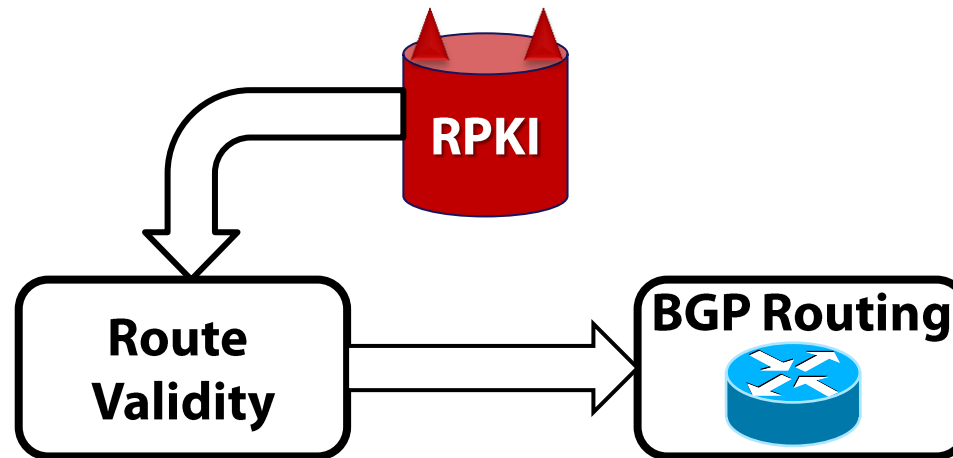


On the Risk of Misbehaving RPKI Authorities




**BOSTON
UNIVERSITY**

Sharon Goldberg

**Danny Cooper, Ethan Heilman,
Kyle Brogle, Leonid Reyzin**

the canadian bitcoin BGP subprefix hijack (feb 3, 2014)



and many other BGP prefix hijacks...

2010

REPORT TO CONGRESS

of the

U.S.-CHINA ECONOMIC AND SECURITY REVIEW COMMISSION



Interception of Internet Traffic

For a brief period in April 2010, a state-owned Chinese telecommunications firm “hijacked” massive volumes of Internet traffic.*¹¹⁴ Evidence related to this incident does not clearly indicate whether it was perpetrated intentionally and, if so, to what ends. However, computer security researchers have noted that the capability could enable severe malicious activities.¹¹⁵

The New York Times



The Lede

The New York Times News Blog

Pakistan Blamed for Worldwide YouTube Break

By MIKE NIZZA FEBRUARY 25, 2008 9:34 AM

If all had gone according to plan, Pakistan would have been the latest government taking part in an unsettling trend from its predecessors, though beyond its borders you couldn't

NetClean

WhiteBox

WIRED

GEAR SCIENCE ENTERTAINMENT BUSINESS SECURITY DESIGN OPINION

Someone's Been Siphoning Data Through a Huge Security Hole in the Internet

BY KIM ZETTER 12.05.13 | 6:30 AM | PERMALINK

Traceroute Path 2: from Denver, CO to Denver, CO via Iceland



Hijacked traffic went all the way to Iceland, where it may have been copied before being released to its intended destination. The green arrows show the path the traffic should have traveled; the red arrows show the path it took. Map courtesy of Renesys

renesys

Con-Ed Steals the 'Net

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

Well, not the whole Internet, but Con Edison (AS27506) hijacked the Internet earlier today, probably by mistake. Earlier this morning, a NANOG mailing list claiming that Con Ed was “stealing”

BGP MON

Hijack event today by Indosat

Posted by Andree Toonk - April 3, 2014 - Hijack, News and Updates - 1 Comment

Today we observed a large-scale ‘hijack’ event that affected many of the prefixes on the Internet.

Symantec. Confidence in a connected world.

United States

Spam and Fraud Activity Trends

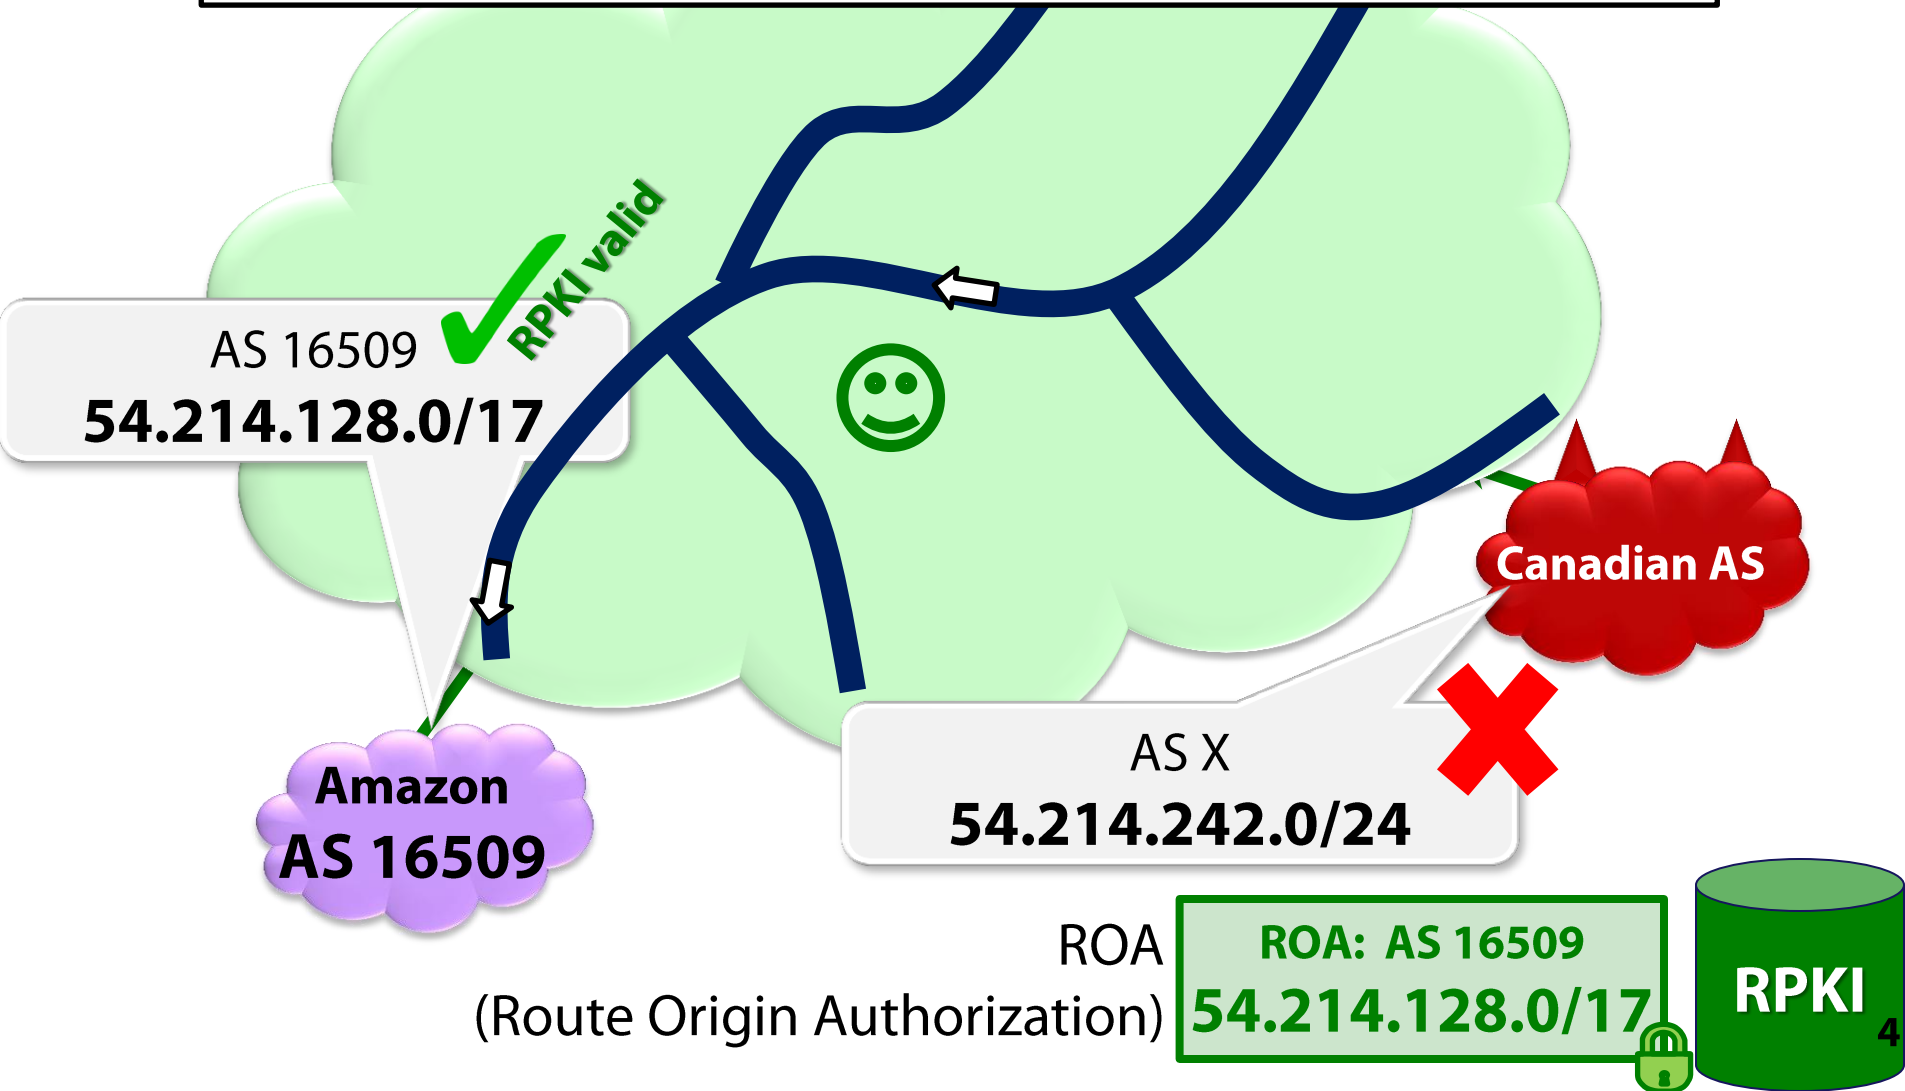
Future Spam Trends: BGP Hijacking
Case Study - Beware of “Fly-by Spammers”

Background

what is the fundamental vulnerability?

Problem: Route origin announcements are not authenticated.

Solution: The RPKI authenticates route origins.




the Resource Public Key Infrastructure (RPKI) [RFC6480]

- **How does the RPKI protect routing?**
 - It prevents prefix & subprefix hijacks caused by common misconfigs
 - Advanced path-validation solutions build on RPKI
 - BGPSEC [L'12] Secure BGP [KLS'99], soBGP [W'03],...
 - Even without path validation, our research [SIGCOMM'10, SIGCOMM'13] shows RPKI is good at limiting advanced BGP attacks
- **What about other routing security solutions?**
 - Anomaly detectors alarm when strange routes appear
 - BGPmon, renesys, pgBGP [KFR'06], PHAS [LMPWZZ'06], ...
 - Prefix filtering with IRRs
 - Requires distant ASes to implement filtering properly
 - Usually performed only on customer edges

traditional threat model for the RPKI

The RPKI is trusted but routing is under attack.




talk outline

Security audit of the RPKI [HotNets'13]

Misbehaving RPKI authorities can blackhole routes in BGP. Why?

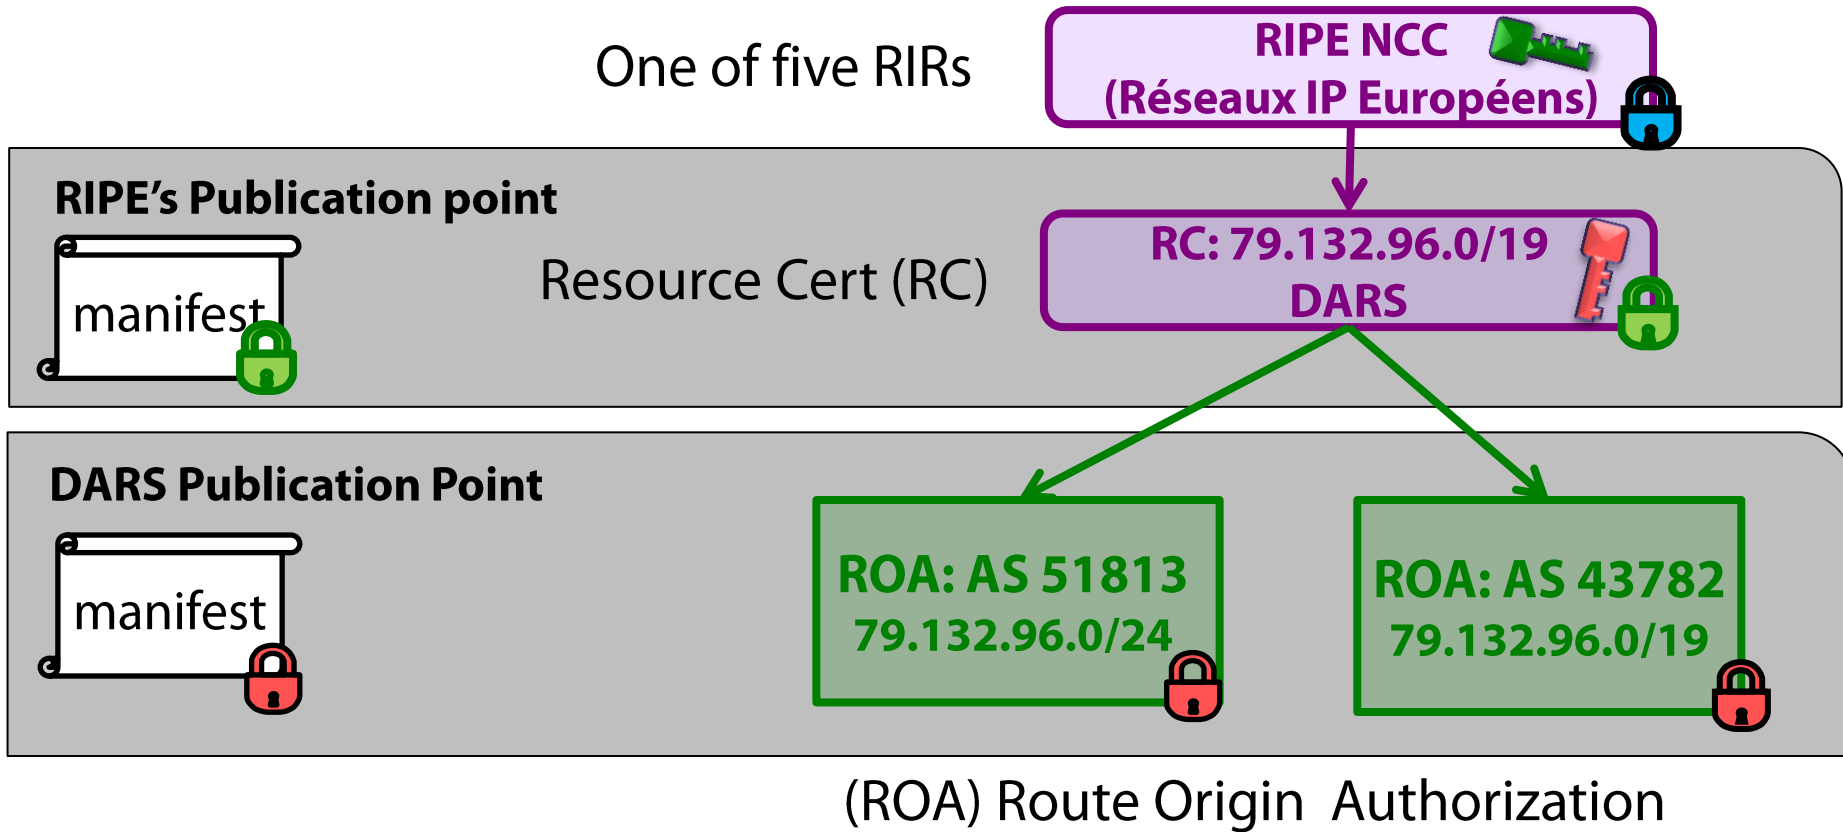
1. RPKI authorities can whack ROAs
2. Whacked ROAs can cause BGP routes to become **invalid**
3. Should drop **invalid** BGP routes to stop **sub**prefix hijacks.



Proposal to require consent to whack objects [SIGCOMM'14]

- There is a draft for similar proposal: [\[draft-kent-sidr-suspenders-02\]](#)

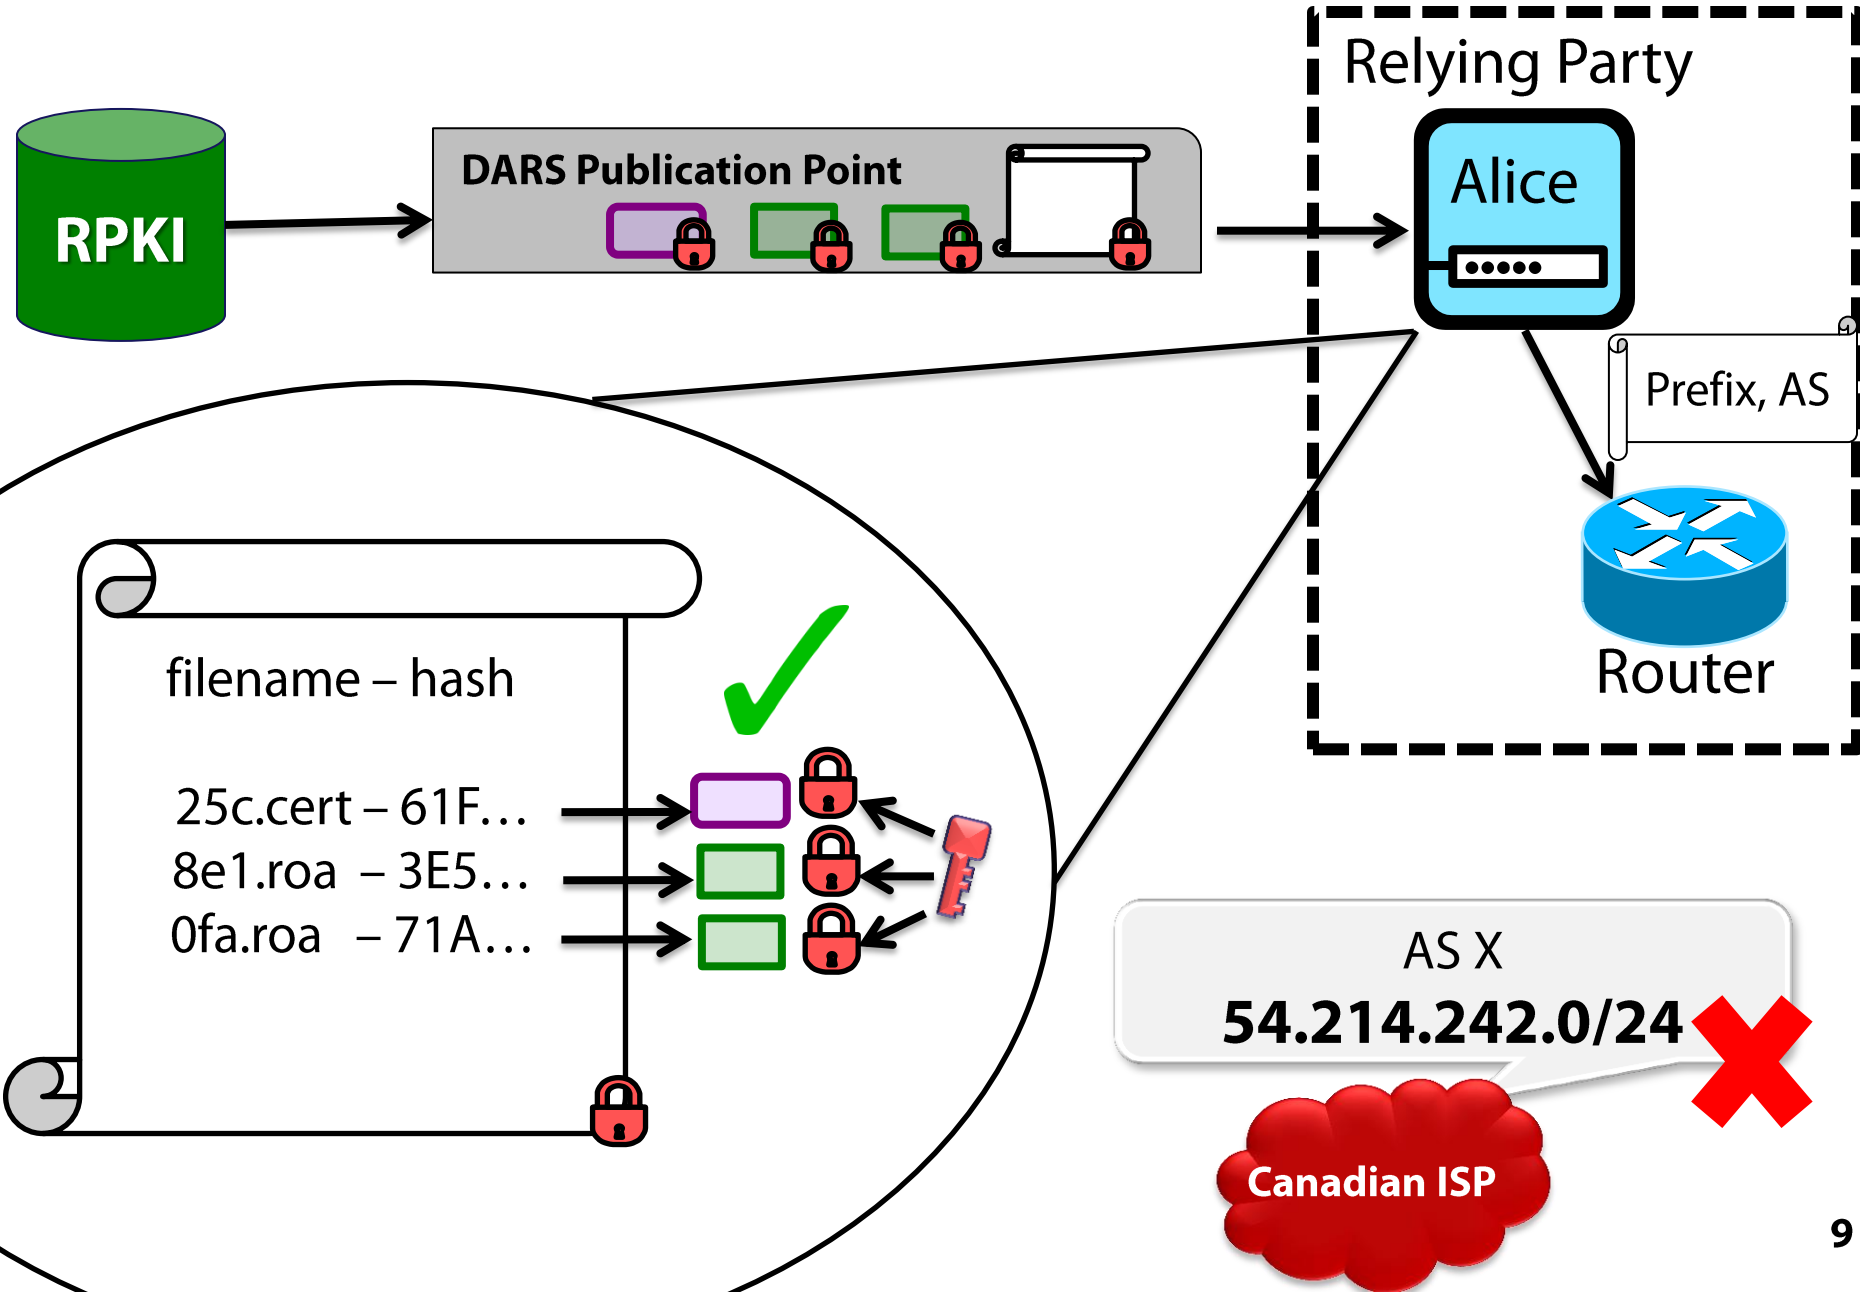
structure of the RPKI [RFC 6480]



