RIPE NCC

RIPE NETWORK COORDINATION CENTRE

BGP Security

RPKI

Massimiliano Stucchi, Ondřej Caletka

RIPE NCC Learning & Development

Agenda

Introduction

ROAs

Demo: Create ROAs

Deploying RPKI Validators

Demo: Running Validators

Validation

Demo: Setting up BGP Origin Validation

Demo: Discarding BGP Invalids

Introduction

Section 1

Routing on the Internet

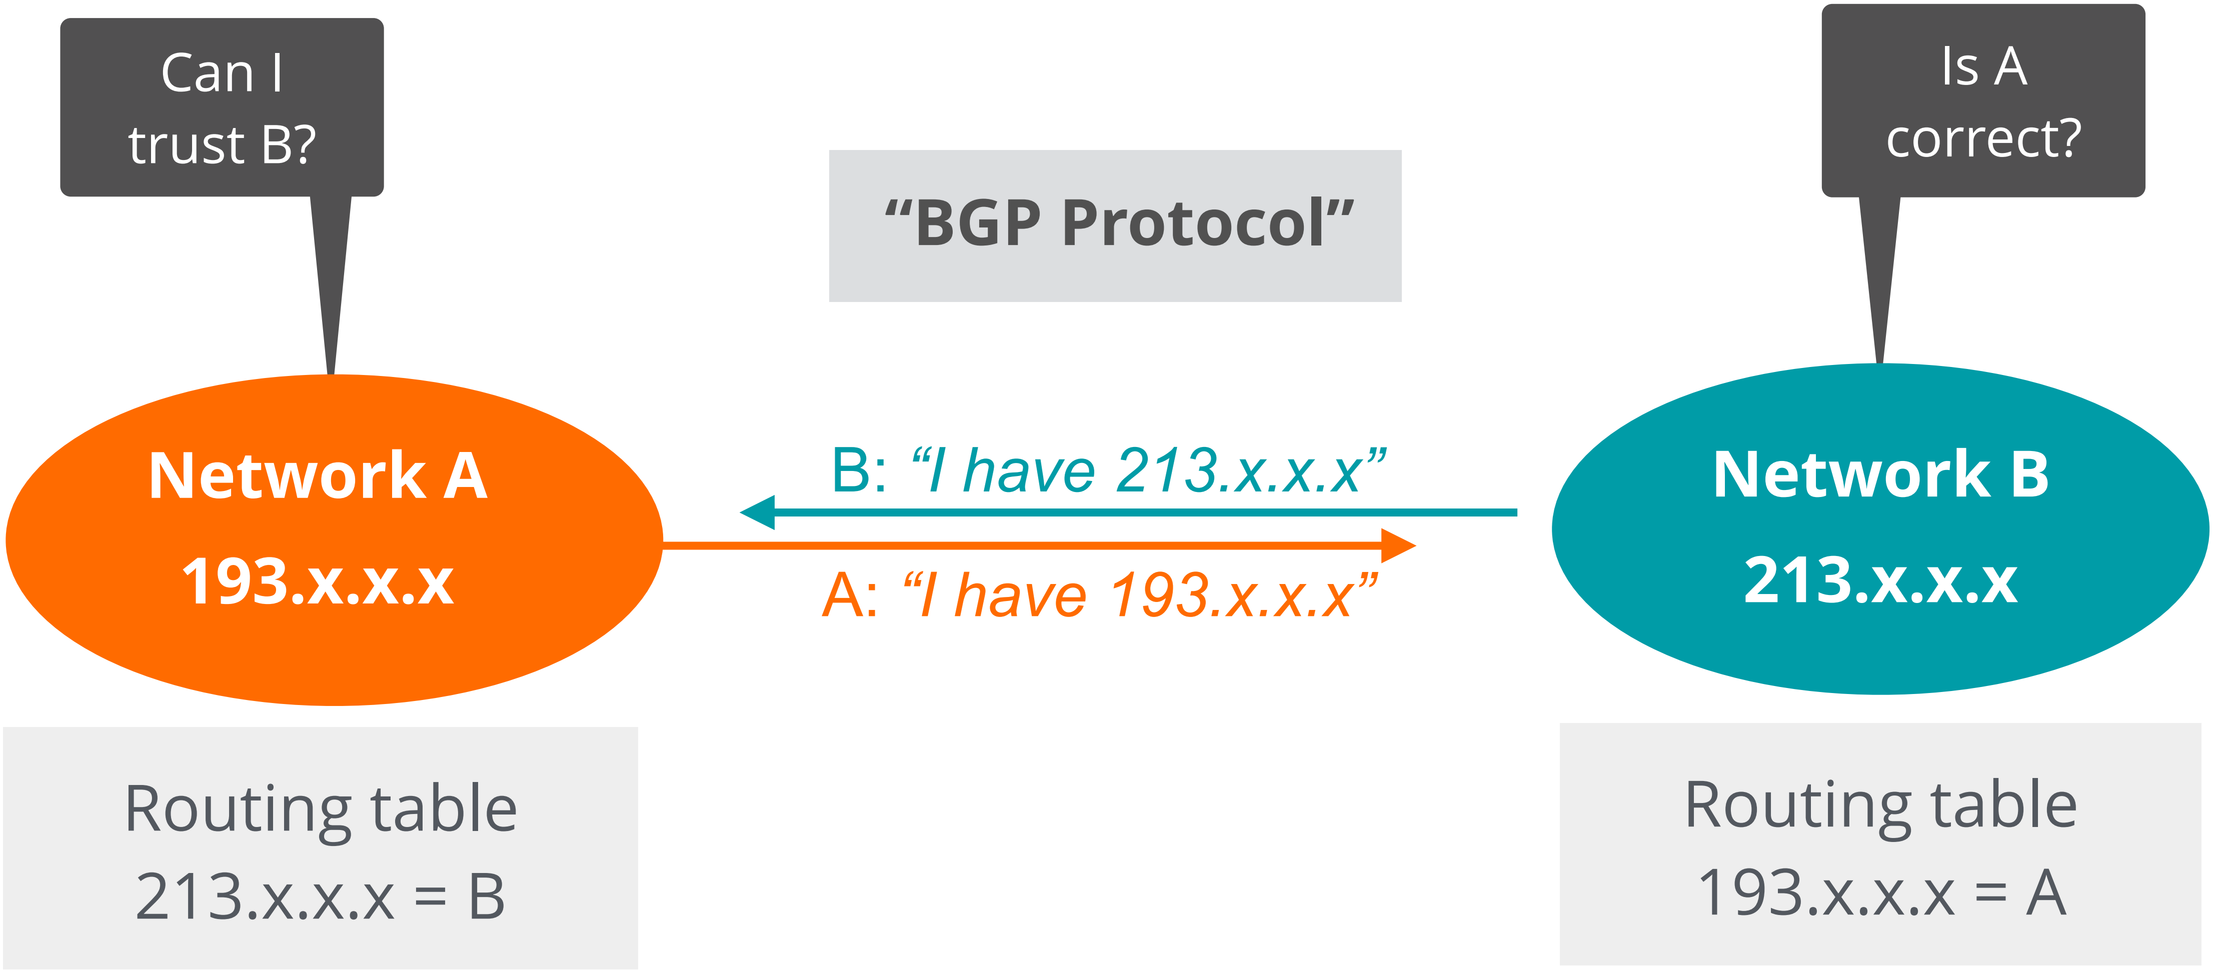

Can I trust B?

“BGP Protocol”

Is A correct?

Network A

193.x.x.x

B: *“I have 213.x.x.x”*

A: *“I have 193.x.x.x”*

Network B

213.x.x.x

Routing table
213.x.x.x = B

Routing table
193.x.x.x = A

How can you have secure routing?

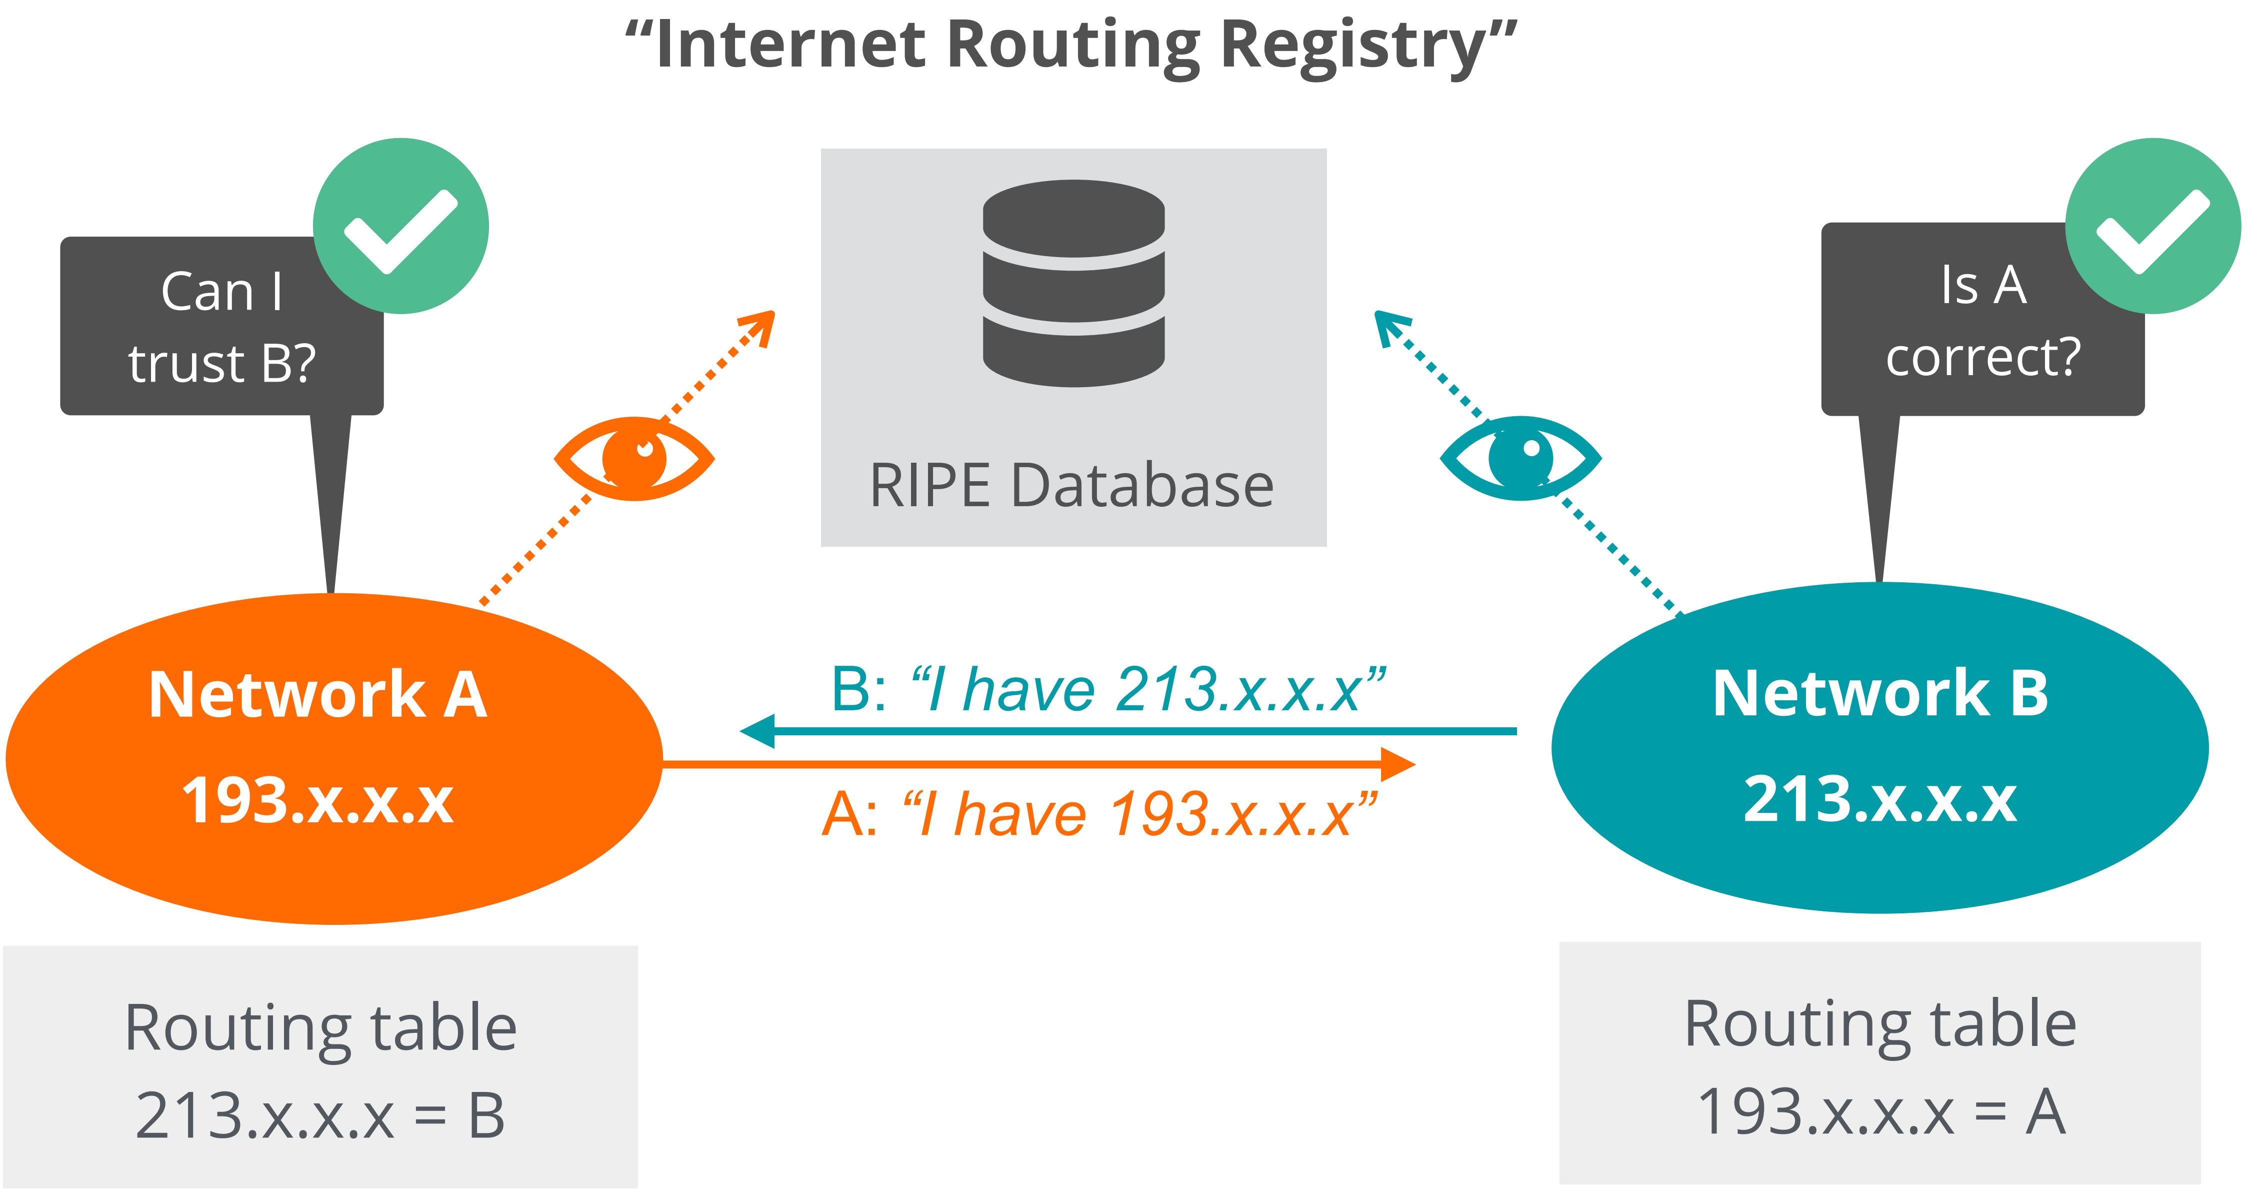

“Internet Routing Registry”

RIPE Database

Can I trust B?

Is A correct?

Network A

193.x.x.x

Network B

213.x.x.x

B: *“I have 213.x.x.x”*

A: *“I have 193.x.x.x”*

Routing table
213.x.x.x = B

Routing table
193.x.x.x = A

Question

Is the **Internet Routing Registry** (IRR) enough for BGP security?

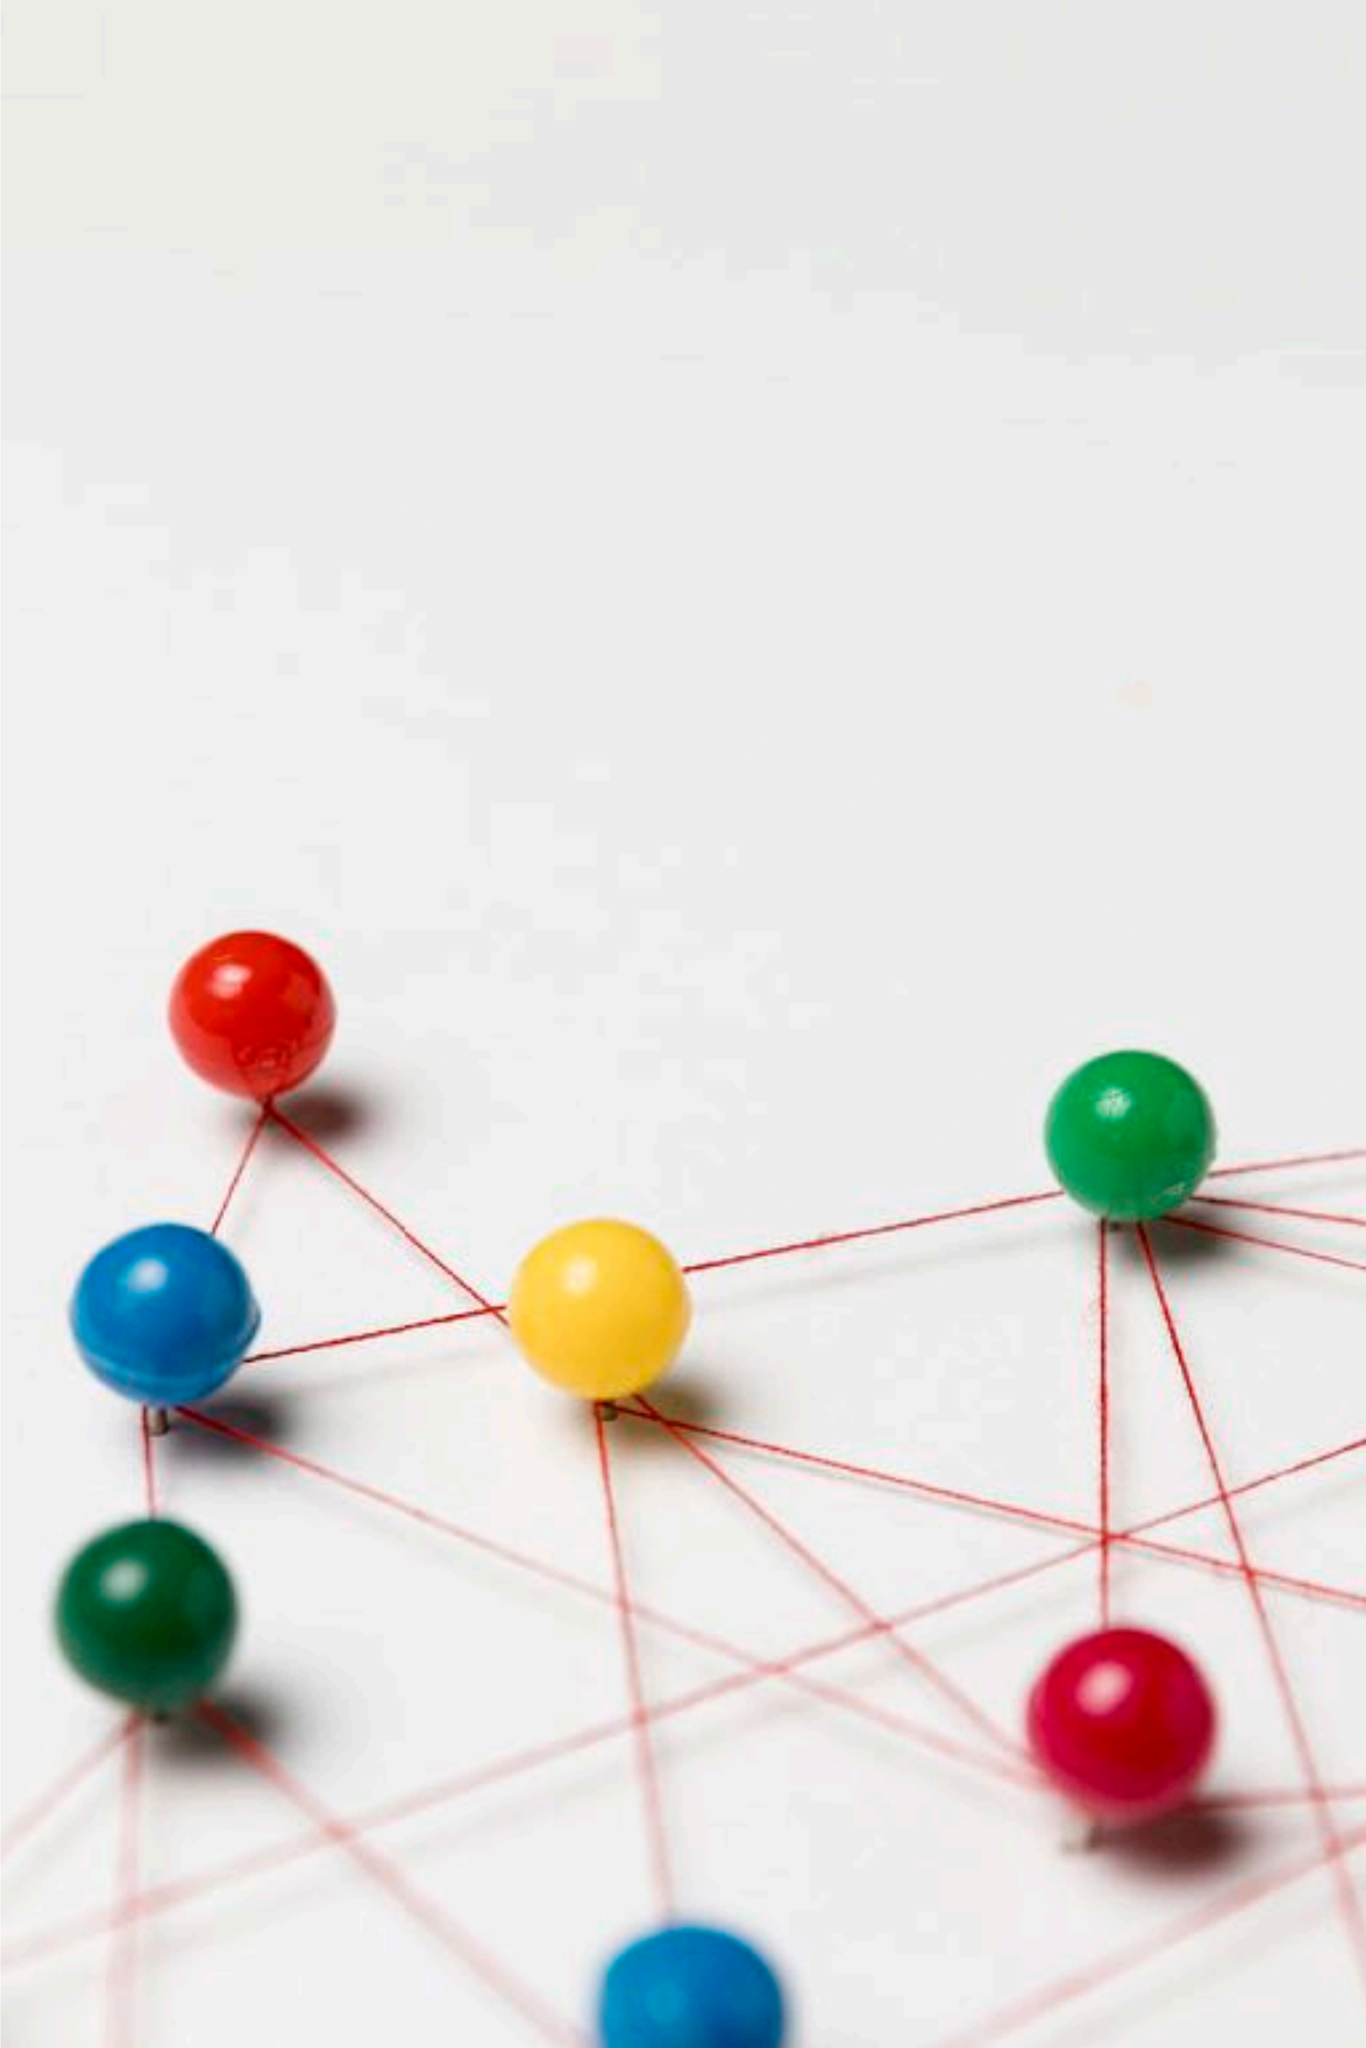

⏱ 1 min.

Problem Statement

- Some IRR data **cannot** 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

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

A Short History

- Operated since 2008 by all RIRs
  - Community-driven standardisation (IETF)

- Adds crypto-security to IP addresses and ASNs
  - Provides **data you can trust**

