RIPE NCC

RIPE NETWORK COORDINATION CENTRE

Routing Security

Lebanese University, Faculty of Sciences

Nathalie Trenaman| 13 April 2021 |

Introduction to BGP routing

Internet building blocks

ASN (Autonomous System Number)

Internet building blocks

Autonomous System

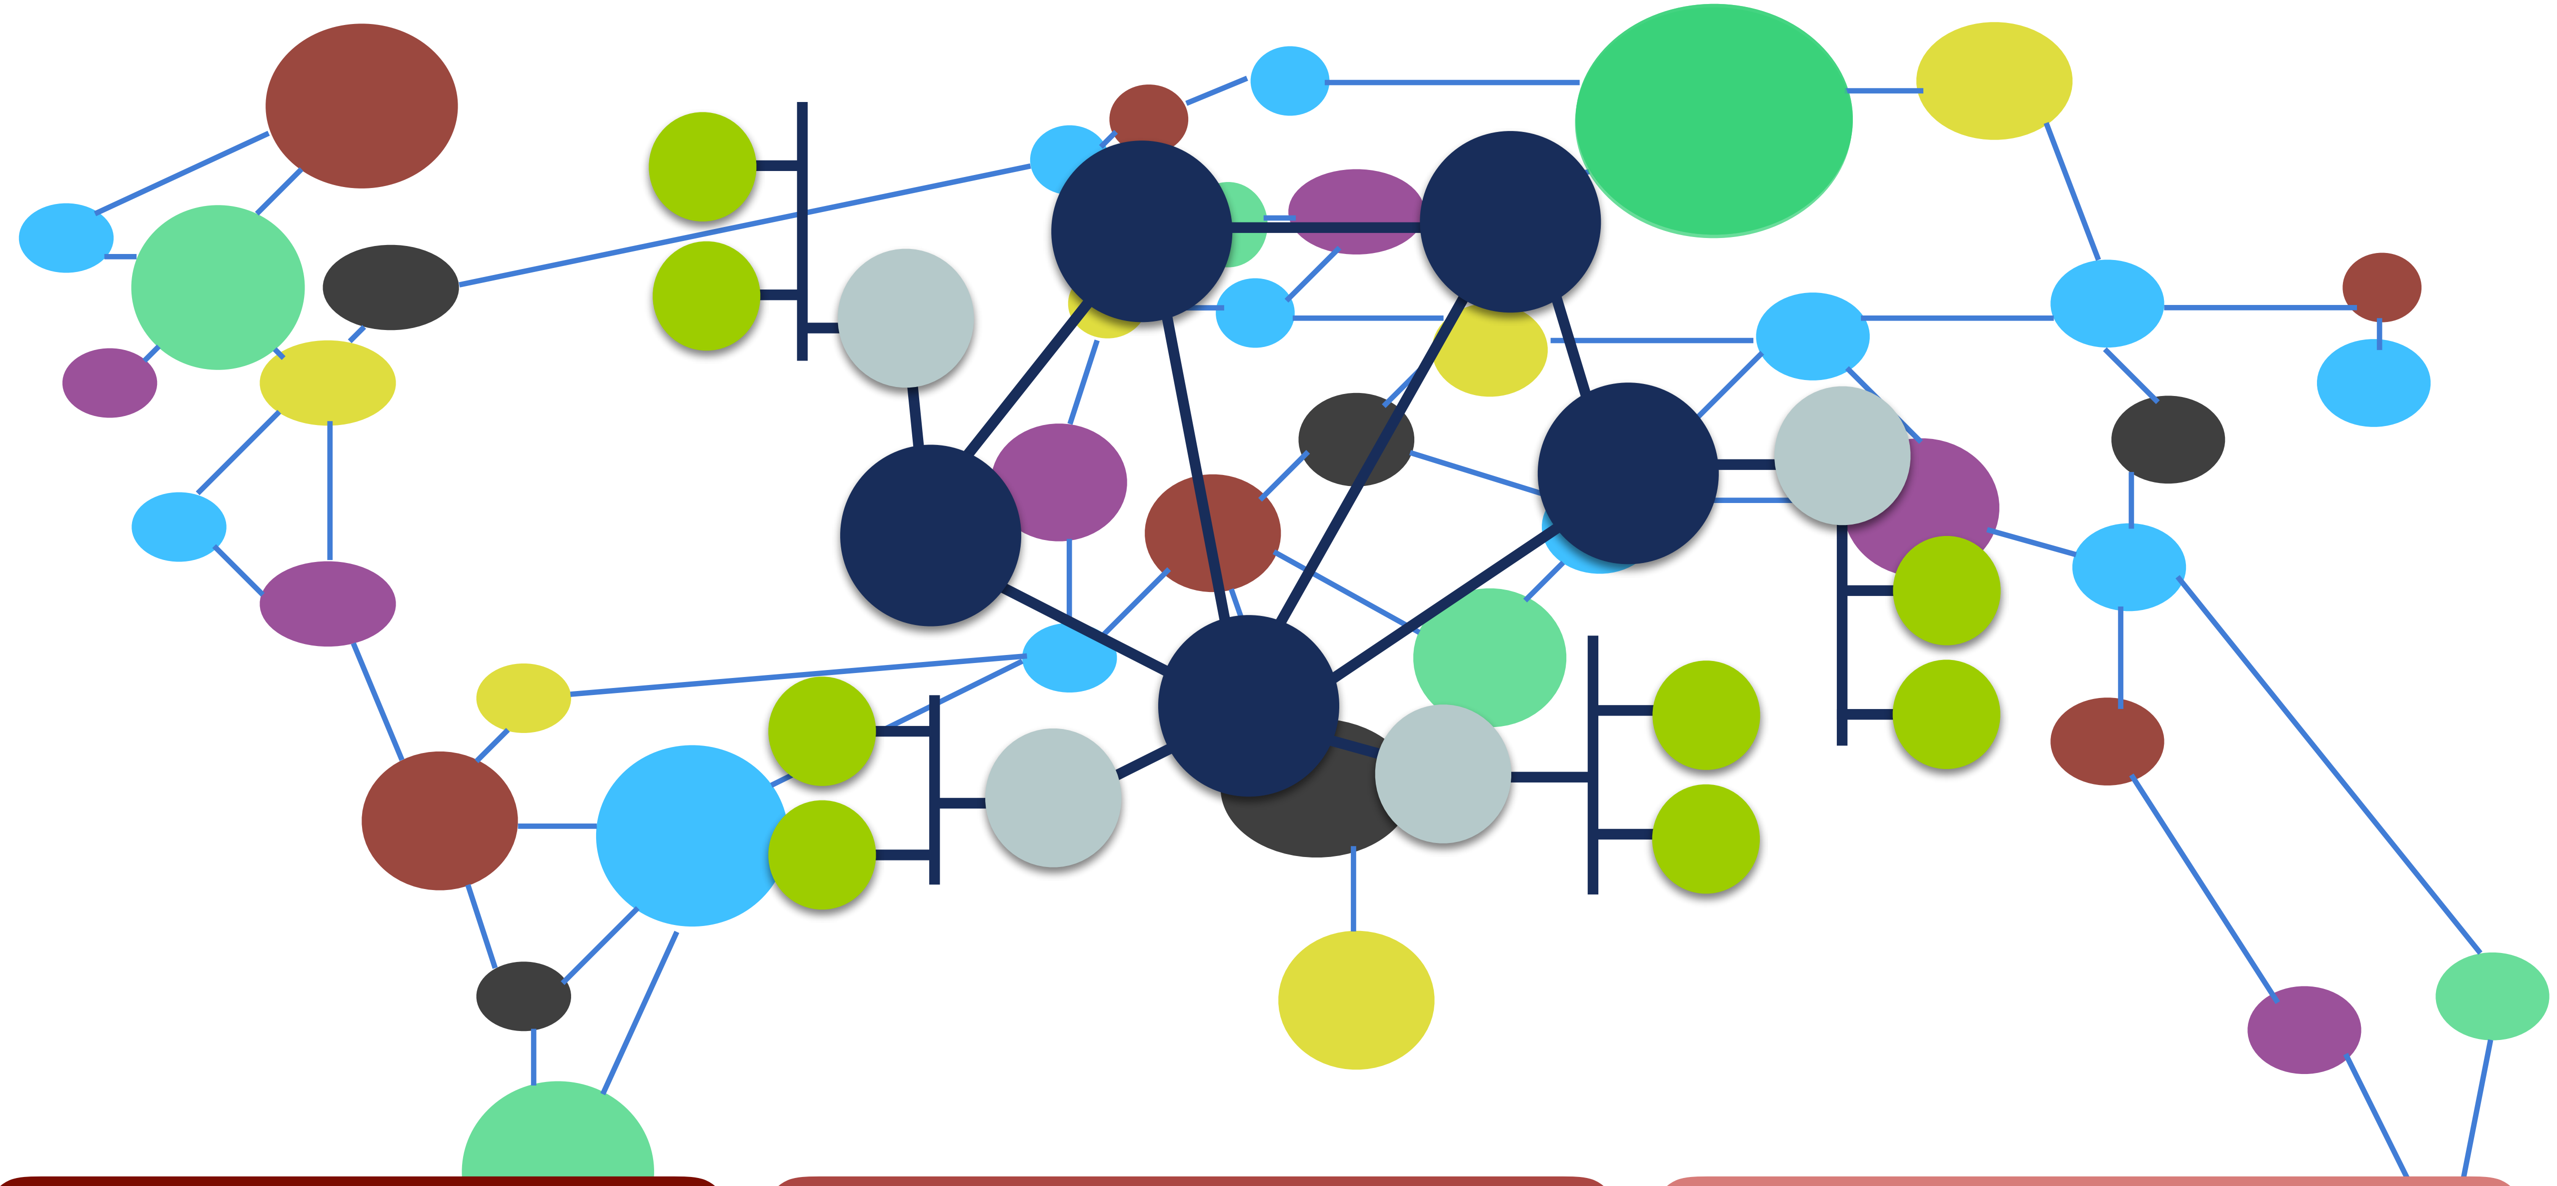

ASN

Addresses

Interconnect

Routing on the Internet

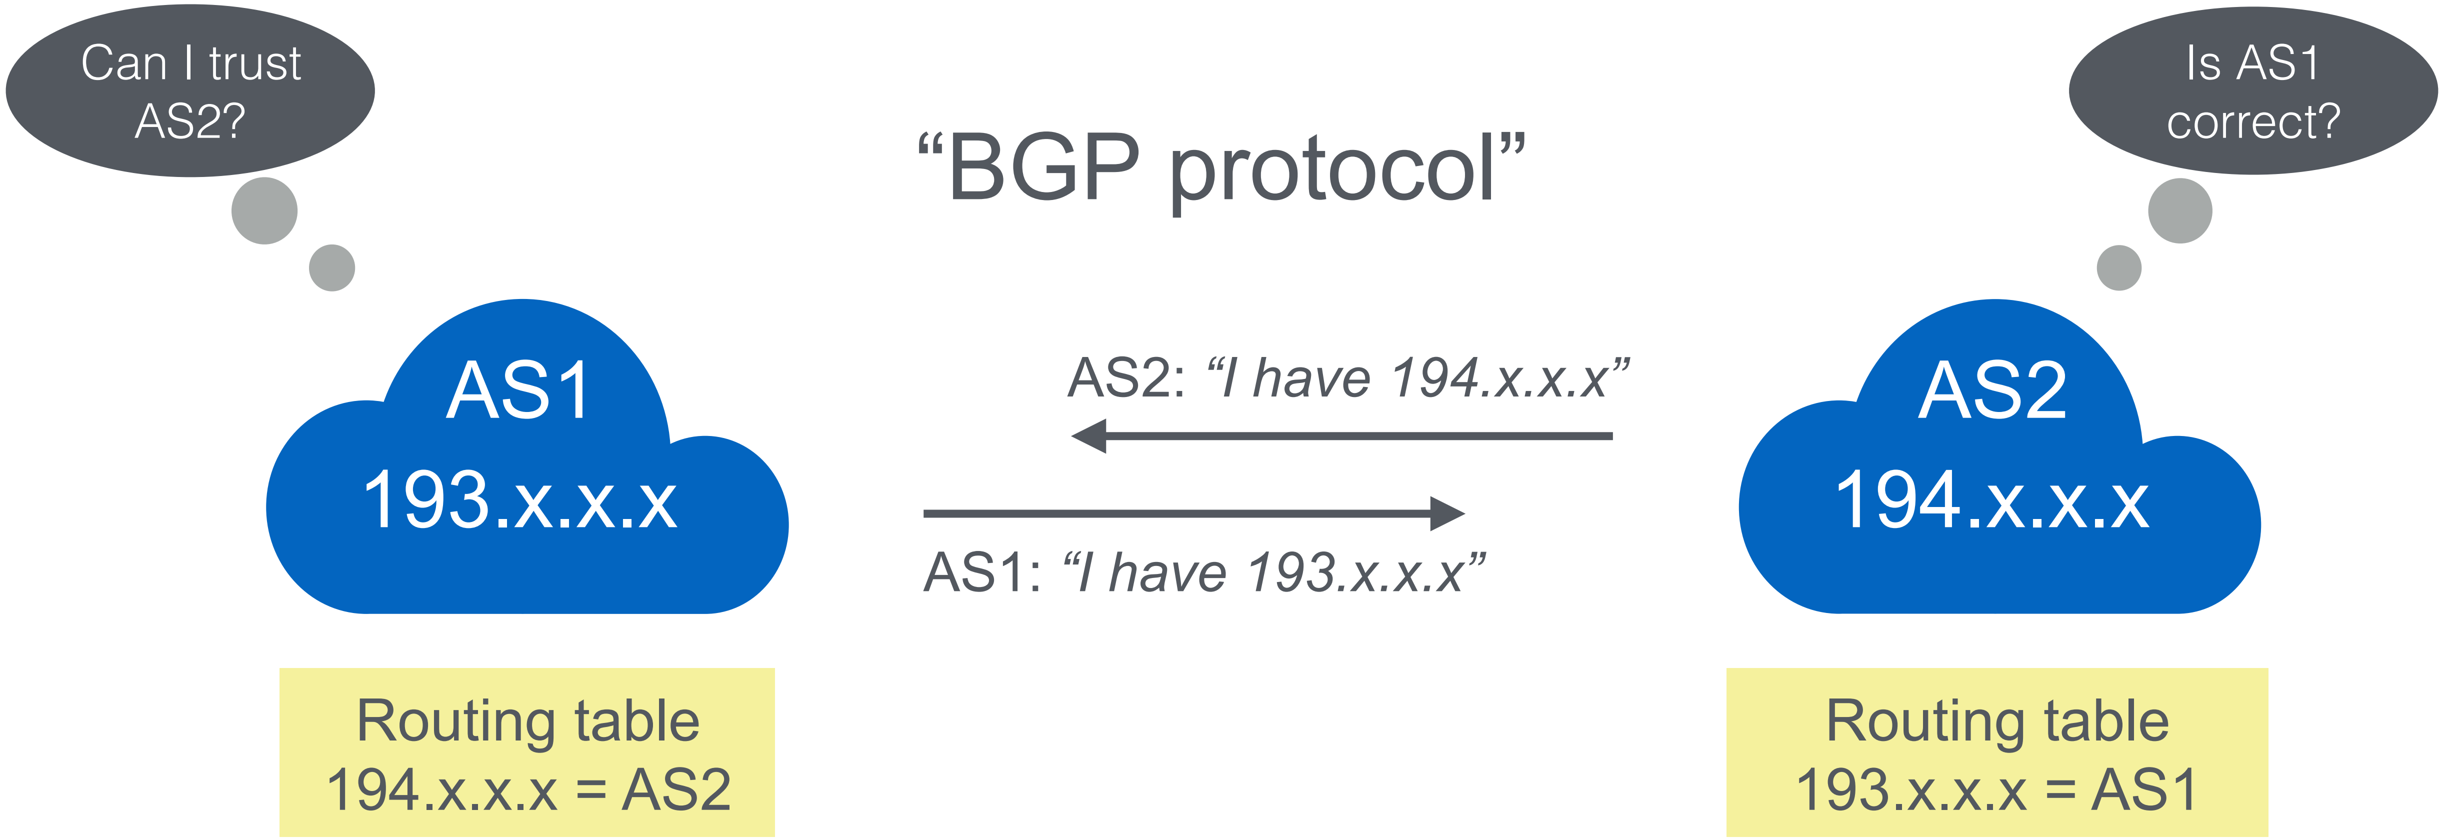

Can I trust AS2?

Is AS1 correct?

“BGP protocol”

AS1

193.x.x.x

AS2: *“I have 194.x.x.x”*

AS1: *“I have 193.x.x.x”*

AS2

194.x.x.x

Routing table
194.x.x.x = AS2

Routing table
193.x.x.x = AS1

Route Propagation

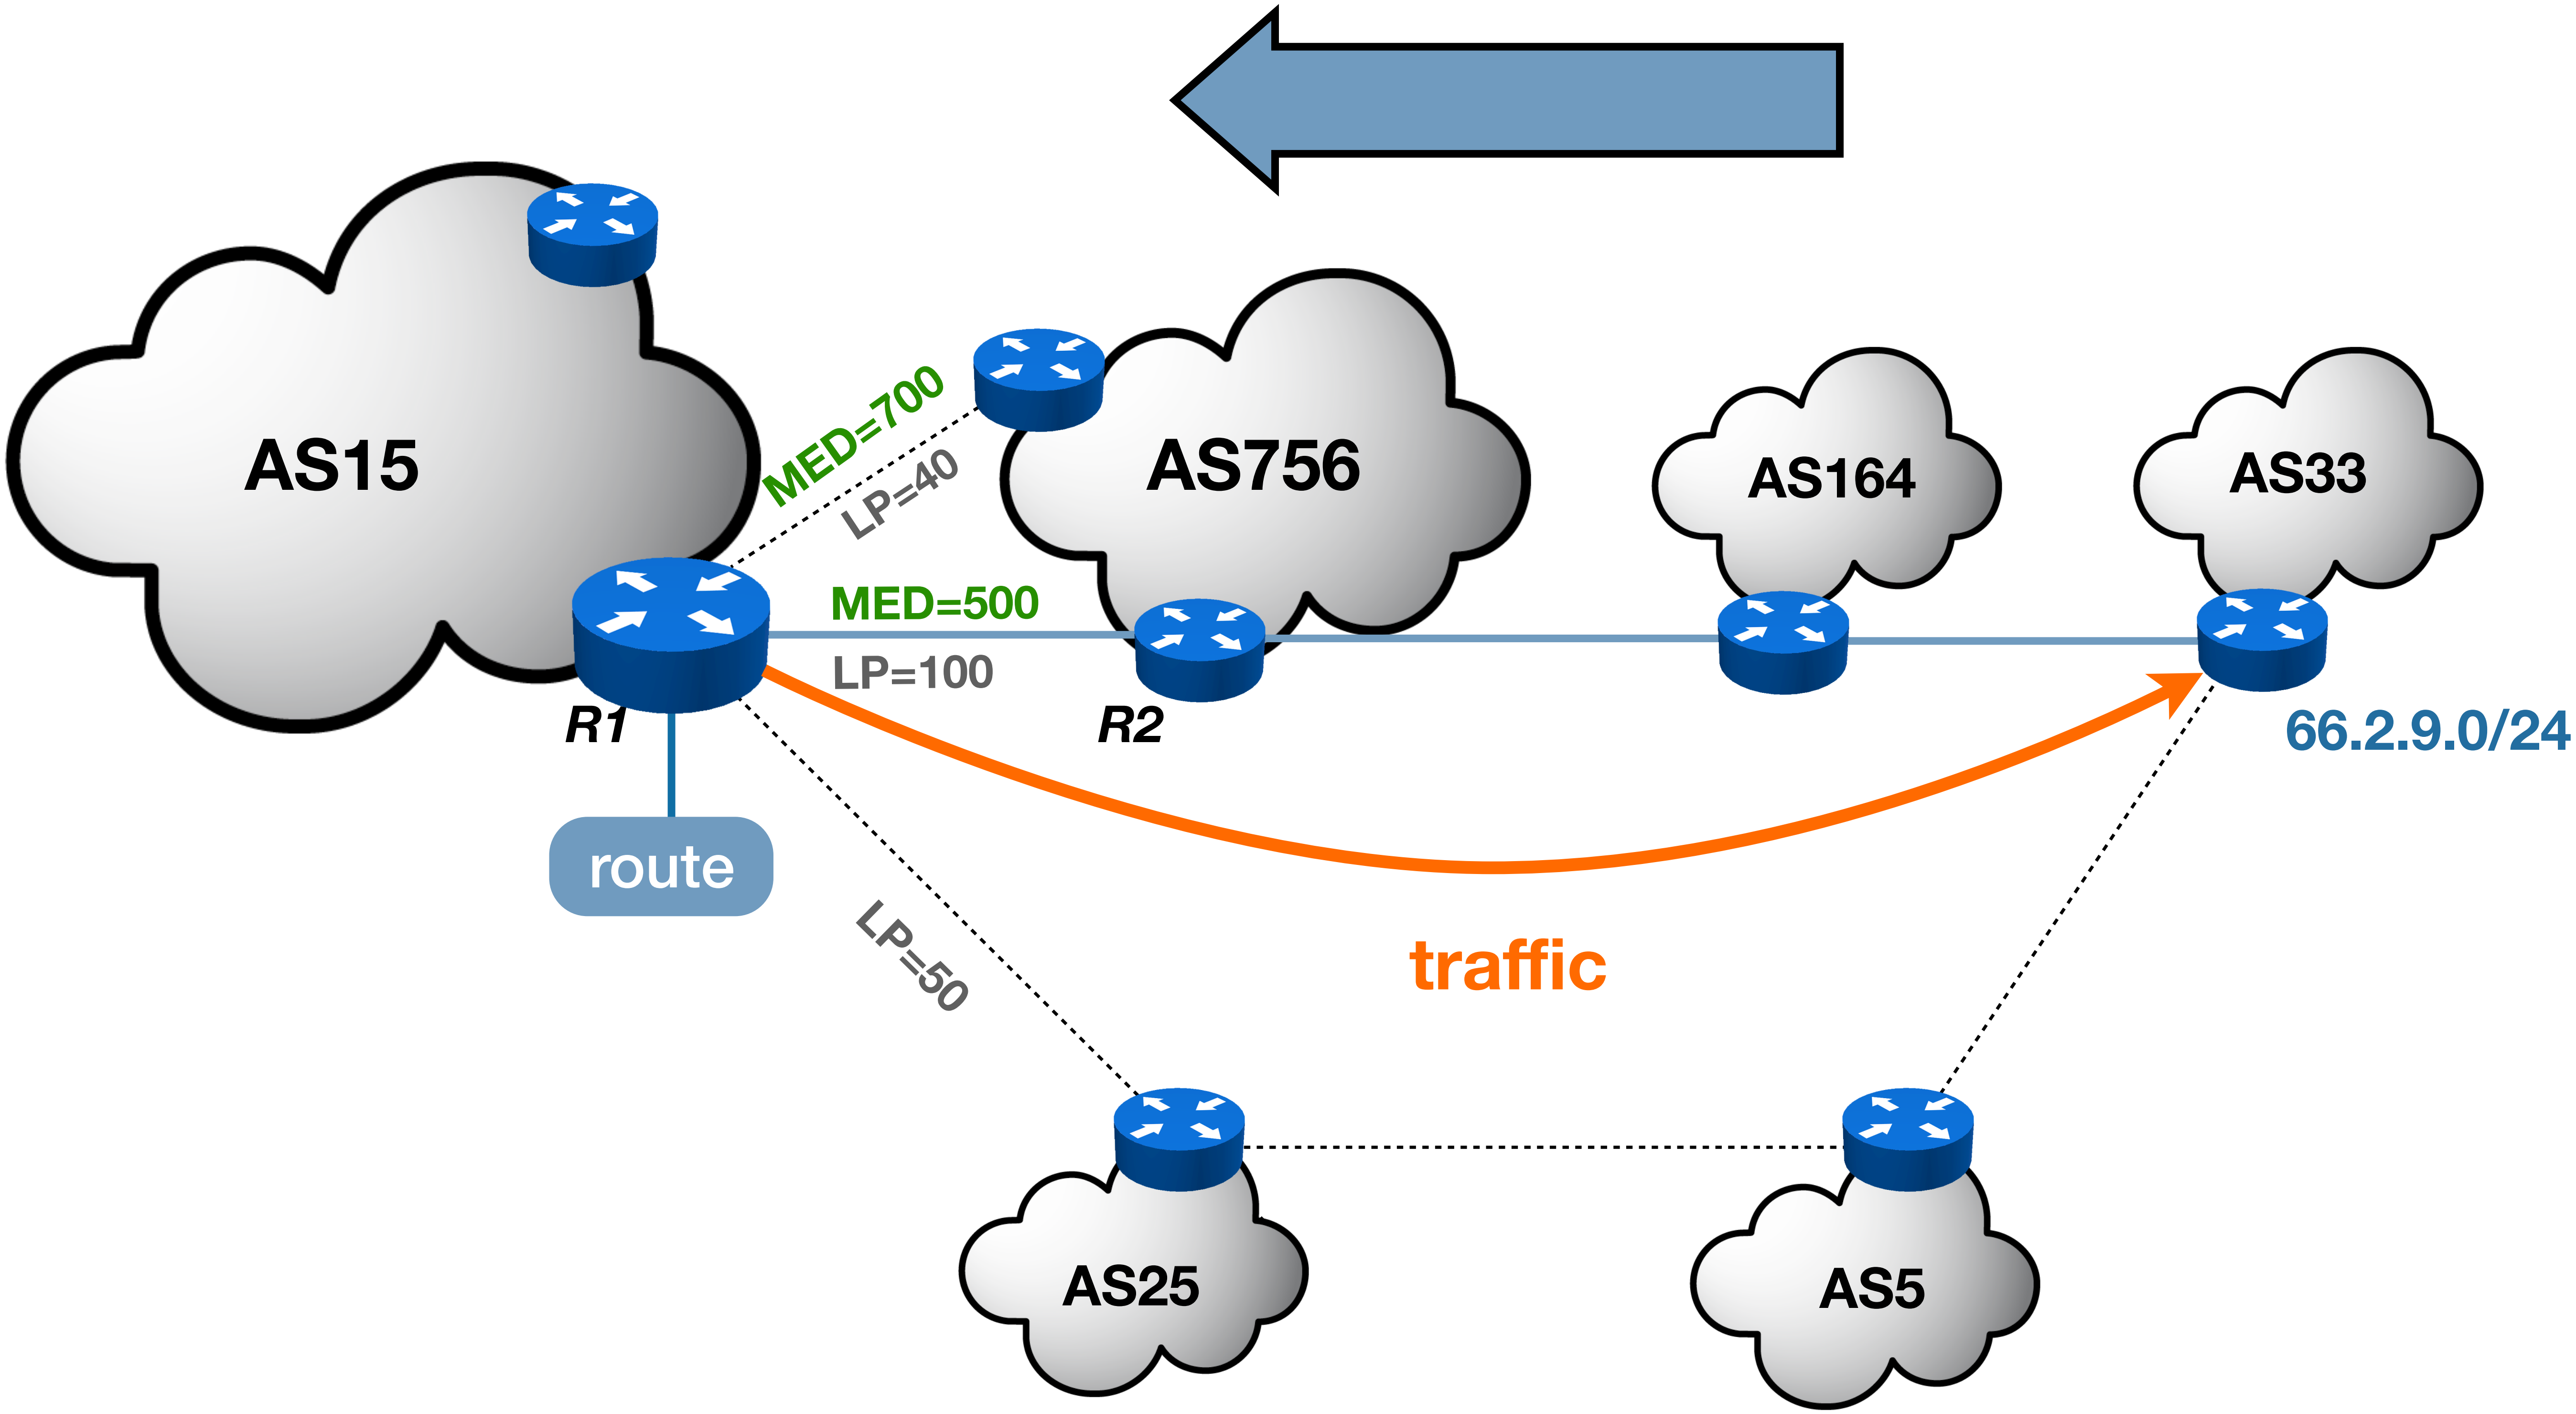

Accidents Happen

- Fat Fingers
  - 2 and 3 are really close on our keyboards….

- Policy Violations (leaks)
  - Oops, we did not want this to go on the public Internet
  - Infamous incident with Pakistan Telecom and YouTube

Incidents Are Common

- 2019 Routing Security Review
  - 12,600 incidents
  - 4,4% of all ASNs affected
  - 3,000 ASNs are victims of at least one incident
  - 1,300 ASNs caused at least one incident

Source: https://bgpstream.com

How Bad Is It?

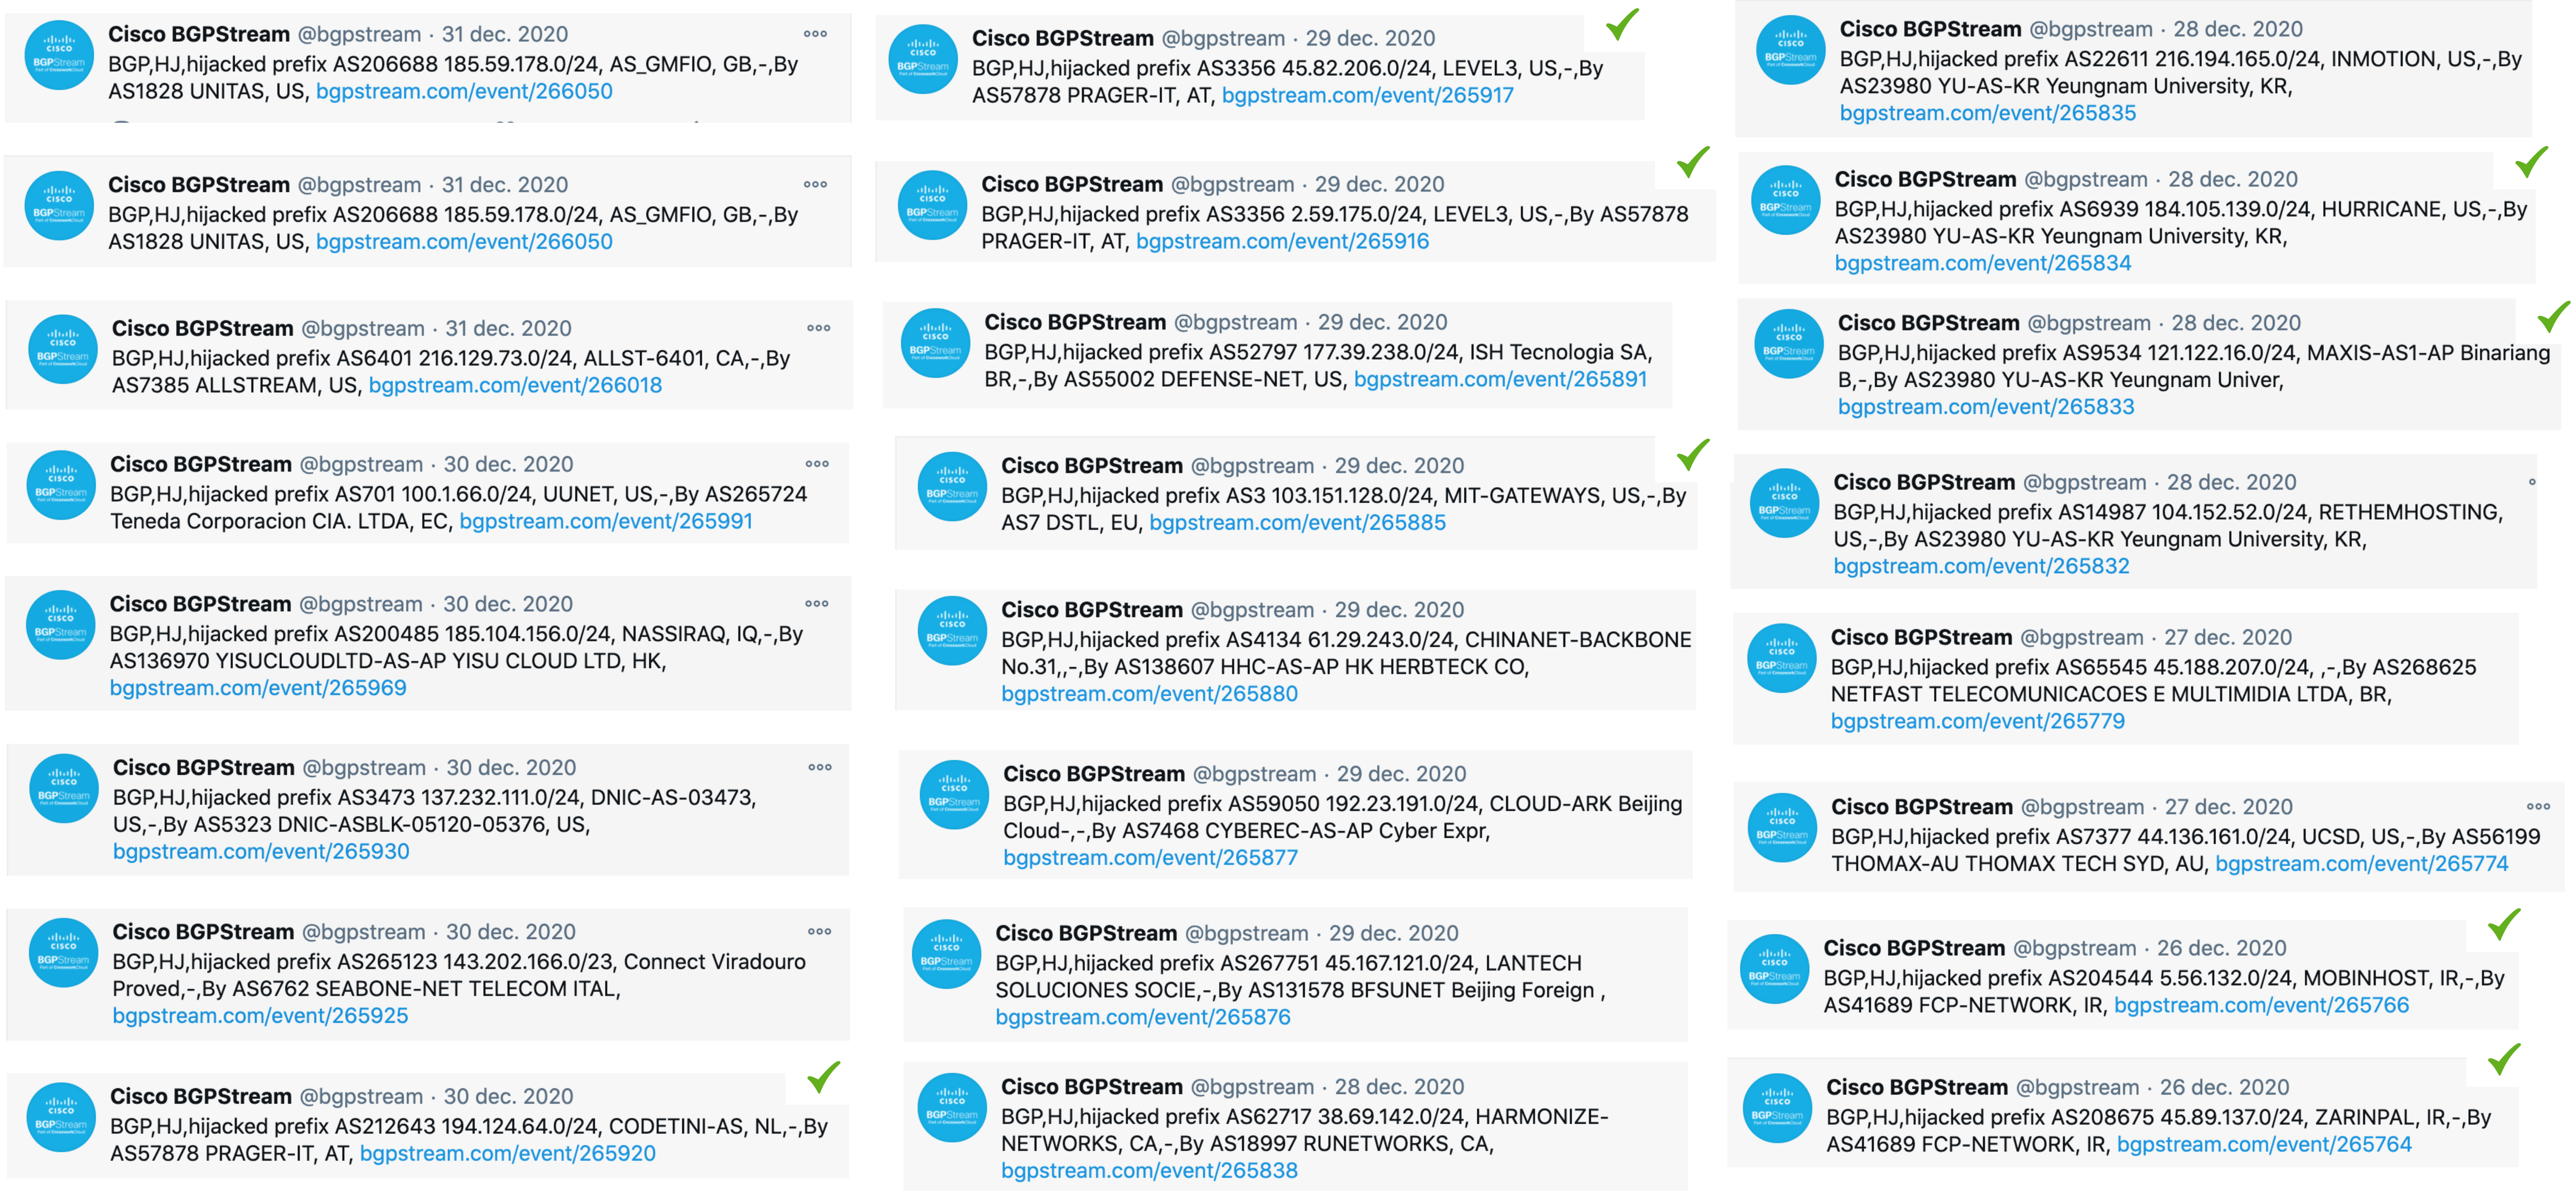

Cisco BGPStream @bgpstream · 31 dec. 2020
BGP,HJ,hijacked prefix AS206688 185.59.178.0/24, AS_GMFIO, GB,-,By AS1828 UNITAS, US, bgpstream.com/event/266050

Cisco BGPStream @bgpstream · 31 dec. 2020
BGP,HJ,hijacked prefix AS206688 185.59.178.0/24, AS_GMFIO, GB,-,By AS1828 UNITAS, US, bgpstream.com/event/266050