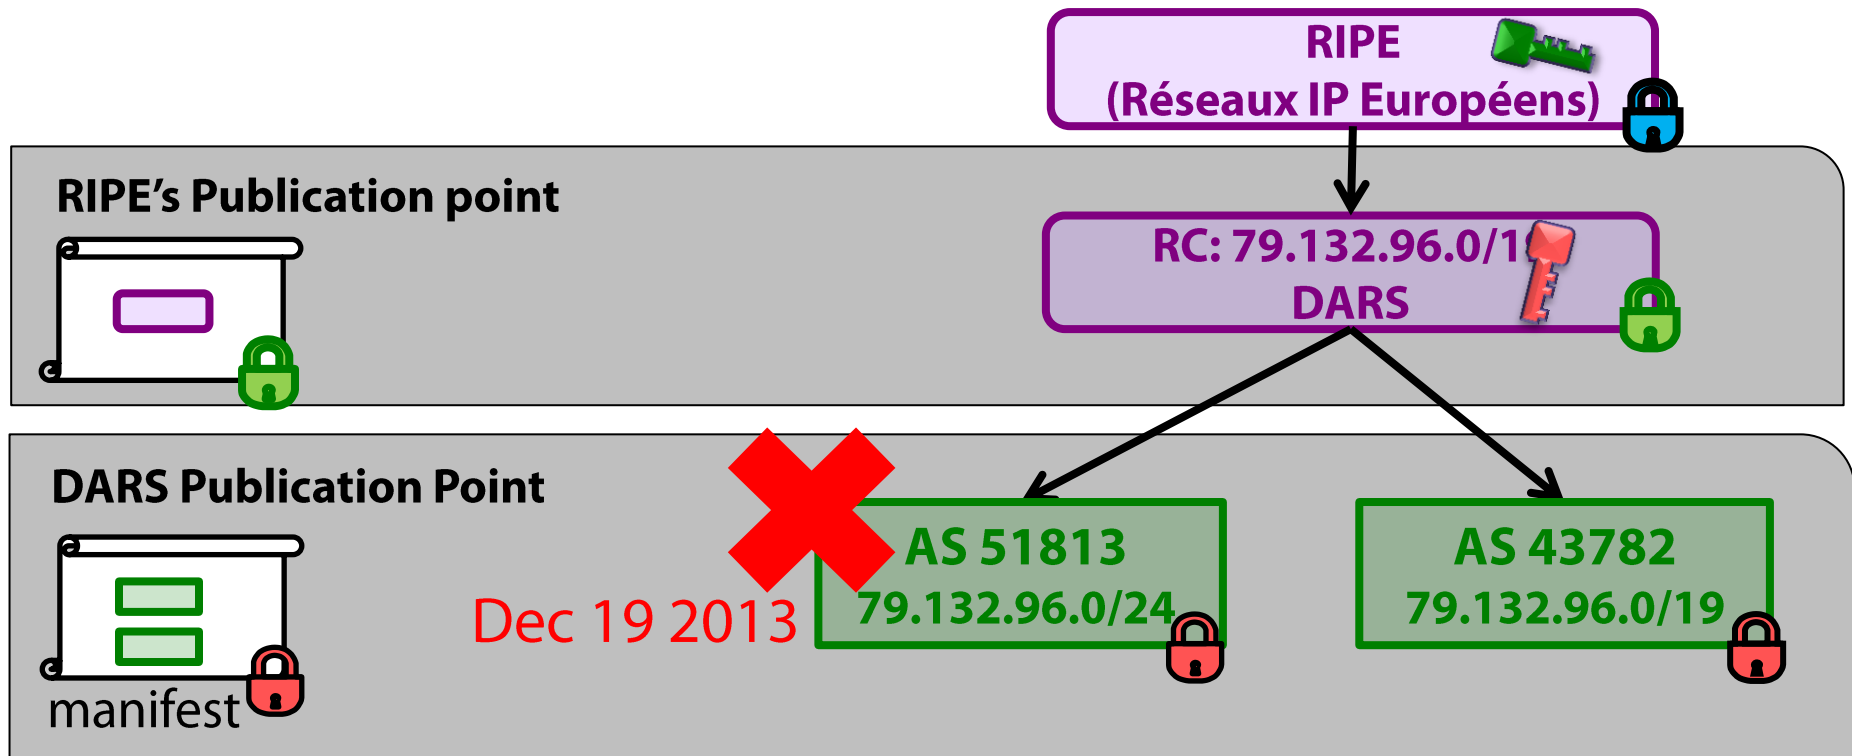
Deployment Status of the RPKI:

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

how relying parties sync to the RPKI [RFC 6480]



issue 1: RPKI authorities can **unilaterally** whack ROAs



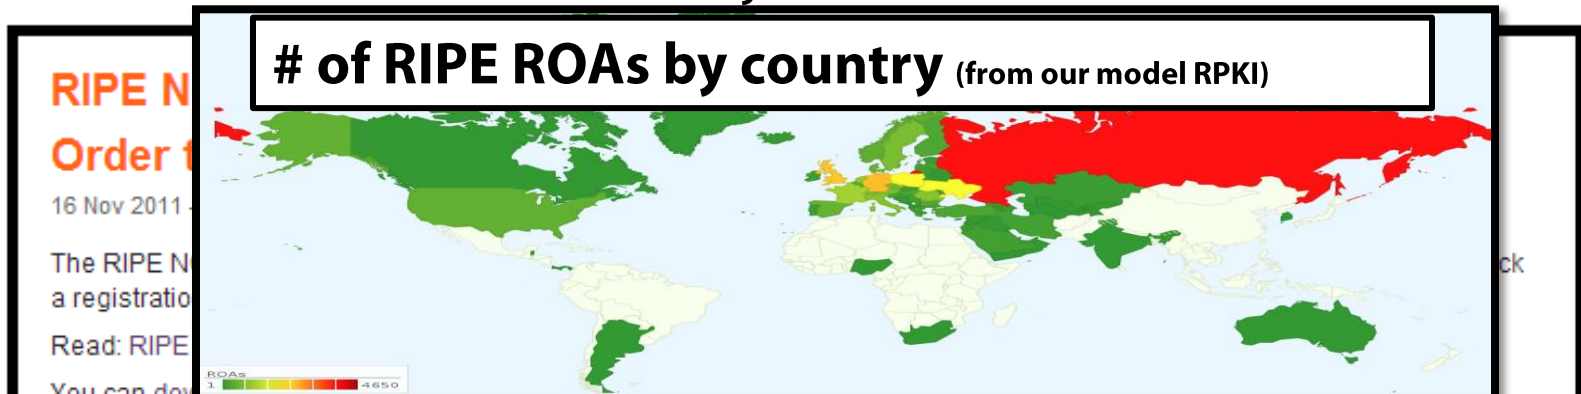
If routers "drop invalid" routes, they could lose connectivity to a legitimate route!

AS 51813
79.132.96.0/24

AS51813

IP prefix takedowns by deleting ROAs?

- Prior to the RPKI, authorities could allocate IPs but not revoke them.
- But RPKI authorities **can** revoke IP 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 lockdown prefixes registration (Nov'11)
 - US court issued a writ of attachment on Iran's IP prefixes (June'14)
 - IP allocation does not reflect jurisdiction.

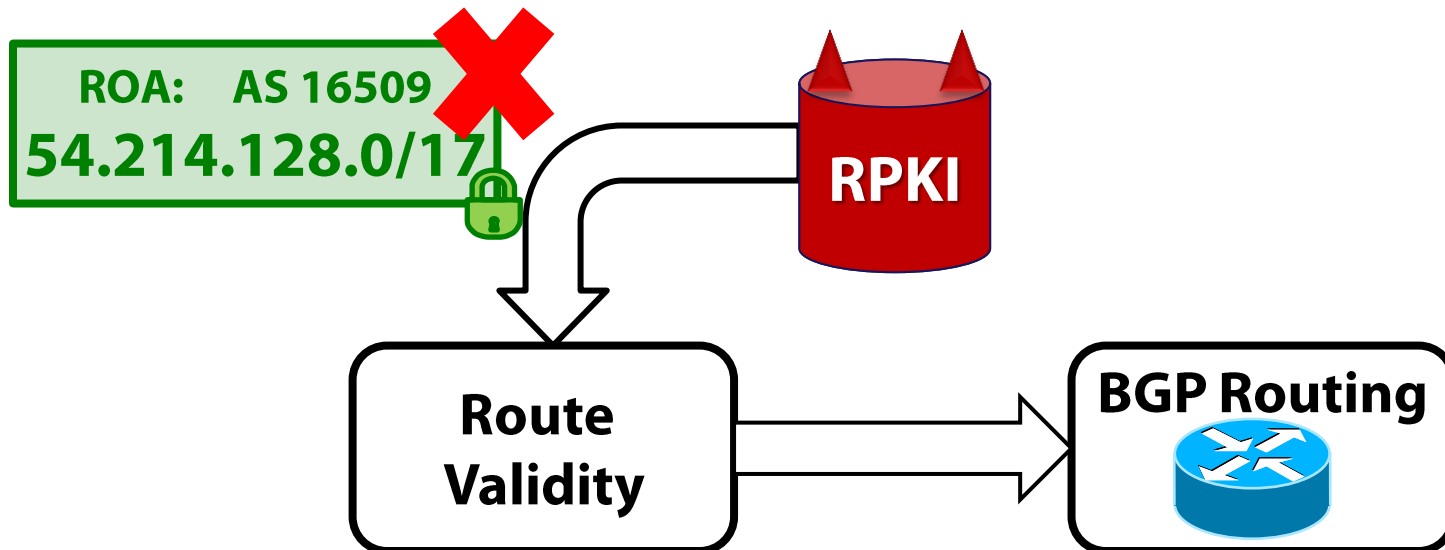


talk outline

Security audit of the RPKI [HotNets'13]

Misbehaving RPKI authorities can blackhole routes in BGP. Why?

1. RPKI authorities can whack ROAs
2. Whacked ROAs can cause BGP routes to become **invalid**



issue 2: whacked ROAs can cause BGP routes to be invalid

- valid BGP route
- invalid BGP route
- unknown BGP route

← "World before RPKI"

Reality: interdependent validity outcomes



AS 16509  valid
54.214.128.0/17

Amazon
AS 16509

AS X 
54.214.242.0/24

Canadian
AS

AS 16509
54.214.128.0/17 

RPKI  13

issue 2: whacked ROAs can cause BGP routes to be invalid



← “World before RPKI”

Reality: interdependent validity outcomes



AS 43782 ✓ valid
79.132.96.0/19

DARS
AS 43782

AS 51813 ✓ valid
79.132.96.0/24

Dartel
AS 51813

AS 43782
79.132.96.0/19

AS 51813
79.132.96.0/24


RPKI 14


issue 2: whacked ROAs can cause BGP routes to be invalid





← "World before RPKI"

Reality: interdependent validity outcomes



AS 43782  valid
79.132.96.0/19
DARS
AS 43782

AS 51813 
79.132.96.0/24
Dartel
AS 51813


AS 43782
79.132.96.0/19 
RPKI
14

issue 2: whacked ROAs can cause BGP routes to be invalid



← “World before RPKI”

Reality: interdependent validity outcomes

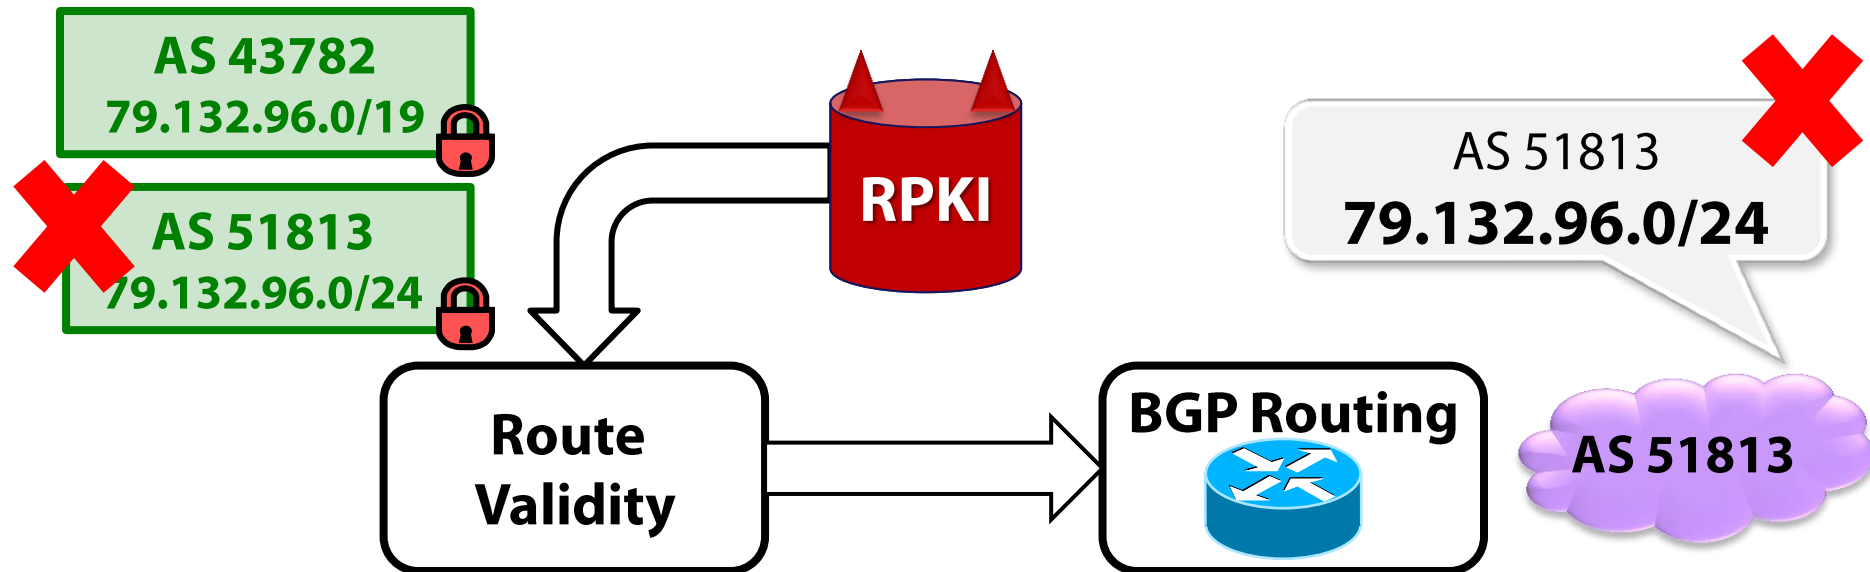


talk outline

Security audit of the RPKI [HotNets'13]

Misbehaving RPKI authorities can blackhole routes in BGP. Why?

1. RPKI authorities can whack ROAs
2. Whacked ROAs can cause BGP routes to become **invalid**
3. Should drop **invalid** BGP routes to stop **sub**prefix hijacks.



Proposal to require consent for whacked objects [SIGCOMM'14]

- There is a draft for similar proposal: [\[draft-kent-sidr-suspenders-02\]](#) 15




proposal: require consent to whack objects [SIGCOMM'14]

- **Design goals:**

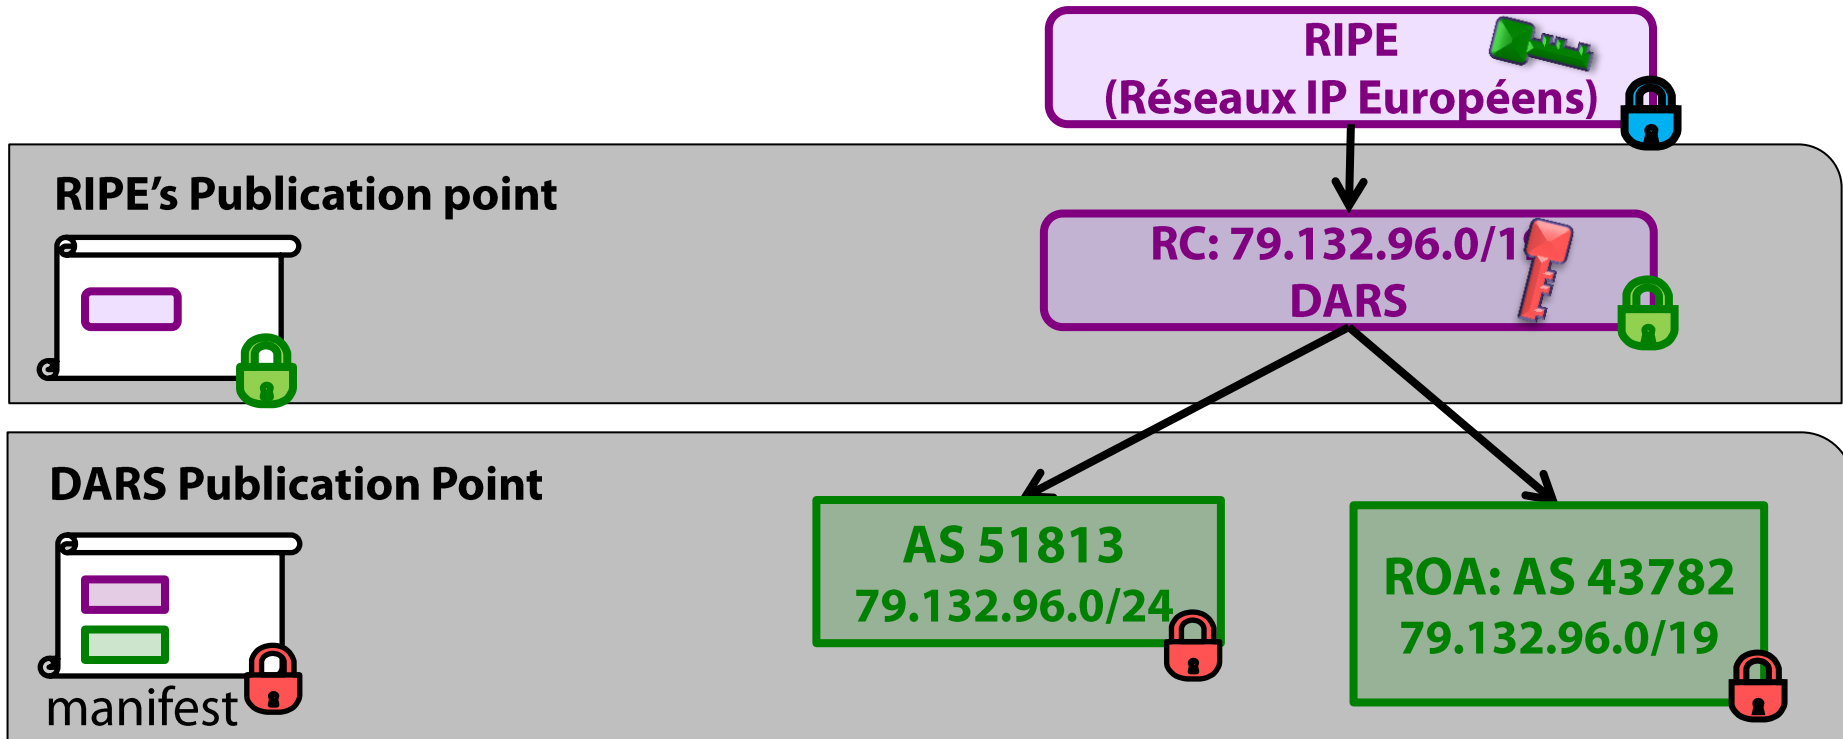
- **Consent:** Resource certs (RCs) must consent to be whacked.
- **Transparency:** Relying parties audit RPKI & alarm on problems.
- **Consistency:** Relying parties have consistent views of the RPKI.