RPKI Chain of Trust

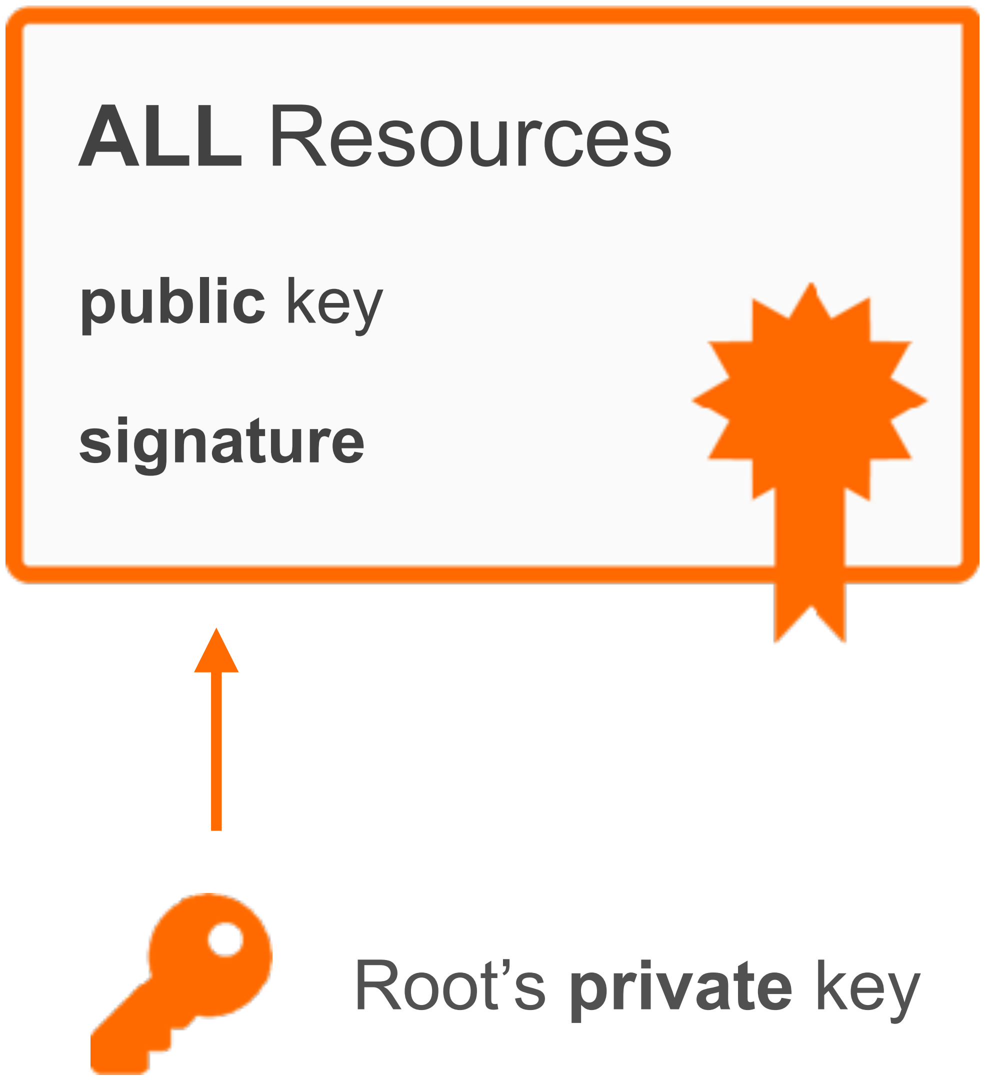

ALL Resources

public key

signature

Root’s **private** key

RIPE NCC Root Certificate

Self-signed

RPKI Chain of Trust

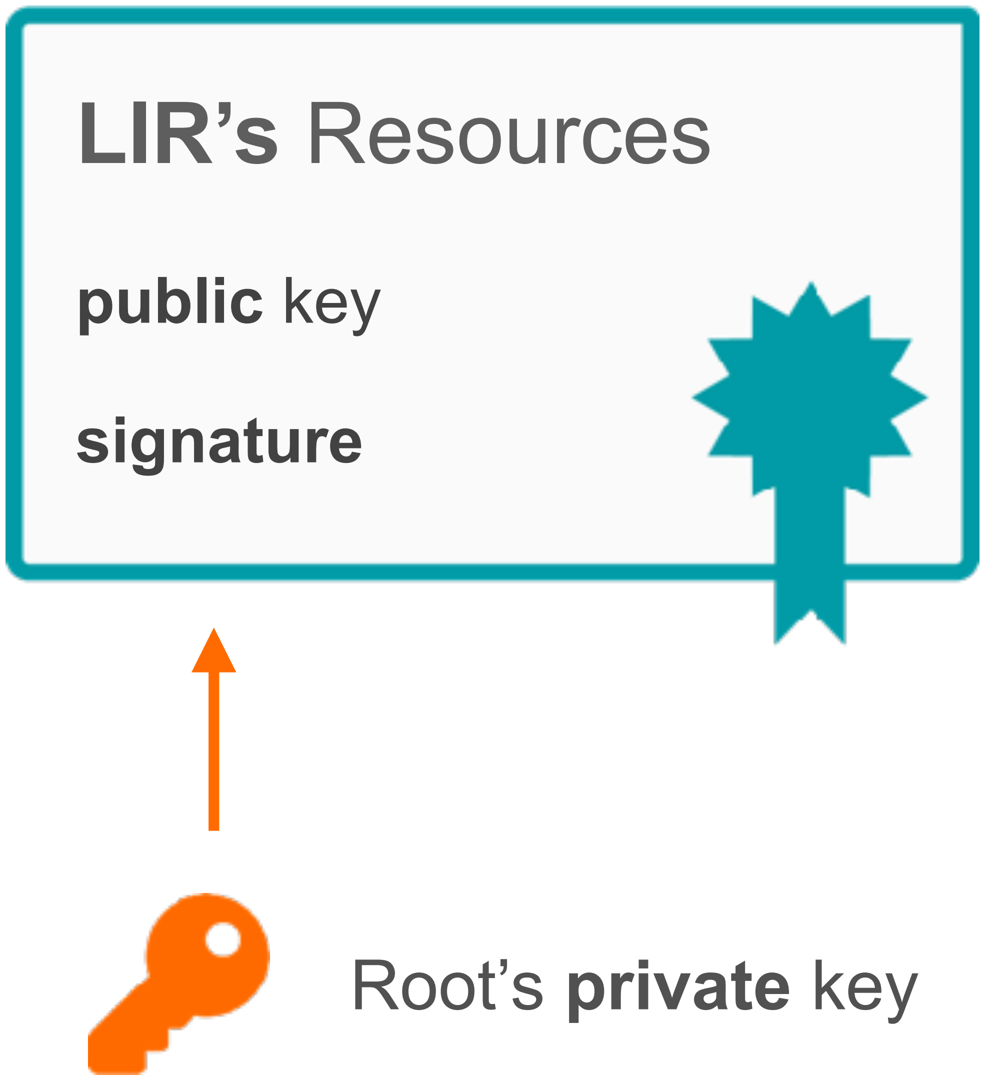

LIR’s Resources

public key

signature

Root’s **private** key

LIR Certificate

Signed by the Root private key

RPKI Chain of Trust

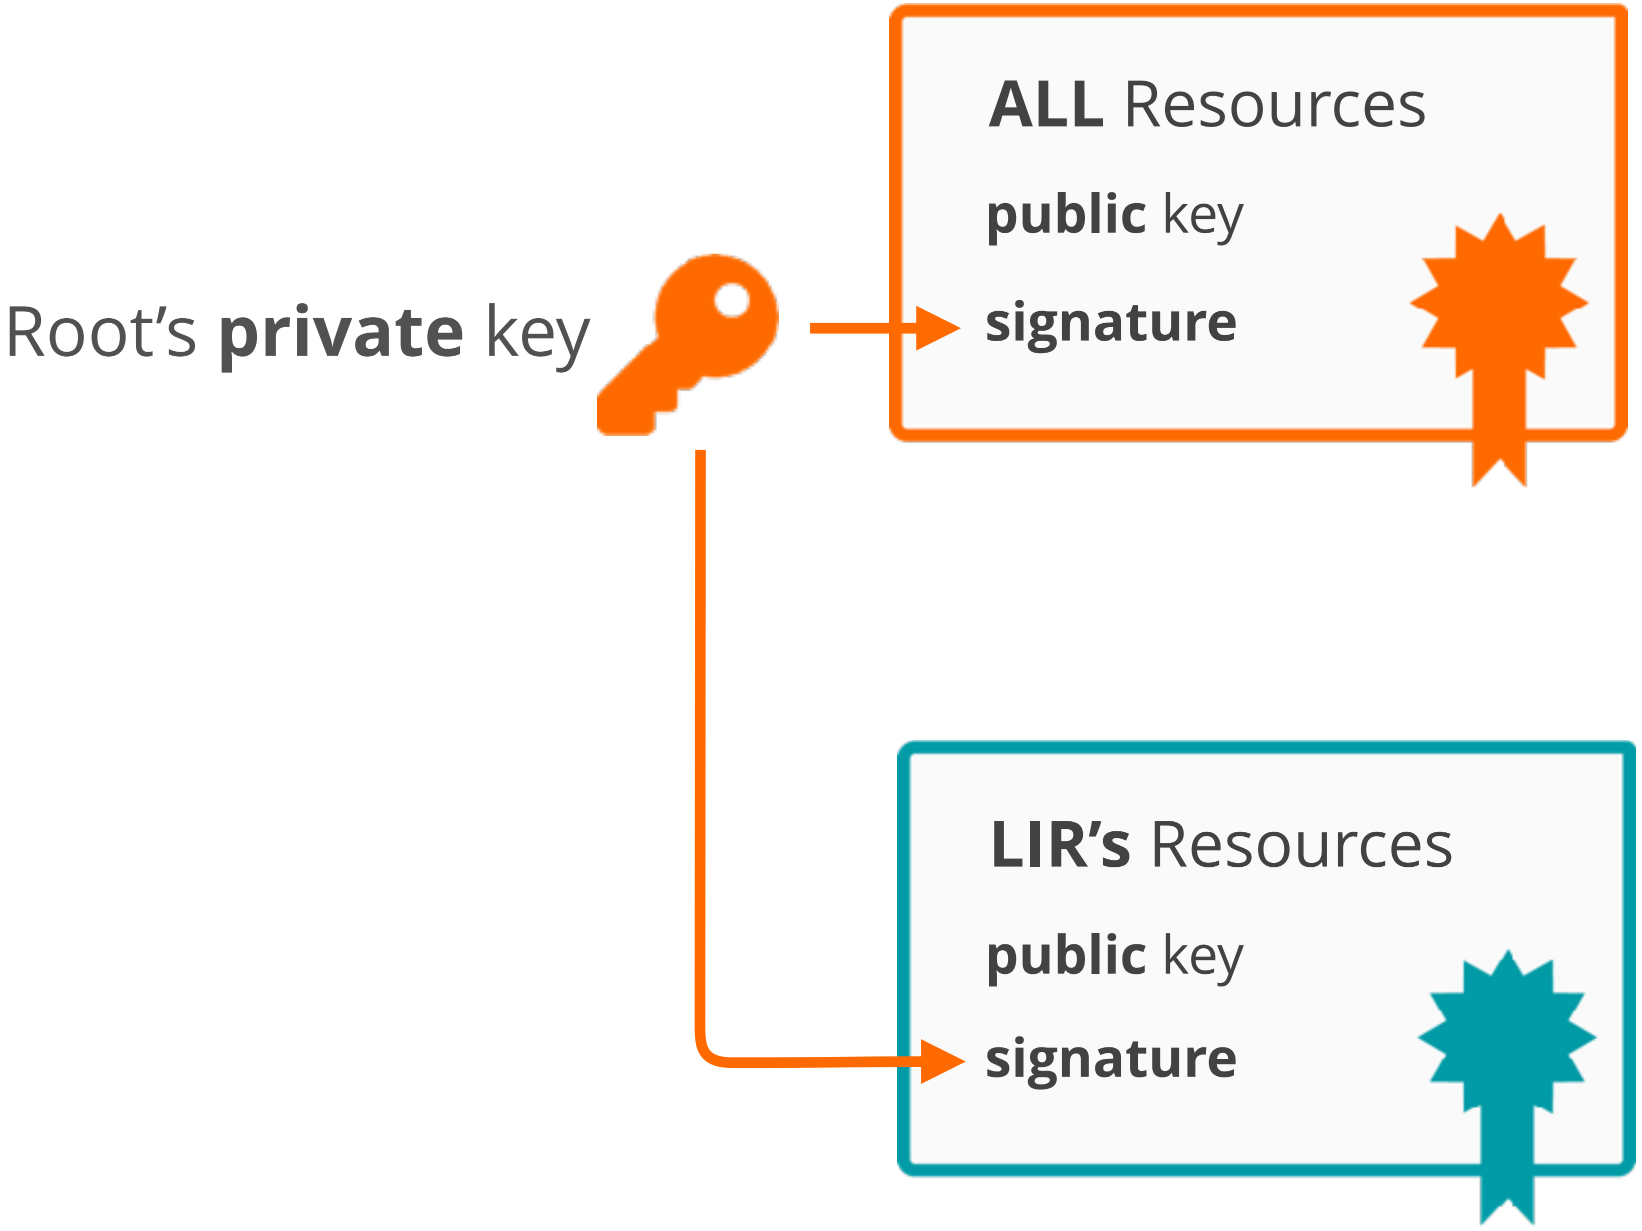

Root's **private** key

ALL Resources

public key

signature

LIR's Resources

public key

signature

ROAs

Section 2

Elements of RPKI

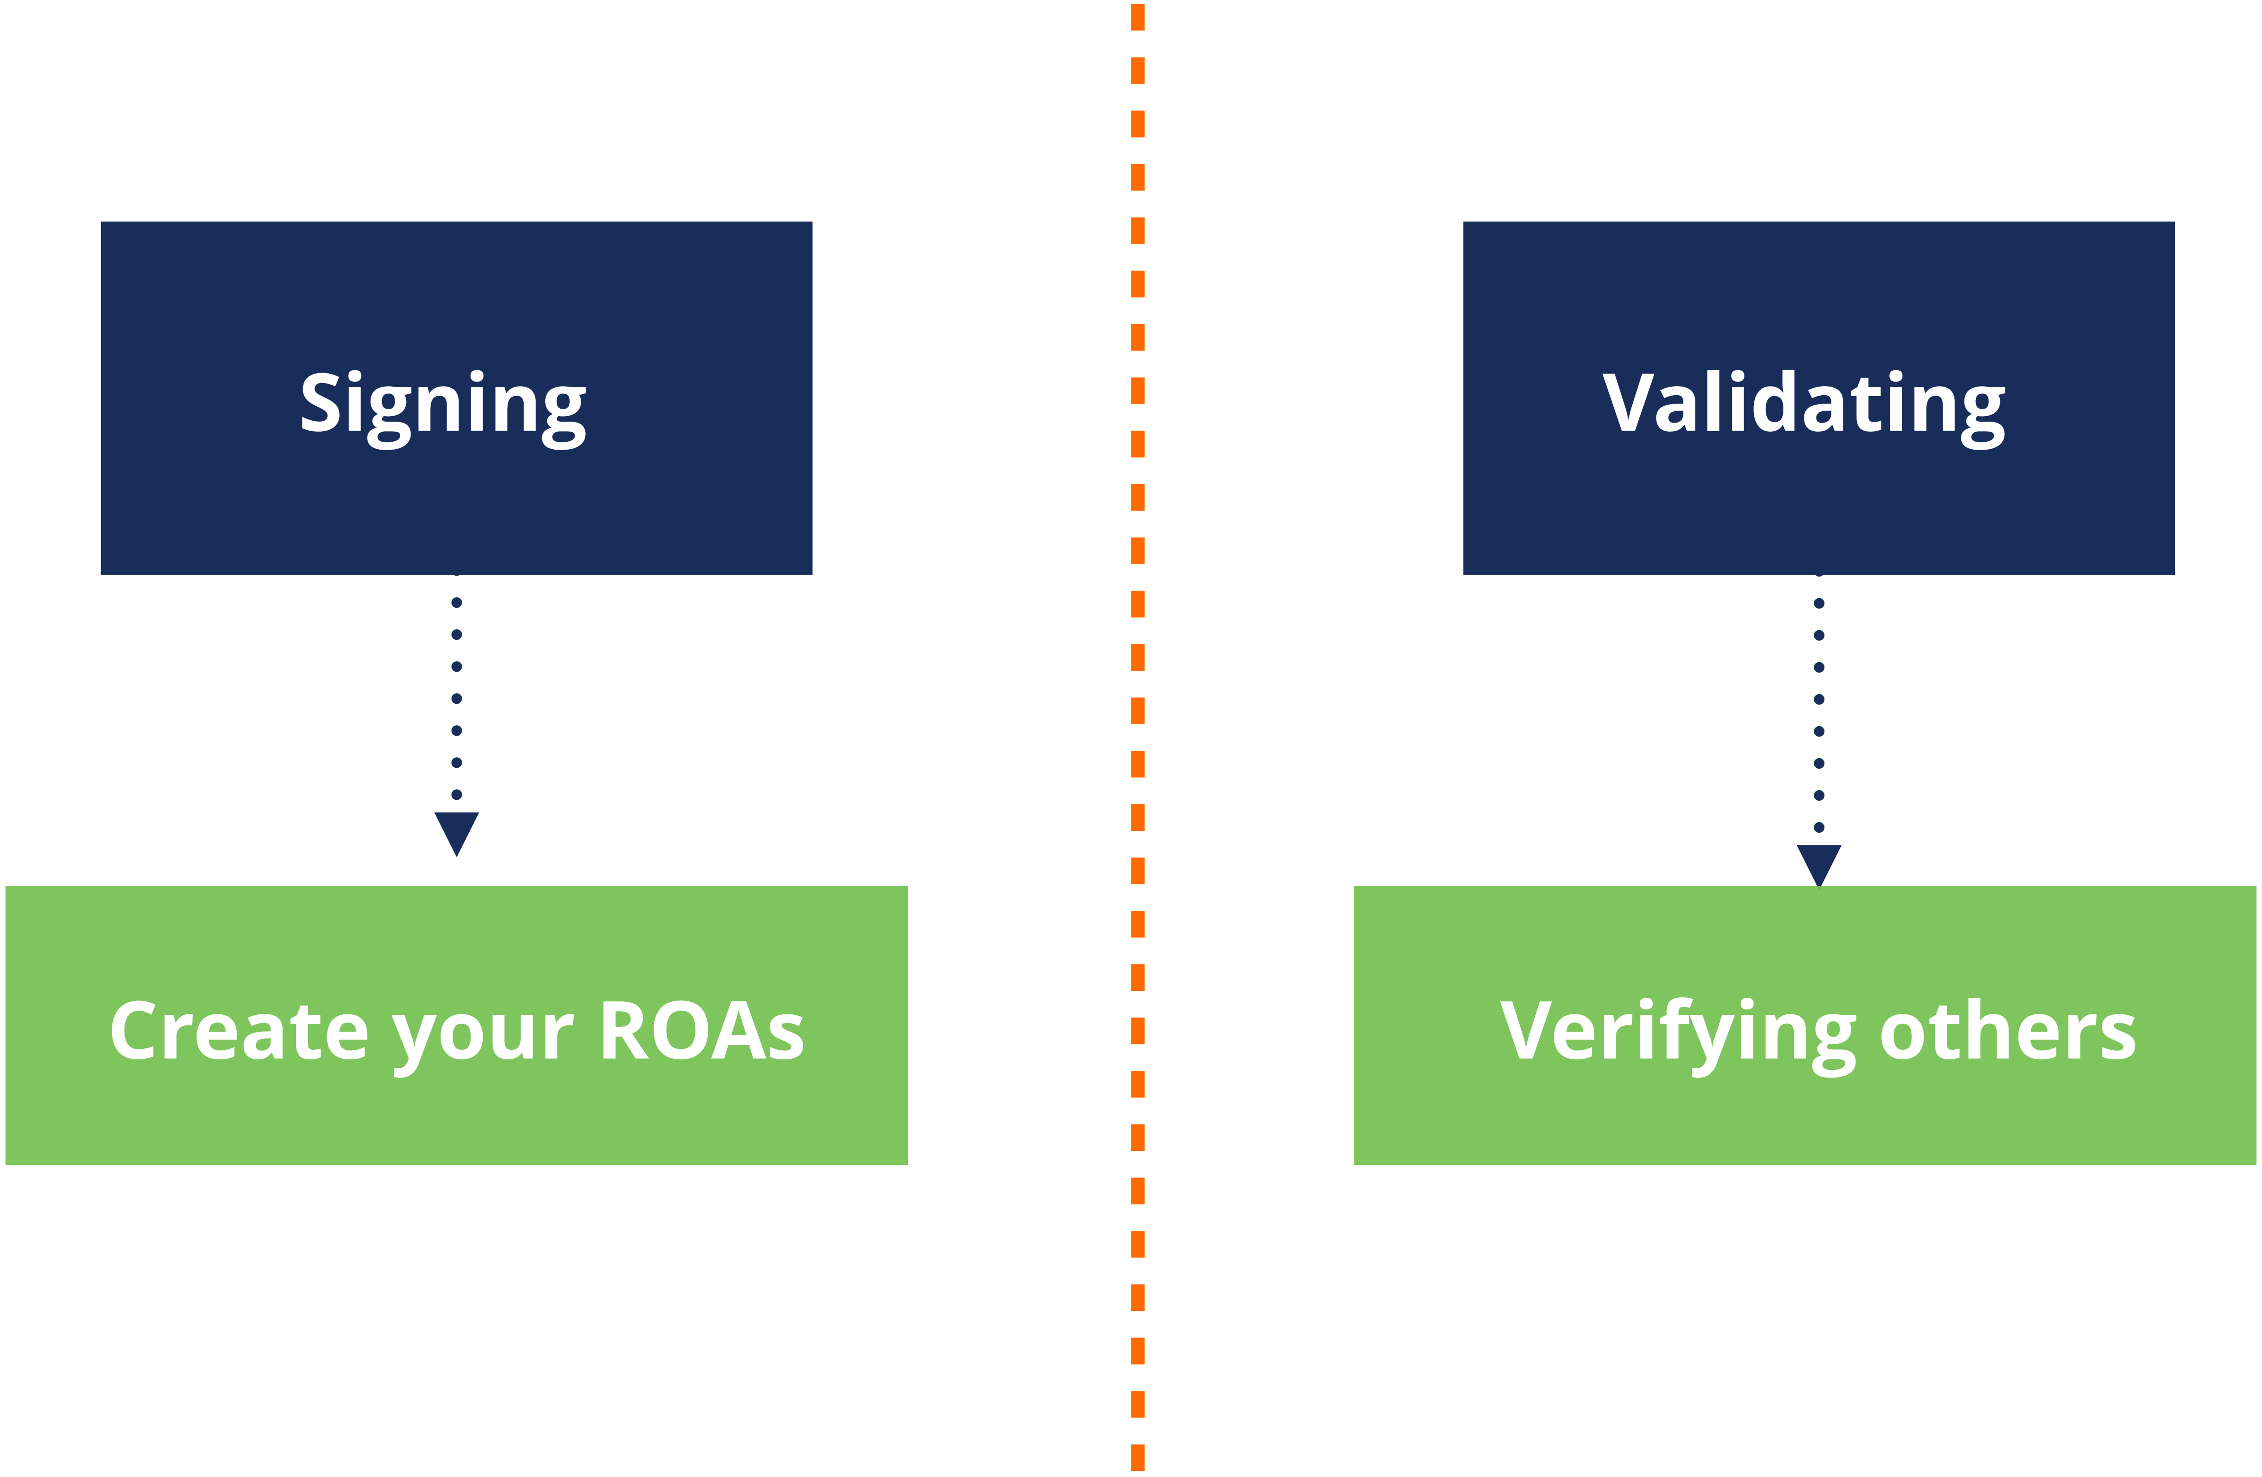

Signing

Create your ROAs

Validating

Verifying others

Elements of RPKI

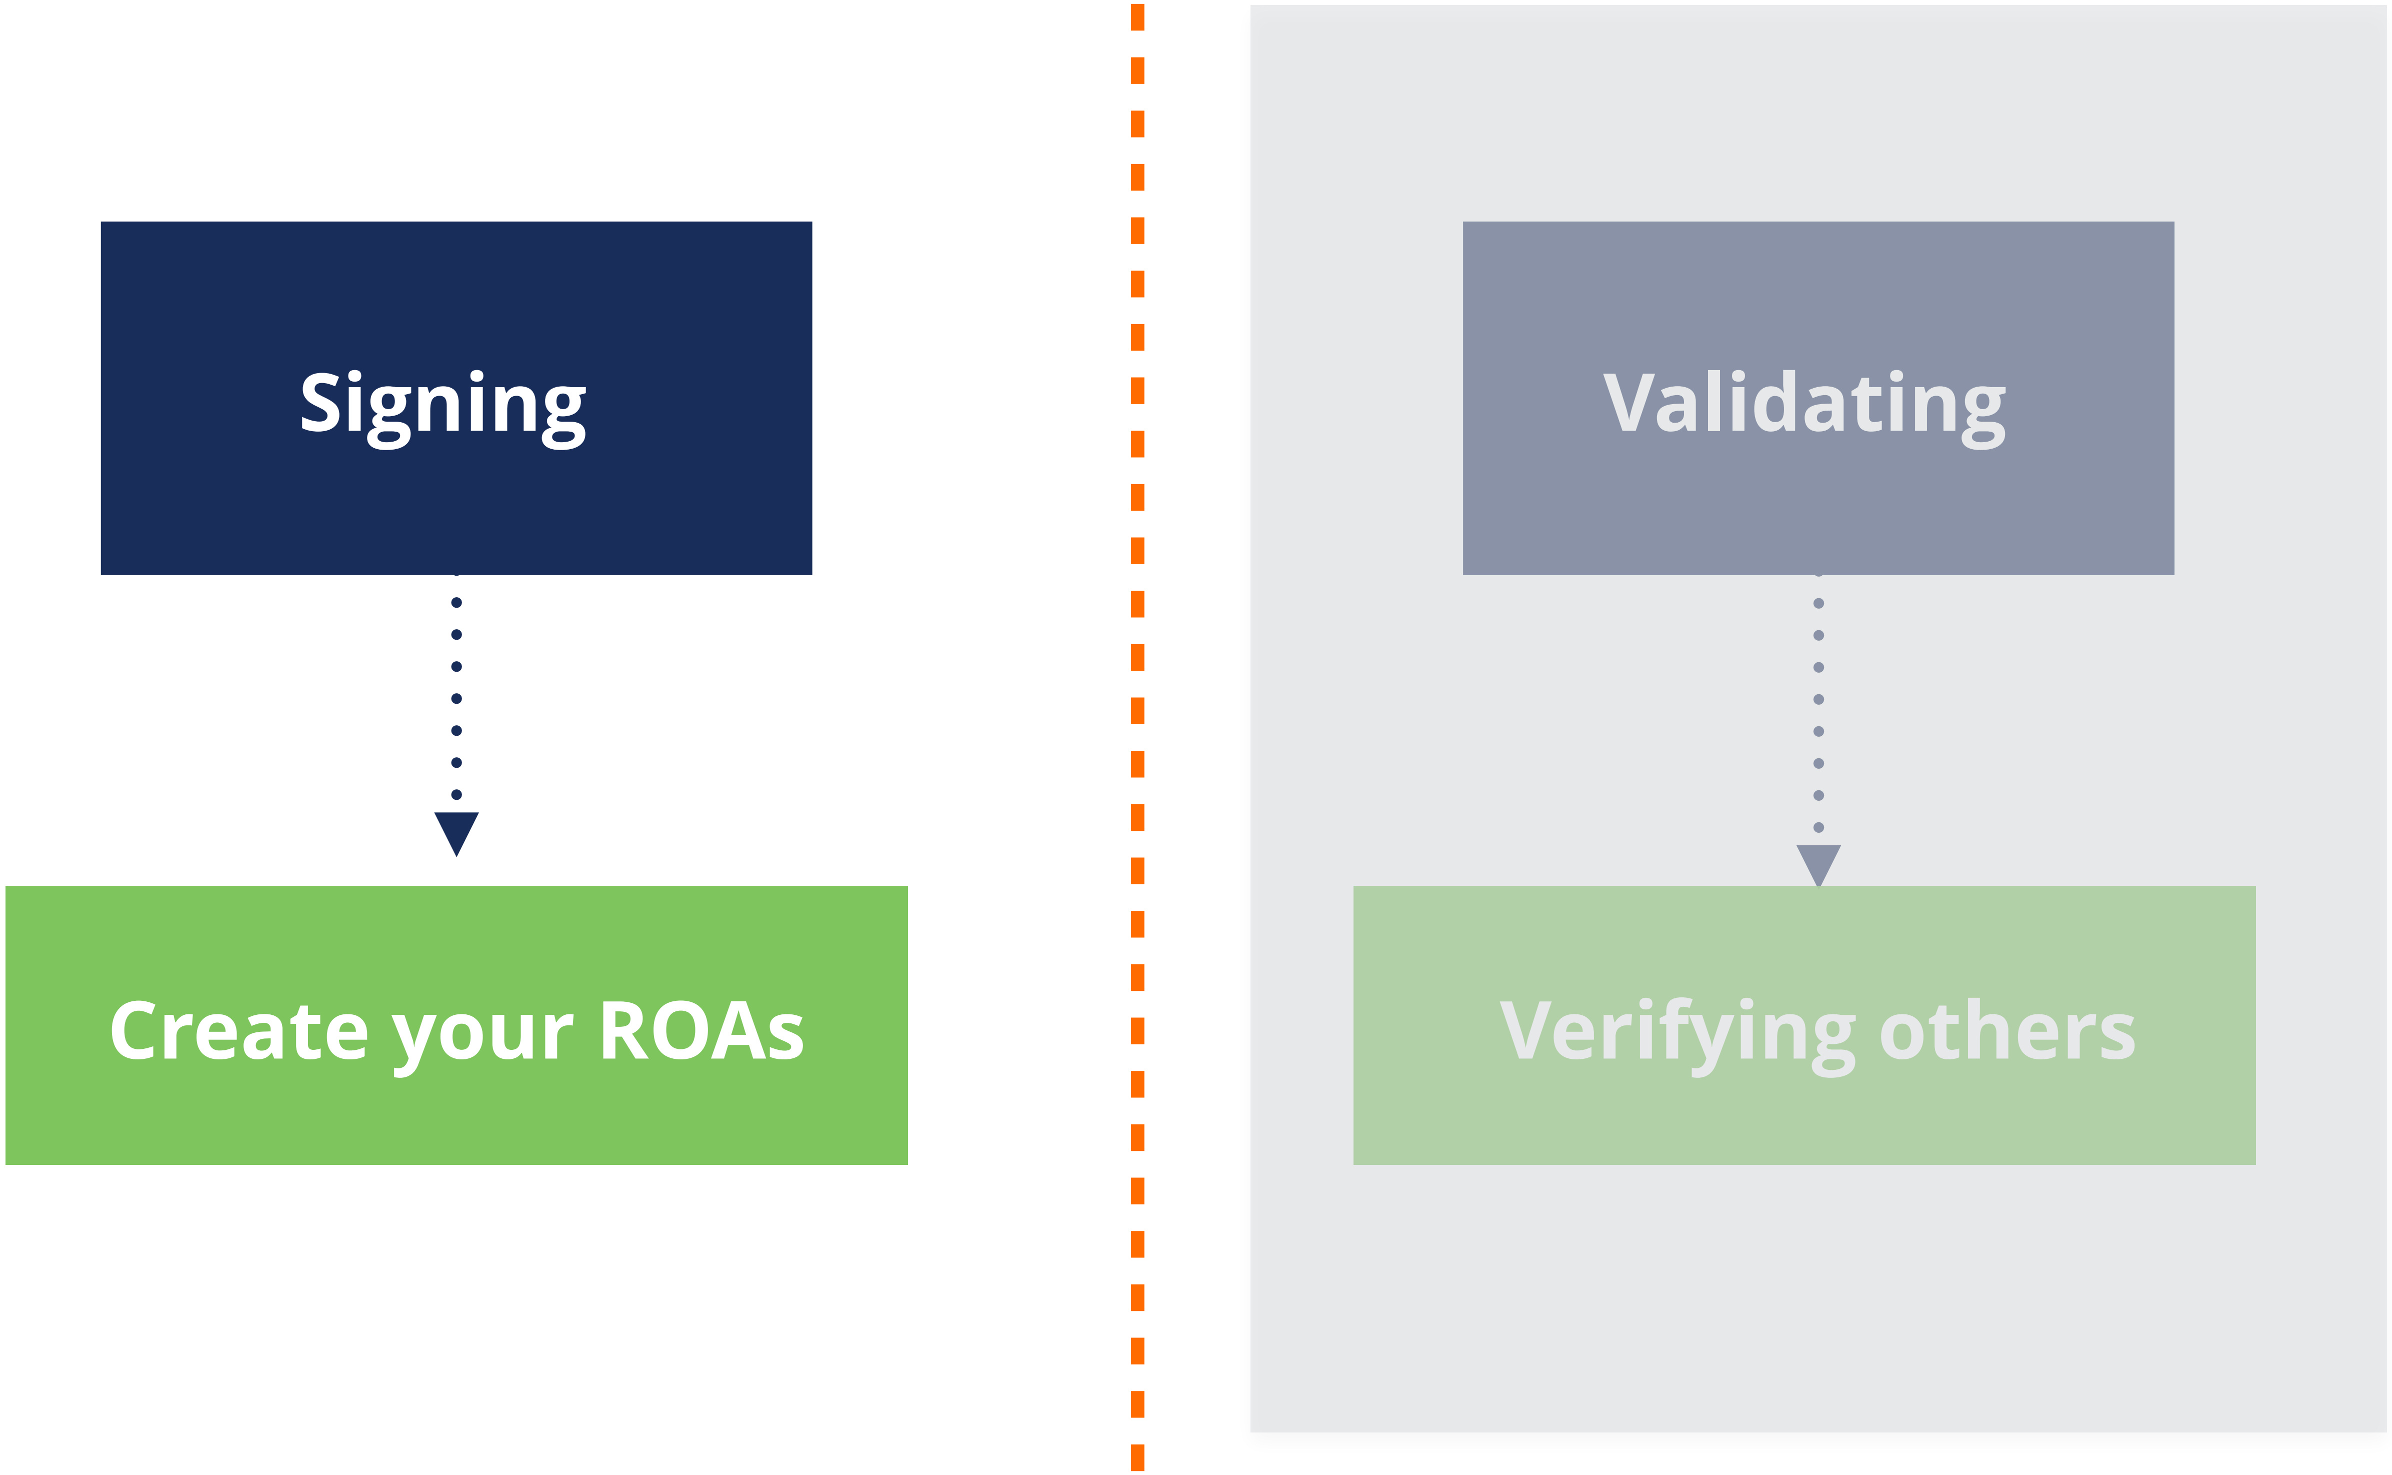

Signing

Create your ROAs

Validating

Verifying others

Question

Have you created **RPKI ROAs** for your prefixes?

⏱ 1 min.

What is a ROA ?

An **authorised statement** from a resource holder

ROA

Prefix

Origin

Prefix
is authorised to be announced
AS Number

- LIRs can create a ROA for their resources
- Multiple ROAs can exist for the same prefix
- ROAs can overlap

What is in a ROA ?

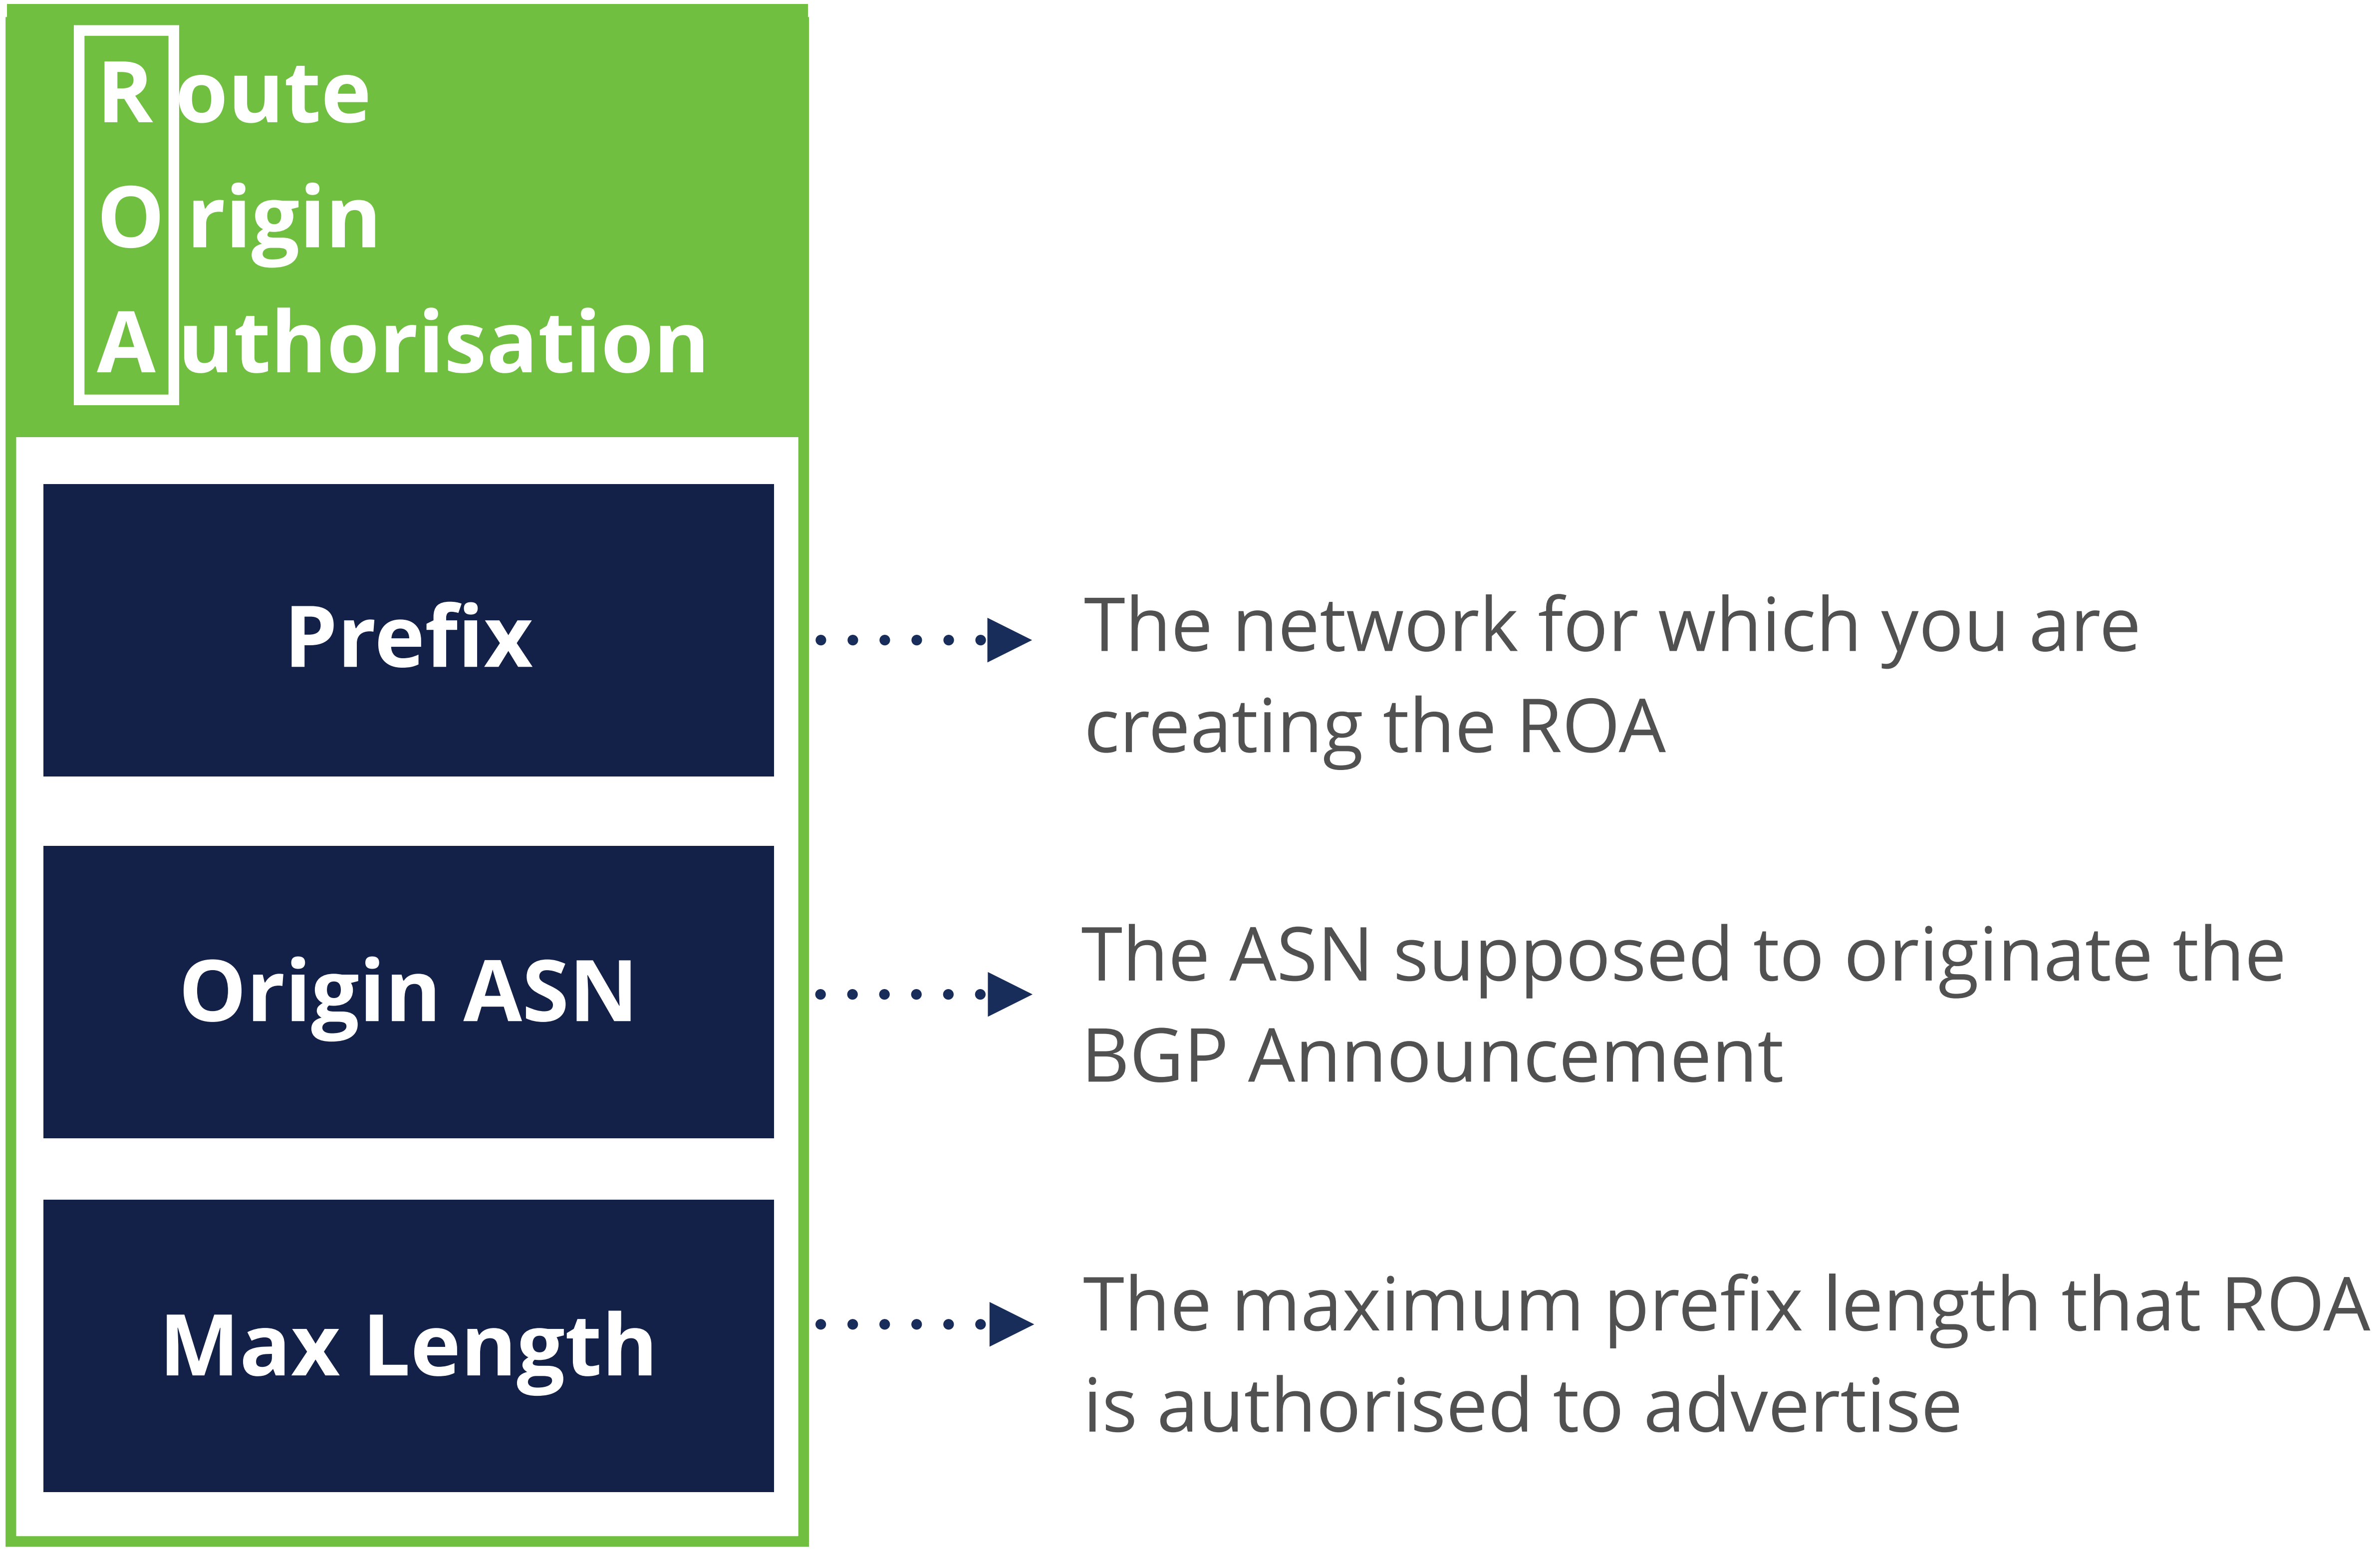

Route
Origin
Authorisation

Prefix → The network for which you are creating the ROA

Origin ASN → The ASN supposed to originate the BGP Announcement

Max Length → The maximum prefix length that ROA is authorised to advertise

What is max-length?

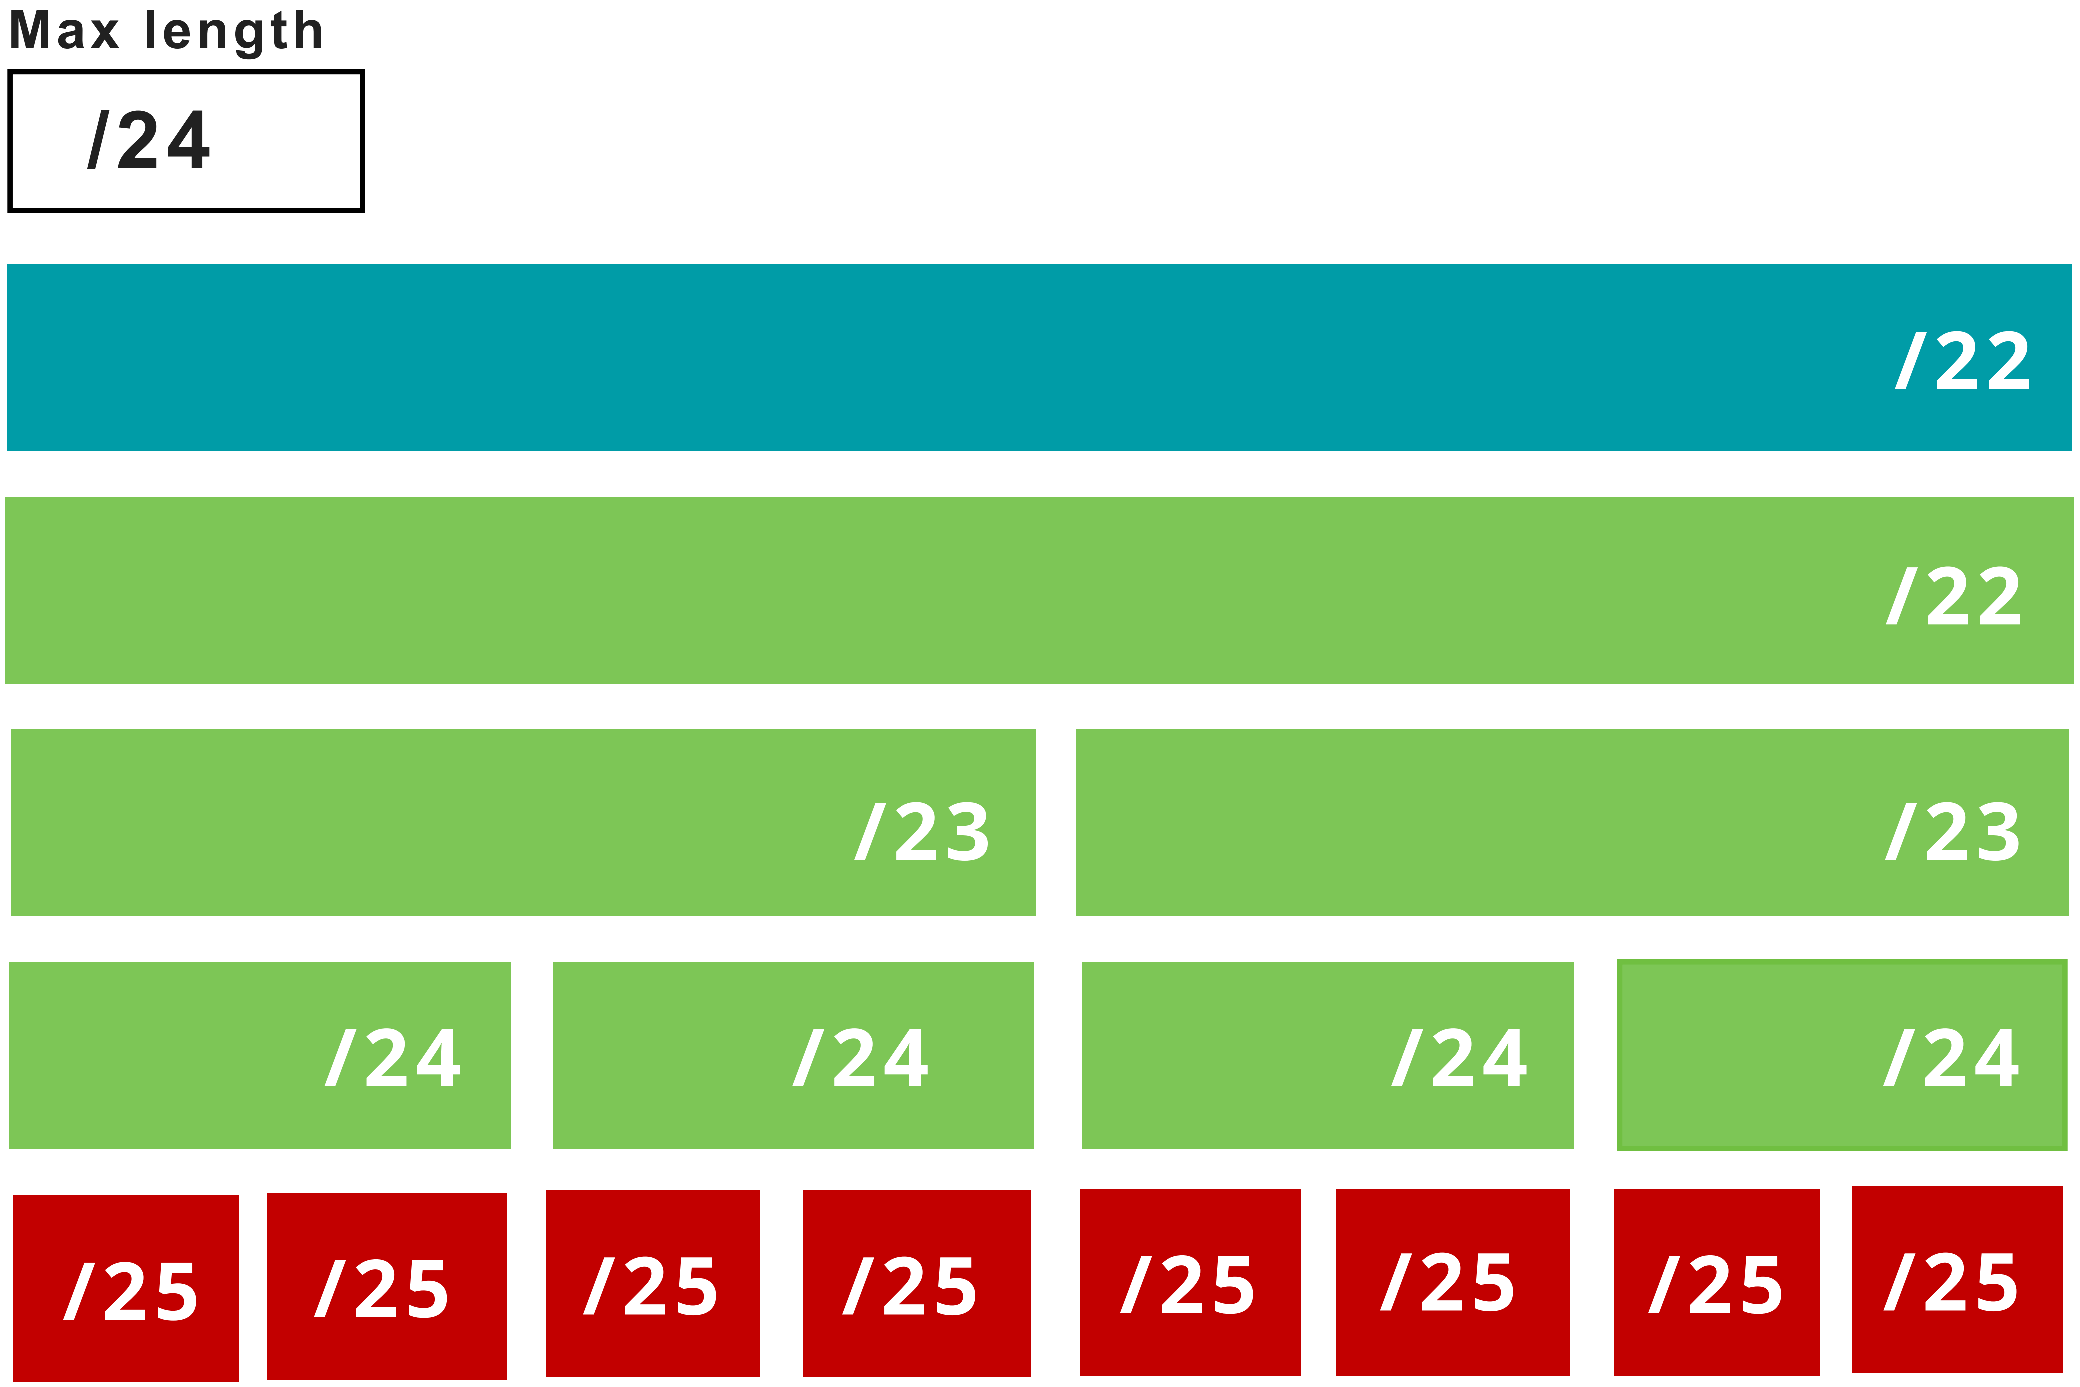

Max length

/24

/22

/22

/23

/23

/24

/24

/24

/24

/25

/25

/25

/25

/25

/25

/25

/25

How should we use max-length?