Cisco BGPStream @bgpstream · 31 dec. 2020
BGP,HJ,hijacked prefix AS6401 216.129.73.0/24, ALLST-6401, CA,-,By AS7385 ALLSTREAM, US, bgpstream.com/event/266018

Cisco BGPStream @bgpstream · 30 dec. 2020
BGP,HJ,hijacked prefix AS701 100.1.66.0/24, UUNET, US,-,By AS265724 Teneda Corporacion CIA. LTDA, EC, bgpstream.com/event/265991

Cisco BGPStream @bgpstream · 30 dec. 2020
BGP,HJ,hijacked prefix AS200485 185.104.156.0/24, NASSIRAQ, IQ,-,By AS136970 YISUCLOUDLTD-AS-AP YISU CLOUD LTD, HK, bgpstream.com/event/265969

Cisco BGPStream @bgpstream · 30 dec. 2020
BGP,HJ,hijacked prefix AS3473 137.232.111.0/24, DNIC-AS-03473, US,-,By AS5323 DNIC-ASBLK-05120-05376, US, bgpstream.com/event/265930

Cisco BGPStream @bgpstream · 30 dec. 2020
BGP,HJ,hijacked prefix AS265123 143.202.166.0/23, Connect Viradouro Proved,-,By AS6762 SEABONE-NET TELECOM ITAL, bgpstream.com/event/265925

✓ **Cisco BGPStream** @bgpstream · 30 dec. 2020
BGP,HJ,hijacked prefix AS212643 194.124.64.0/24, CODETINI-AS, NL,-,By AS57878 PRAGER-IT, AT, bgpstream.com/event/265920

✓ **Cisco BGPStream** @bgpstream · 29 dec. 2020
BGP,HJ,hijacked prefix AS3356 45.82.206.0/24, LEVEL3, US,-,By AS57878 PRAGER-IT, AT, bgpstream.com/event/265917

✓ **Cisco BGPStream** @bgpstream · 29 dec. 2020
BGP,HJ,hijacked prefix AS3356 2.59.175.0/24, LEVEL3, US,-,By AS57878 PRAGER-IT, AT, bgpstream.com/event/265916

Cisco BGPStream @bgpstream · 29 dec. 2020
BGP,HJ,hijacked prefix AS52797 177.39.238.0/24, ISH Tecnologia SA, BR,-,By AS55002 DEFENSE-NET, US, bgpstream.com/event/265891

✓ **Cisco BGPStream** @bgpstream · 29 dec. 2020
BGP,HJ,hijacked prefix AS3 103.151.128.0/24, MIT-GATEWAYS, US,-,By AS7 DSTL, EU, bgpstream.com/event/265885

Cisco BGPStream @bgpstream · 29 dec. 2020
BGP,HJ,hijacked prefix AS4134 61.29.243.0/24, CHINANET-BACKBONE No.31,,-,By AS138607 HHC-AS-AP HK HERBTECK CO, bgpstream.com/event/265880

Cisco BGPStream @bgpstream · 29 dec. 2020
BGP,HJ,hijacked prefix AS59050 192.23.191.0/24, CLOUD-ARK Beijing Cloud-,-,By AS7468 CYBEREC-AS-AP Cyber Expr, bgpstream.com/event/265877

Cisco BGPStream @bgpstream · 29 dec. 2020
BGP,HJ,hijacked prefix AS267751 45.167.121.0/24, LANTECH SOLUCIONES SOCIE,-,By AS131578 BFSUNET Beijing Foreign , bgpstream.com/event/265876

Cisco BGPStream @bgpstream · 28 dec. 2020
BGP,HJ,hijacked prefix AS62717 38.69.142.0/24, HARMONIZE-NETWORKS, CA,-,By AS18997 RUNETWORKS, CA, bgpstream.com/event/265838

Cisco BGPStream @bgpstream · 28 dec. 2020
BGP,HJ,hijacked prefix AS22611 216.194.165.0/24, INMOTION, US,-,By AS23980 YU-AS-KR Yeungnam University, KR, bgpstream.com/event/265835

✓ **Cisco BGPStream** @bgpstream · 28 dec. 2020
BGP,HJ,hijacked prefix AS6939 184.105.139.0/24, HURRICANE, US,-,By AS23980 YU-AS-KR Yeungnam University, KR, bgpstream.com/event/265834

✓ **Cisco BGPStream** @bgpstream · 28 dec. 2020
BGP,HJ,hijacked prefix AS9534 121.122.16.0/24, MAXIS-AS1-AP Binariang B,-,By AS23980 YU-AS-KR Yeungnam Univer, bgpstream.com/event/265833

Cisco BGPStream @bgpstream · 28 dec. 2020
BGP,HJ,hijacked prefix AS14987 104.152.52.0/24, RETHEMHOSTING, US,-,By AS23980 YU-AS-KR Yeungnam University, KR, bgpstream.com/event/265832

Cisco BGPStream @bgpstream · 27 dec. 2020
BGP,HJ,hijacked prefix AS65545 45.188.207.0/24, ,-,By AS268625 NETFAST TELECOMUNICACOES E MULTIMIDIA LTDA, BR, bgpstream.com/event/265779

Cisco BGPStream @bgpstream · 27 dec. 2020
BGP,HJ,hijacked prefix AS7377 44.136.161.0/24, UCSD, US,-,By AS56199 THOMAX-AU THOMAX TECH SYD, AU, bgpstream.com/event/265774

✓ **Cisco BGPStream** @bgpstream · 26 dec. 2020
BGP,HJ,hijacked prefix AS204544 5.56.132.0/24, MOBINHOST, IR,-,By AS41689 FCP-NETWORK, IR, bgpstream.com/event/265766

✓ **Cisco BGPStream** @bgpstream · 26 dec. 2020
BGP,HJ,hijacked prefix AS208675 45.89.137.0/24, ZARINPAL, IR,-,By AS41689 FCP-NETWORK, IR, bgpstream.com/event/265764

Routing on the Internet

“Internet Routing Registry”

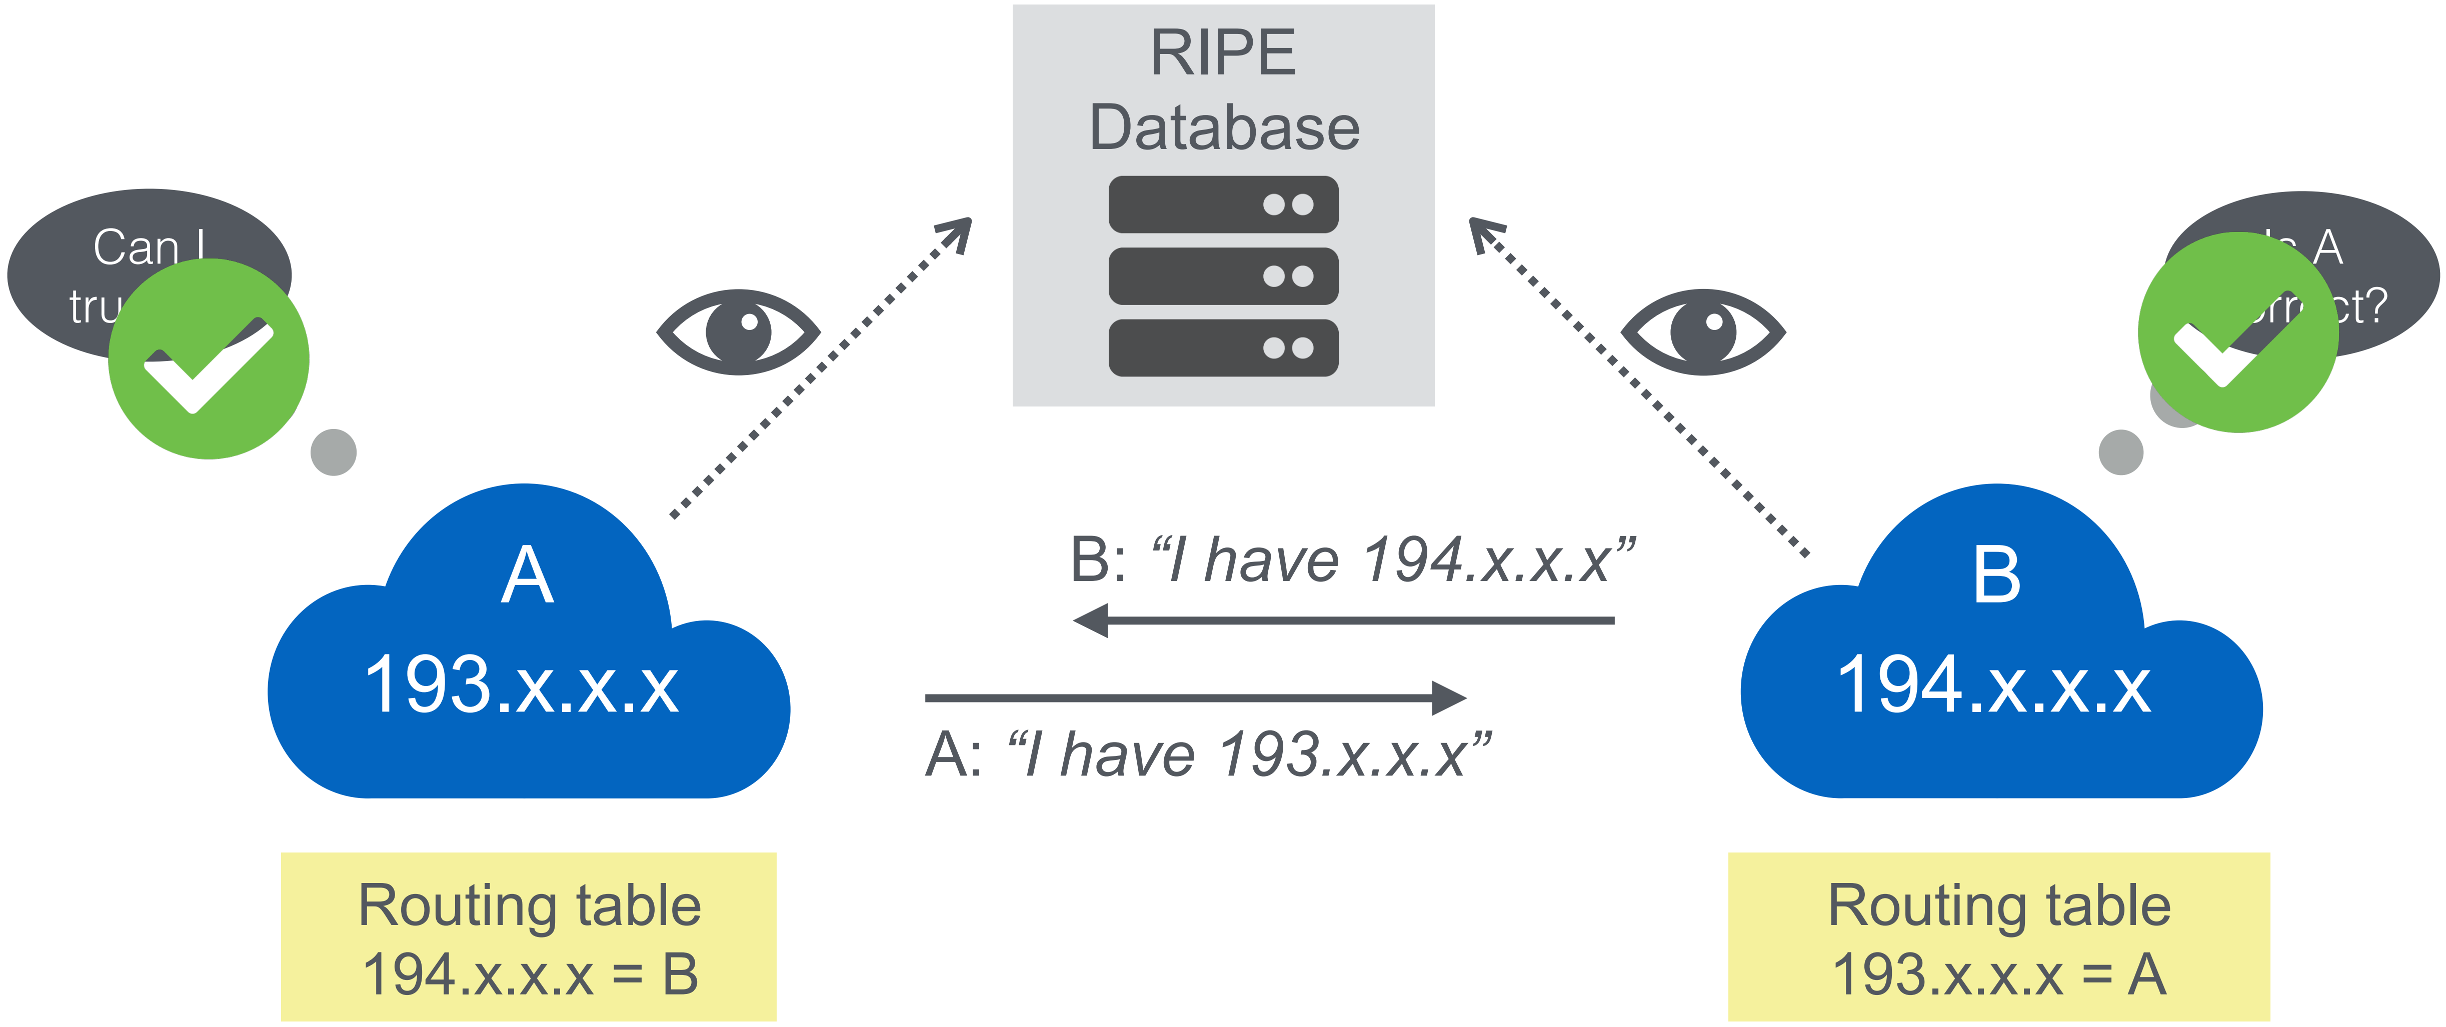

RIPE Database

A

193.x.x.x

B

194.x.x.x

B: *“I have 194.x.x.x”*

A: *“I have 193.x.x.x”*

Routing table
194.x.x.x = B

Routing table
193.x.x.x = A

Problem Statement

- Some IRR data can not be fully trusted
  - Accuracy
  - Incomplete data
  - Lack of maintenance
- Not every RIR has an IRR
  - Third party databases need to be used
  - No verification of who holds IPs/ASNs

Problem Statement

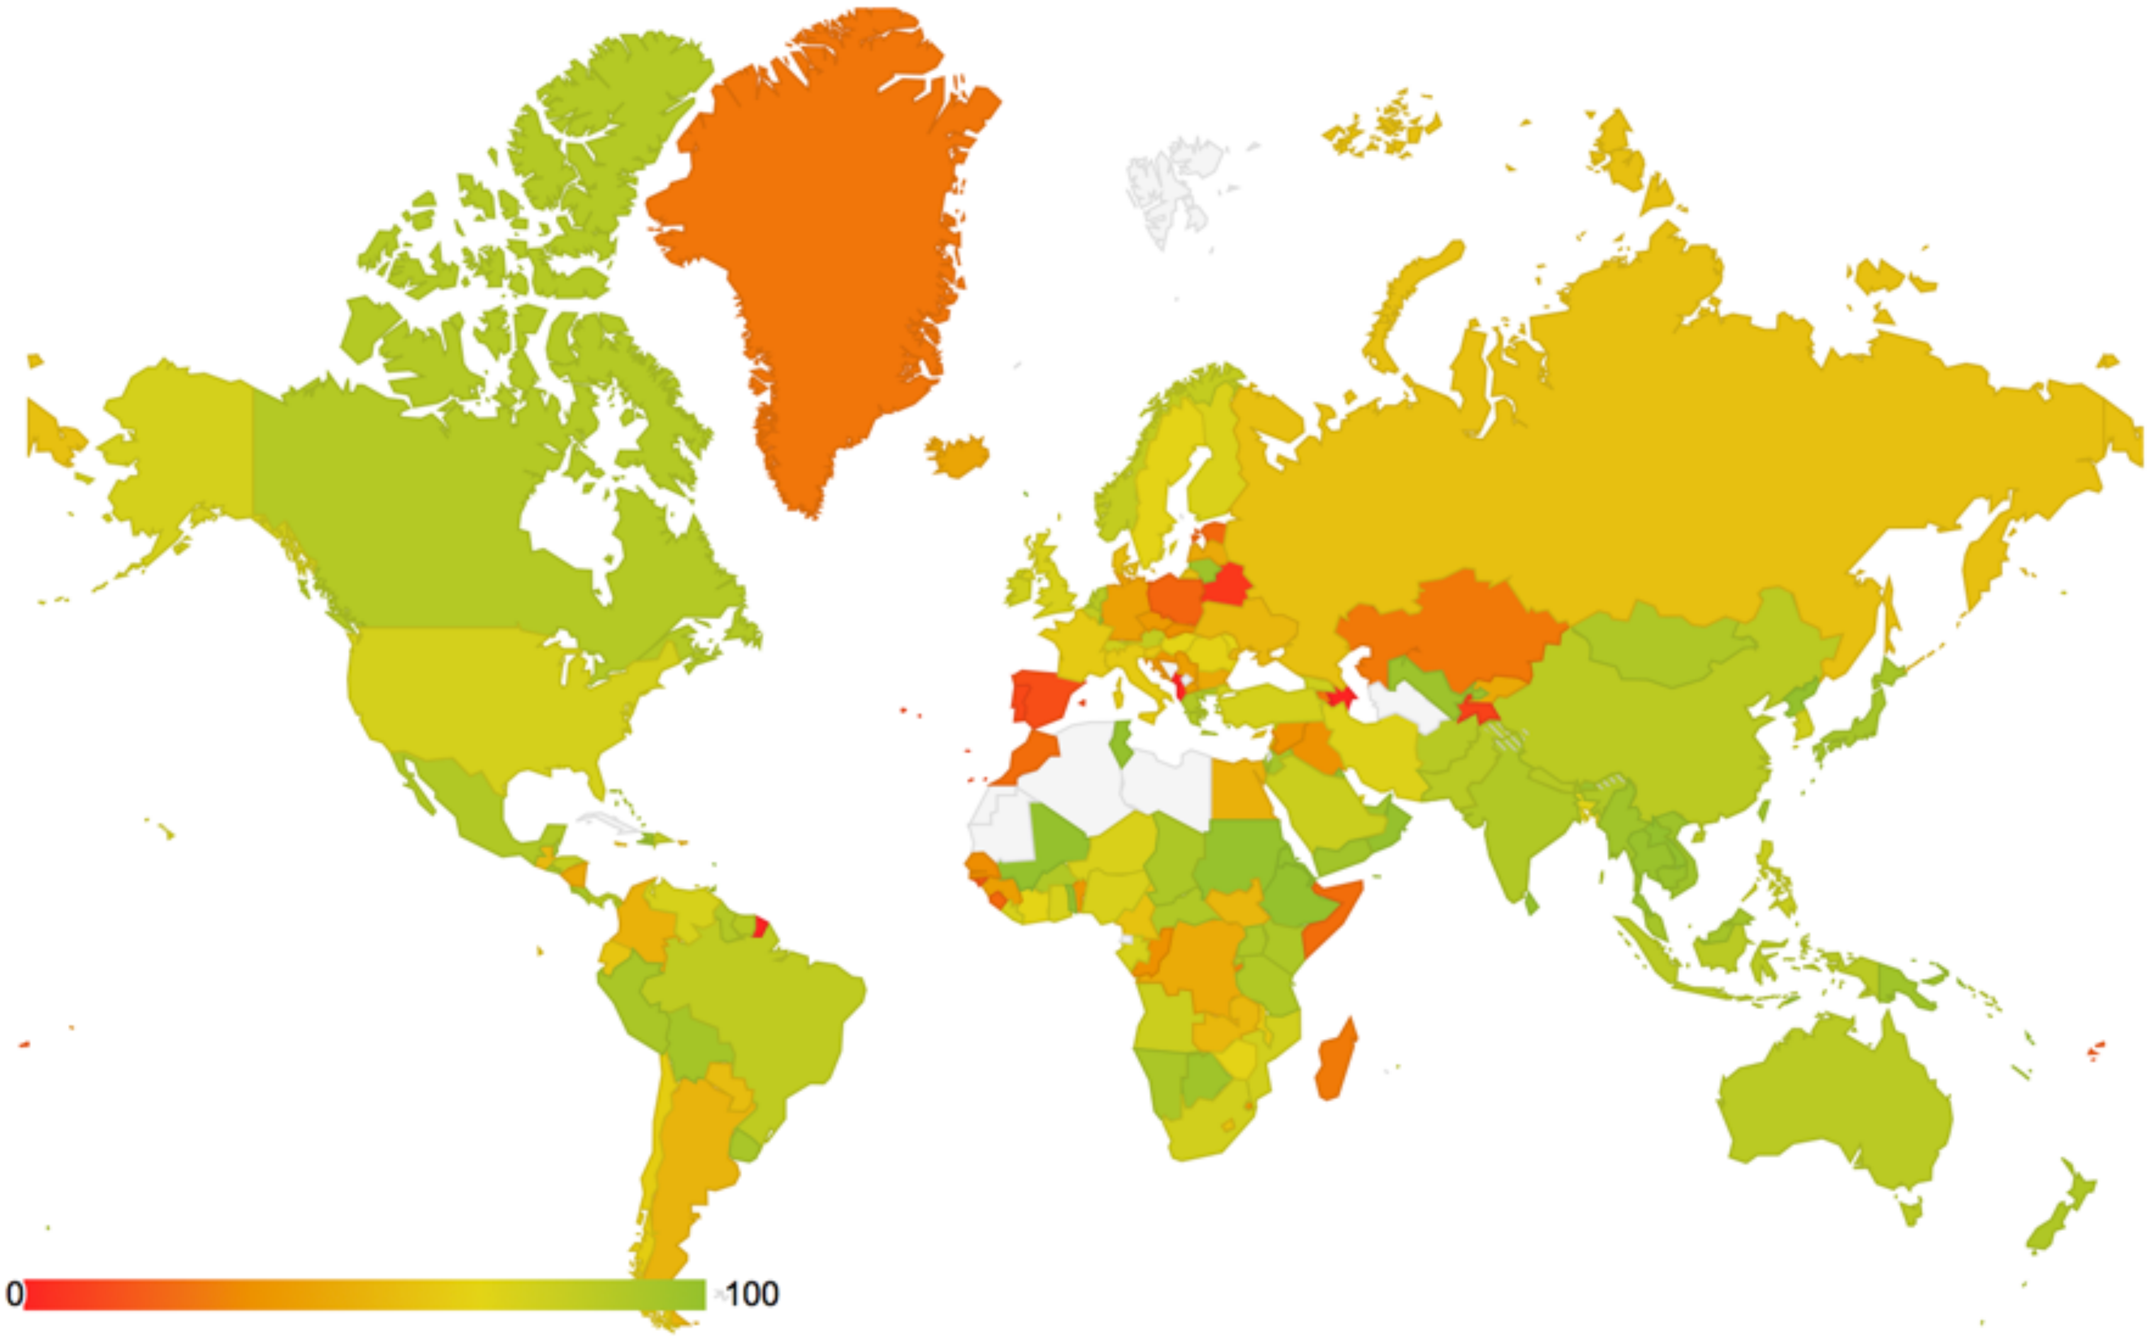

•

Internet Routing Registry

- Many exist, most widely used
  - RIPE Database
  - RADB
- Verification of holdership over resources
  - RIPE Database for RIPE Region resources only
  - RADB allows paying customers to create any object
  - Lots of the other IRRs do not formally verify holdership

Introduction to RPKI

Resource Public Key Infrastructure

- Ties IP addresses and ASNs to public keys
- Follows the hierarchy of the registries
- Authorised statements from resource holders
  - “ASN X is authorised to announce my Prefix Y”
  - Signed, holder of Y