- **Threat Model:**

- Similar to certificate transparency [RFC 6962]
- Relying parties honestly audit the RPKI
- Everyone else (incl. RPKI authorities) is untrusted



how to consent? introducing .dead objects

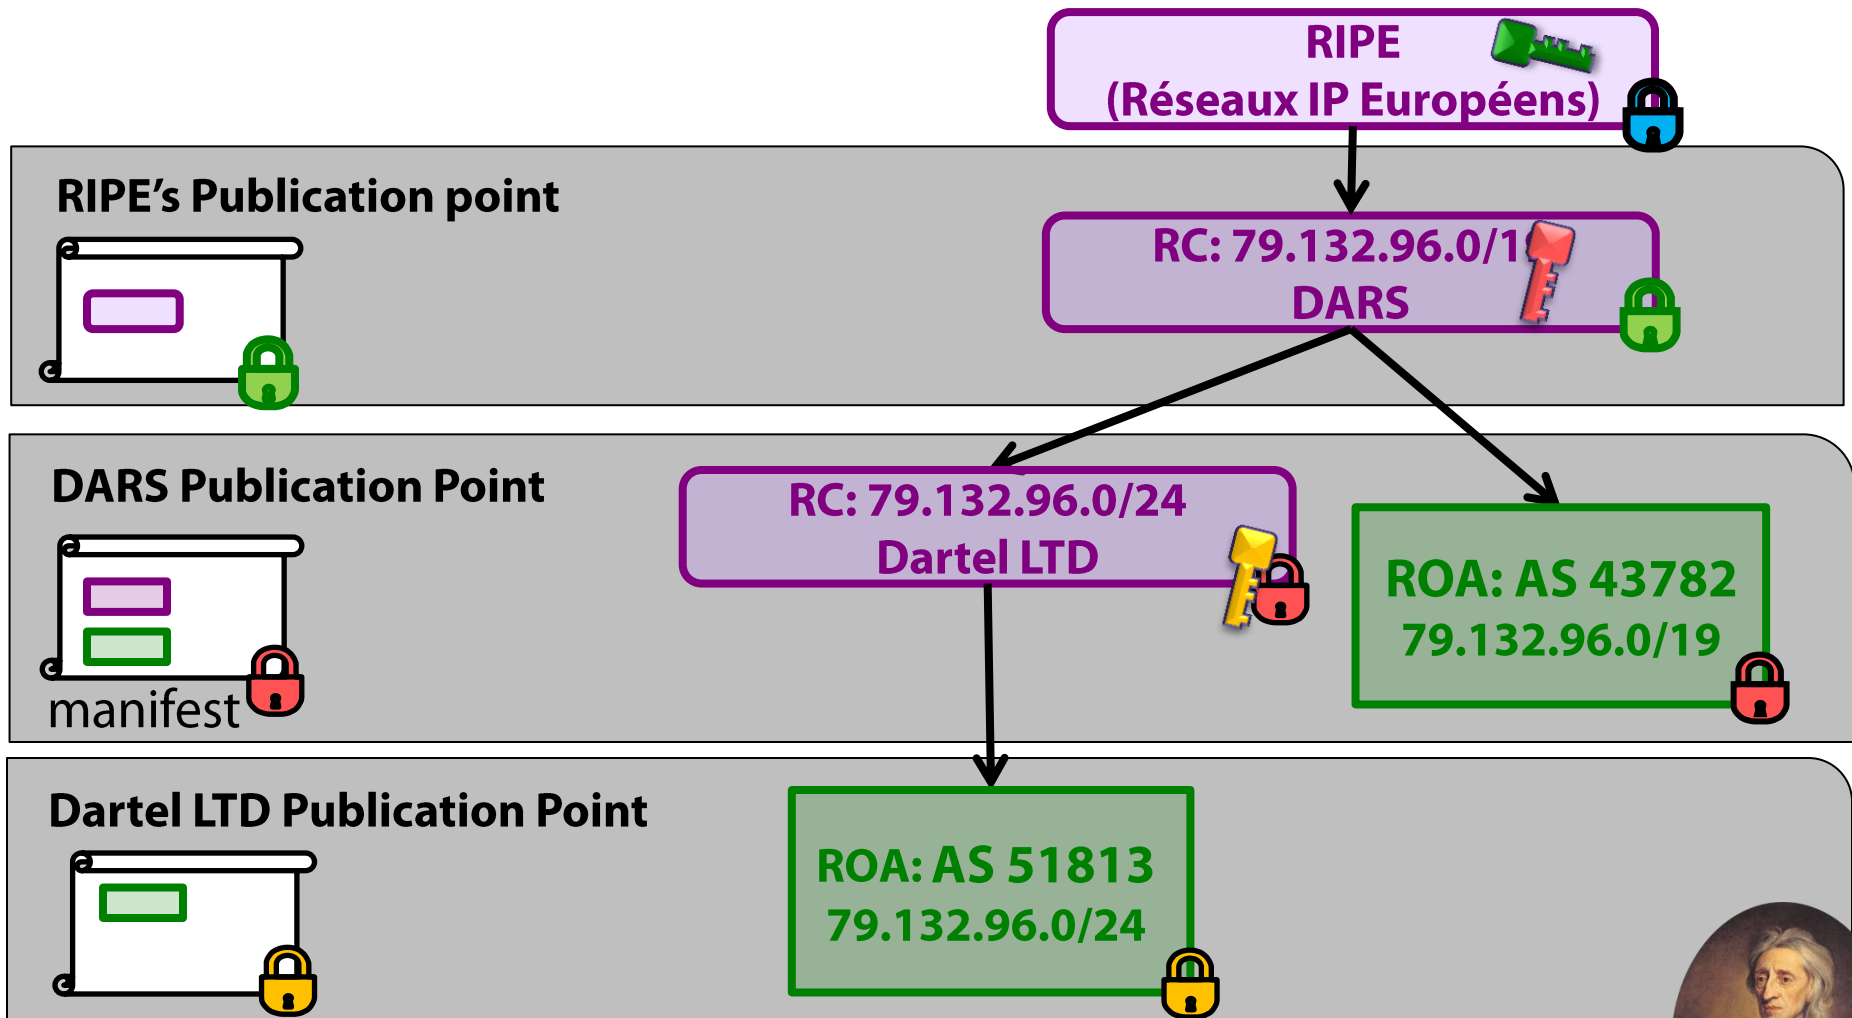


If an authority wants to revoke IP prefixes from a child RC, it needs consent from that child & its impacted* descendant RCs.

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




how to consent? introducing .dead objects

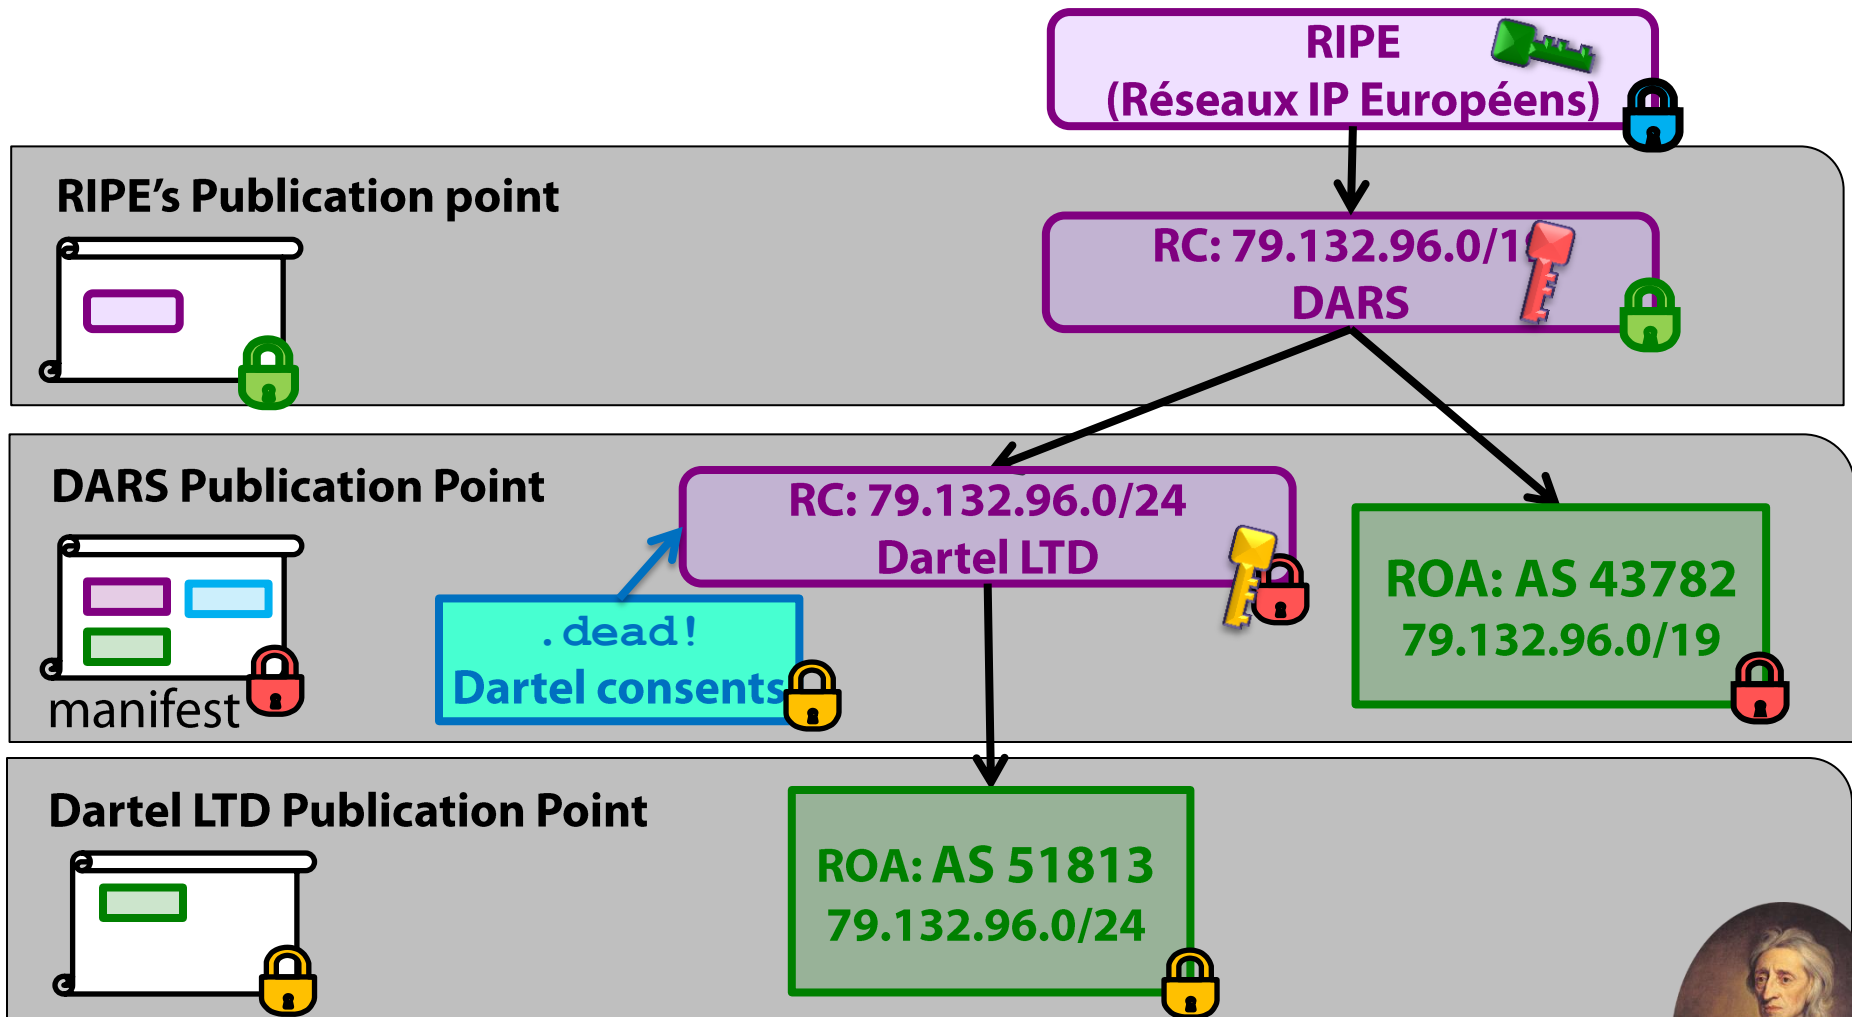


If an authority wants to revoke IP prefixes from a child RC, it needs consent from that child & its impacted* descendant RCs.

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




how to consent? introducing **.dead** objects

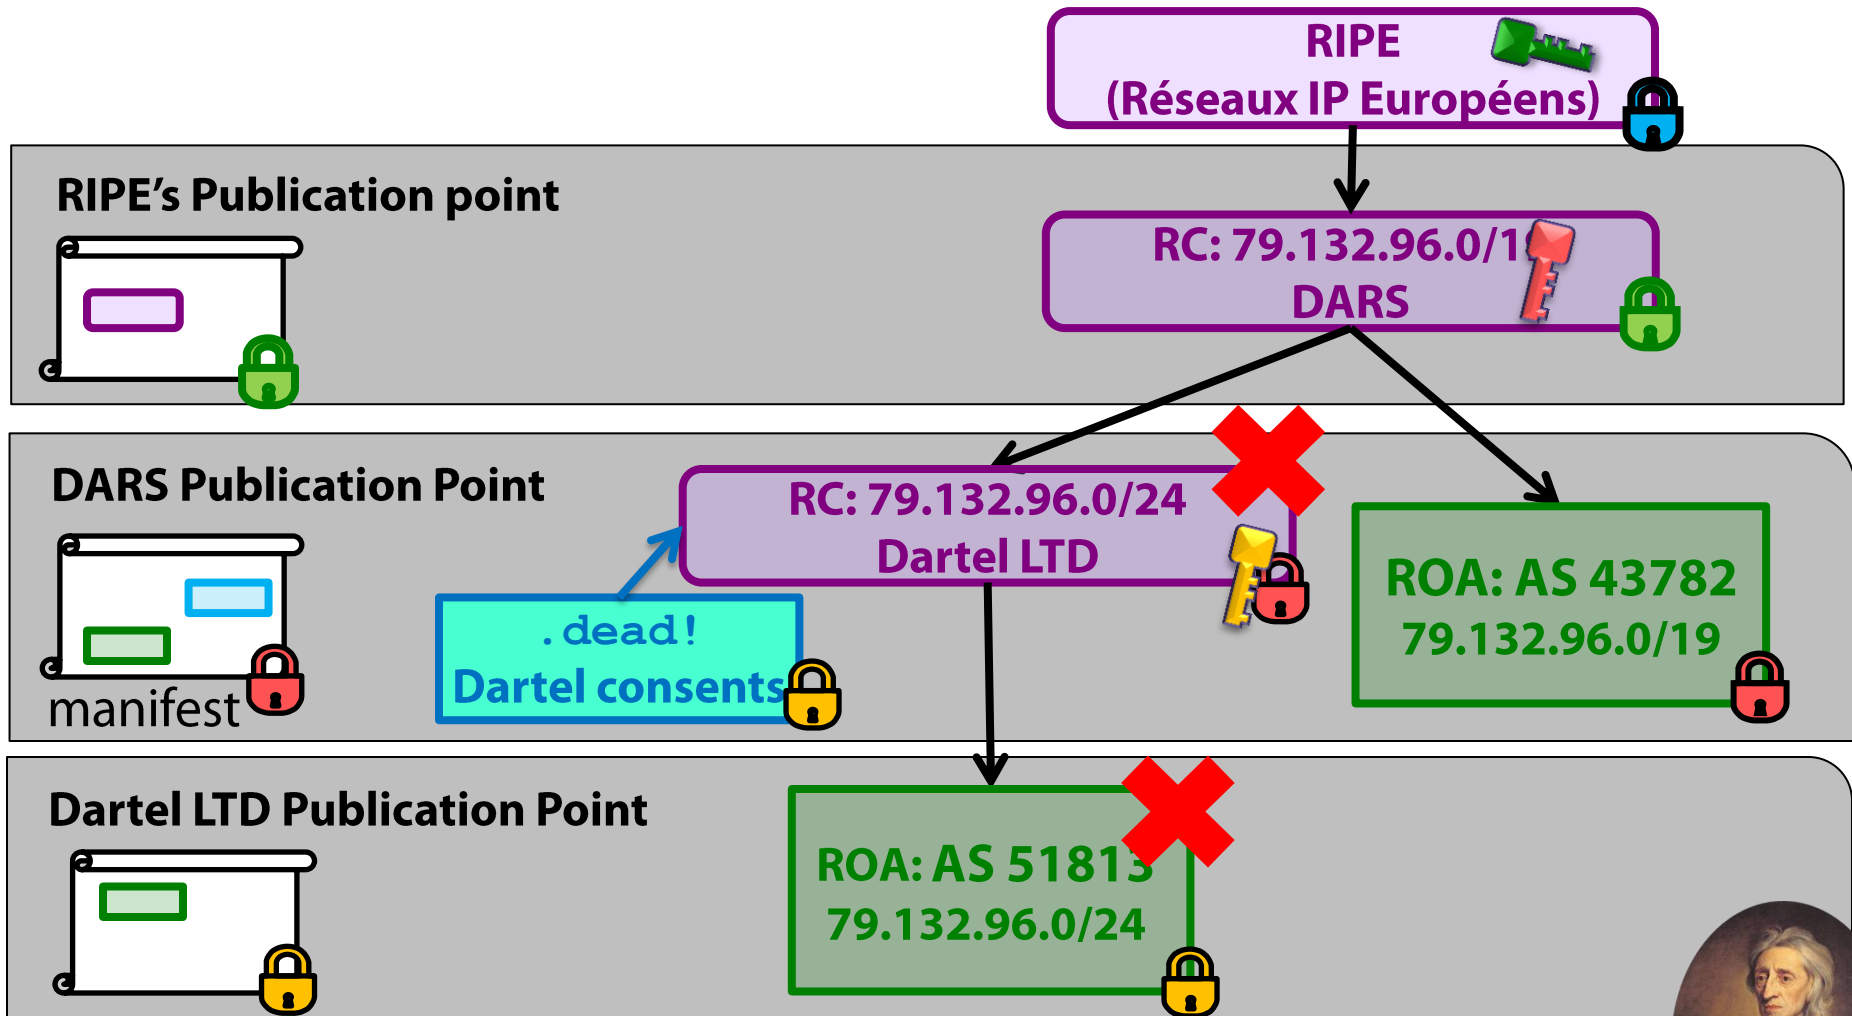


If an authority wants to revoke IP prefixes from a child RC, it needs consent from that child & its impacted* descendant RCs.

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




how to consent? introducing .dead objects

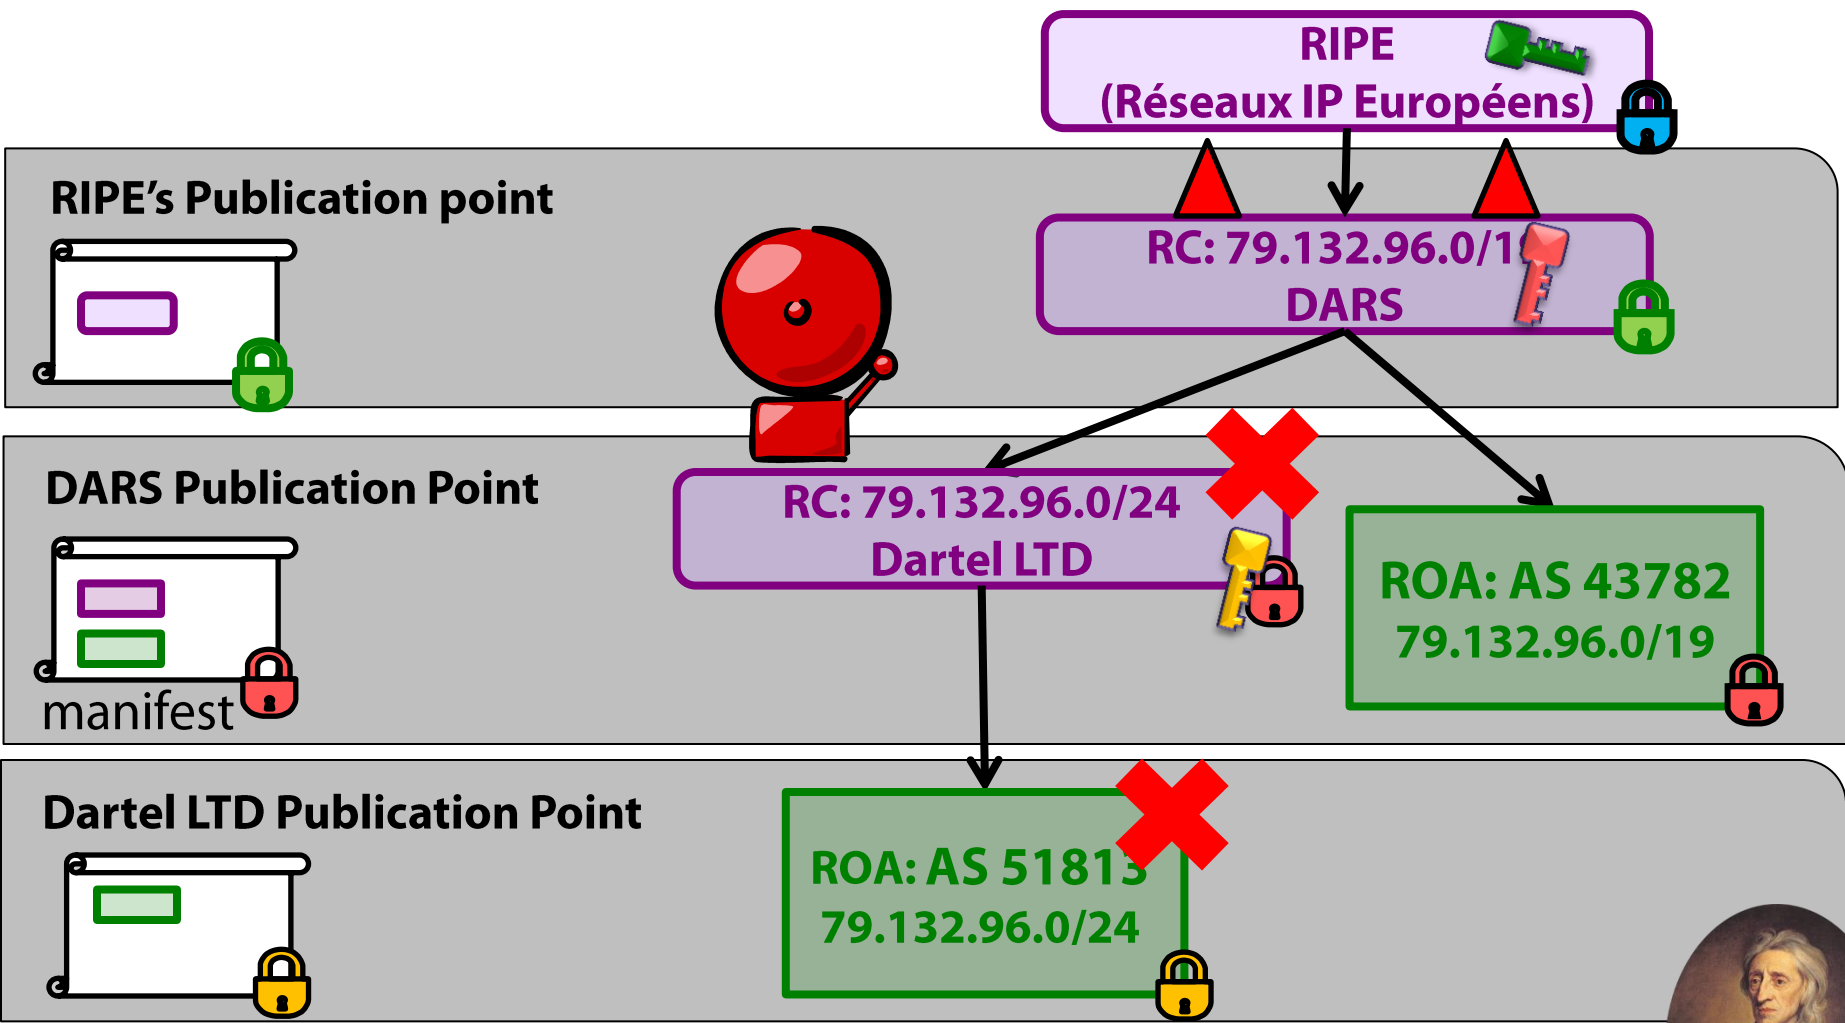


If an authority wants to revoke IP prefixes from a child RC, it needs consent from that child & its impacted* descendant RCs.

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




how to consent? introducing .dead objects



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 morning & misses violations between syncs!

Why does Alice need to catch violations between syncs?