You created a single ROA authorising the entire /22

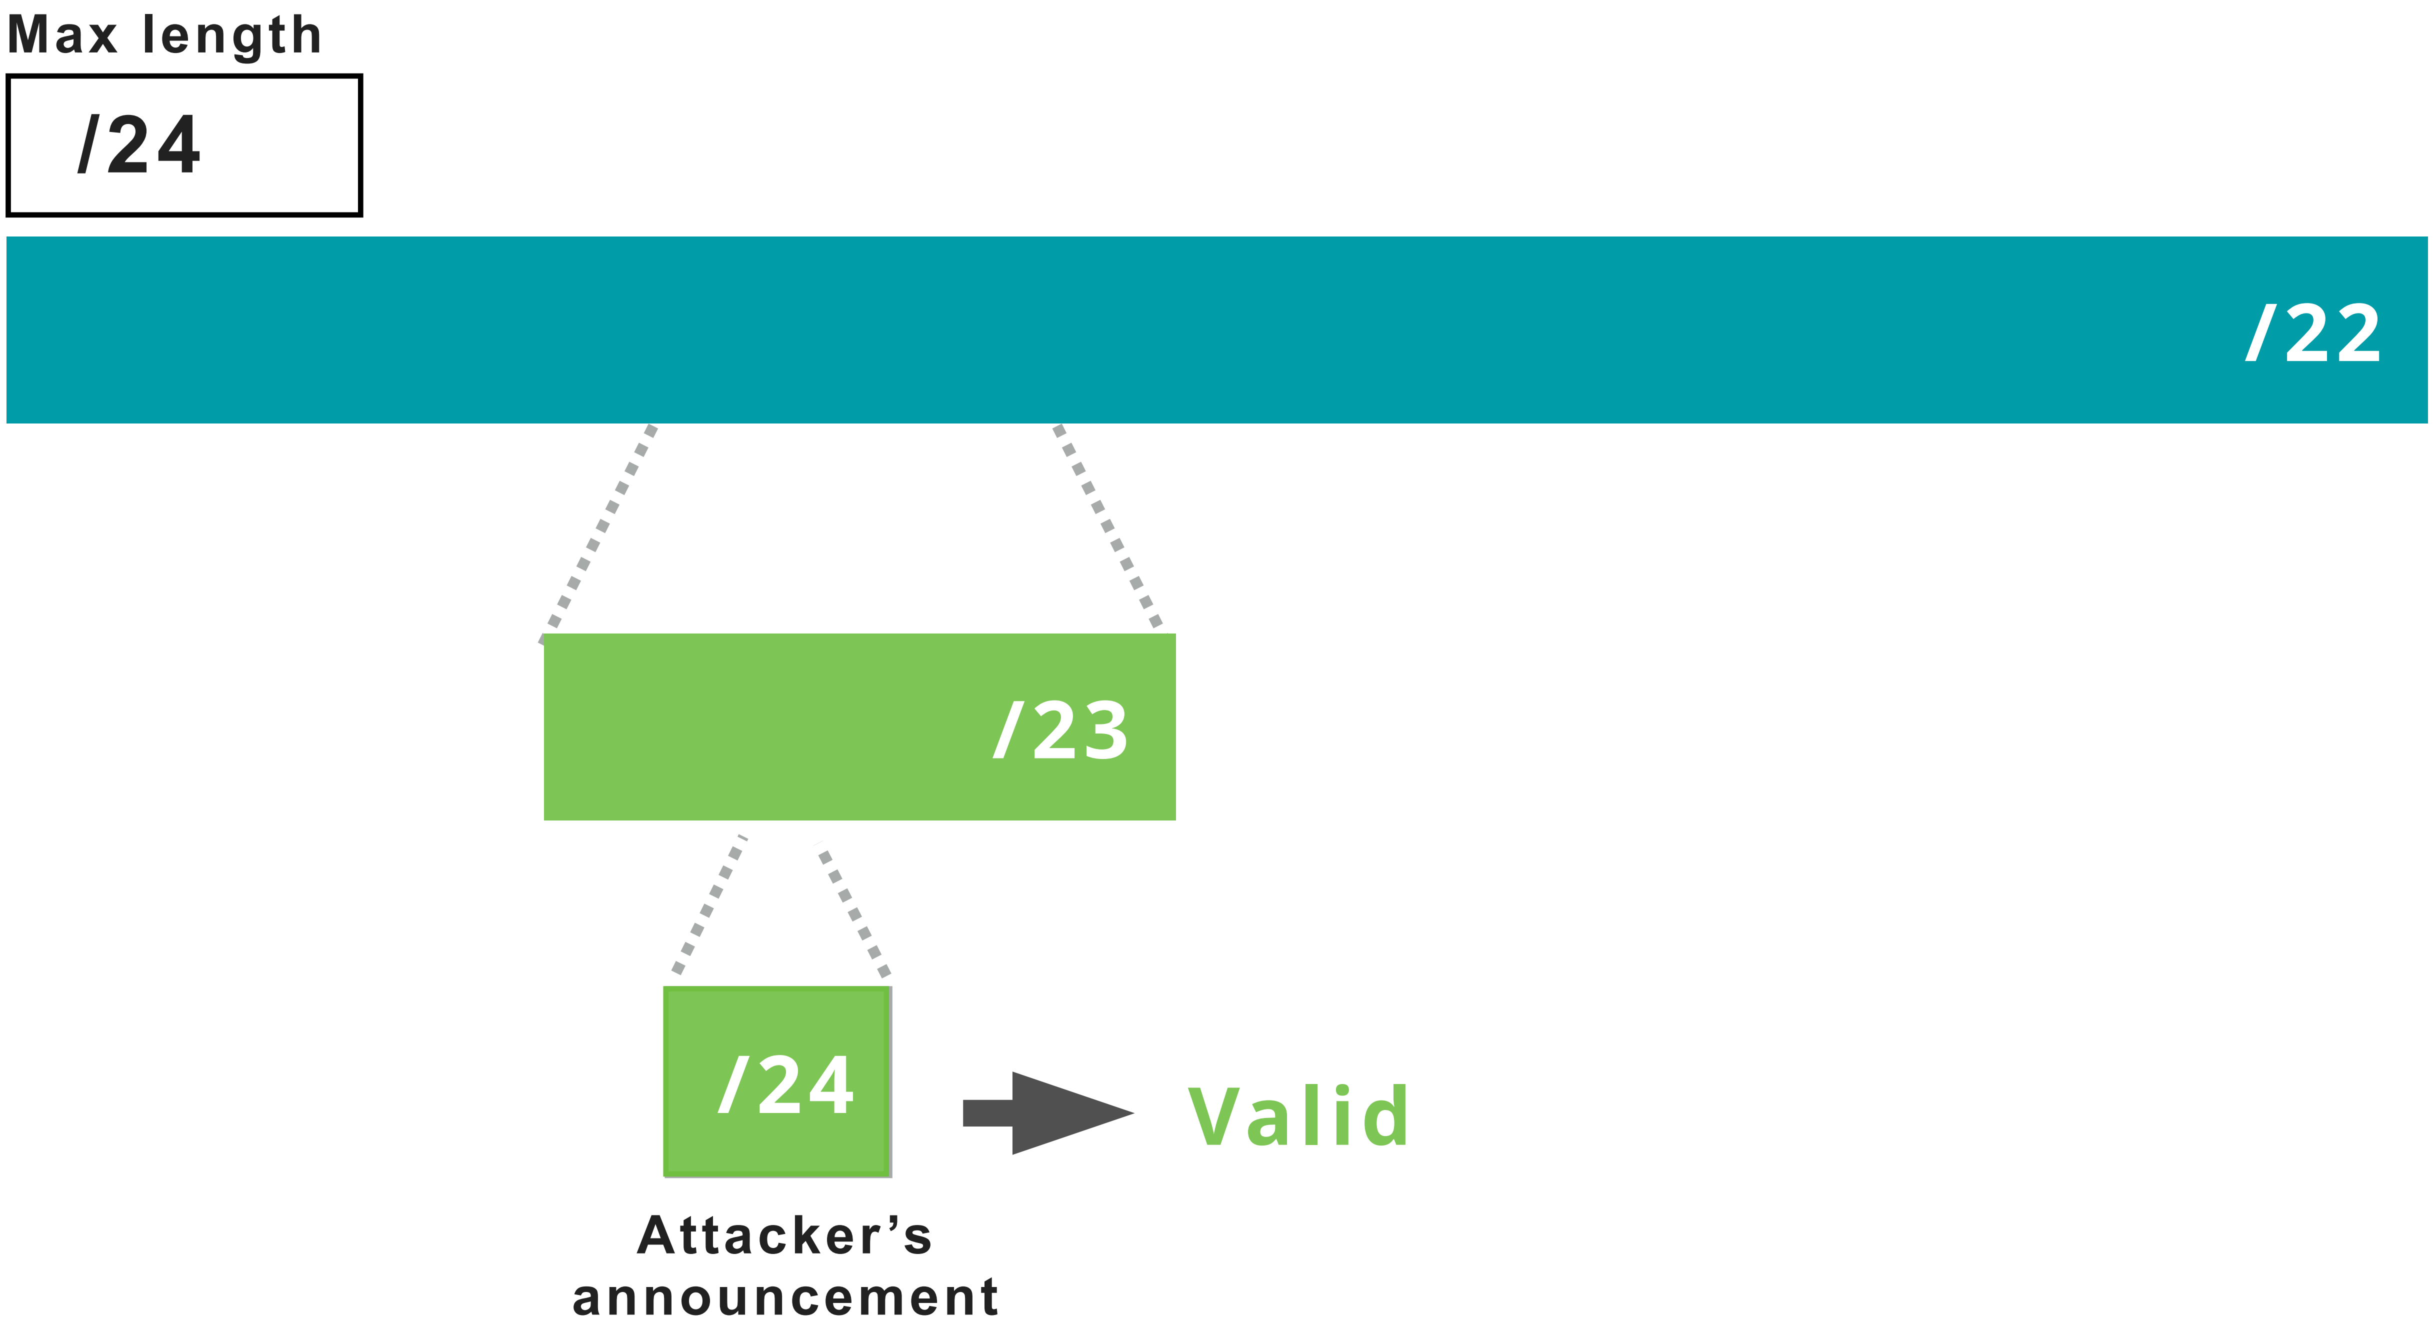

How should we use max-length?

Create ROAs for BGP announcements only

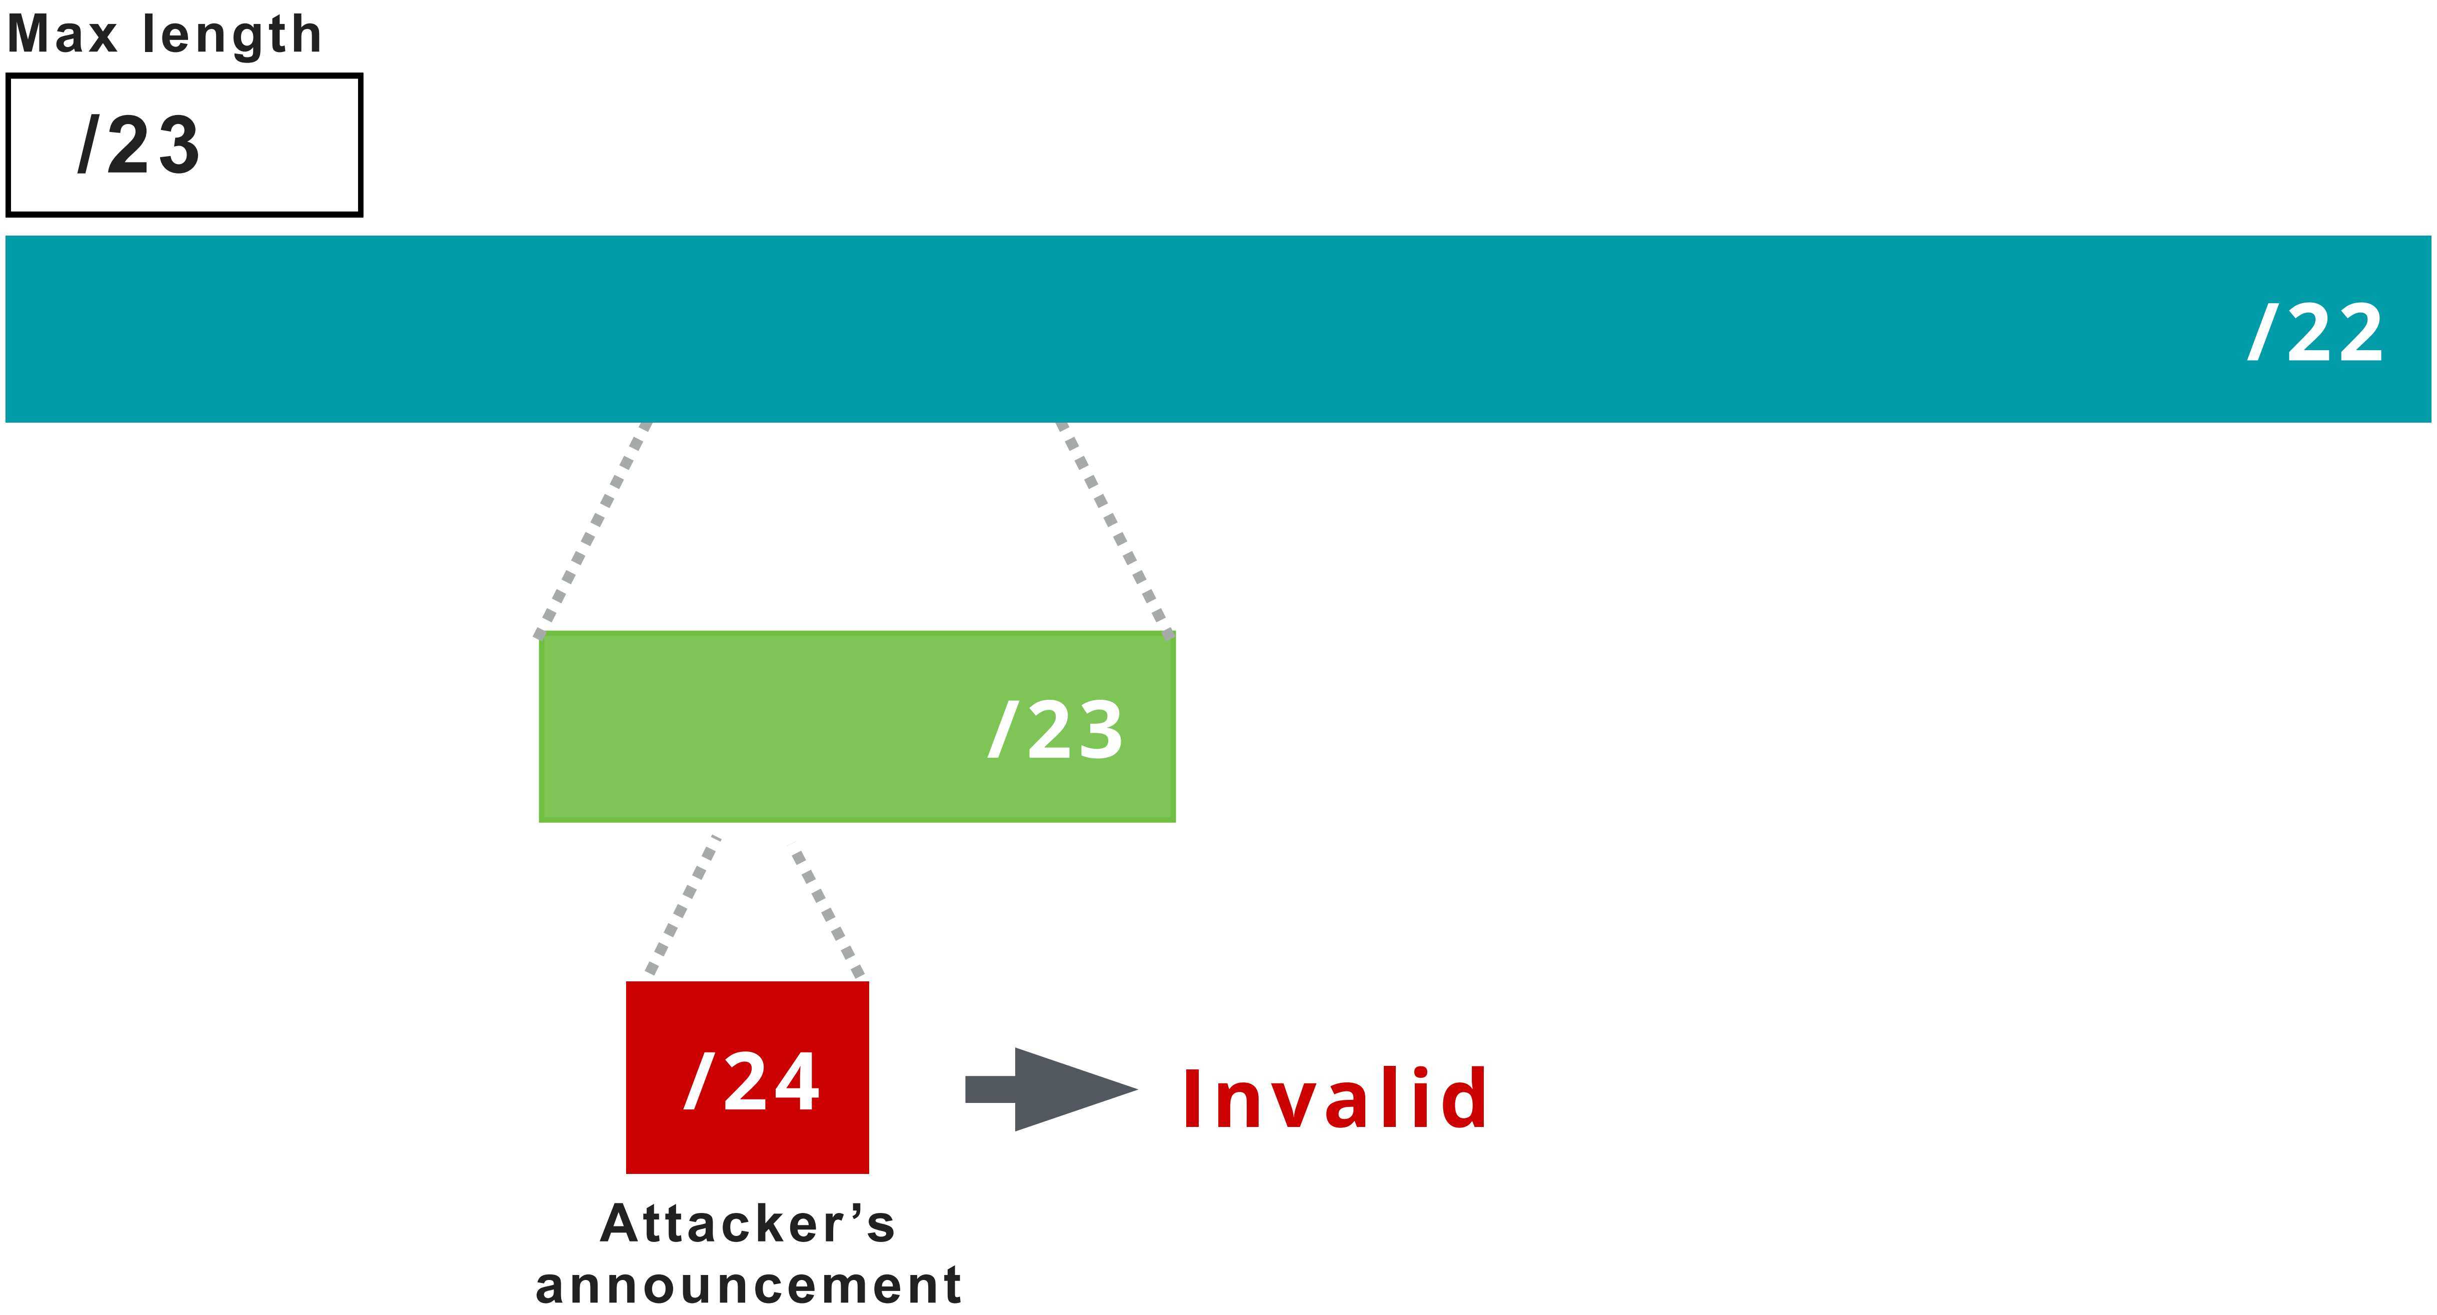

Max length

/23

/22

/23

/24

Invalid

Attacker’s announcement

Quiz time!

Which information is correct about **max-length?**

A. It is an optional field

B. It is a mandatory field, you cannot leave it empty

C. It is the maximum prefix-length a ROA is authorized to advertise

D. It is the maximum prefix length you can announce in BGP

1 min.

Quiz time!

According to this ROA, which announcements will be considered as **valid** and **accepted** by the router?

A. 193.0.24.0/22

B. 193.0.24.0/23

C. 193.0.26.0/24

D. 193.0.24.0/24

E. 193.0.25.0/24

ROA

Prefix: 193.0.24.0/23
Origin: AS65530
Max-length: /24

1 min.

ROA Signature

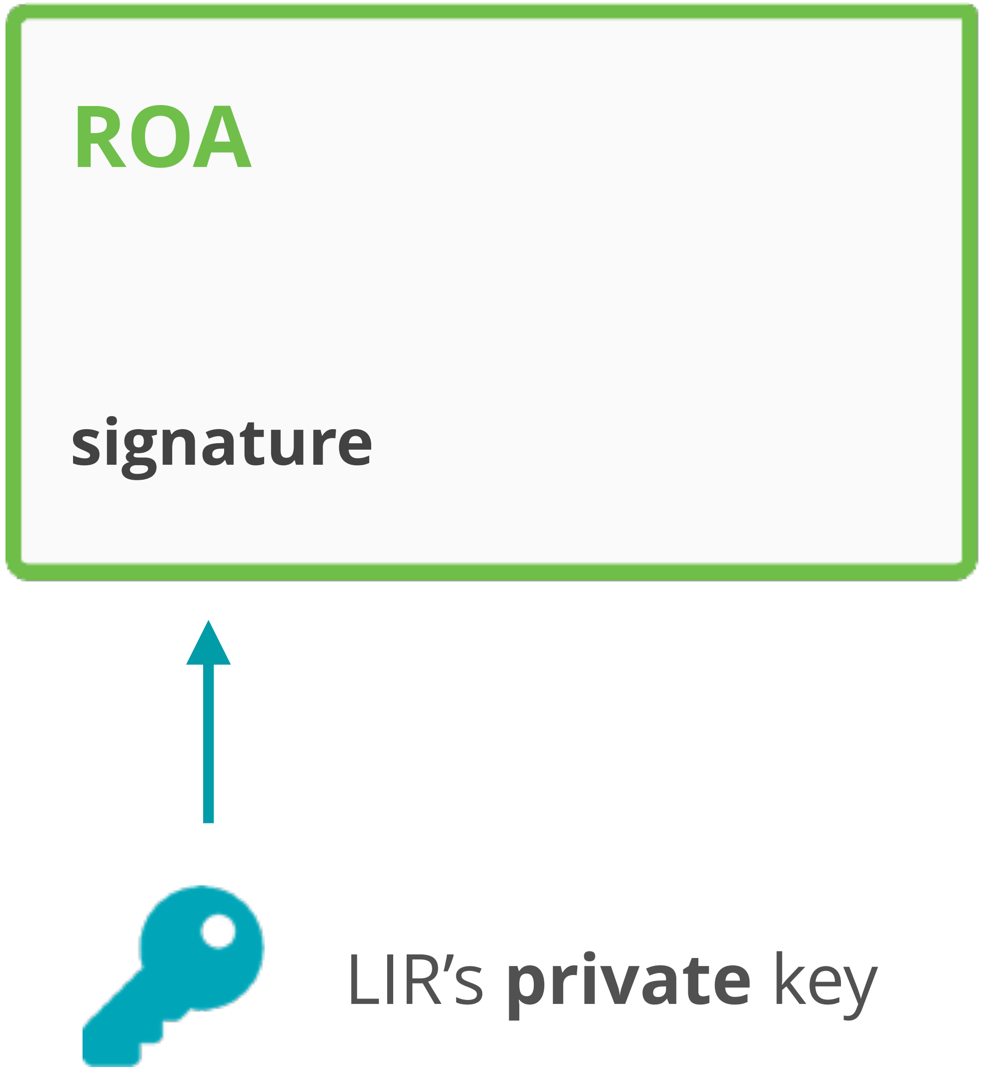

ROA

signature

LIR’s **private** key

Prefix
is authorised to be announced by
AS Number

RPKI Chain of Trust

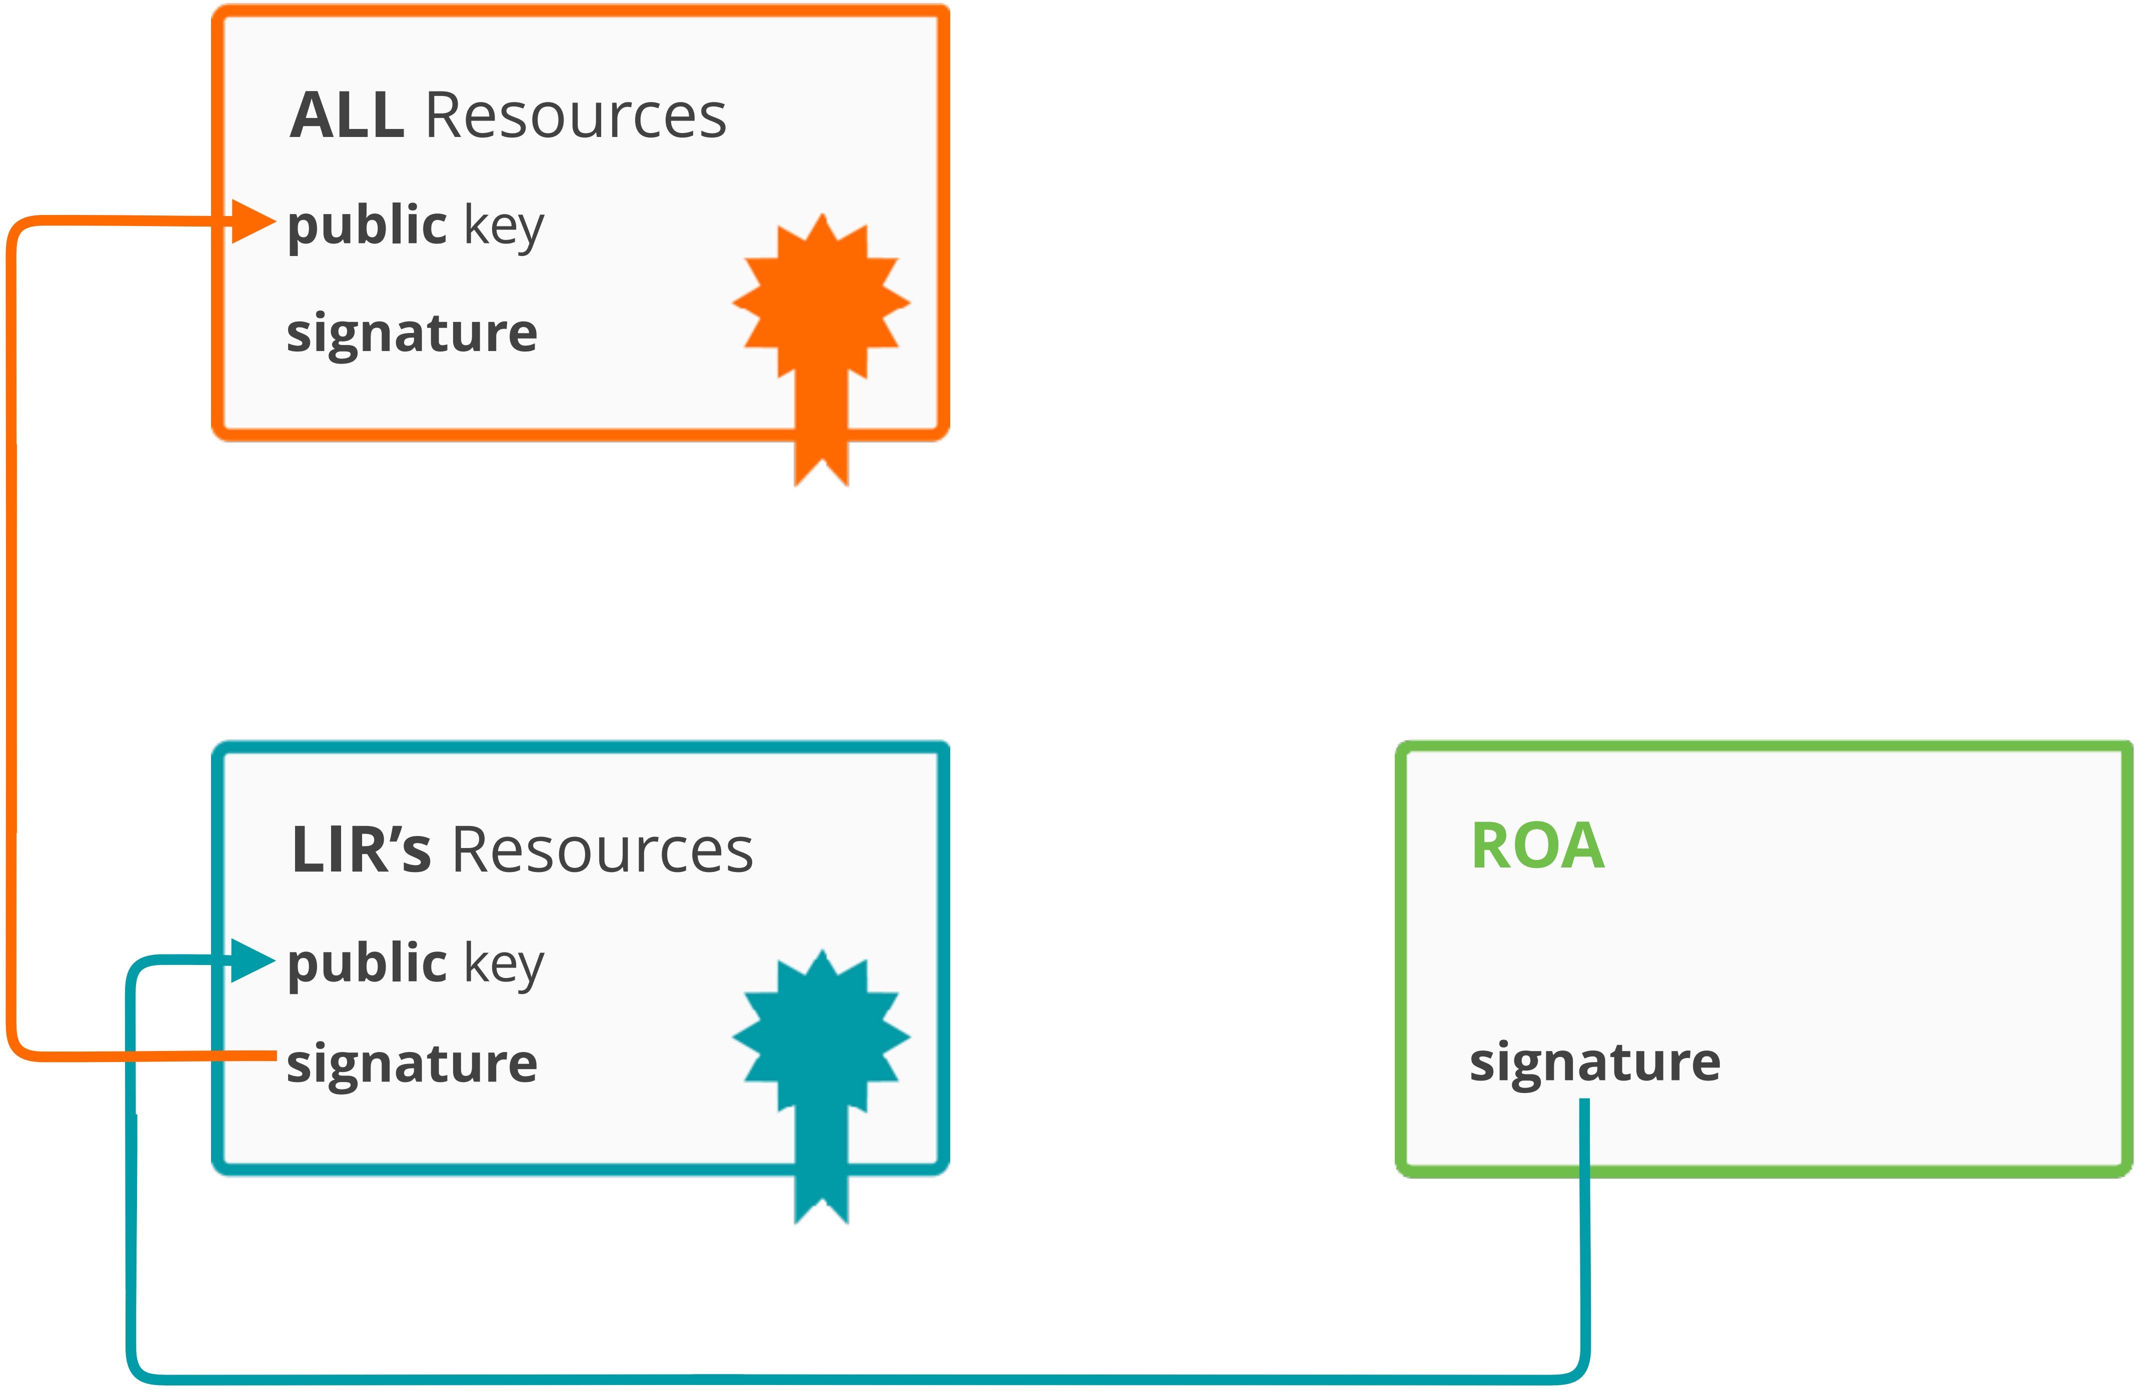

ALL Resources

public key

signature

LIR's Resources

public key

signature

ROA

signature

RPKI Certificate Structure

Certificate hierarchy follows allocation hierarchy

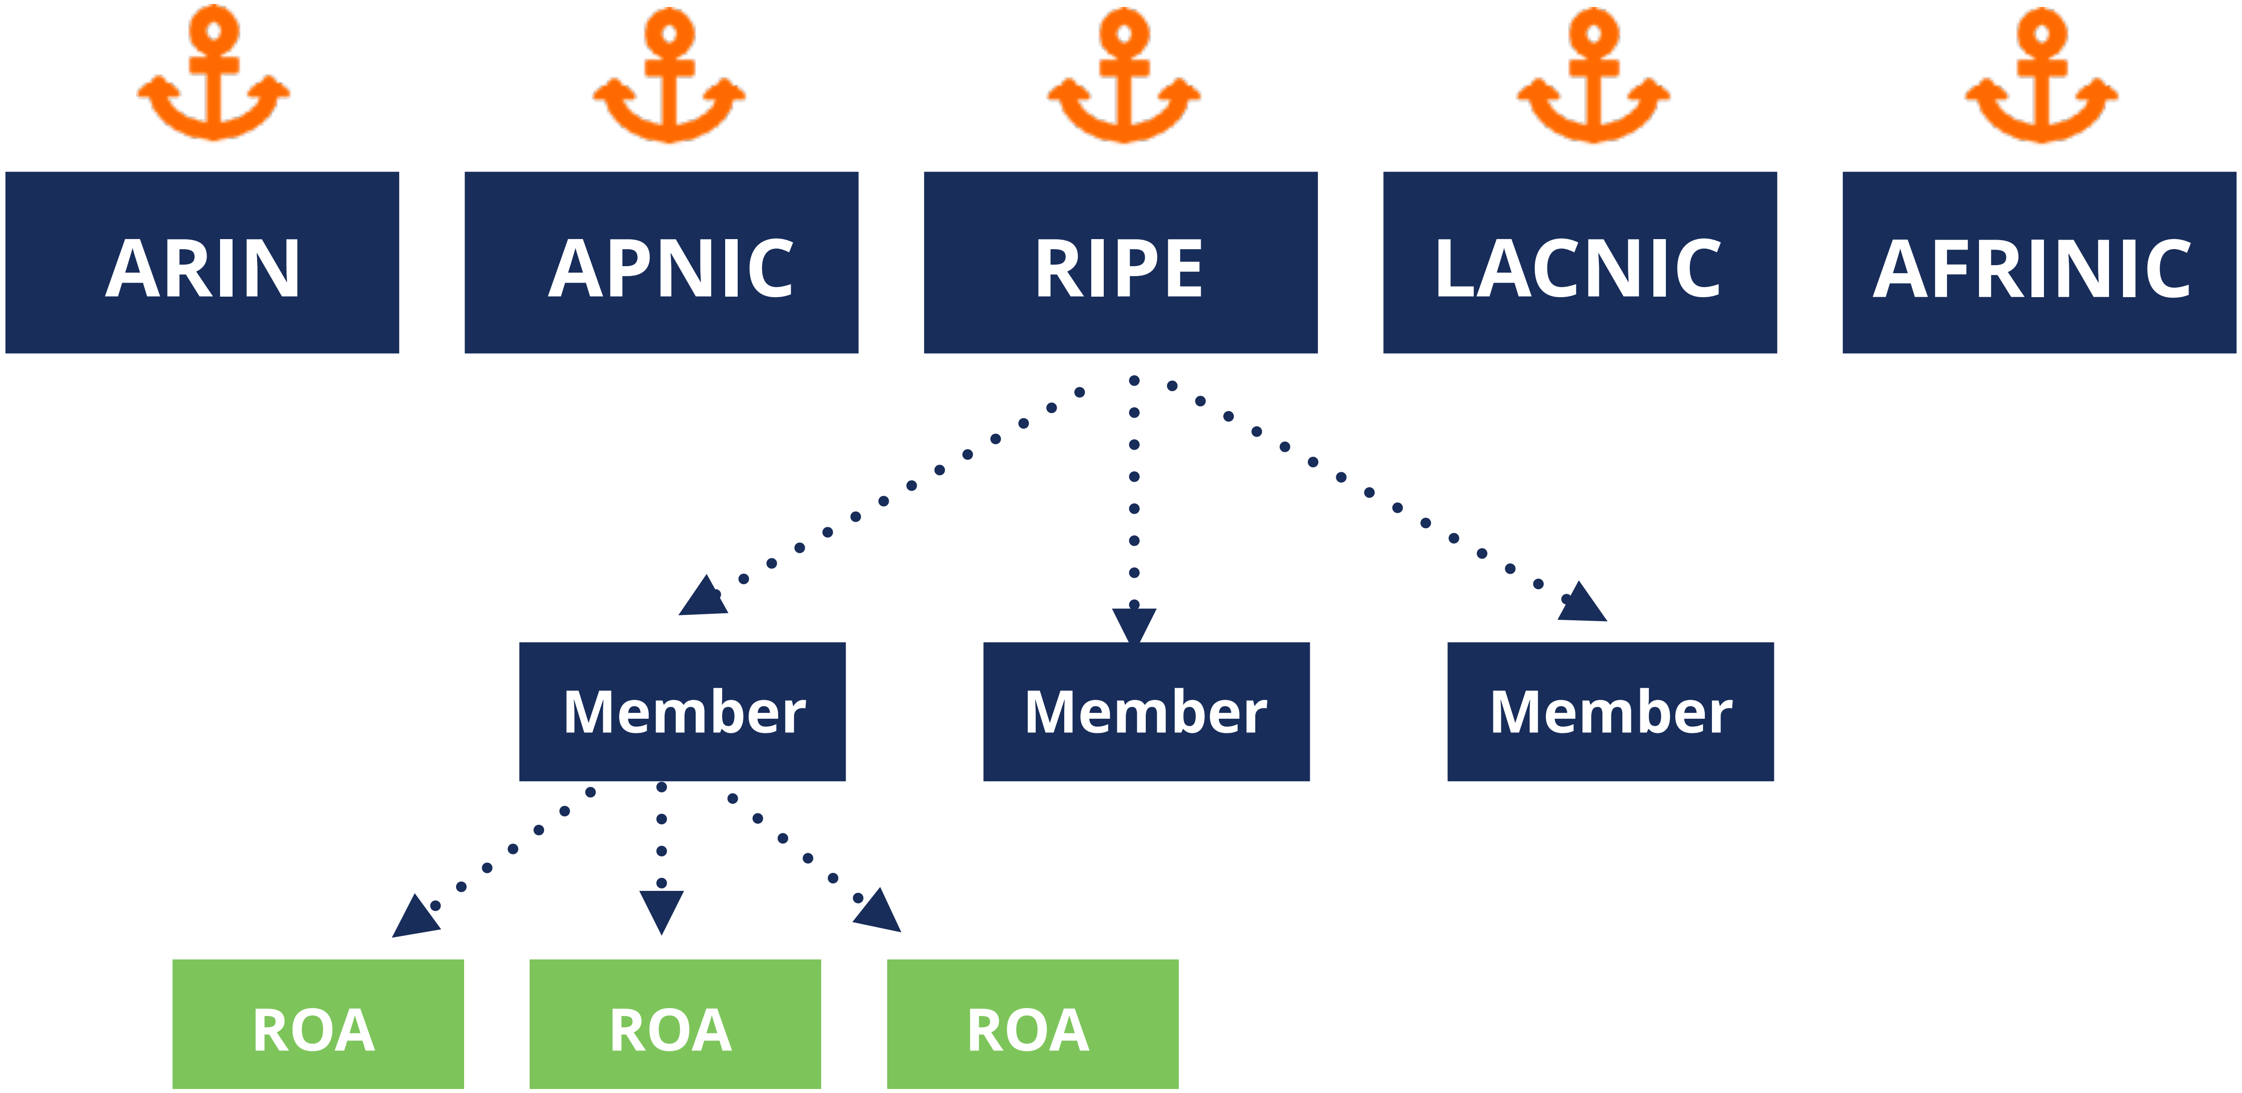

Hosted or Delegated RPKI

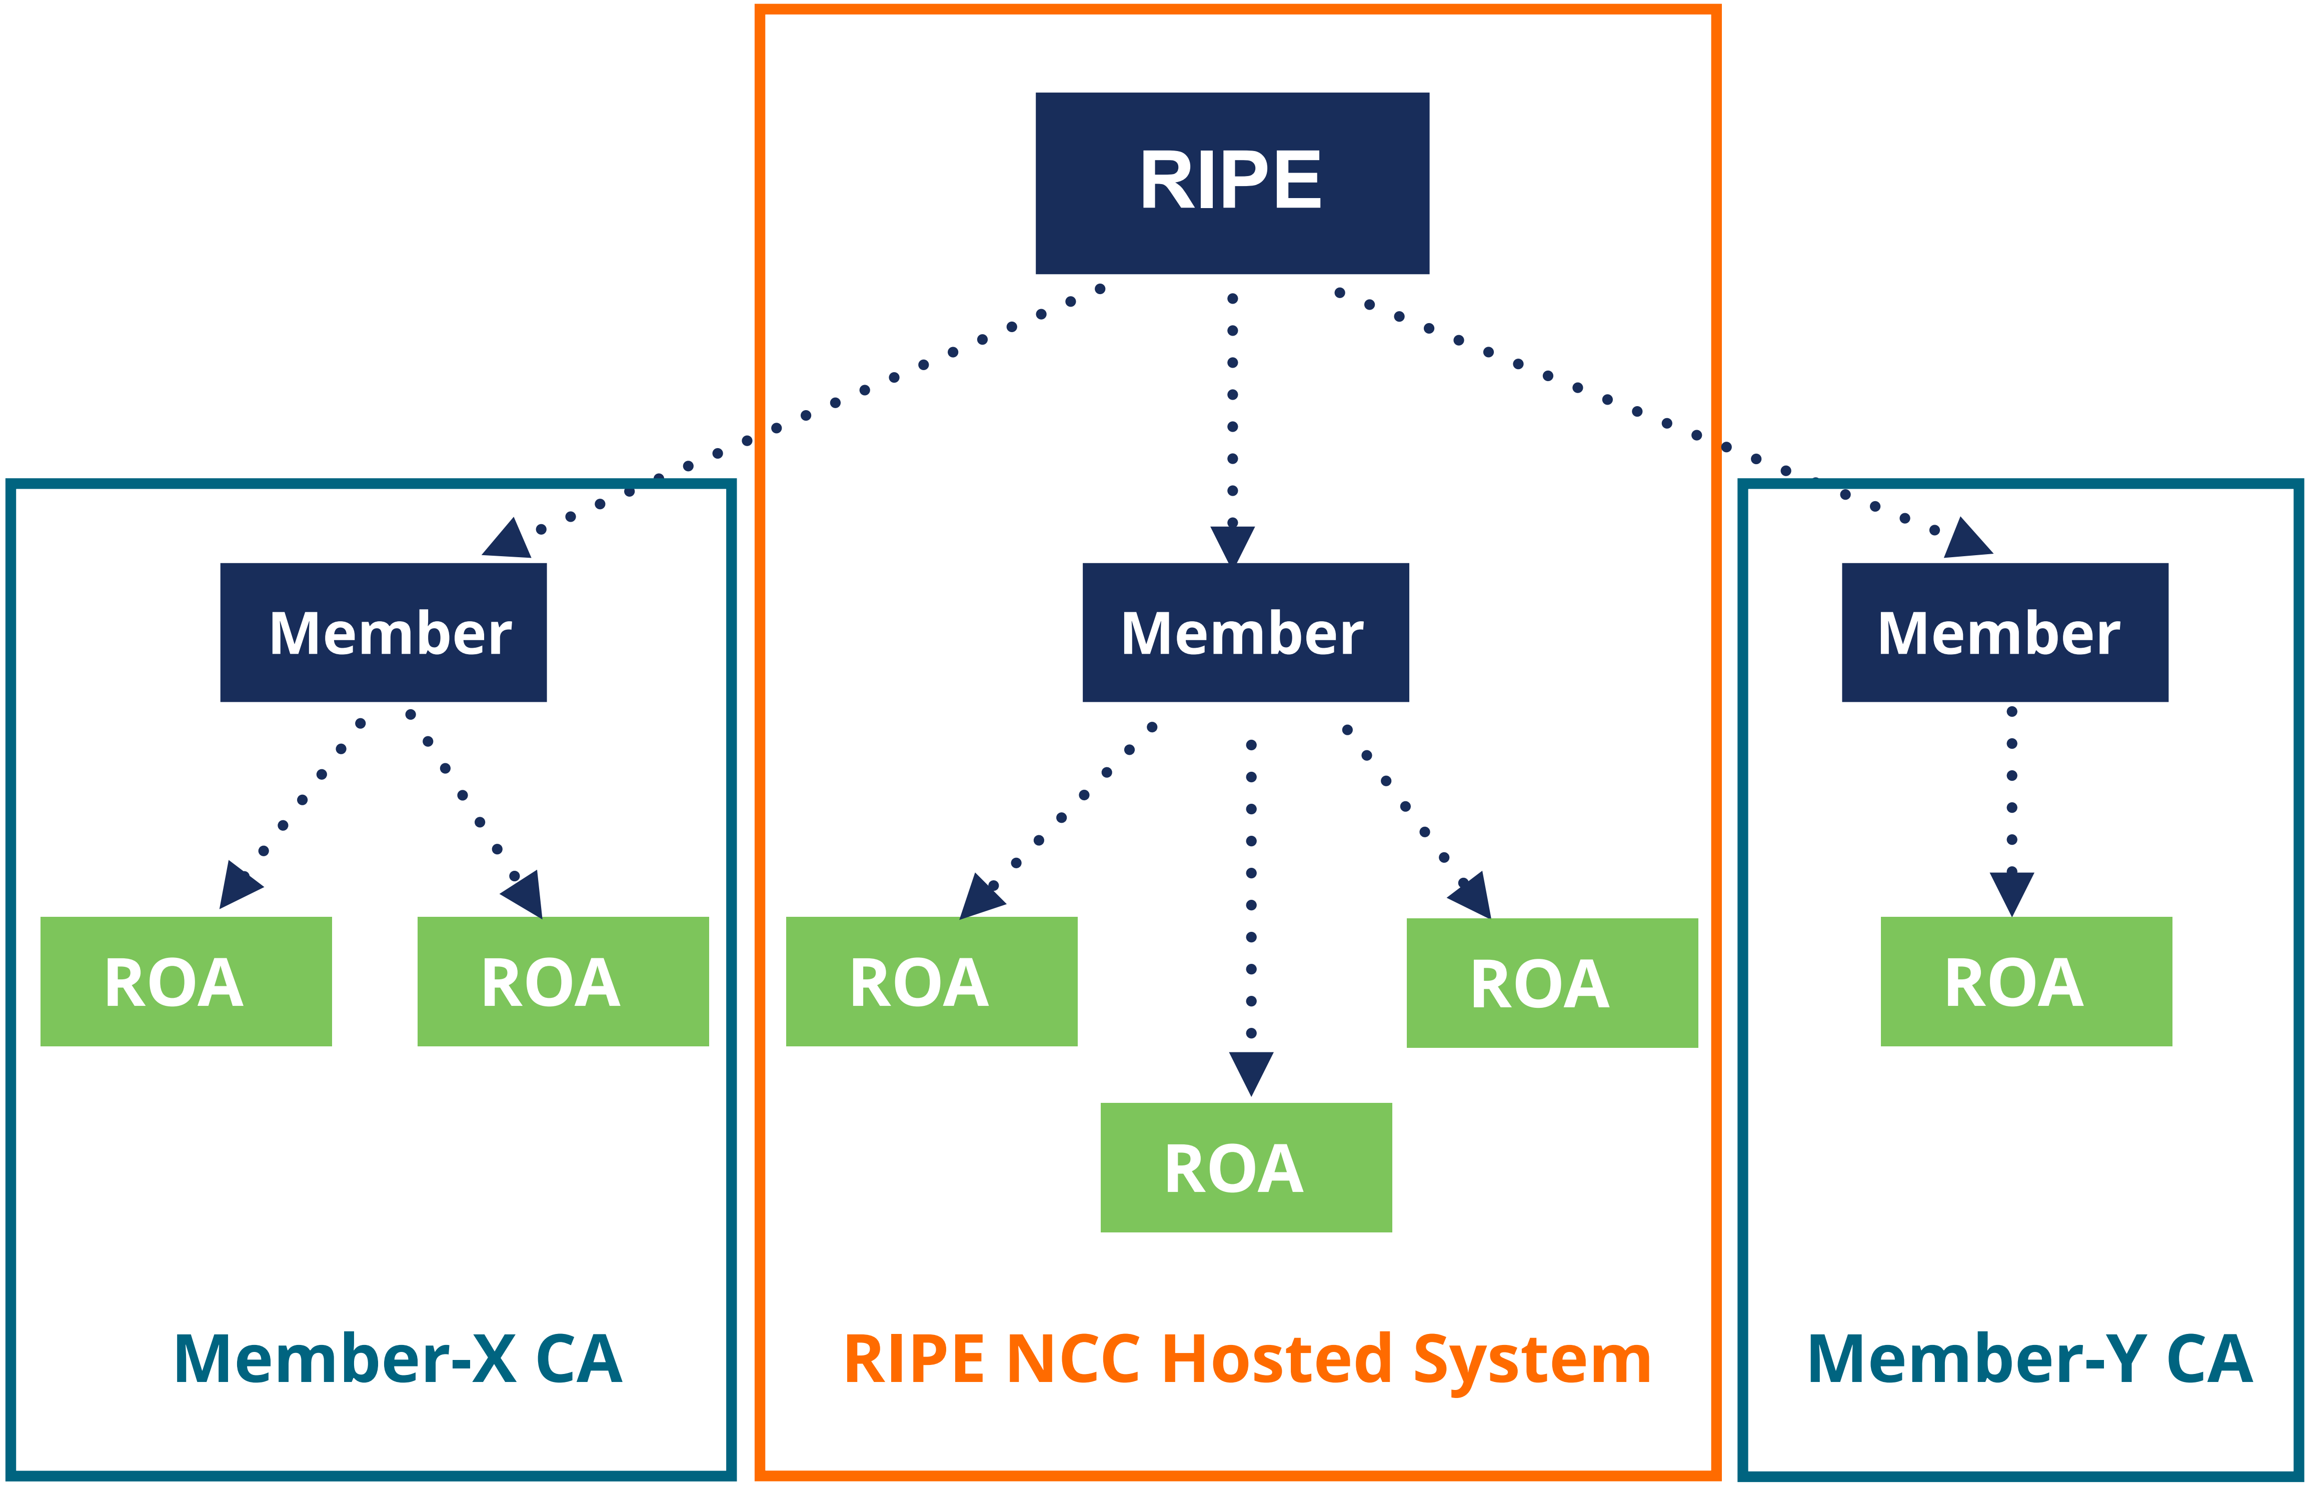

Hosted RPKI

- RIR hosts a CA and signs all ROAs
- Automate signing and key rollovers
- Allows you focus on creating and publishing ROAs

