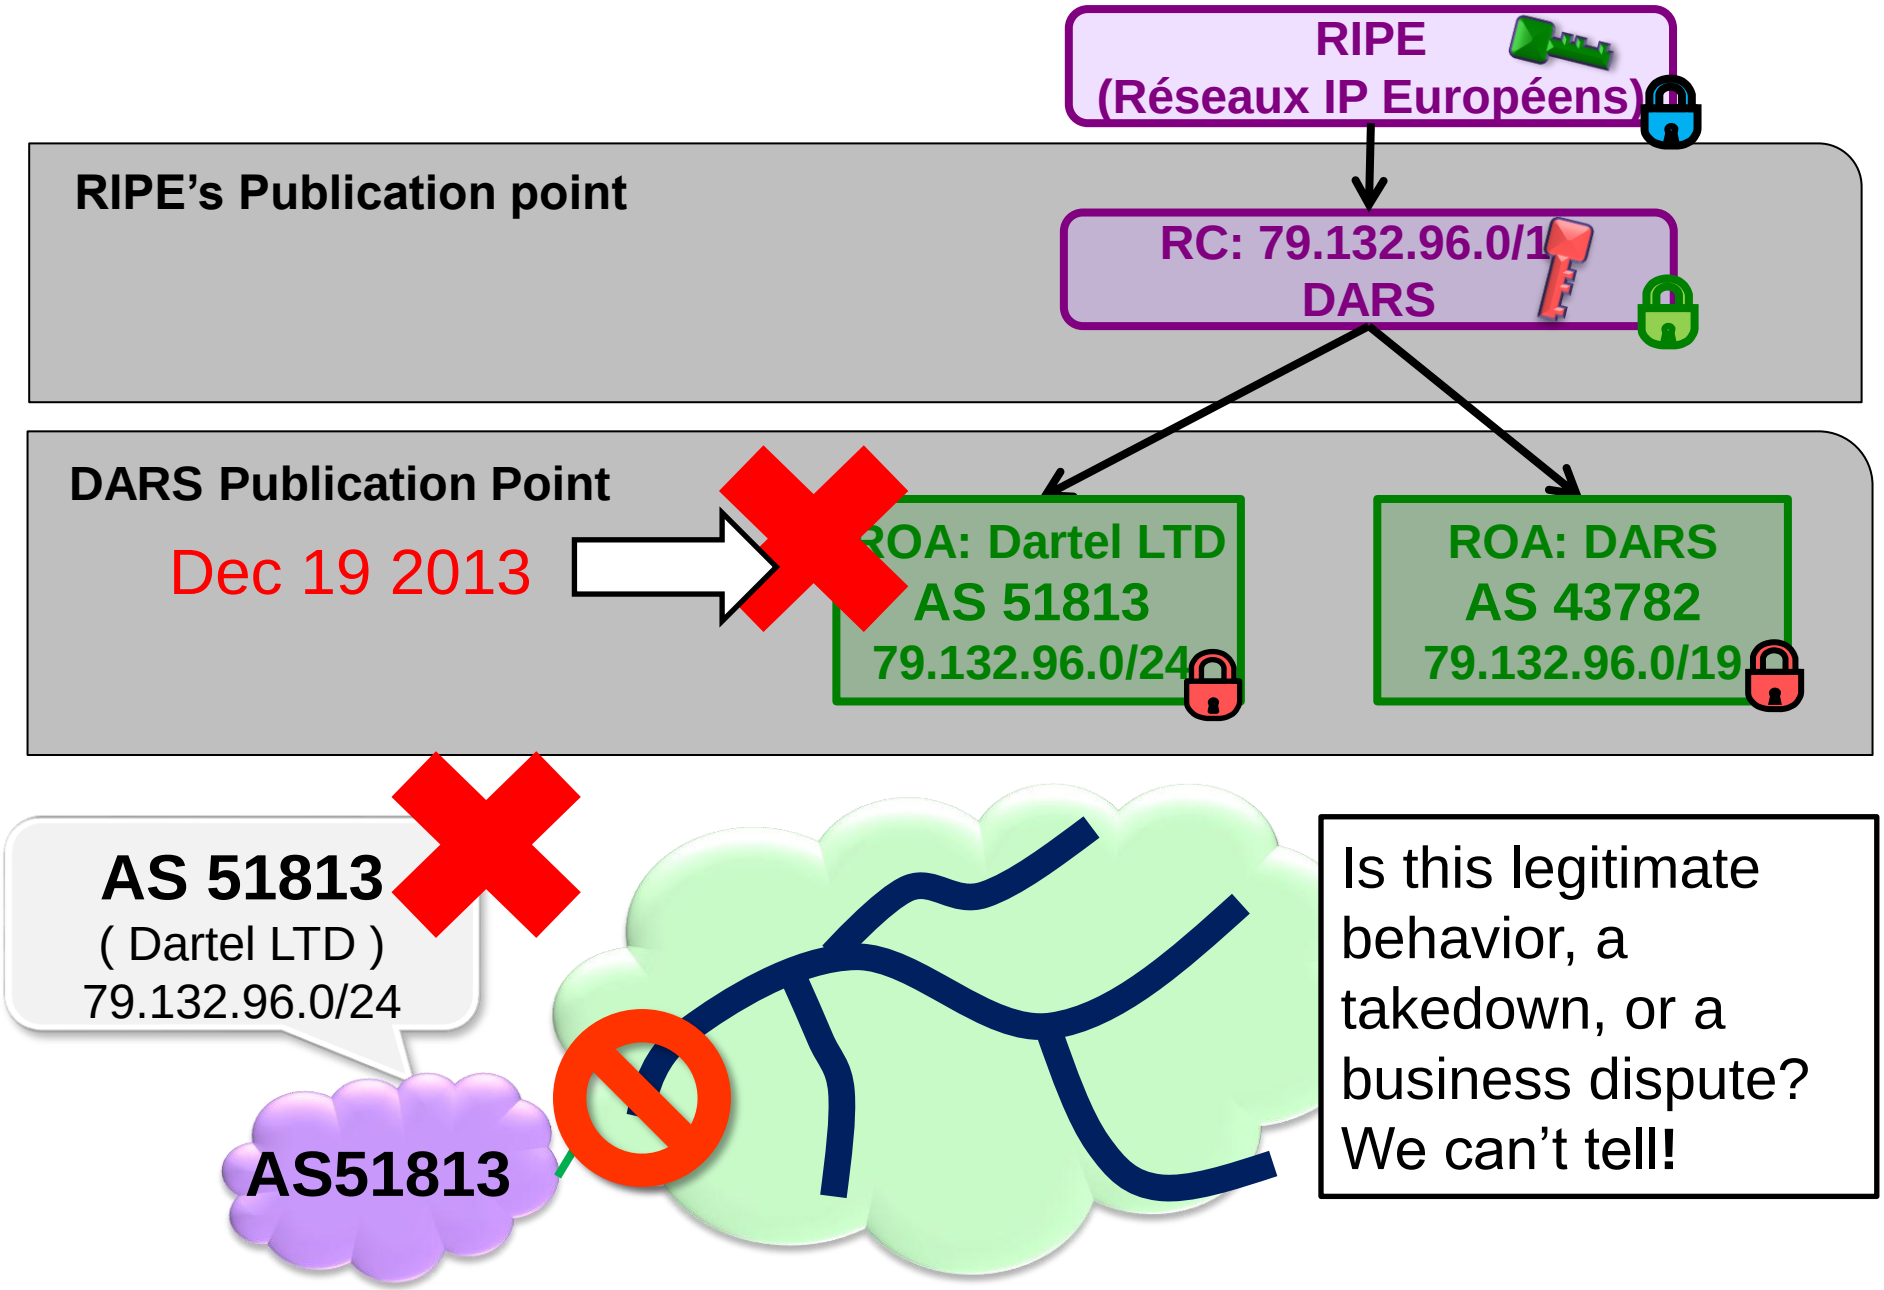

Proposed changes to the RPKI

- **Design Goals:**
  - **Transparency:** Relying parties audit the RPKI & alarm on problems.
  - **Consent**: RCs can indicate their consent to be revoked. Alarms are raised for revocations without consent.
  - **Consistency**: Relying parties have the same view of the RPKI.

- **Our Threat Model:**
  - Similar to the threat model used in certificate transparency **[RFC 6962]**
  - Relying parties are honest

  - Everyone else (including RPKI authorities) is untrusted

How consent works.