- So Alice can audit the RPKI
- So we can have consistency (explained later)

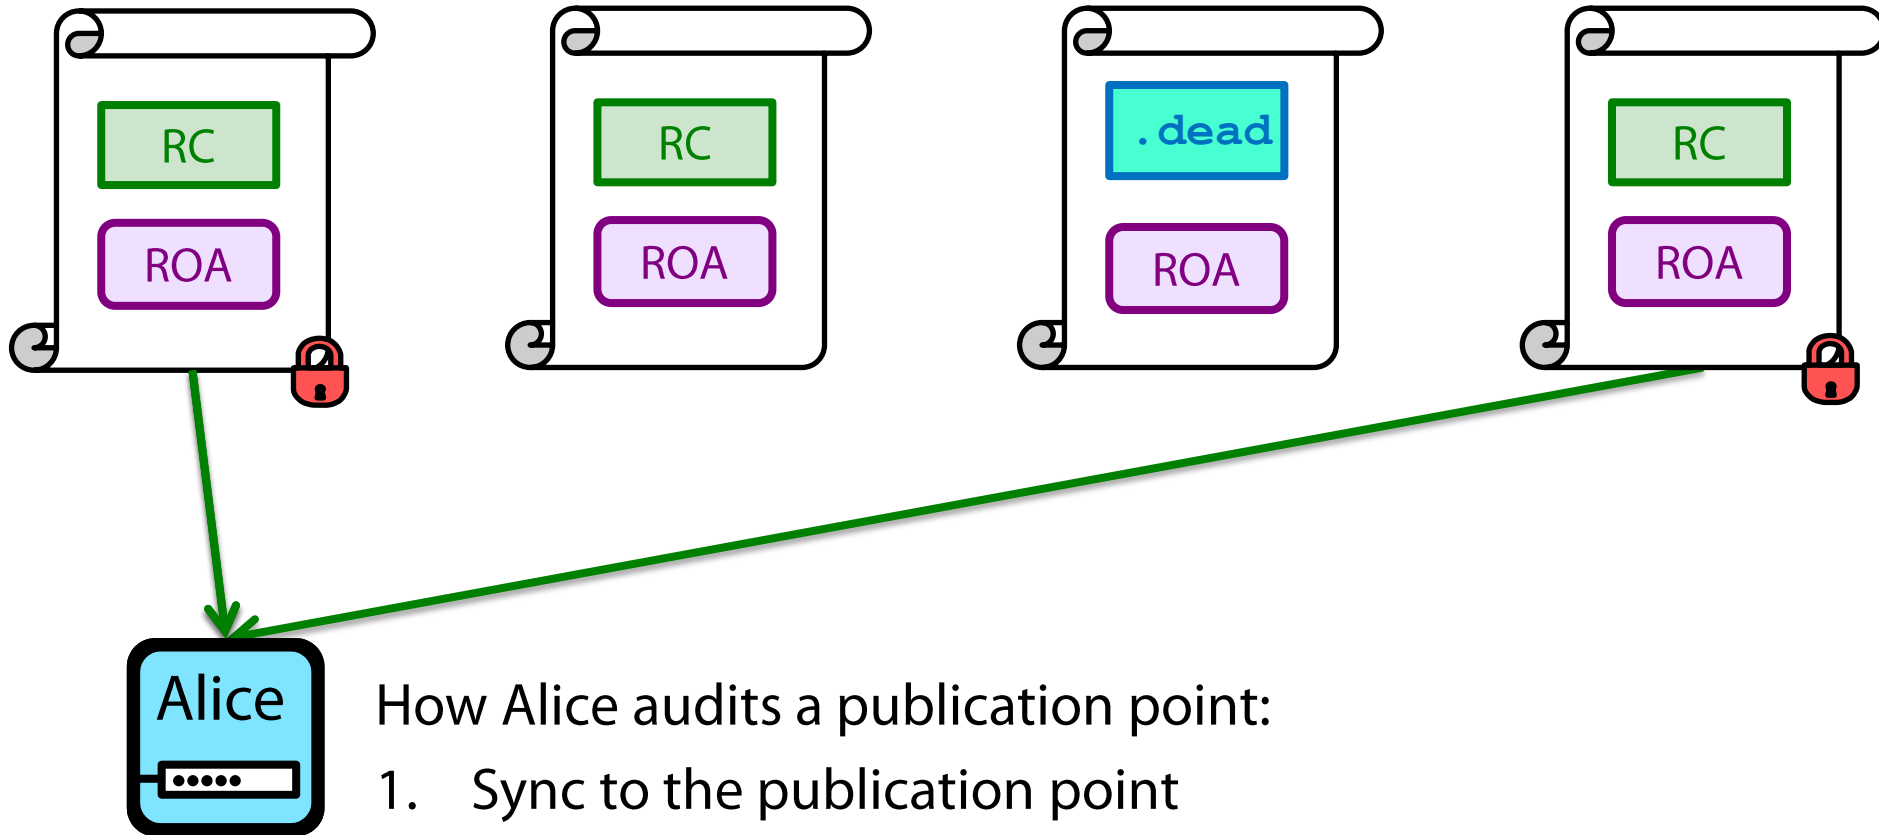
catching alarms between syncs!



How Alice audits a publication point:

1. Sync to the publication point
2. Use change log to reconstruct intermediate manifests


catching alarms between syncs!



How Alice audits a publication point:

1. Sync to the publication point
2. Use change log to reconstruct intermediate manifests


catching alarms between syncs!



How Alice audits a publication point:

1. Sync to the publication point
2. Use change log 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 audits a publication point:

1. Sync to the publication point
2. Use change log 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 audits a publication point:


1. Sync to the publication point
2. Use change log to reconstruct intermediate manifests

Valid Remains Valid. Our auditing algorithm makes sure that once a relying party has seen a valid resource cert (RC), that RC remains valid until it consents to be deleted/modified.

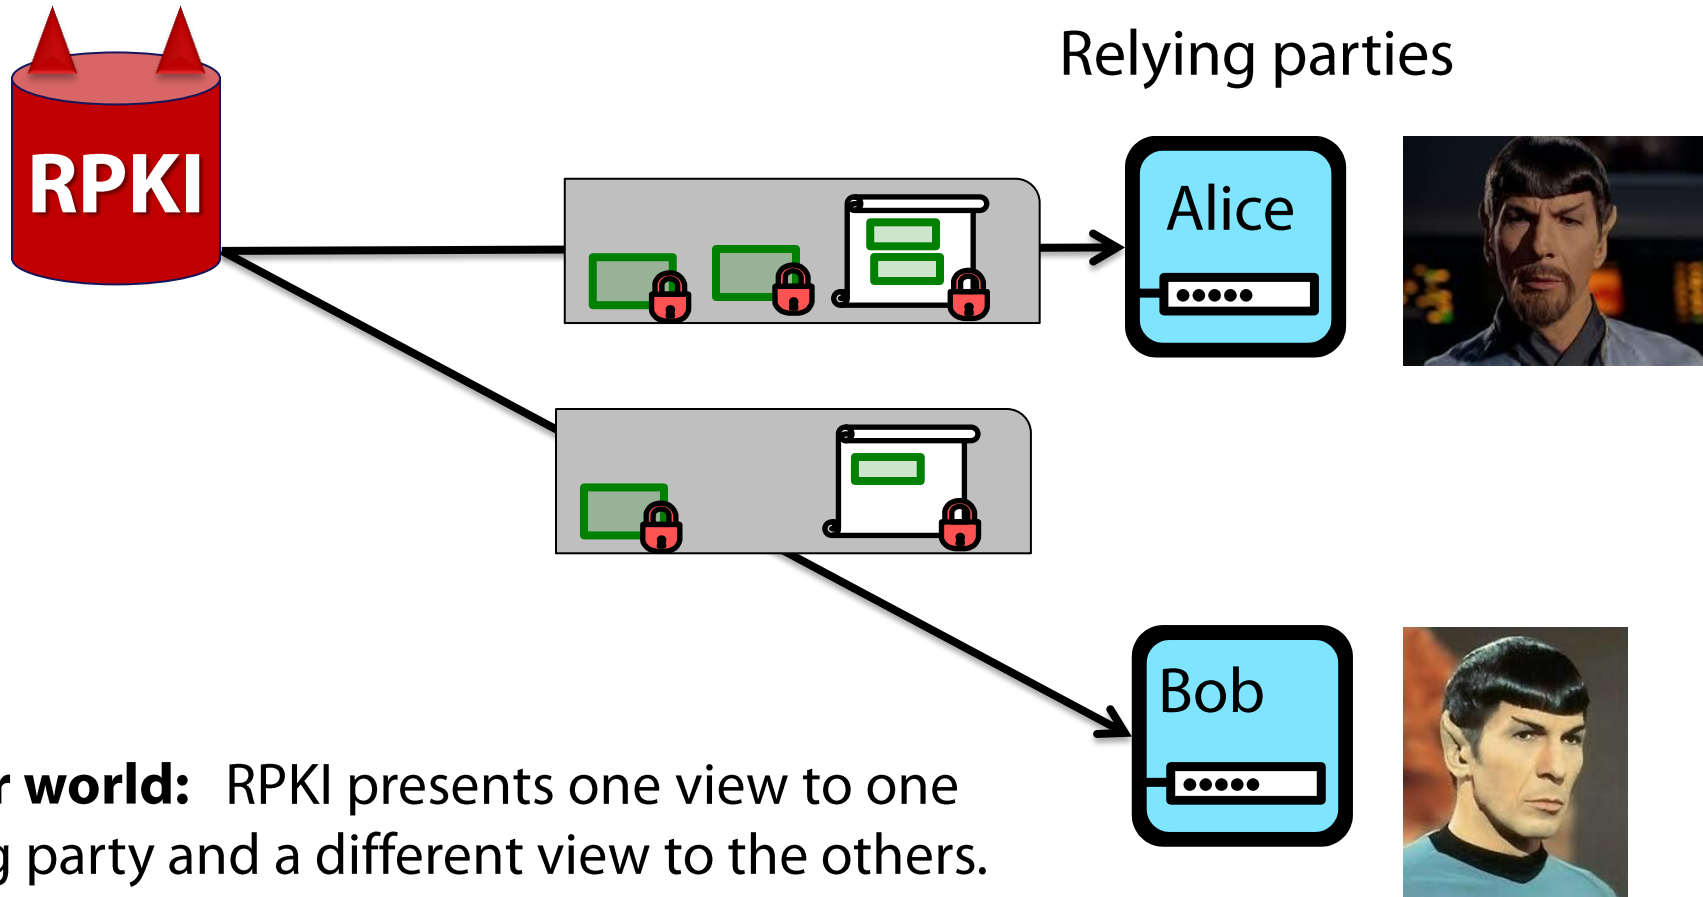
proposal : require consent to delete objects [SIGCOMM'14]

- **Design goals:**

- ✓ – **Consent:** **.dead** objects indicate consent to whack resource certs (RCs)
 - **Consistency:** Relying parties have consistent views of the RPKI.
- ✓ – **Transparency:** Relying parties audit RPKI & alarm on problems.
 - “Drop invalid” for prefixes that are not part of an alarm
 - Manually audit prefixes that are part of an alarm.



mirror worlds: inconsistent views of the RPKI






Mirror world: RPKI presents one view to one relying party and a different view to the others.

Why do we care?

- Auditing is less meaningful if Alice's view is different from everyone else's.
- Eg. Suppose Alice audits the RPKI to make sure her own ROAs are OK.

detecting mirror worlds using manifest hash chains



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

No mirror worlds!

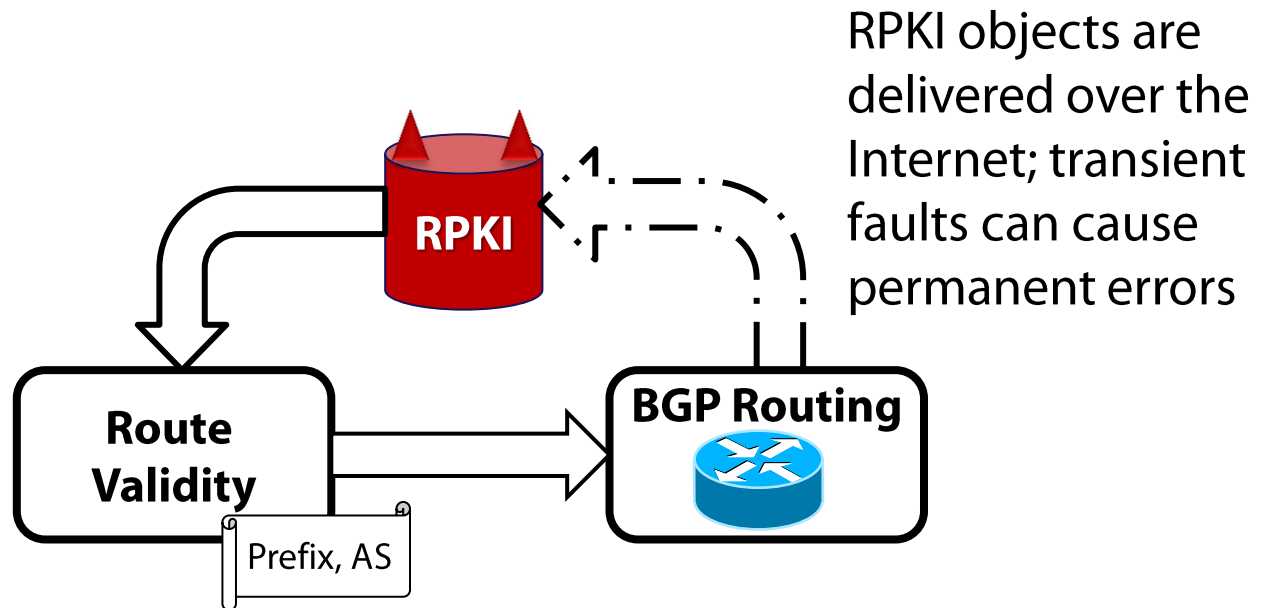
If the consistency check passes, relying parties saw the same valid objects.

outline

✓ Part 1: security audit of the RPKI [HotNets'13]

we need to harden the RPKI's delivery mechanism!

- eg. expired [draft-ietf-sidr-multiple-publication-points-01]




✓ Part 2: proposal to improve RPKI transparency [SIGCOMM'14]


- There is a draft for similar proposal: [draft-kent-sidr-suspenders-02]

conclusion: more work needed

1. Robust delivery of RPKI objects from repos. **[To do!]**
2. Prevent misconfigured ROAs using config tools:
 1. RIPE RPKI management UI: <http://localcert.ripe.net:8088/>
 2. NIST RPKI deployment monitor: <http://rpki-monitor.antd.nist.gov/>
 3. RTRlib: <http://page.mi.fu-berlin.de/waehl/publications/whss-roslr-13.html>
 4. LACNIC RPKI looking glass: www.labs.lacnic.net/rpkitools/looking_glass/
 5. rcynic web interface
3. Limit risk of unilateral RPKI takedowns. **[To do!]**
 1. Our proposal **[SIGCOMM'14]**
 2. **[draft-kent-sidr-suspenders-02]**
4. React to RPKI alarms with nuanced routing policies. **[To do!]**

Thanks! Project page:
<http://www.cs.bu.edu/~goldbe/papers/RPKImanip.html>



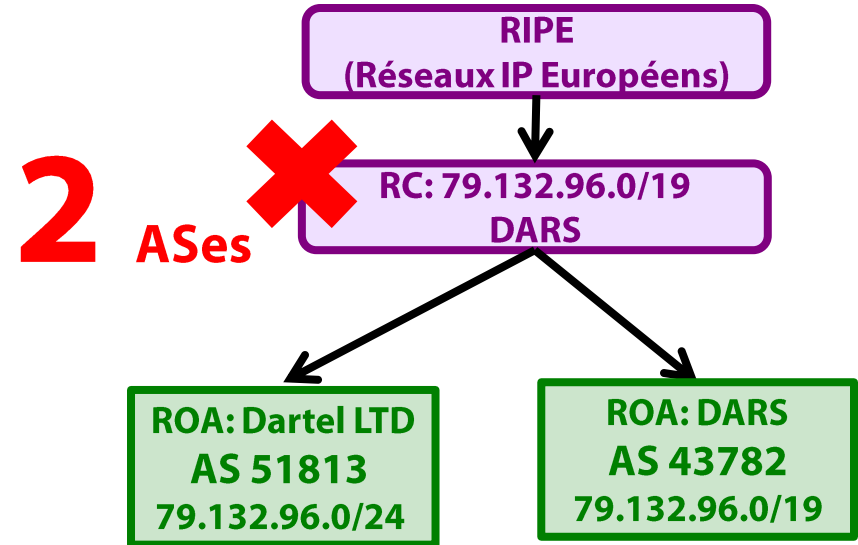


questions



how many parties need to consent?

- How many ASes need to be involved when a leaf resource cert 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 often does would the RPKI need .deads?

Doesn't require a .dead
(874 objects)

Required participation of all impacted ASes
(3,336 objects)

7,779 objects altered in total *



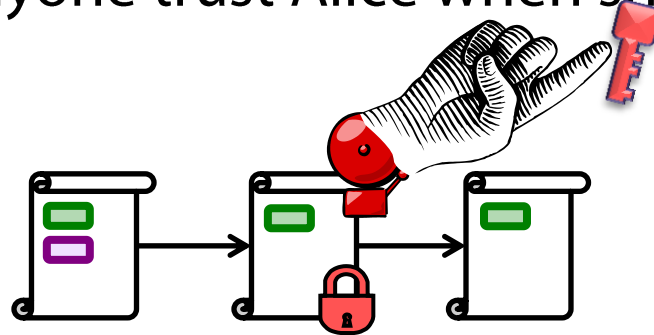
Not needed in our design
(3,569 objects)

Excluding the RIPE restructuring,
only **5%** of cases (230 objects) required a .dead.

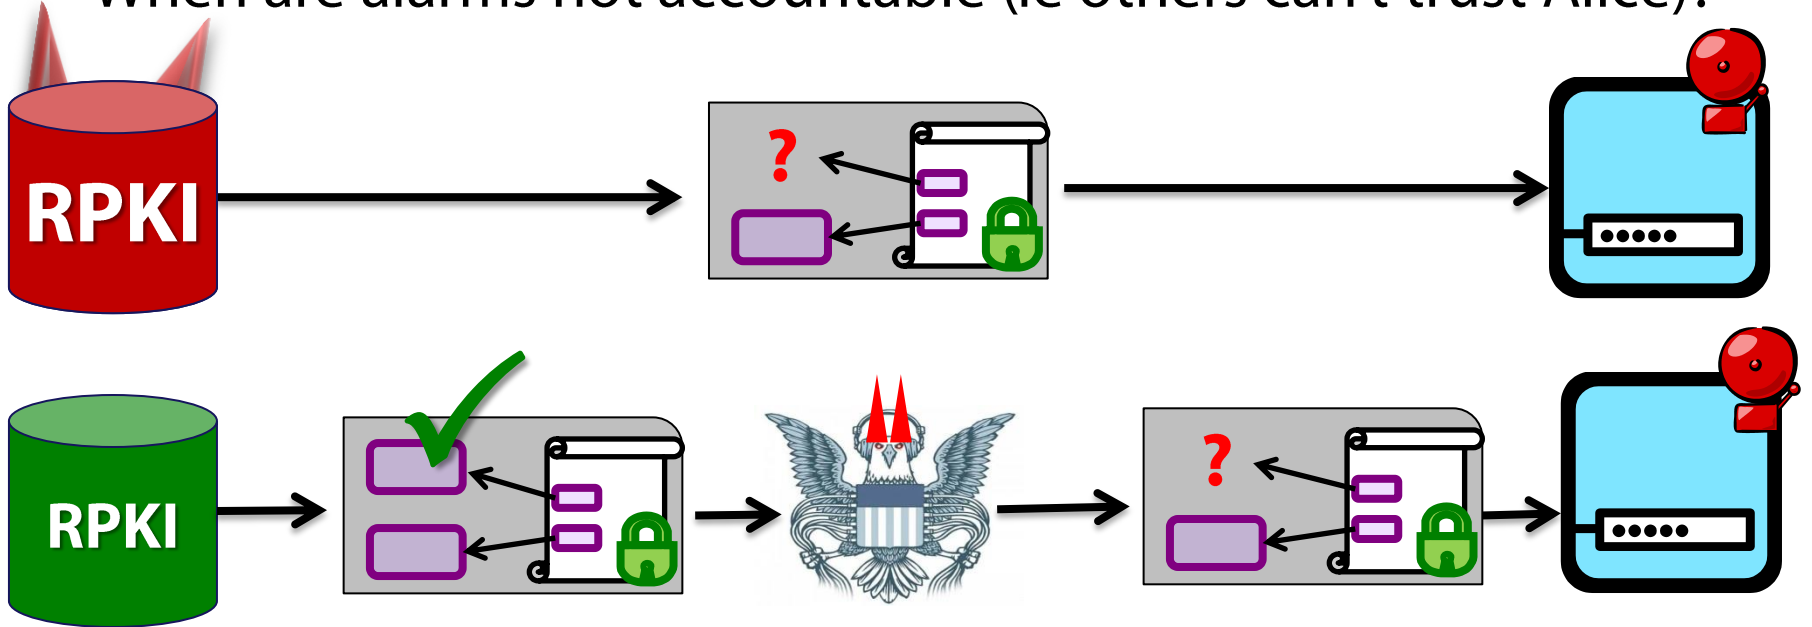
* all data from a ~3 month trace of the taken RPKI 2013/10/23 to 2014/01/21

Blaming authorities with accountable alarms.

- Why should anyone trust Alice when she raises an alarm?

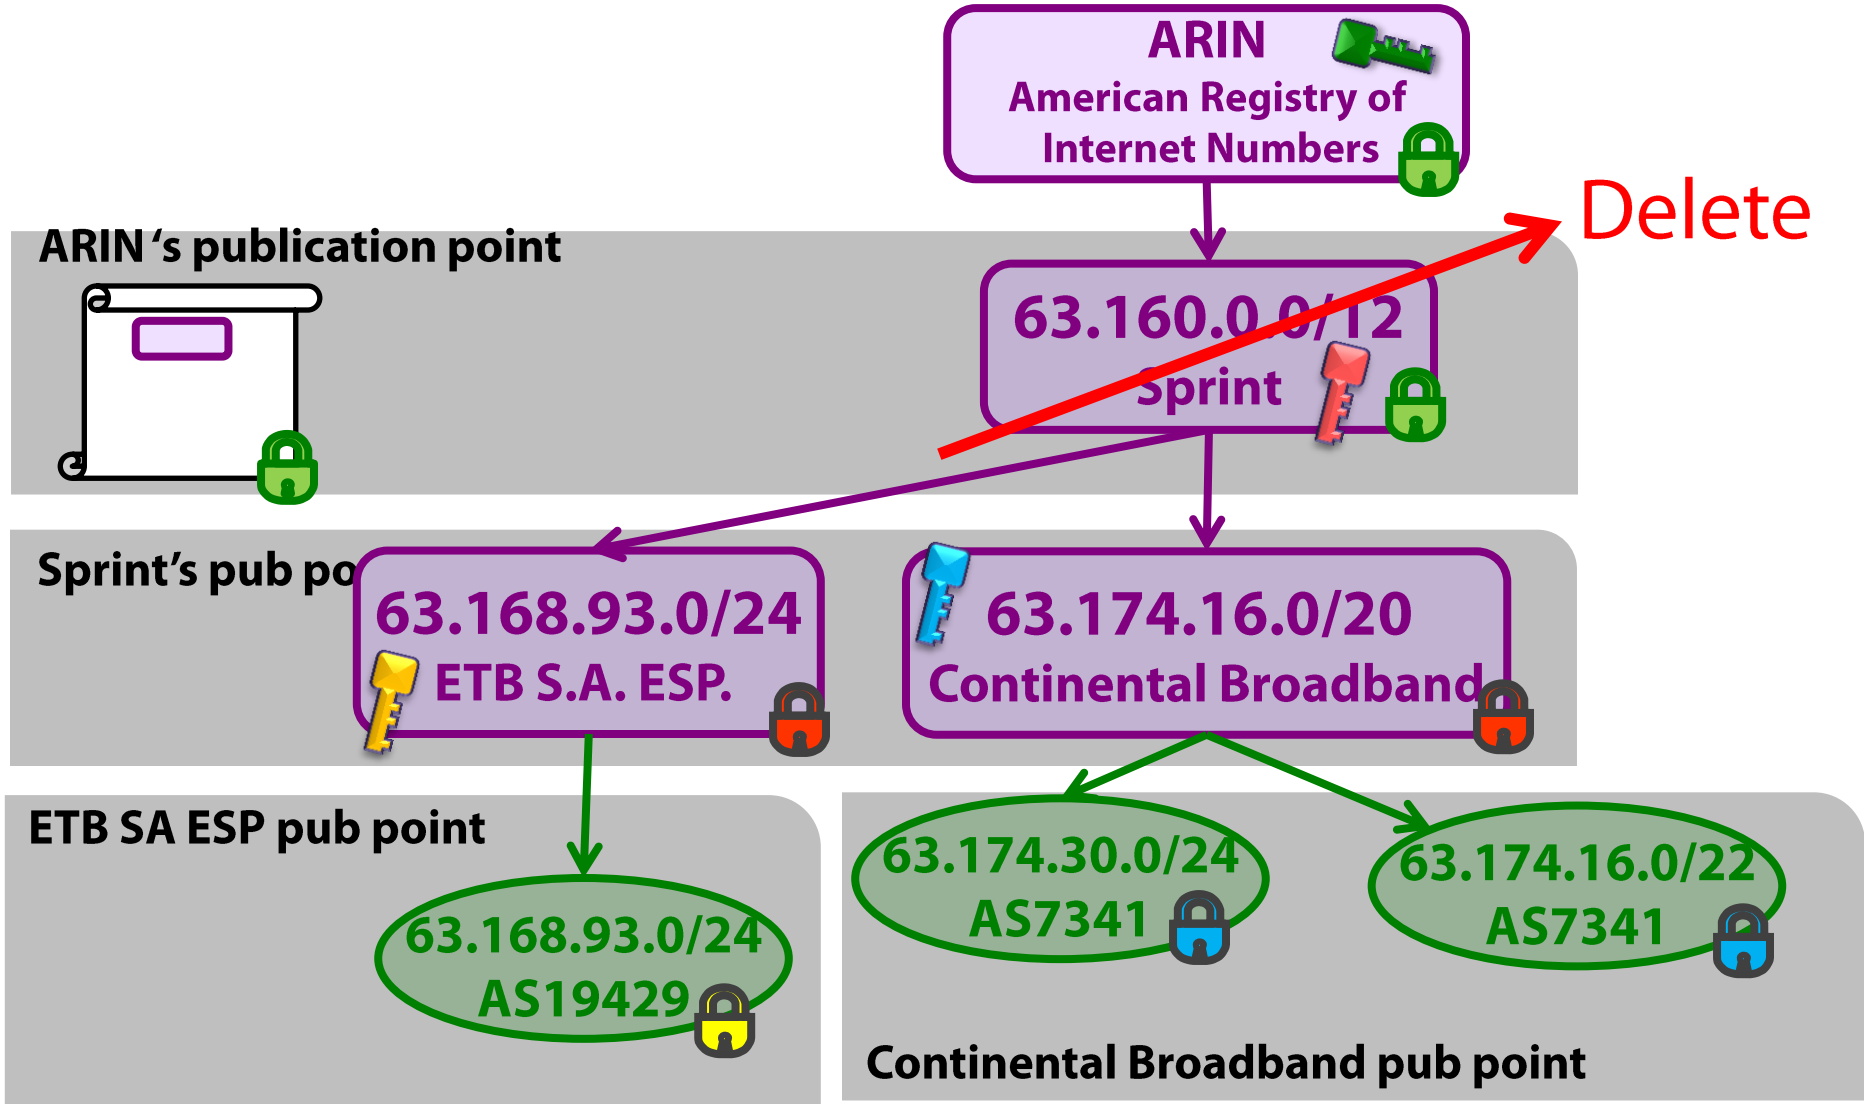


- When are alarms not accountable (ie others can't trust Alice)?