Delegated RPKI

- Run your own Certificate Authority software
  - Dragon Research Labs, RPKI Toolkit
  - NLnet Labs, Krill

- Setup connection with RIPE NCC CA
- Generate your LIR certificate and get it signed by parent CA

Logging in to the RPKI Dashboard

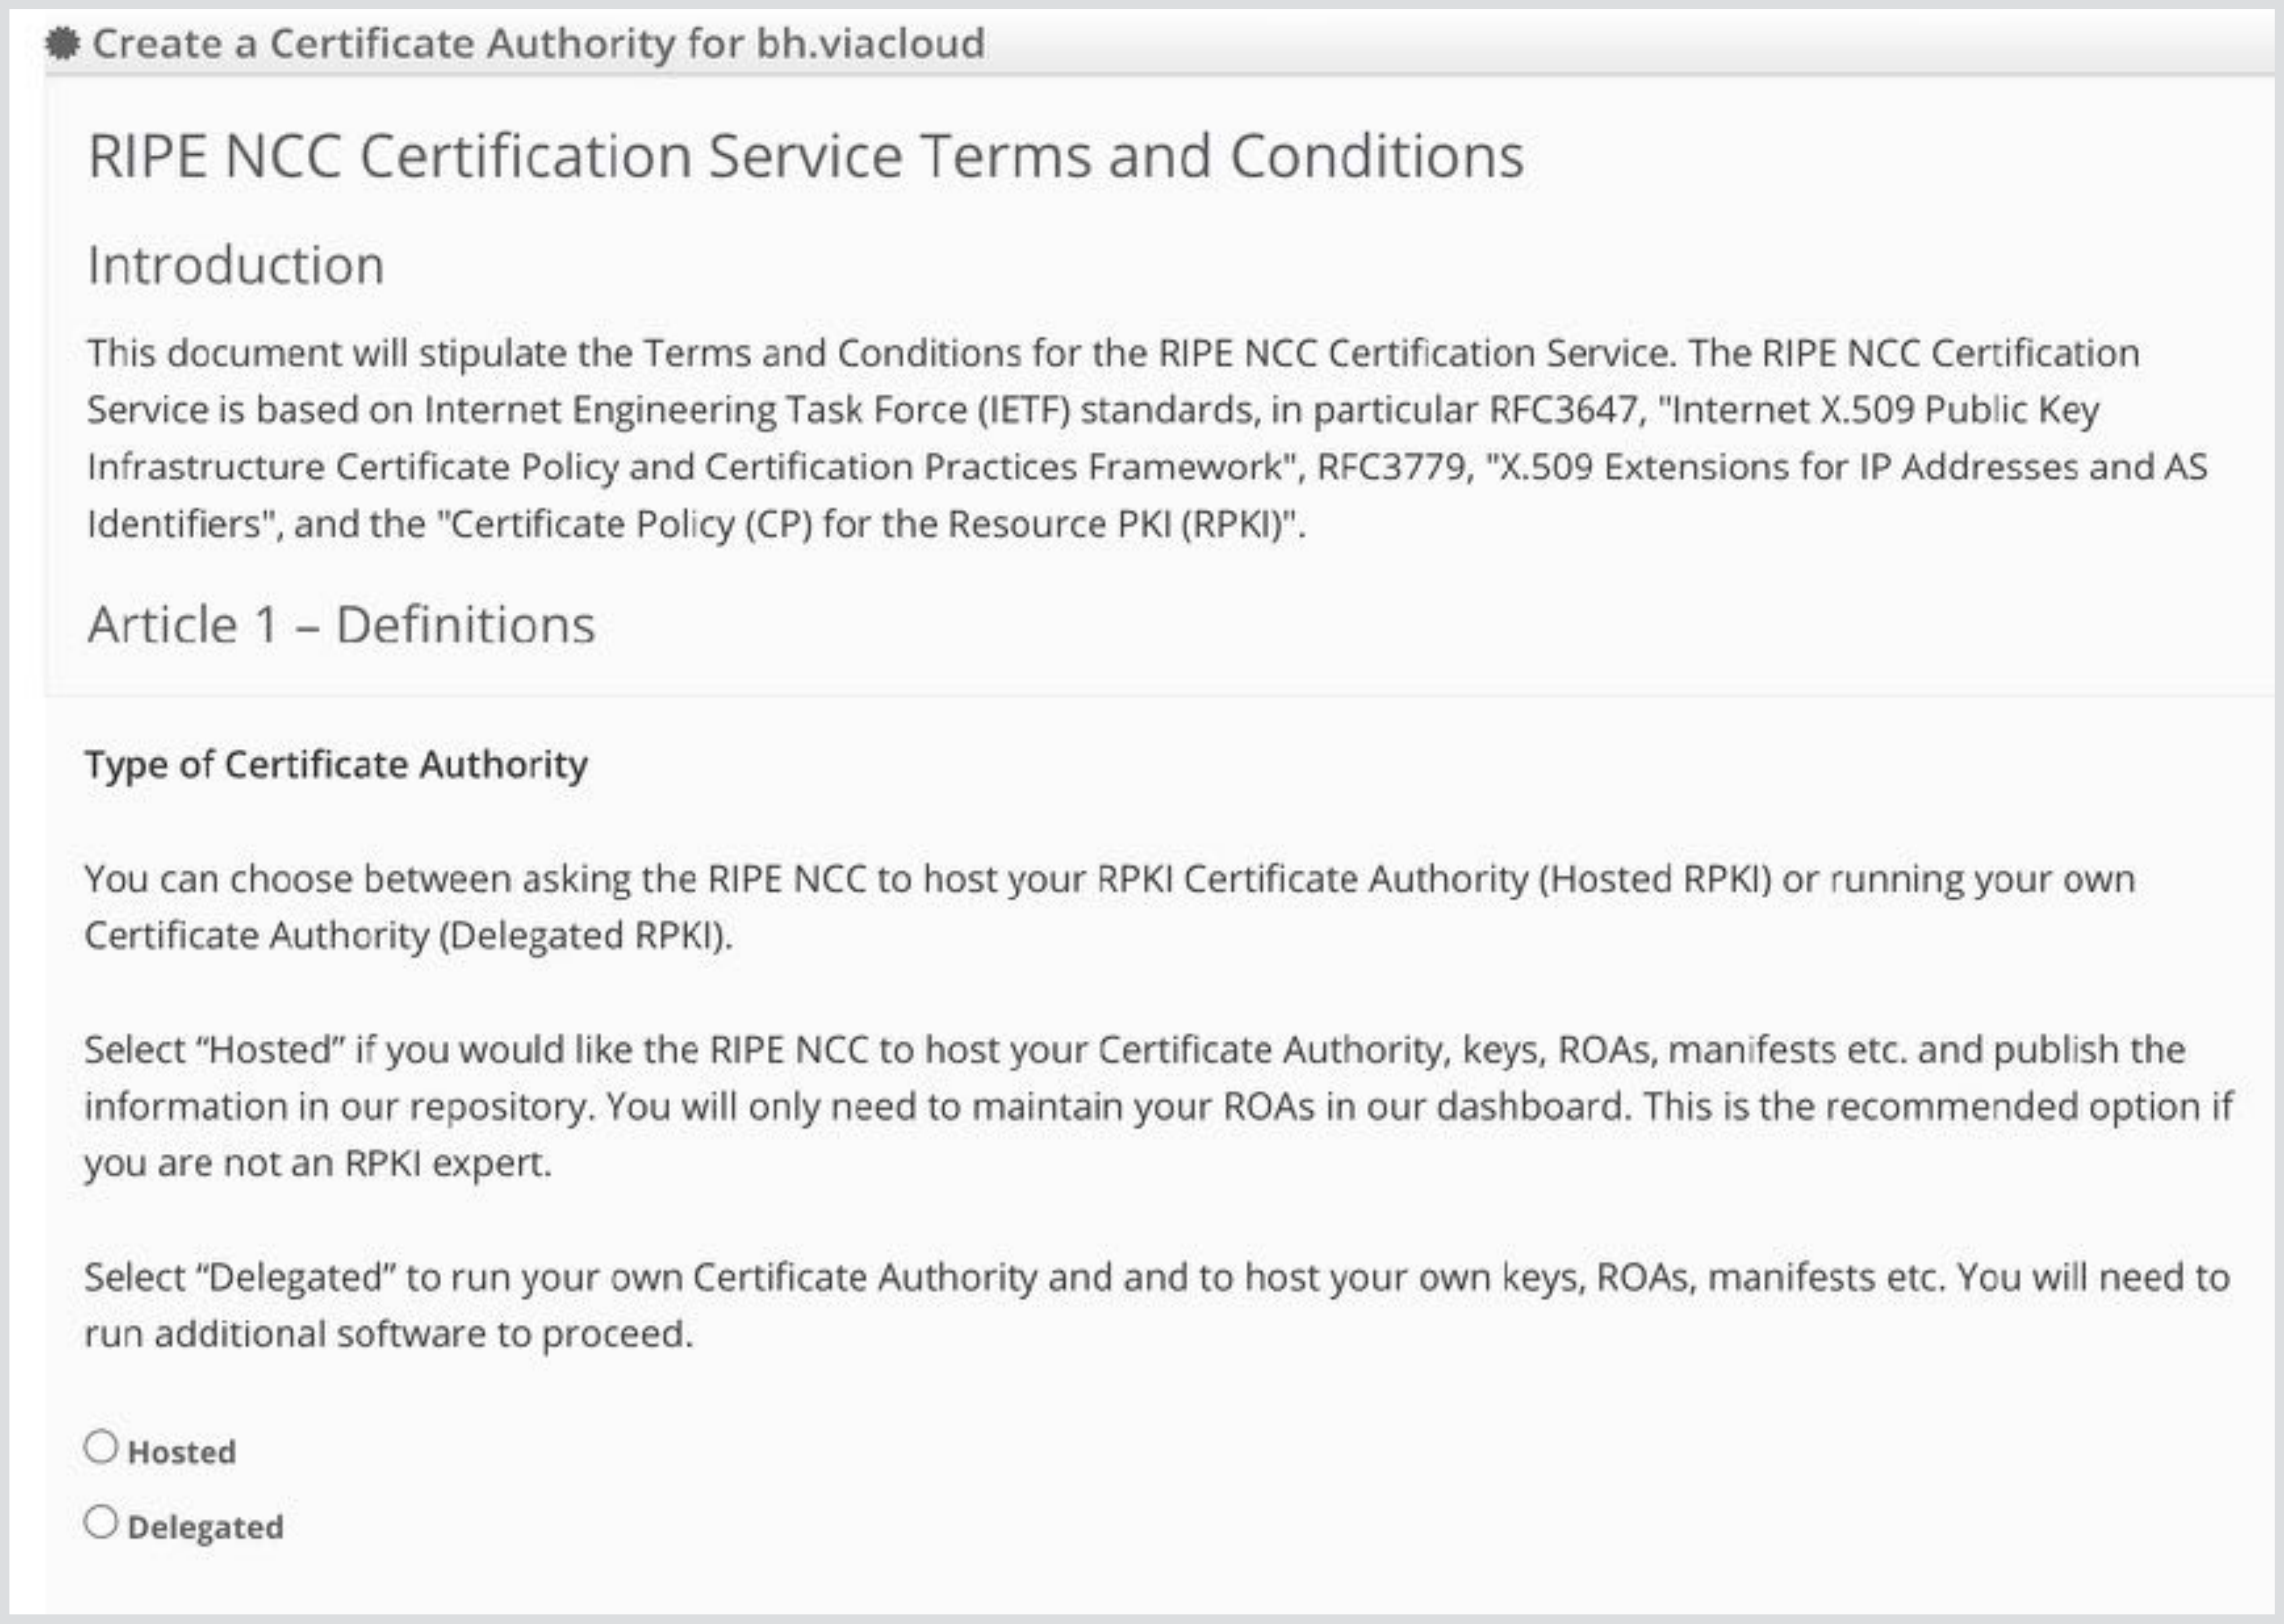

Create a Certificate Authority for bh.viacloud

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**

RPKI Dashboard

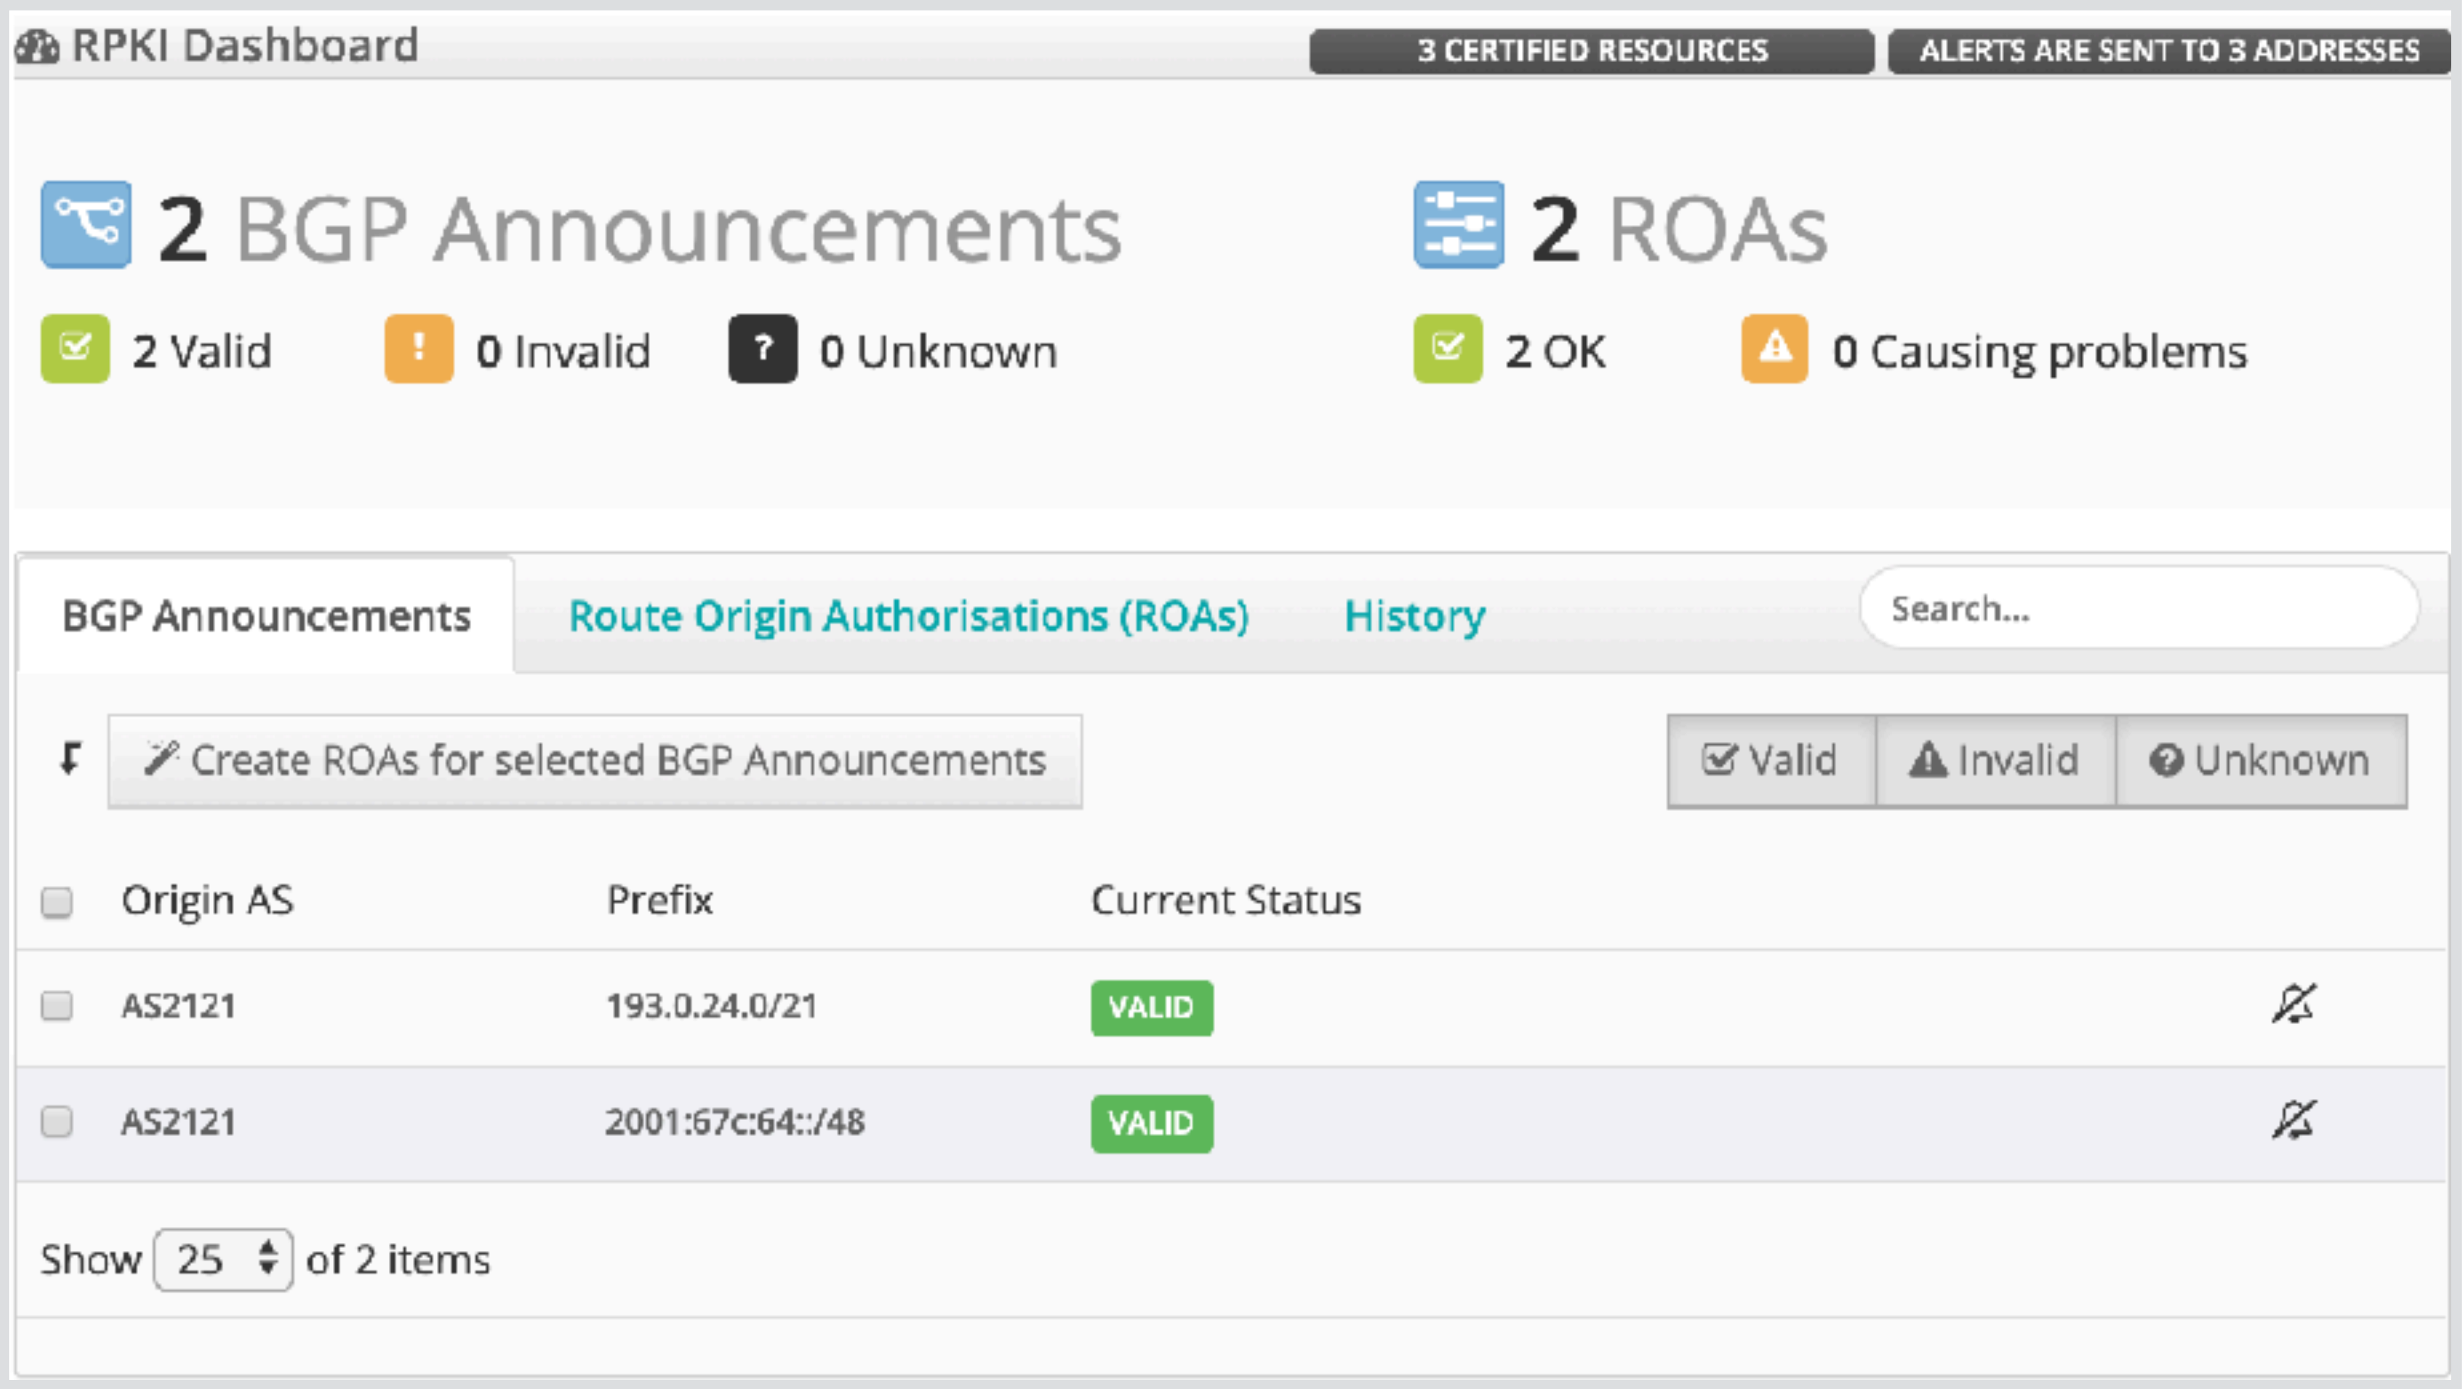

RPKI Dashboard

3 CERTIFIED RESOURCES | ALERTS ARE SENT TO 3 ADDRESSES

2 BGP Announcements

✓ 2 Valid ! 0 Invalid ? 0 Unknown

2 ROAs

✓ 2 OK ⚠ 0 Causing problems

BGP Announcements | Route Origin Authorisations (ROAs) | History | Search...

Create ROAs for selected BGP Announcements

☑ Valid ⚠ Invalid ? Unknown

□	Origin AS	Prefix	Current Status
□	AS2121	193.0.24.0/21	VALID
□	AS2121	2001:67c:64::/48	VALID

Show 25 of 2 items

Revoke hosted CA

Certifying PI Resources

Requested and managed by PI End User or by Sponsoring LIR

1. Complete the wizard successfully

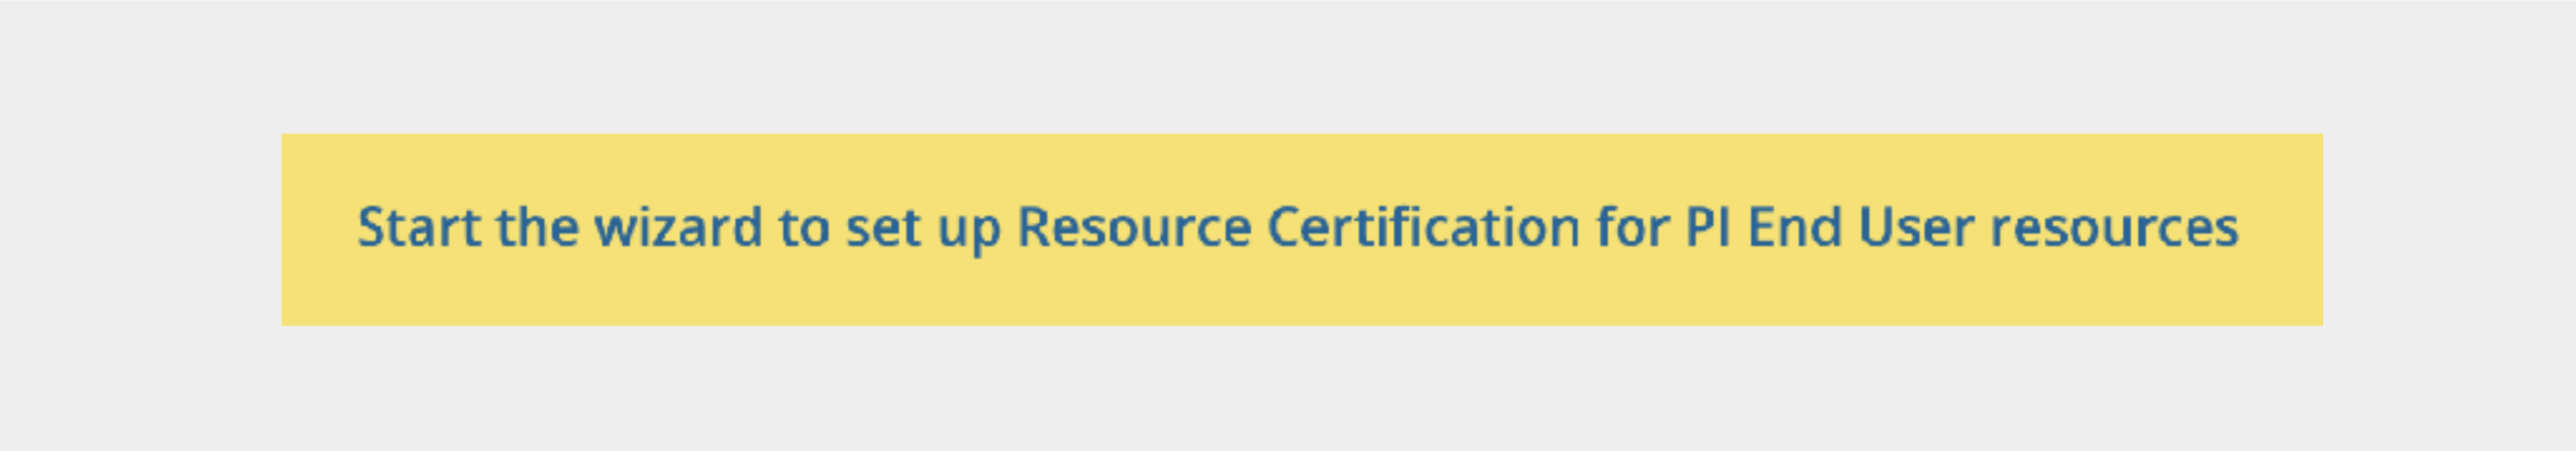

Start the wizard to set up Resource Certification for PI End User resources

2. Login to https://my.ripe.net and request a certificate
   - Sign in with your RIPE NCC Access account
3. Manage your ROAs

Questions

Demo!

Creating ROAs

It’s time to try this yourself!

3 min.

Connect to Localcert:
https://localcert.ripe.net/#/

Deploying RPKI Validators

Section 3

Elements of RPKI

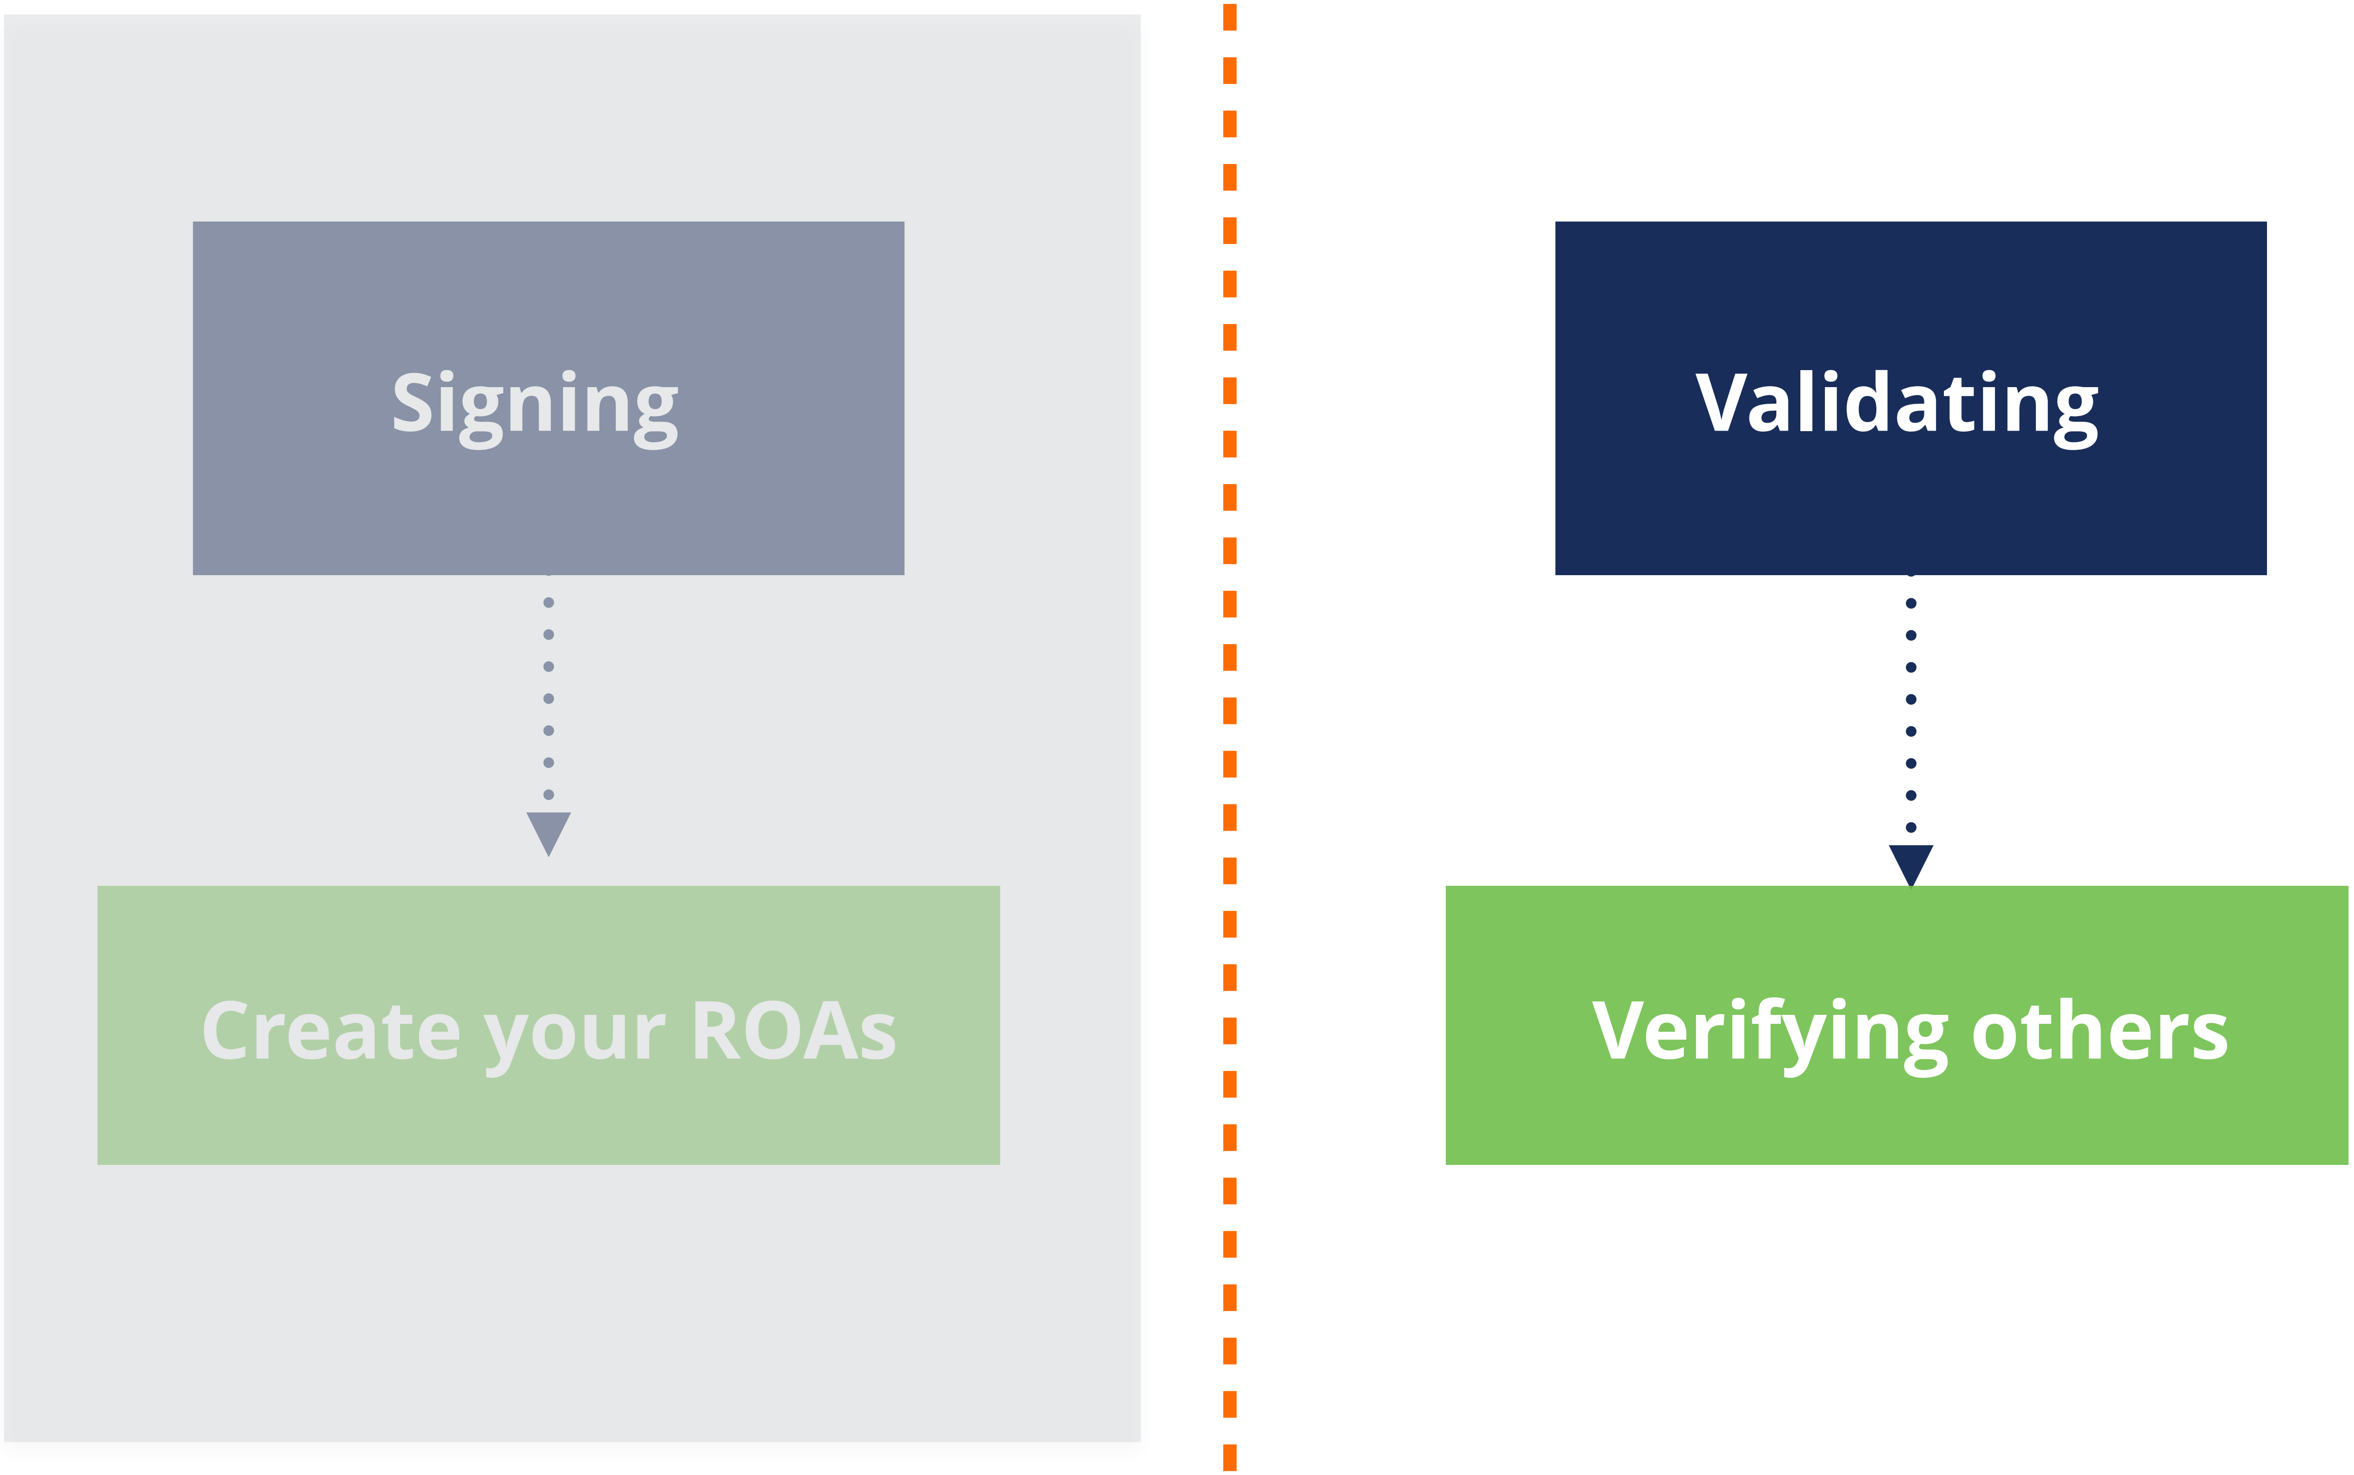

Signing

Create your ROAs

Validating

Verifying others

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 routers using the RPKI-RTR Protocol