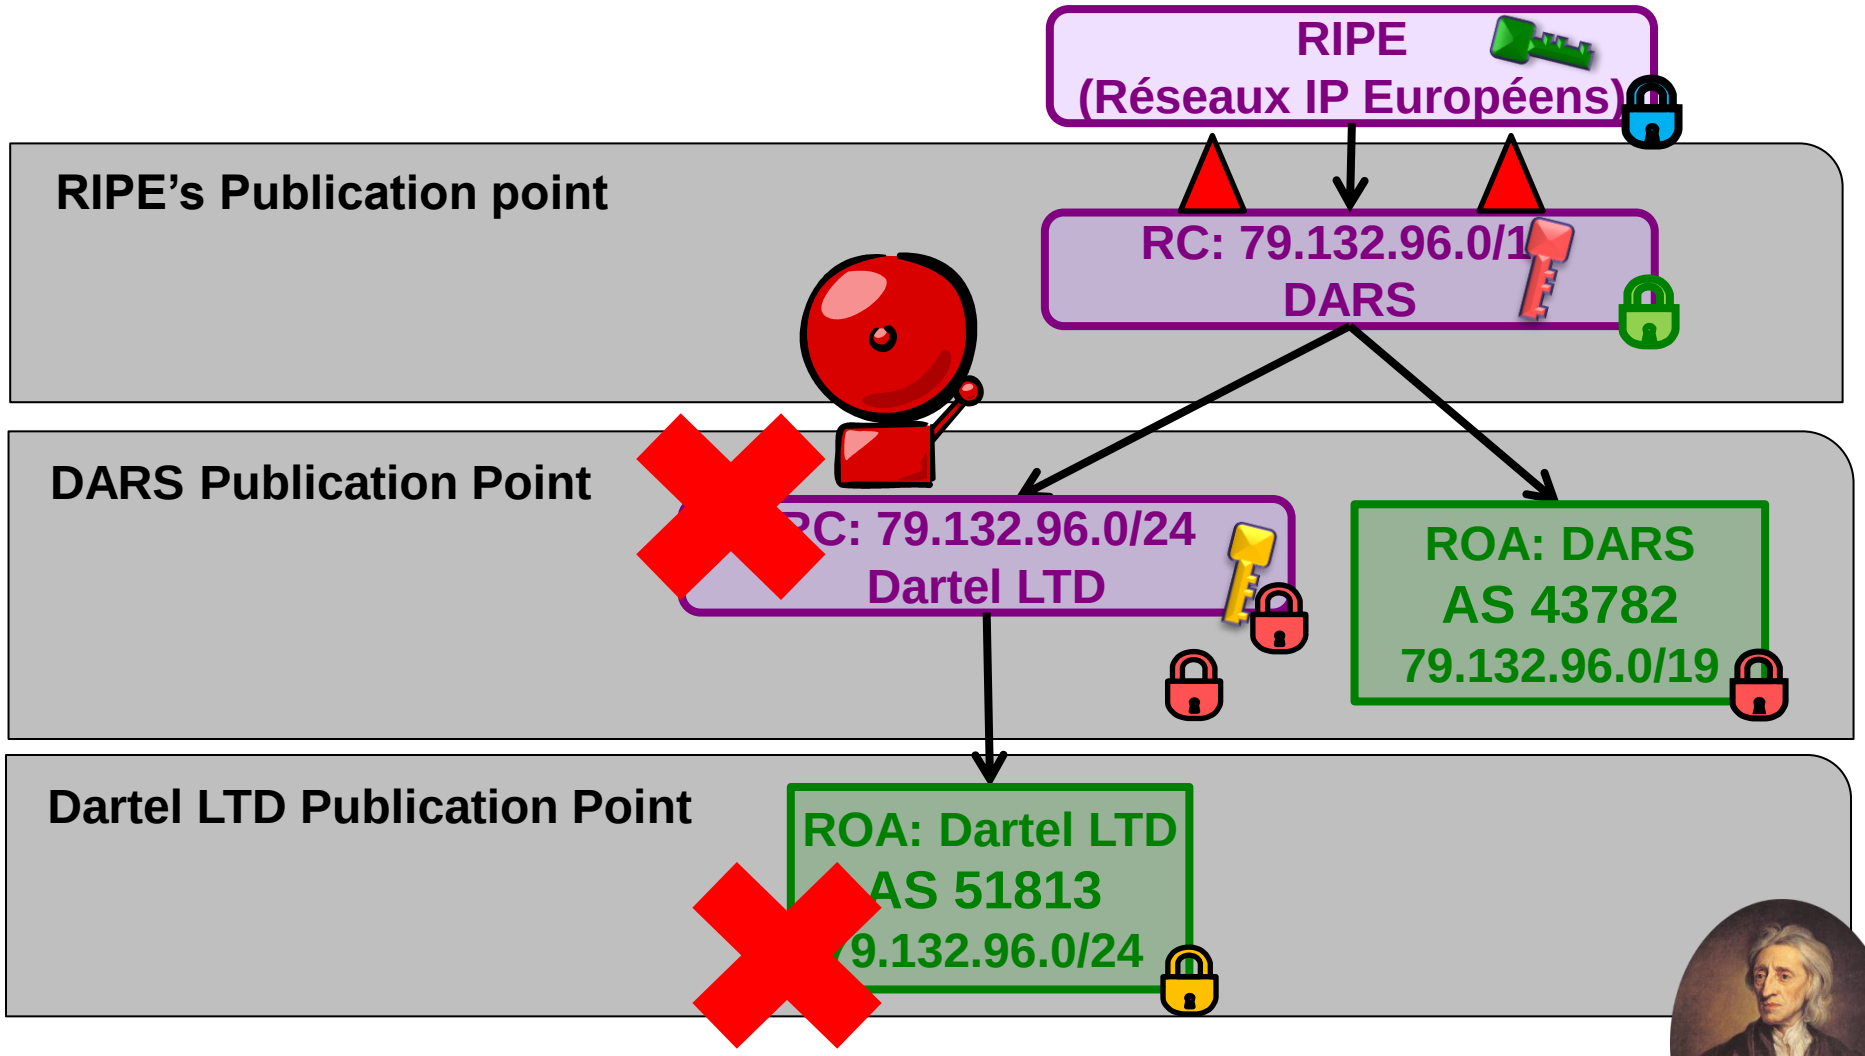

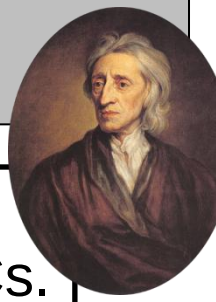

If an authority wants to revoke IP prefixes from a child RC,
it needs consent from that child & its impacted* descendant RCs.

*Descendants aren't always impacted by changes to the parent; ask me why later!

How consent works.