RPKI Certificate Structure

Certificate hierarchy follows allocation hierarchy

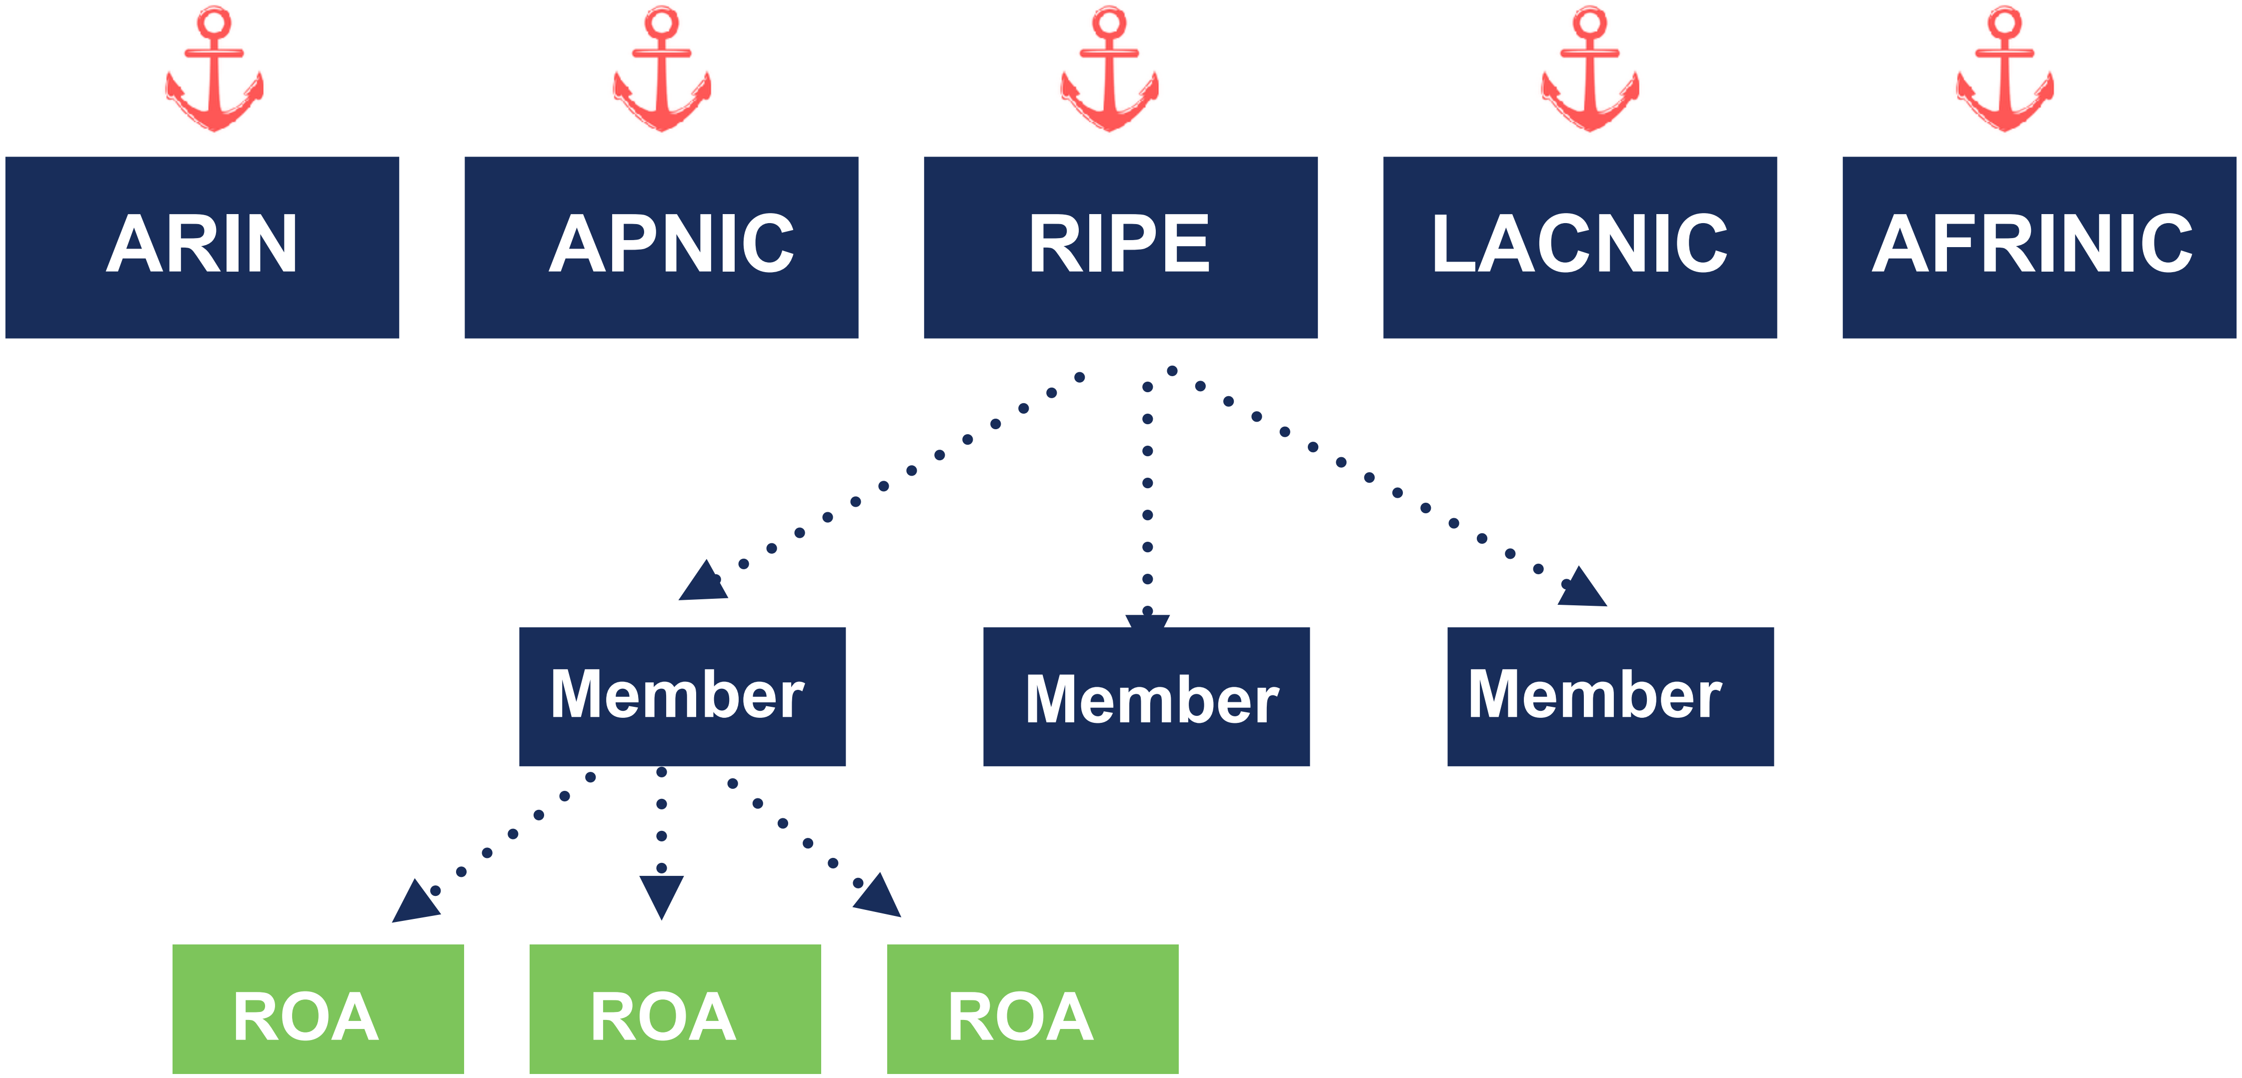

RPKI Chain of Trust

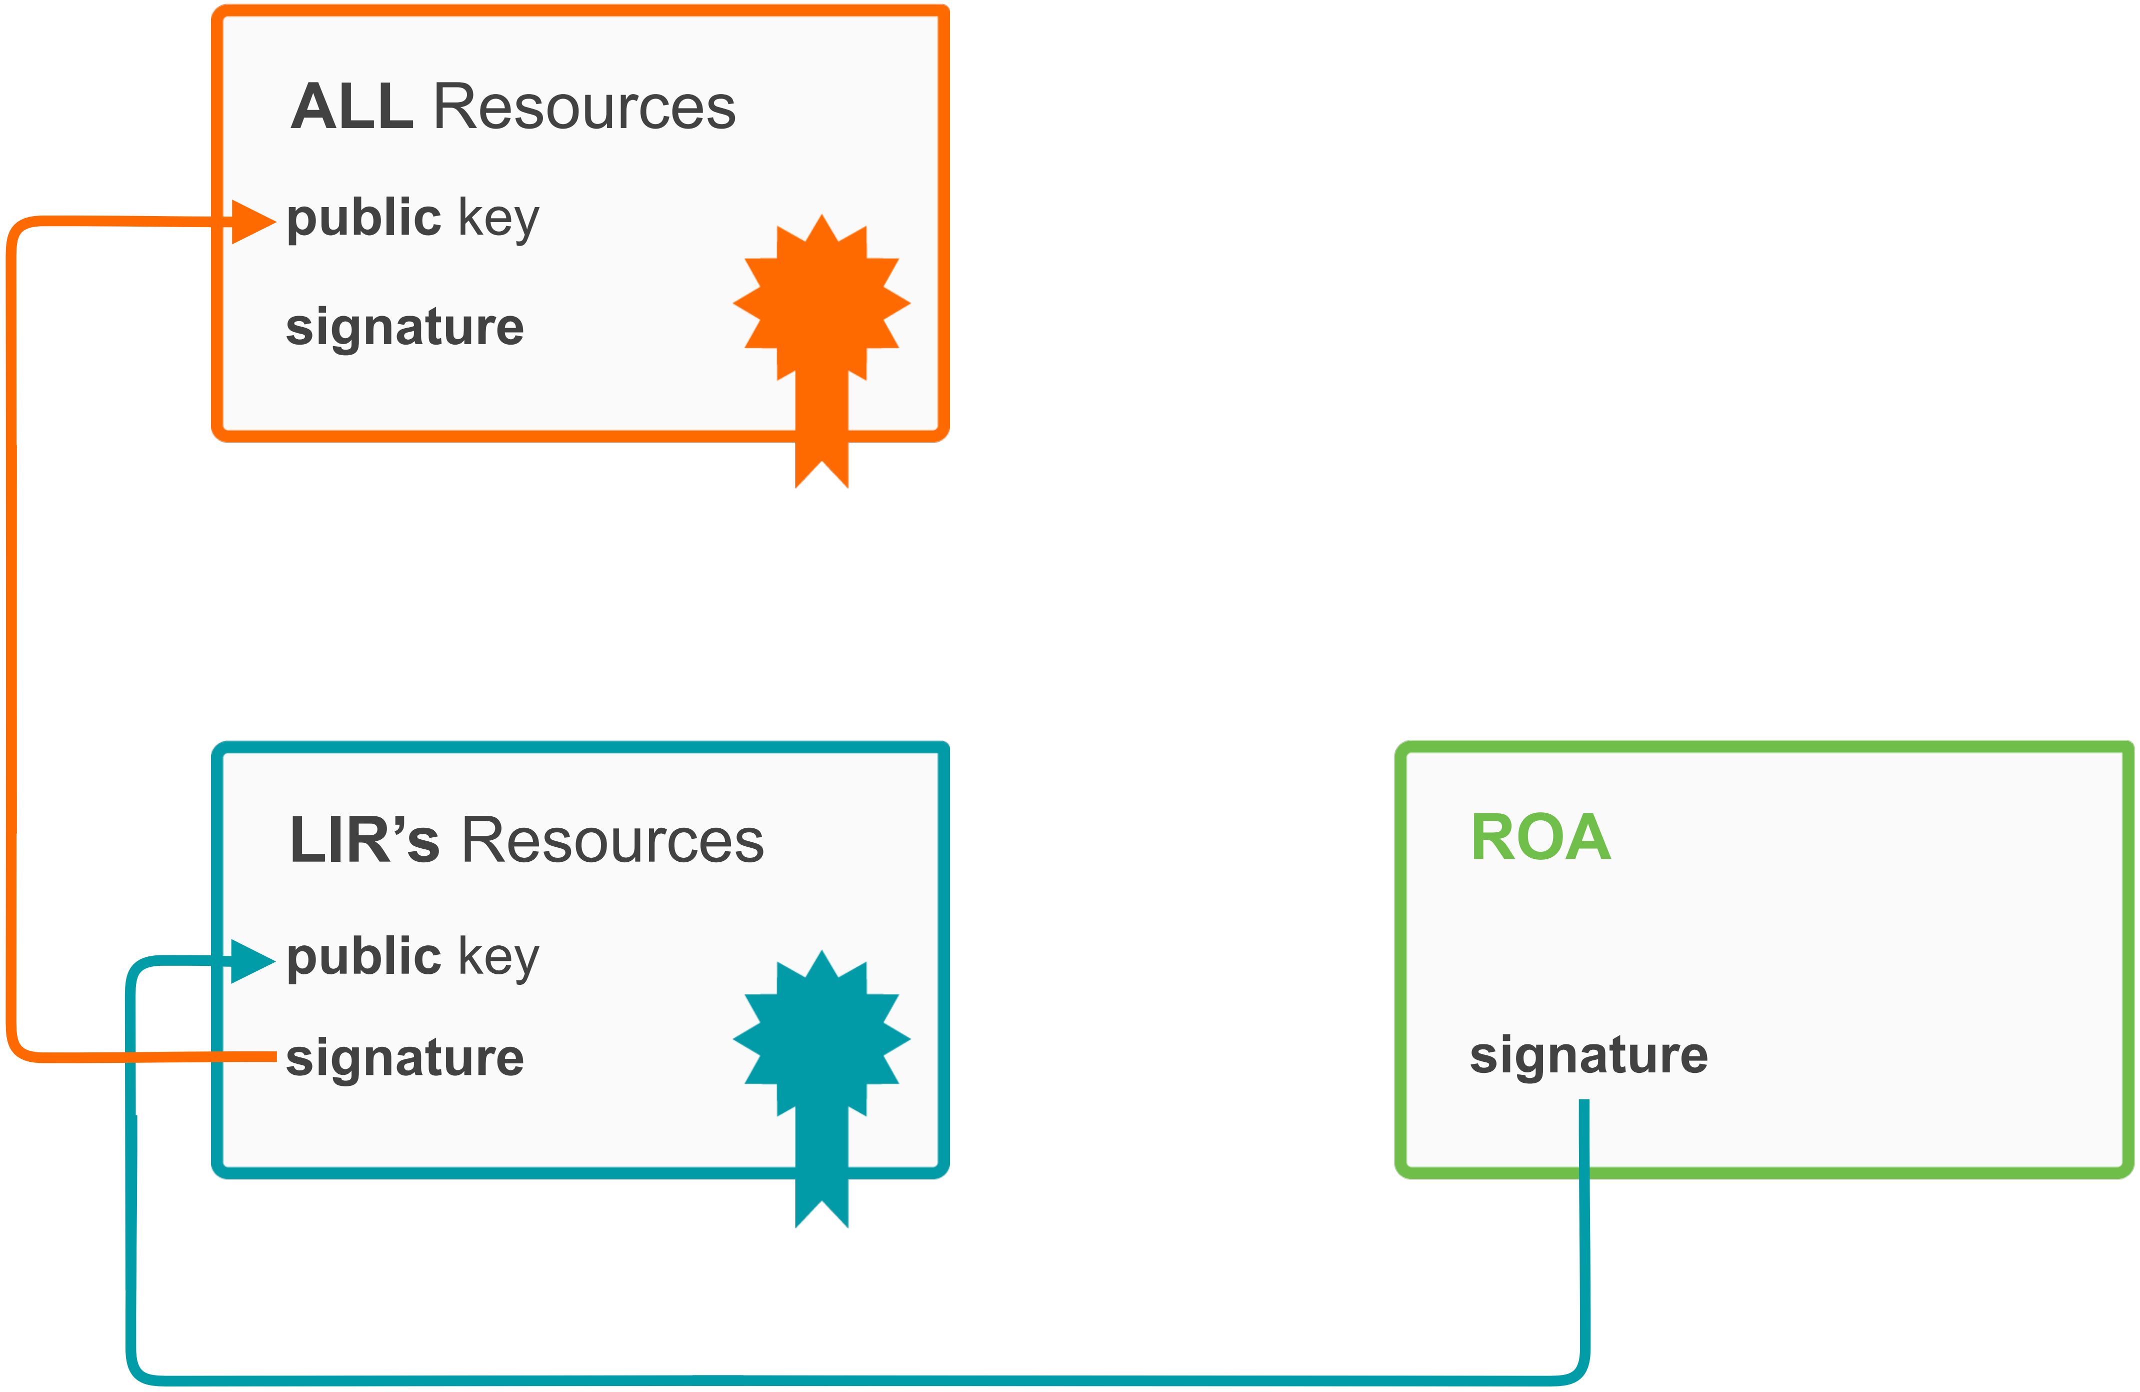

ALL Resources

public key

signature

LIR’s Resources

public key

signature

ROA

signature

RPKI Chain of Trust

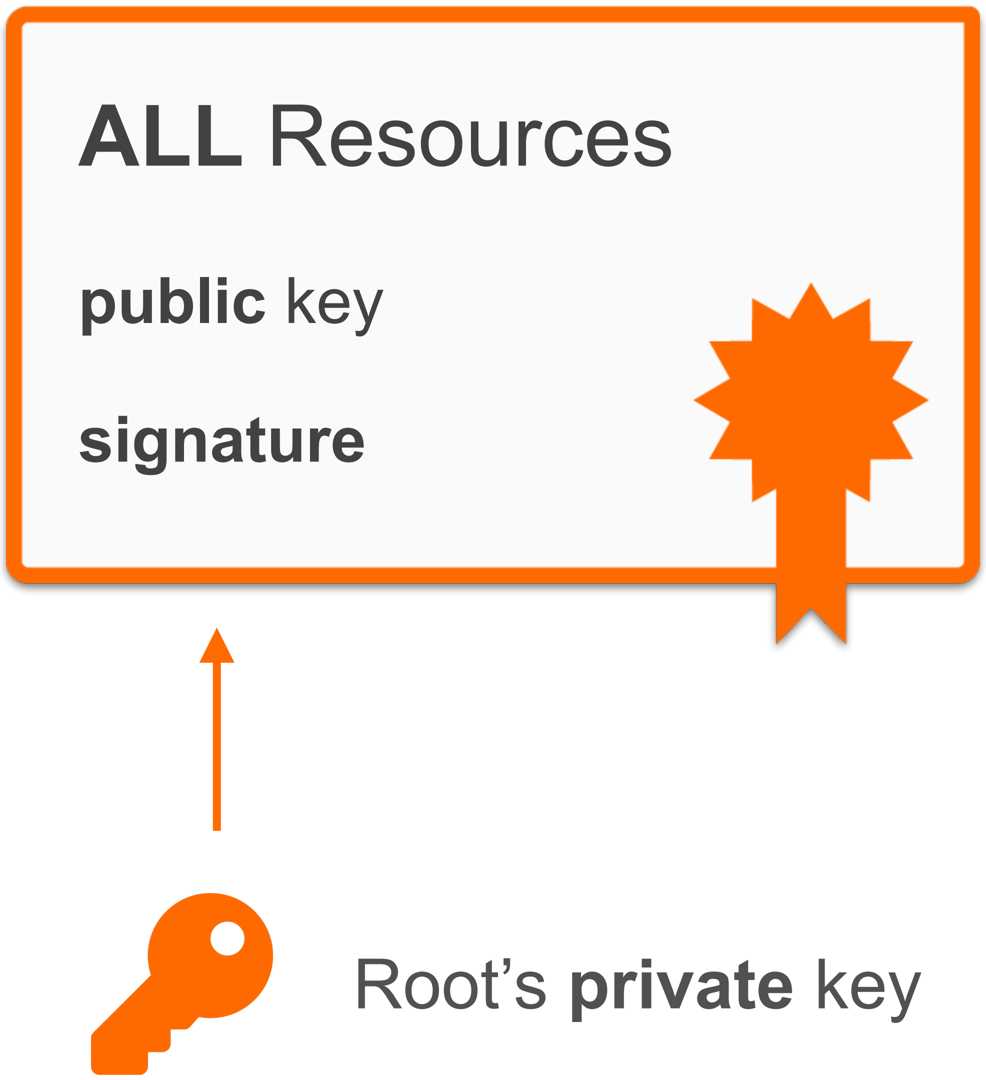

ALL Resources

public key

signature

Root’s **private** key

RIPE NCC Root Certificate

Self-signed

RPKI Chain of Trust

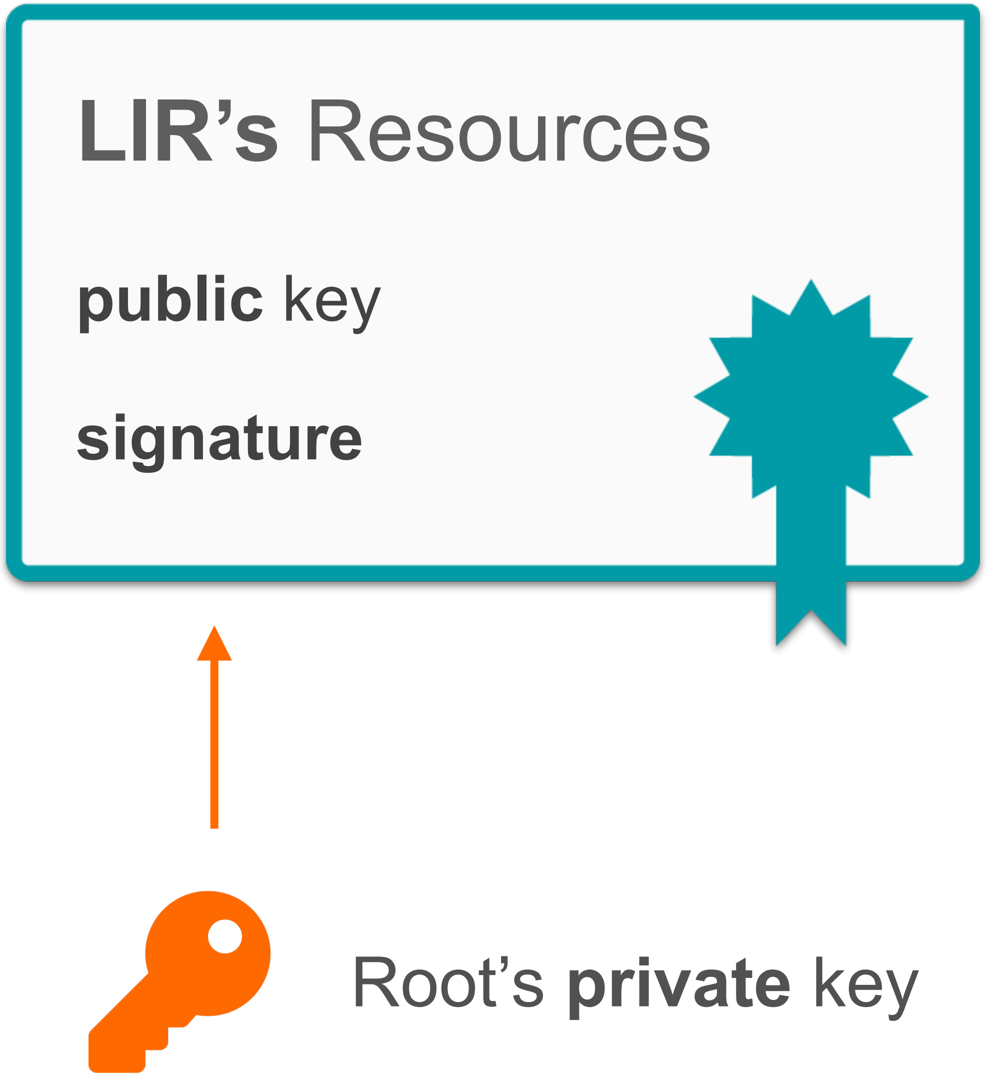

LIR Certificate

Signed by the Root private key

RPKI Adoption

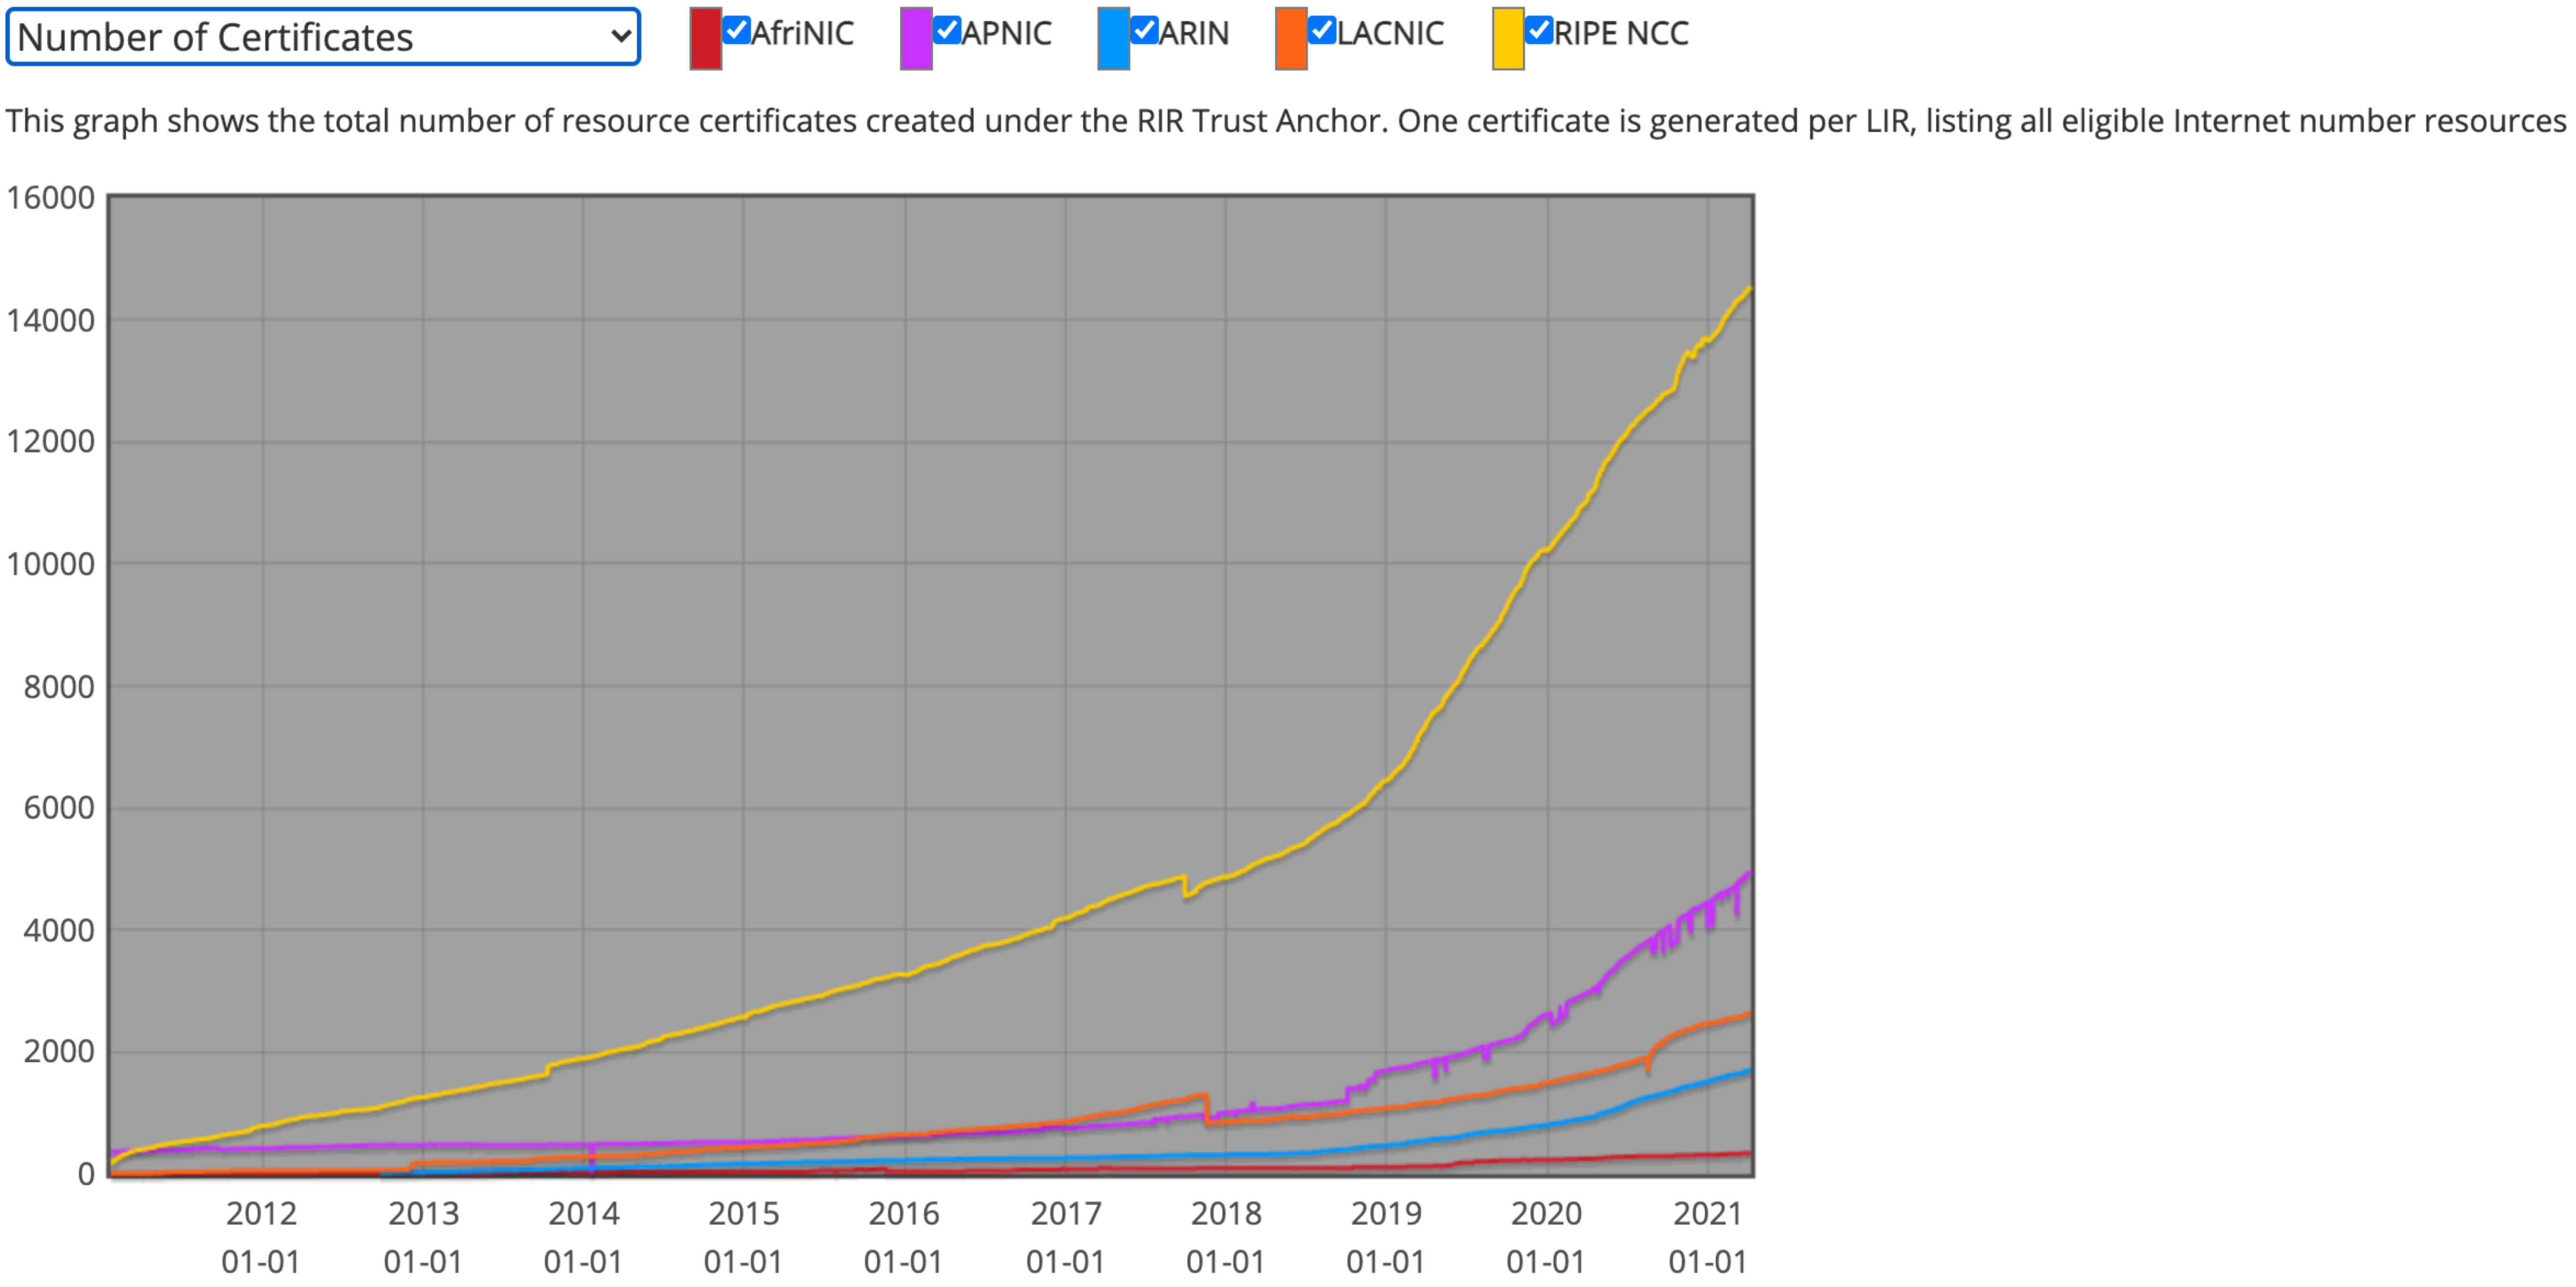

Number of Certificates ⌄

🟥 ☑ AfriNIC 🟪 ☑ APNIC 🟦 ☑ ARIN 🟧 ☑ LACNIC 🟨 ☑ RIPE NCC

This graph shows the total number of resource certificates created under the RIR Trust Anchor. One certificate is generated per LIR, listing all eligible Internet number resources

Two elements of RPKI

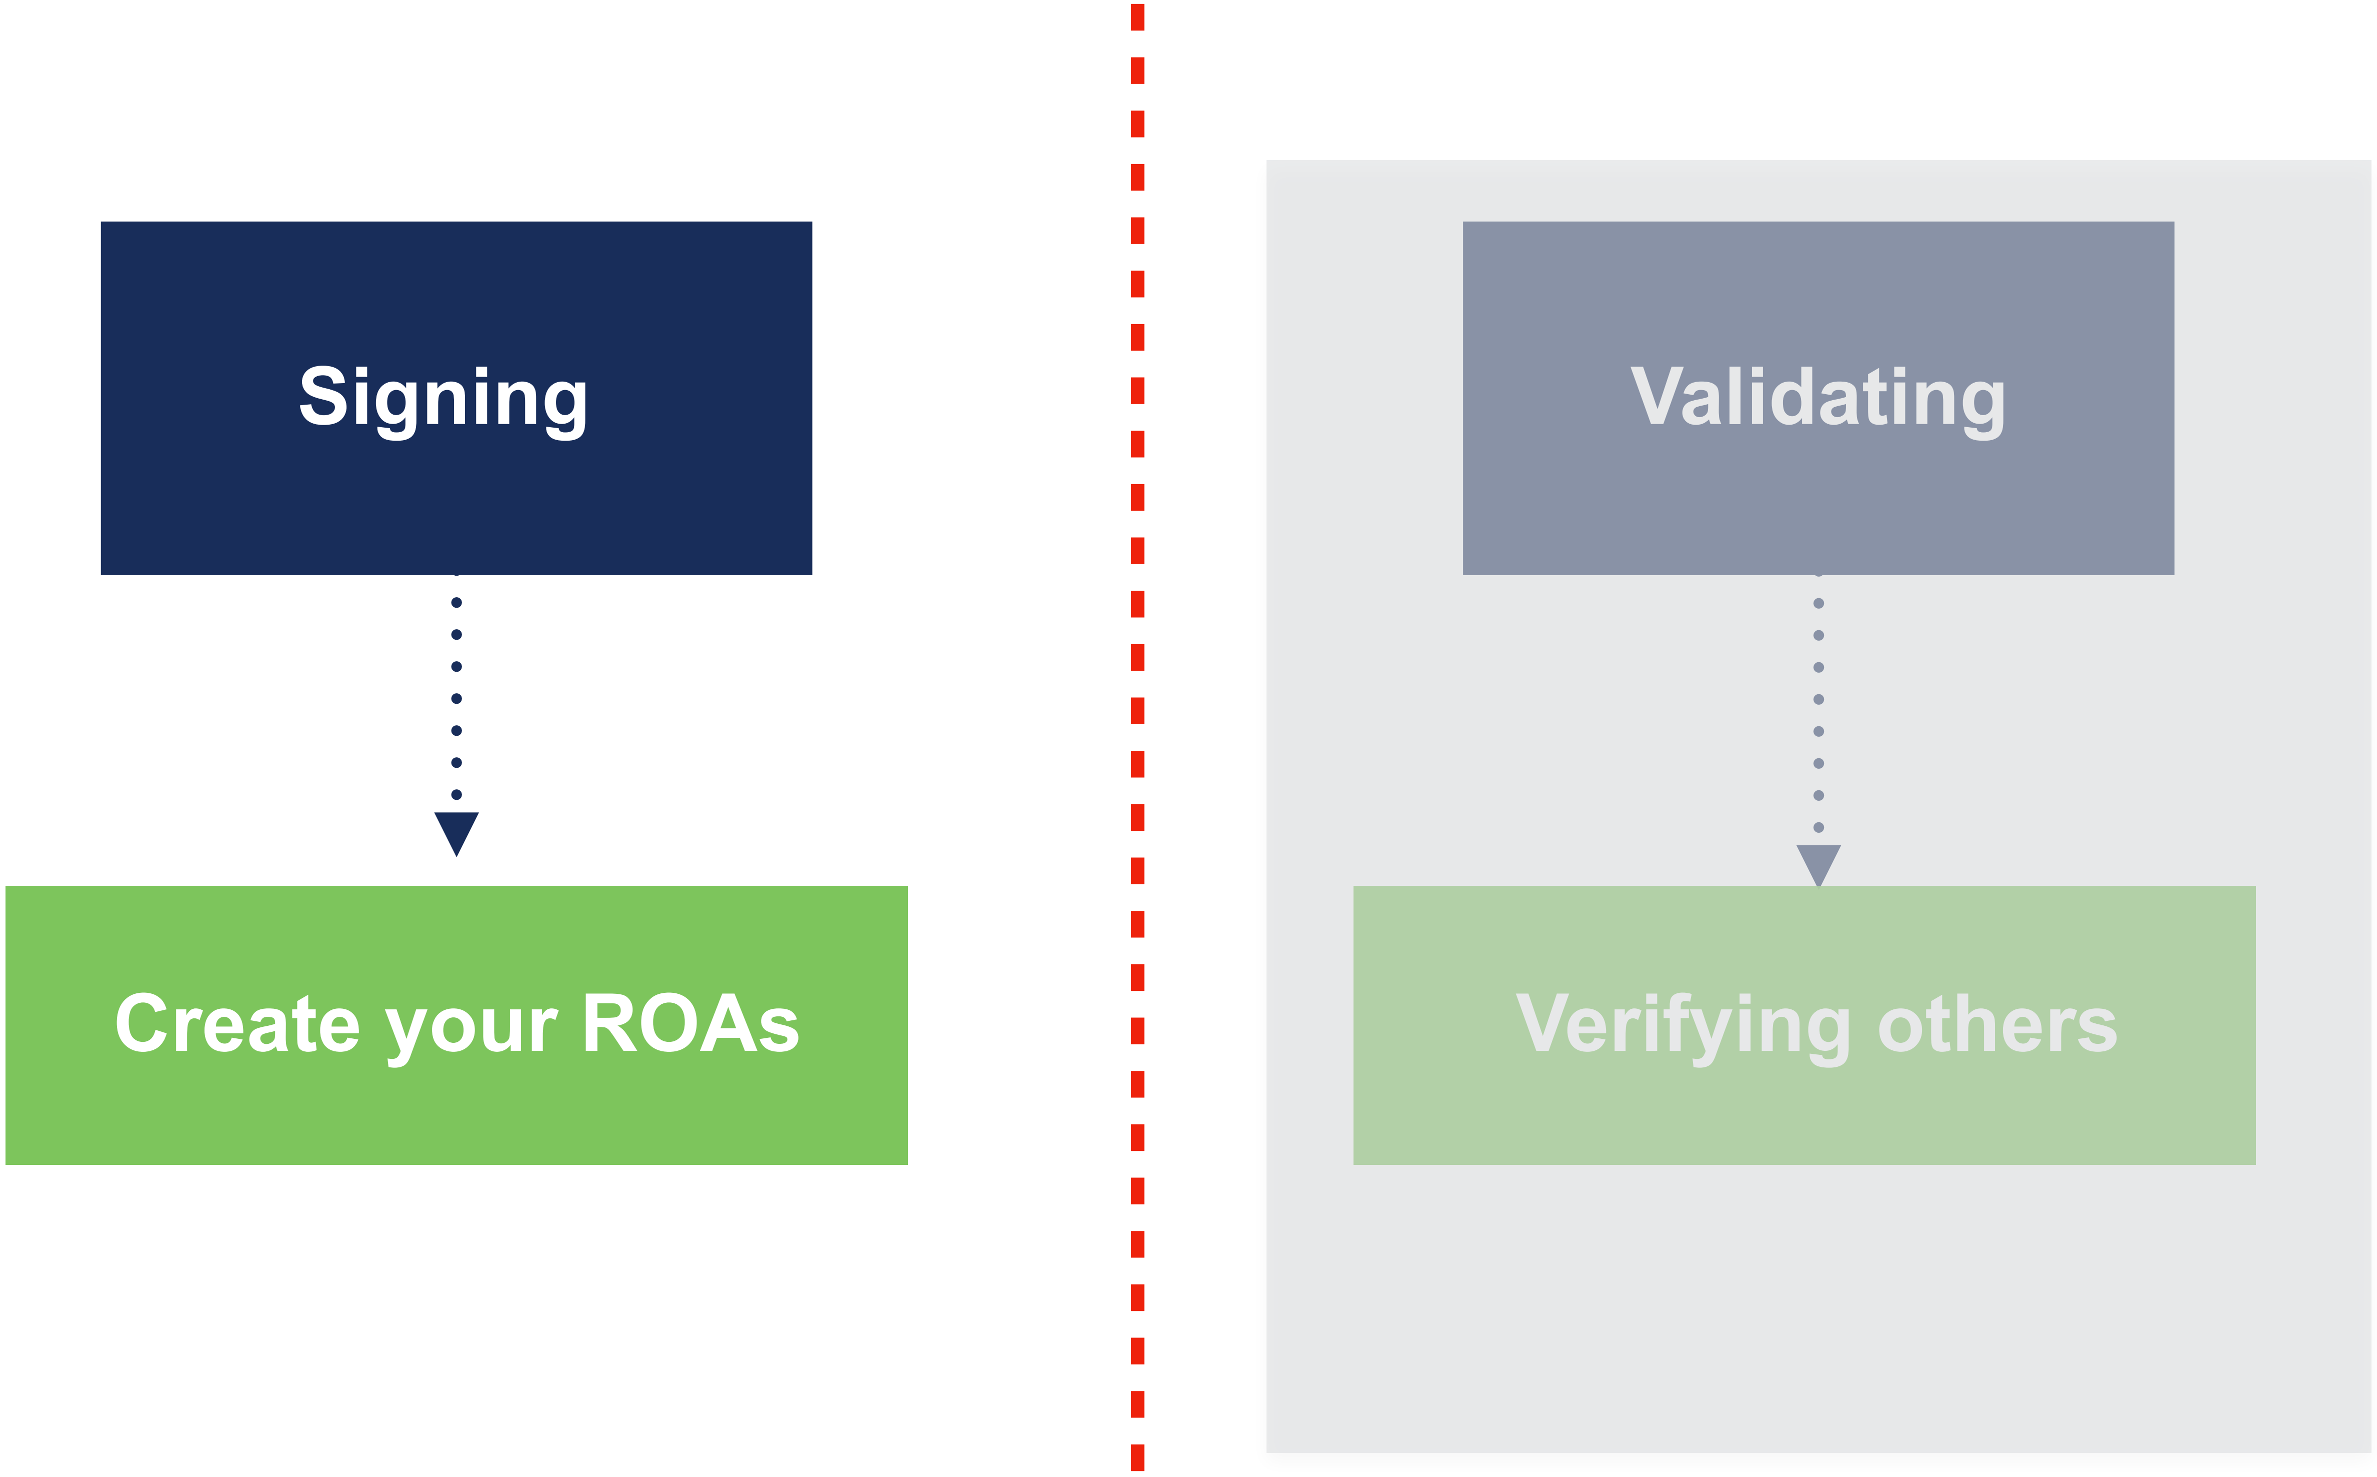

ROAs

ROA (Route Origin Authorisation)

- A ROA is…
- LIRs can create a ROA for each one of their resources (IP address ranges)
- Multiple ROAs can be created for an IP range
- ROAs can overlap

What is in a ROA ?