Trust Anchor Locator (TAL)

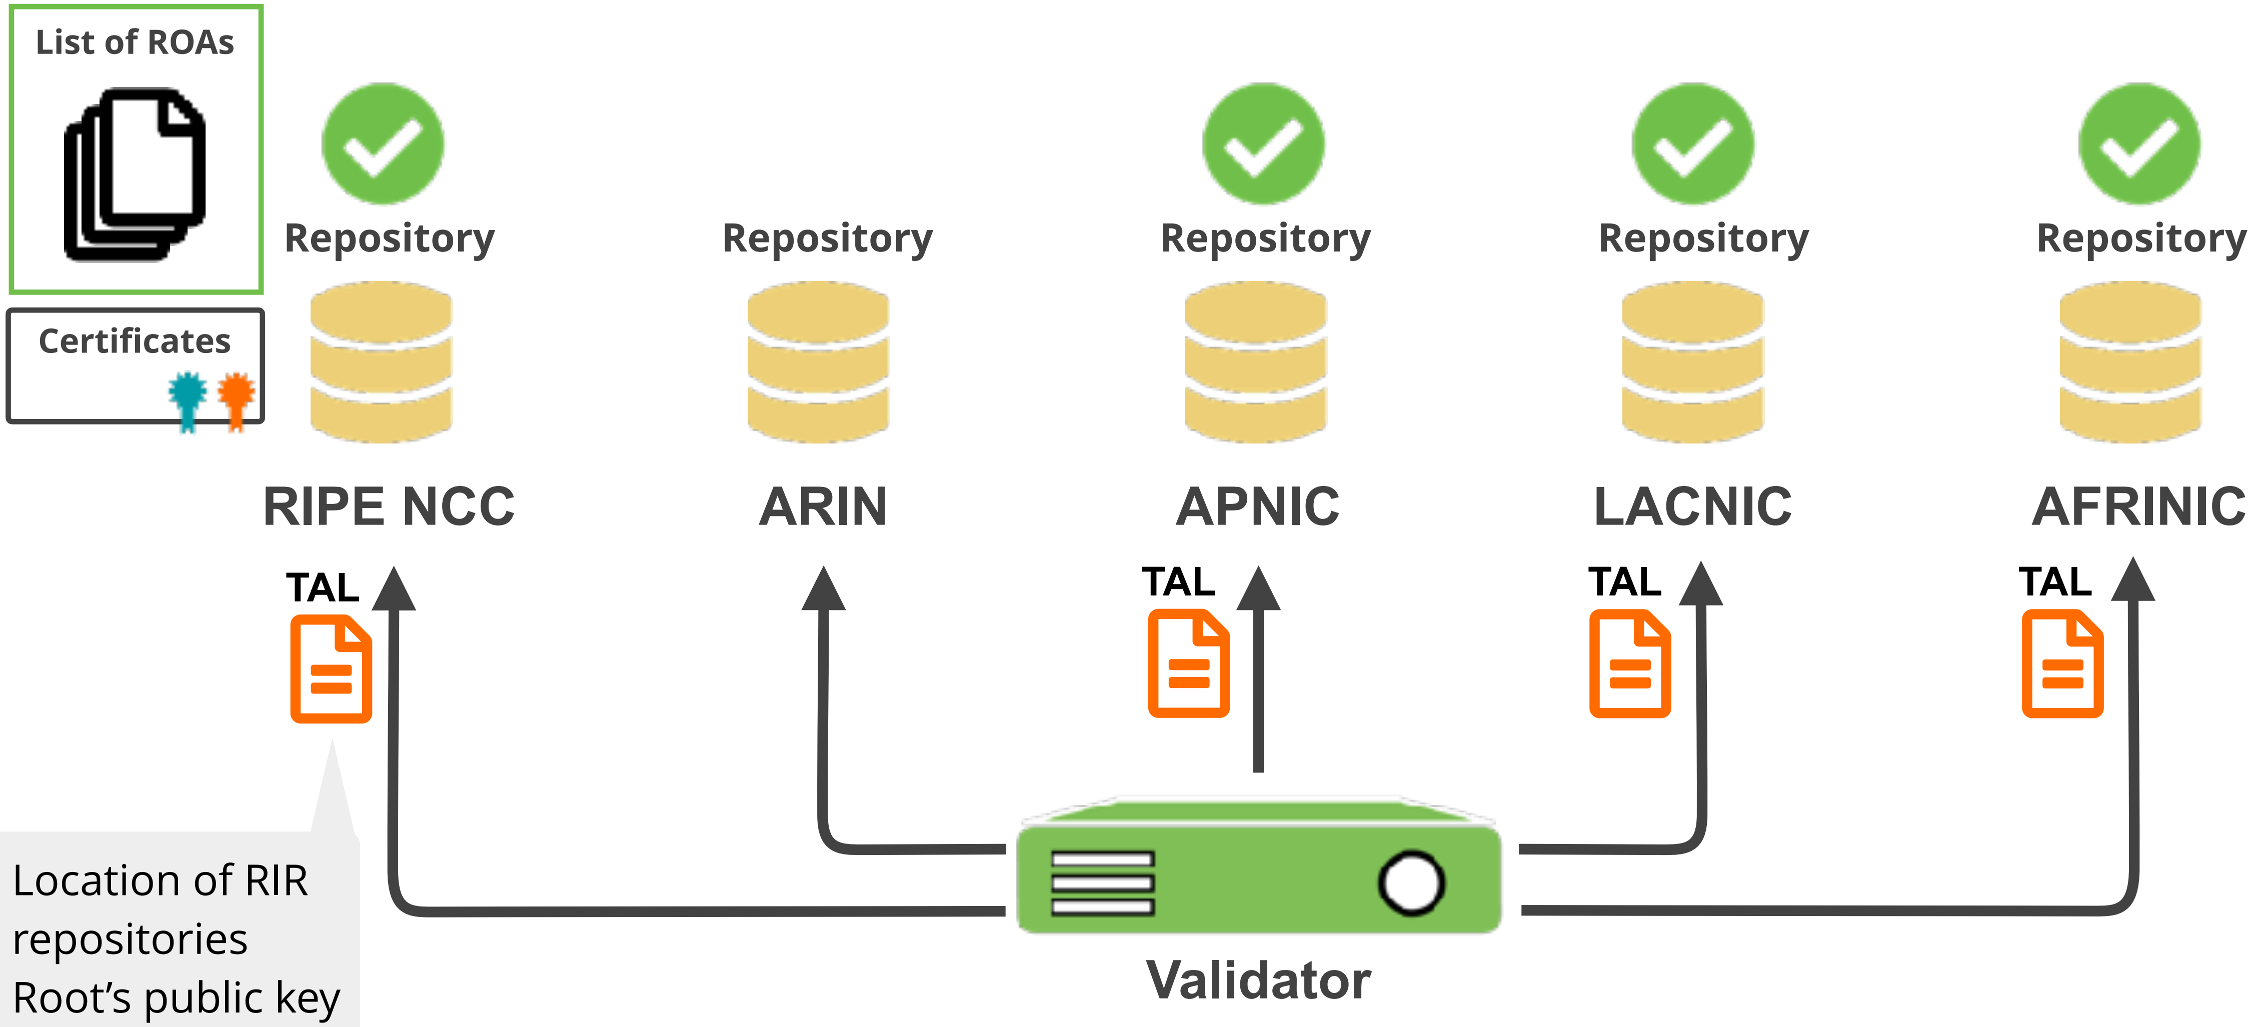

- Location of RIR repositories
- Root’s public key

Relying Party

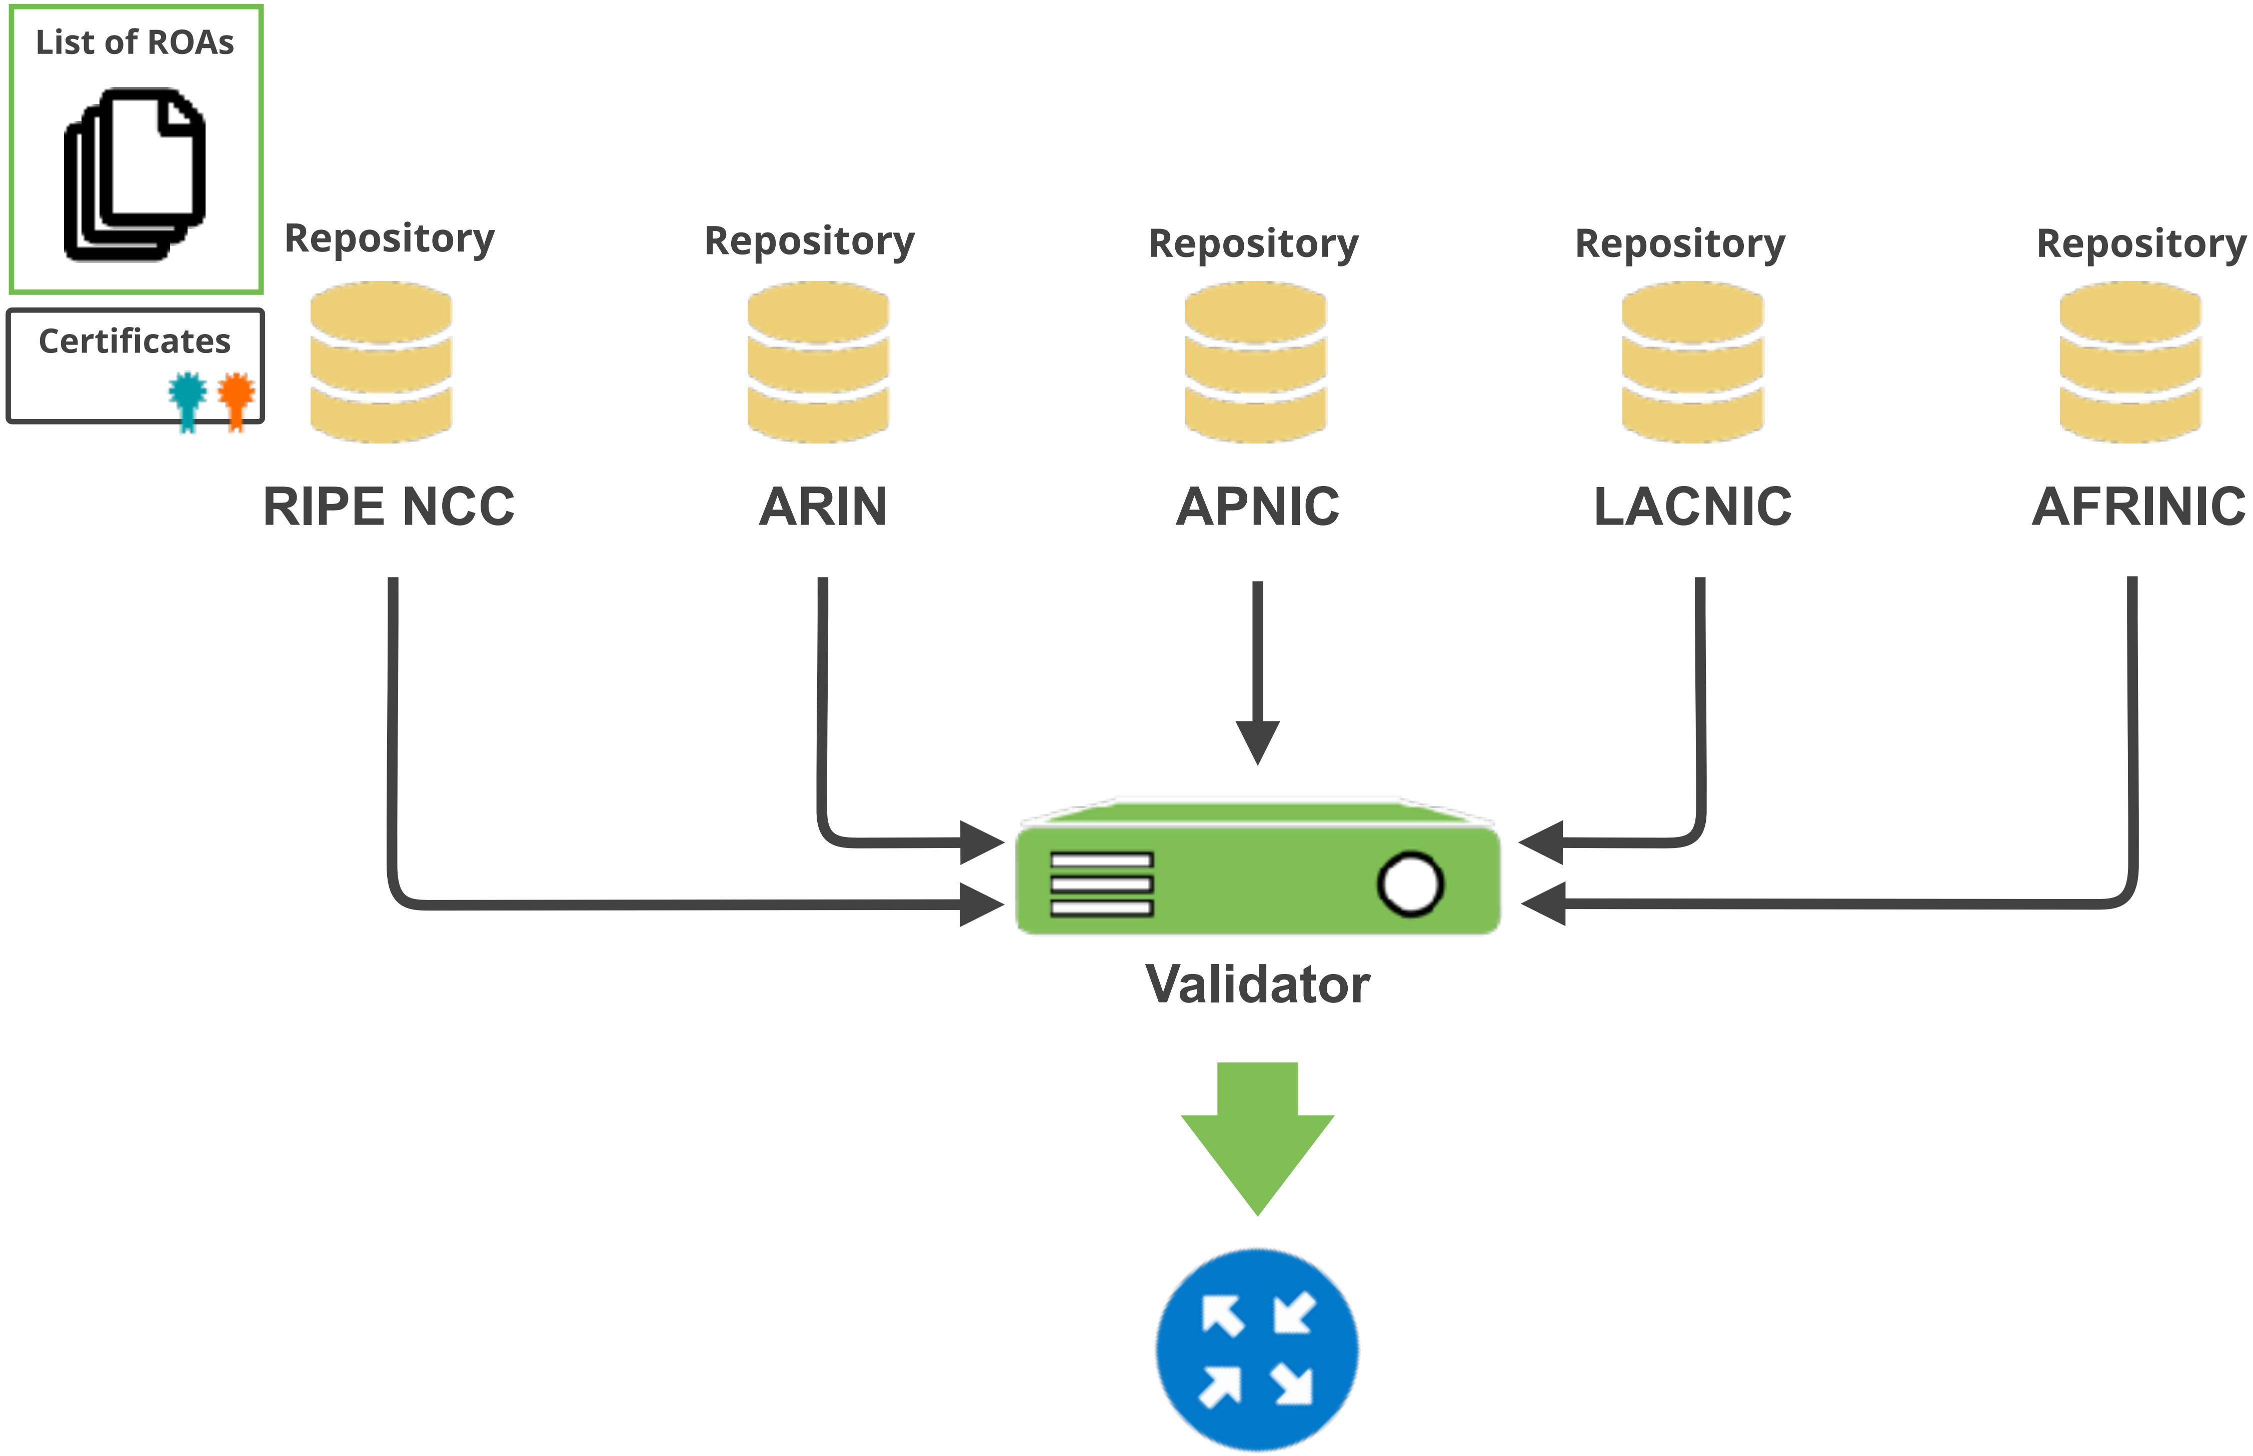

Relying Party

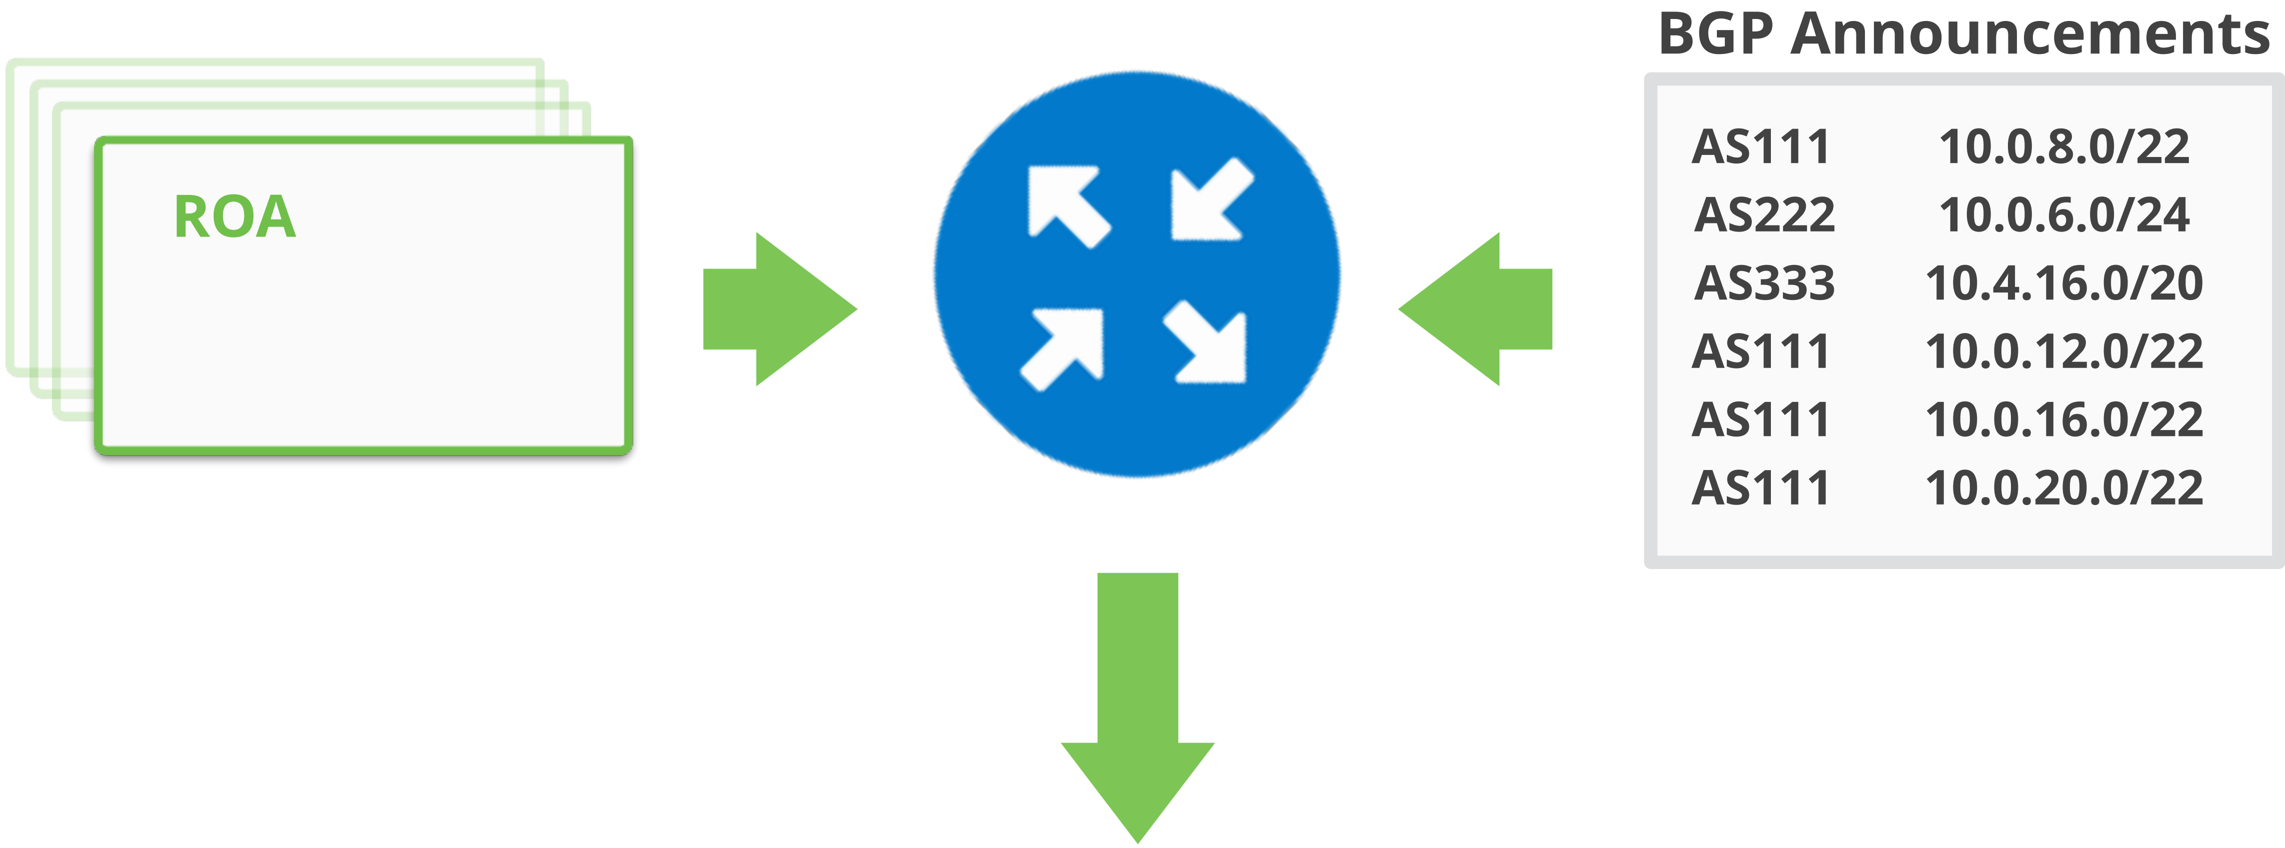

ROA

BGP Announcements

AS111	10.0.8.0/22
AS222	10.0.6.0/24
AS333	10.4.16.0/20
AS111	10.0.12.0/22
AS111	10.0.16.0/22
AS111	10.0.20.0/22

BETTER ROUTING DECISIONS

RPKI Validator Options

- **Routinator**
  - Built with Rust, built by NLNetlabs
- **rpki-client**
  - Part of OpenBSD project, written in C
- **OctoRPKI**
  - Cloudflare’s Relying Party software, written in the Go
- **FORT**
  - Open source RPKI validator, Written in C

Links for Validators

RPKI Validators:

https://github.com/NLnetLabs/routinator.git

https://rpki-client.org/

https://github.com/cloudflare/cfrpki#octorpki

https://github.com/NICMx/FORT-validator/

For more info...

https://rpki.readthedocs.io

Demo!

Running Validators

How to Configure Validators

- Run at least **two** validators
  - Routinator (0.8.2)
  - FORT (1.4.2)

- Configure the **correct TALs**
  - They have already been downloaded
  - ARIN TAL needs to be installed separately

Start the Routinator

On the Server:

```
routinator server --rtr 100.64.1.1:3323
```

➡ TAL directory is **missing!**

➡ We need to initialize via **init command!**

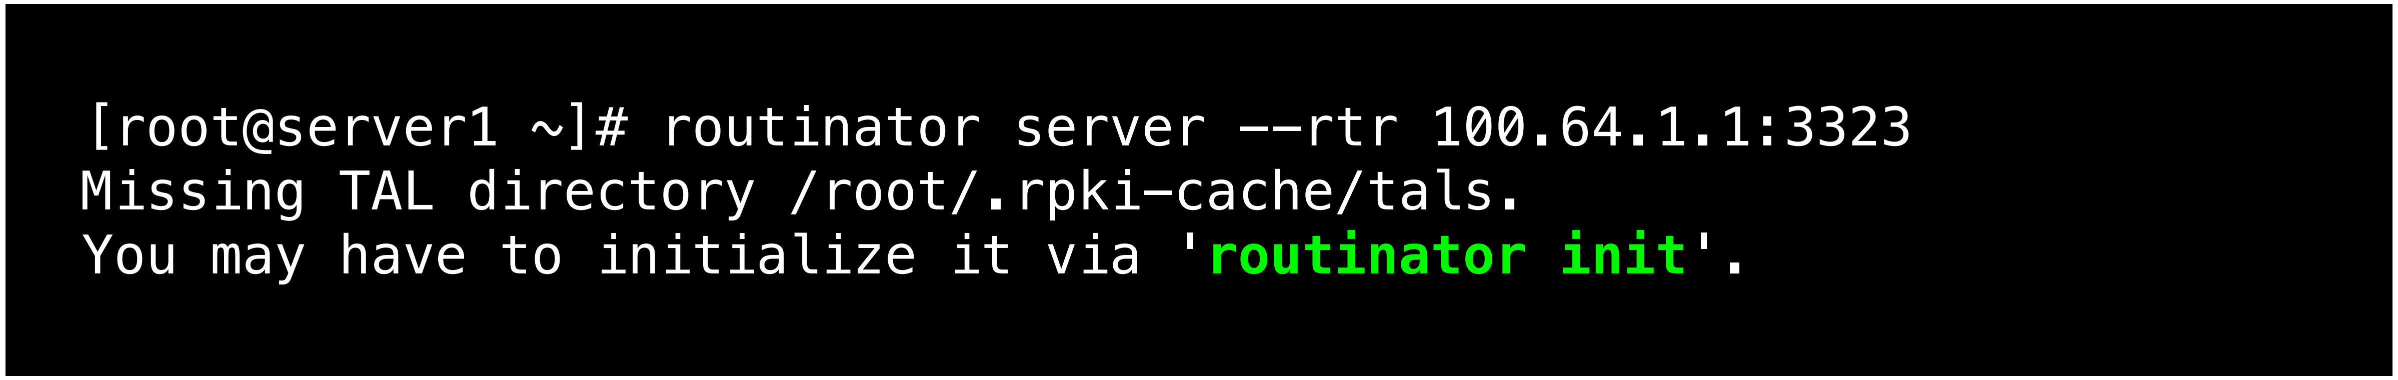

```
[root@server1 ~]# routinator server --rtr 100.64.1.1:3323
Missing TAL directory /root/.rpki-cache/tals.
You may have to initialize it via 'routinator init'.
```

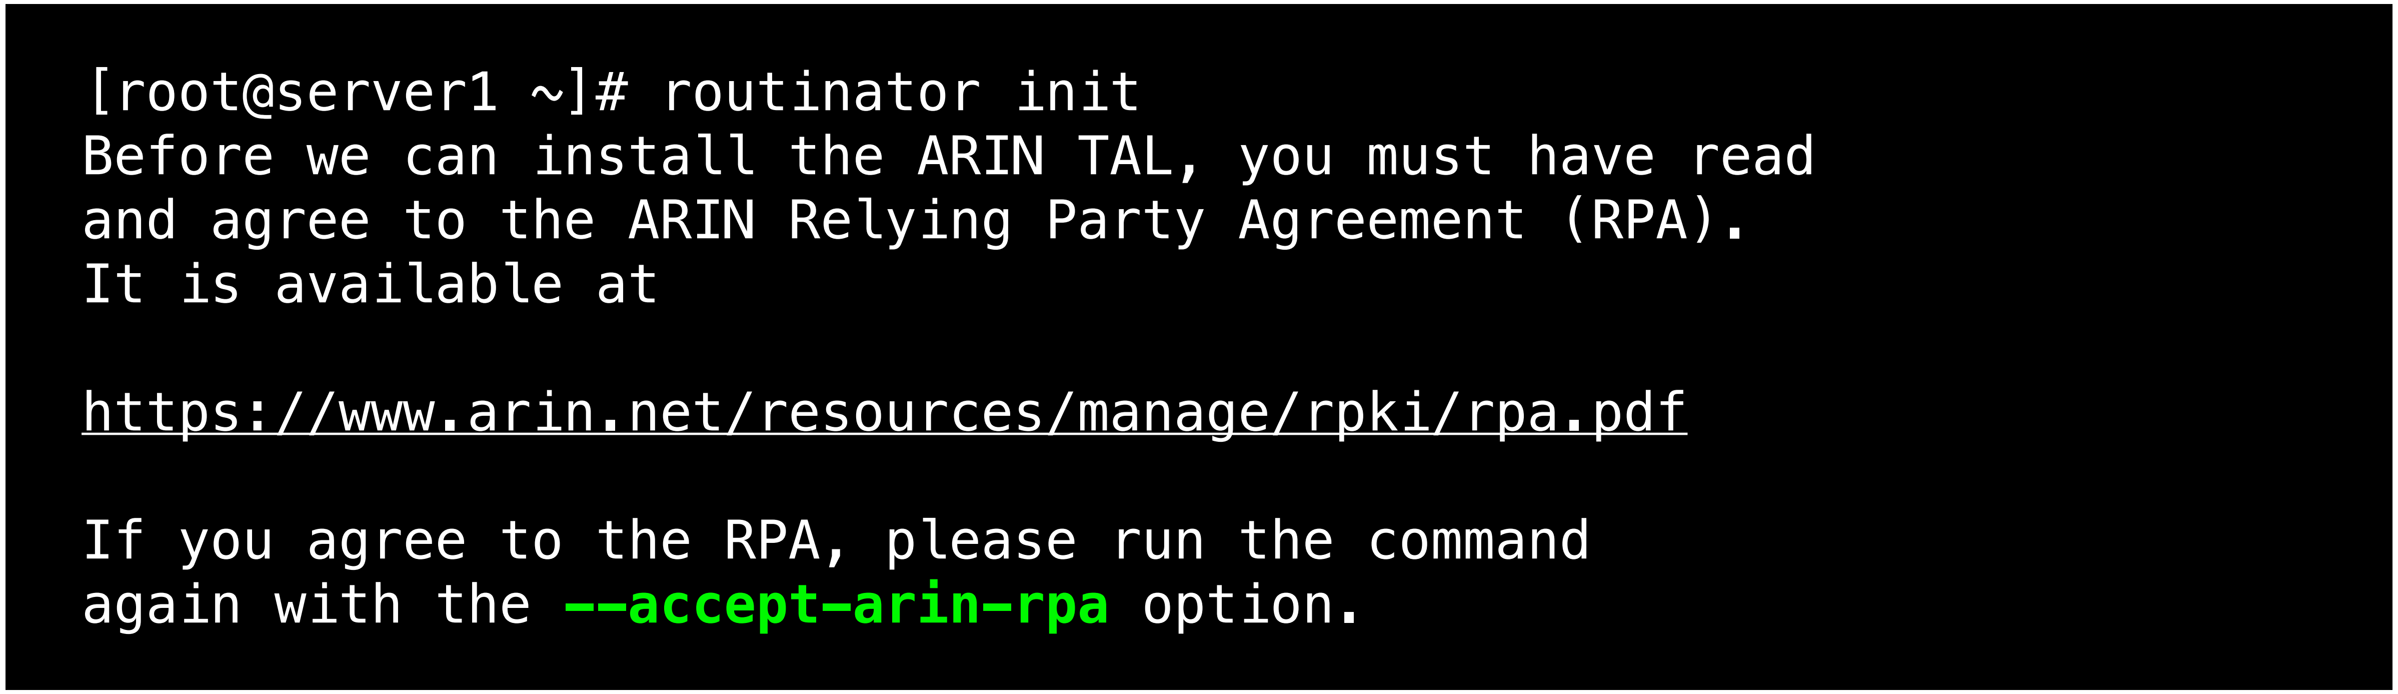

```
[root@server1 ~]# routinator init
Before we can install the ARIN TAL, you must have read
and agree to the ARIN Relying Party Agreement (RPA).
It is available at

https://www.arin.net/resources/manage/rpki/rpa.pdf

If you agree to the RPA, please run the command
again with the --accept-arin-rpa option.
```

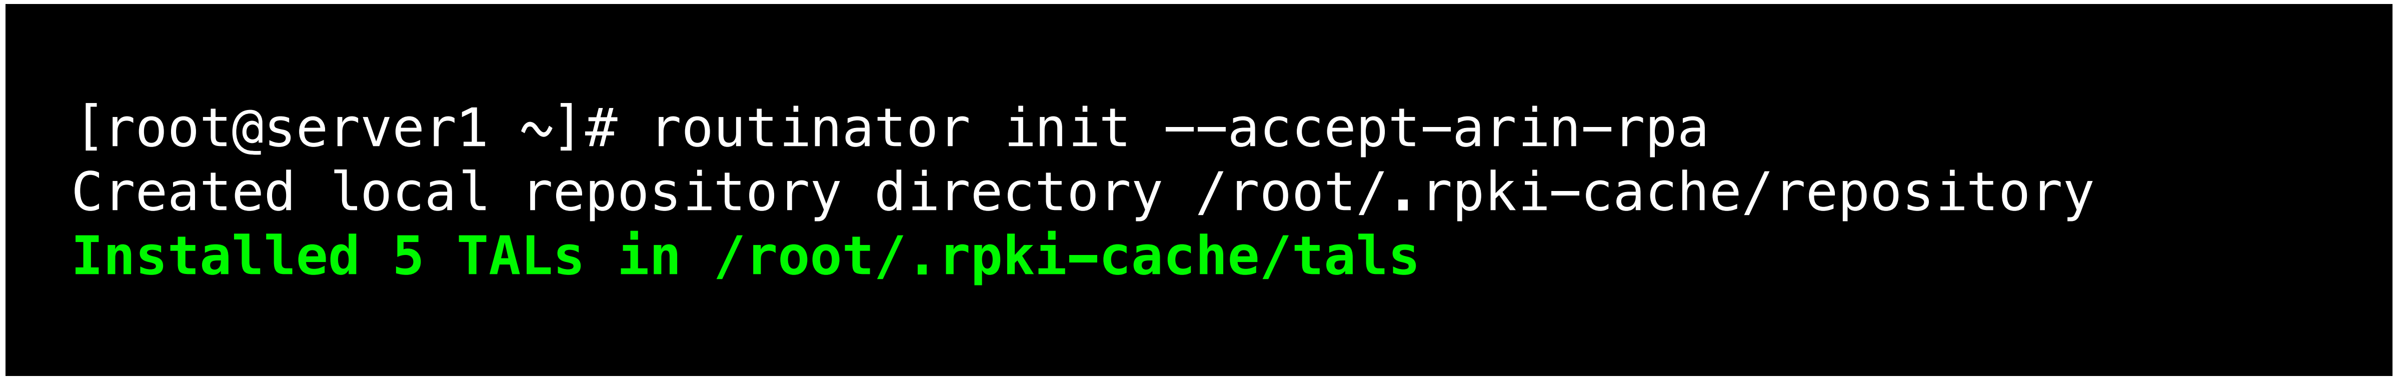

```
[root@server1 ~]# routinator init --accept-arin-rpa
Created local repository directory /root/.rpki-cache/repository
Installed 5 TALs in /root/.rpki-cache/tals
```