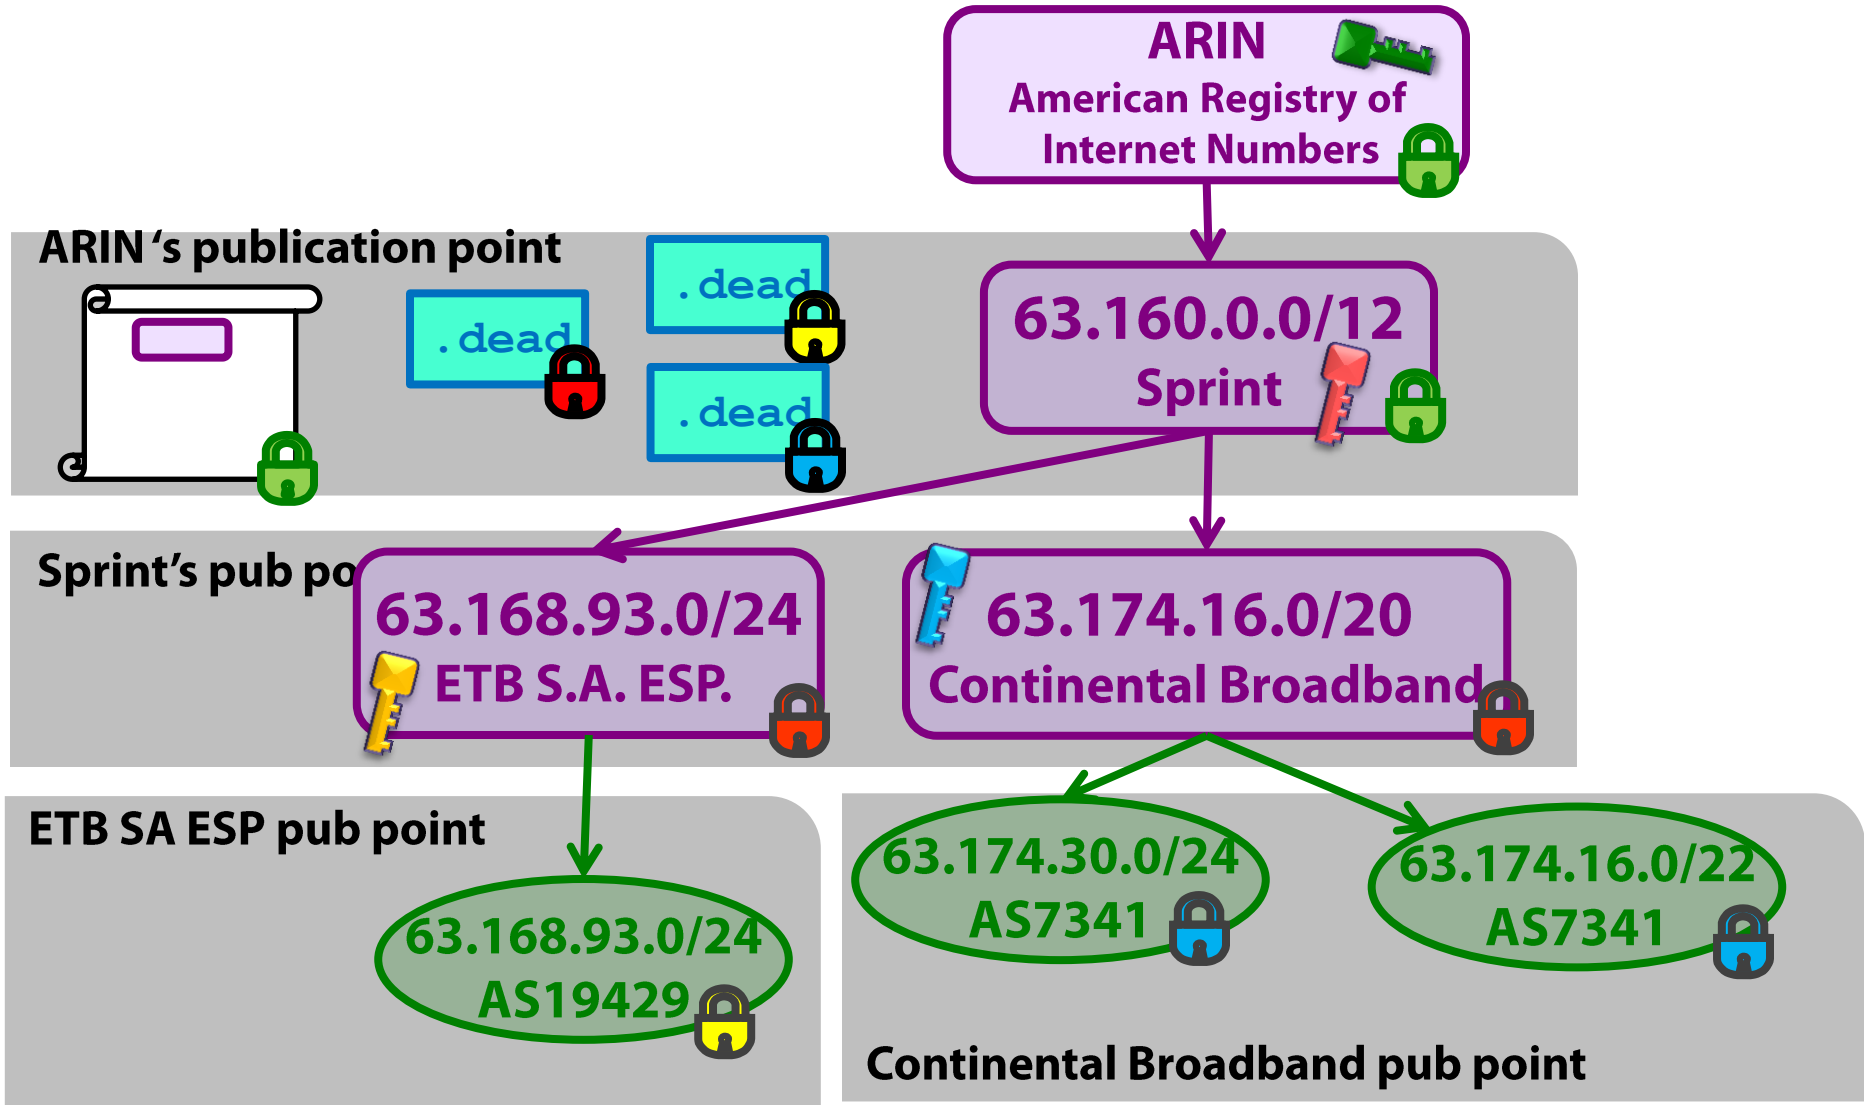


Alarms are accountable in every circumstance other than missing information. 38

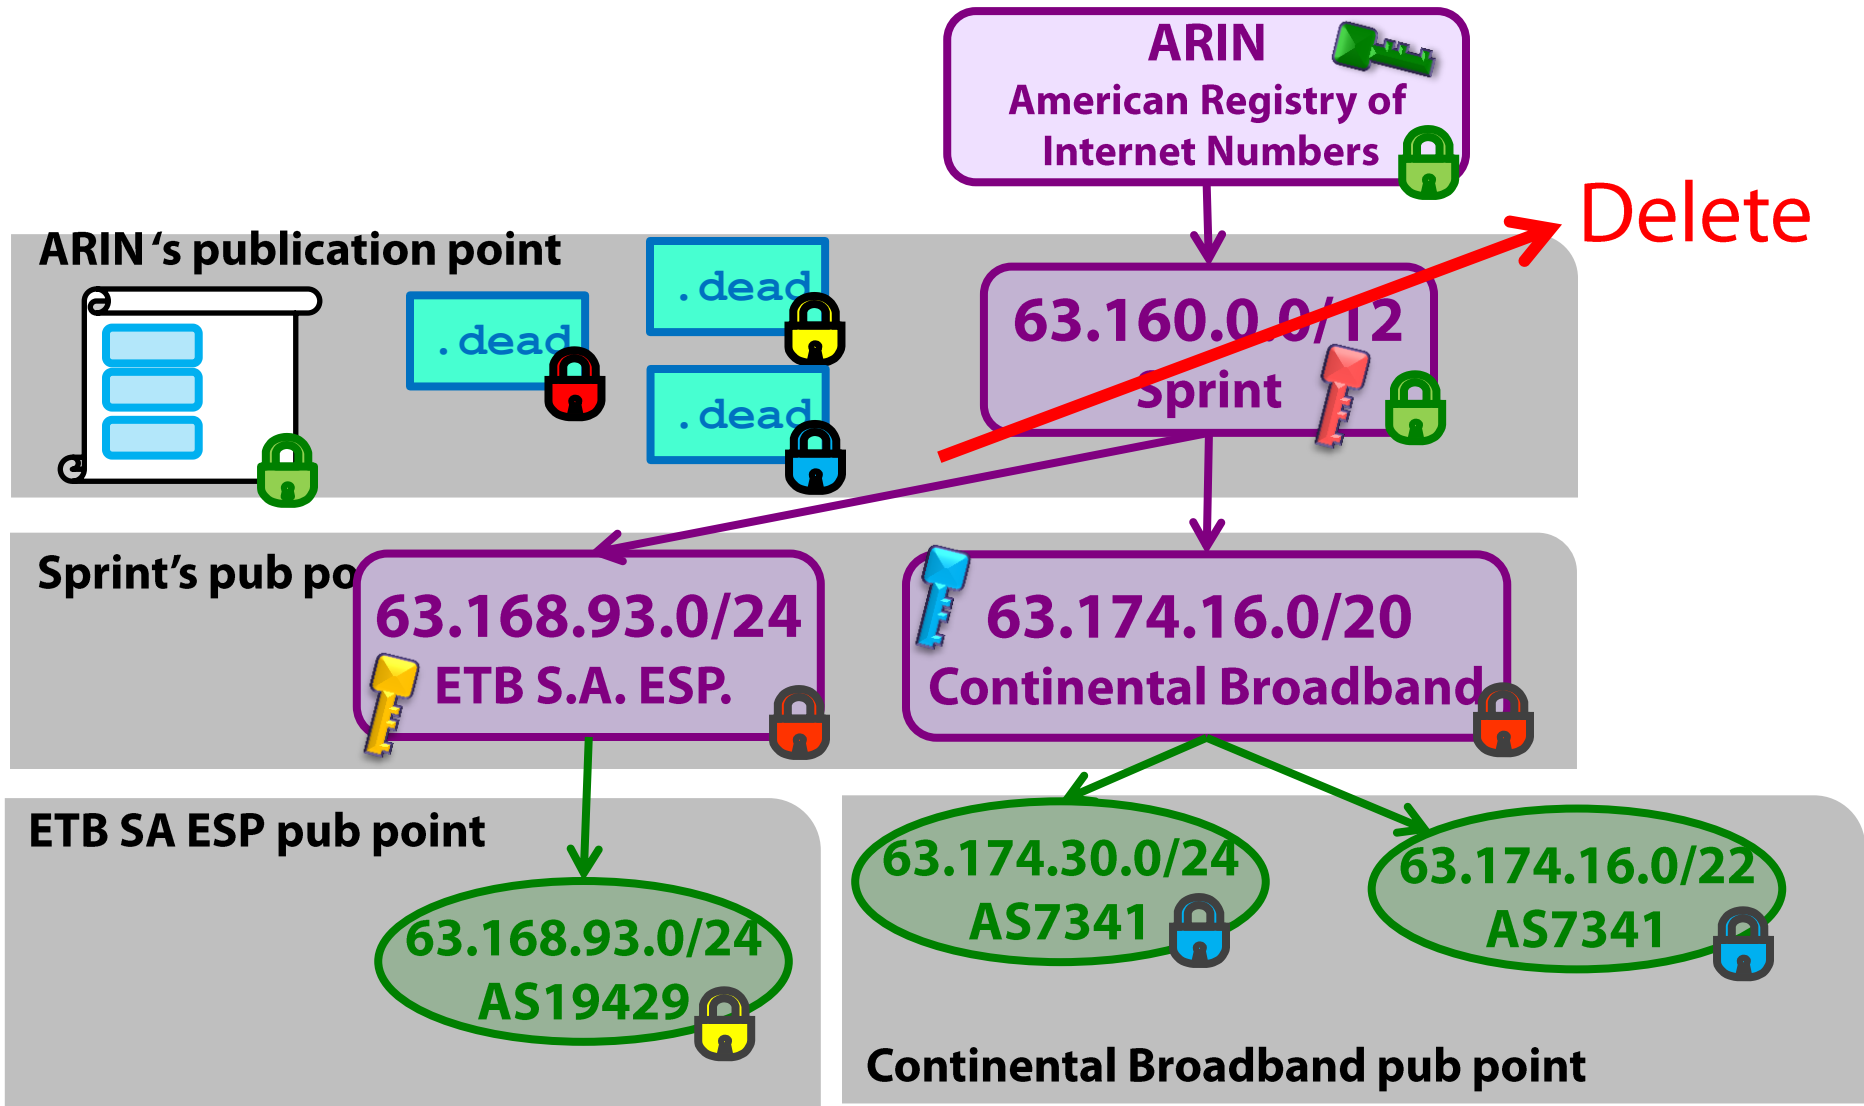
consent in a deep hierarchy: **deletion**



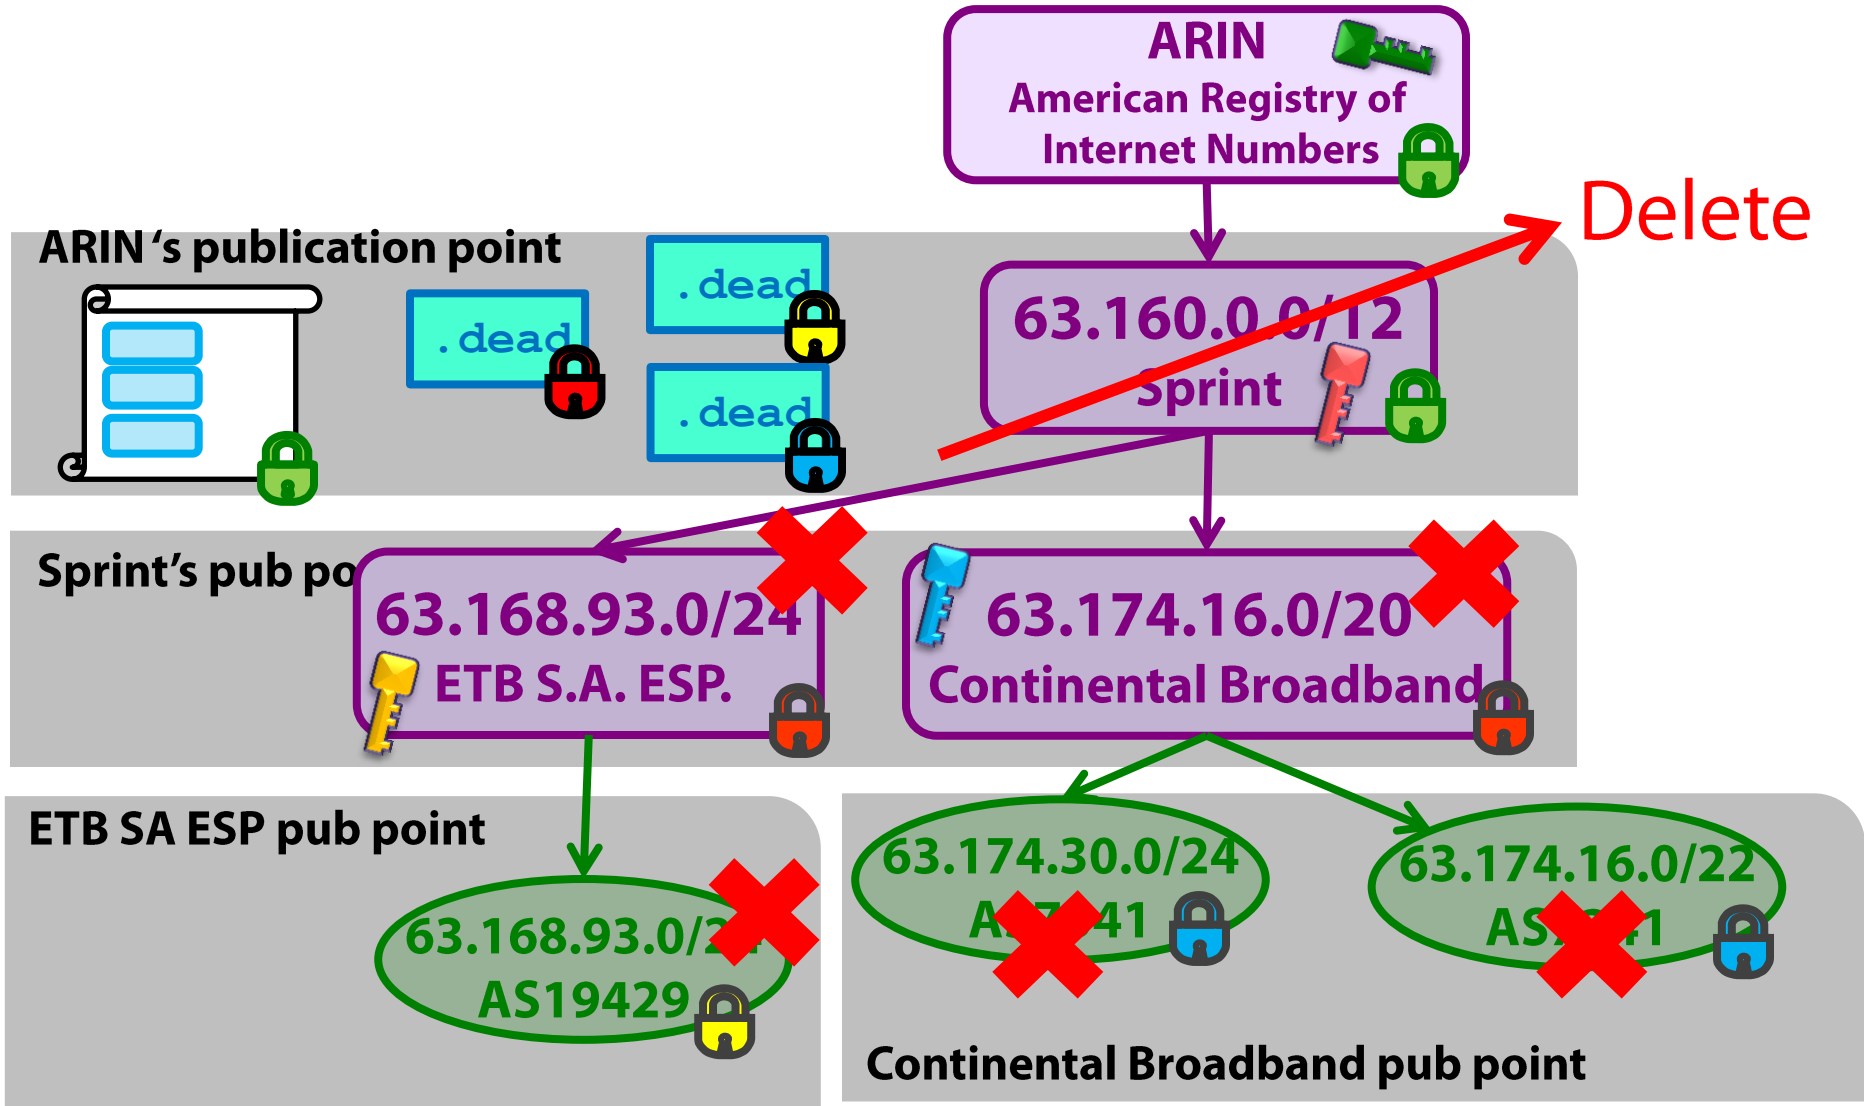
consent in a deep hierarchy: **deletion**