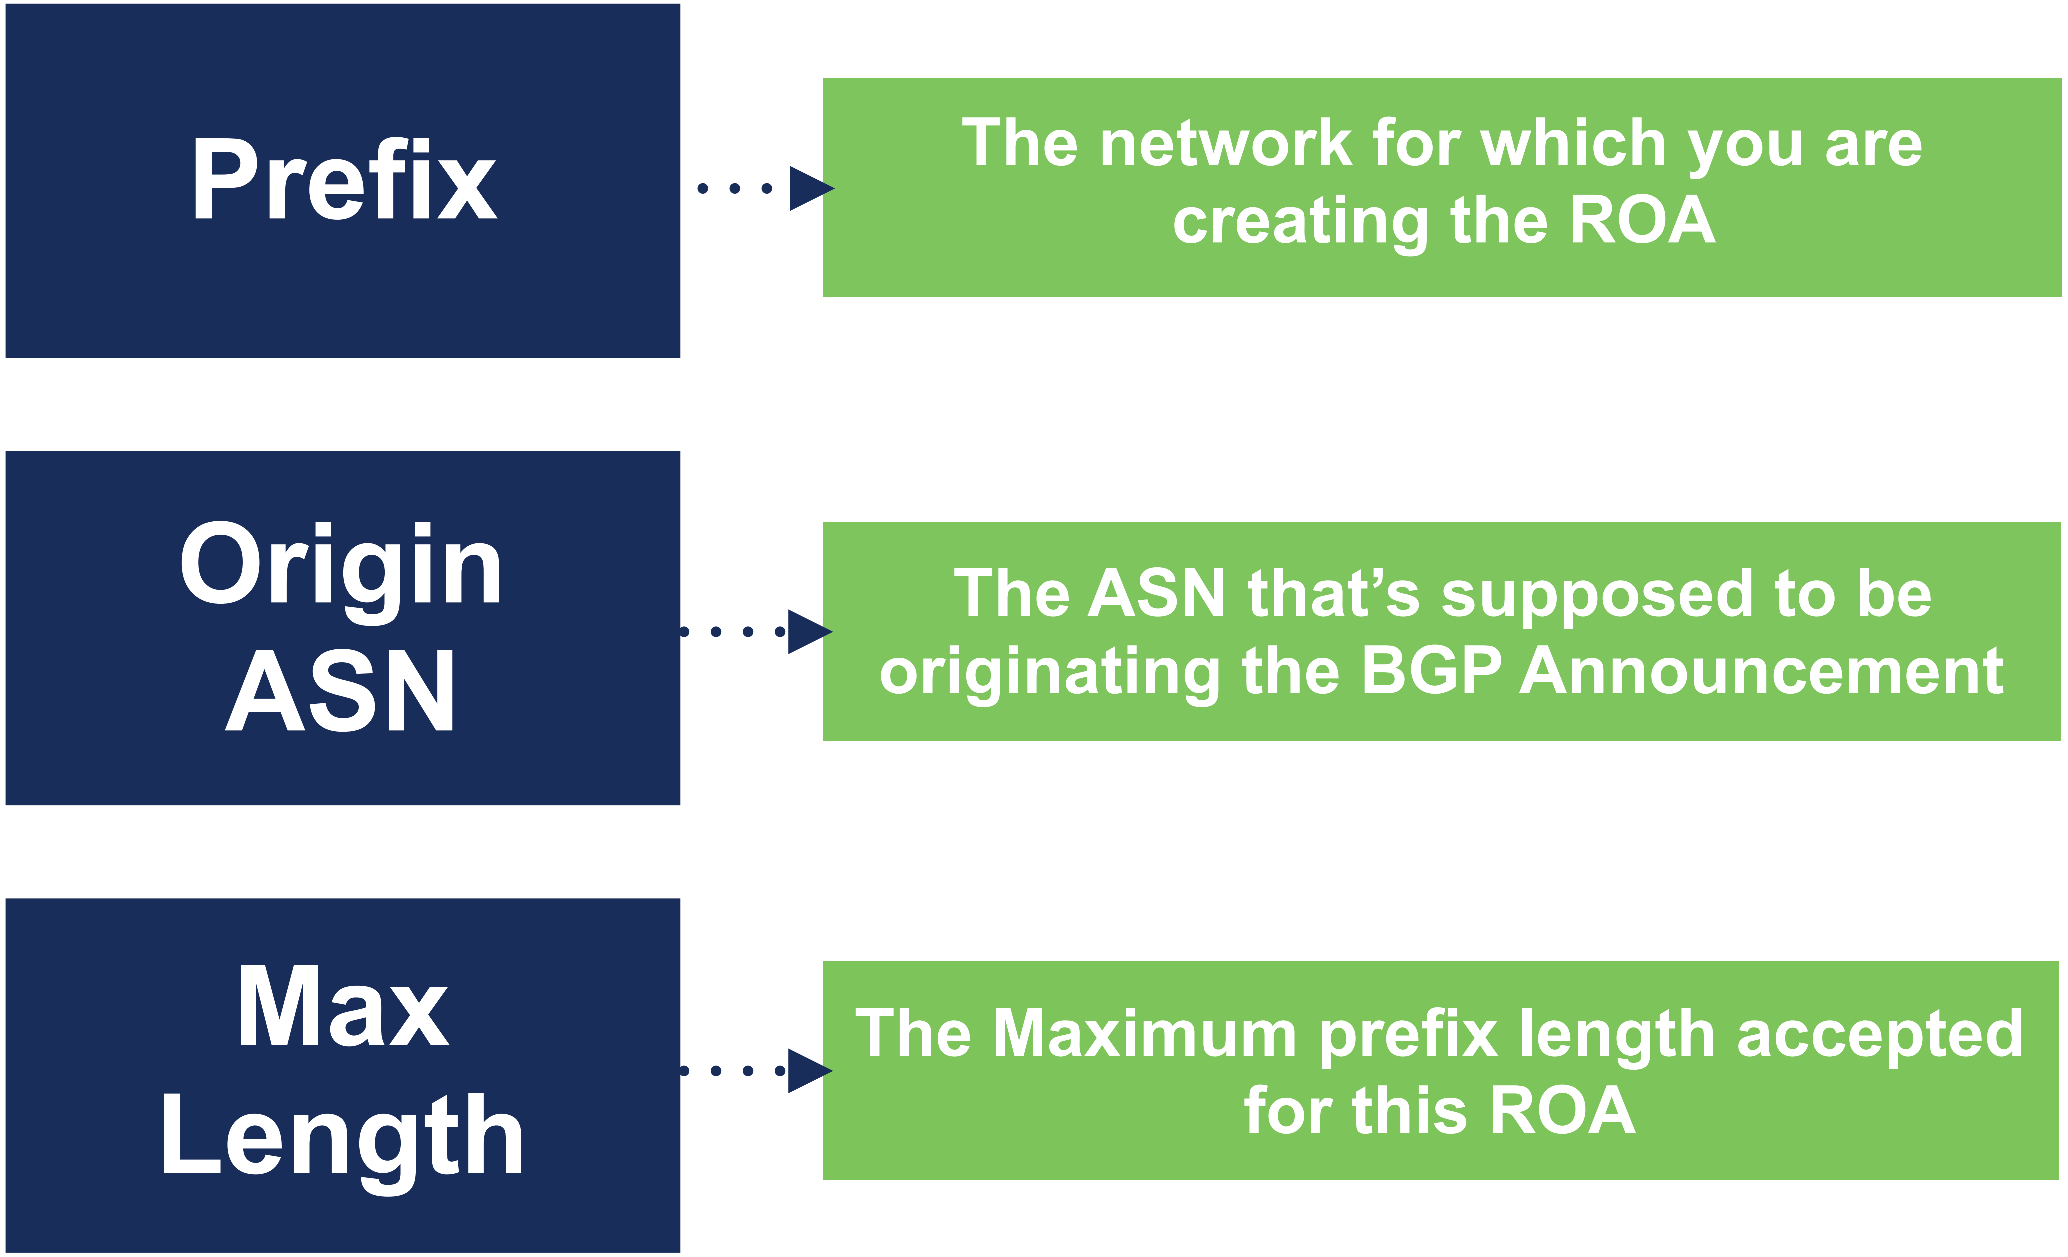

Prefix → The network for which you are creating the ROA

Origin ASN → The ASN that’s supposed to be originating the BGP Announcement

Max Length → The Maximum prefix length accepted for this ROA

RPKI Chain of Trust

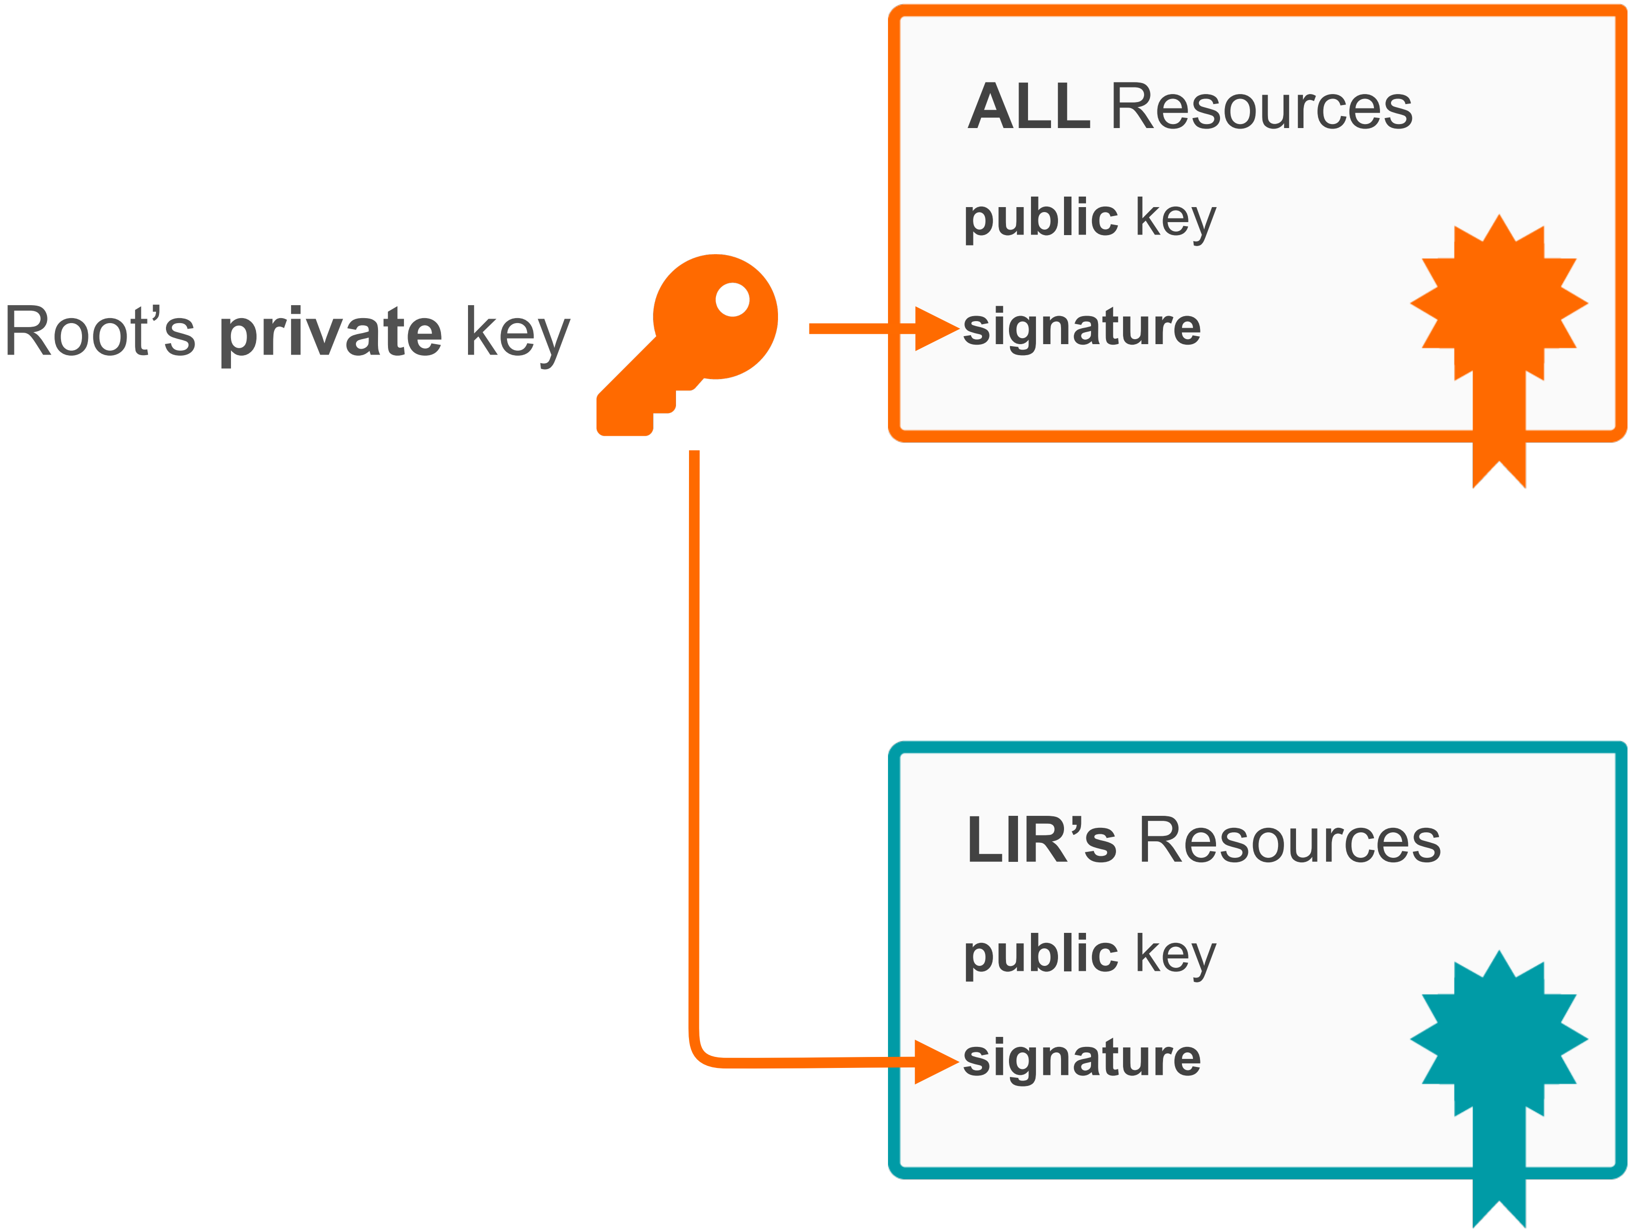

Root's **private** key

ALL Resources

public key

signature

LIR's Resources

public key

signature

Route Origin Authorisation

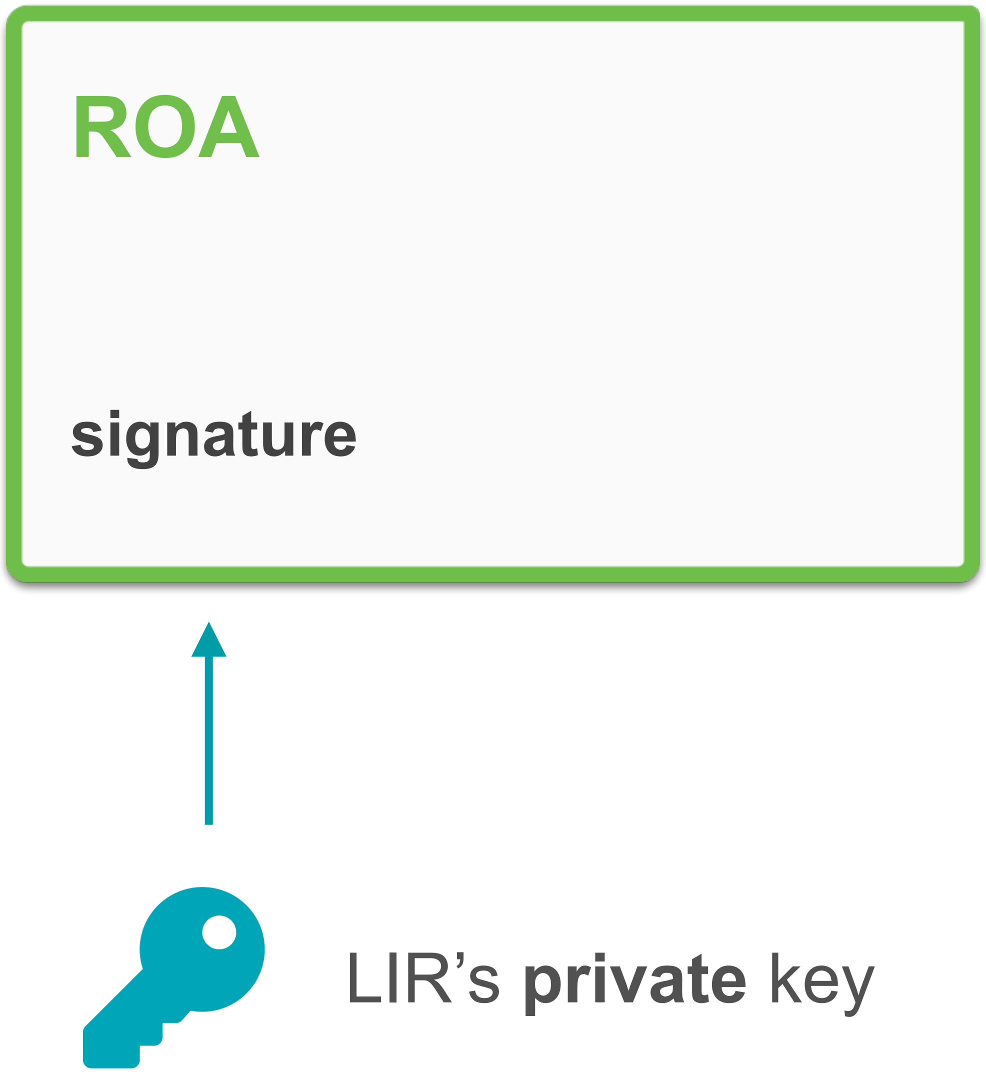

ROA

signature

LIR's **private** key

Prefix

is authorised to be announced by

AS Number

RPKI Chain of Trust

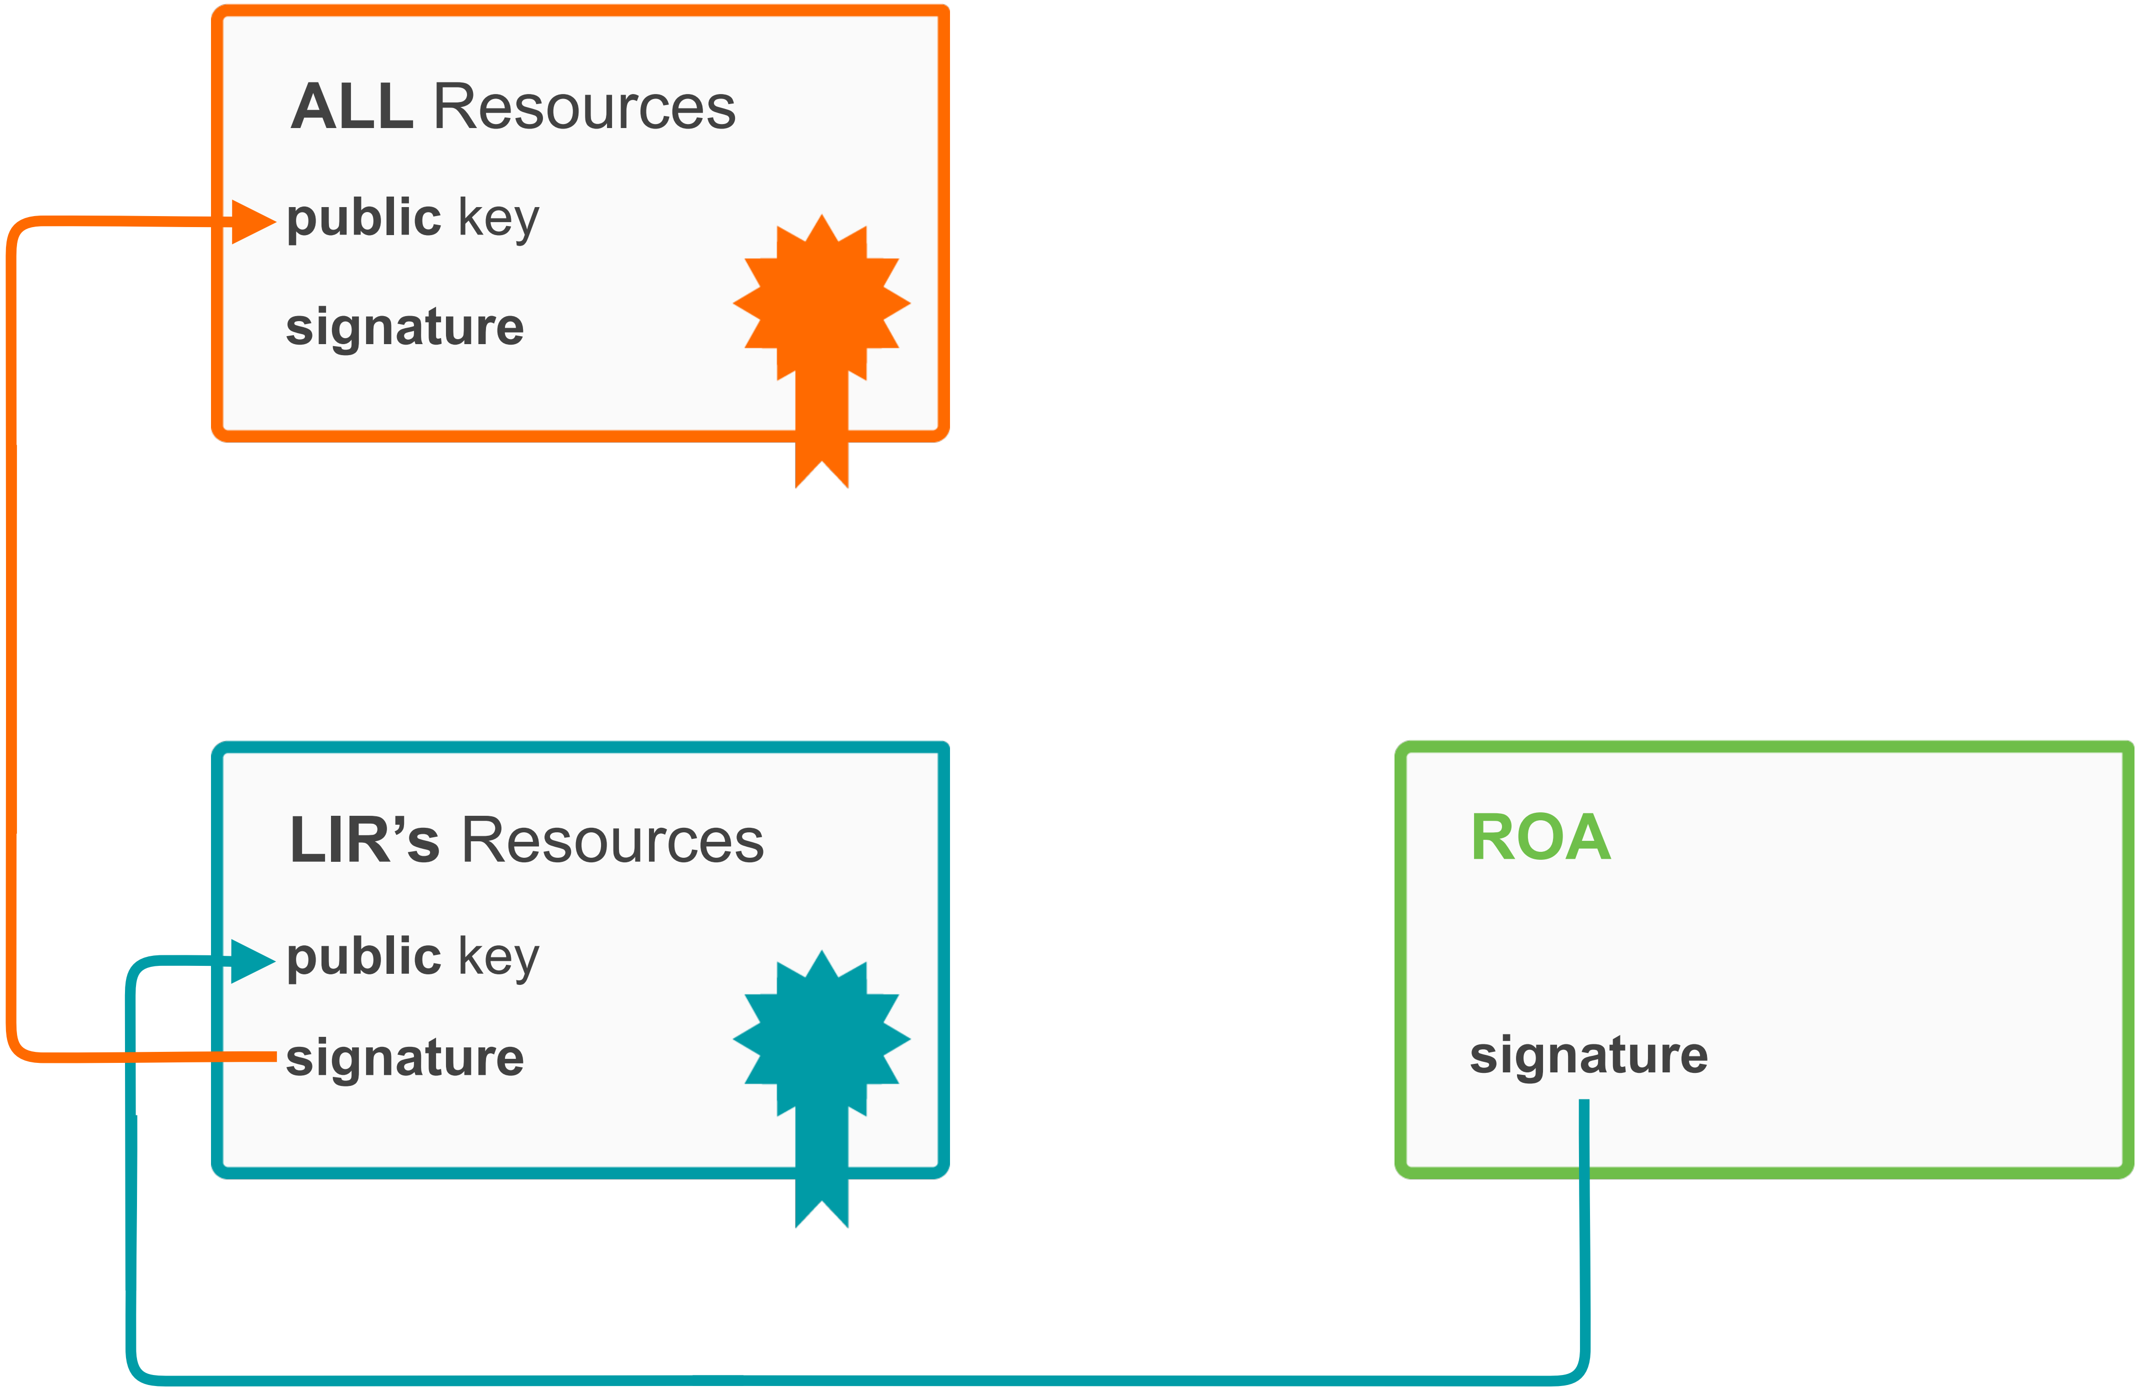

Hosted or Delegated RPKI

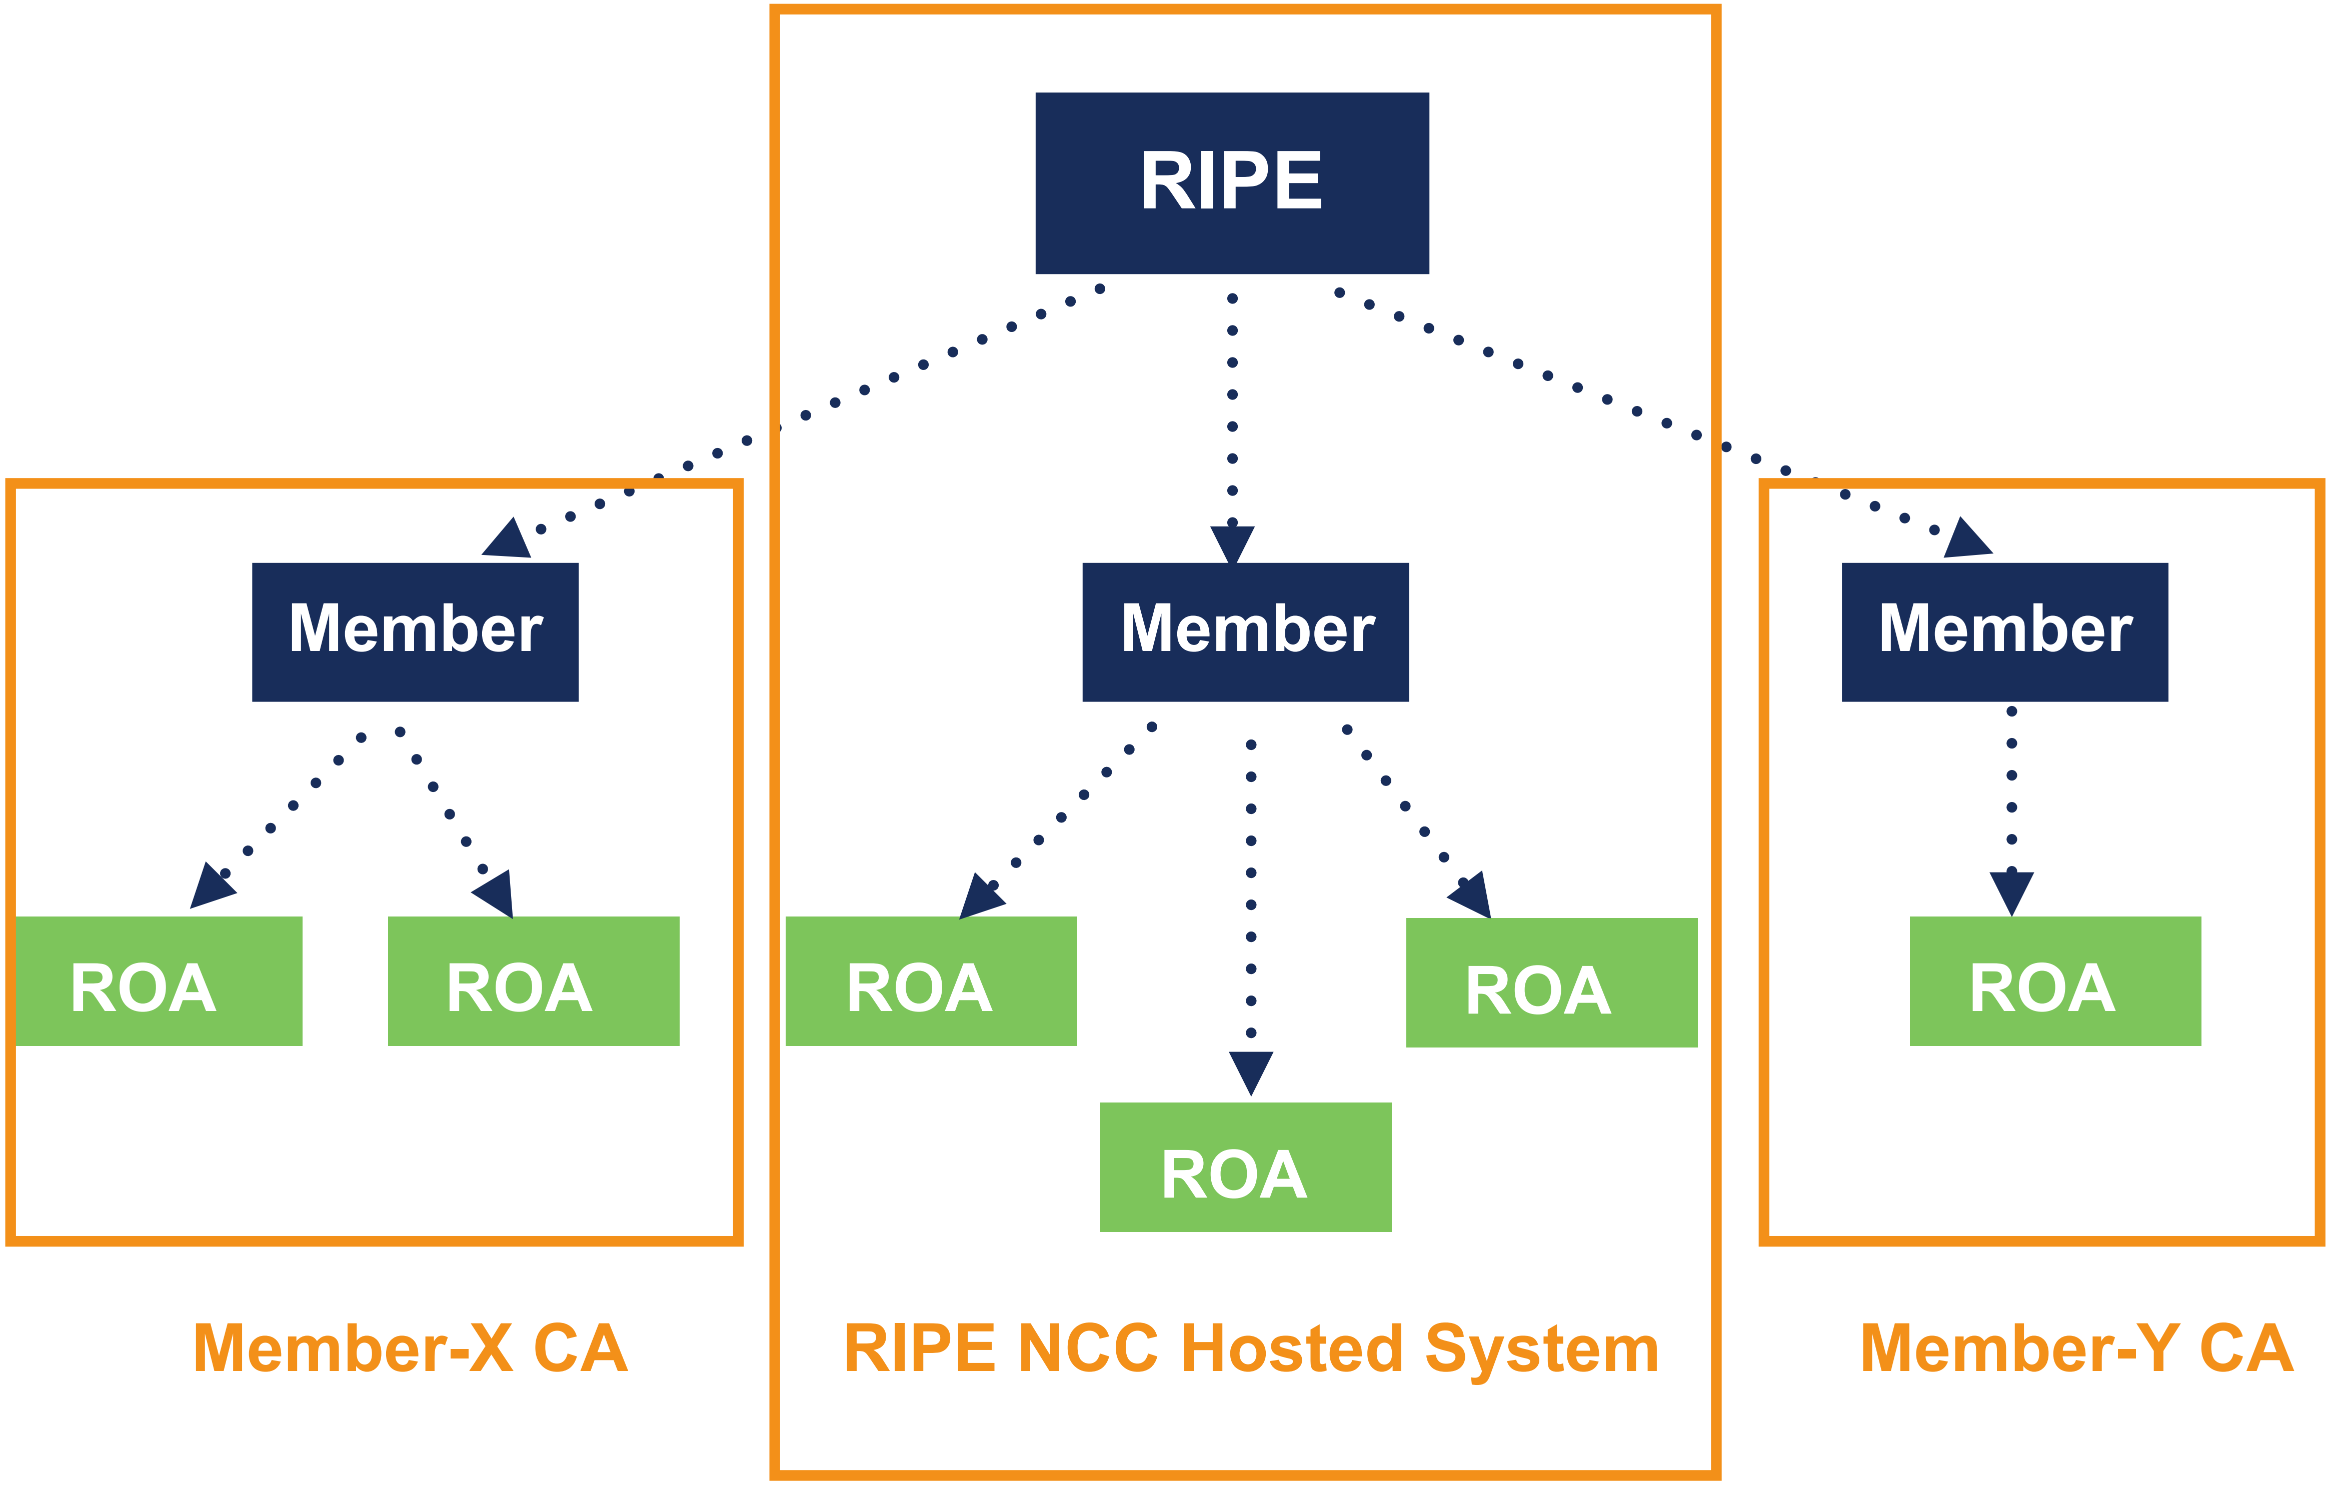

Hosted RPKI

- Automatic signing and key roll overs
  - One click setup of resource certificate
  - User has a valid and published certificate for as long as they are the holder of the resources
  - All the complexity is handled by the hosted system

- Lets you focus on creating and publishing ROAs
  - Match your intended BGP configuration

Delegated RPKI

- Run your own Certification Authority software
  - Dragon research Labs, RPKI toolkit
  - NLNetLabs, Krill
- Setup connection with RIPE NCC CA
- Generate a certificate and get it signed by the parent CA
- Run your own repository

First login to the dashboard

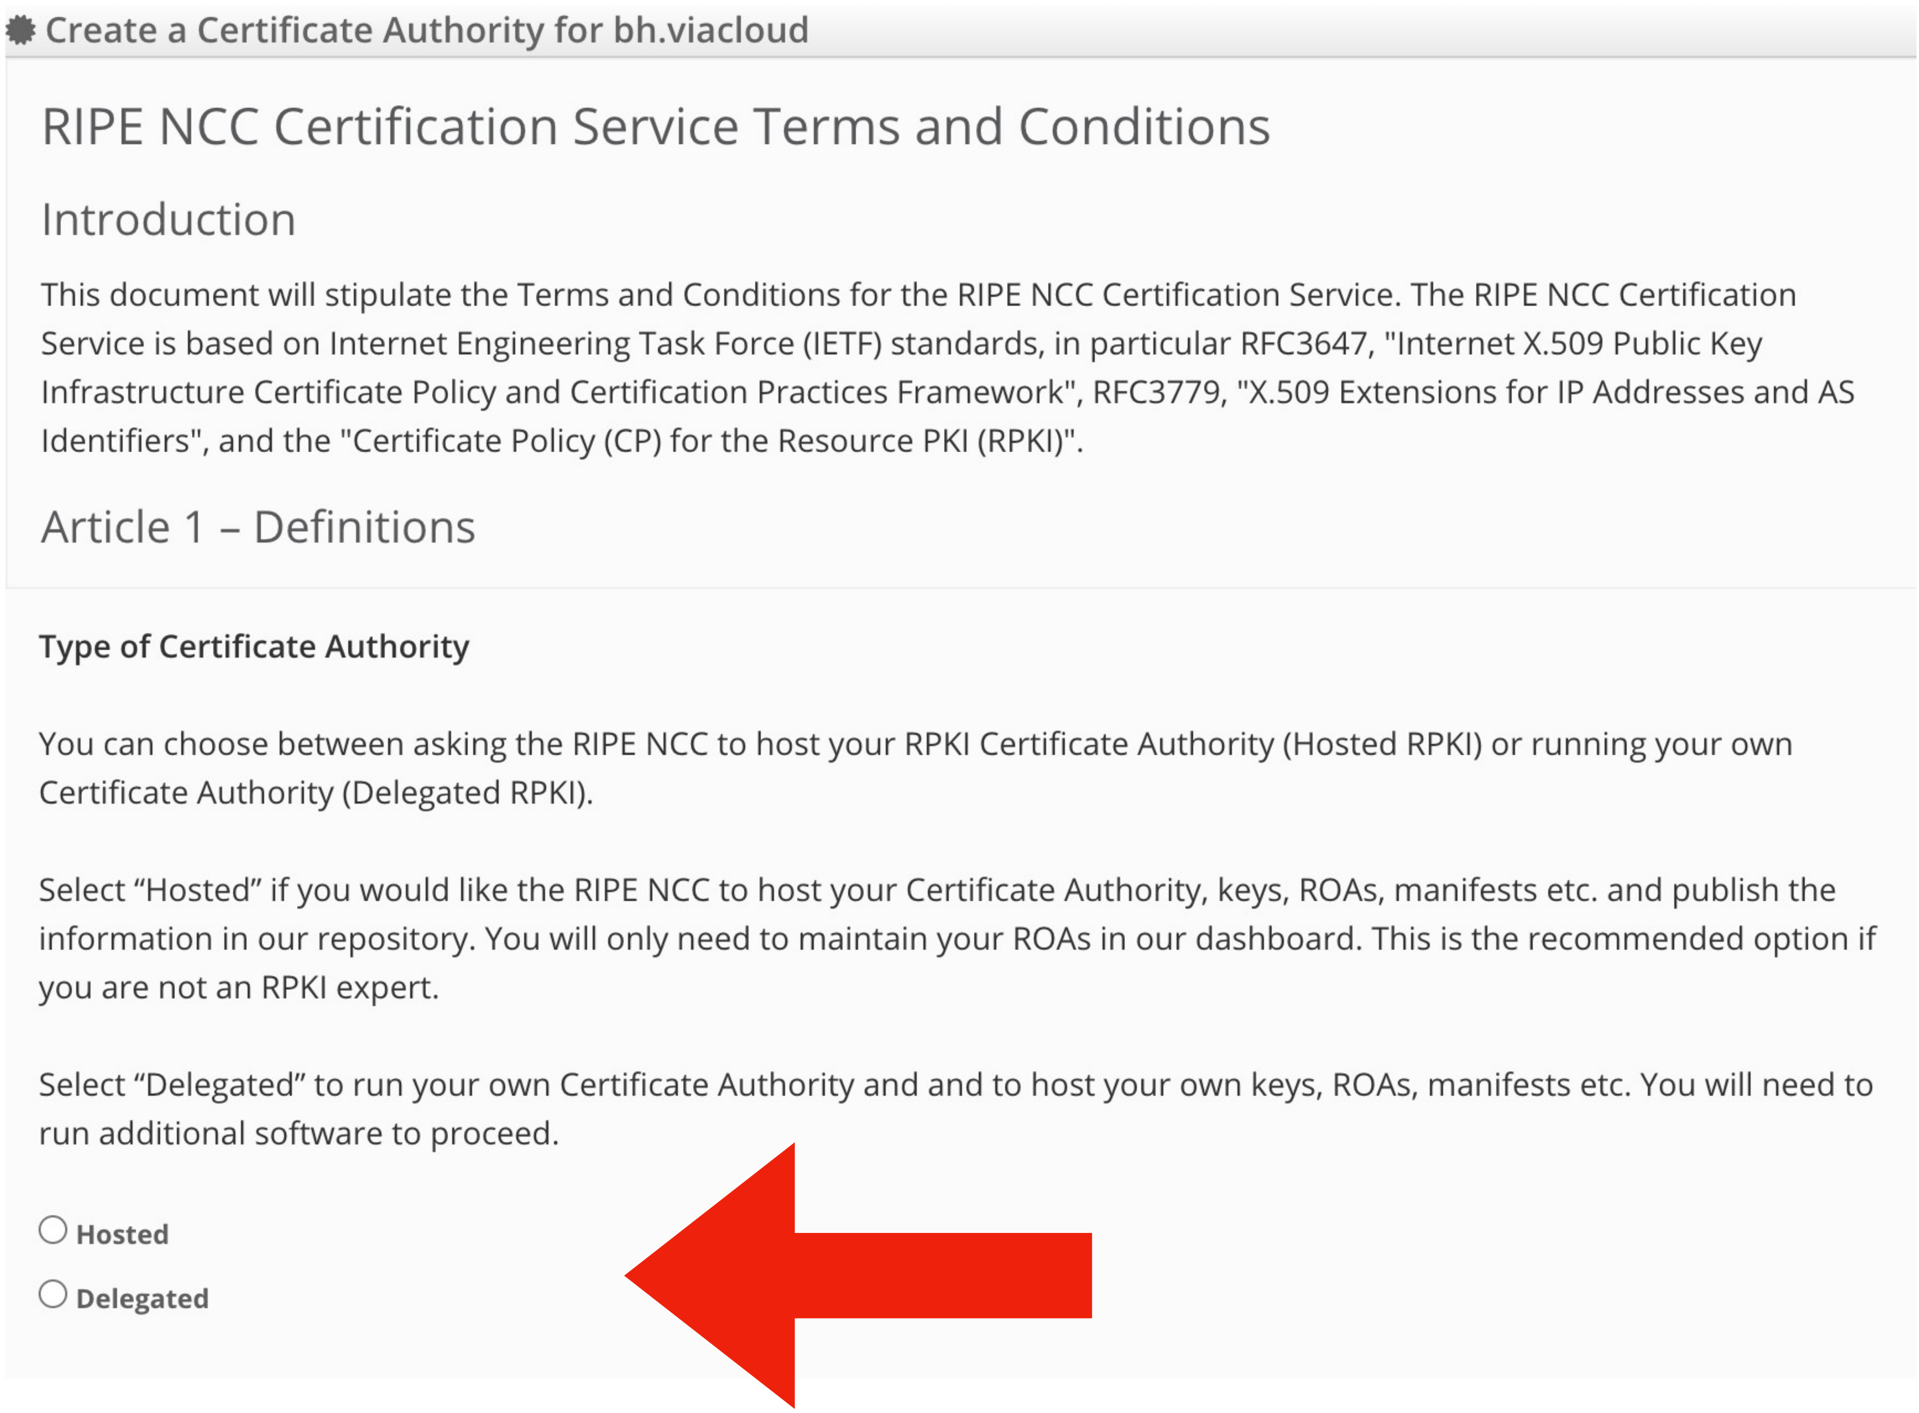

Create a Certificate Authority for bh.viacloud

RIPE NCC Certification Service Terms and Conditions

Introduction

This document will stipulate the Terms and Conditions for the RIPE NCC Certification Service. The RIPE NCC Certification Service is based on Internet Engineering Task Force (IETF) standards, in particular RFC3647, "Internet X.509 Public Key Infrastructure Certificate Policy and Certification Practices Framework", RFC3779, "X.509 Extensions for IP Addresses and AS Identifiers", and the "Certificate Policy (CP) for the Resource PKI (RPKI)".