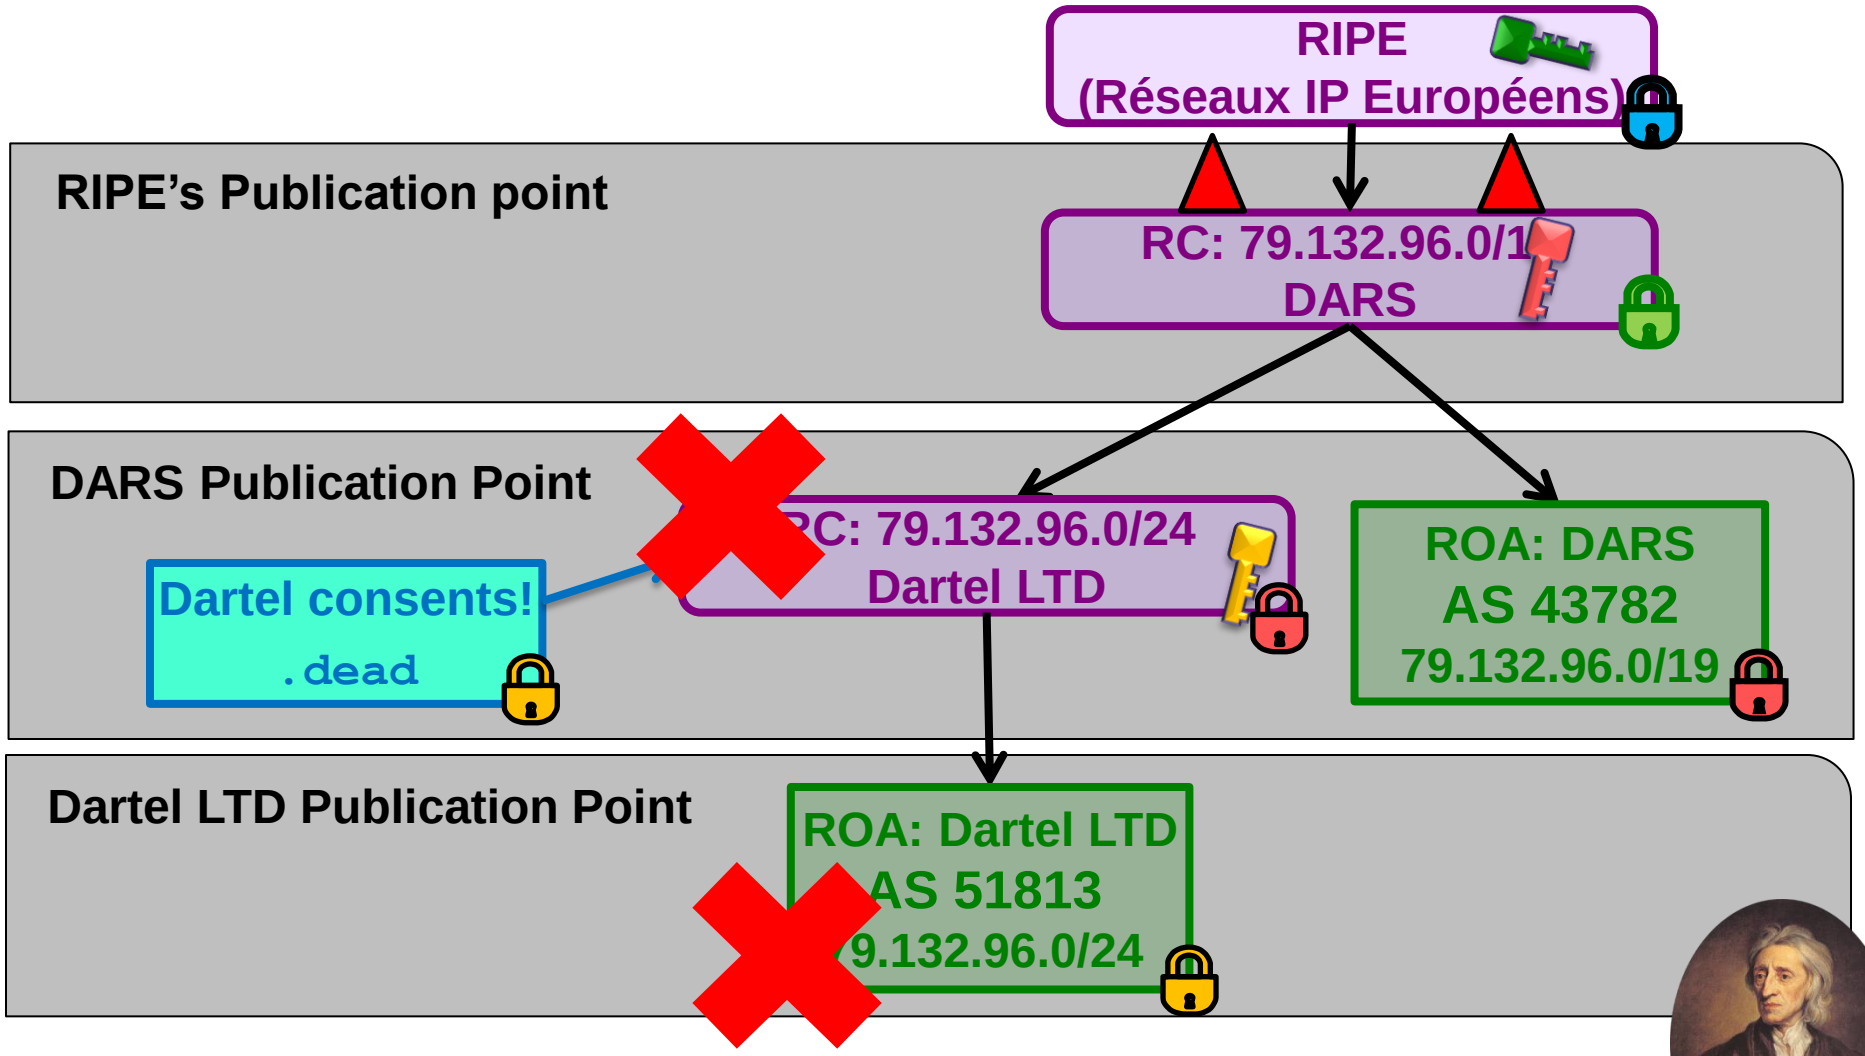

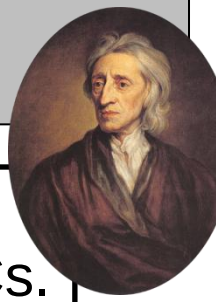

If an authority wants to revoke IP prefixes from a child RC,
it needs consent from that child & its impacted* descendant RCs.

*Descendants aren't always impacted by changes to the parent; ask me why later!

What about alarms between syncs?

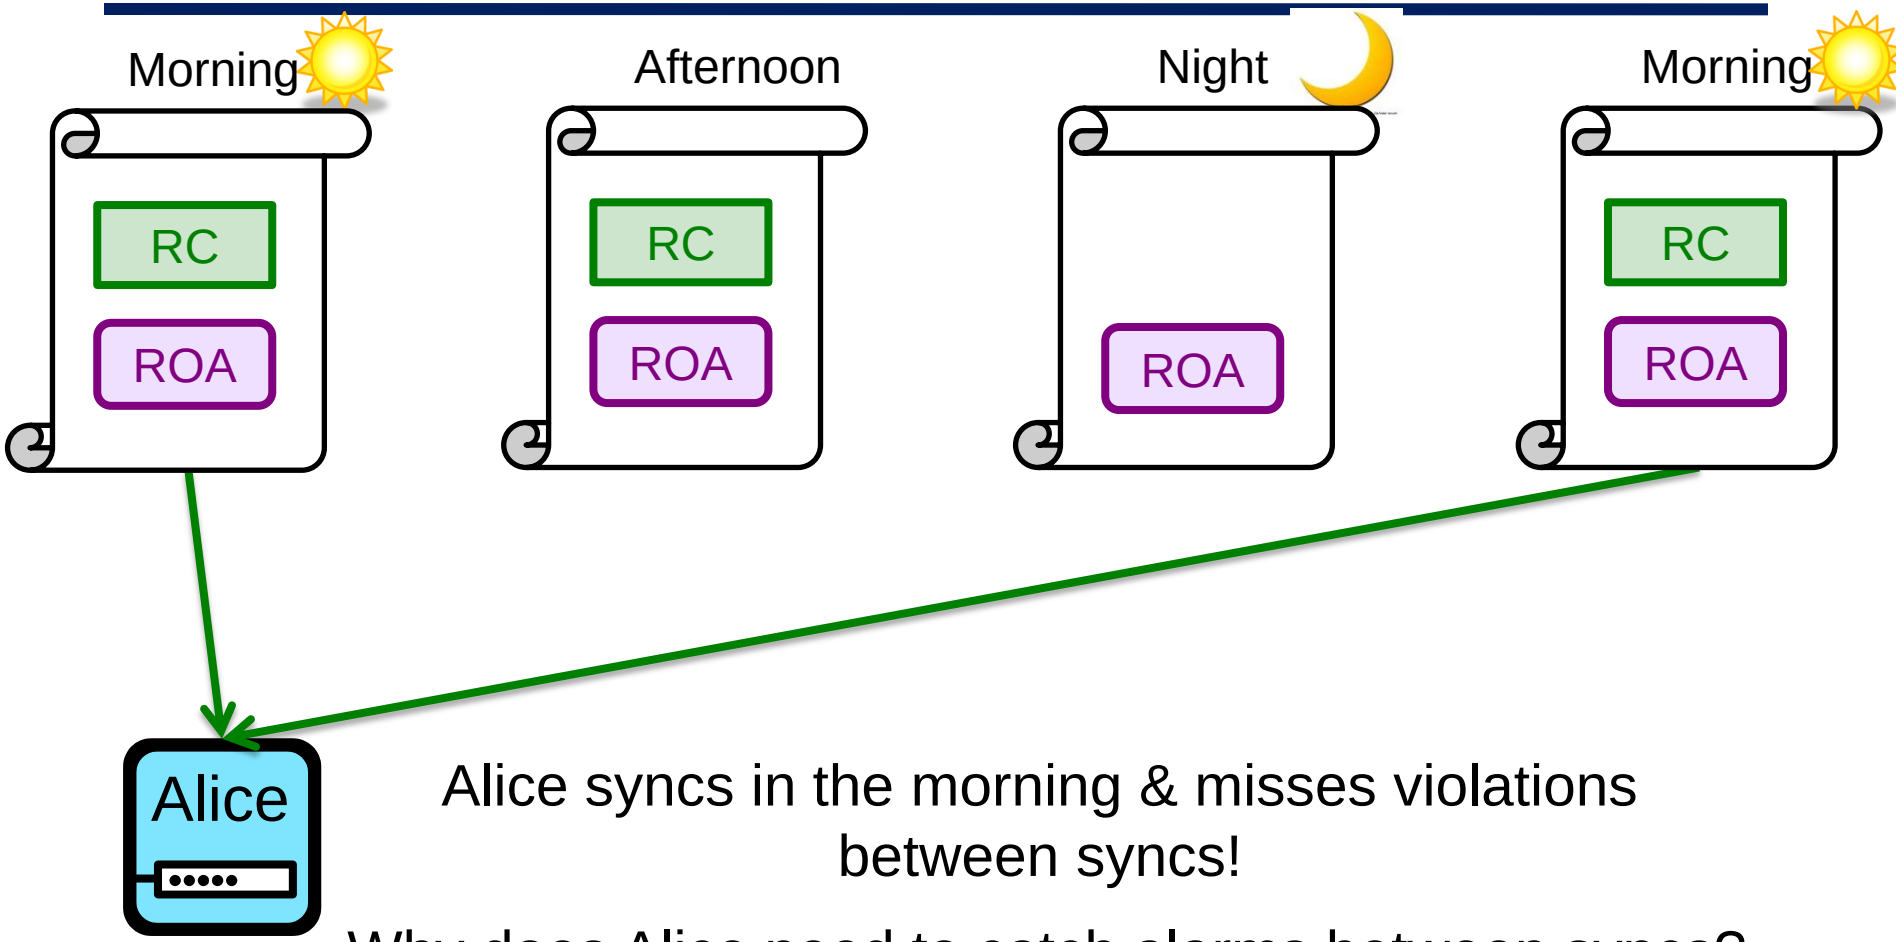

Alice syncs in the morning & misses violations between syncs!

Why does Alice need to catch alarms between syncs?

1) So relying parties can audit the RPKI
2) So we can have consistency (explained later)

Catching alarms between syncs.