Start the Routinator

On the Server:

```
routinator server --rtr 100.64.1.1:3323
```

Check if it’s running

```
ps aux | grep routinator
```

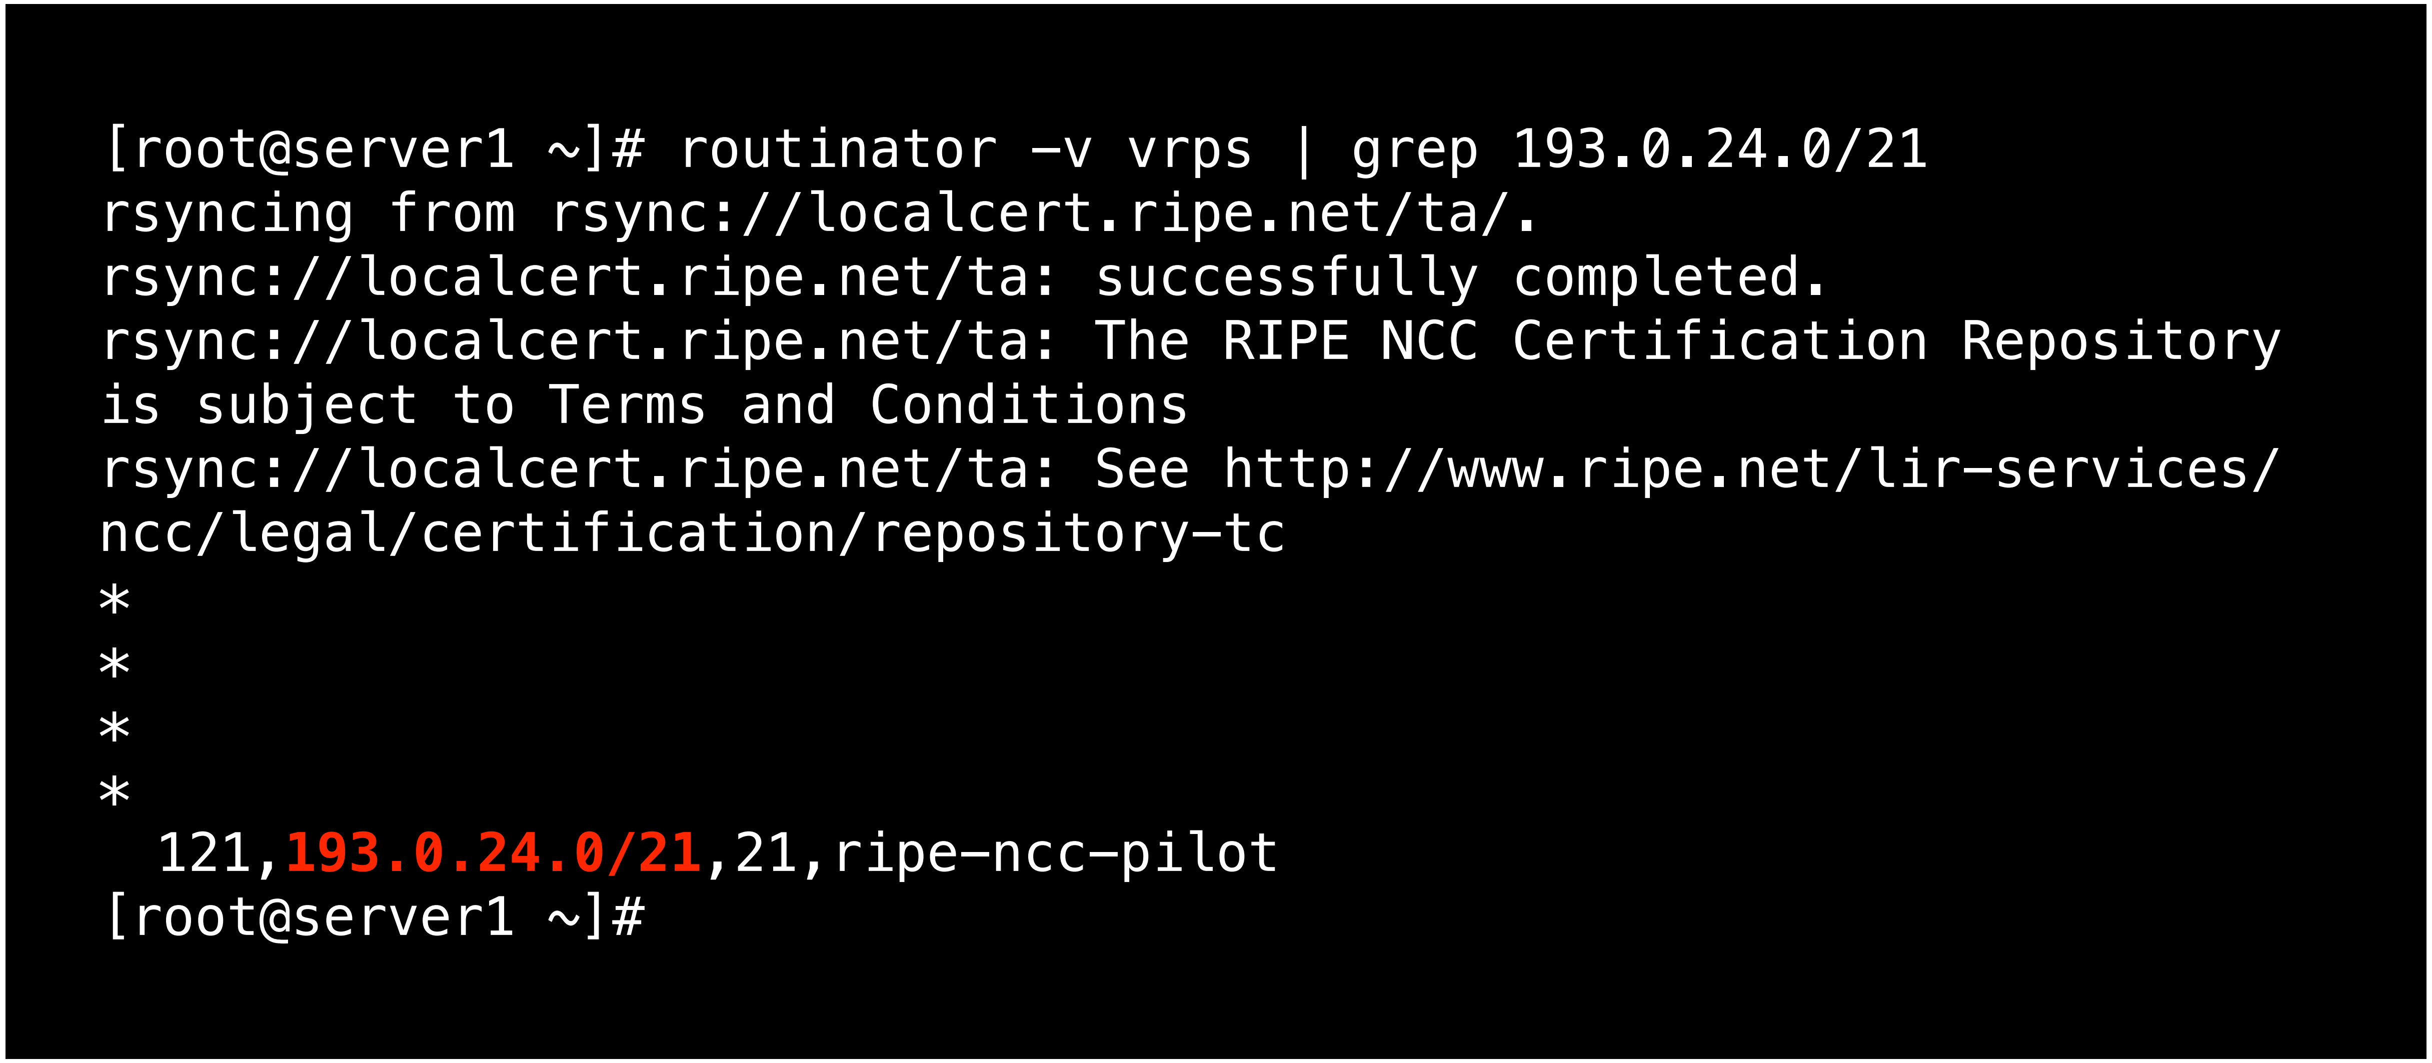

```
[root@server1 ~]# routinator -v vrps | grep 193.0.24.0/21
rsyncing from rsync://localcert.ripe.net/ta/.
rsync://localcert.ripe.net/ta: successfully completed.
rsync://localcert.ripe.net/ta: The RIPE NCC Certification Repository
is subject to Terms and Conditions
rsync://localcert.ripe.net/ta: See http://www.ripe.net/lir-services/
ncc/legal/certification/repository-tc
*
*
*
*
  121,193.0.24.0/21,21,ripe-ncc-pilot
[root@server1 ~]#
```

Start FORT validator

On the Server:

```
fort --init-tals -tal=/etc/fort/tal
```

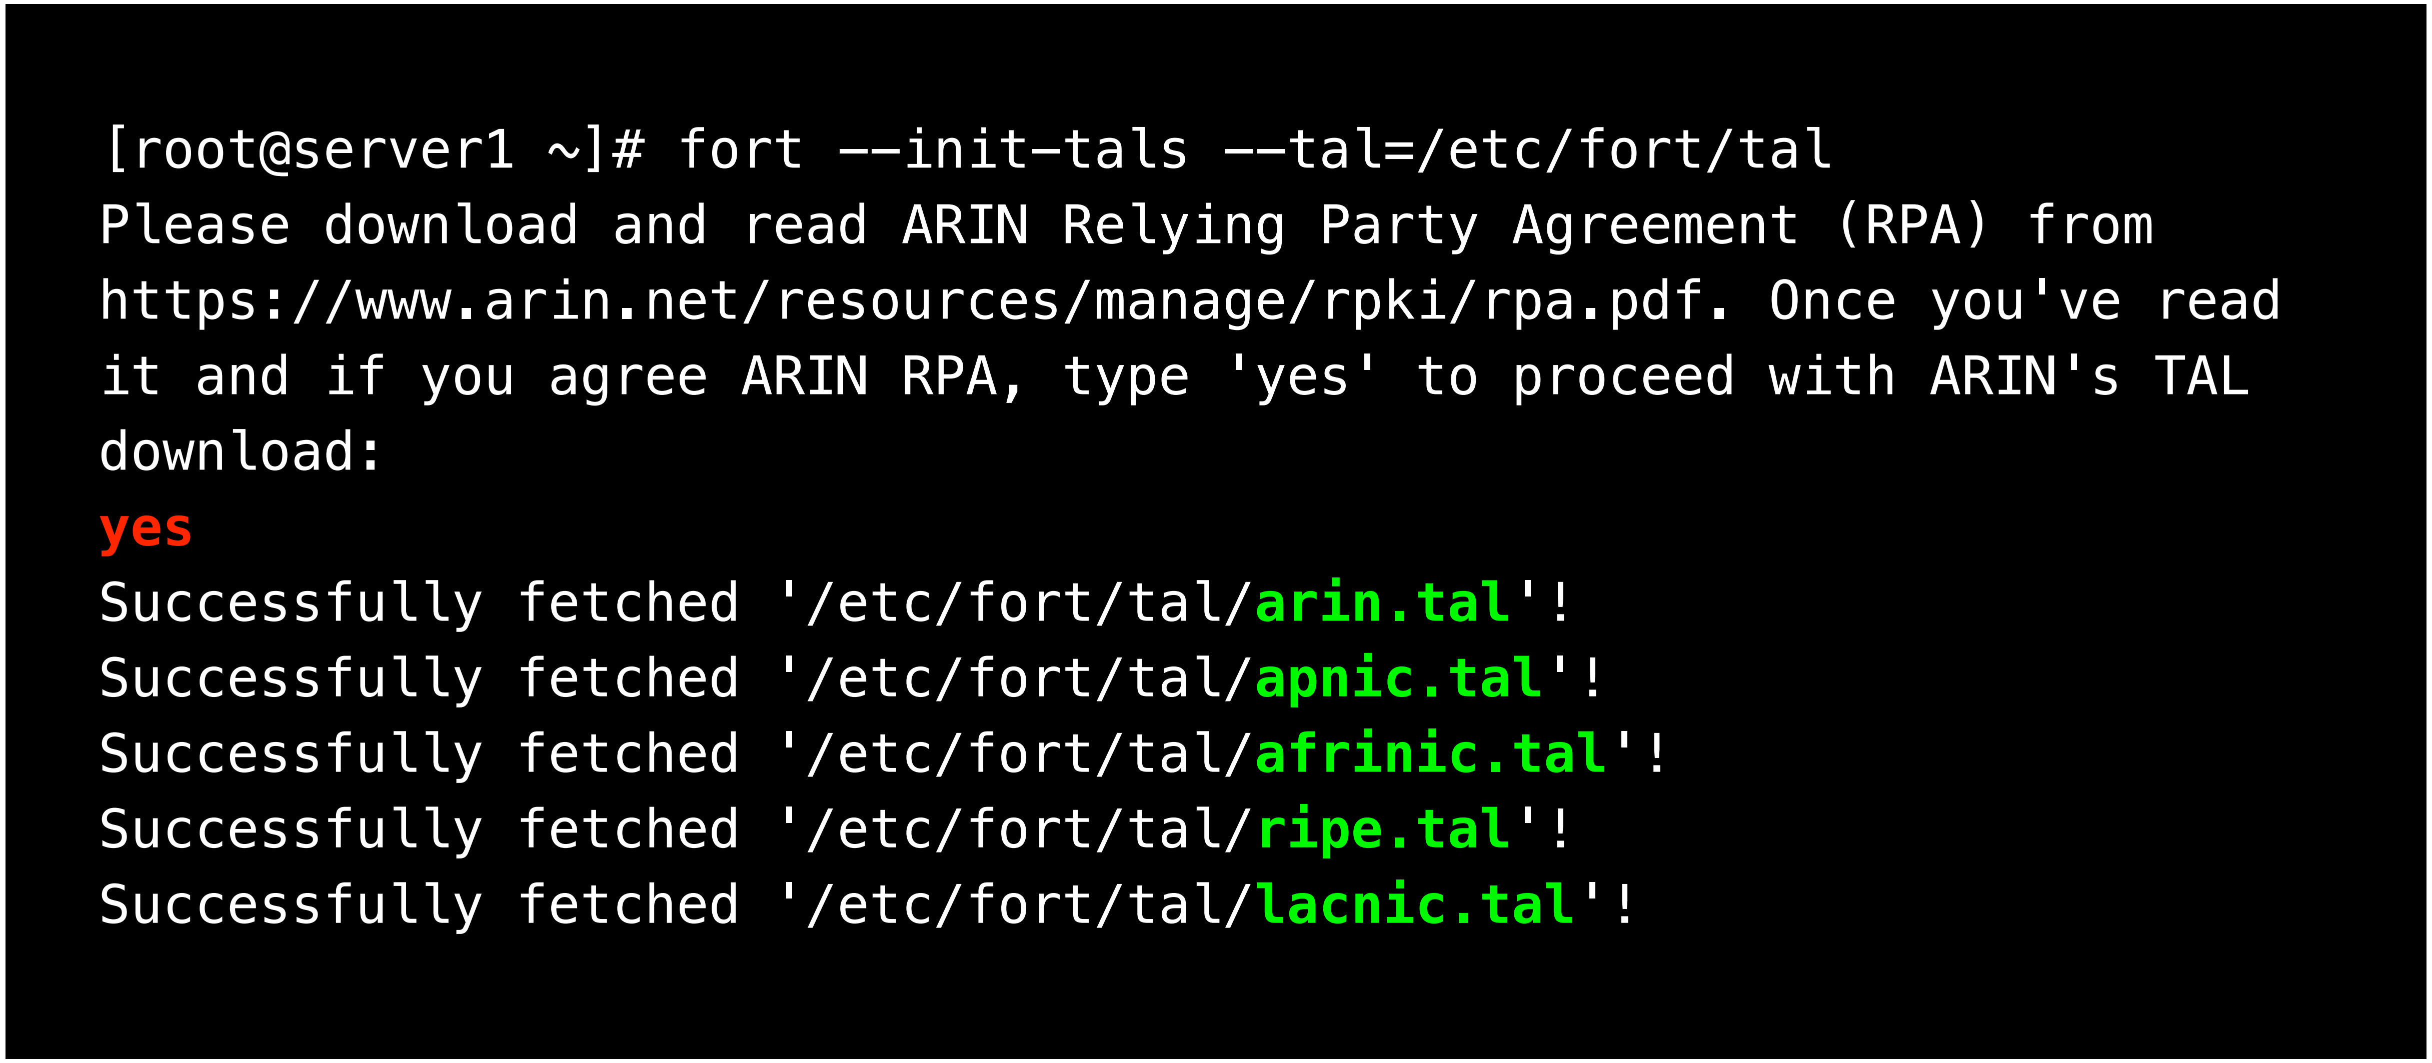

```
[root@server1 ~]# fort --init-tals --tal=/etc/fort/tal
Please download and read ARIN Relying Party Agreement (RPA) from
https://www.arin.net/resources/manage/rpki/rpa.pdf. Once you've read
it and if you agree ARIN RPA, type 'yes' to proceed with ARIN's TAL
download:
yes
Successfully fetched '/etc/fort/tal/arin.tal'!
Successfully fetched '/etc/fort/tal/apnic.tal'!
Successfully fetched '/etc/fort/tal/afrinic.tal'!
Successfully fetched '/etc/fort/tal/ripe.tal'!
Successfully fetched '/etc/fort/tal/lacnic.tal'!
```

Start FORT validator

On the Server:

```
systemctl start fort
```

Check if it is running and the logs (exit with ctrl-c):

```
Systemctl status fort

journalctl –u fort
```

- FORT will not start RTR server before it does the validation for the first time.
- It listens on port **323** by default.
- Configuration is in **`/etc/fort/config.json`**

- To check whether FORT is listening

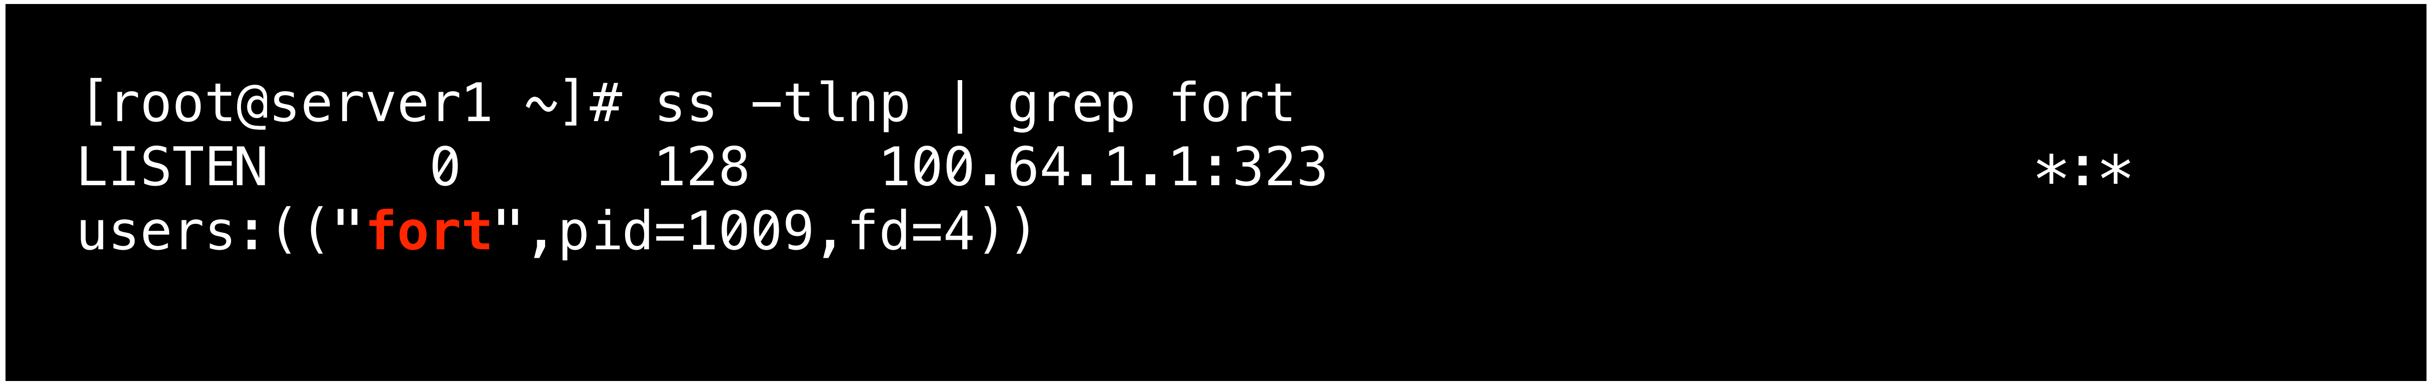

```
[root@server1 ~]# ss -tlnp | grep fort
LISTEN     0      128    100.64.1.1:323                      *:*
users:(("fort",pid=1009,fd=4))
```

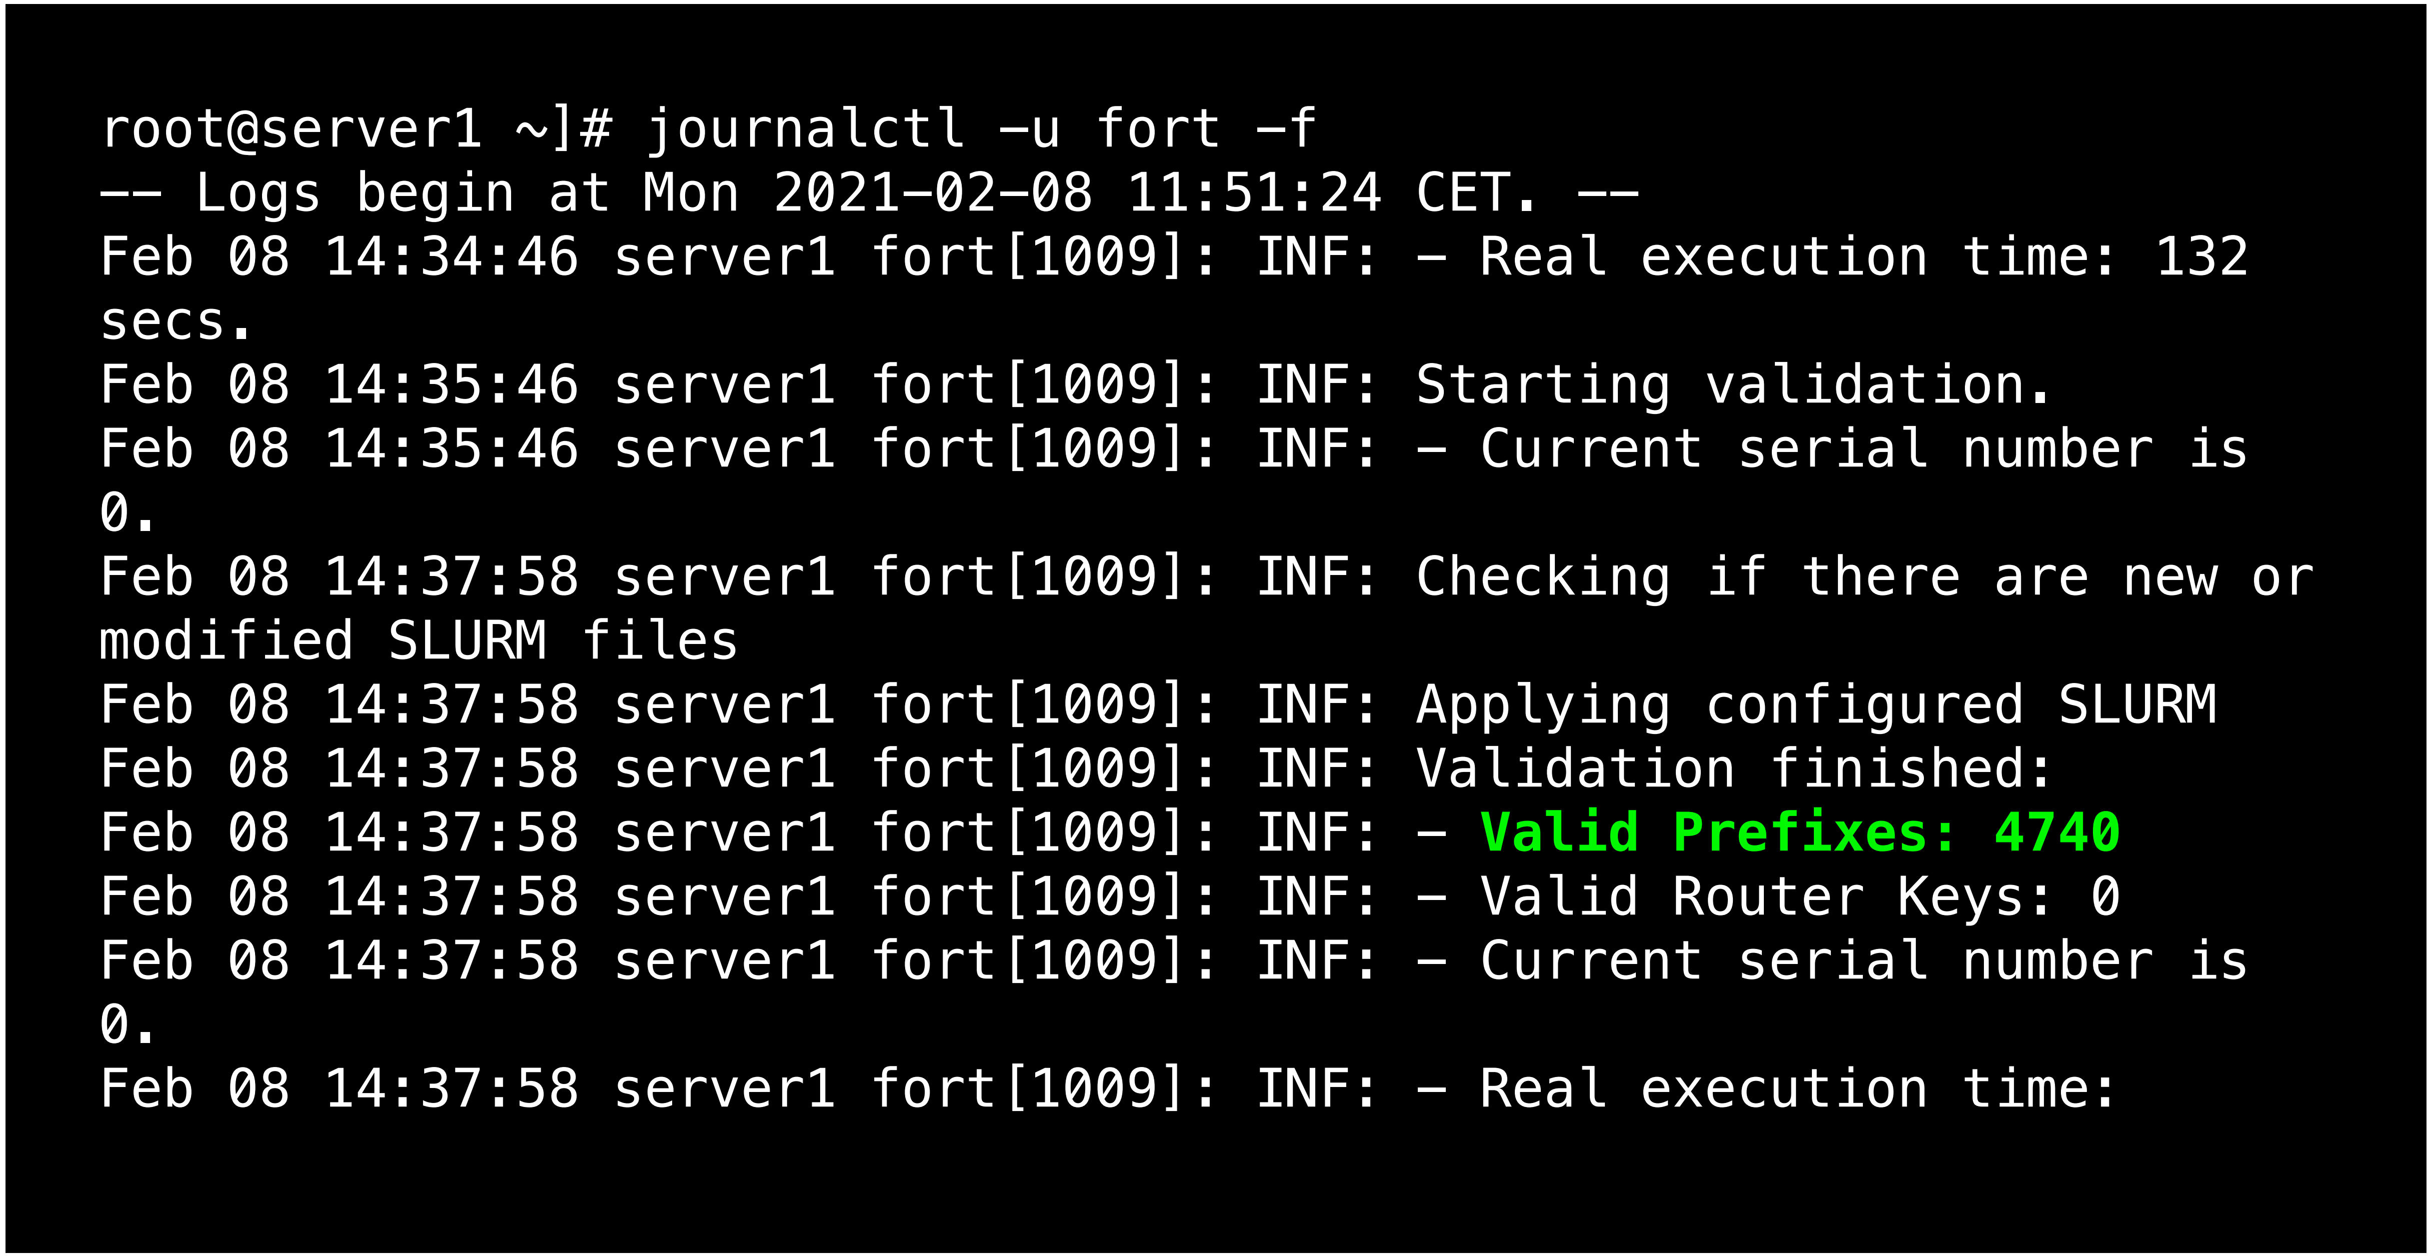

```
root@server1 ~]# journalctl -u fort -f
-- Logs begin at Mon 2021-02-08 11:51:24 CET. --
Feb 08 14:34:46 server1 fort[1009]: INF: - Real execution time: 132
secs.
Feb 08 14:35:46 server1 fort[1009]: INF: Starting validation.
Feb 08 14:35:46 server1 fort[1009]: INF: - Current serial number is
0.
Feb 08 14:37:58 server1 fort[1009]: INF: Checking if there are new or
modified SLURM files
Feb 08 14:37:58 server1 fort[1009]: INF: Applying configured SLURM
Feb 08 14:37:58 server1 fort[1009]: INF: Validation finished:
Feb 08 14:37:58 server1 fort[1009]: INF: - Valid Prefixes: 4740
Feb 08 14:37:58 server1 fort[1009]: INF: - Valid Router Keys: 0
Feb 08 14:37:58 server1 fort[1009]: INF: - Current serial number is
0.
Feb 08 14:37:58 server1 fort[1009]: INF: - Real execution time:
```

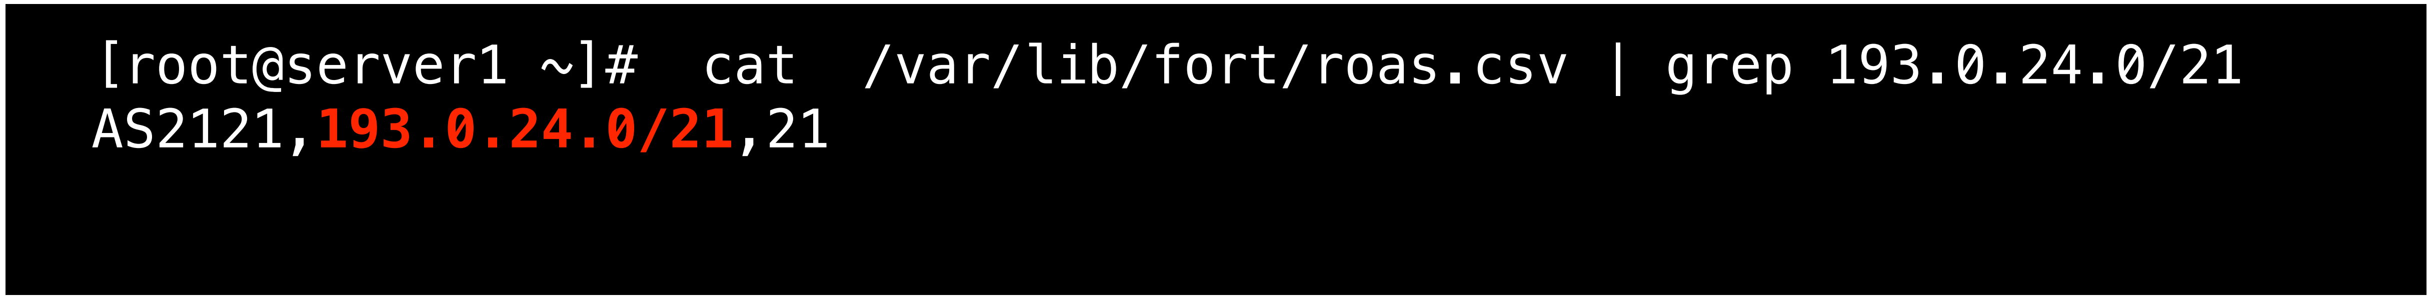

```
[root@server1 ~]#  cat  /var/lib/fort/roas.csv | grep 193.0.24.0/21
AS2121,193.0.24.0/21,21
```

Questions ?

Validation

Section 4

Validation

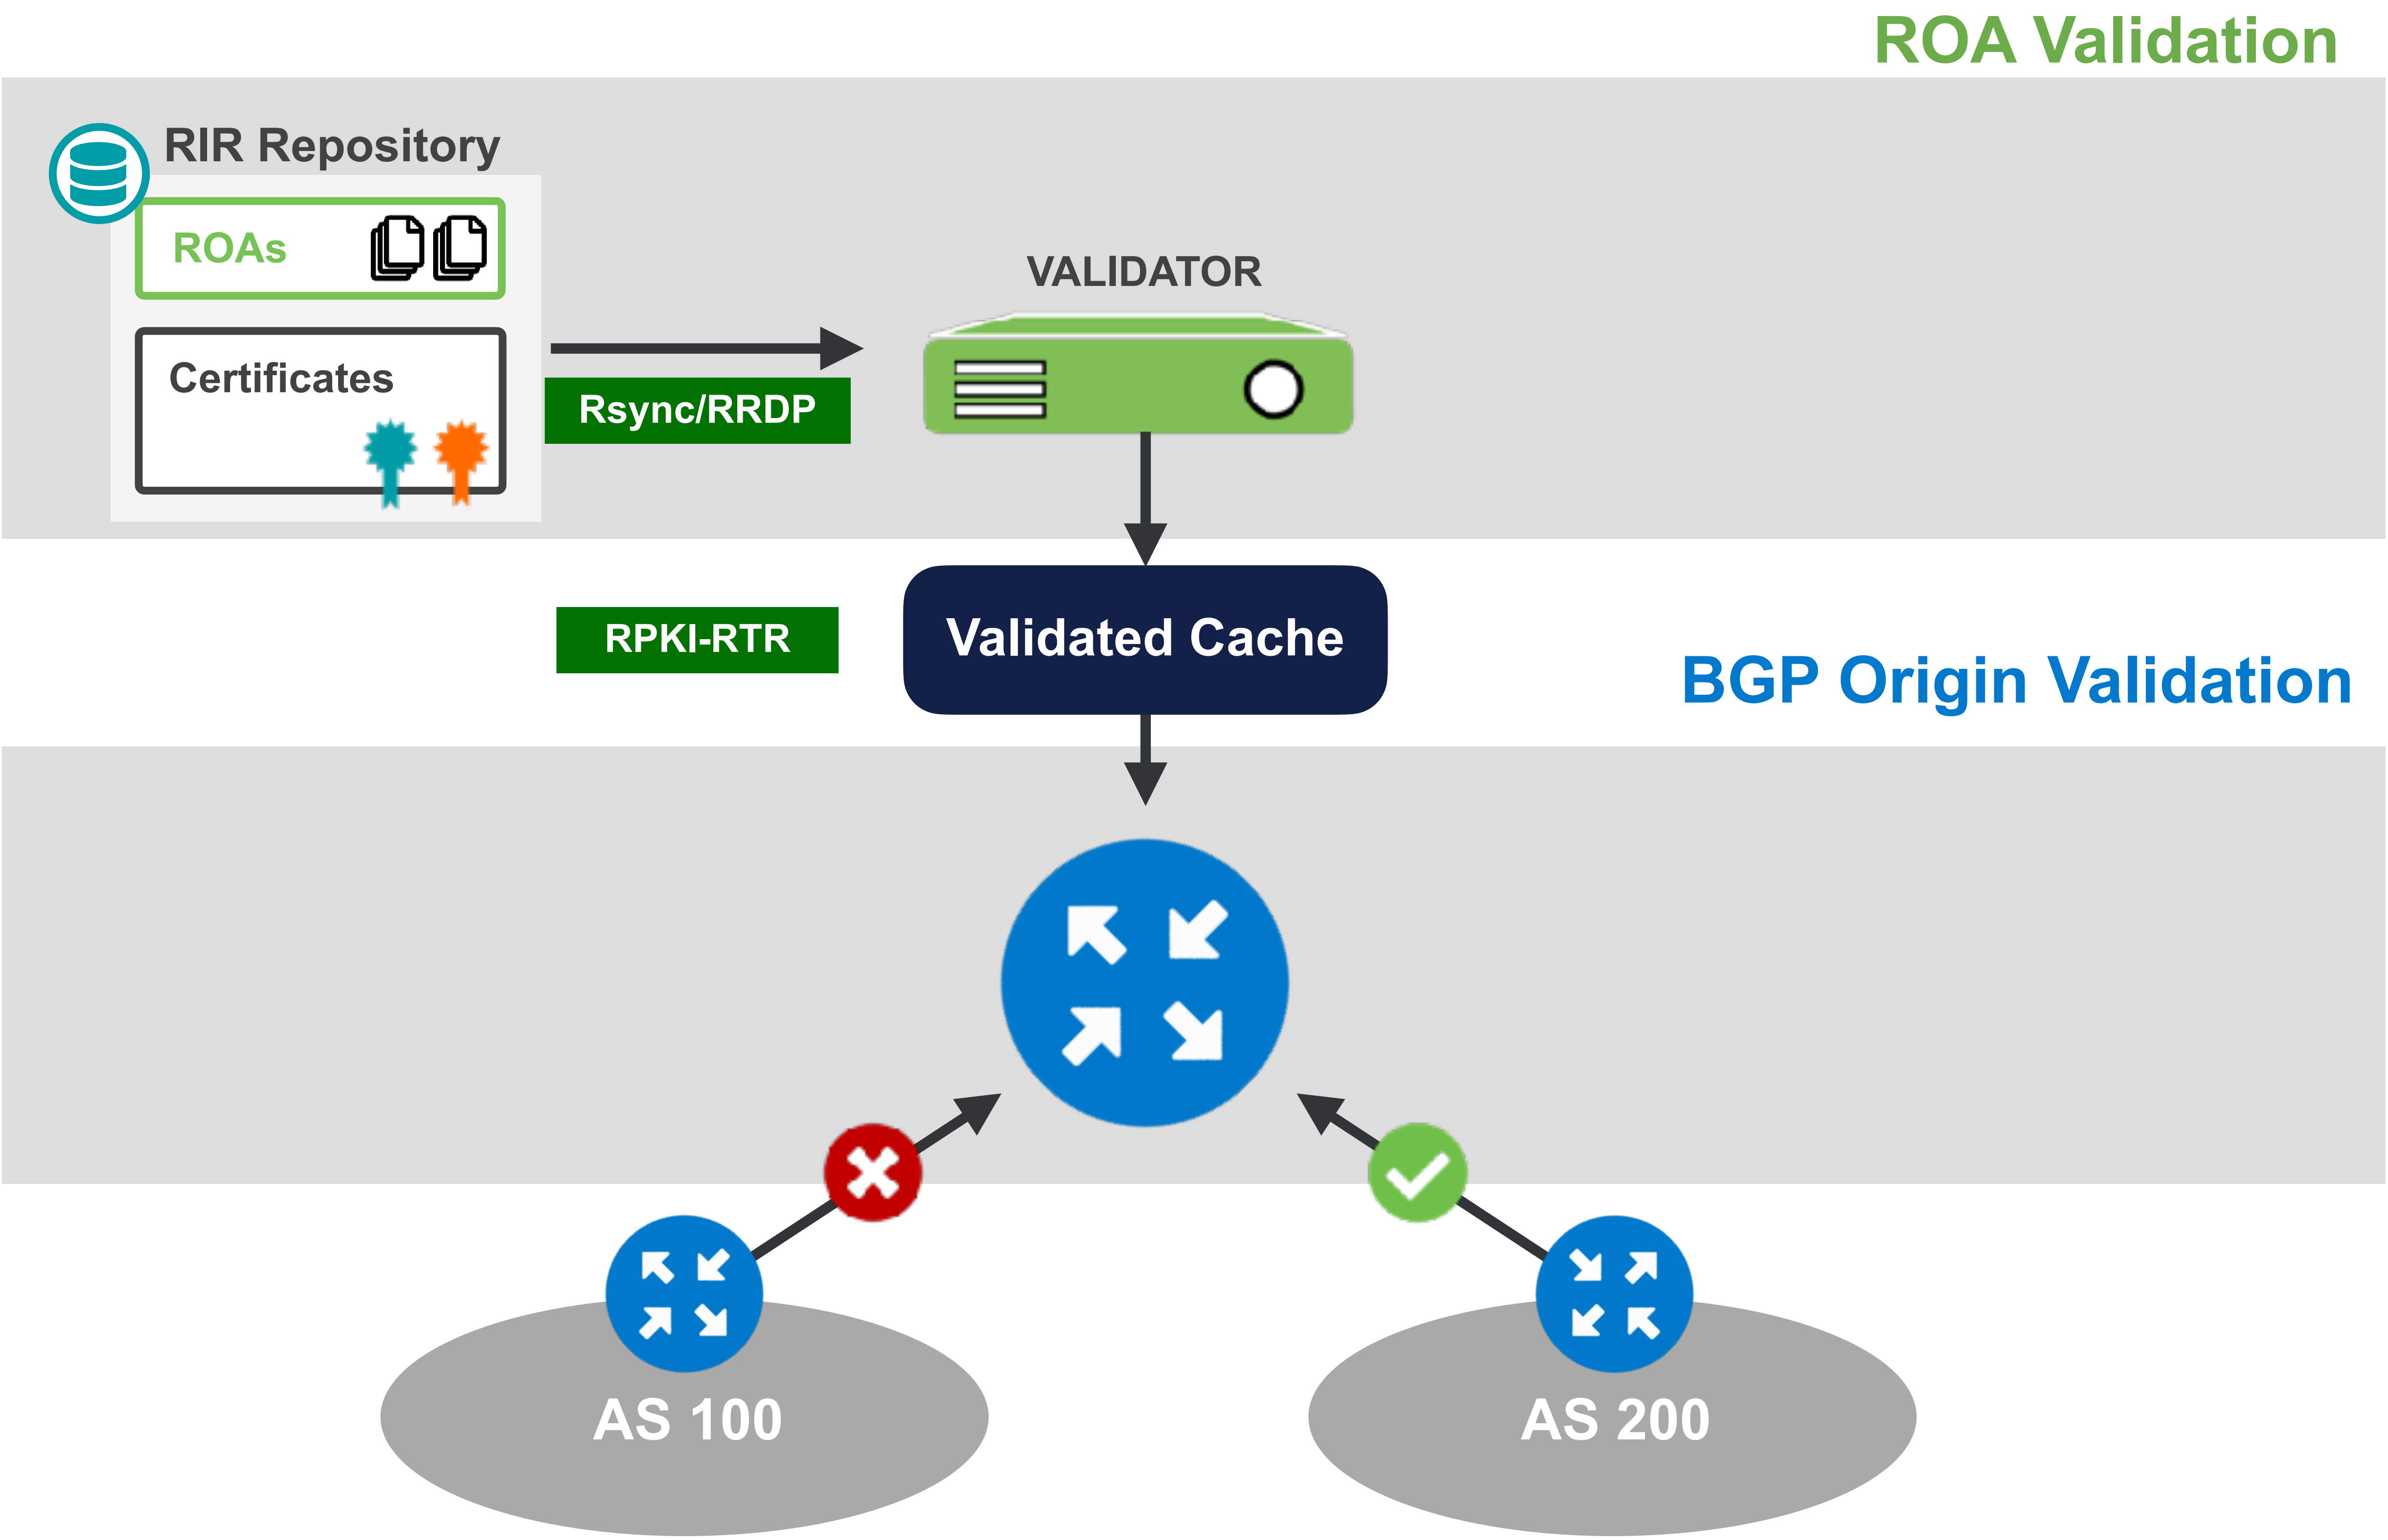
ROA Validation
RIR Repository
ROAs
Certificates
VALIDATOR
Rsync/RRDP
Validated Cache
RPKI-RTR
BGP Origin Validation
AS 100
AS 200