Article 1 – Definitions

Type of Certificate Authority

You can choose between asking the RIPE NCC to host your RPKI Certificate Authority (Hosted RPKI) or running your own Certificate Authority (Delegated RPKI).

Select “Hosted” if you would like the RIPE NCC to host your Certificate Authority, keys, ROAs, manifests etc. and publish the information in our repository. You will only need to maintain your ROAs in our dashboard. This is the recommended option if you are not an RPKI expert.

Select “Delegated” to run your own Certificate Authority and and to host your own keys, ROAs, manifests etc. You will need to run additional software to proceed.

○ **Hosted**

○ **Delegated**

Creating ROAs

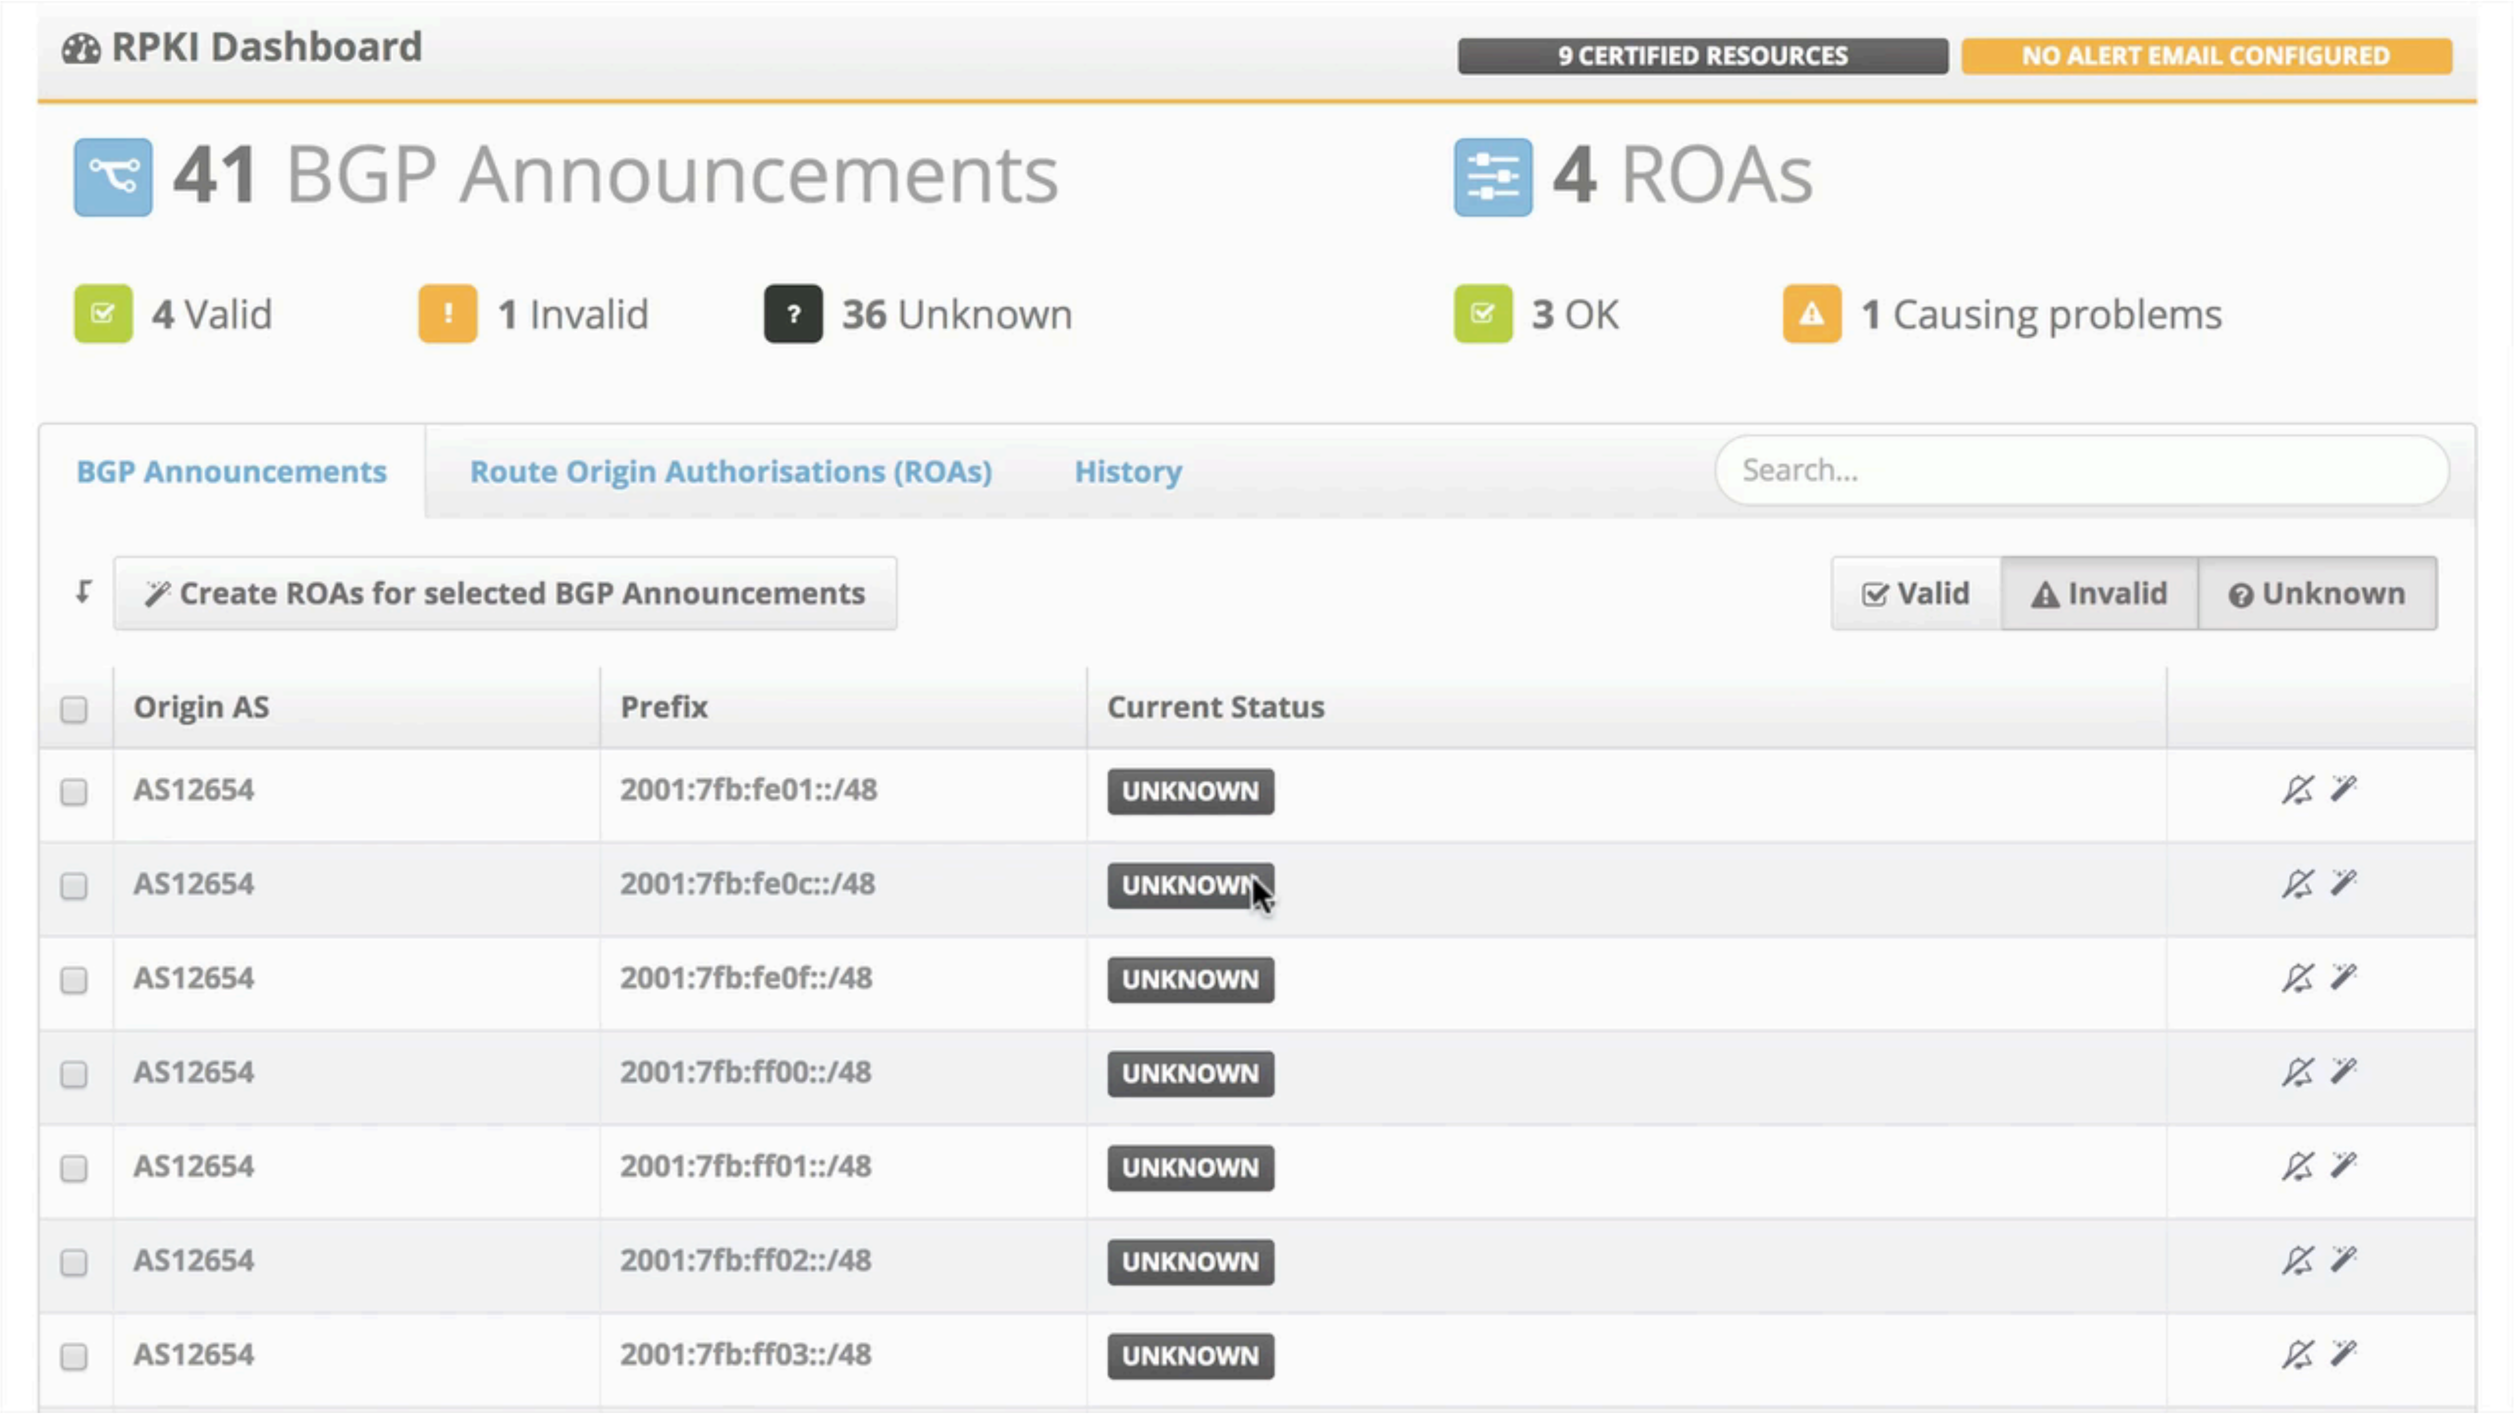
RPKI Dashboard
9 CERTIFIED RESOURCES
NO ALERT EMAIL CONFIGURED
BGP Announcements
Route Origin Authorisations (ROAs)
History
Search...
Create ROAs for selected BGP Announcements
Valid
Invalid
Unknown

41 BGP Announcements

4 Valid 1 Invalid 36 Unknown

4 ROAs

3 OK 1 Causing problems

□	Origin AS	Prefix	Current Status
□	AS12654	2001:7fb:fe01::/48	UNKNOWN
□	AS12654	2001:7fb:fe0c::/48	UNKNOWN
□	AS12654	2001:7fb:fe0f::/48	UNKNOWN
□	AS12654	2001:7fb:ff00::/48	UNKNOWN
□	AS12654	2001:7fb:ff01::/48	UNKNOWN
□	AS12654	2001:7fb:ff02::/48	UNKNOWN
□	AS12654	2001:7fb:ff03::/48	UNKNOWN

Reviewing changes

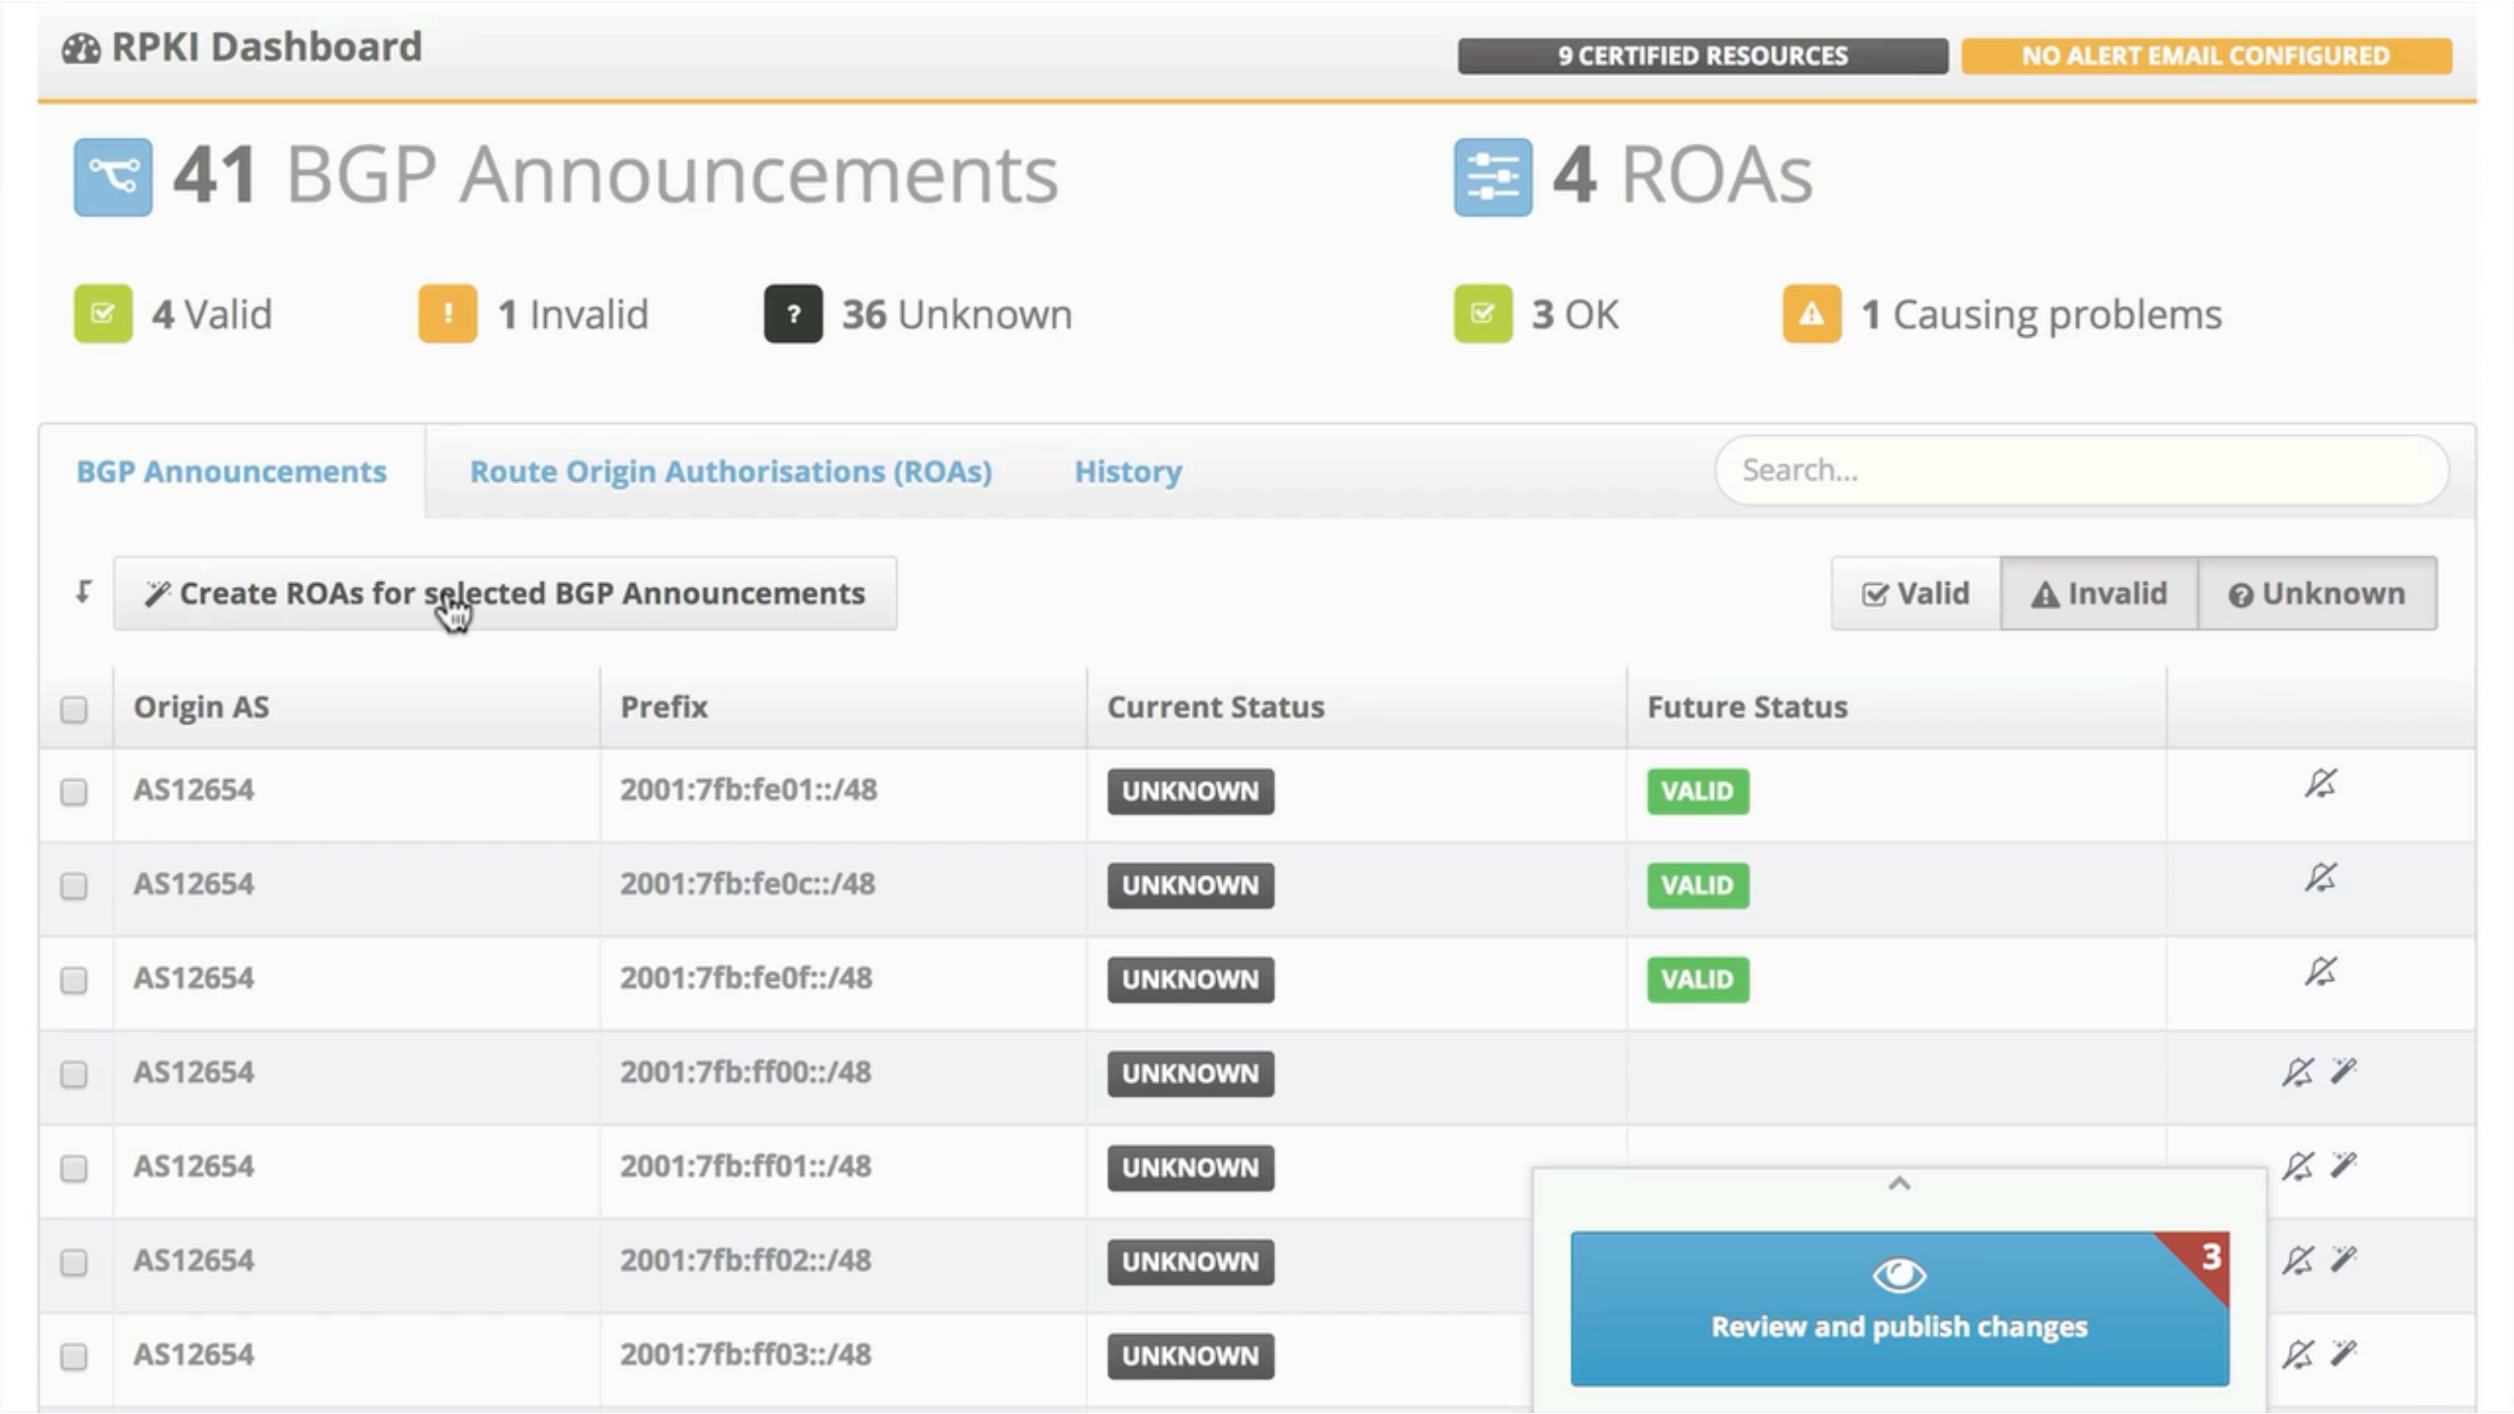

RPKI Dashboard

9 CERTIFIED RESOURCES

NO ALERT EMAIL CONFIGURED

41 BGP Announcements

4 Valid **1** Invalid **36** Unknown

4 ROAs

3 OK **1** Causing problems

BGP Announcements
Route Origin Authorisations (ROAs)
History
Search...
Create ROAs for selected BGP Announcements
Valid
Invalid
Unknown

□	Origin AS	Prefix	Current Status	Future Status
□	AS12654	2001:7fb:fe01::/48	UNKNOWN	VALID
□	AS12654	2001:7fb:fe0c::/48	UNKNOWN	VALID
□	AS12654	2001:7fb:fe0f::/48	UNKNOWN	VALID
□	AS12654	2001:7fb:ff00::/48	UNKNOWN	
□	AS12654	2001:7fb:ff01::/48	UNKNOWN	
□	AS12654	2001:7fb:ff02::/48	UNKNOWN	
□	AS12654	2001:7fb:ff03::/48	UNKNOWN	

Checking the effects

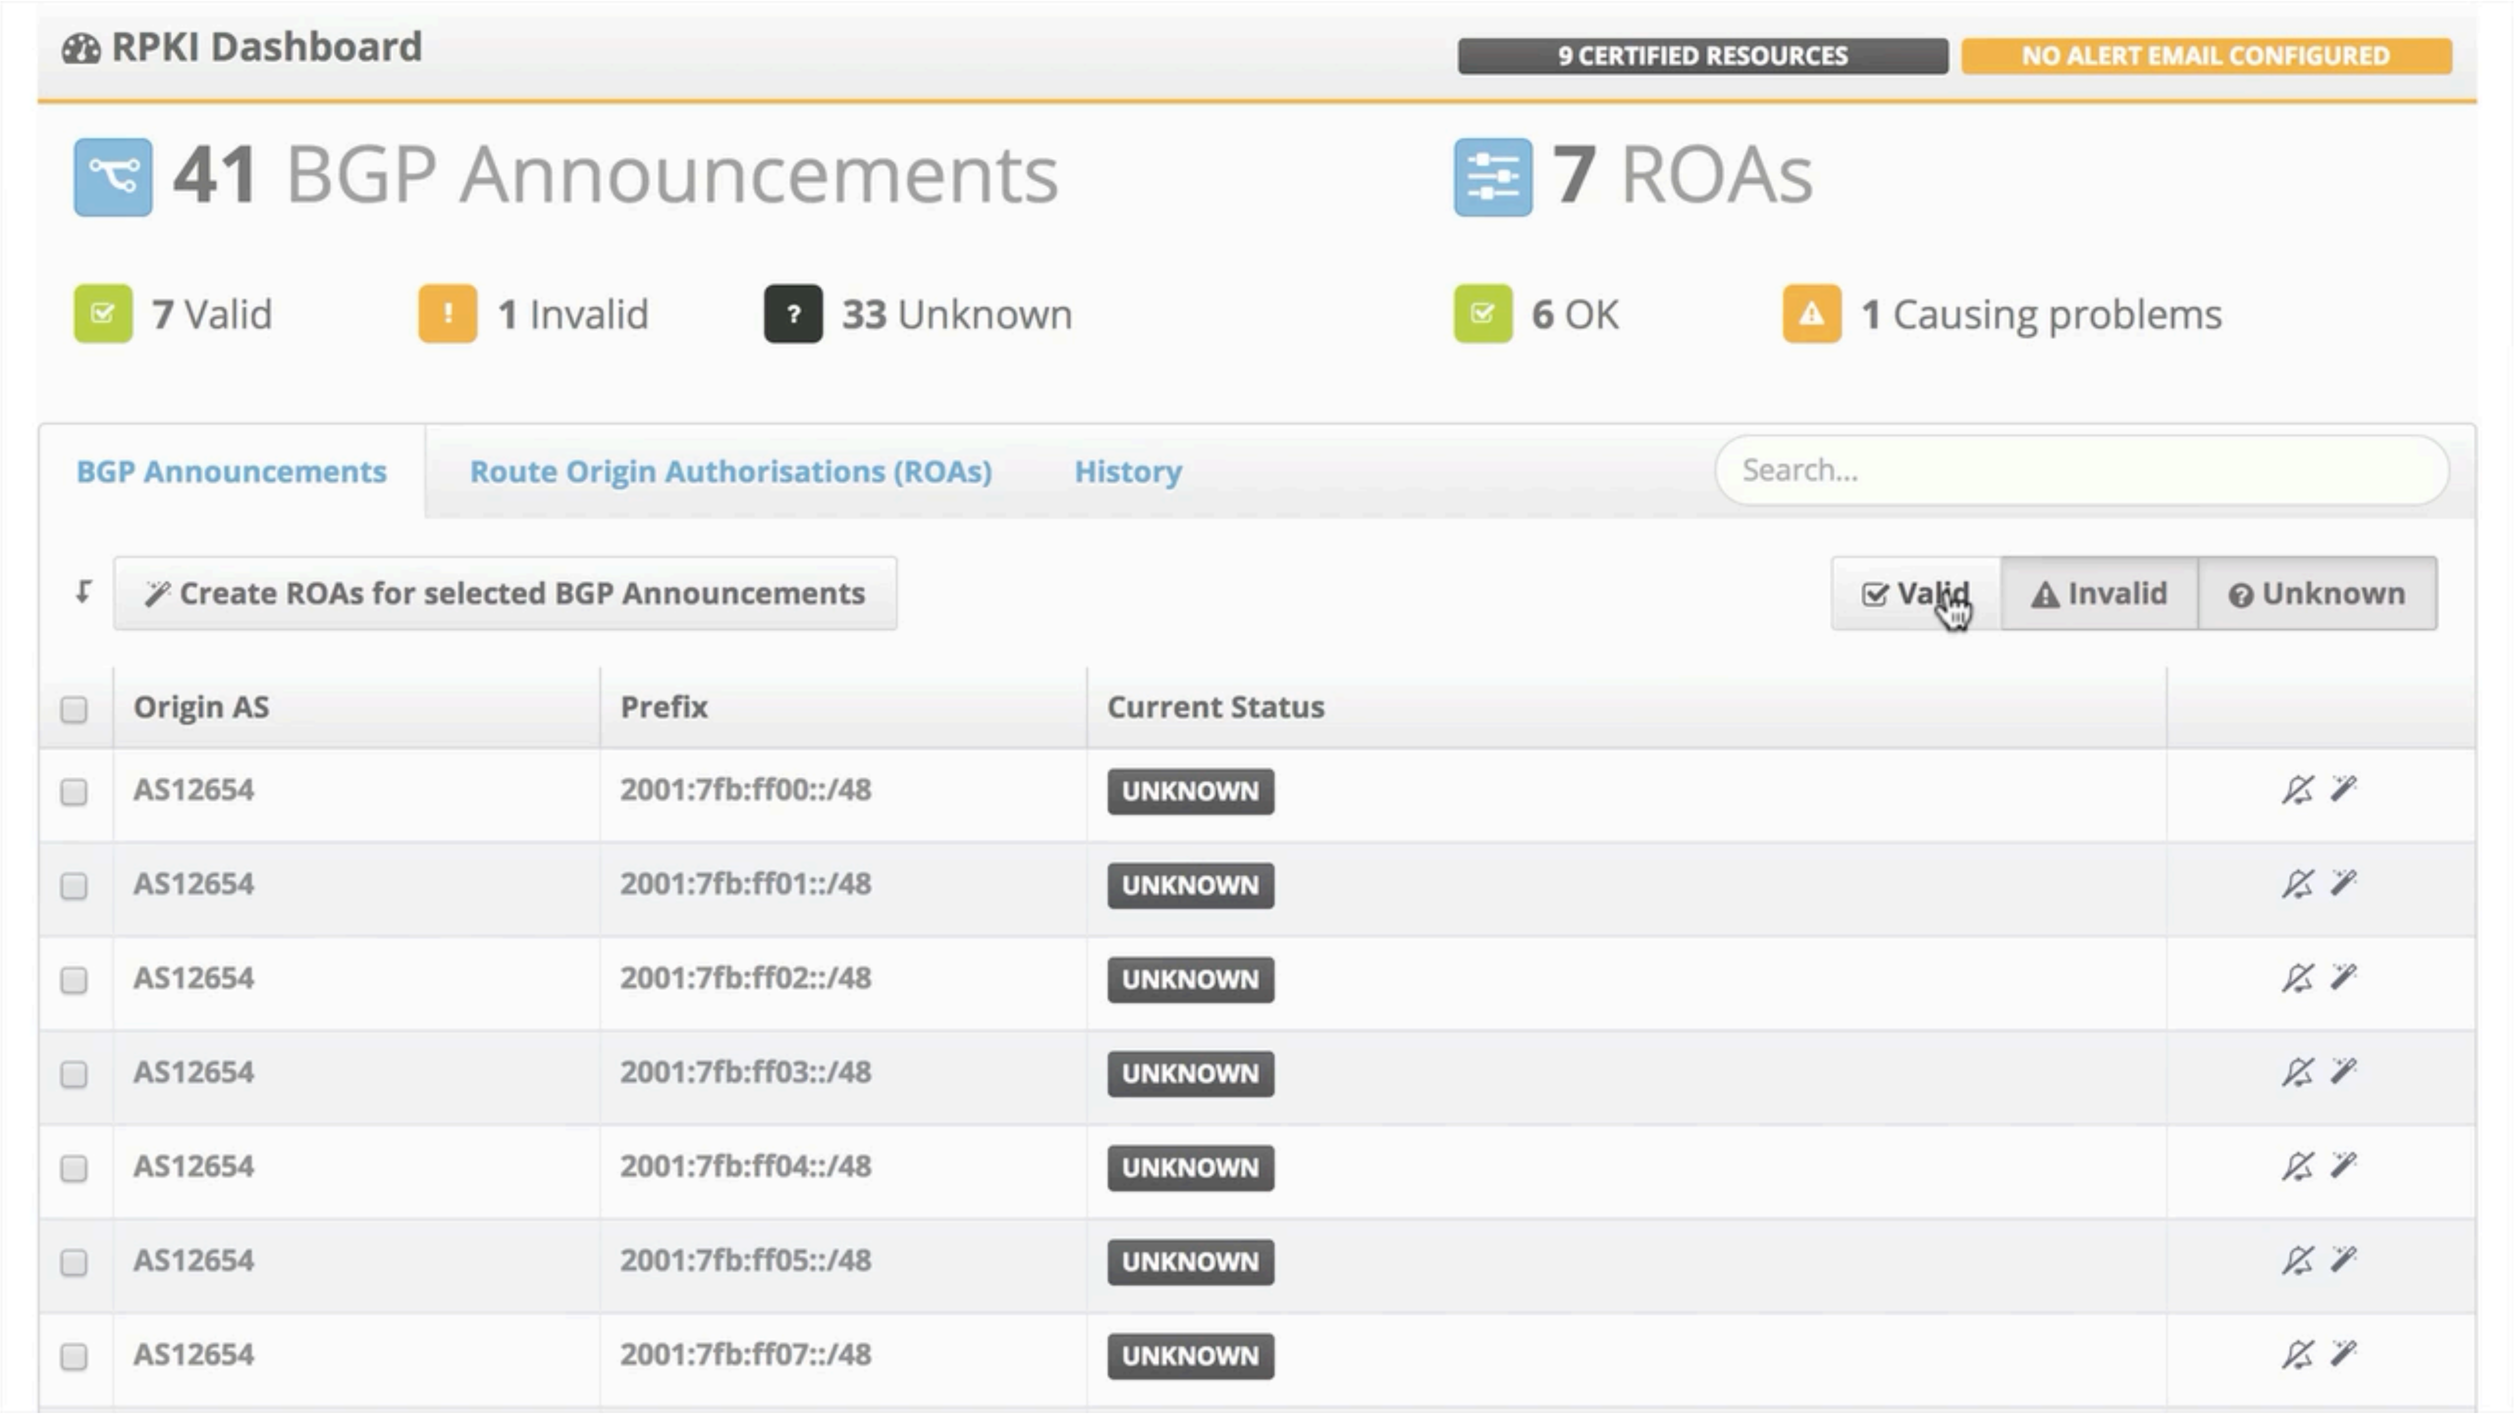

RPKI Dashboard

9 CERTIFIED RESOURCES

NO ALERT EMAIL CONFIGURED

41 BGP Announcements

7 Valid

1 Invalid

33 Unknown

7 ROAs

6 OK

1 Causing problems

BGP Announcements | Route Origin Authorisations (ROAs) | History

Search...