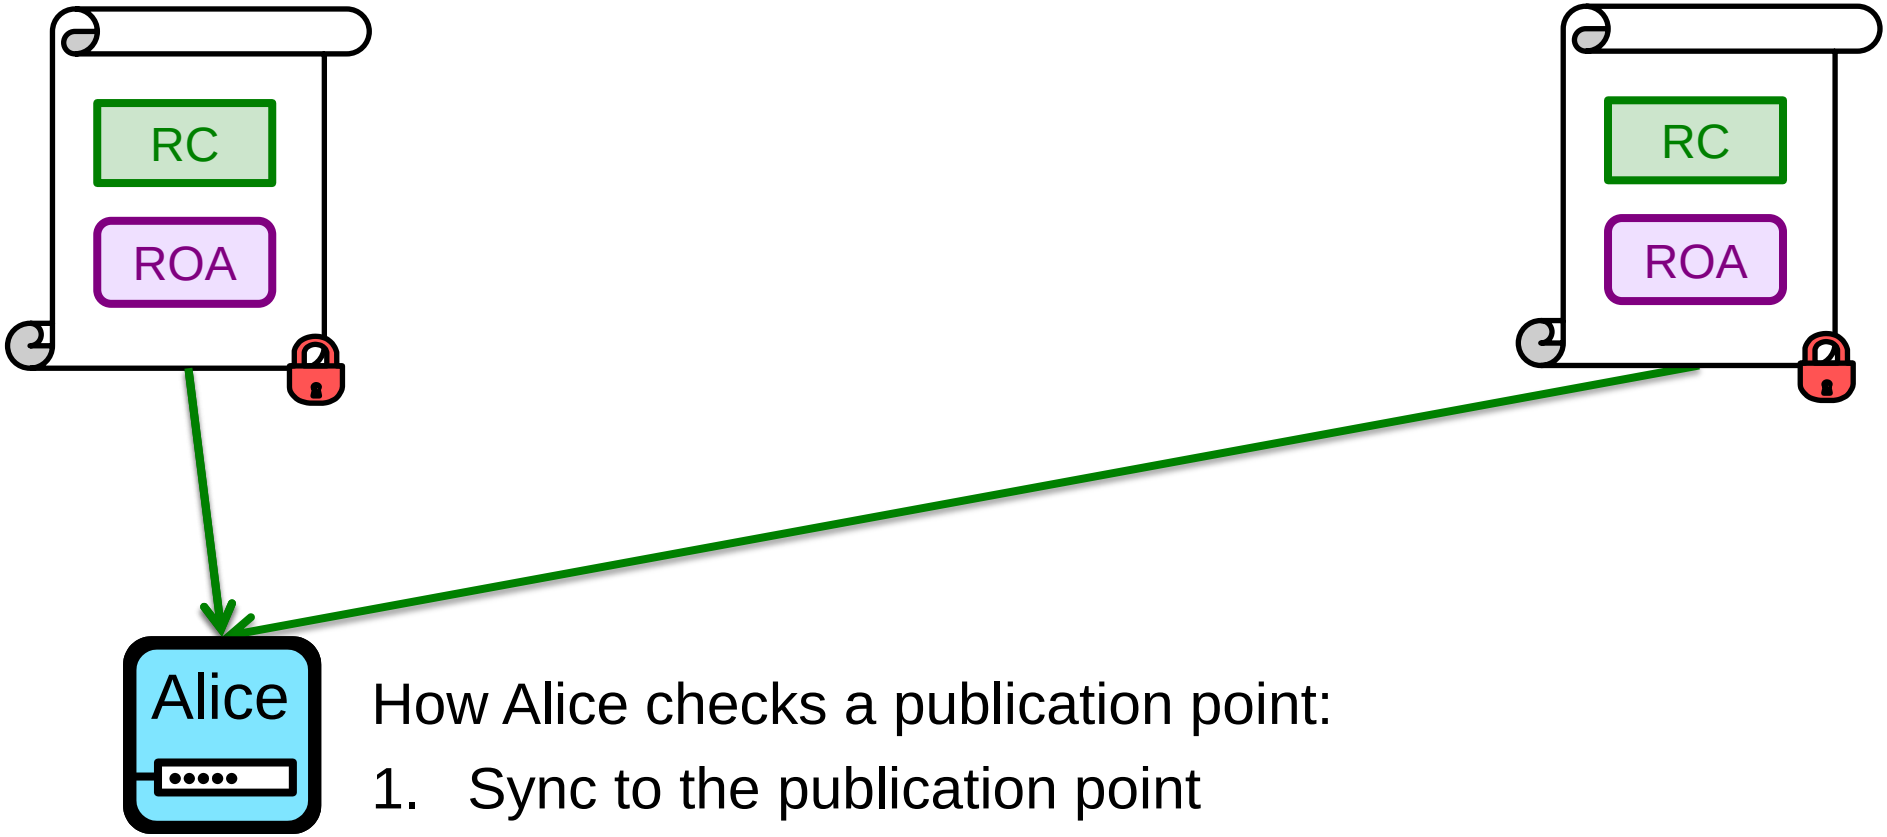

How Alice checks a publication point:

1. Sync to the publication point
2. Use hints file to reconstruct intermediate manifests
3. Verify the hash chain & signature of the latest manifest
4. Alarm if a consent violation is detected.

Catching alarms between syncs.

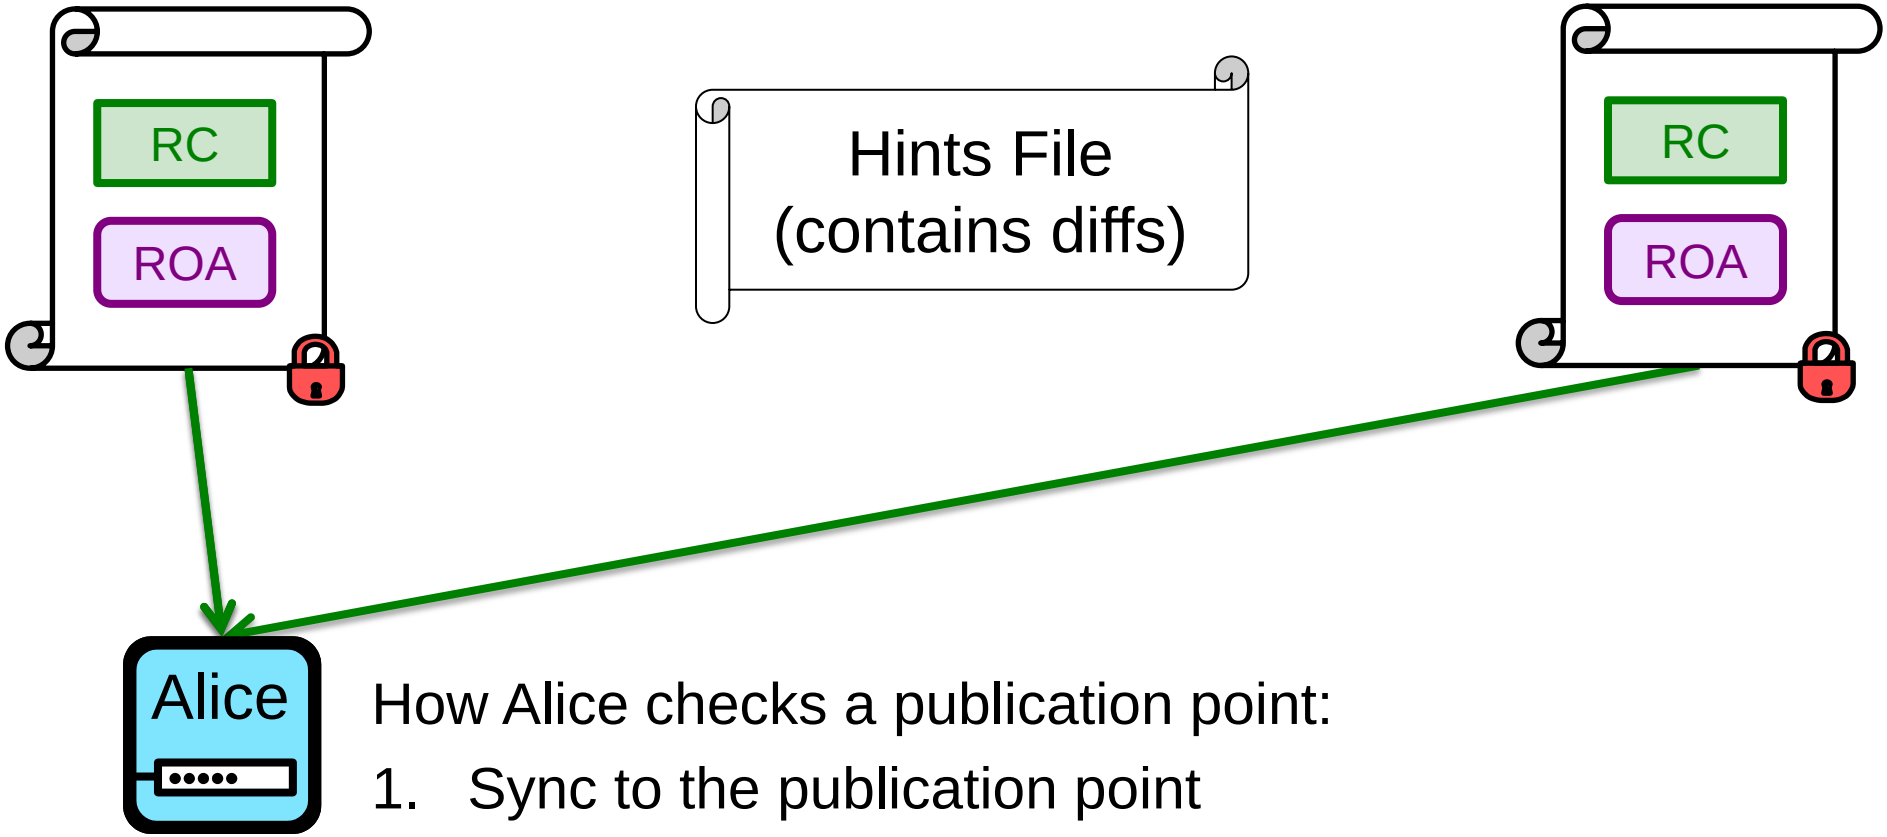

How Alice checks a publication point:

1. Sync to the publication point
2. Use hints file to reconstruct intermediate manifests
3. Verify the hash chain & signature of the latest manifest
4. Alarm if a consent violation is detected.

Catching alarms between syncs.