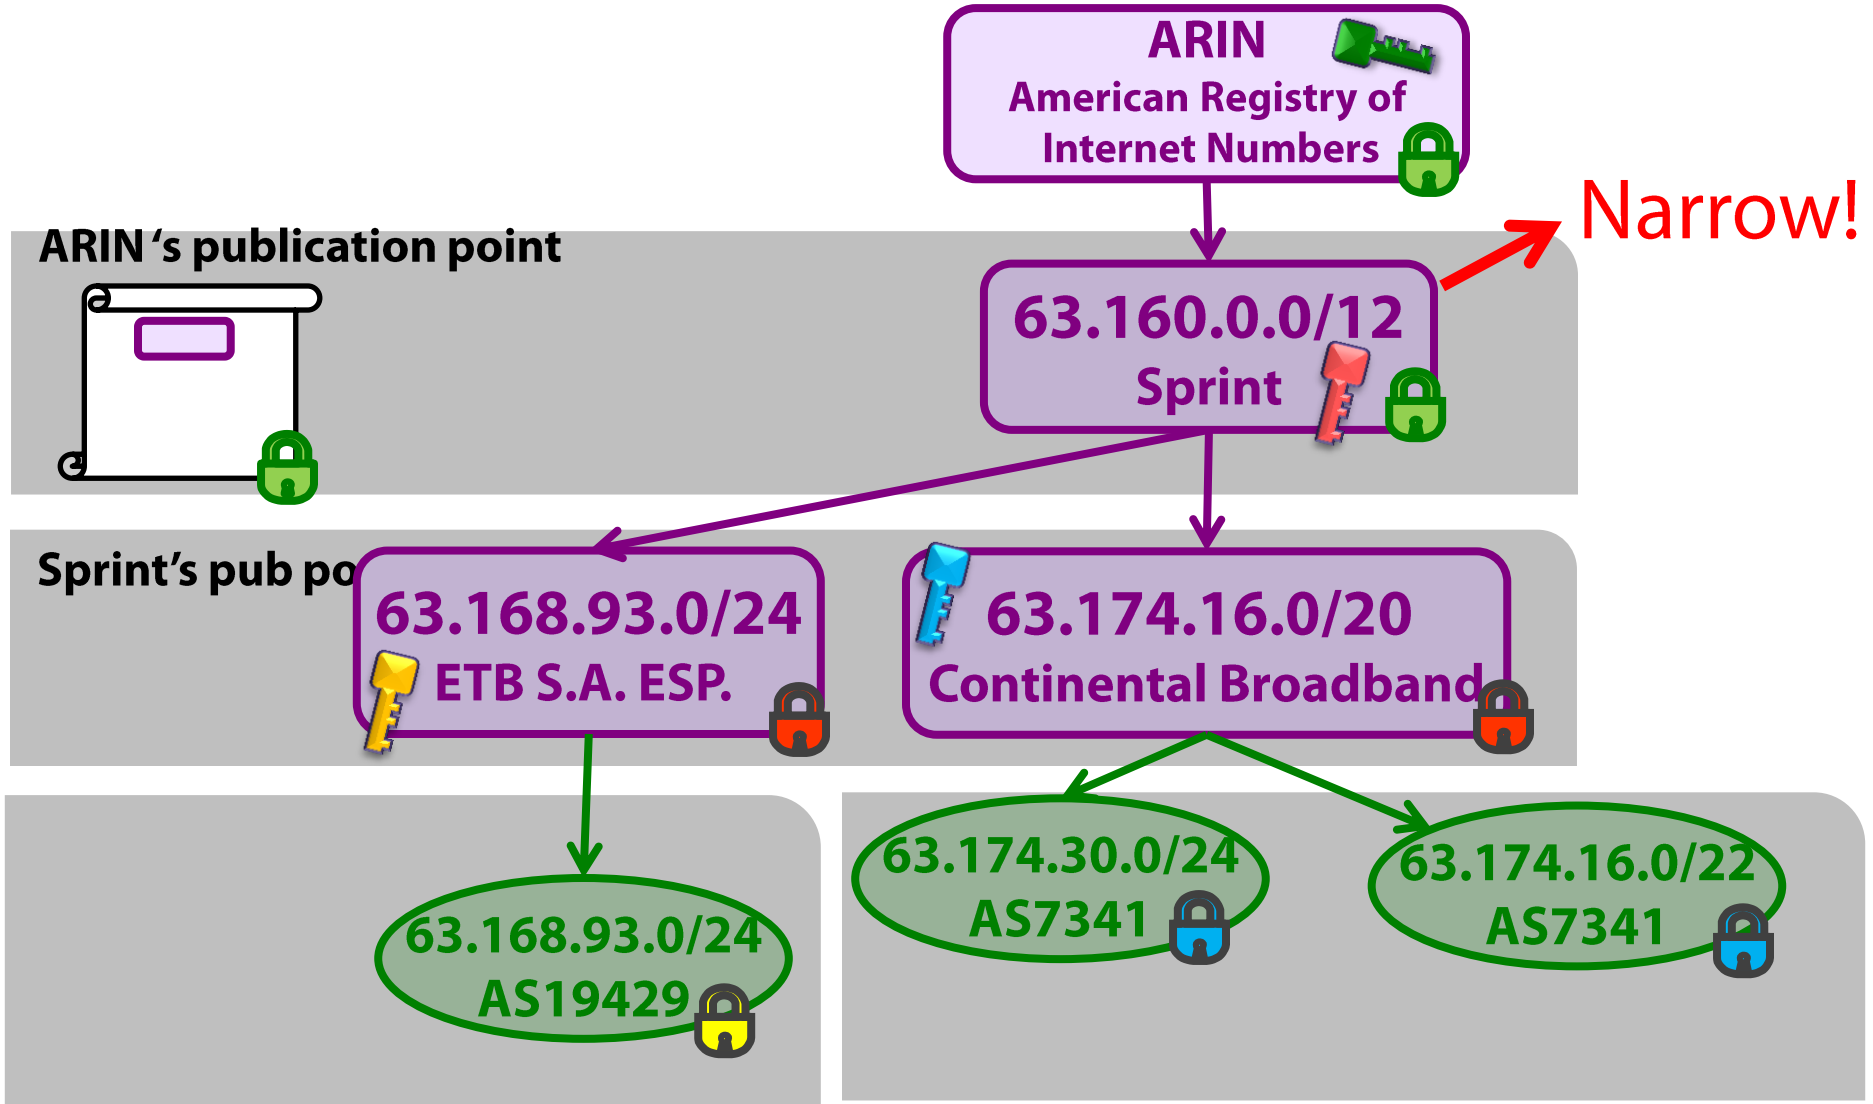
consent in a deep hierarchy: **deletion**



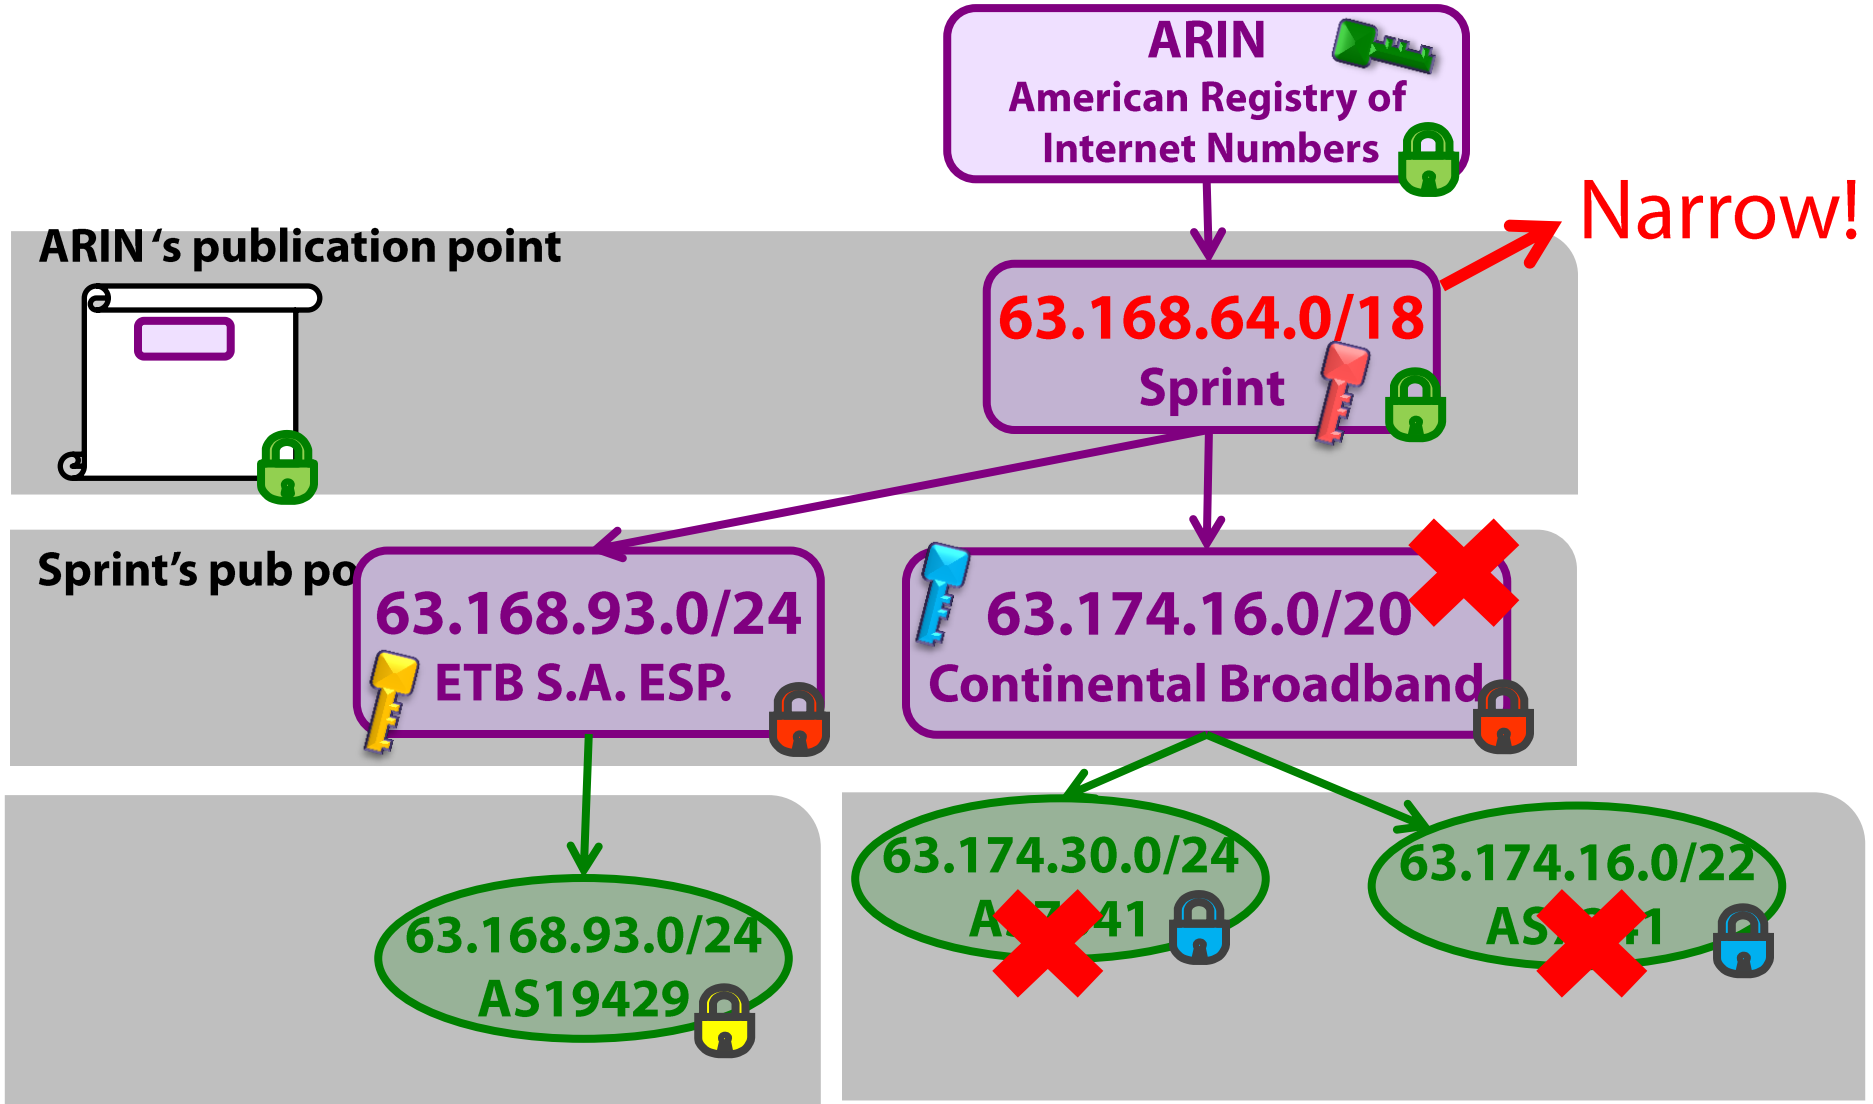
consent in a deep hierarchy: **deletion**




consent in a deep hierarchy: "address block narrowing"



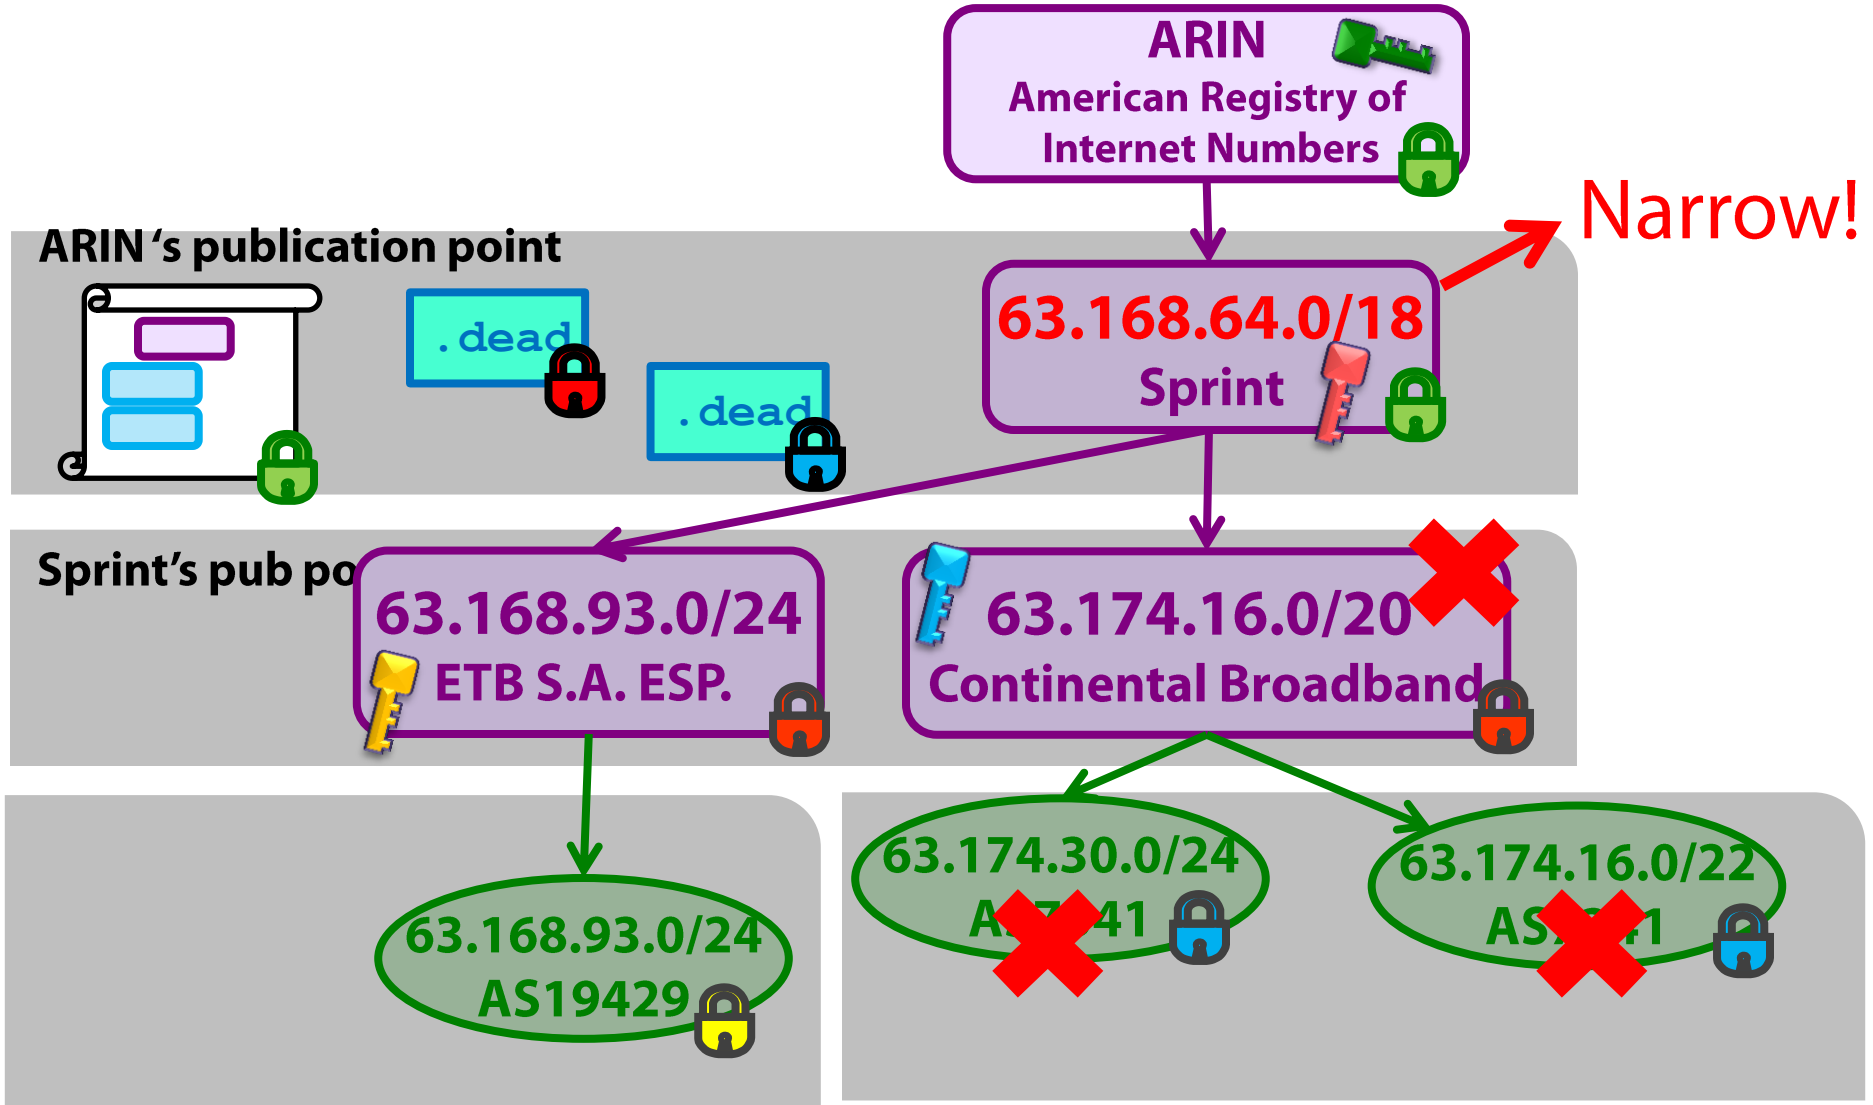
consent in a deep hierarchy: "address block narrowing"




consent in a deep hierarchy: "address block narrowing"




consent in a deep hierarchy: "address block narrowing"