Create ROAs for selected BGP Announcements

☑ Valid | Invalid | Unknown

□	Origin AS	Prefix	Current Status
□	AS12654	2001:7fb:ff00::/48	UNKNOWN
□	AS12654	2001:7fb:ff01::/48	UNKNOWN
□	AS12654	2001:7fb:ff02::/48	UNKNOWN
□	AS12654	2001:7fb:ff03::/48	UNKNOWN
□	AS12654	2001:7fb:ff04::/48	UNKNOWN
□	AS12654	2001:7fb:ff05::/48	UNKNOWN
□	AS12654	2001:7fb:ff07::/48	UNKNOWN

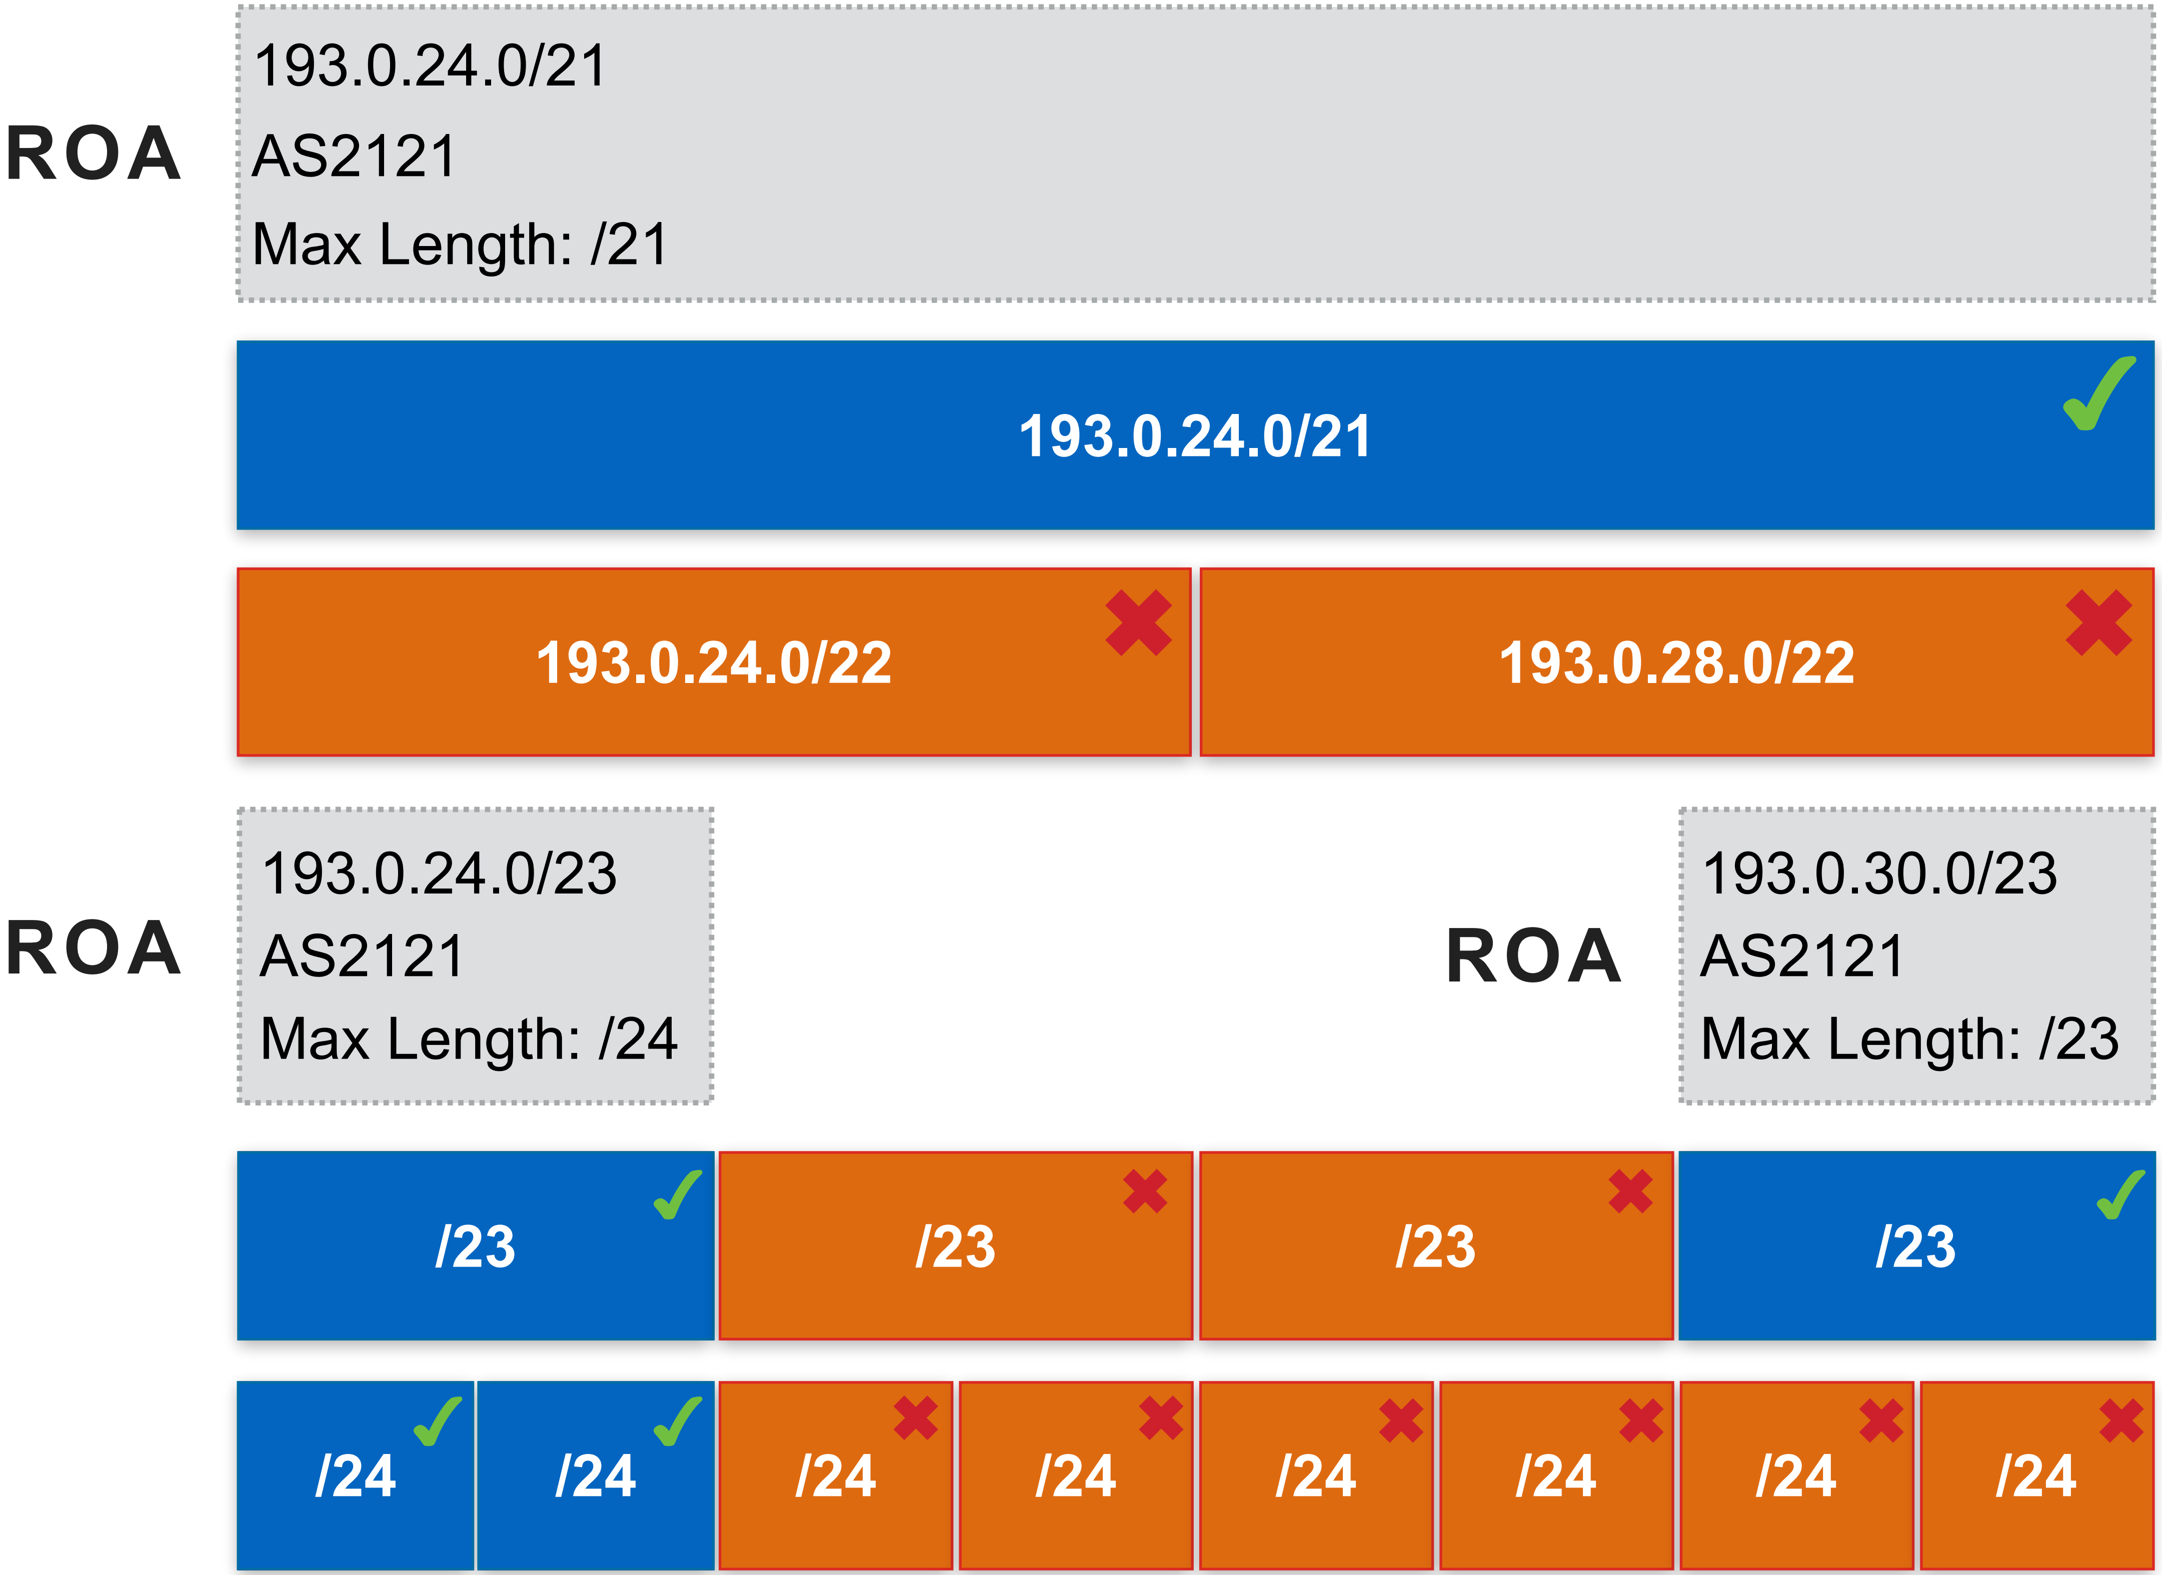

ROA

193.0.24.0/21
AS2121
Max Length: /21

193.0.24.0/21 ✓

193.0.24.0/22 ✗

193.0.28.0/22 ✗

ROA

193.0.24.0/23
AS2121
Max Length: /24

ROA

193.0.30.0/23
AS2121
Max Length: /23

/23 ✓ | /23 ✗ | /23 ✗ | /23 ✓

/24 ✓ | /24 ✓ | /24 ✗ | /24 ✗ | /24 ✗ | /24 ✗ | /24 ✗ | /24 ✗

RPKI Adoption

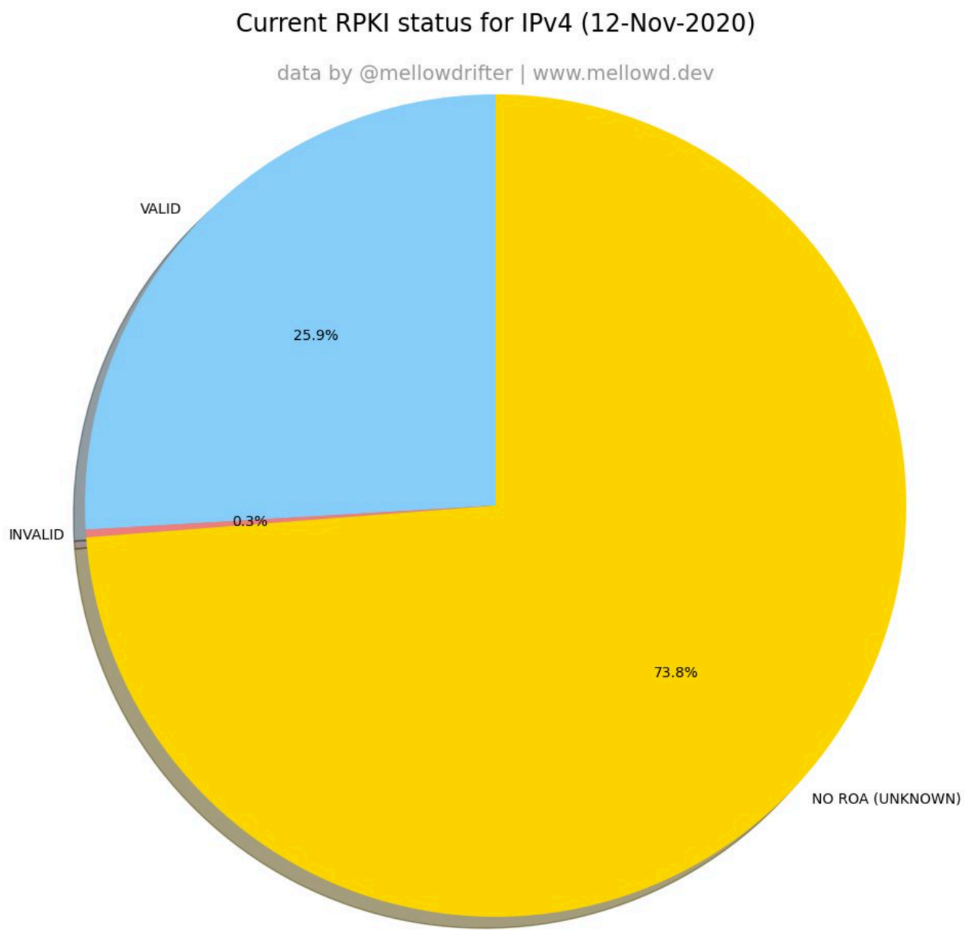

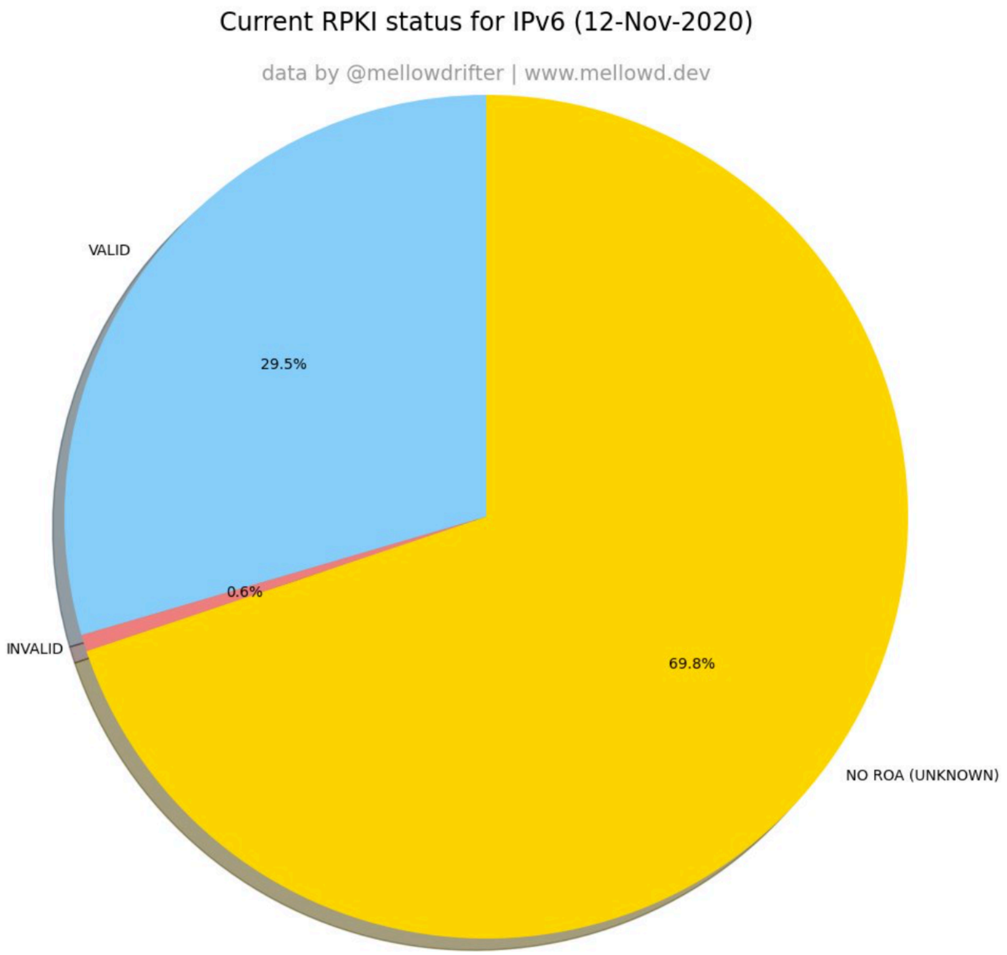

ROA Adoption

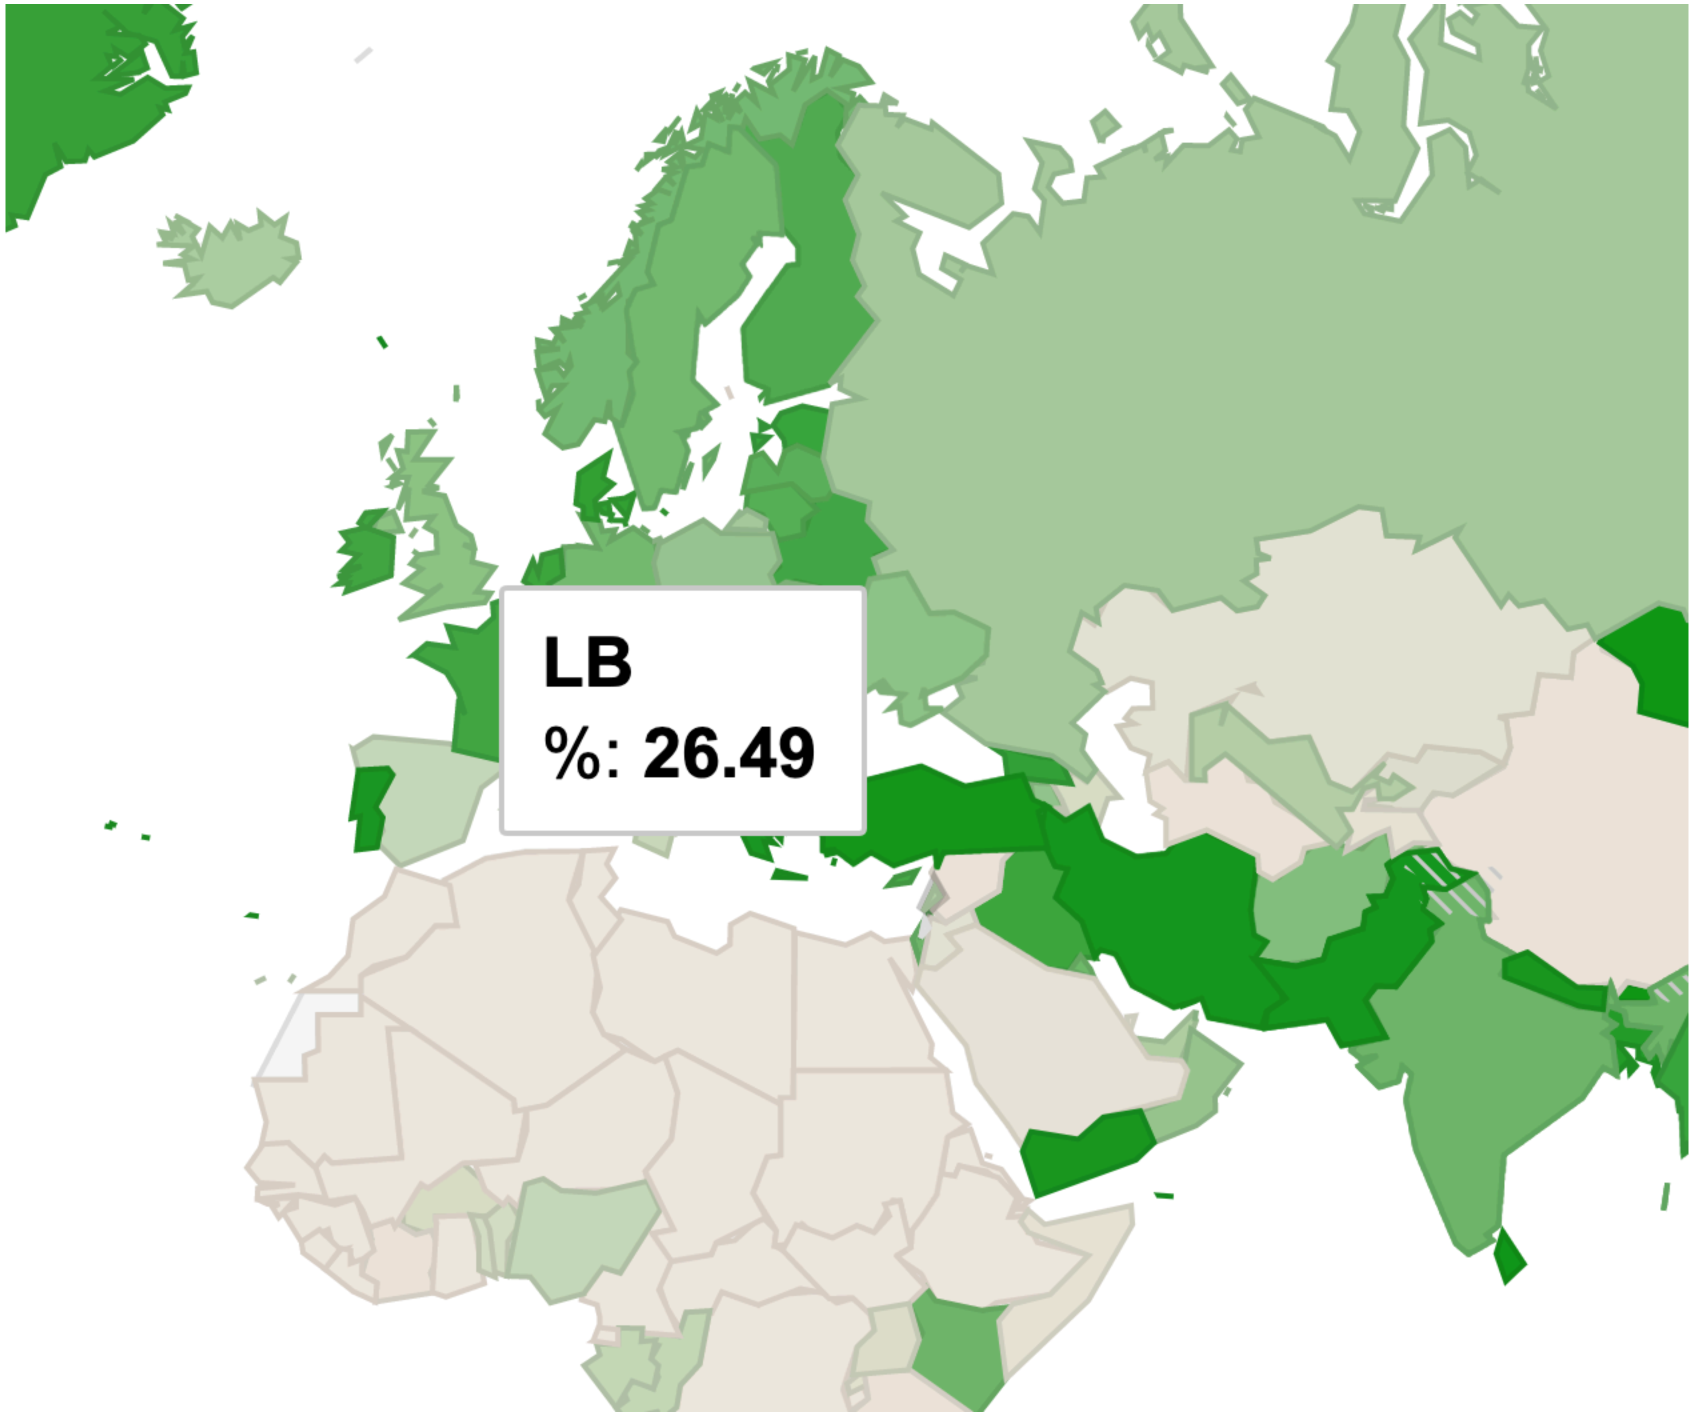

ROA Accuracy

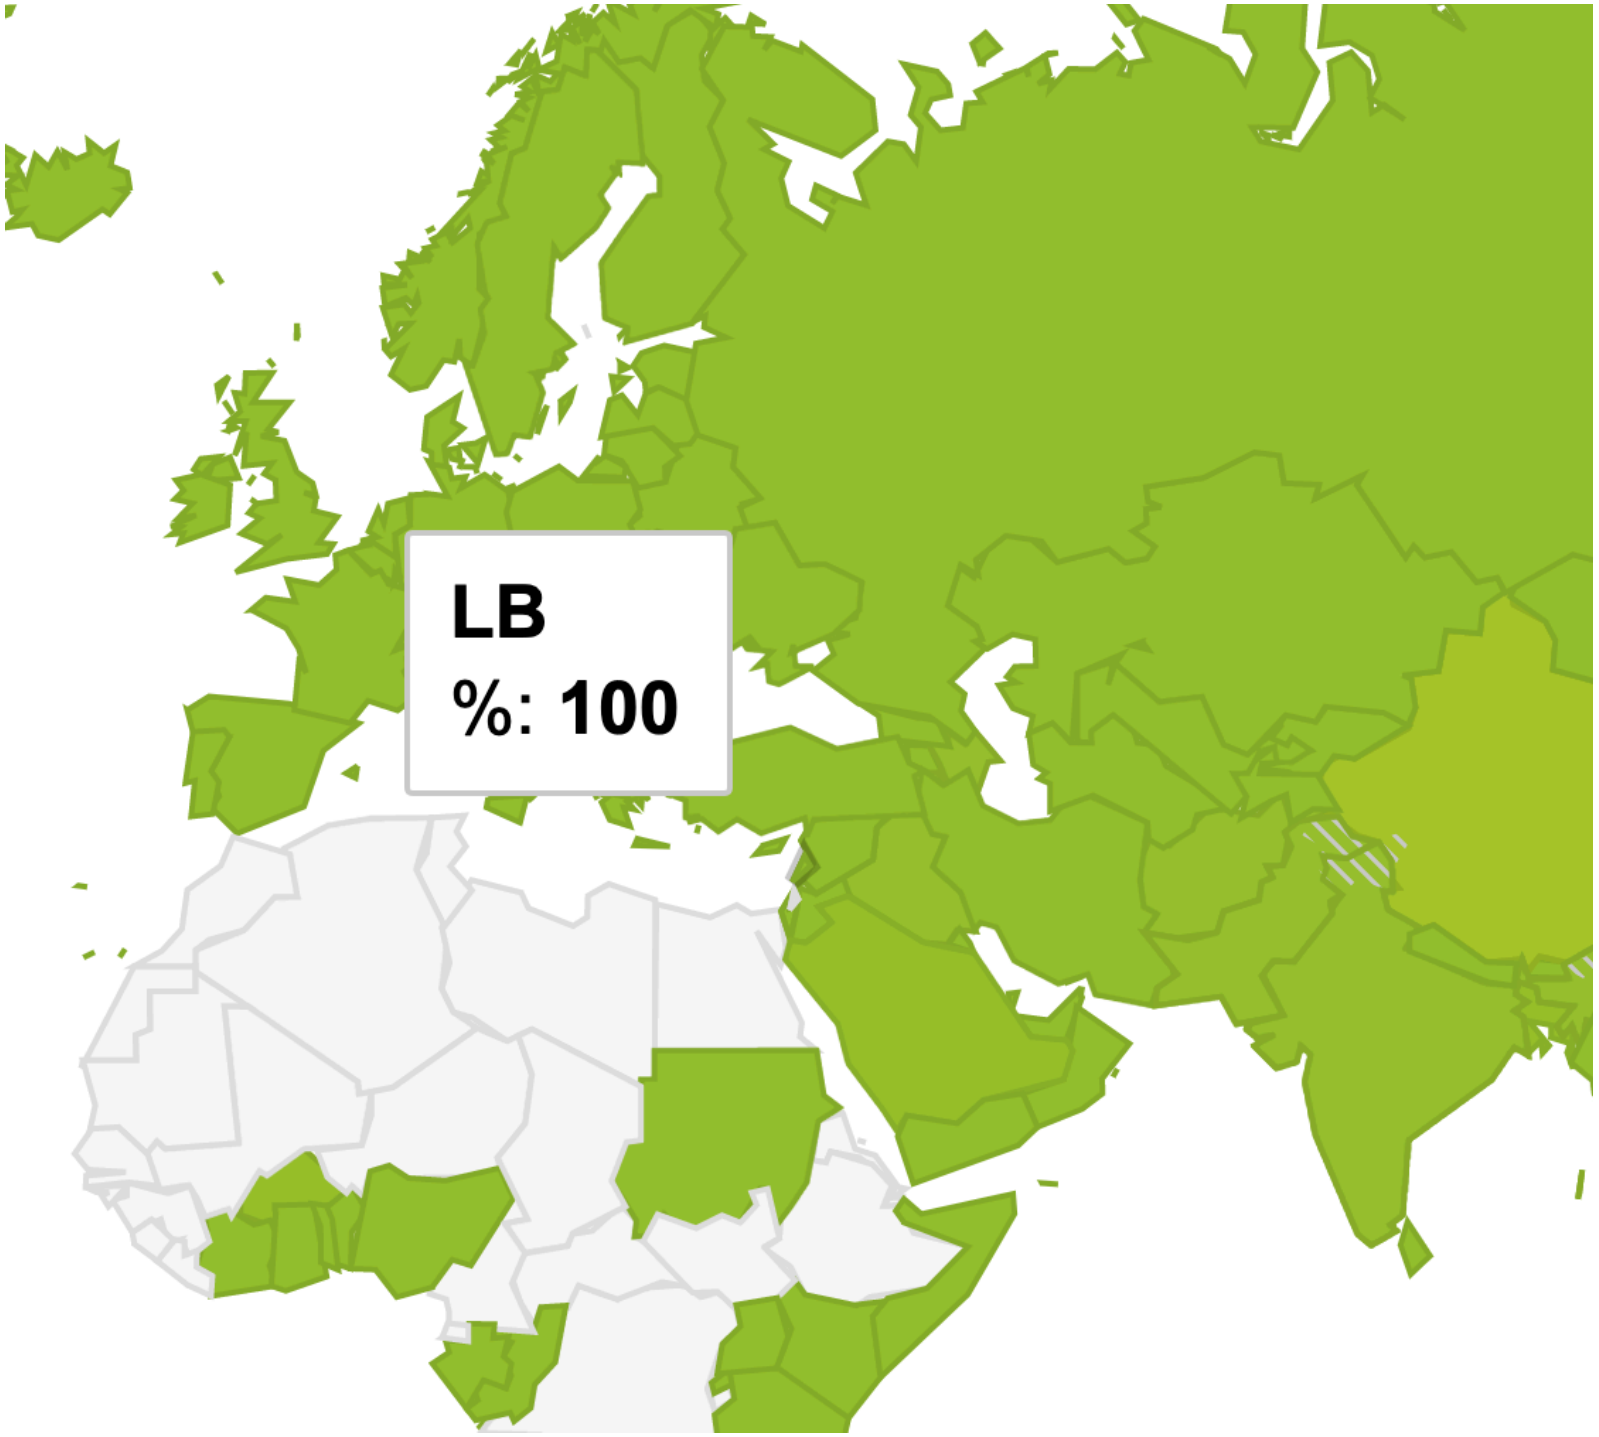

Validation Tools

Two elements of RPKI

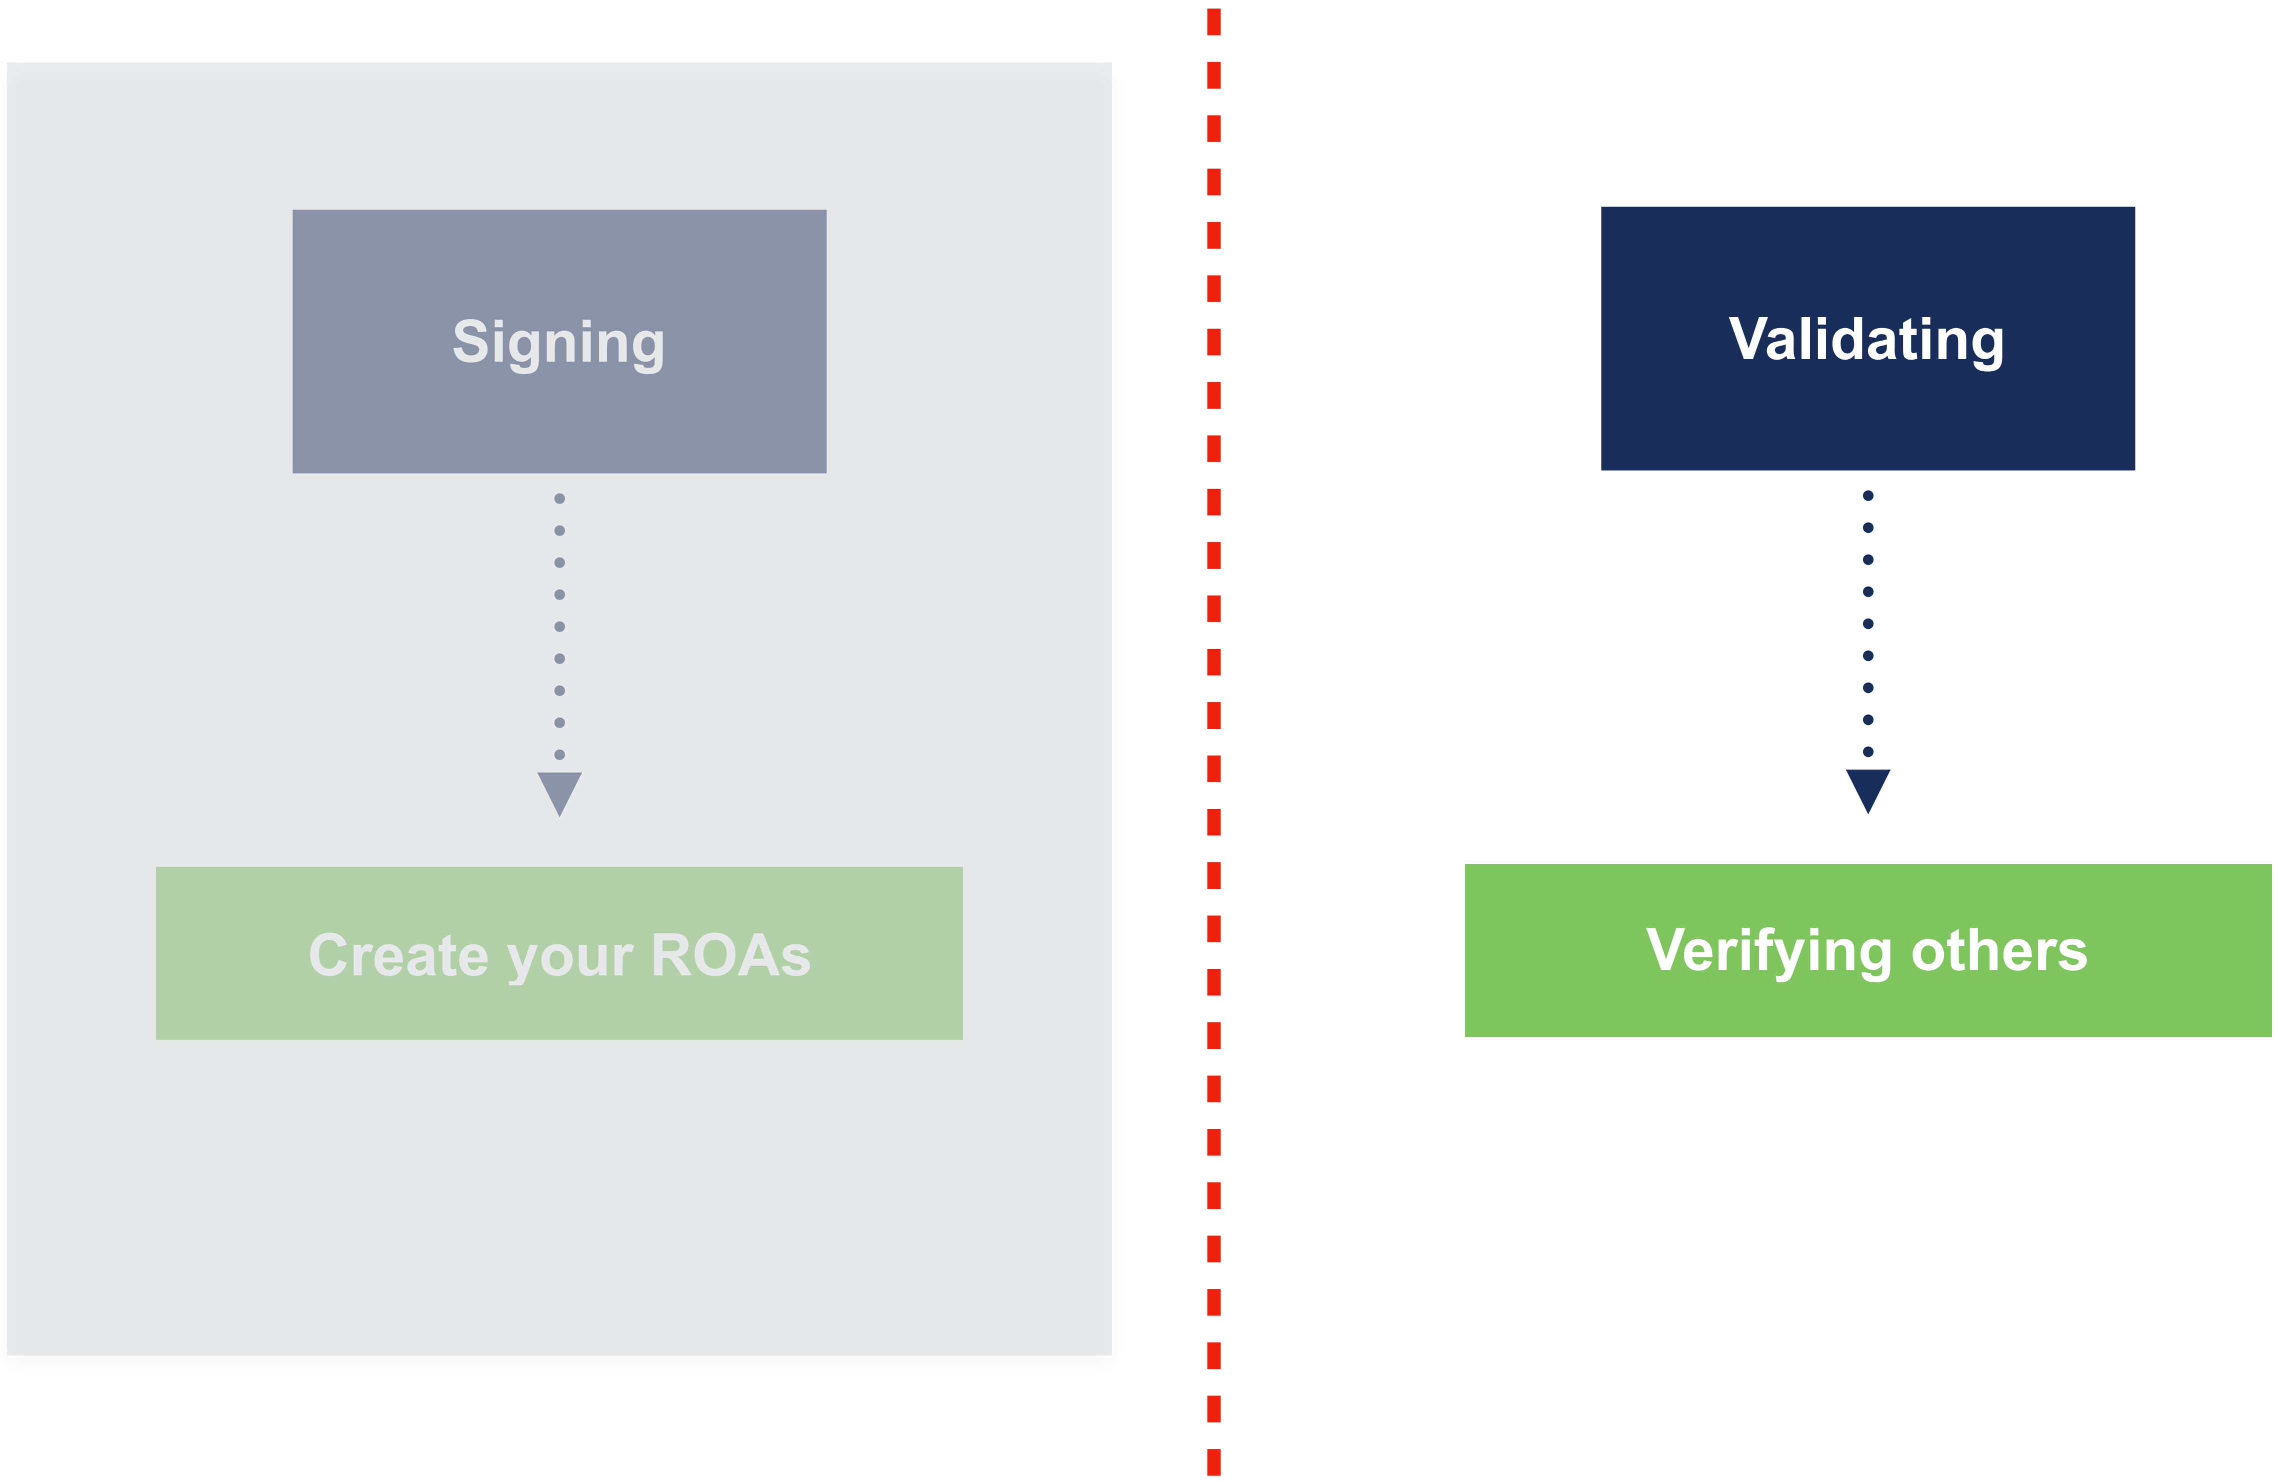

Routing on the Internet

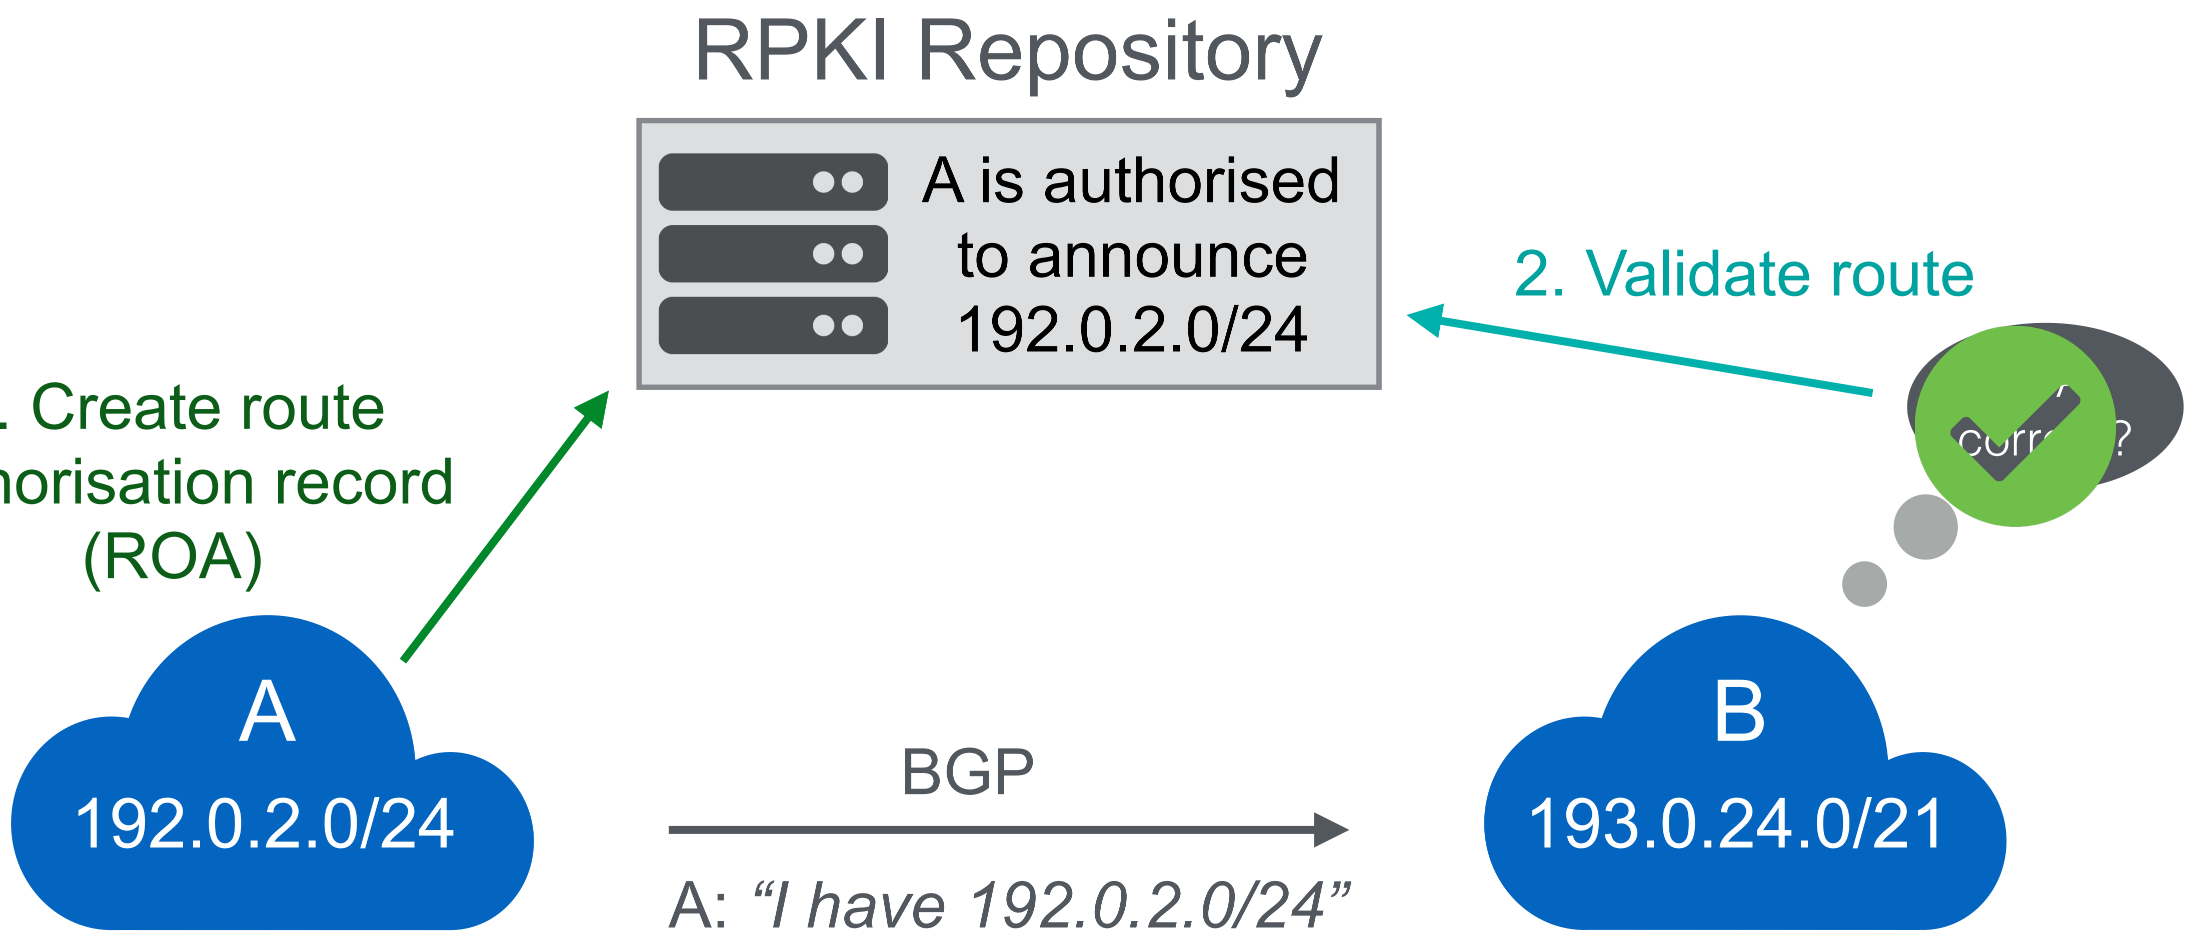

RPKI Repository

A is authorised to announce 192.0.2.0/24