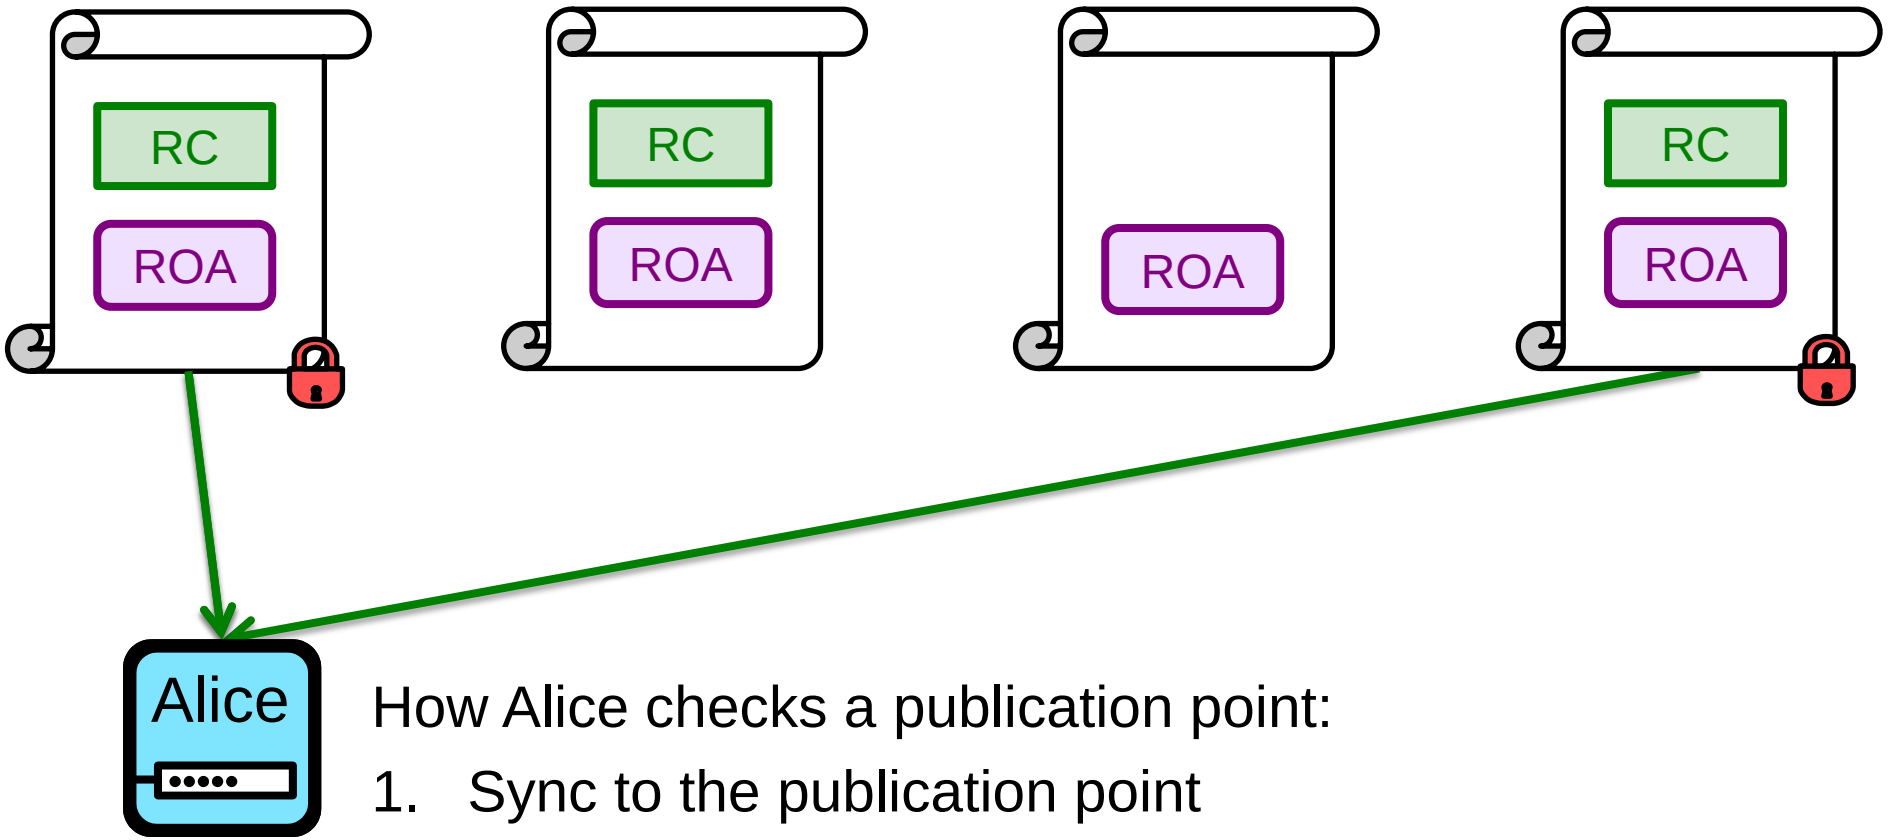

How Alice checks a publication point:

1. Sync to the publication point
2. Use hints file to reconstruct intermediate manifests
3. Verify the hash chain & signature of the latest manifest
4. Alarm if a consent violation is detected.

Catching alarms between syncs.

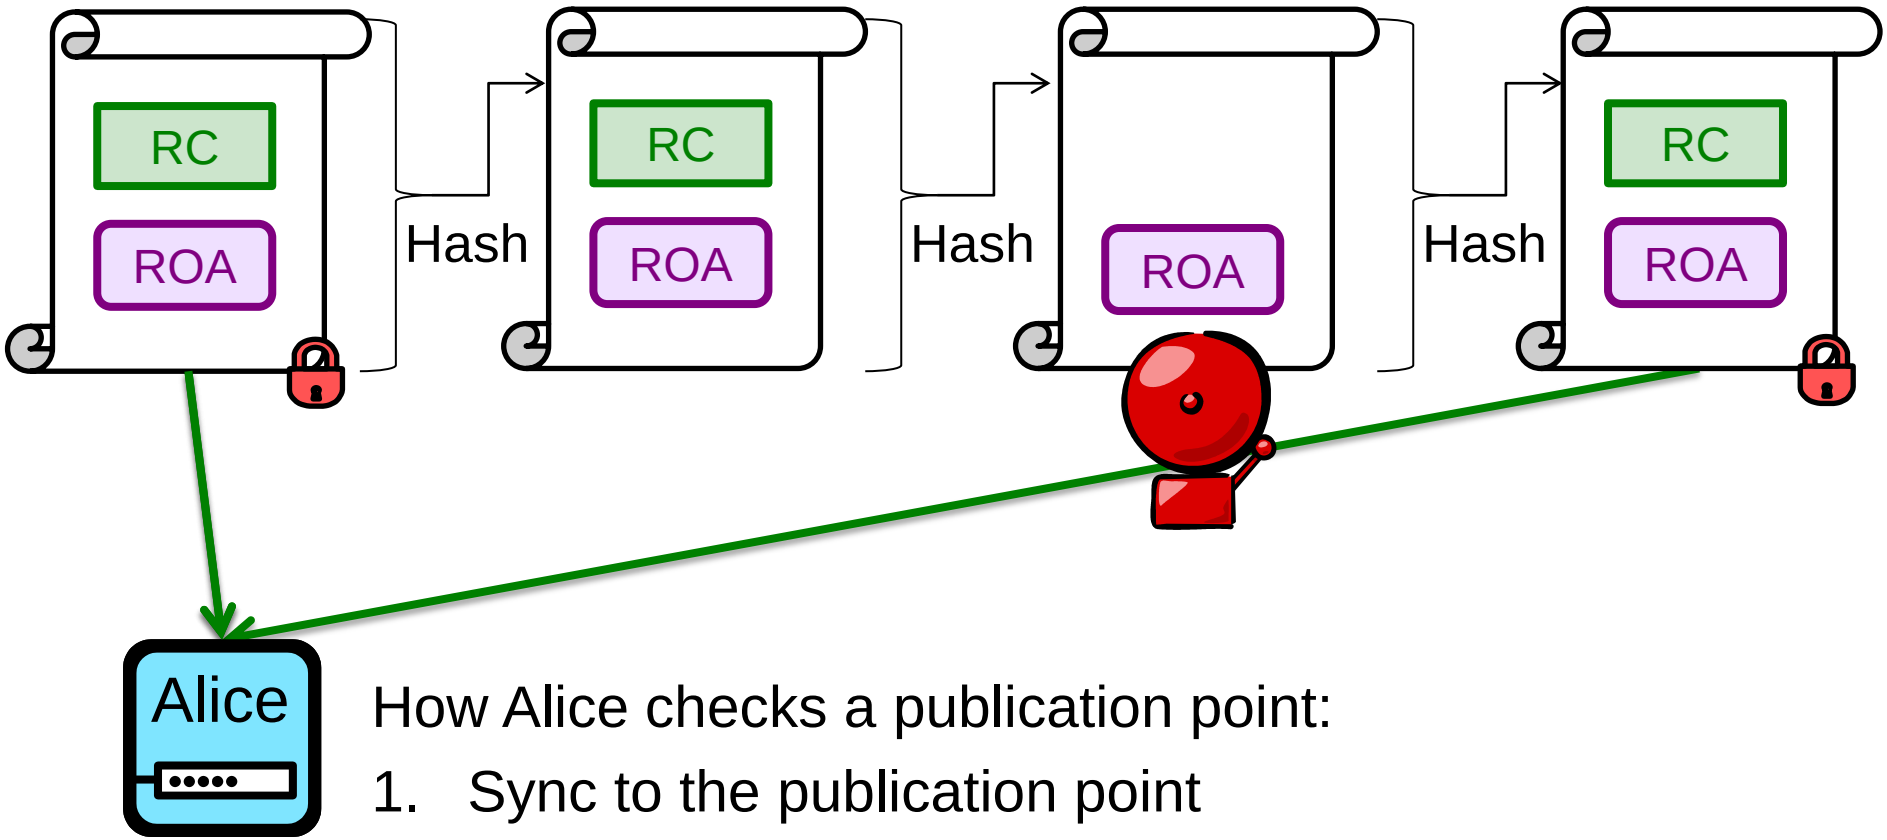

How Alice checks a publication point:

1. Sync to the publication point
2. Use hints file to reconstruct intermediate manifests
3. Verify the hash chain & signature of the latest manifest
4. Alarm if a consent violation is detected.

Catching alarms between syncs.