ROA Validation

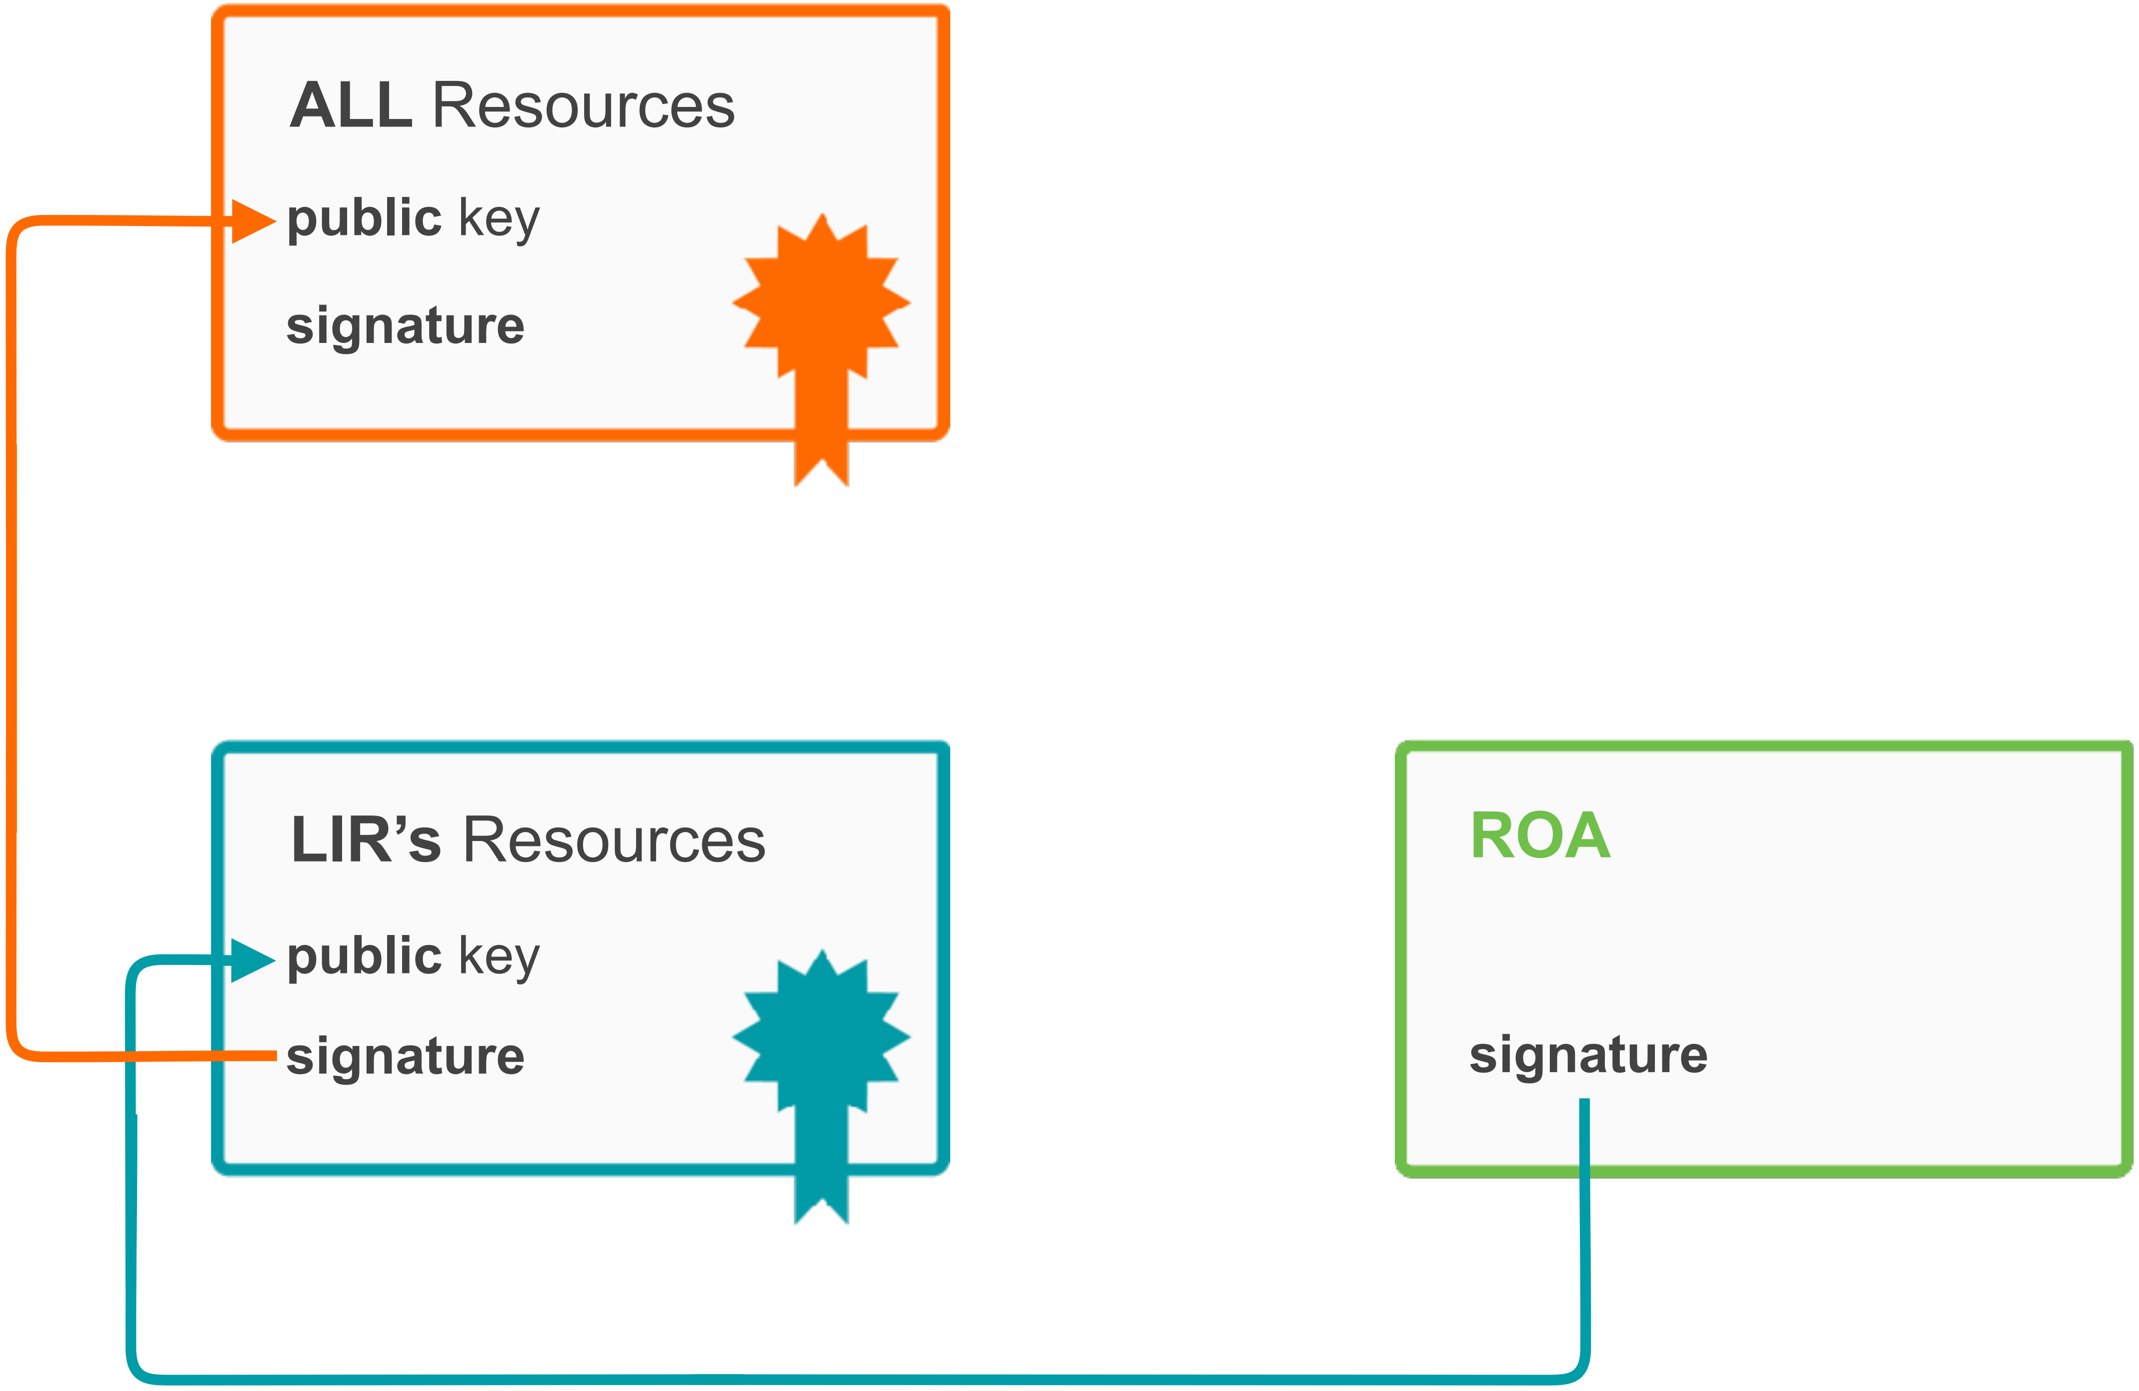

ALL Resources

public key

signature

LIR's Resources

public key

signature

ROA

signature

Quiz time!

What does it mean if a ROA is **“invalid”**?

A. There is no ROA for that specific prefix

B. Validity period of the LIR certificate expired

C. A ROA exists for the prefix but max-length or ASN does not match.

D. Chain of trust fails and the ROA can not be validated.

1 min.

BGP Prefix Origin Validation-RFC6811

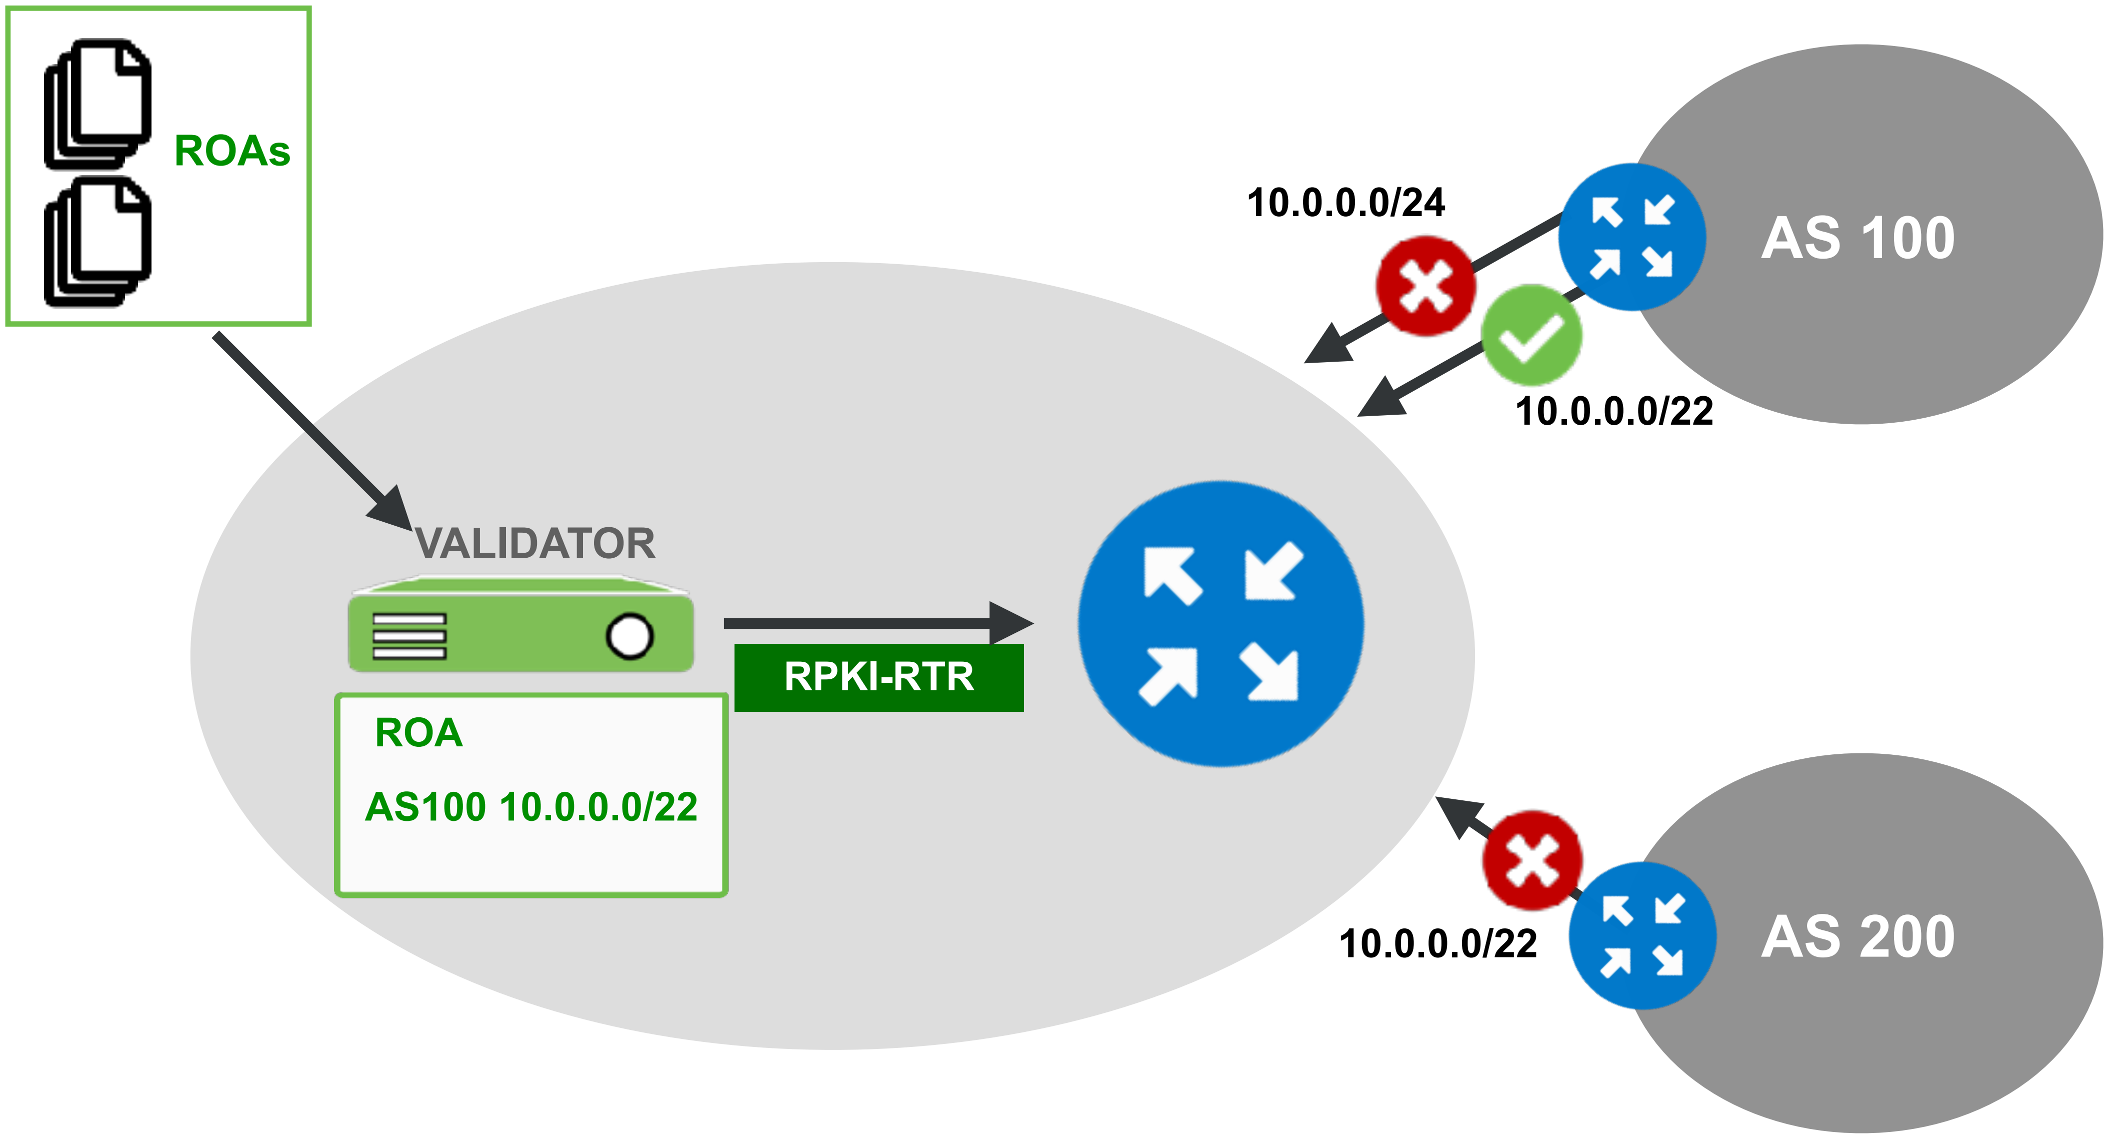

RPKI Validation States

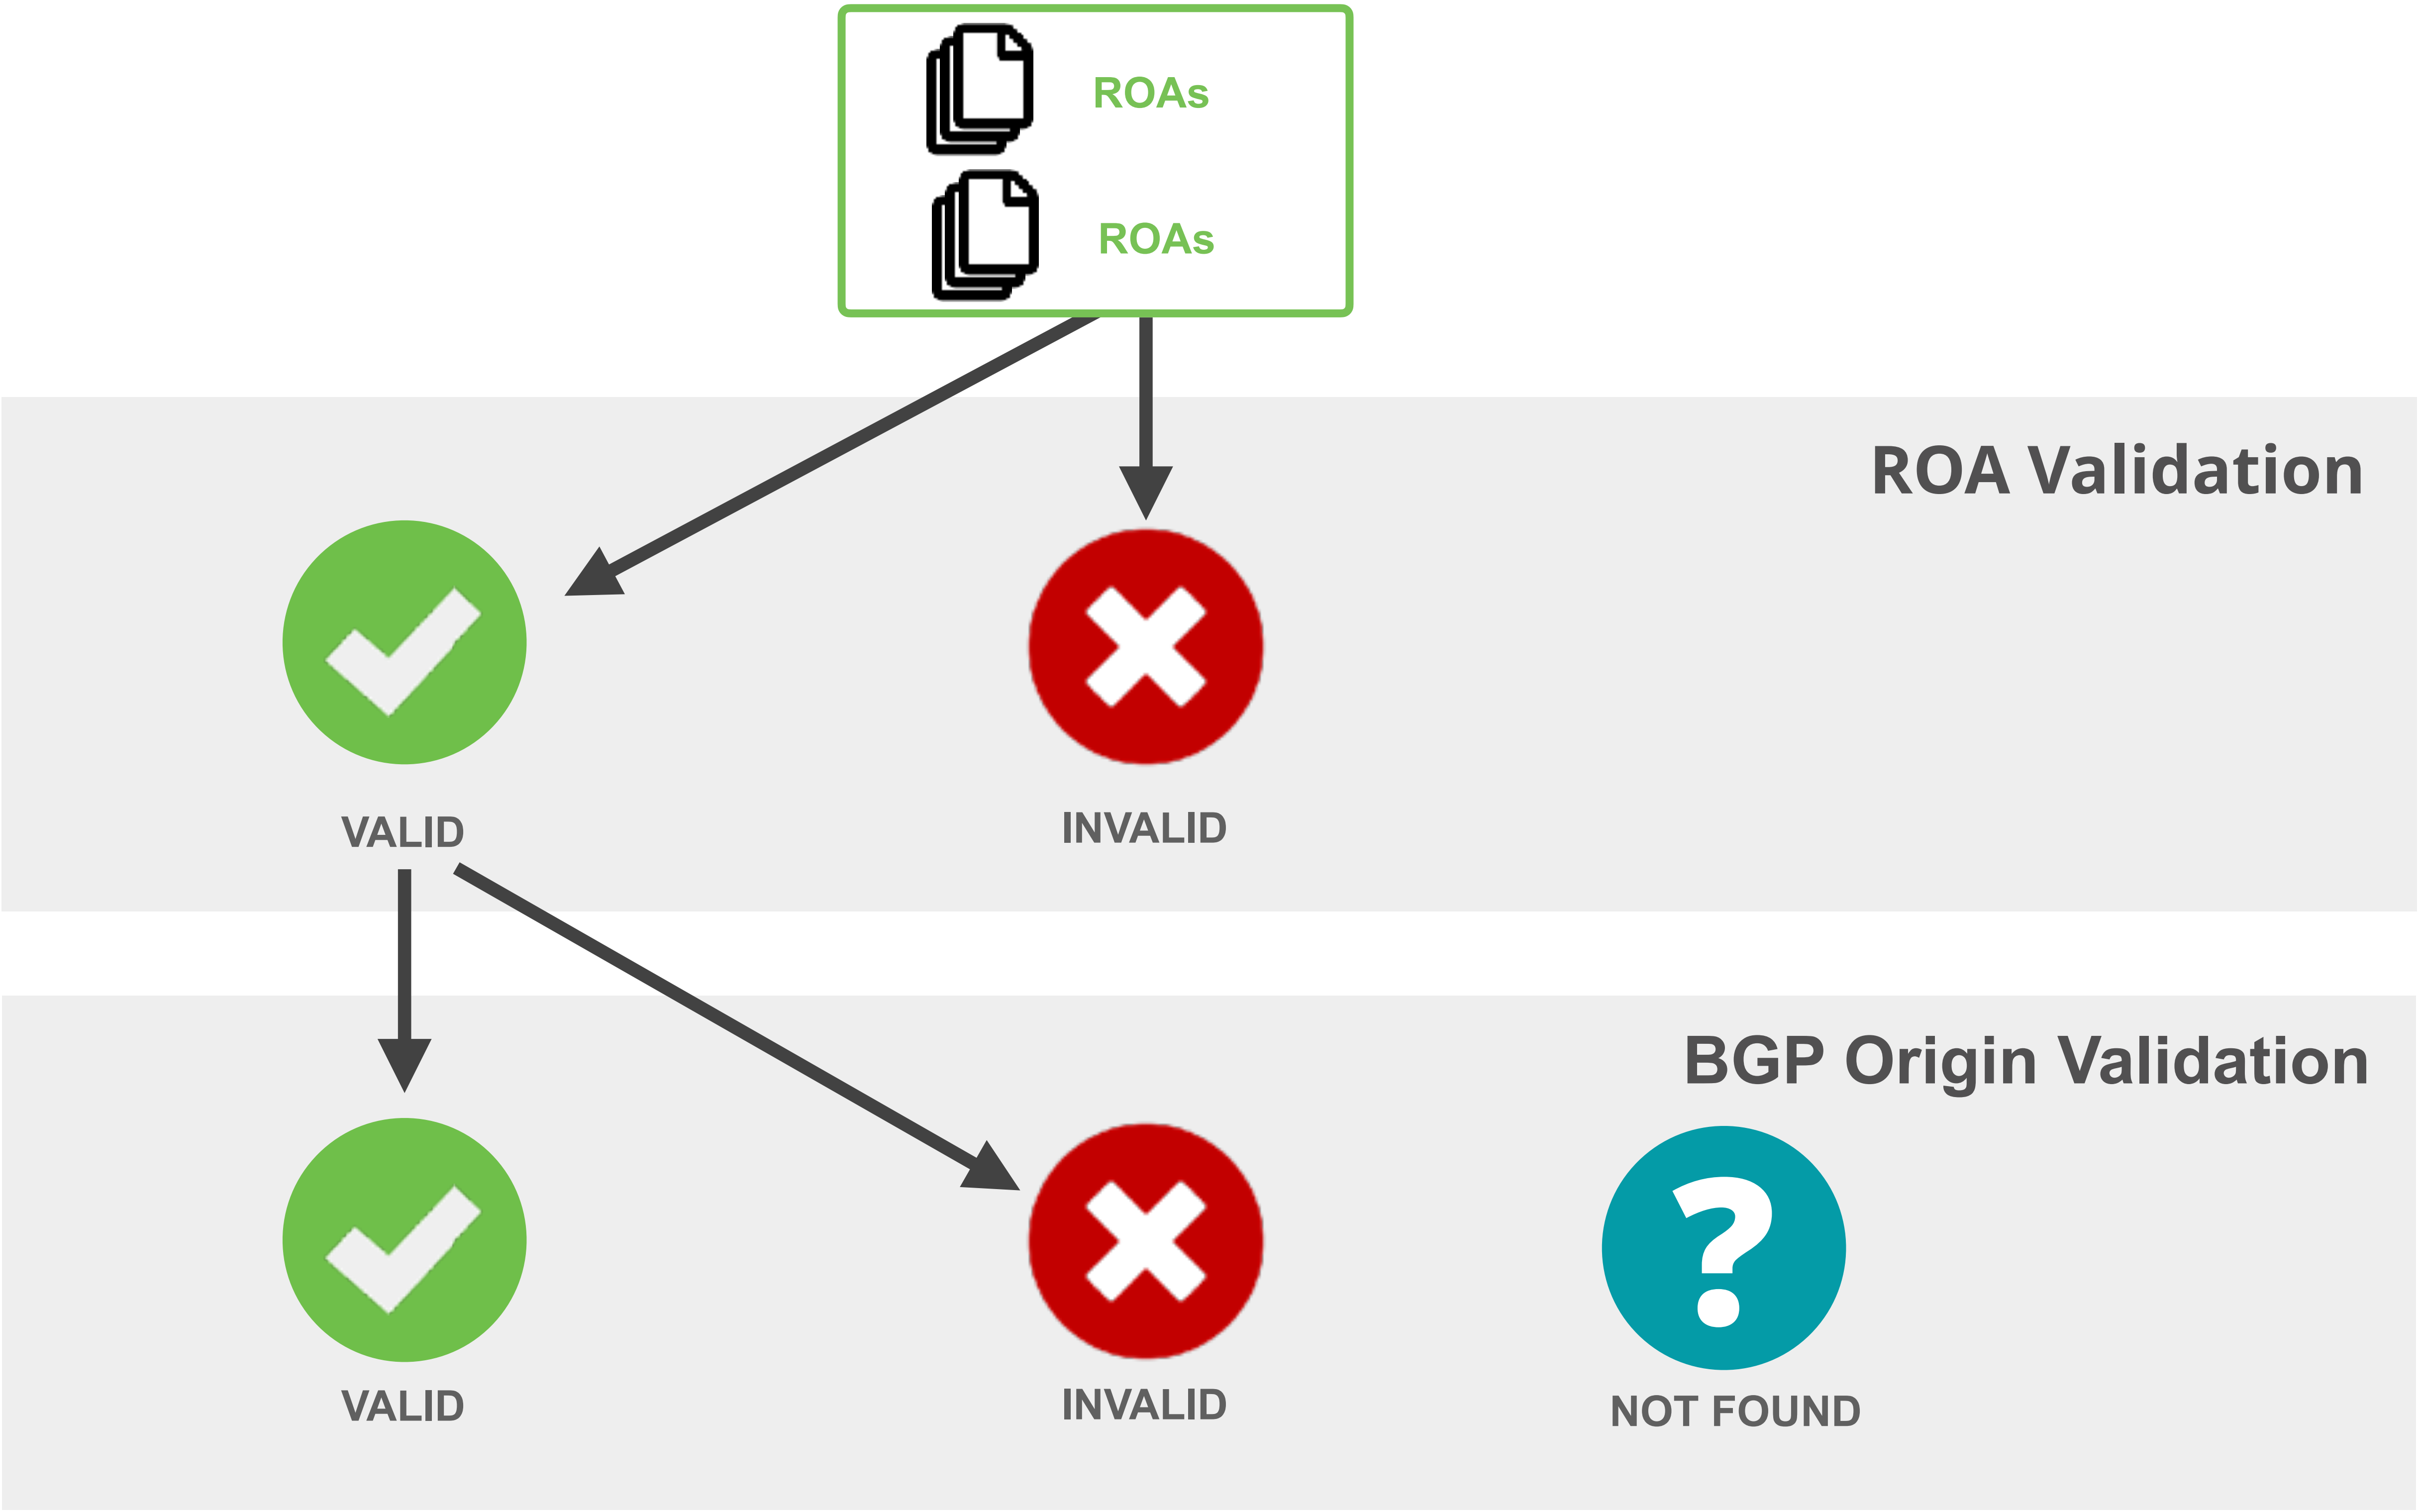

Take the poll!

The RPKI status of a specific prefix in the BGP table is shown as **“Invalid”**.

What does this mean?

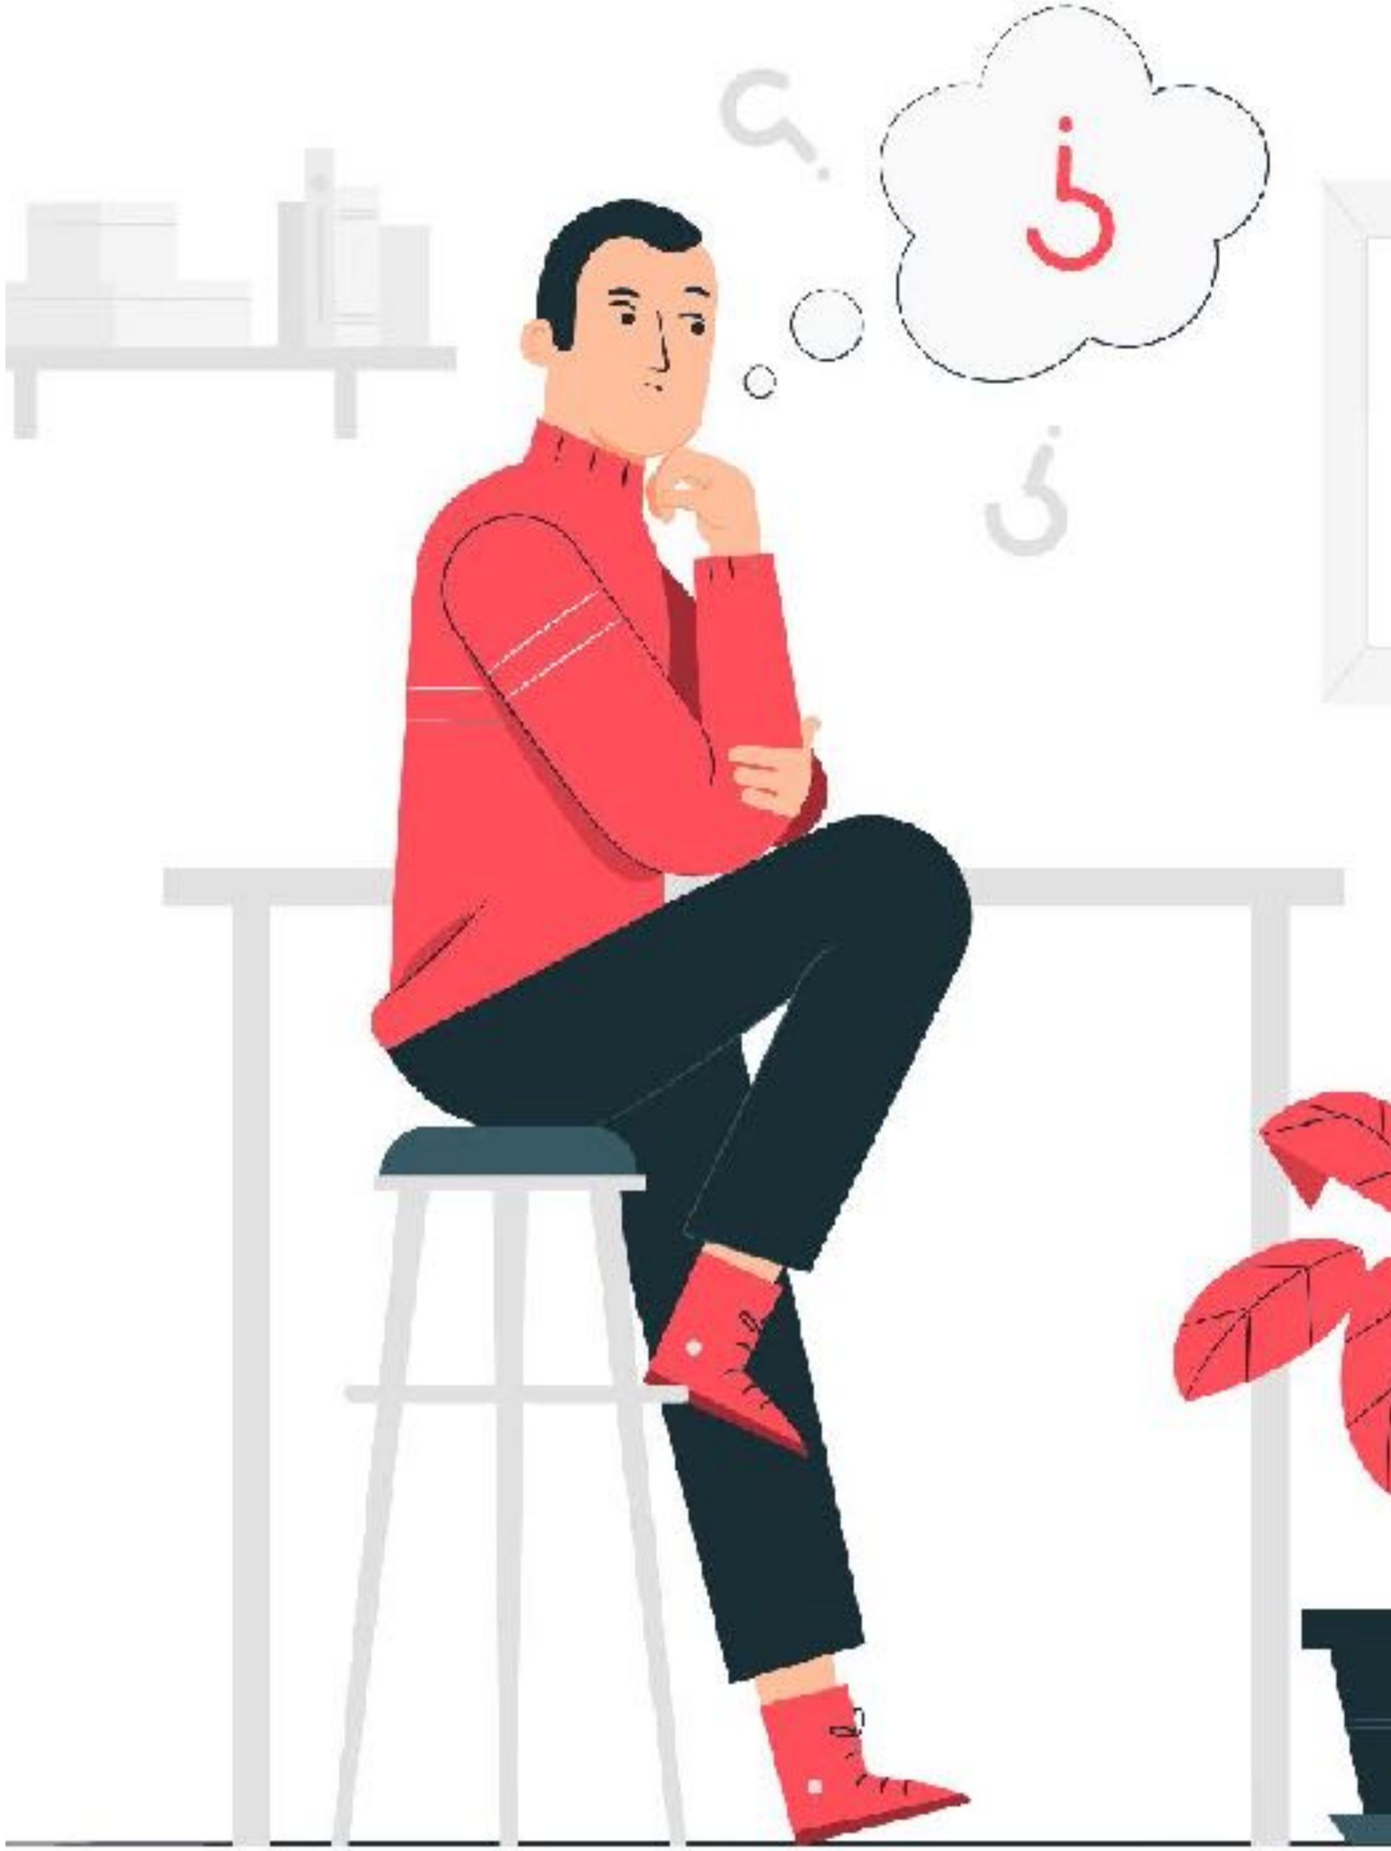

1 min.