1. Create route authorisation record (ROA)

2. Validate route

A

192.0.2.0/24

BGP

A: *“I have 192.0.2.0/24”*

B

193.0.24.0/21

Trust Anchor Locator (TAL)

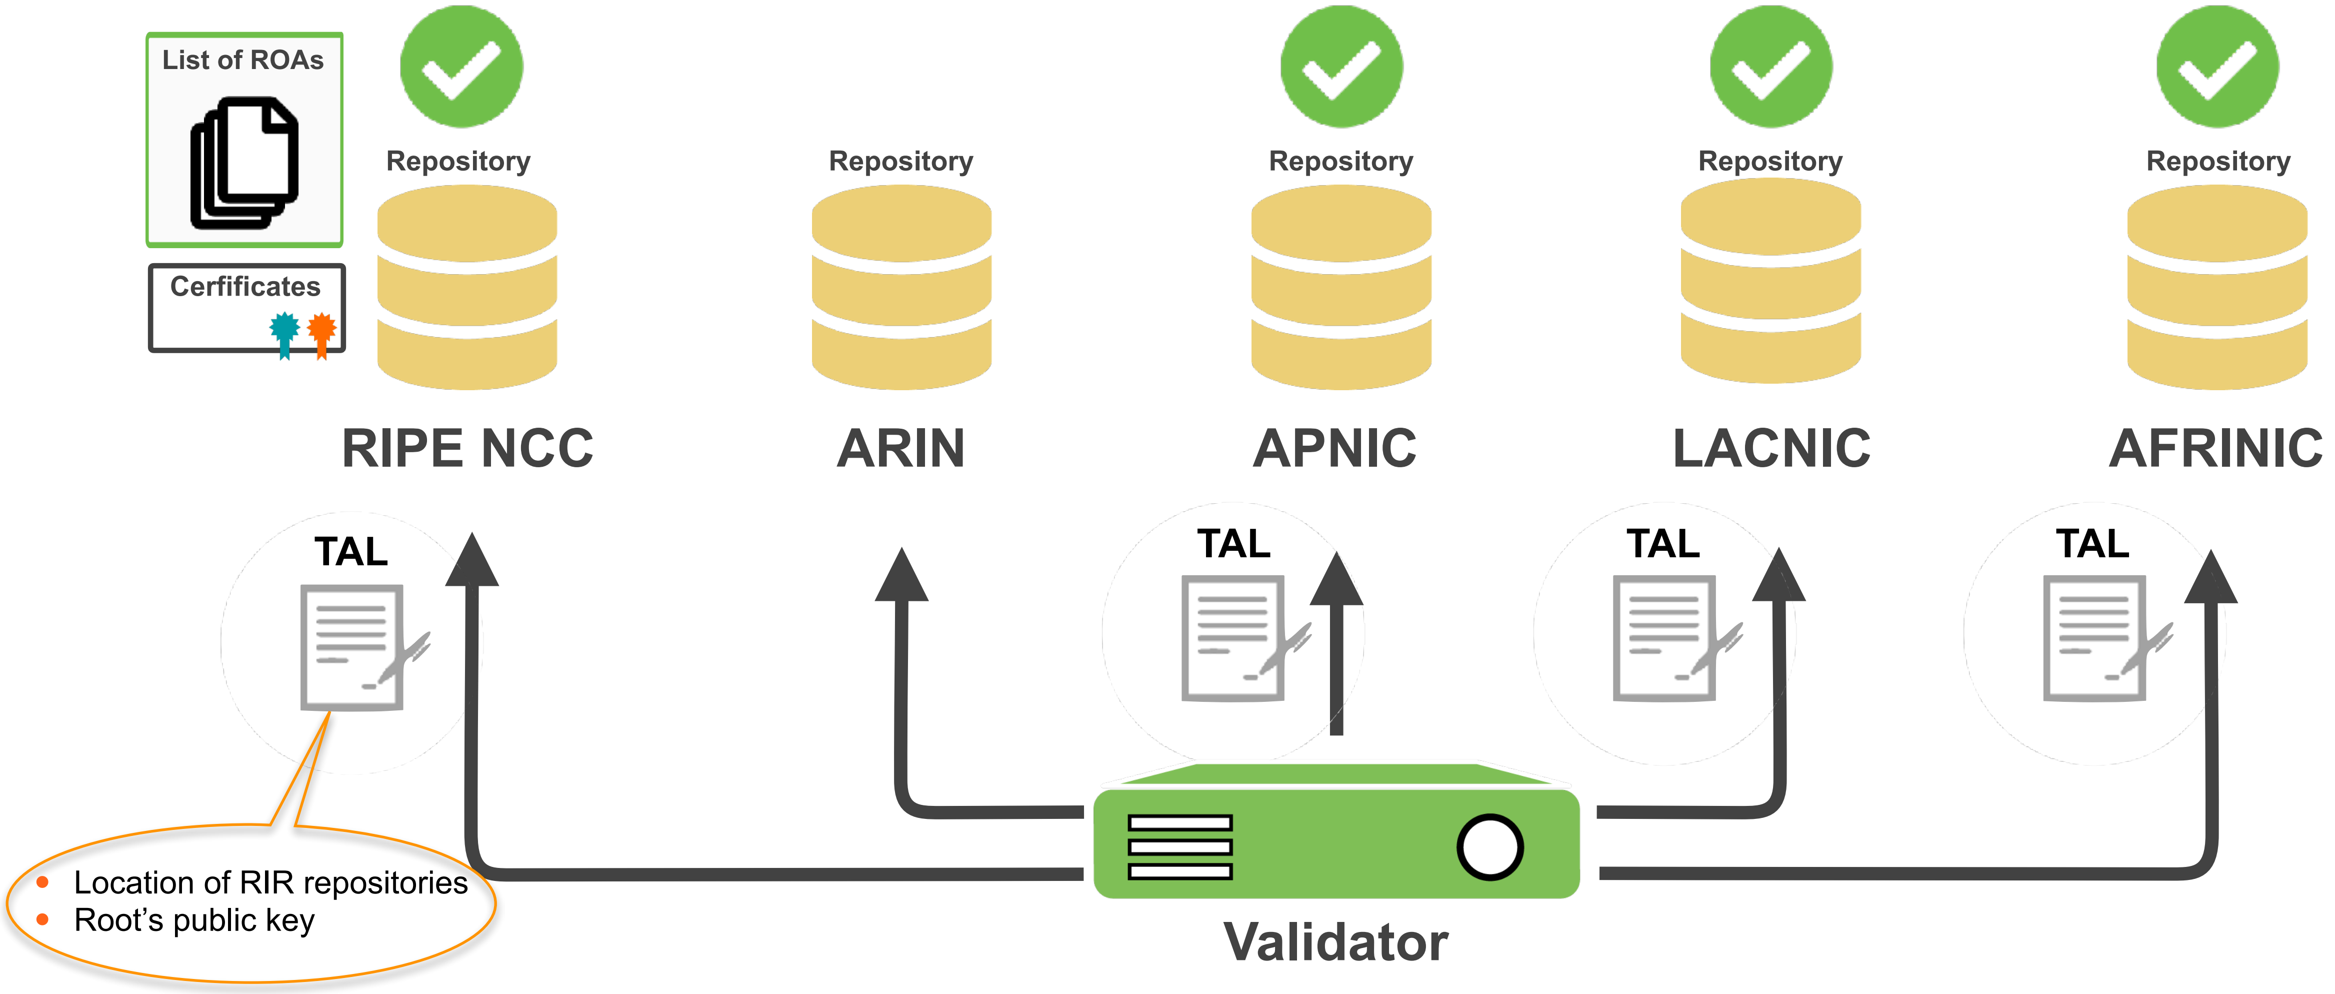

- Location of RIR repositories
- Root's public key

RPKI Validators

- Software that creates a local “validated cache” with all the valid ROAs
  - Downloads the RPKI repository from the RIRs
  - Validates the chain of trust of all the ROAs and associated CAs
  - Talks to your routers using the RPKI-RTR Protocol

Relying Party

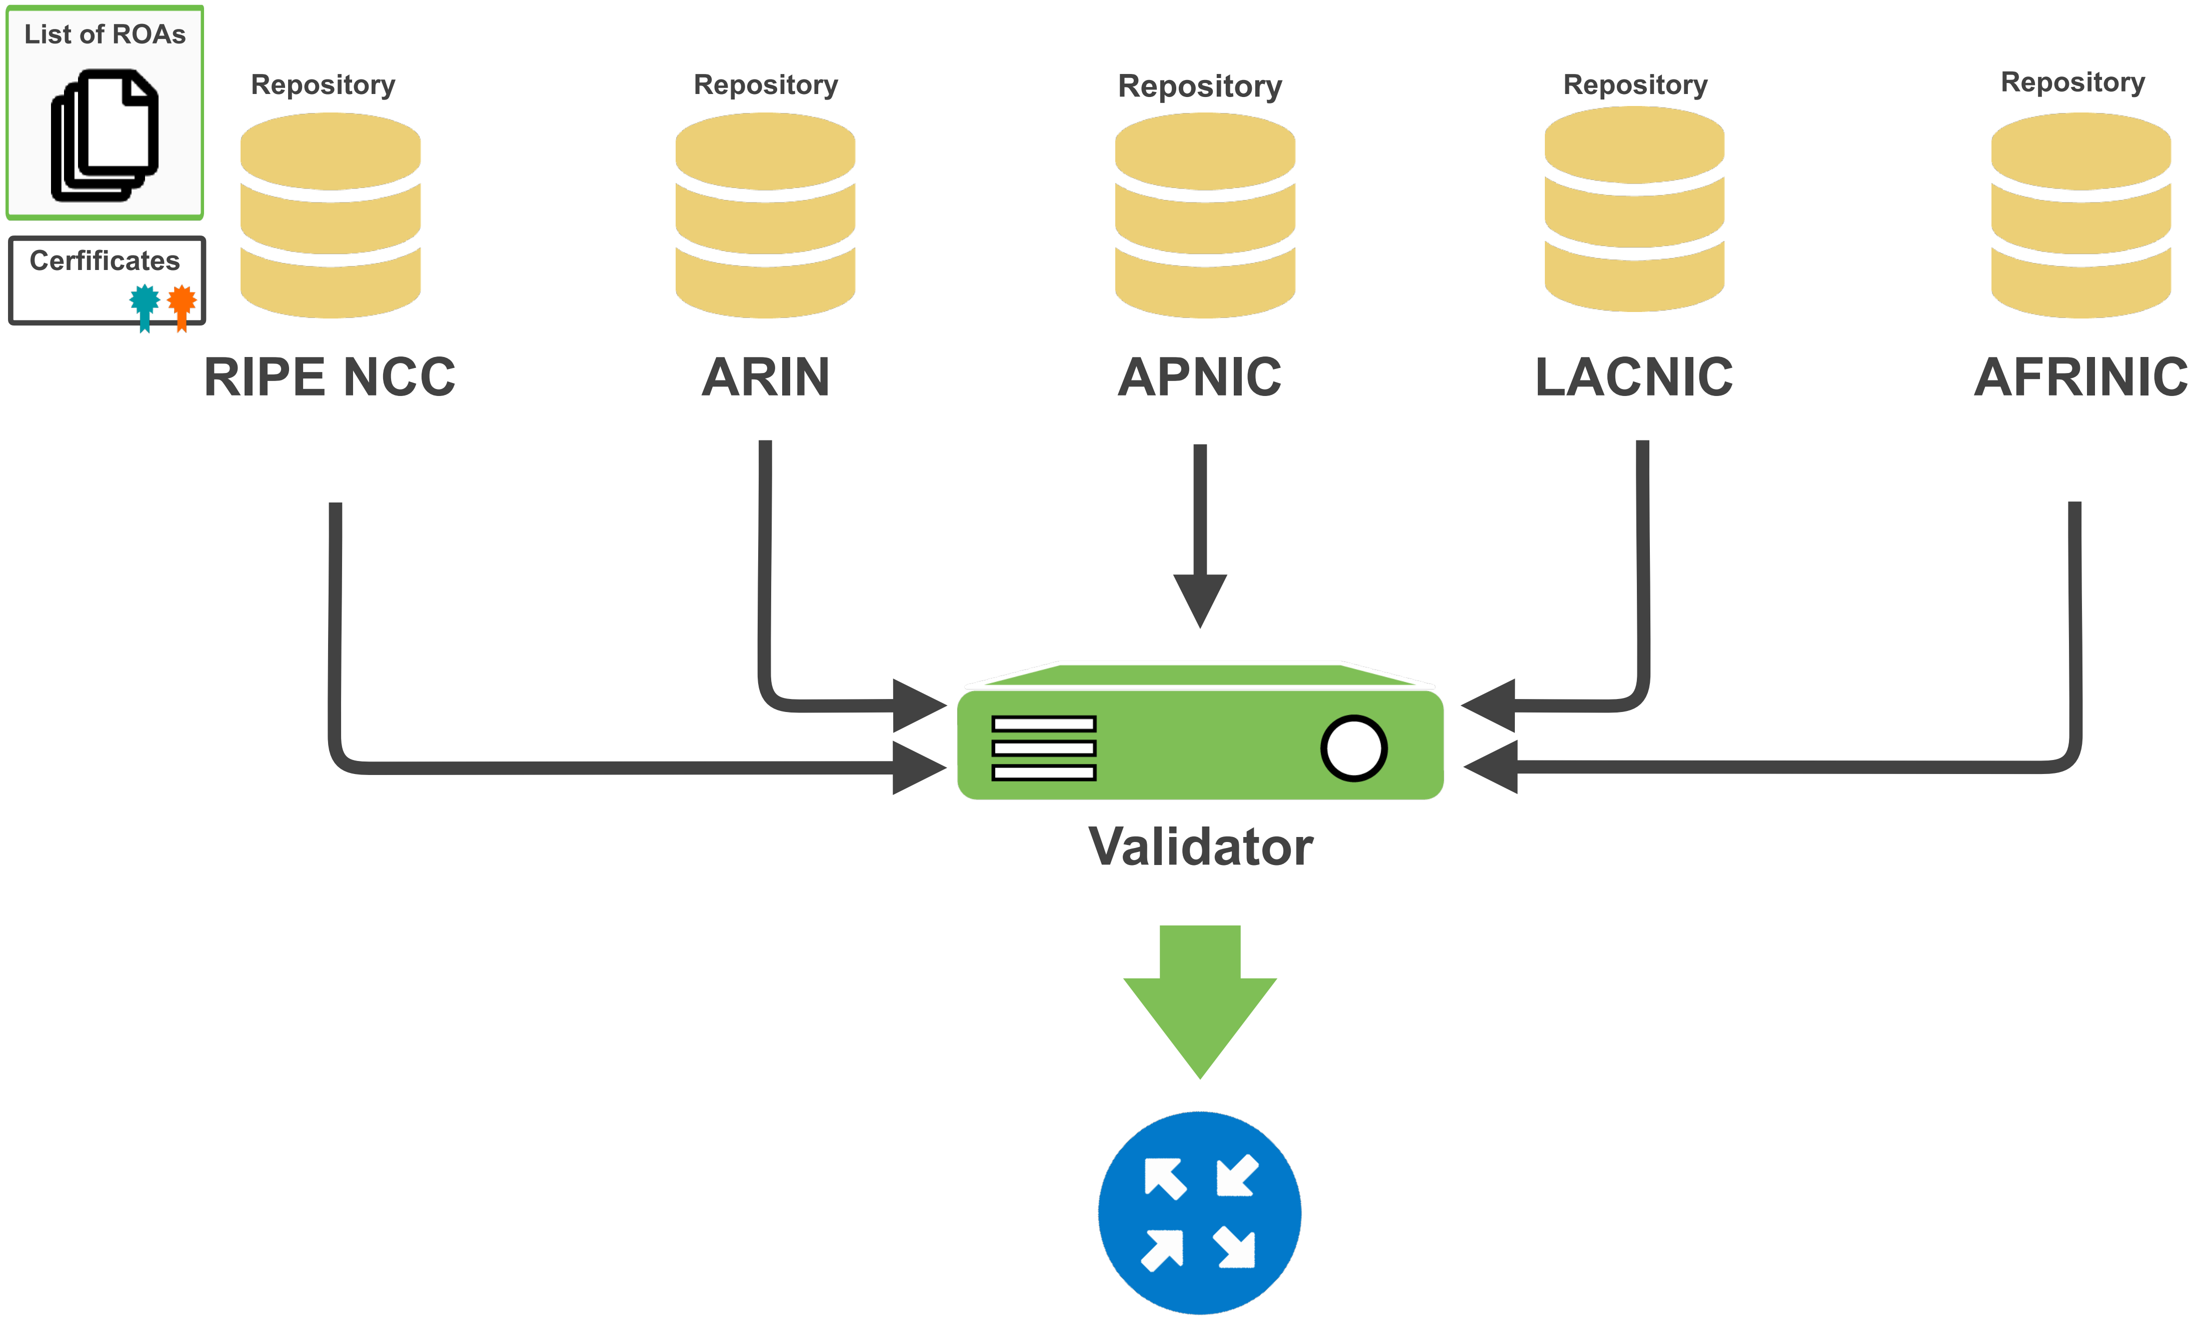

RPKI-RTR

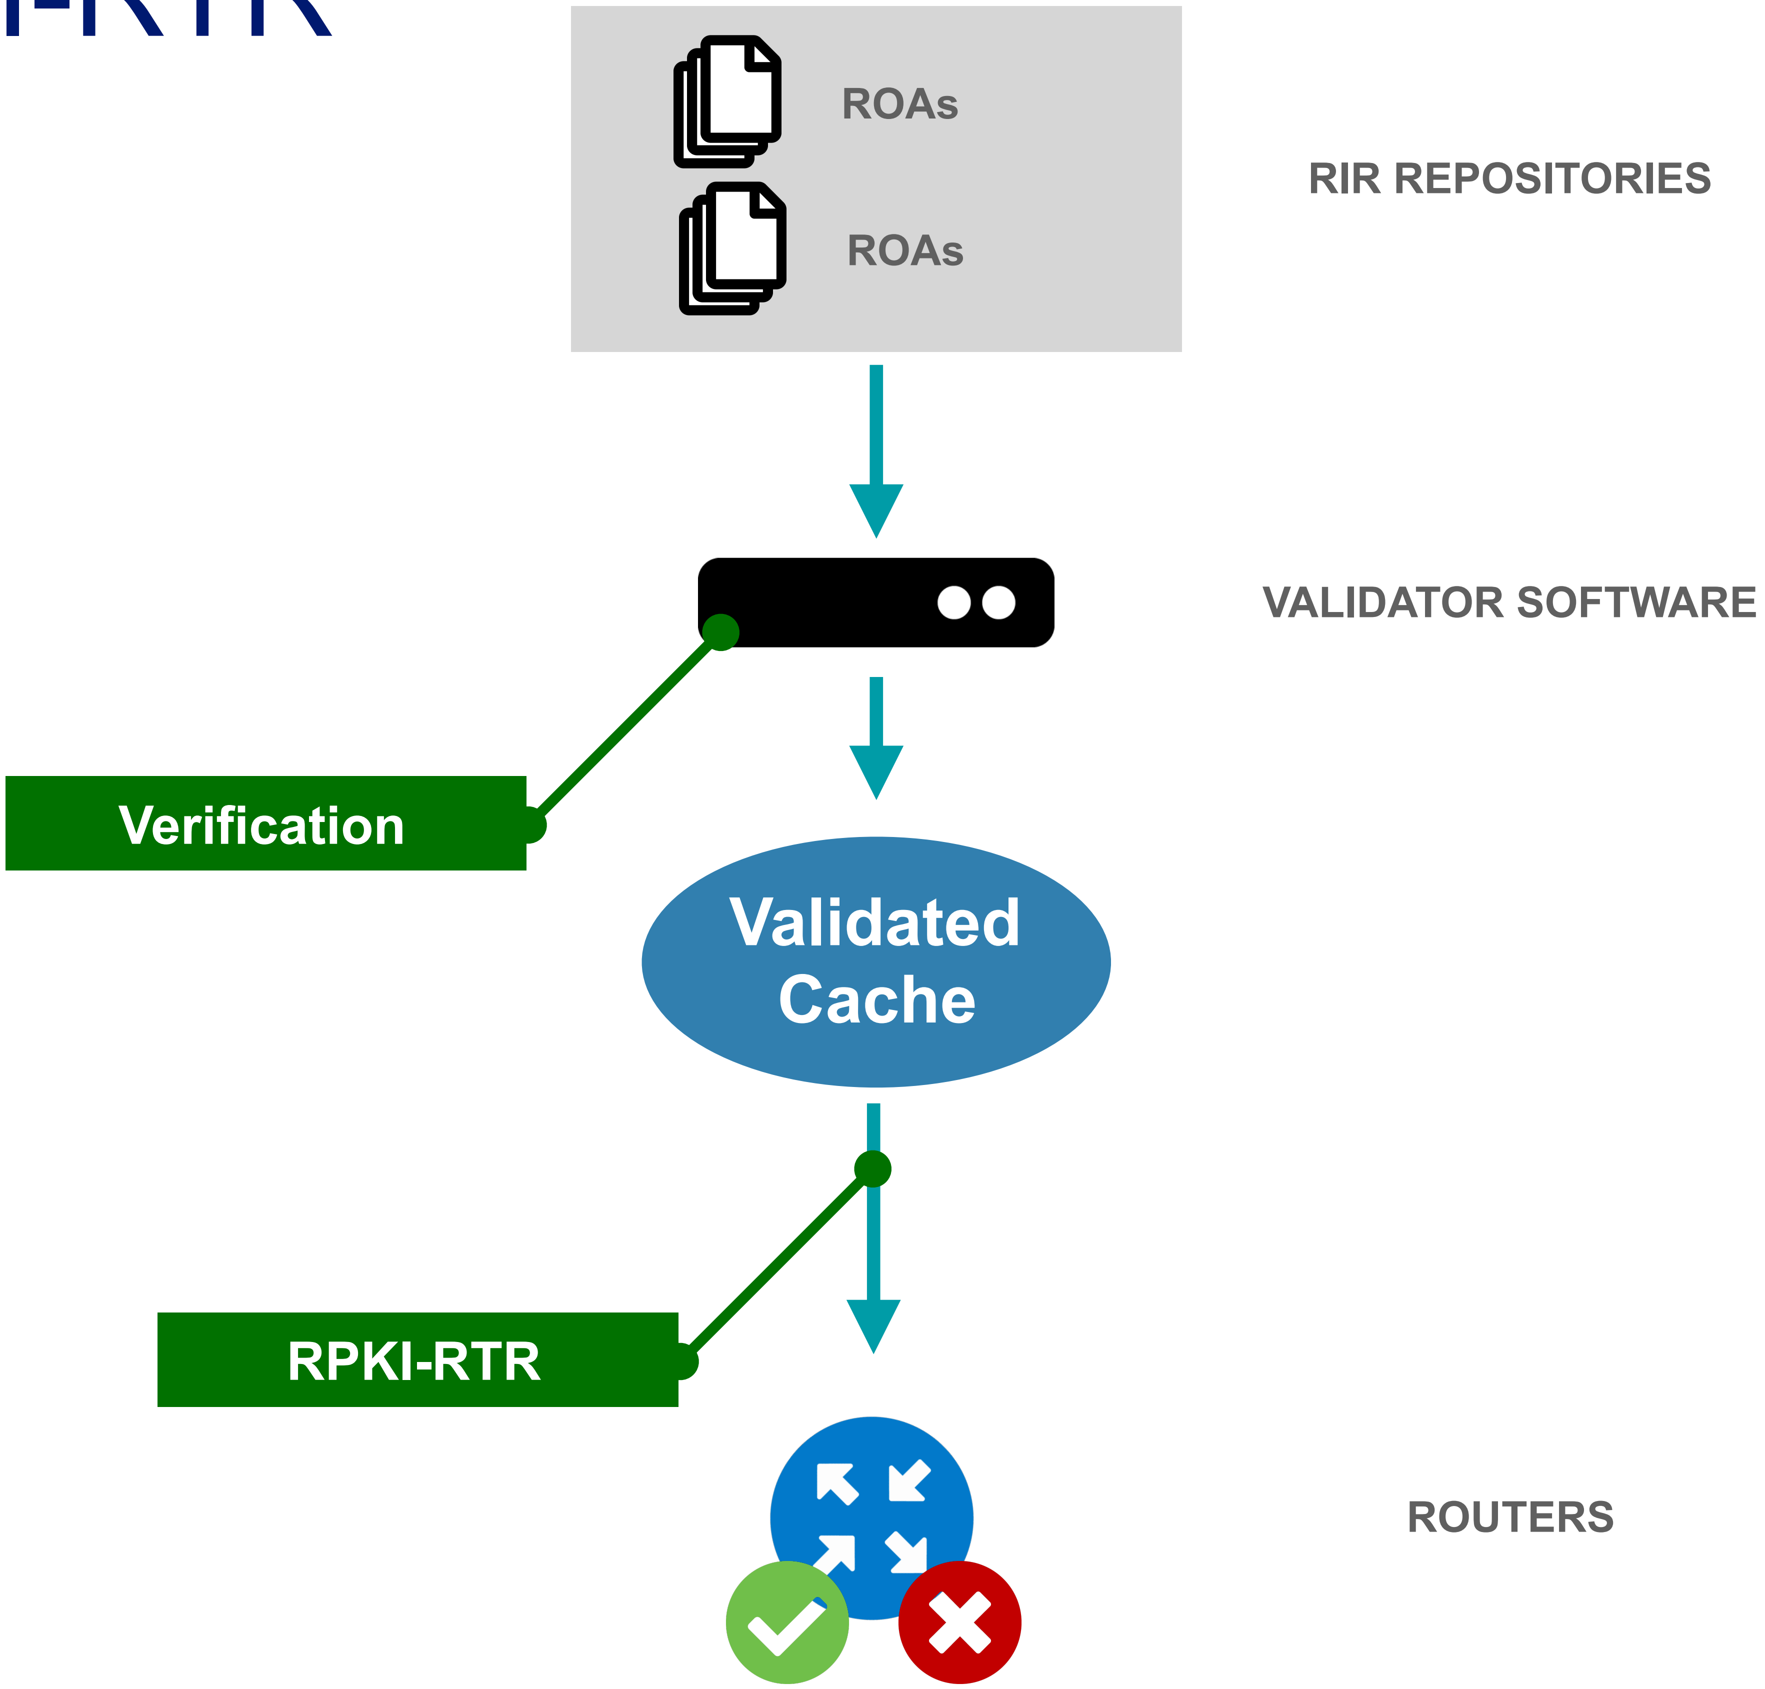

ROAs

ROAs

RIR REPOSITORIES

VALIDATOR SOFTWARE

Verification

Validated Cache

RPKI-RTR

ROUTERS

Relying Party

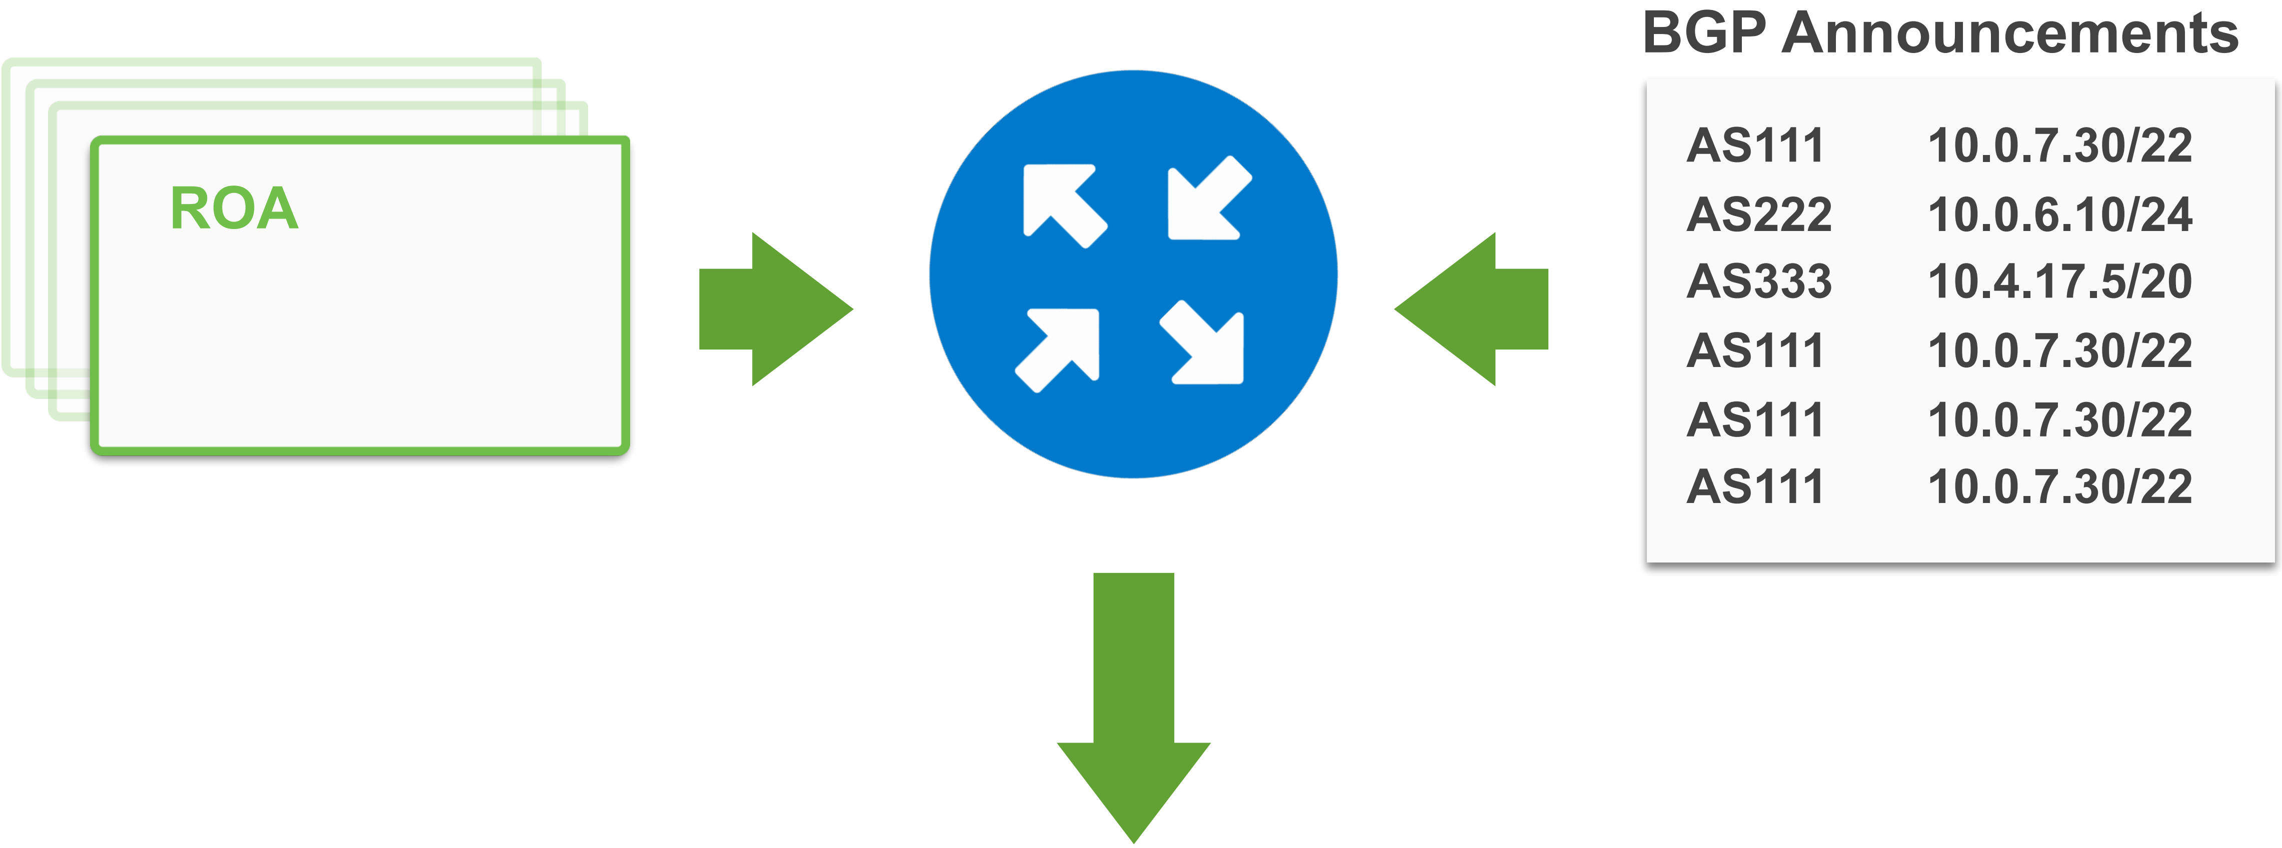

ROA

BGP Announcements

AS111	10.0.7.30/22
AS222	10.0.6.10/24
AS333	10.4.17.5/20
AS111	10.0.7.30/22
AS111	10.0.7.30/22
AS111	10.0.7.30/22

BETTER ROUTING DECISIONS

RIPE NCC Validator

- https://github.com/RIPE-NCC/rpki-validator
- Version 3.1
- Java-based, web interface, white-list functionality
- Can speak RPKI-RTR