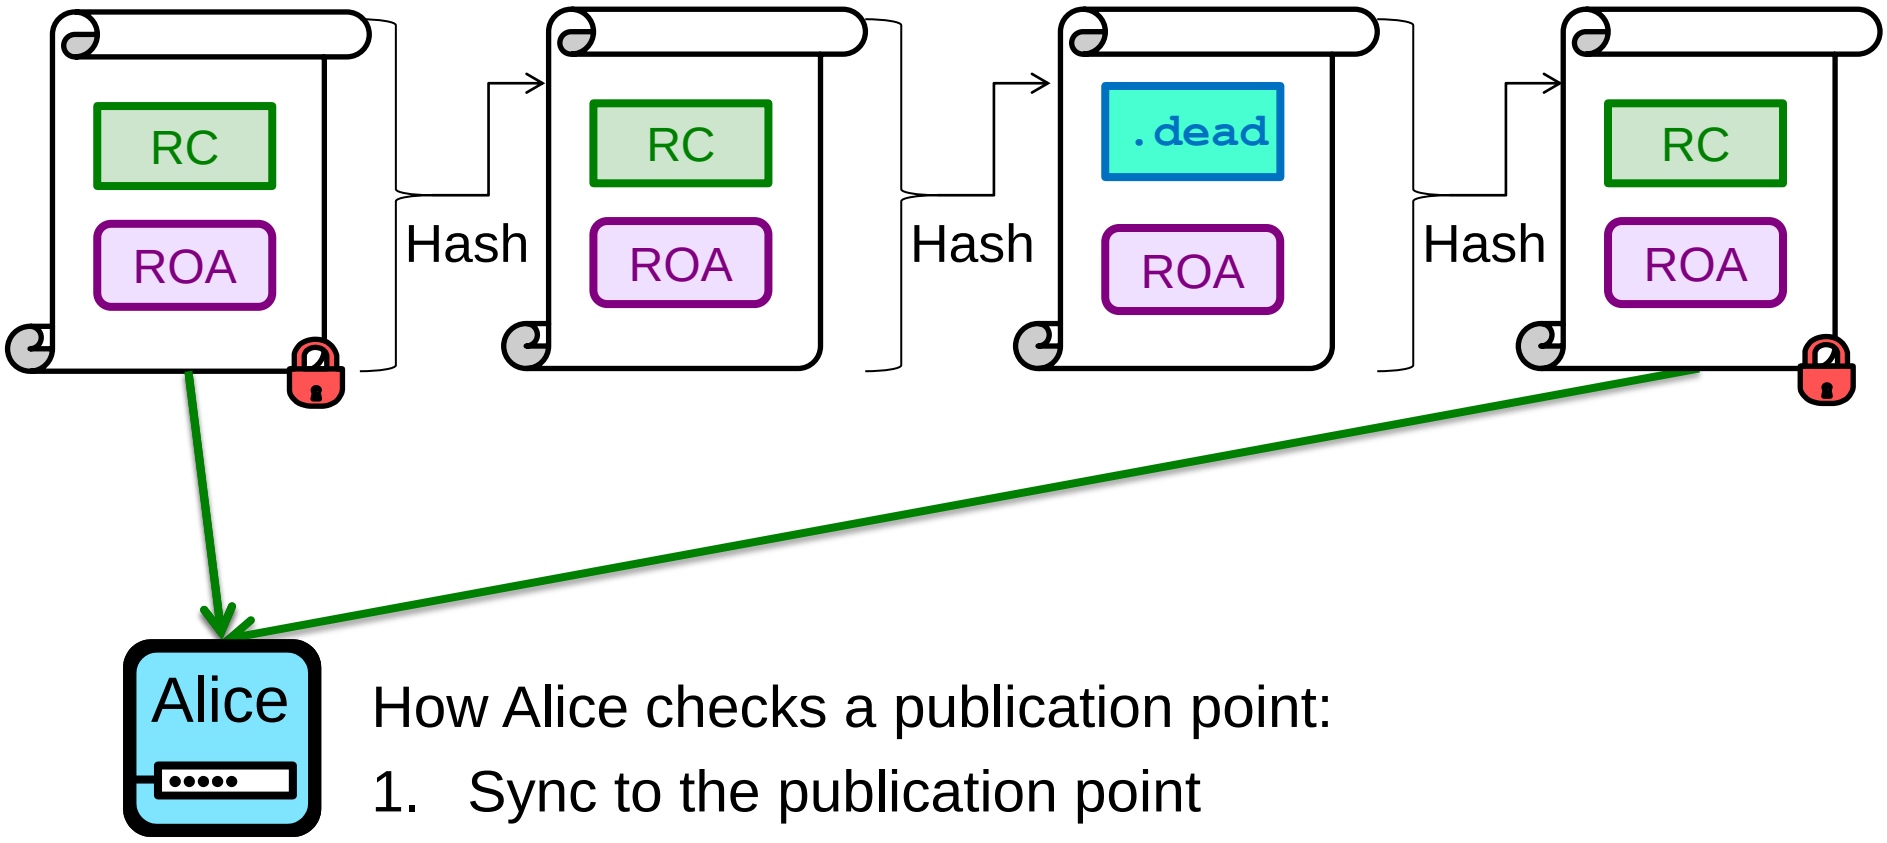

How Alice checks a publication point:

1. Sync to the publication point
2. Use hints file to reconstruct intermediate manifests

Theorem: Valid Remains Valid.
Once a relying party has seen a valid RC,
that RC remains valid until it consents to be deleted/modified.

How many parties need to consent?

- How many ASes need to be involved when an RC is revoked?

- Production RPKI
  - average **1.5** ASes / leaf RC

- Model fully-deployed RPKI
  - average **1.6** ASes / leaf RC
  - **99.3%** need **<10** ASes / leaf RC
  - **0.02%** need **>100** ASes / leaf RC

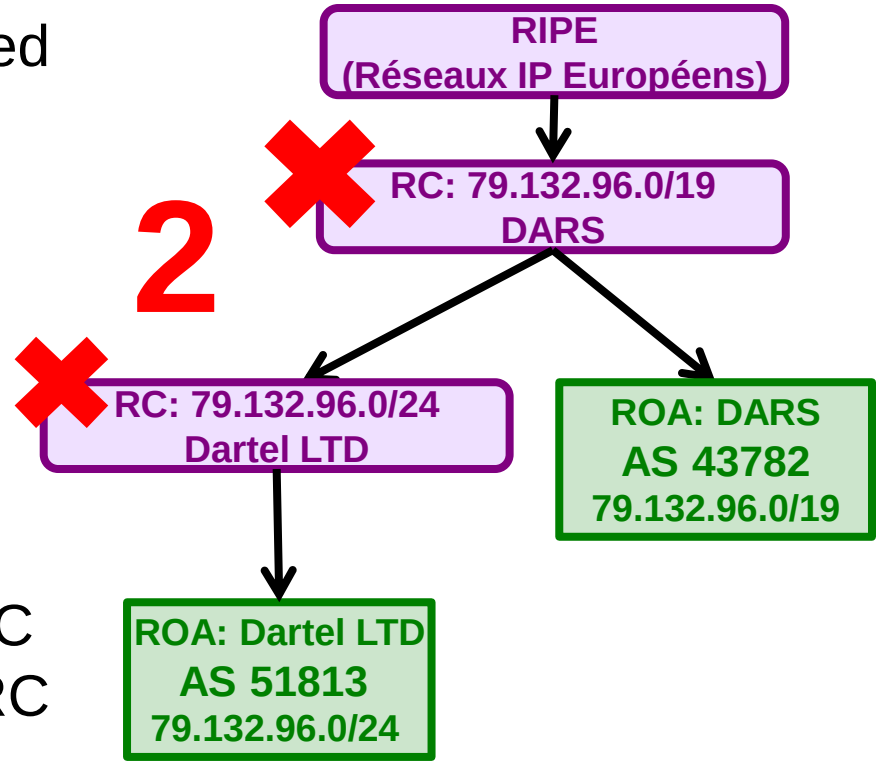

Results: production RPKI

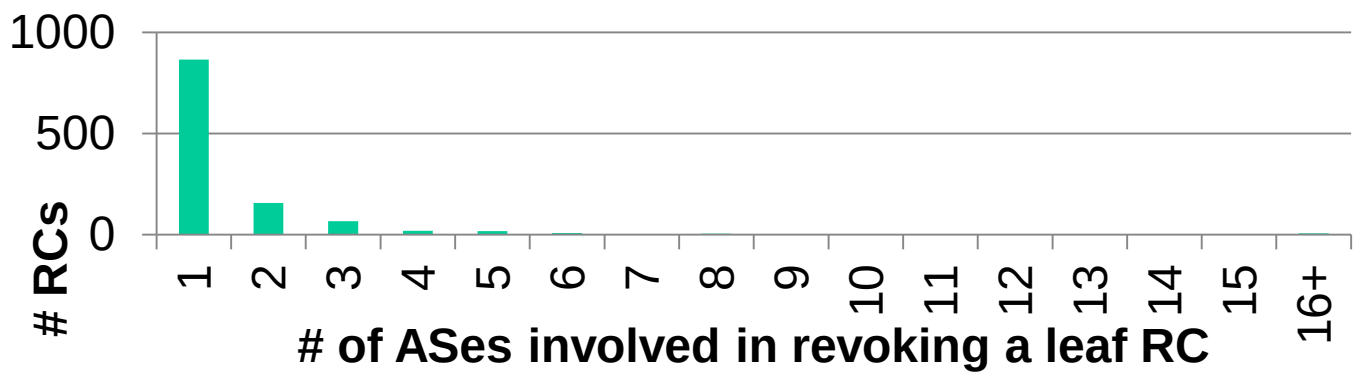

How many parties need to consent?

- How many ASes need to be involved when an RC is revoked?