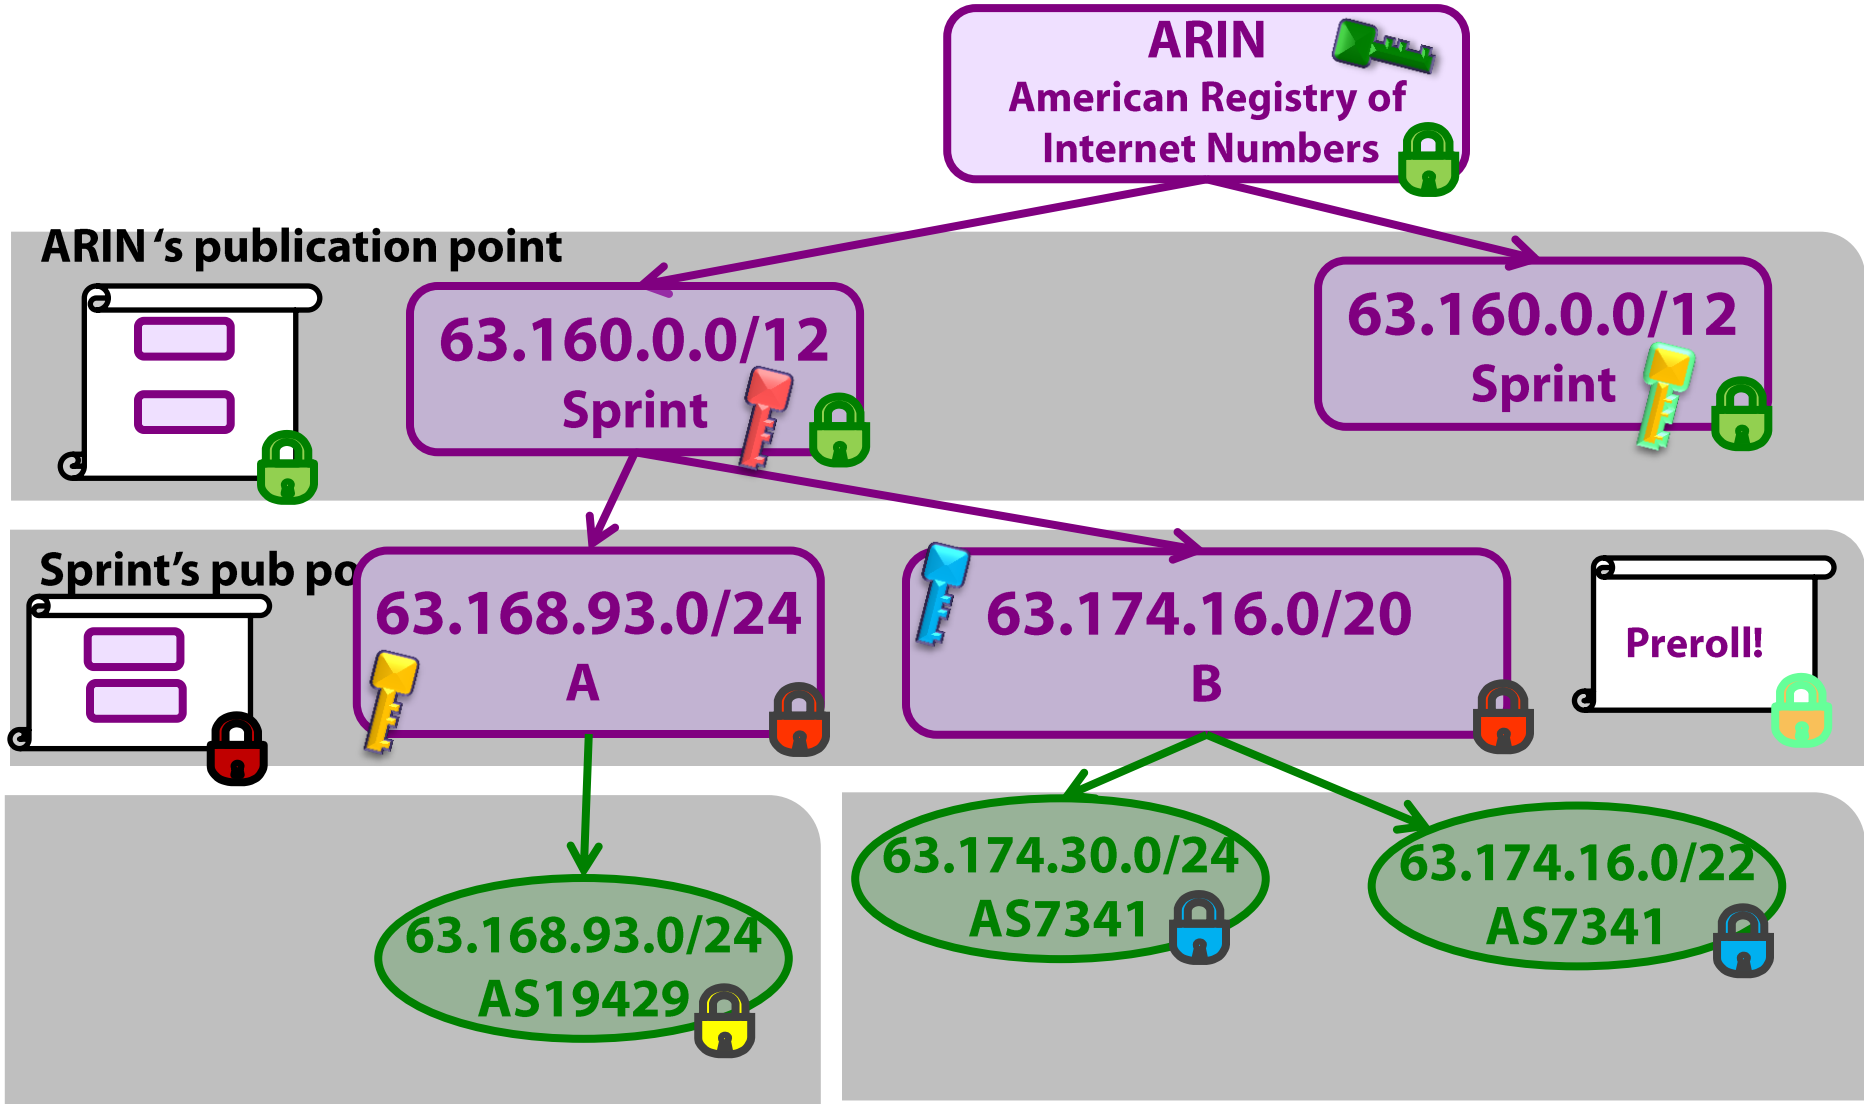
key rollover



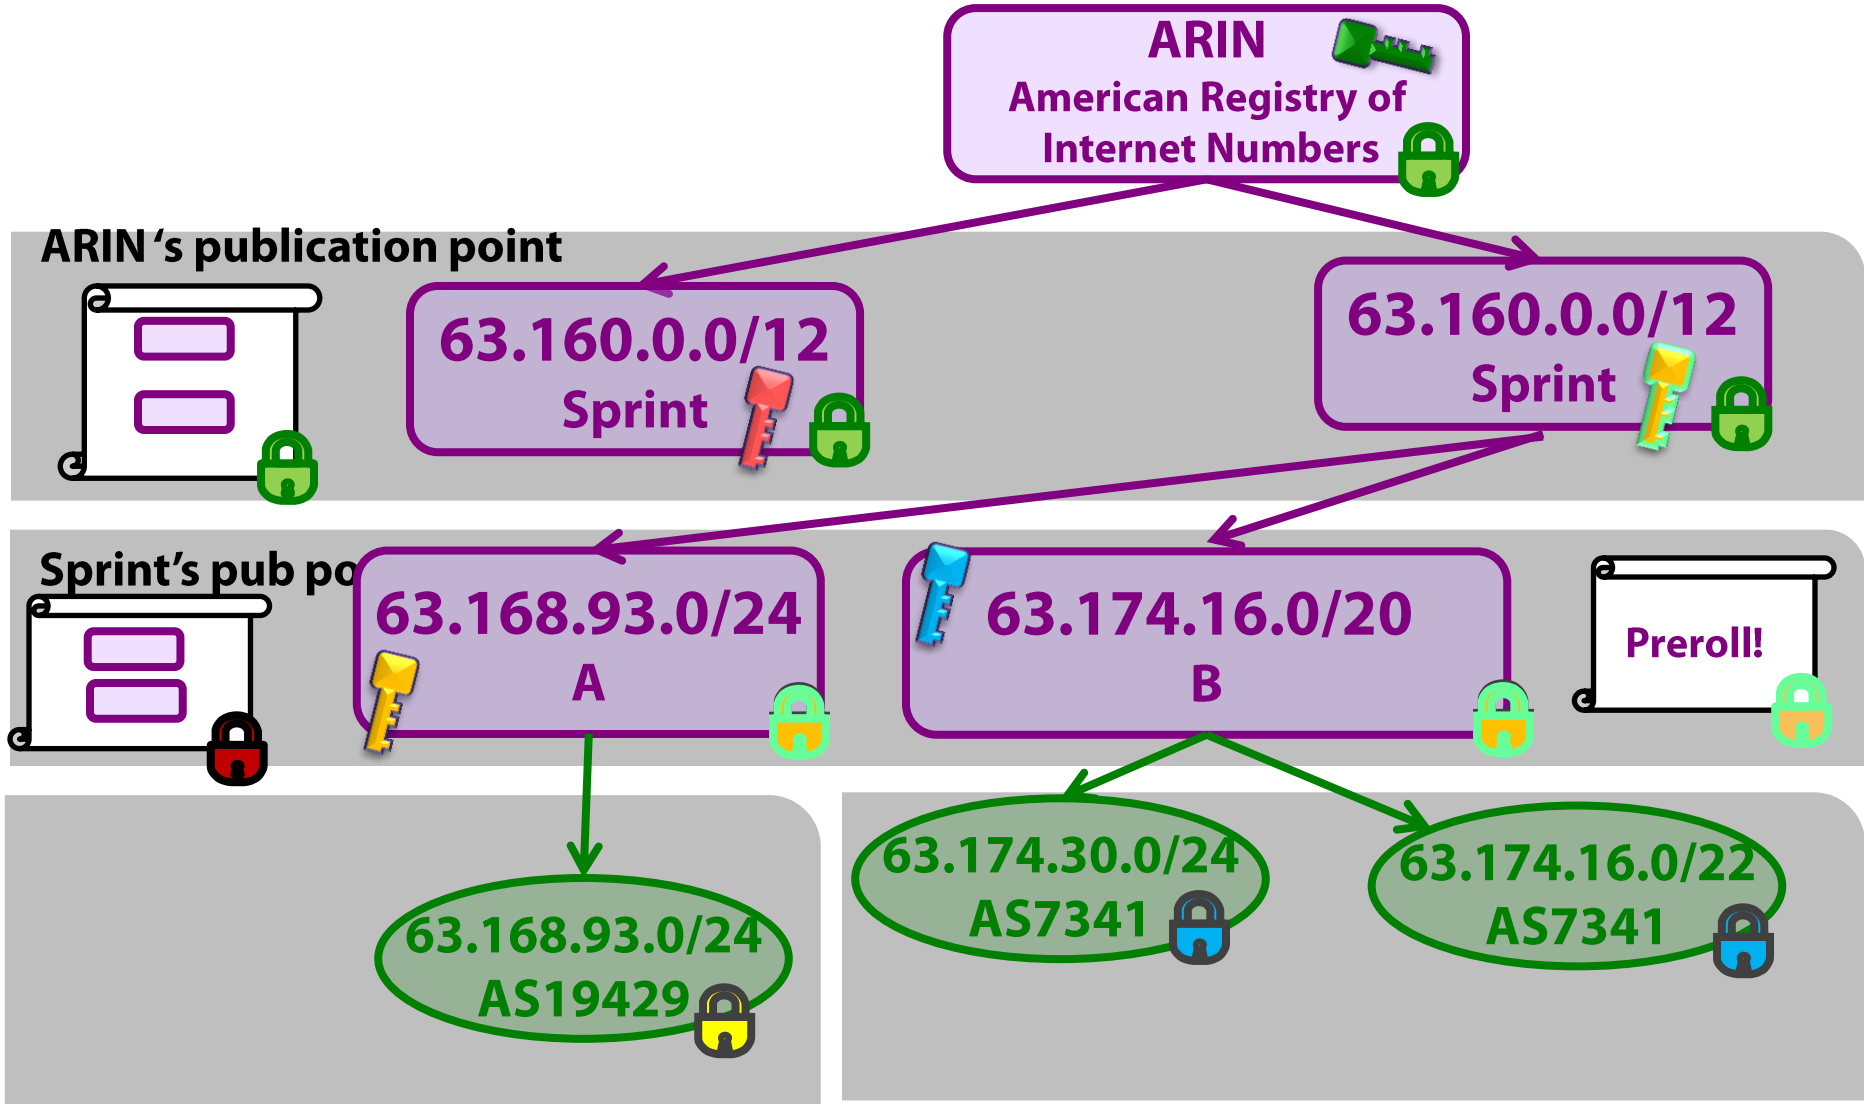
key rollover (step 0)




key rollover (step 1)



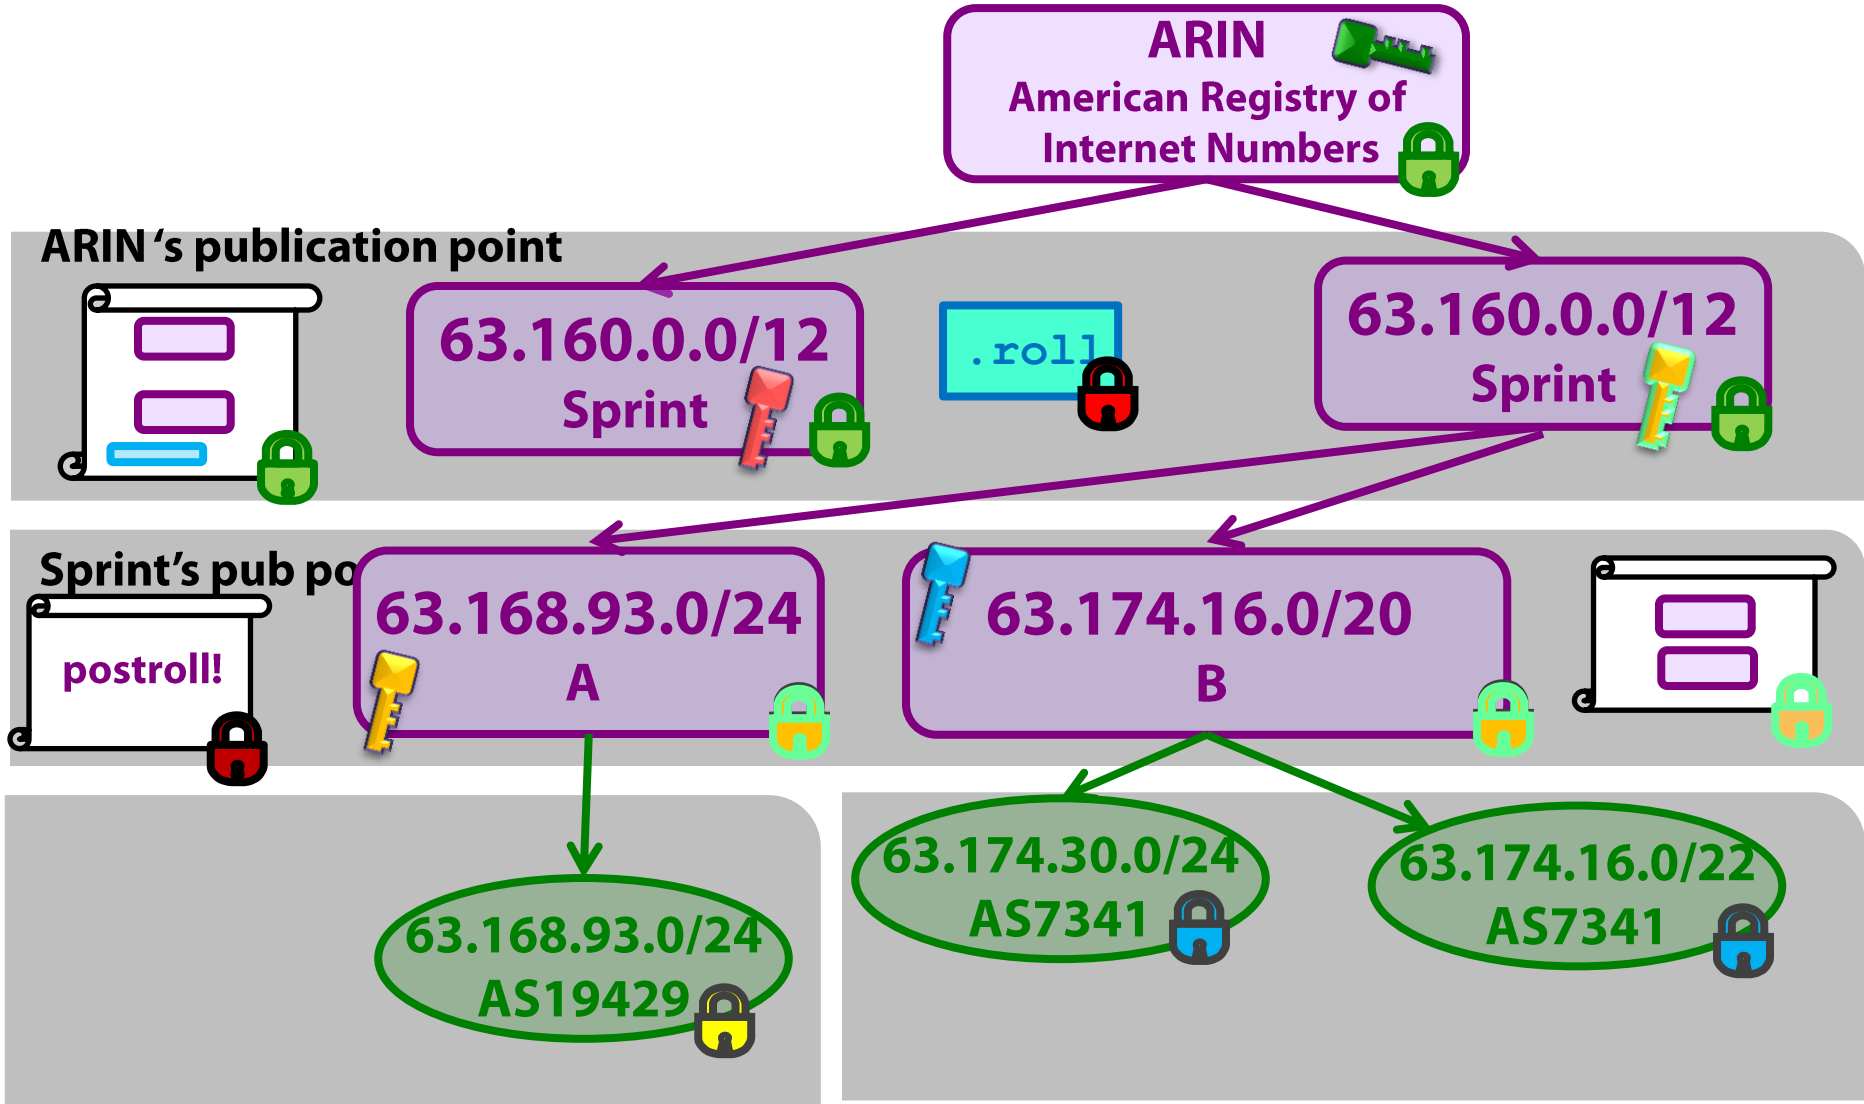
key rollover (step 2)




key rollover (step 2)




key rollover (step 3)



key rollover (step 3)

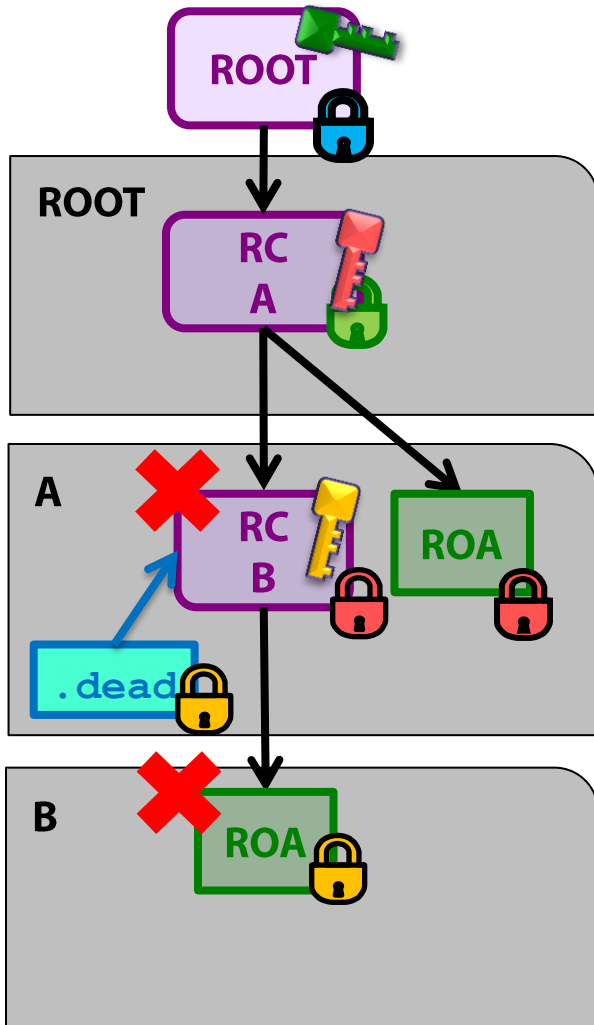


key rollover (step 3)



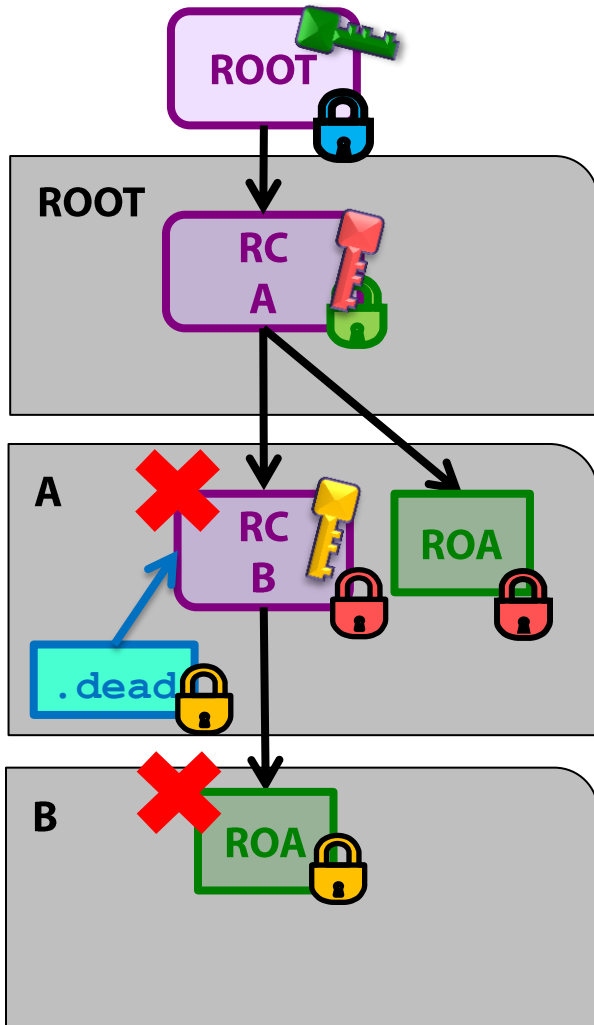
our proposal vs suspenders

our proposal
[SIGCOMM'14]

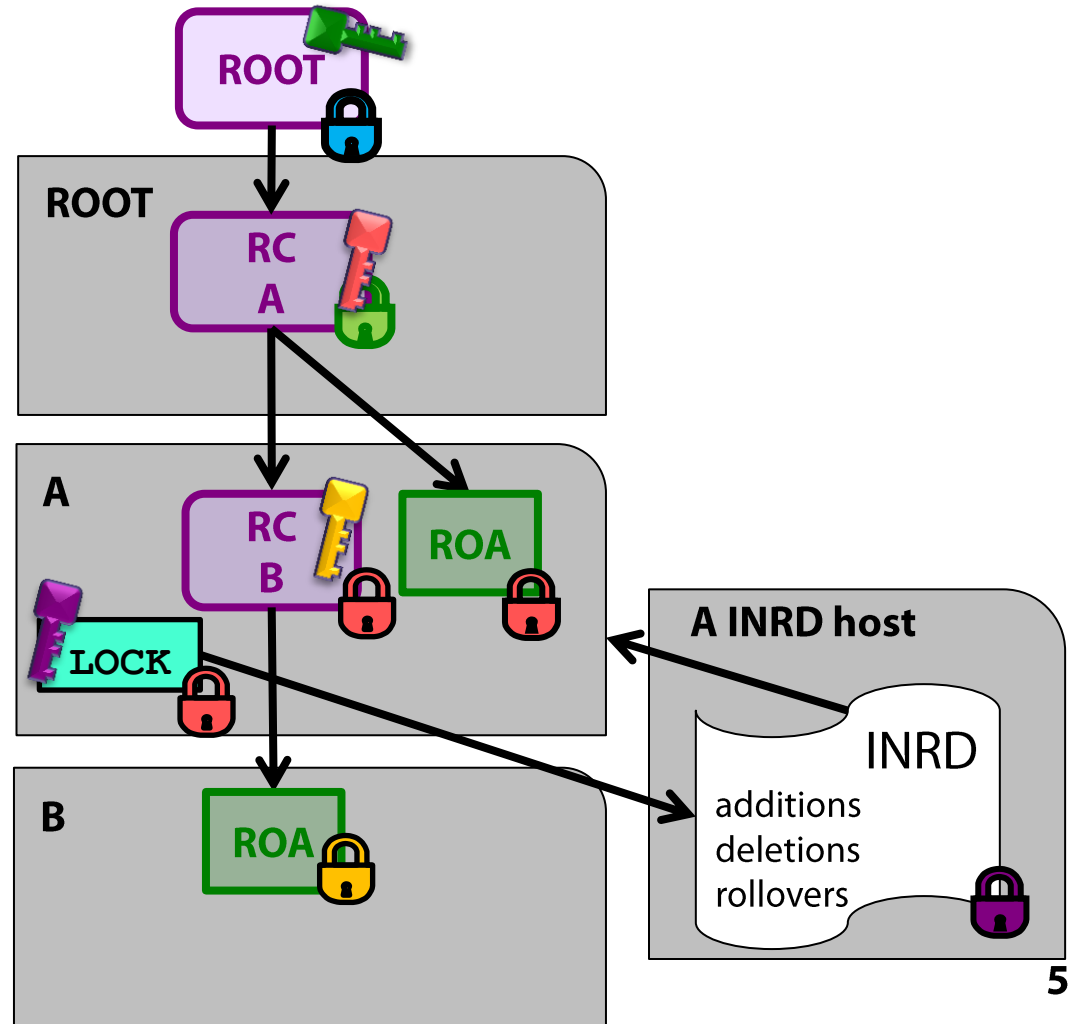


our proposal vs suspenders

our proposal [SIGCOMM'14]