DEPRECATED

Alternatives

- All are open source:
  - Routinator - https://github.com/NLnetLabs/routinator/
  - FORT - https://github.com/NICMx/FORT-validator/
  - OctoRPKI - https://github.com/cloudflare/cfrpki
  - RPKI-client - https://rpki-client.org/
  - Prover - https://github.com/lolepezy/rpki-prover
  - Rpstir2 - https://github.com/bgpsecurity/rpstir2

ROA Validation

Two elements of RPKI

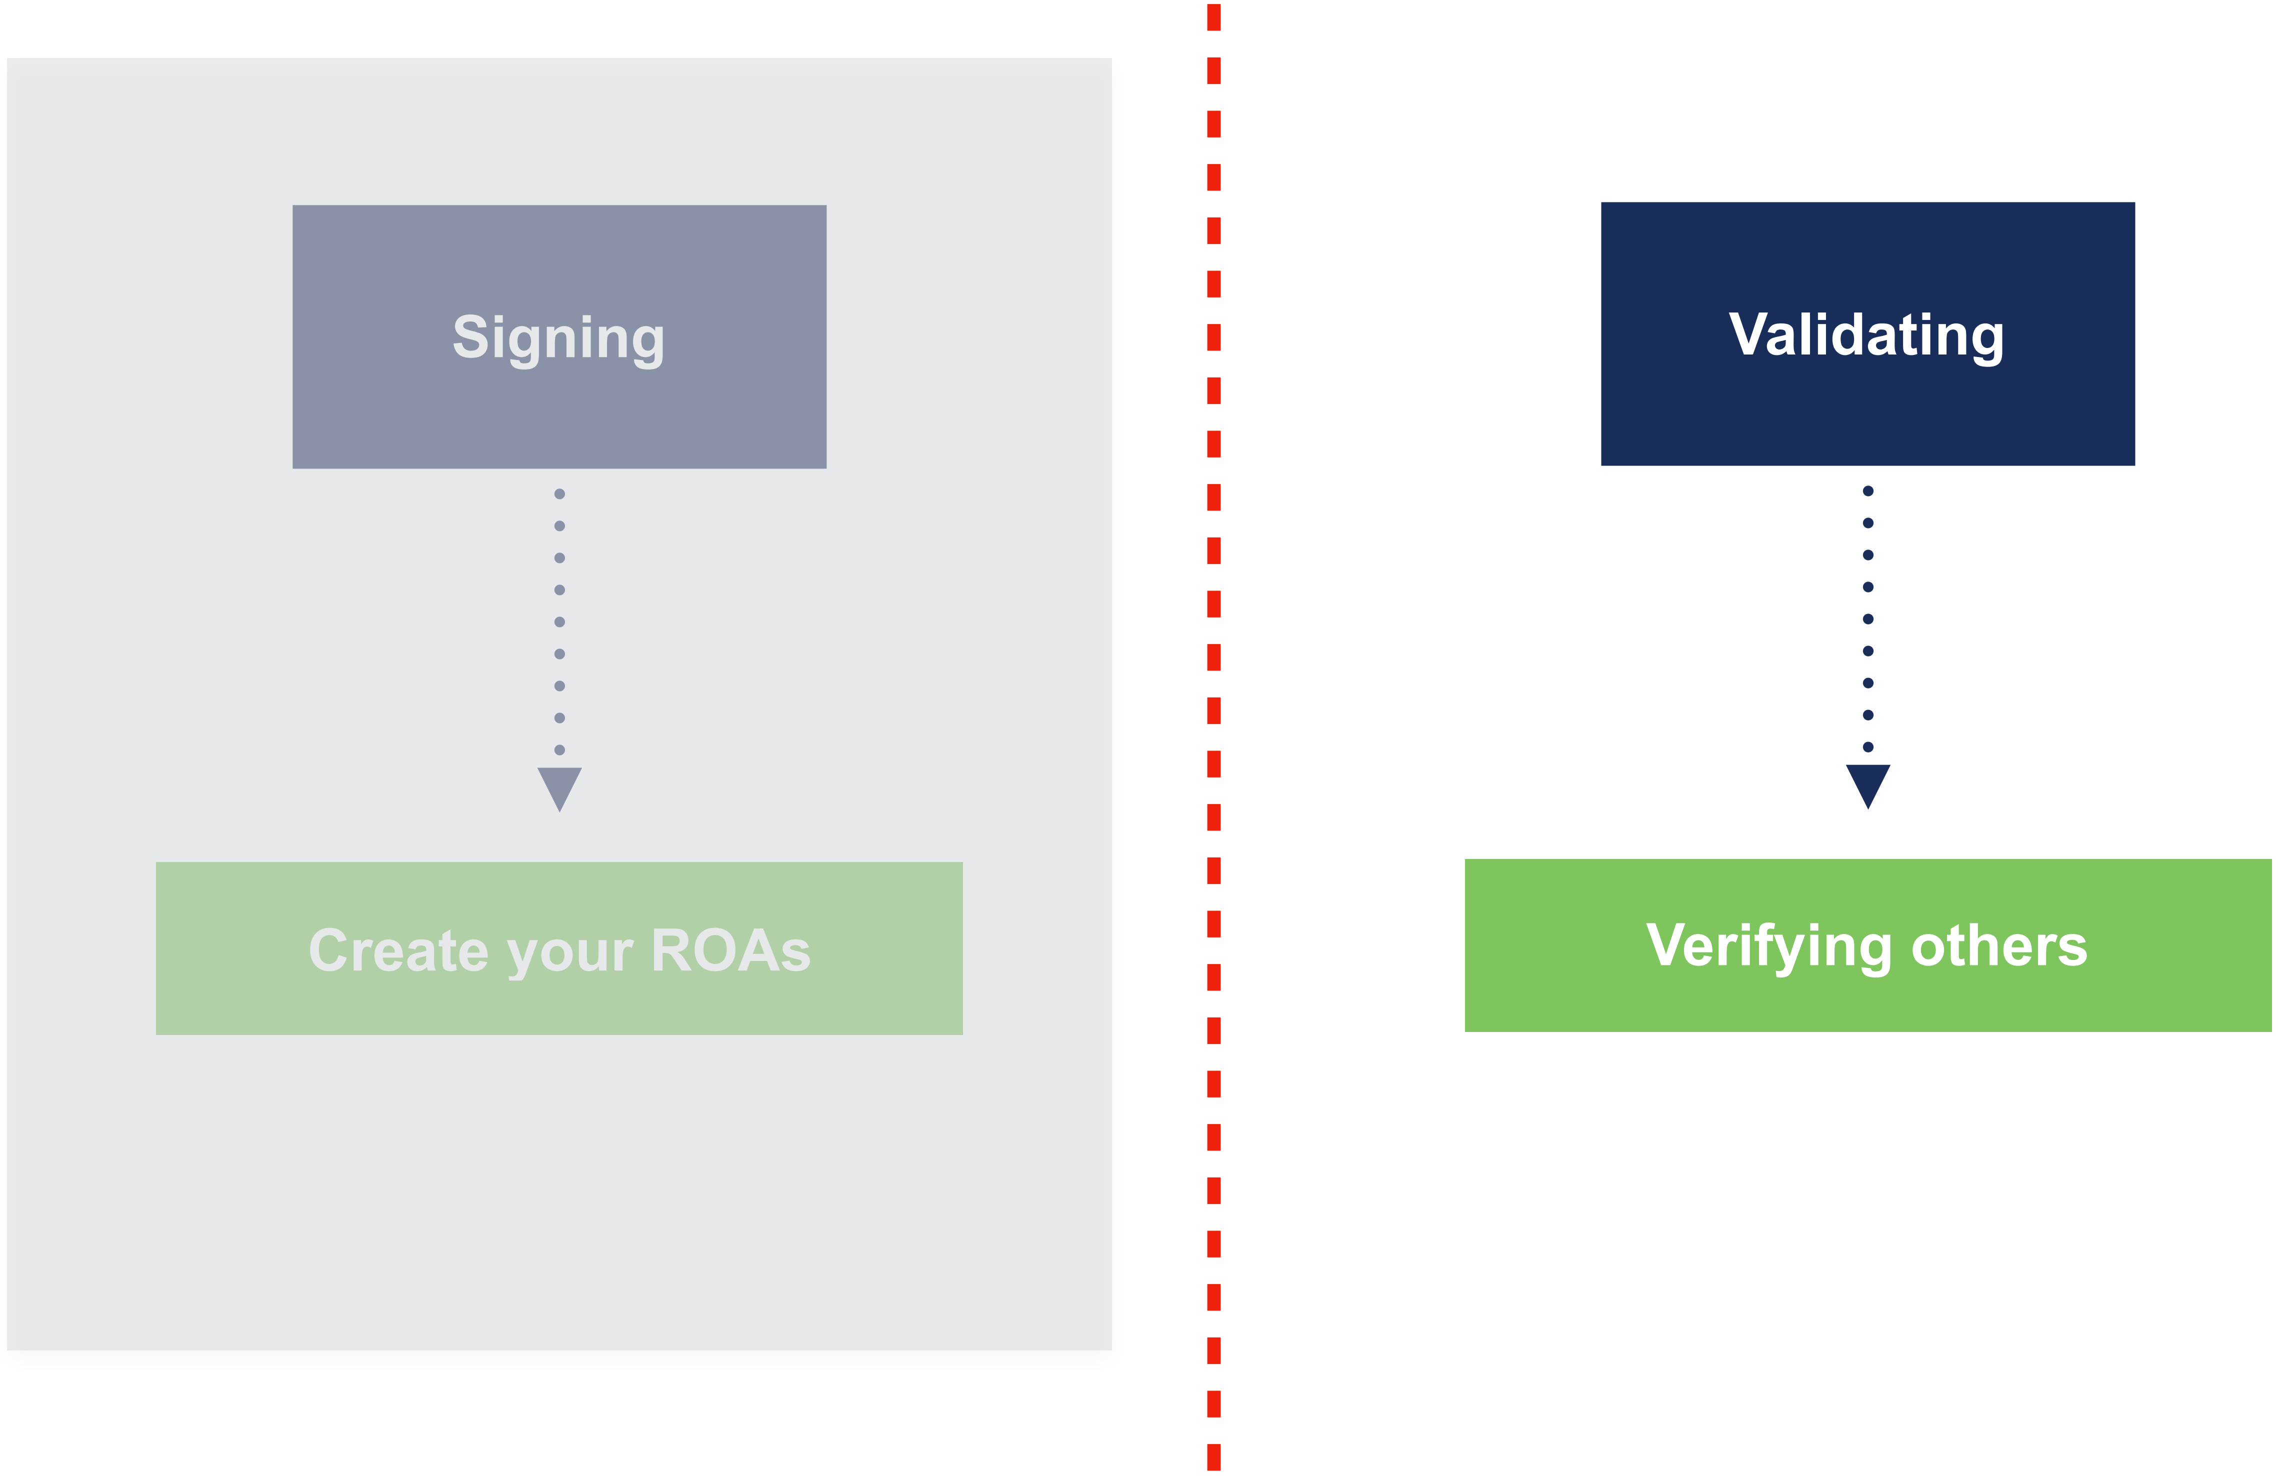

ROA Validation

- Routers receive data from the validated cache via RPKI-RTR

- Based on this and on BGP announcements, you have to make decisions
  - Accept or discard the BGP Announcement
  - As temporary measure, you could influence other attributes, such as Local Preference

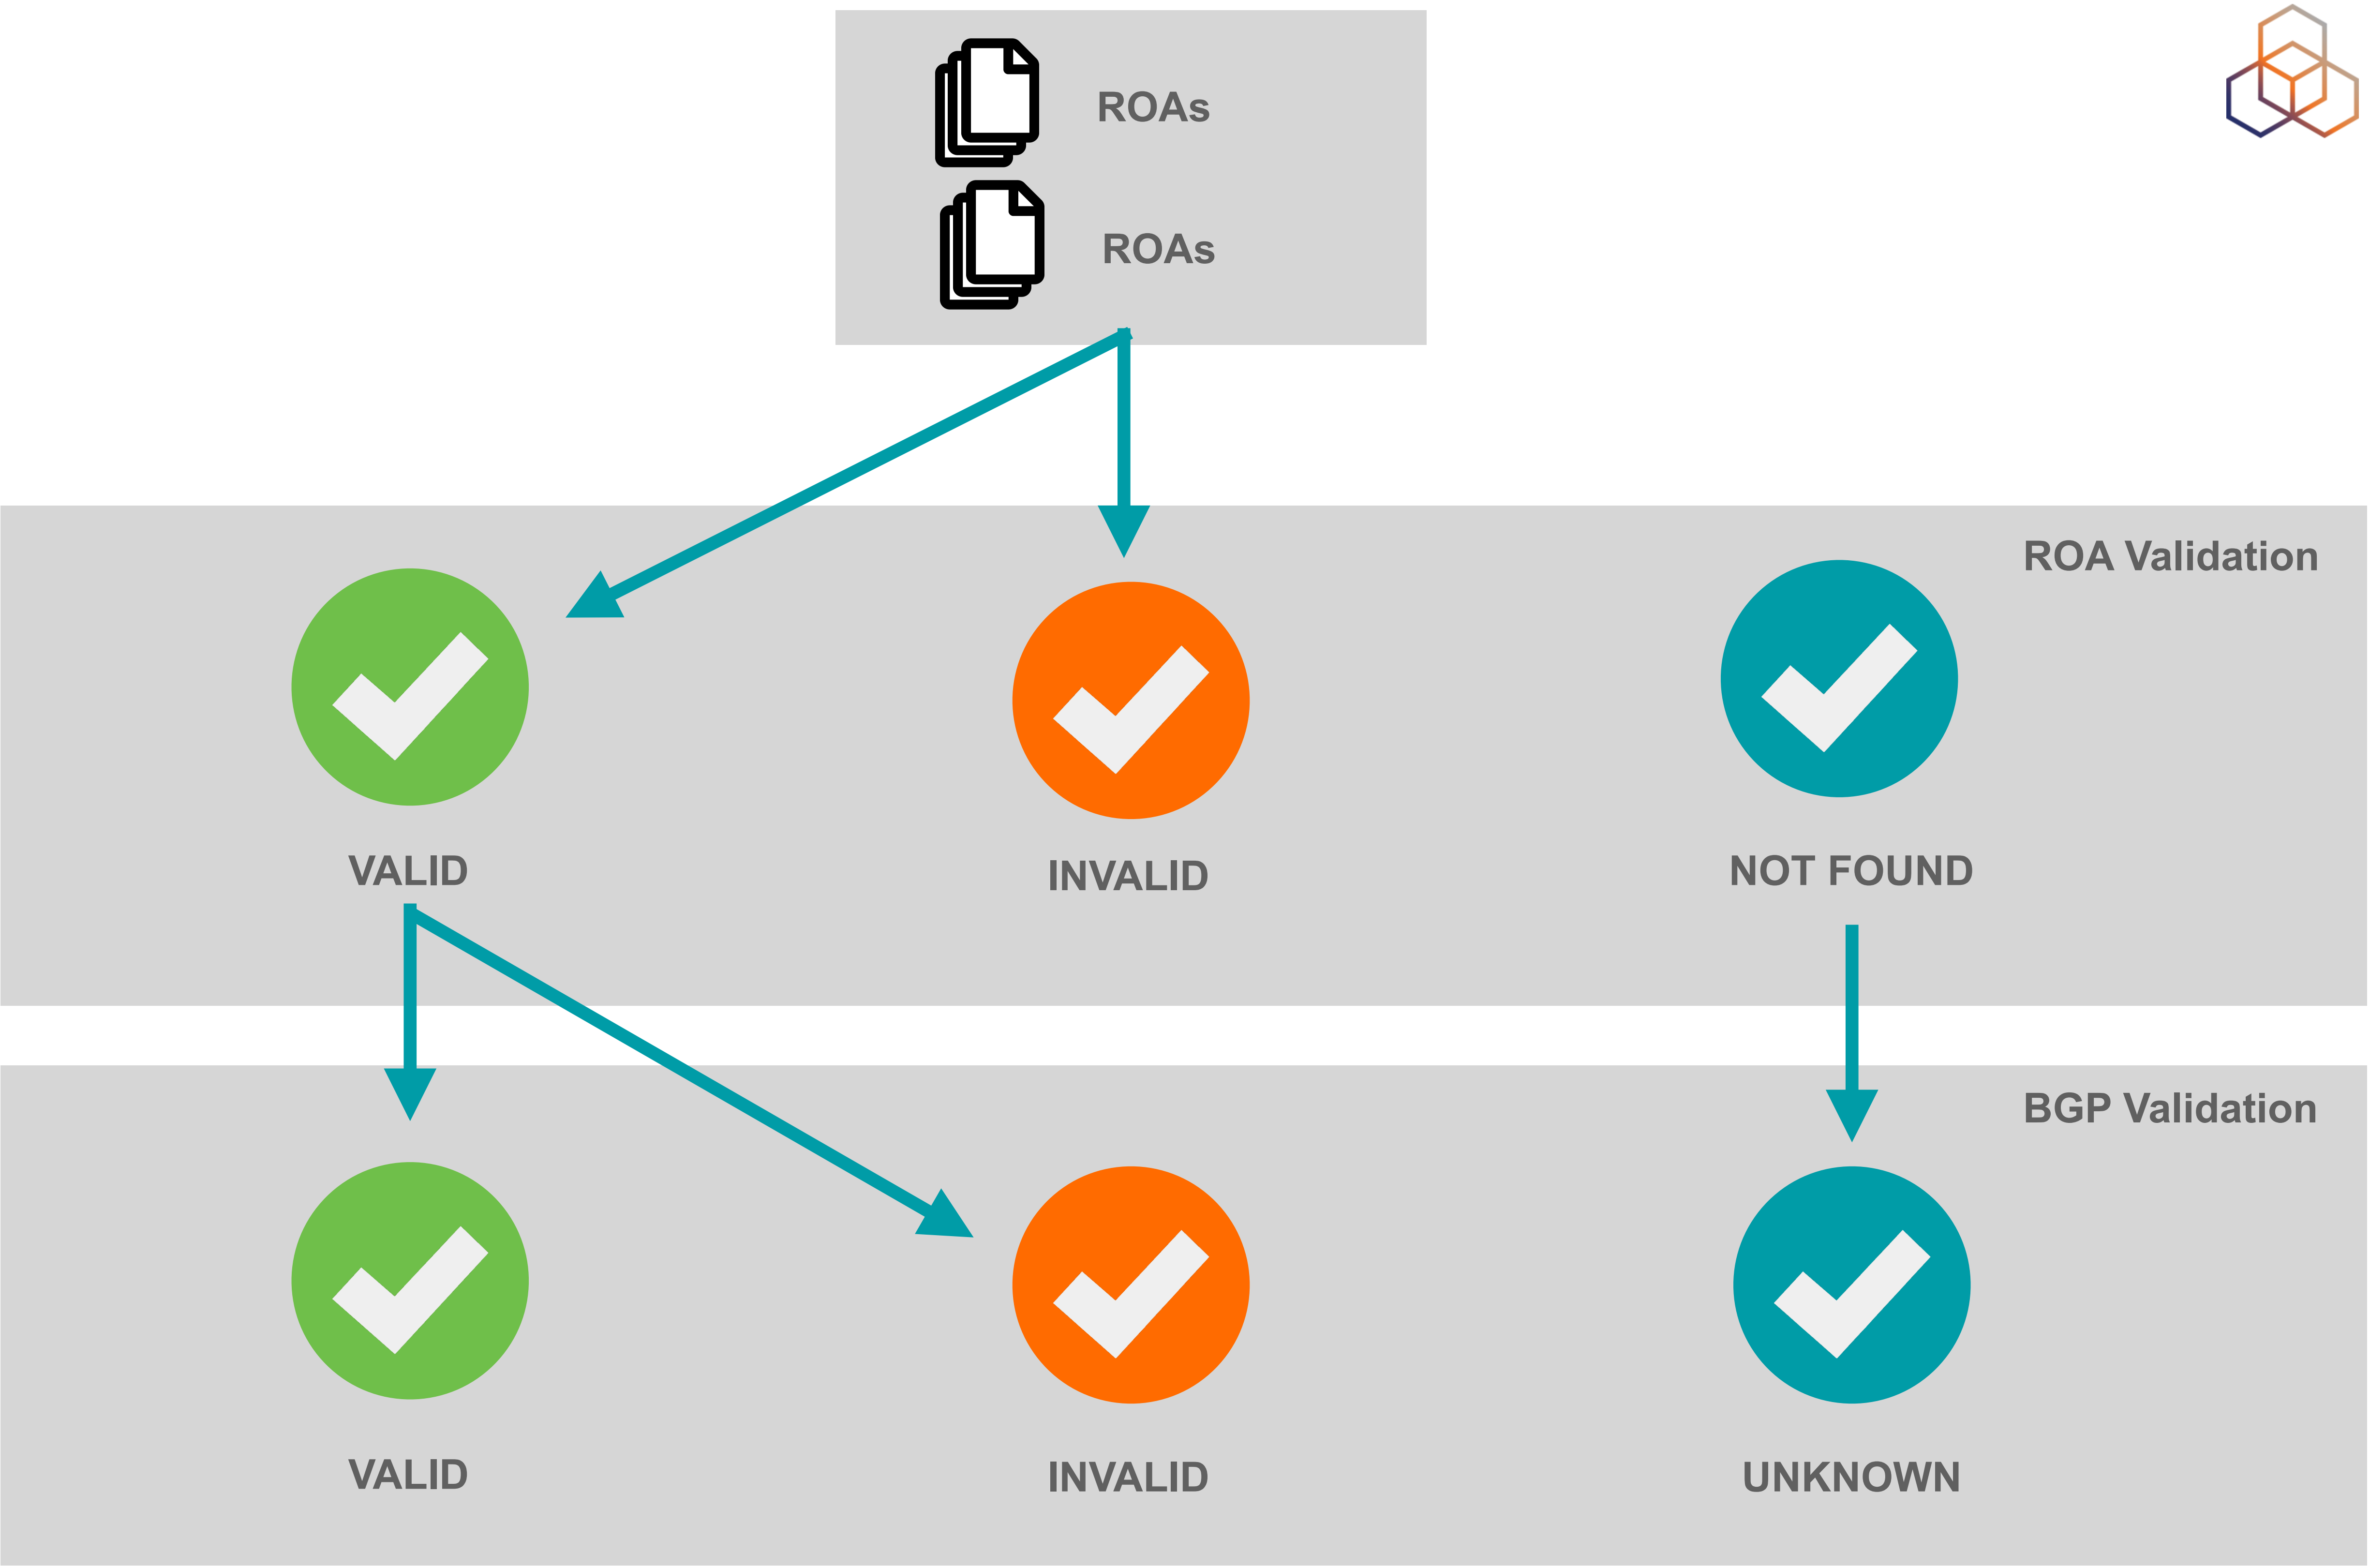
ROAs
ROAs
ROA Validation
VALID
INVALID
NOT FOUND
BGP Validation
VALID
INVALID
UNKNOWN

Invalid ROA

- Invalid ROA
  - The ROA in the repository cannot be validated by the client (ISP) so it is not included in the validated cache
- Invalid BGP announcement
  - There is a ROA in validated cache for that prefix but for a different AS.
  - Or the max length doesn’t match.
- If no ROA in the cache then announcement is “unknown”

Whitelisting

- If there is an invalid ROA for a network that’s important for you or your customers, you can whitelist it
- This is done on your local validator software
  - It creates a “fake” ROA for the resources you want
- It allows you to contact the operator to fix their ROA
  - Think of e-mail, contact forms, etc…

Take the Poll!

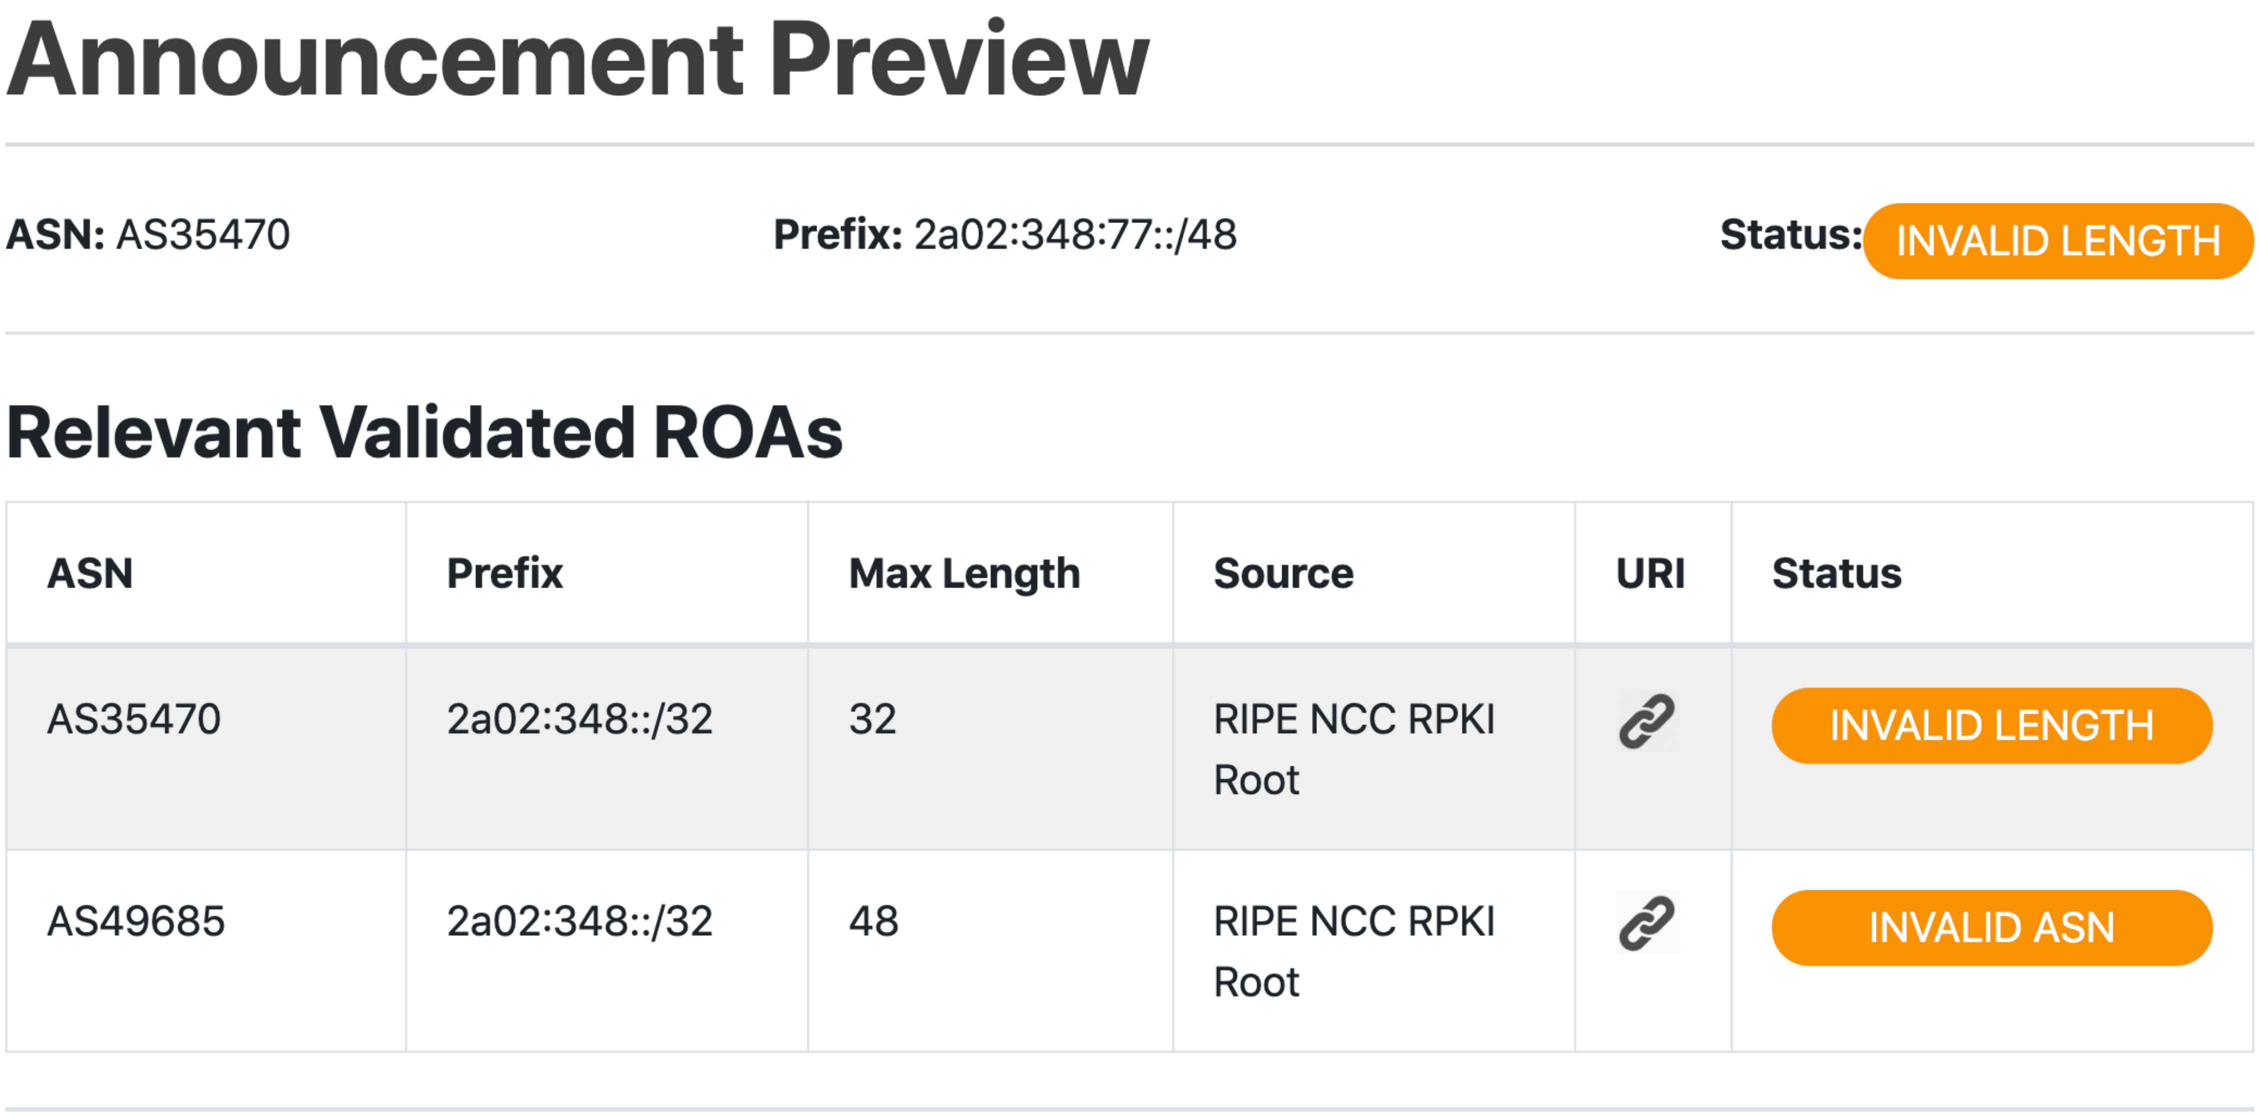

Announcement Preview

ASN: AS35470 **Prefix:** 2a02:348:77::/48 **Status:** INVALID LENGTH

Relevant Validated ROAs

ASN	Prefix	Max Length	Source	URI	Status
AS35470	2a02:348::/32	32	RIPE NCC RPKI Root	🔗	INVALID LENGTH
AS49685	2a02:348::/32	48	RIPE NCC RPKI Root	🔗	INVALID ASN

Status of RPKI ROV

Name	Type	Details	Status
Telia	Transit	Signed & Filtering	Safe
Cogent	Transit	Signed & Filtering	Safe
GTT	Transit	Signed & Filtering	Safe
NTT	Transit	Signed & Filtering	Safe
Hurricane Electric	Transit	Signed & Filtering	Safe
Tata	Transit	Signed & Filtering	Safe
PCCW	Transit	Signed & Filtering	Safe
RETN	Transit	Partially Signed &	Safe
Cloudflare	Cloud	Signed & Filtering	Safe
Amazon	Cloud	Signed & Filtering	Safe
Netflix	Cloud	Signed & Filtering	Safe
Wikimedia	Cloud	Signed & Filtering	Safe
Scaleway	Cloud	Signed & Filtering	Safe

- Source: isbgpsafeyet.com

Where do we go from here ?

- RPKI is only one of the steps towards full BGP Validation
  - Paths are not validated

- We need more building blocks
  - BGPSec (RFC)
  - ASPA (draft)
  - AS-Cones (draft)

Questions

nathalie@ripe.net

rpki@ripe.net