- Production RPKI
  - average **1.5** ASes / leaf RC

- Model fully-deployed RPKI
  - average **1.6** ASes / leaf RC
  - **99.3%** need **<10** ASes / leaf RC
  - **0.02%** need **>100** ASes / leaf RC

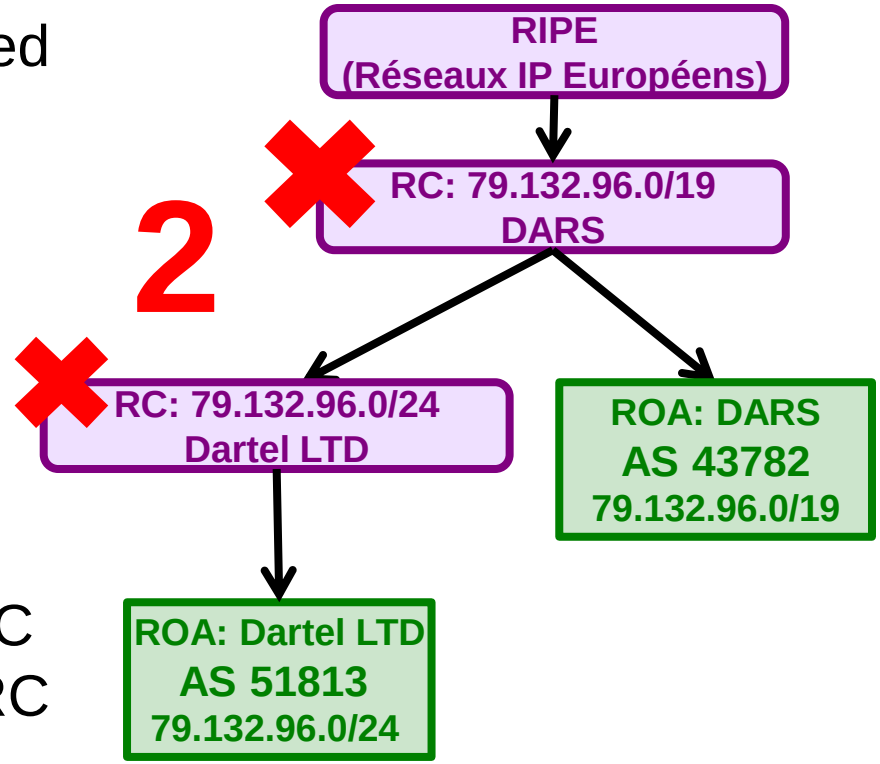

"With great power comes great responsibility"

- Voltaire, Spiderman

Proposed changes to the RPKI

- **Design Goals:**
  - **Transparency:** Relying parties audit the RPKI through alarms.
  - **Consent**: If an authority wants to revoke IP prefixes from a child RC, it needs consent from the child RC & its impacted descendant RCs.
  - **Consistency**: Relying parties have the same view of the RPKI.

Mirror world attacks.

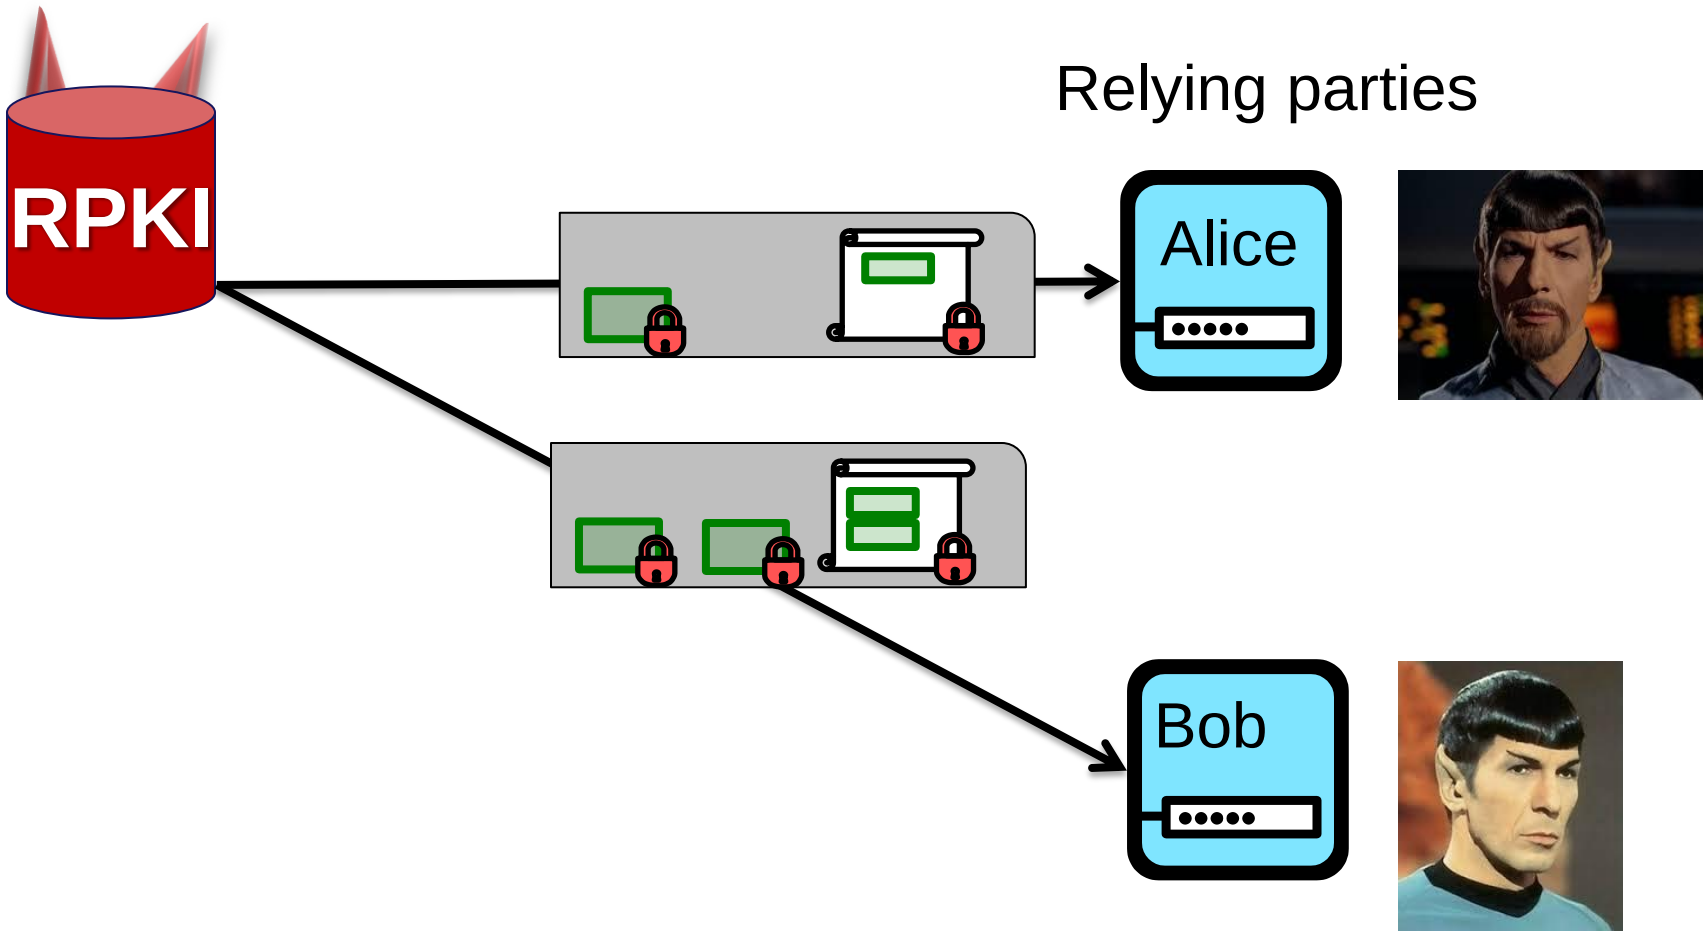

Mirror world attack: RPKI Authority presents one view to a relying parties and a different view to others.

Detecting mirror worlds using manifest hash chains

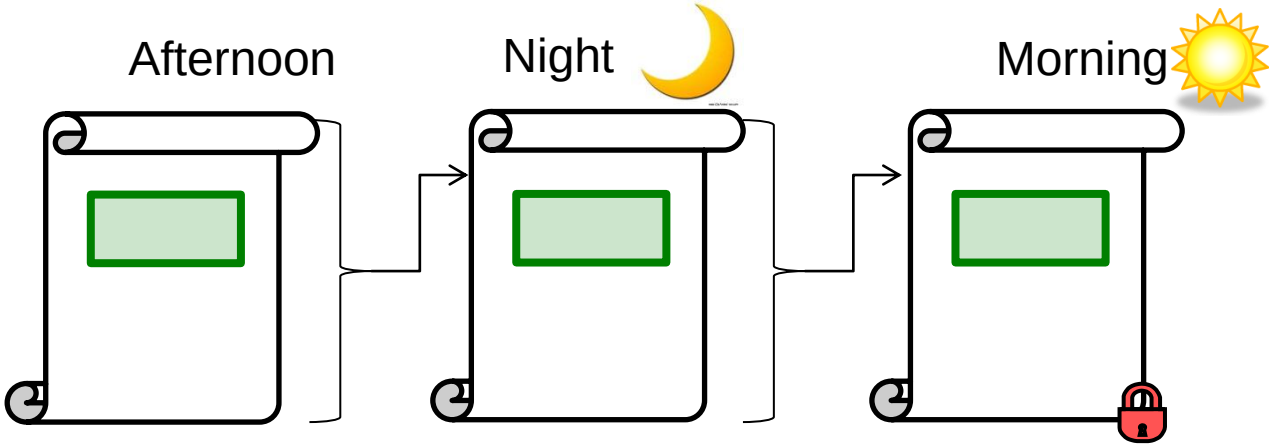

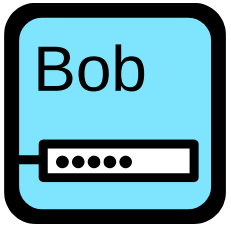

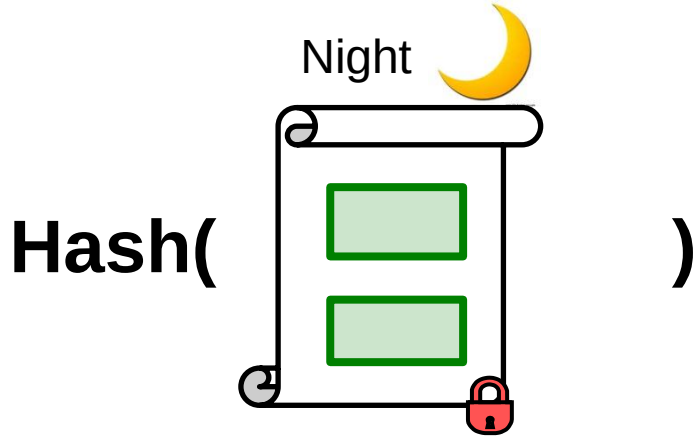

Bob sends a hash of his latest manifest & Alice finds it in her hashchain.