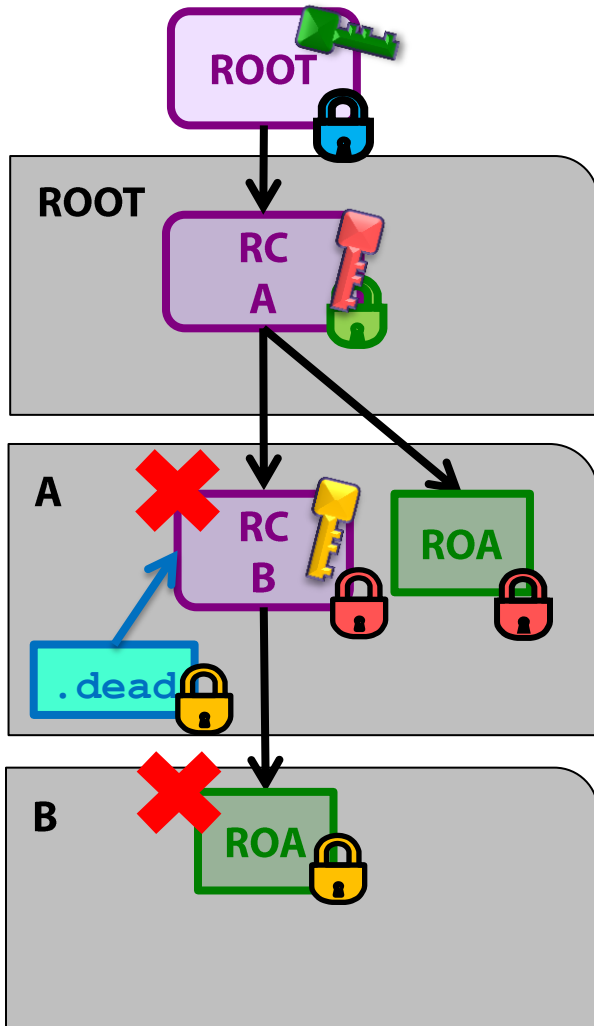


suspenders [draft-kent-sidr-suspenders-02]

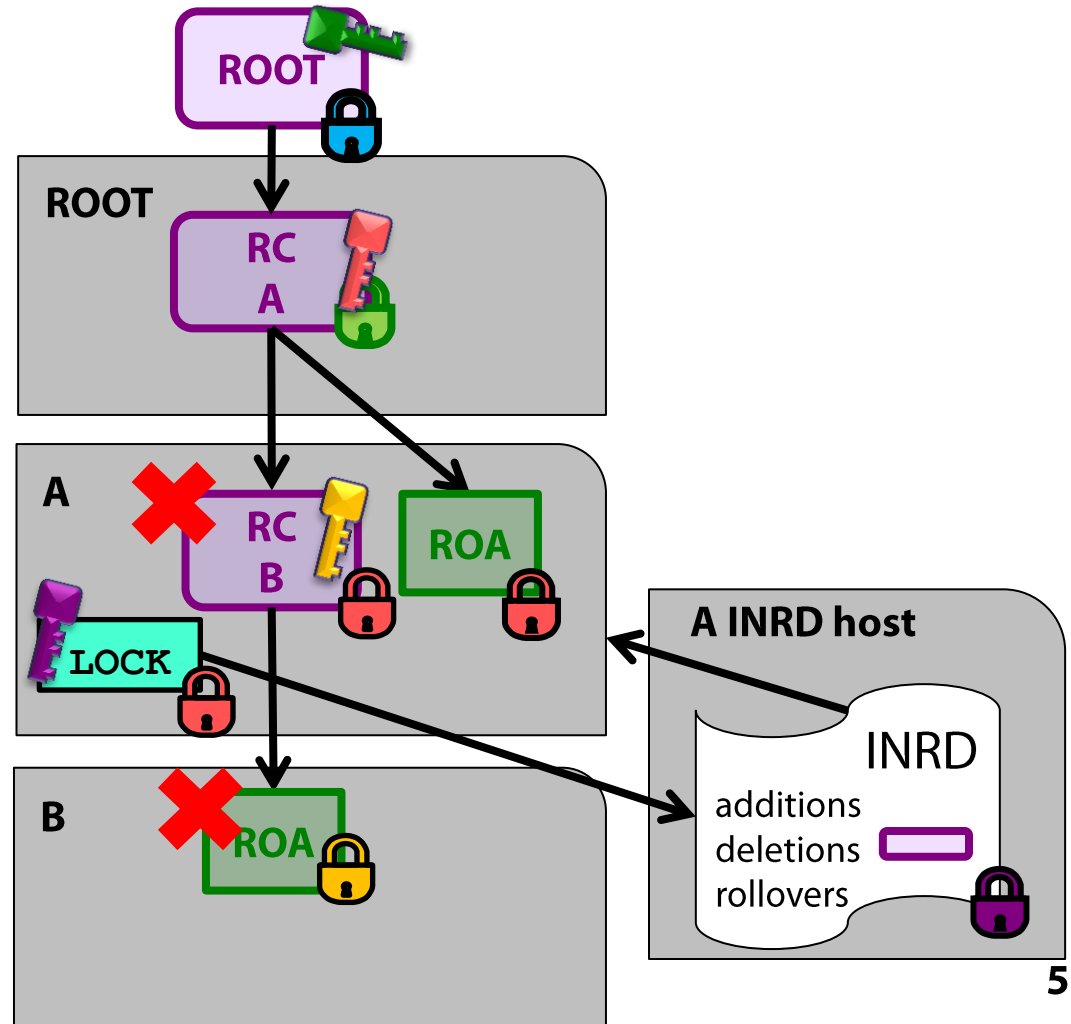


our proposal vs suspenders

our proposal [SIGCOMM'14]



suspenders [draft-kent-sidr-suspenders-02]



our proposal vs suspenders

		Our proposal	Suspenders
	Auditor:	Any Relying Party	
	Consent for whacking?	Yes: RCs	Yes: RCs & ROAs
	"Consent" for "ROA competition"?	No	Yes
	Consistency?	Yes	No
Require	Limited non-repudiation?	Yes	No?

our proposal vs suspenders

		Our proposal	Suspenders
	Auditor:	Any Relying Party	
	Consent for whacking?	Yes: RCs	Yes: RCs & ROAs
	"Consent" for "ROA competition"?	No	Yes
	Consistency?	Yes	No
Require	Limited non-repudiation?	Yes	No?
	New RPKI objects:	.dead .roll change logs	LOCK INRD
Design	Requires changes to manifests?	Yes	No

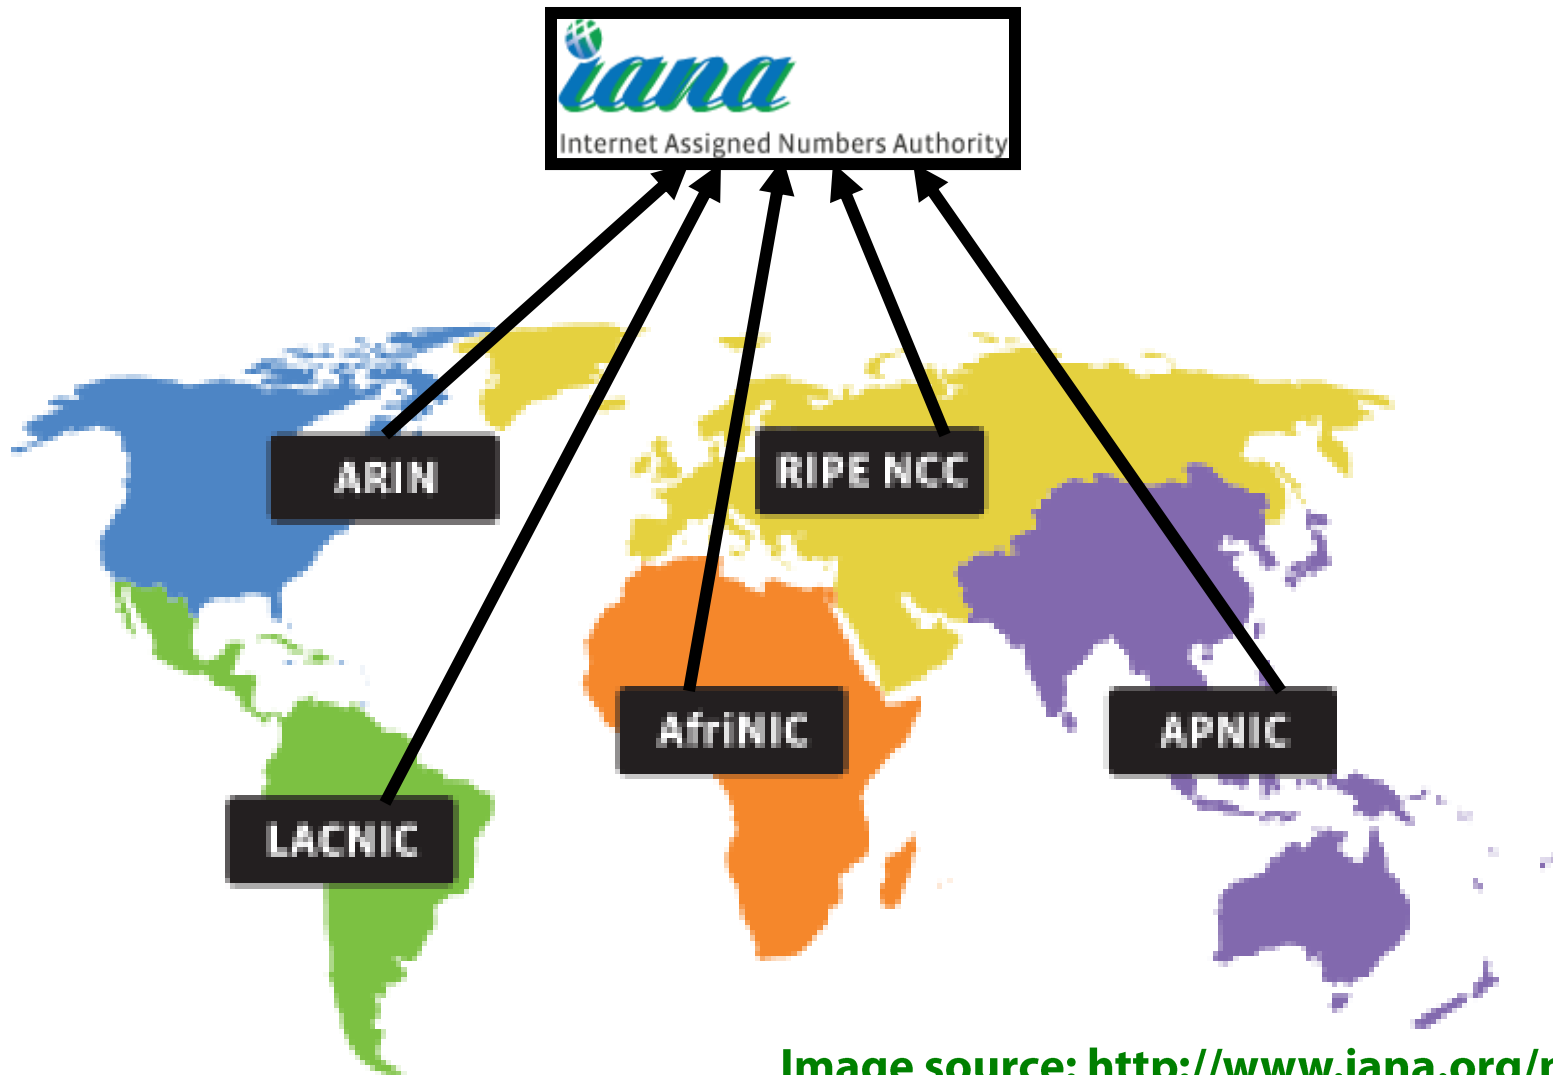
our proposal vs suspenders

		Our proposal	Suspenders
	Auditor:	Any Relying Party	
	Consent for whacking?	Yes: RCs	Yes: RCs & ROAs
	“Consent” for “ROA competition”?	No	Yes
	Consistency?	Yes	No
Require	Limited non-repudiation?	Yes	No?
	New RPKI objects:	.dead .roll change logs	LOCK INRD
Design	Requires changes to manifests?	Yes	No
	“Out of band” publication points?	Yes	No
	“Consenting” subjects need keys?	Yes	Yes
	Proofs of security goals:	Yes	No

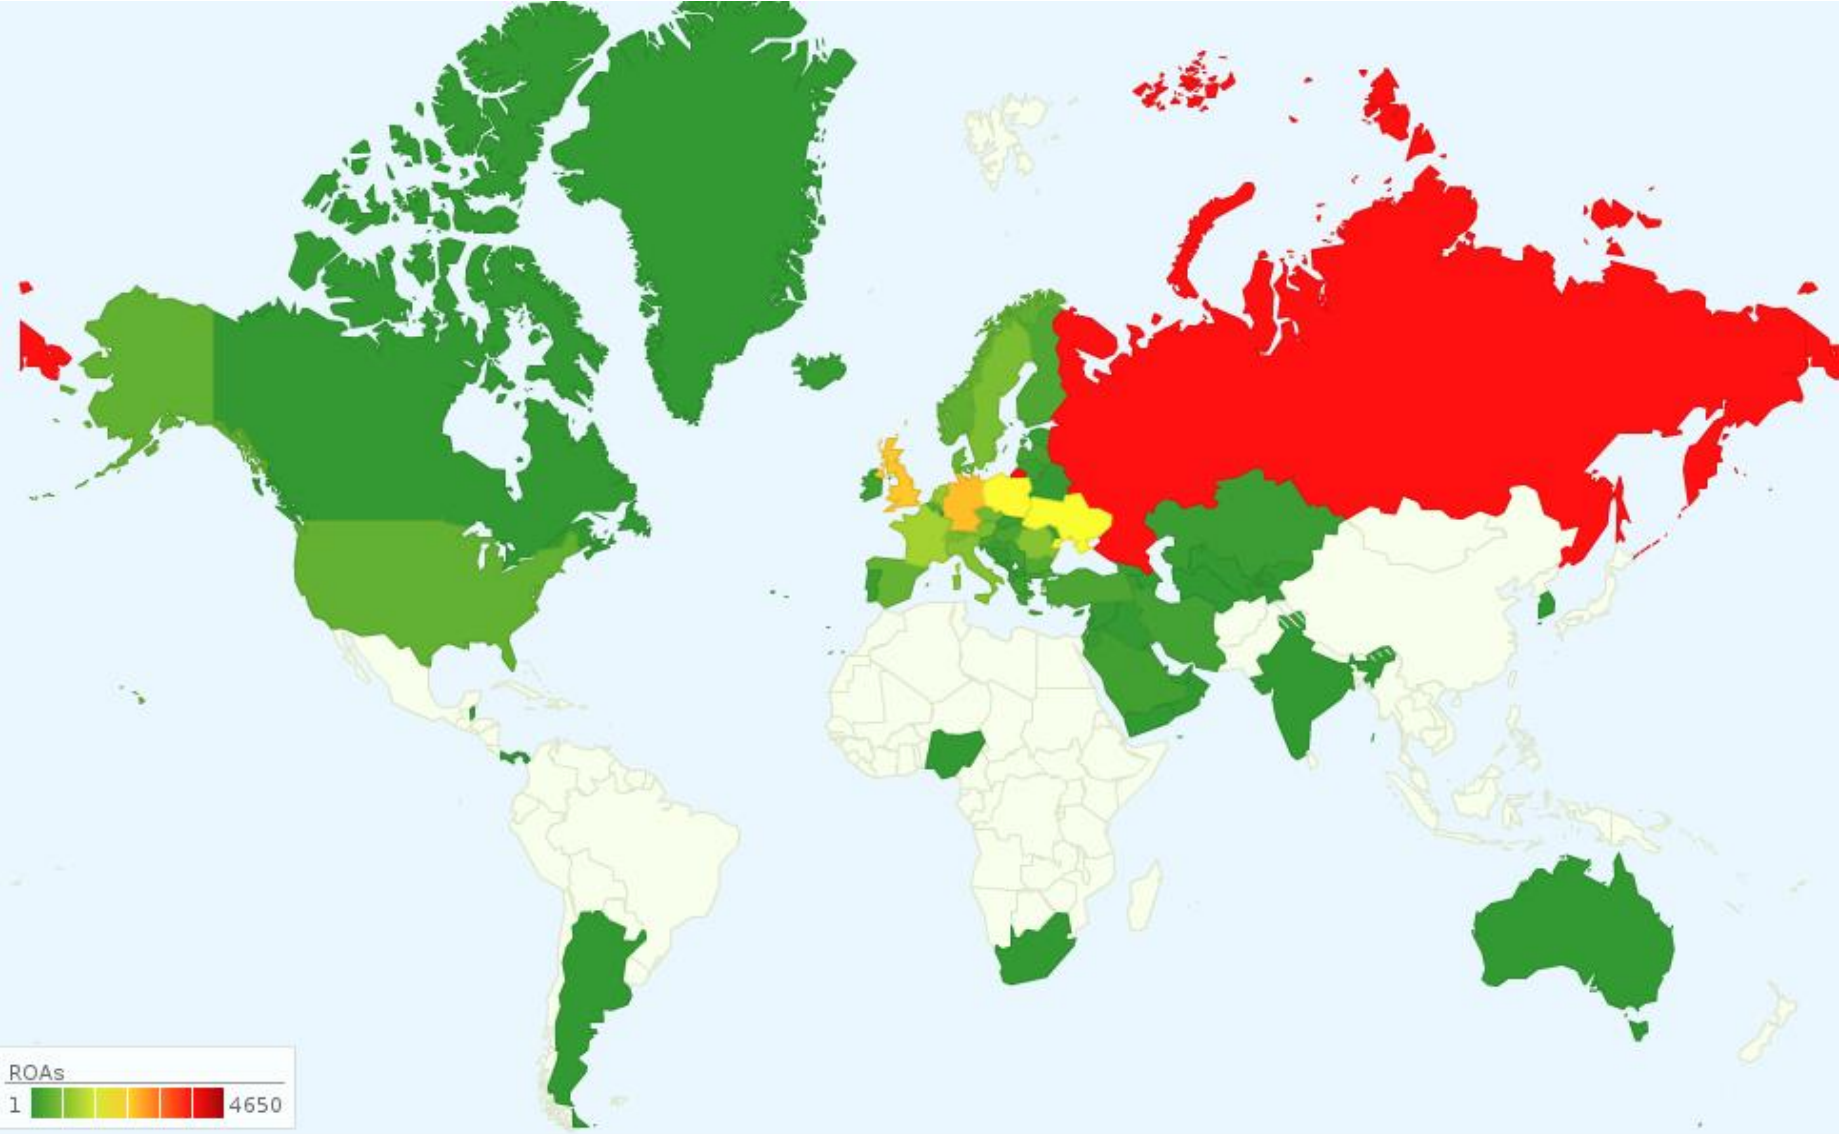
Question for the room: What is the right set of requirements?

Who controls the root?

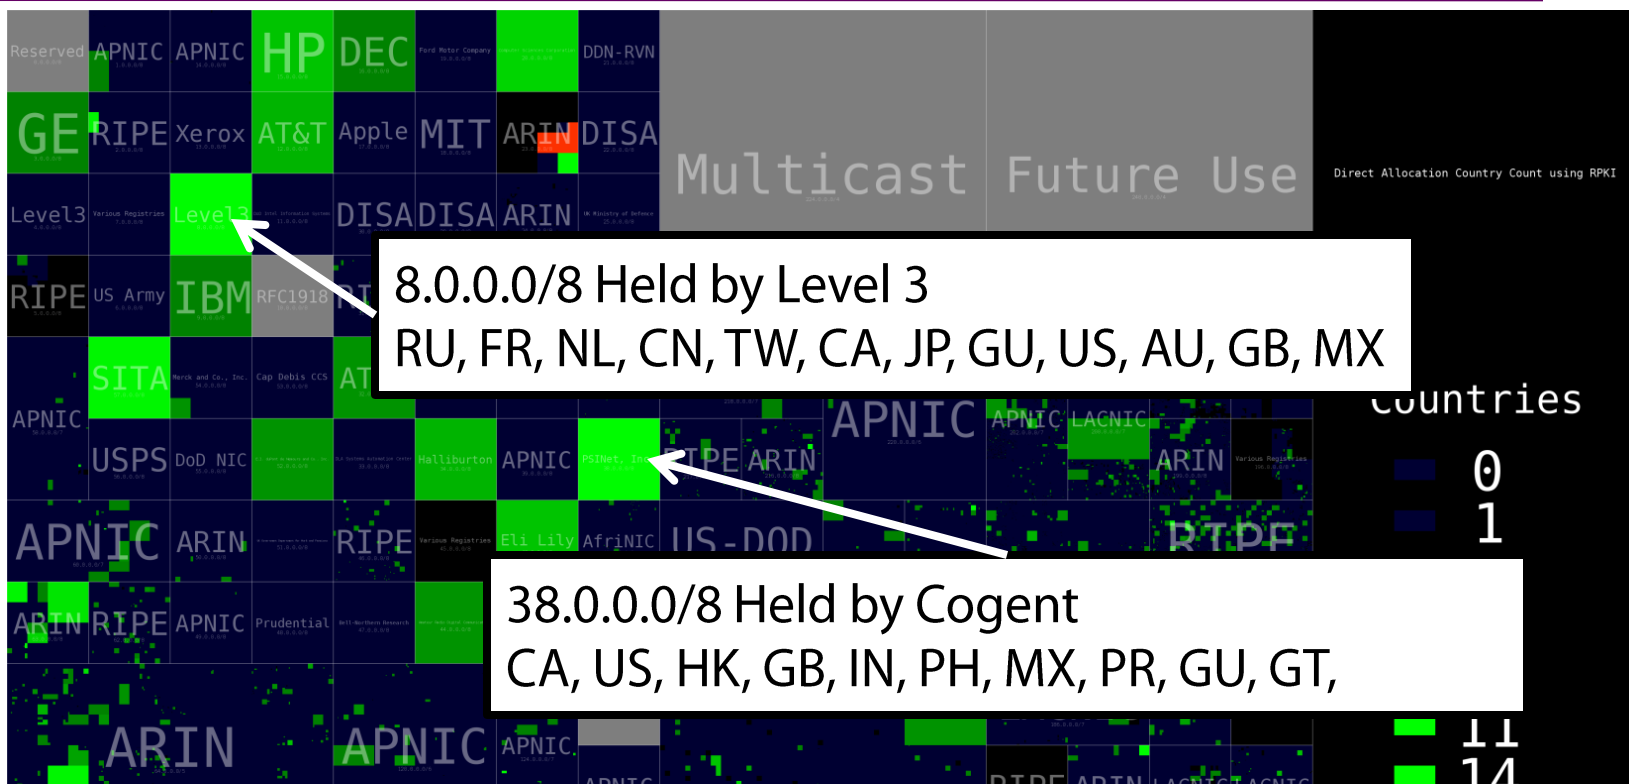
Is there a single root of trust? Unclear; IAB says yes.
Right now there are 25.



Countries covered by RIPE




IPv4 address allocation does not reflect jurisdiction




Data-driven model of the RPKI (today's RPKI is too small)

- ✧ Using RIR direct allocations, routeviews, BGP table dumps
- ✧ RIRs and their direct allocations get RCs, other (prefix,origin AS) pairs in the table dumps get a ROA
- ✧ ASes mapped to countries using RIR data

Number of ROAs issued by each direct allocation



Depth of the RPKI



Depth	ROAs
3	118,028
4	108,043
5	10,863
6	293
7	9


Route Validity Depends on More Than a Signature Chain

What if we add a ROA?

63.160.0.0/12
AS 1239

Legend



- Valid
- Invalid
- Unknown



Adding a ROA Can Invalidate Routes!

What if we add a ROA?

63.160.0.0/12
AS 1239



Adding a ROA Can Invalidate Routes!

Why does this happen?

Otherwise, we can still subprefix hijack!

(more on this in a moment!)