Theorem: No mirror worlds.
If the consistency check passes,
relying parties saw the same valid objects.

The challenge of asynchronous validity changes.

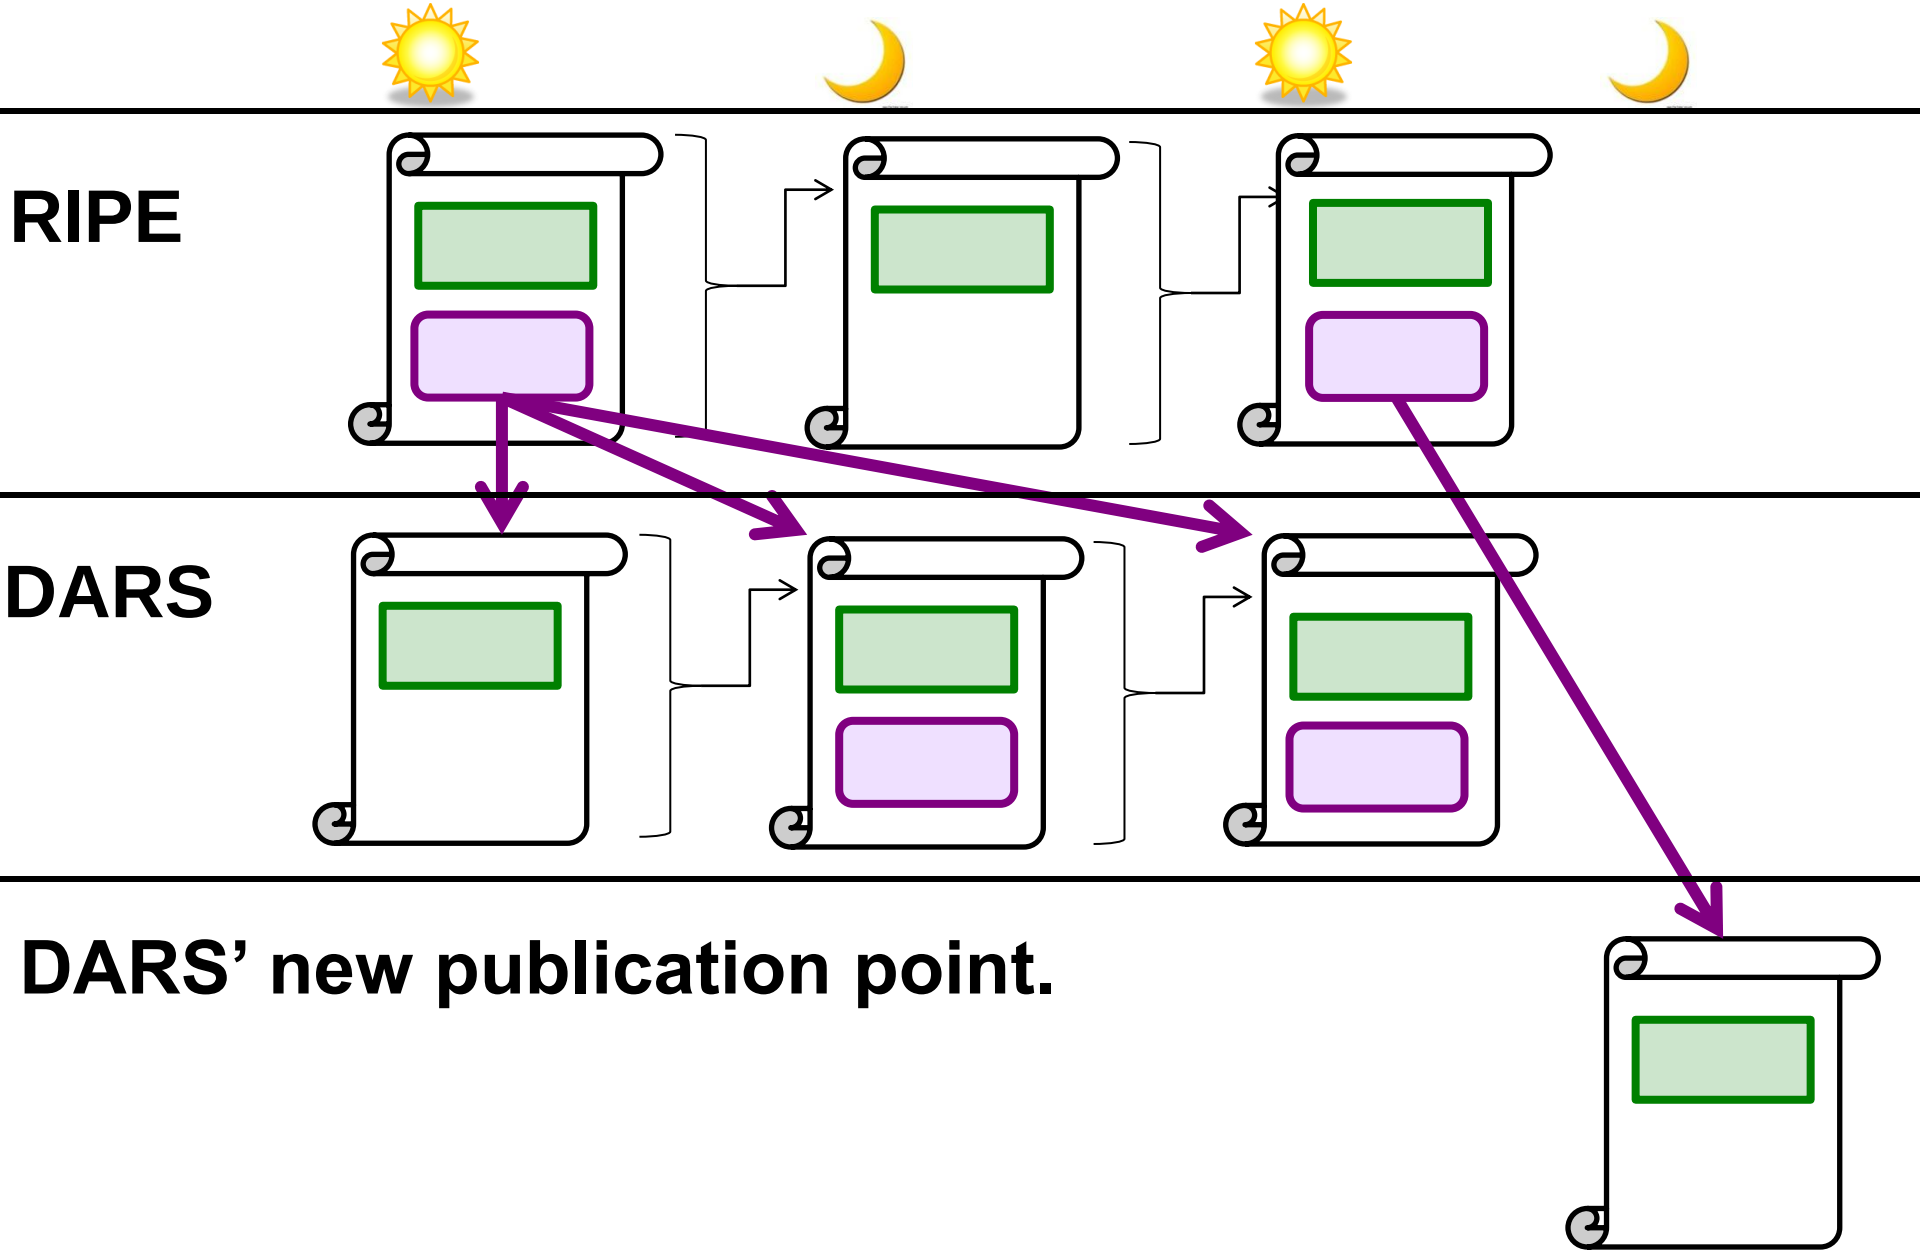

Summary.

Motivation: RPKI secures interdomain routing,
but creates a new danger of misbehaving authorities.

- **Our proposed changes:**

  Consent through .dead objects.

  Consistency through via hints files, hash-chained manifests, & checks between relying parties.

- **Our changes are practical and effective:**
  - We extend existing mechanisms within the RPKI.
  - Consent requires minimal work for ASes (see paper for details).

- **Window of opportunity to influence RPKI design:**
  - Changes being still being made to RPKI specification.
  - Concurrent to our work, IETF is drafting misbehavior defenses **[draft-kent-sidr-suspenders-01].**

check out the full version at

http://cs-people.bu.edu/heilman/sigRPKI.pdf

1. Measurements of revocations in production RPKI
2. Tools for detecting & visualizing revocations and downgrades
3. Details of our proposed changes to the RPKI

download our detector at

https://github.com/BUSEC/RPKI_Downgrade_Detector

Ask questions on twitter: **@Ethan_Heilman #consentRPKI**