Whitelisting

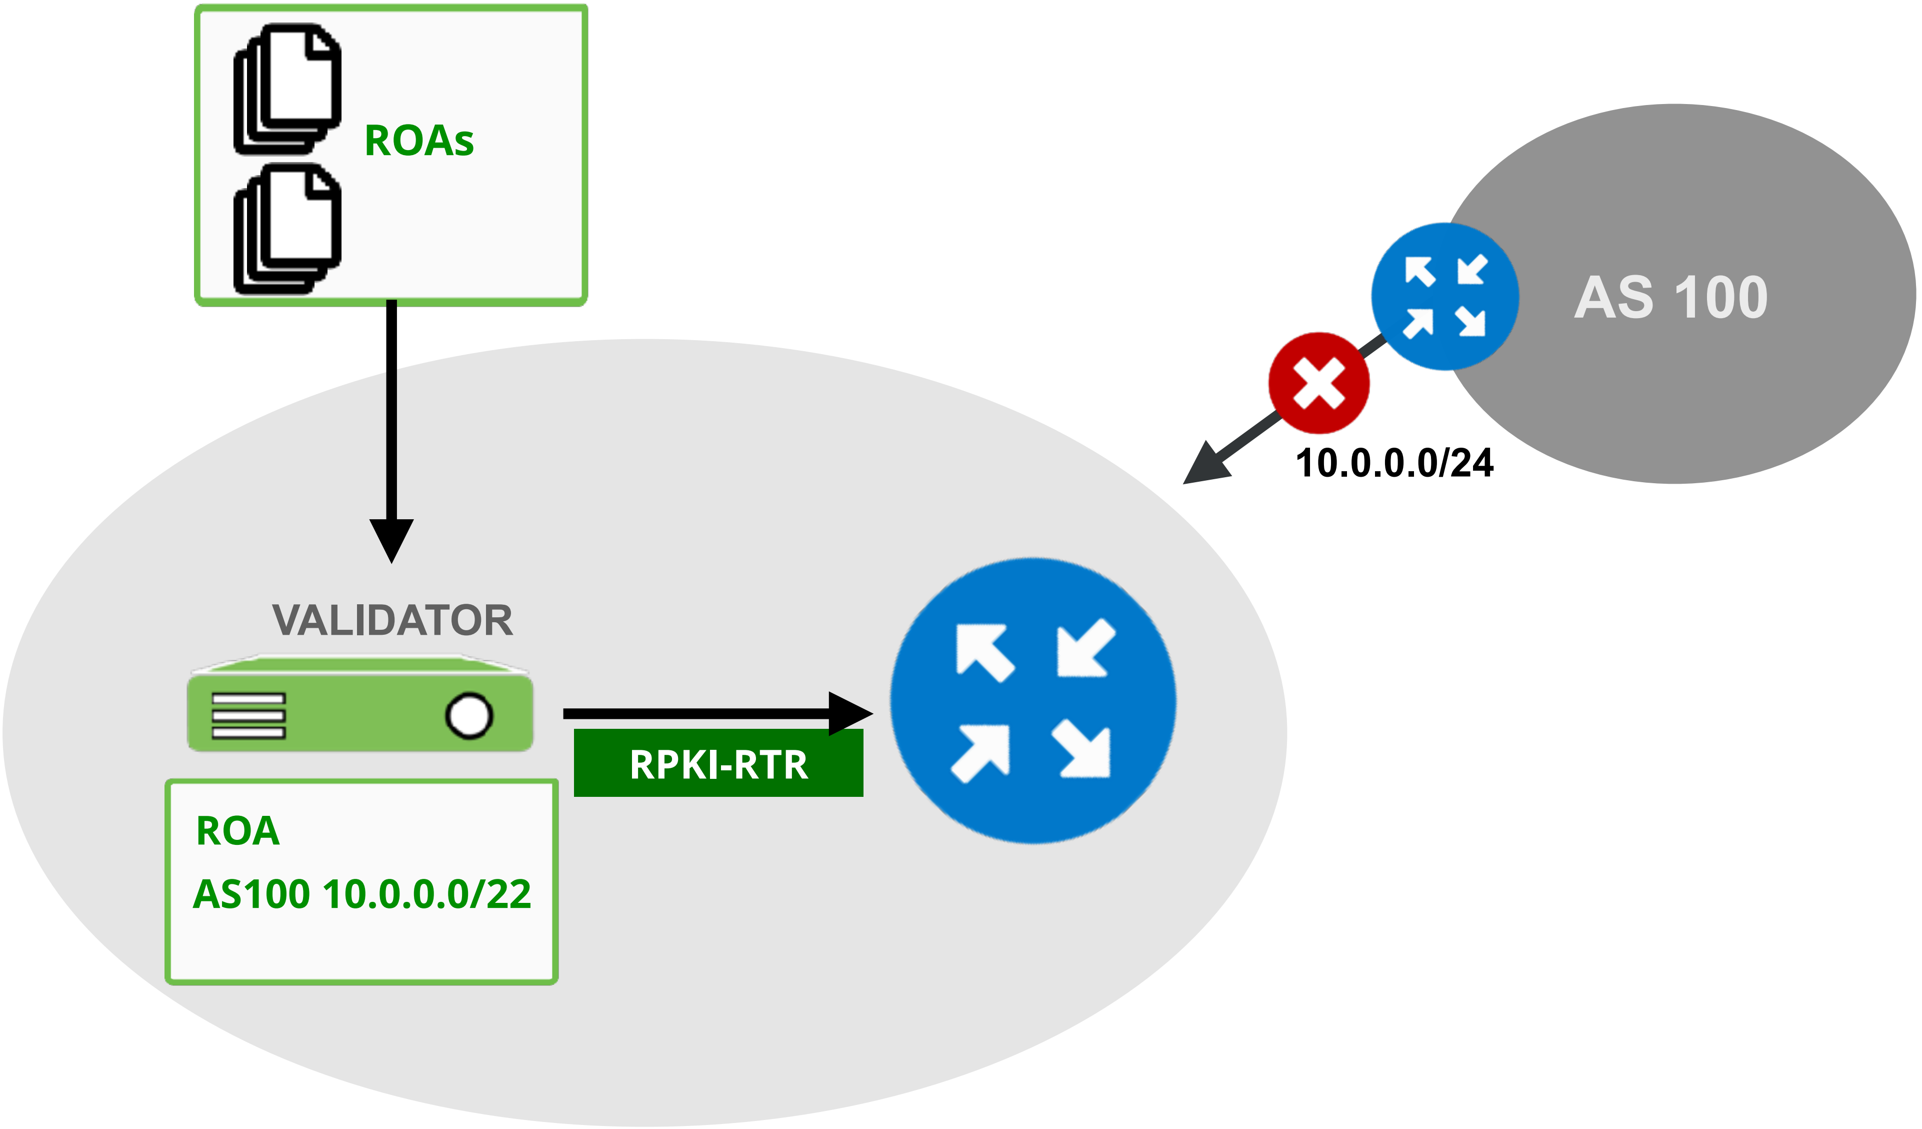

Whitelisting

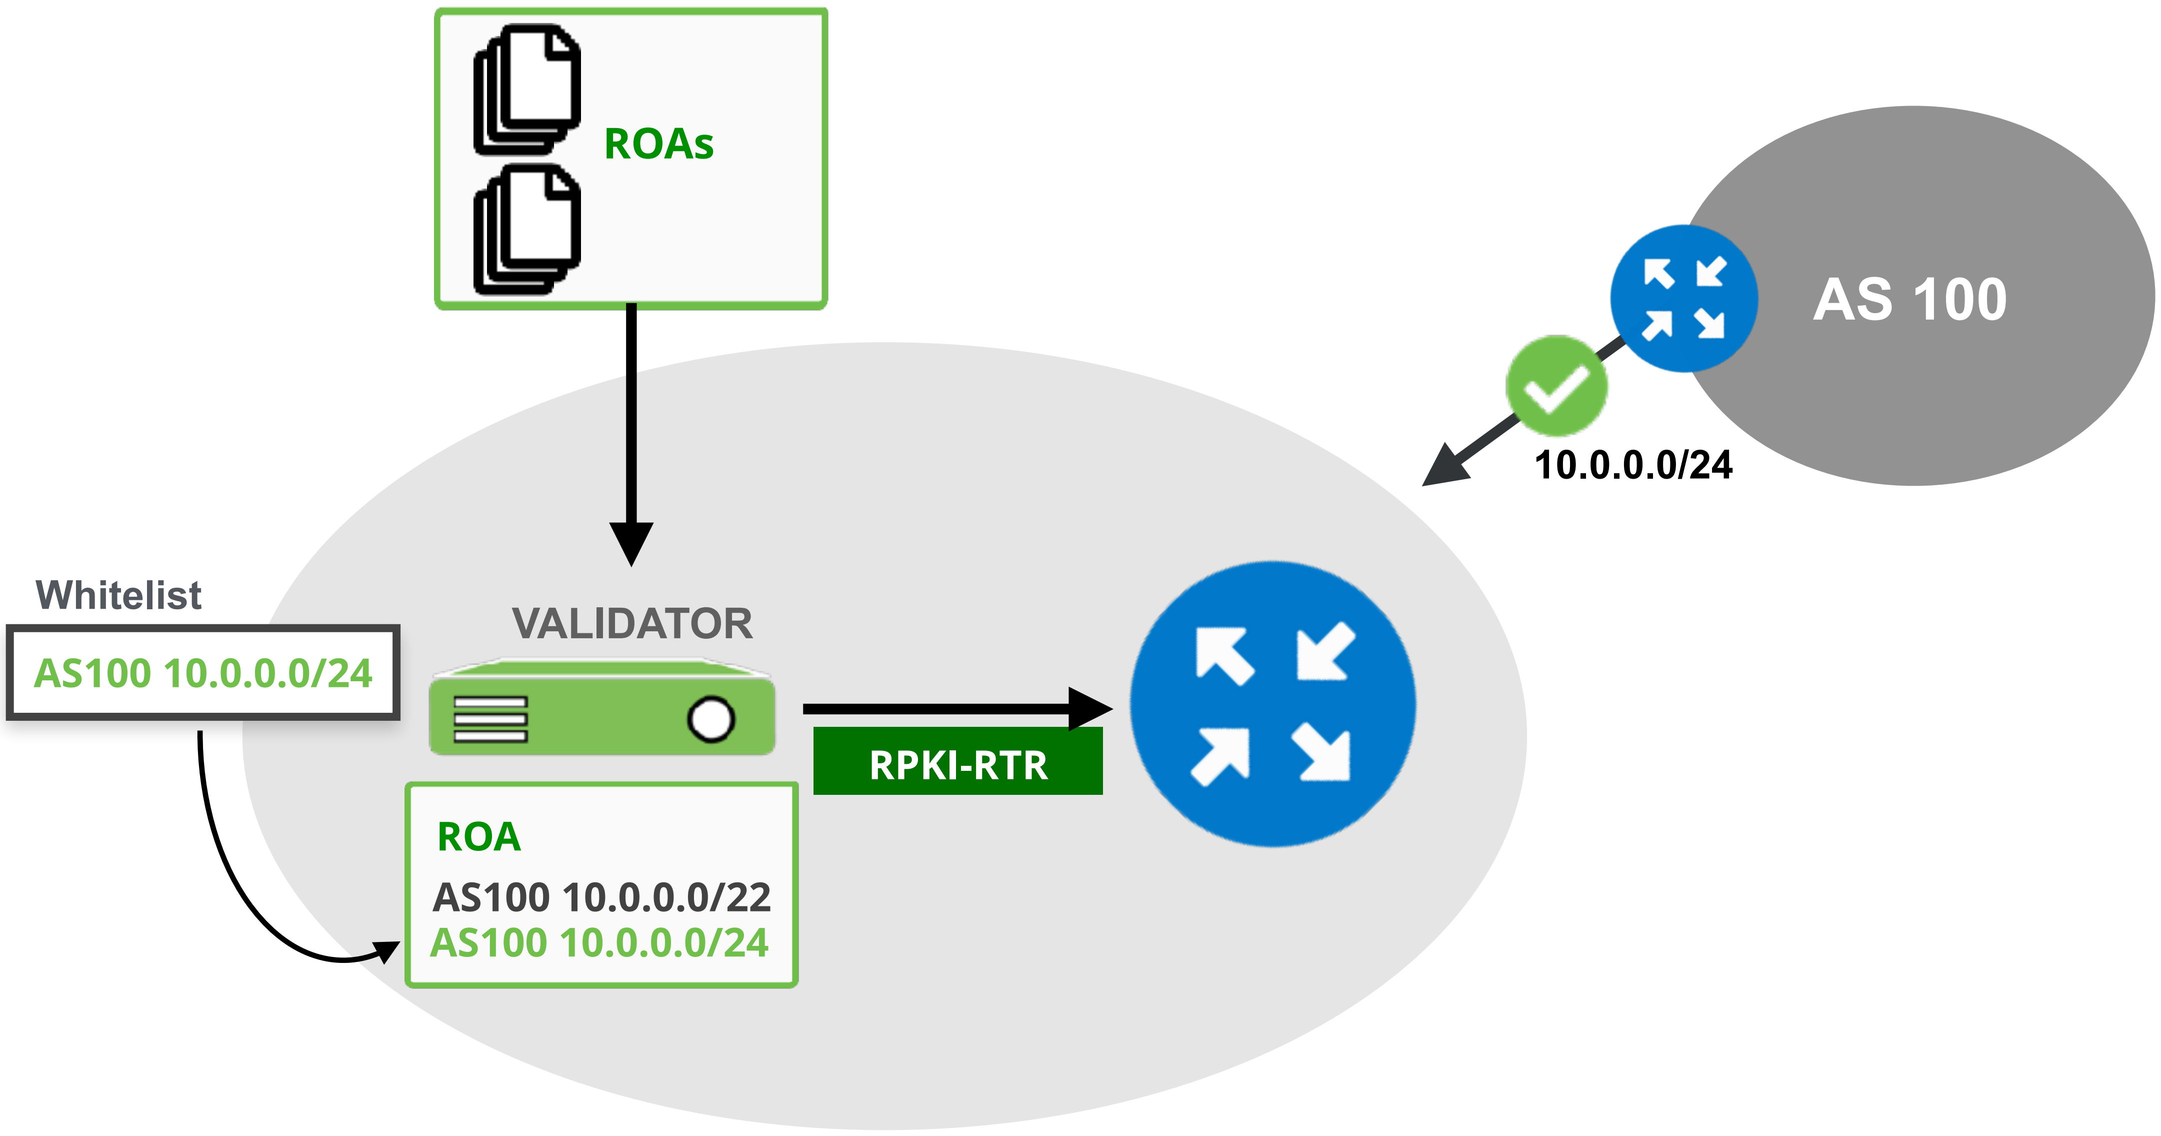

RPKI Status RIPE

Validation results for unique Prefix-Origin pairs in Region RIPE (IPv4)

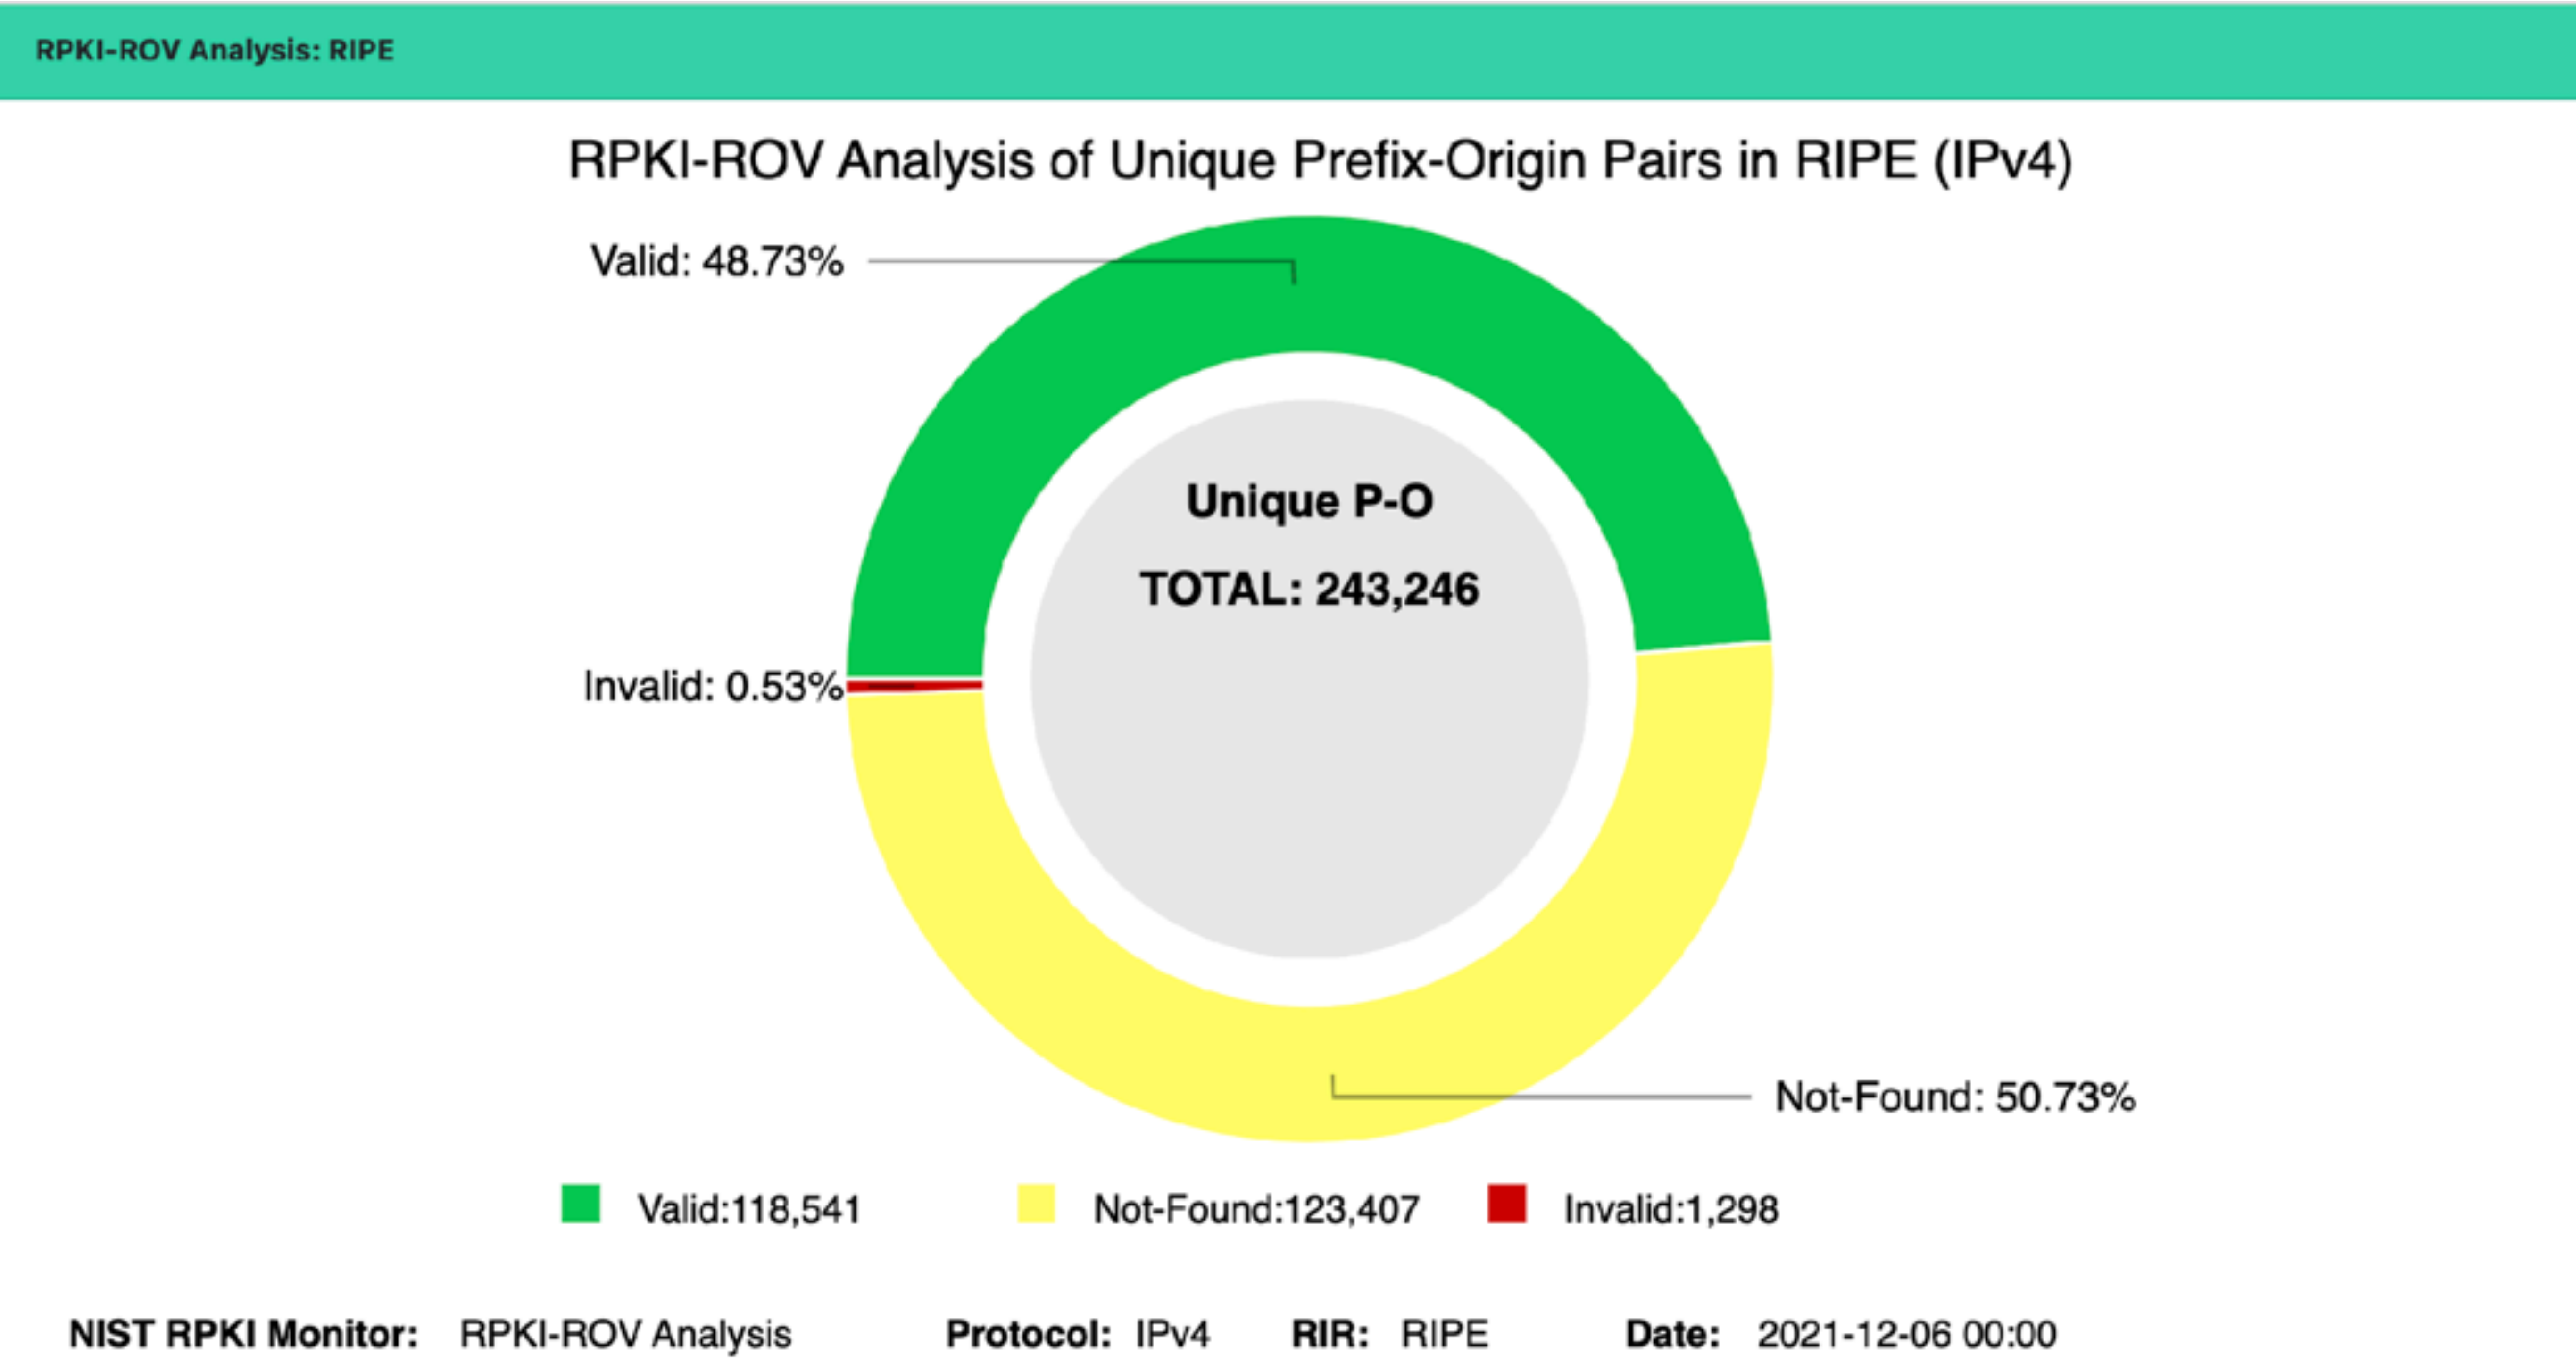

RPKI-ROV Analysis: RIPE

RPKI-ROV Analysis of Unique Prefix-Origin Pairs in RIPE (IPv4)

NIST RPKI Monitor: RPKI-ROV Analysis **Protocol:** IPv4 **RIR:** RIPE **Date:** 2021-12-06 00:00

Demo!

Setting up BGP Origin Validation

Demo Setup

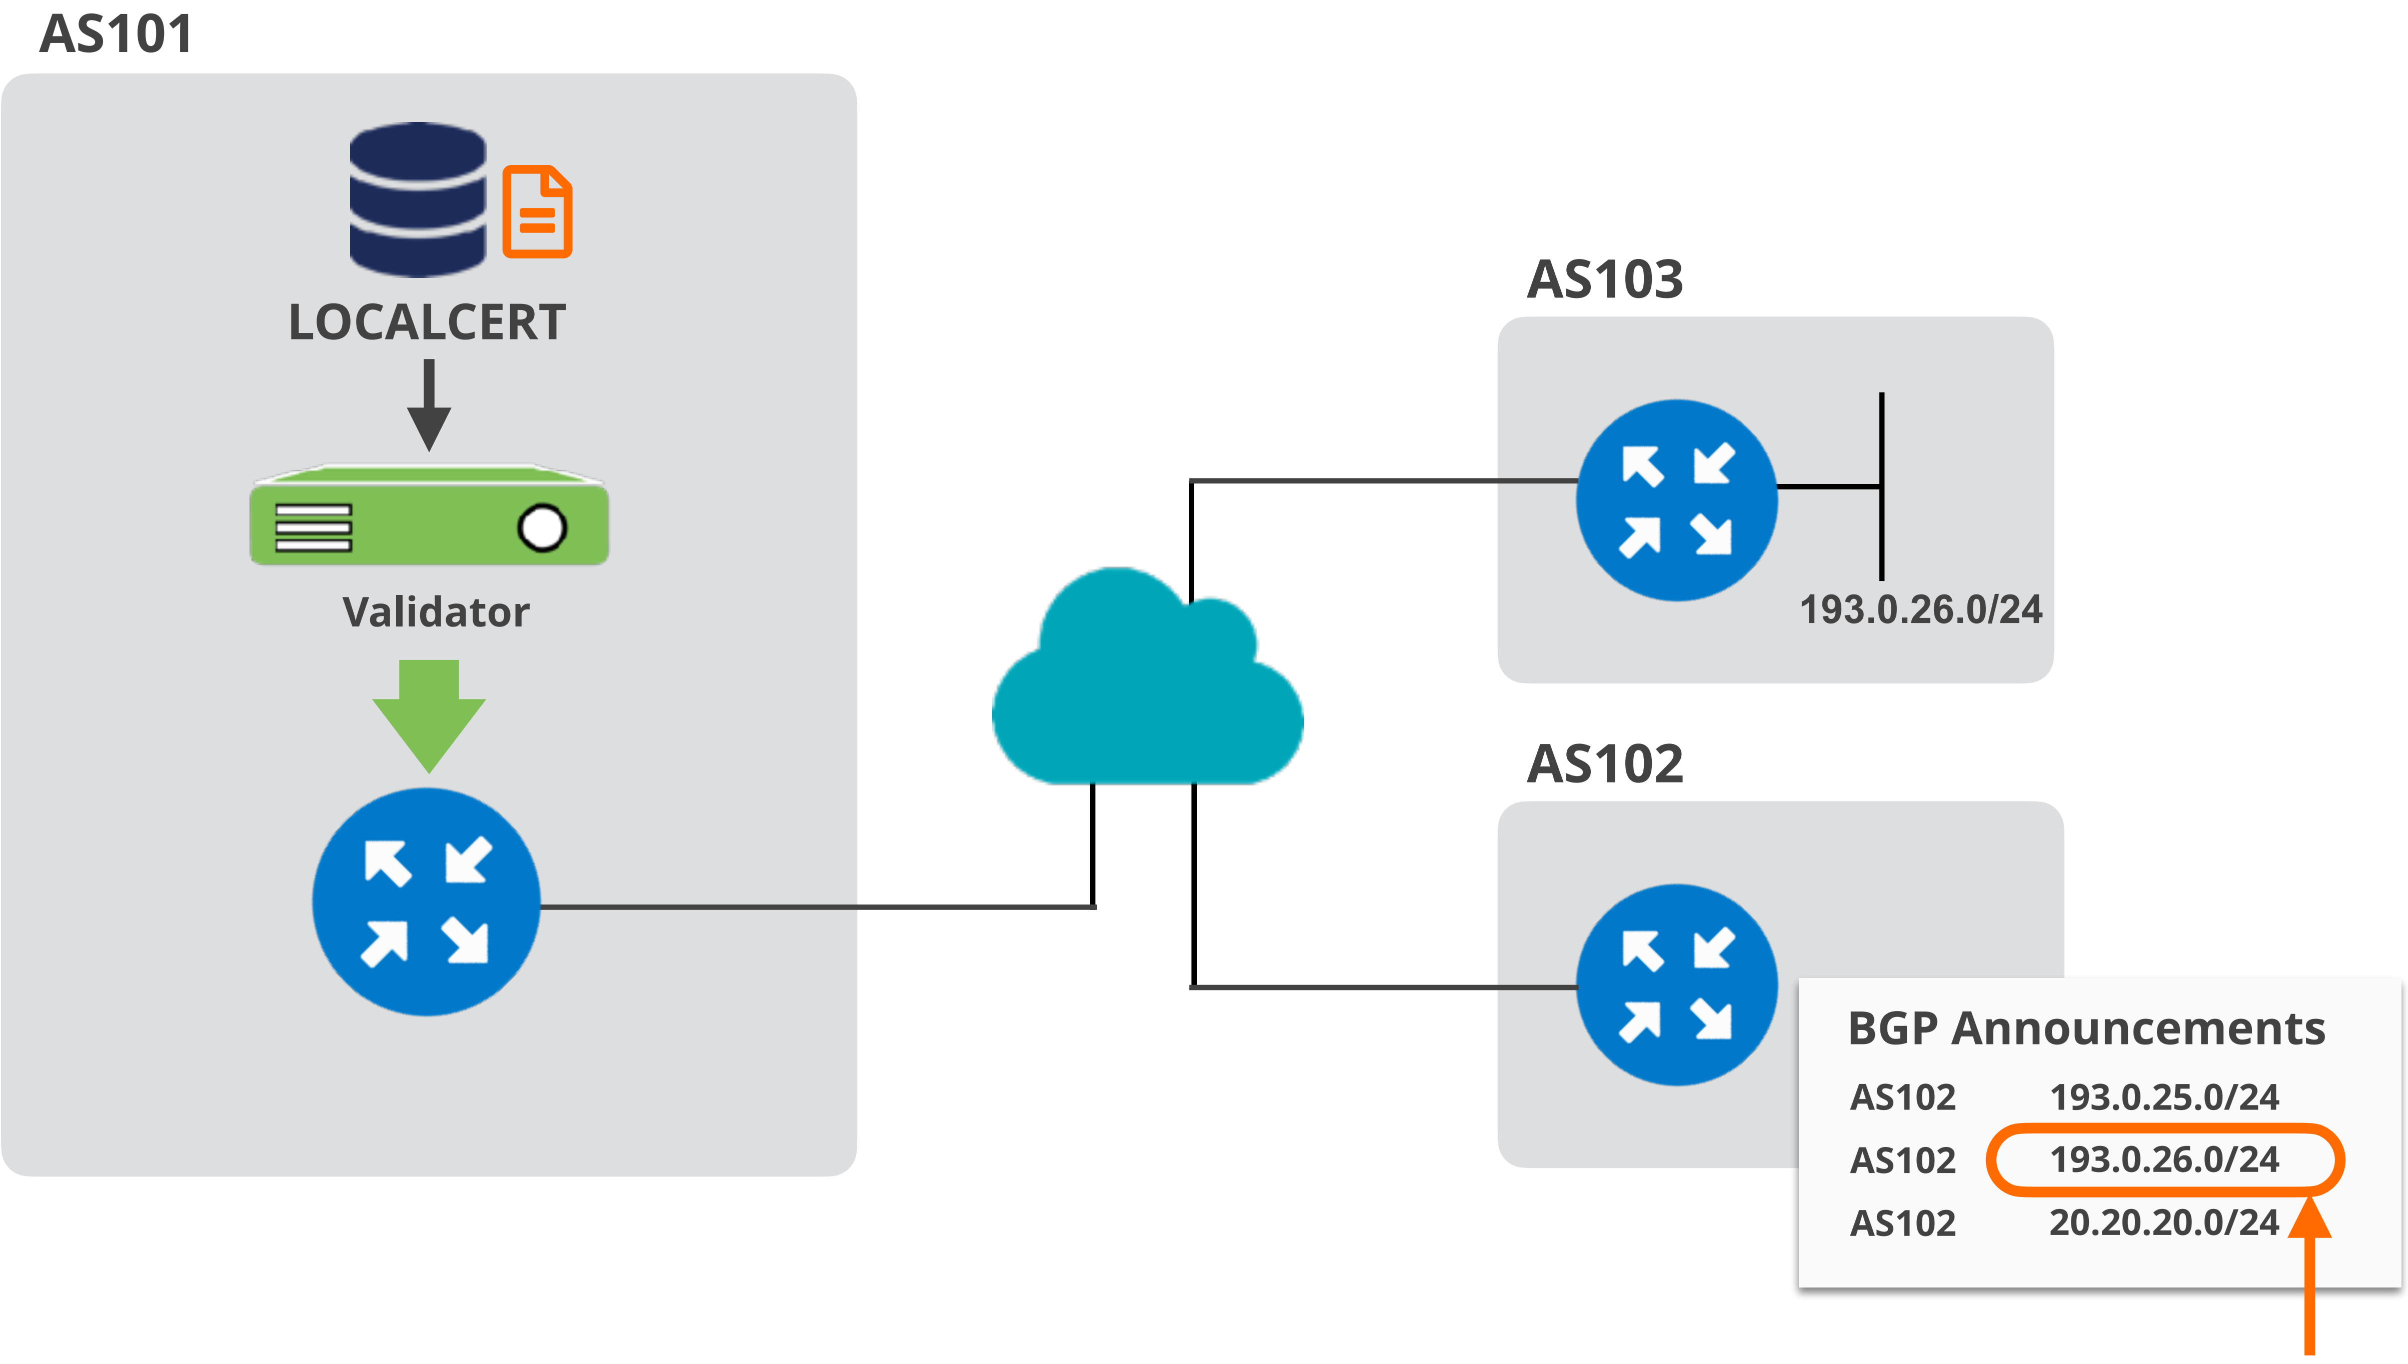

AS101

LOCALCERT

Validator

AS103

193.0.26.0/24

AS102

BGP Announcements

AS102	193.0.25.0/24
AS102	193.0.26.0/24
AS102	20.20.20.0/24

Prefix belongs to AS103

Setup Origin Validation in AS101

- We are using **FORT** and **Routinator** validator options
- Validators are preconfigured
- RPKI-RTR needs to be configured on **AS101 router**
- **AS102 router** will be configured to announce both its networks and **AS103 prefixes**

ROAs Created in the First Demo

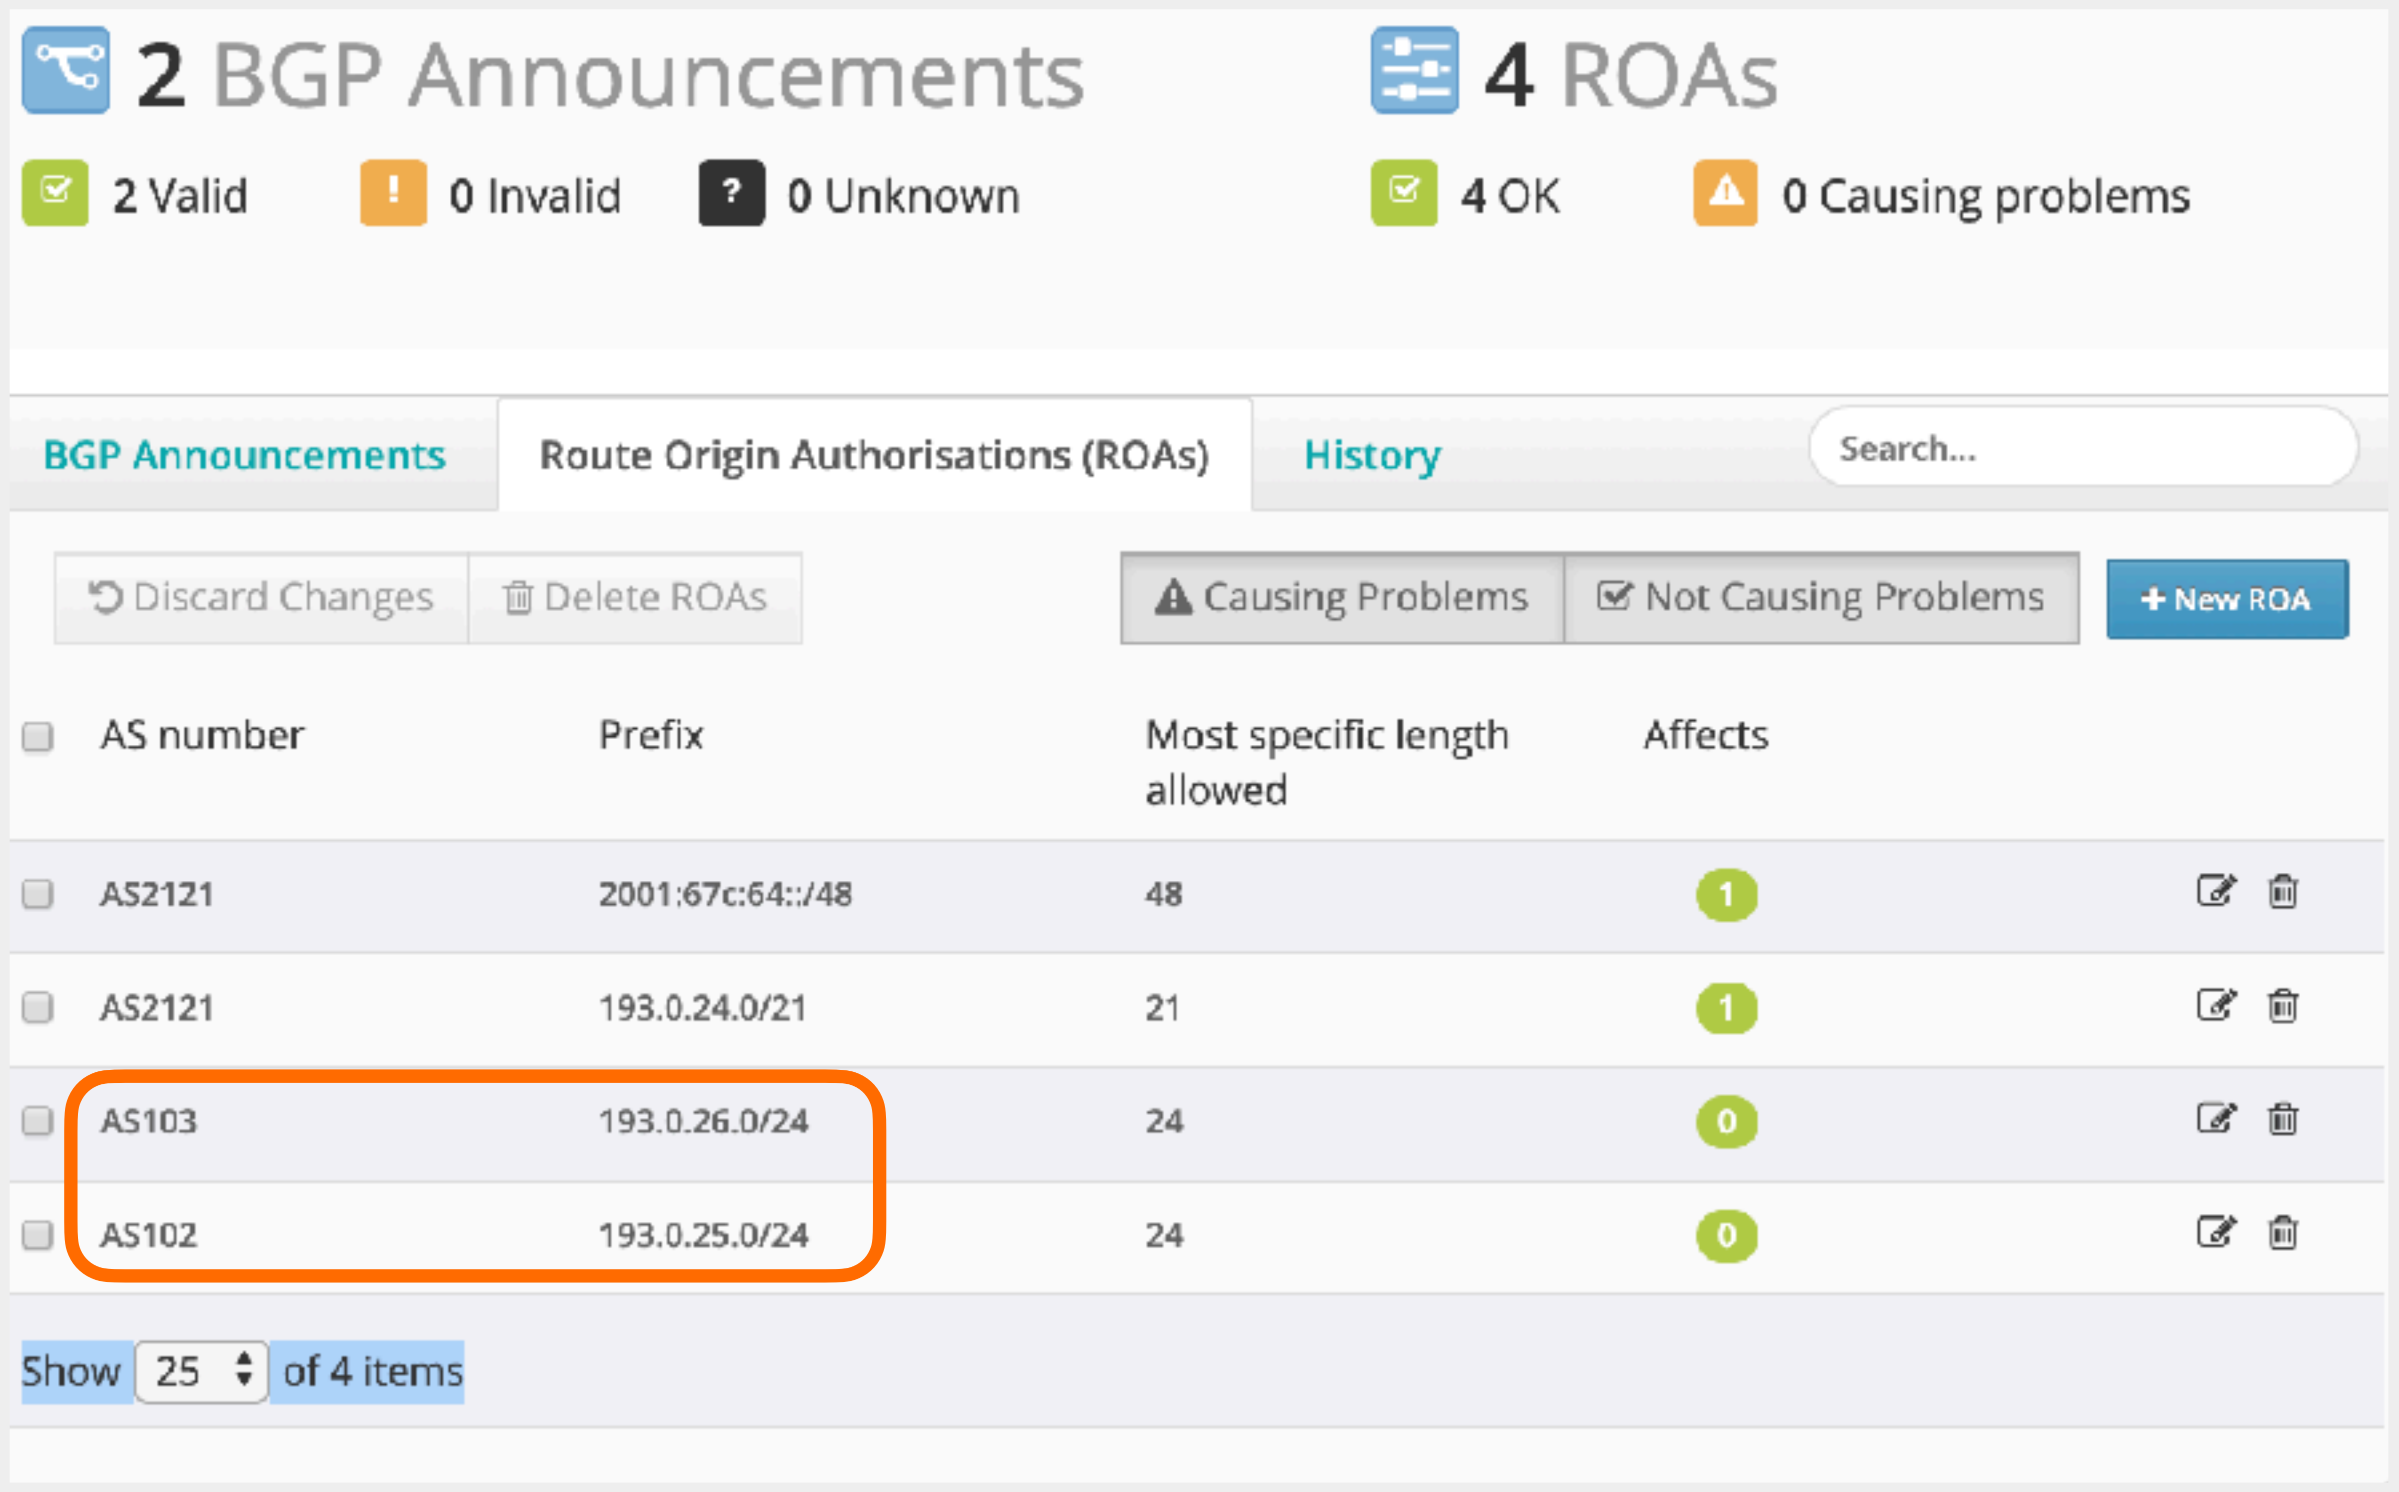

2 BGP Announcements

2 Valid · 0 Invalid · 0 Unknown

4 ROAs

4 OK · 0 Causing problems

□	AS number	Prefix	Most specific length allowed	Affects
□	AS2121	2001:67c:64::/48	48	1
□	AS2121	193.0.24.0/21	21	1
□	AS103	193.0.26.0/24	24	0
□	AS102	193.0.25.0/24	24	0

Show 25 of 4 items

Configure Validator Connection

On AS101 router:

```
(config)# conf t
(config)# router bgp 101
(config-router)# bgp rpki server tcp 100.64.1.1 port 3323 refresh 300
(config-router)# bgp rpki server tcp 100.64.1.1 port 323 refresh 300
```

and check it

```
# show ip bgp rpki servers | i ESTAB
# show ip bgp rpki table
```

Let’s Check How We’re Doing...

```
U1_Router#show ip bgp rpki servers | i ESTAB

Connection state is ESTAB, I/O status: 1, unread input bytes: 0
Connection state is ESTAB, I/O status: 1, unread input bytes: 0
```

```
U1_Router#sho ip bgp rpki table
1547 BGP sovc network entries using 247520 bytes of memory

3851 BGP sovc record entries using 123232 bytes of memory

Network                 Maxlen  Origin-AS  Source  Neighbor

5.32.168.0/21           21      15836      0       100.64.1.1/323
5.32.168.0/21           21      15836      0       100.64.1.1/3323
5.35.224.0/19           24      8972       0       100.64.1.1/323
5.35.224.0/19           24      8972       0       100.64.1.1/3323
5.35.224.0/19           24      29066      0       100.64.1.1/323
5.35.224.0/19           24      29066      0       100.64.1.1/3323
```

FORT

Routinator

Configure BGP announcements

- Let’s configure Router in AS102 to announce prefixes!
- Check origin validation on AS101 router!

```
(config)# router bgp 102
(config-router)# address-family ipv4
(config-router)# network 20.20.20.0 mask 255.255.255.0
(config-router)# network 193.0.25.0
(config-router)# network 193.0.26.0

(config-router)# ip route 20.20.20.0 255.255.255.0 null0
(config-router)# ip route 193.0.25.0 255.255.255.0 null0
(config-router)# ip route 193.0.26.0 255.255.255.0 null0
```

No ROA for this one!

Prefix belongs to AS103!

RPKI Valid

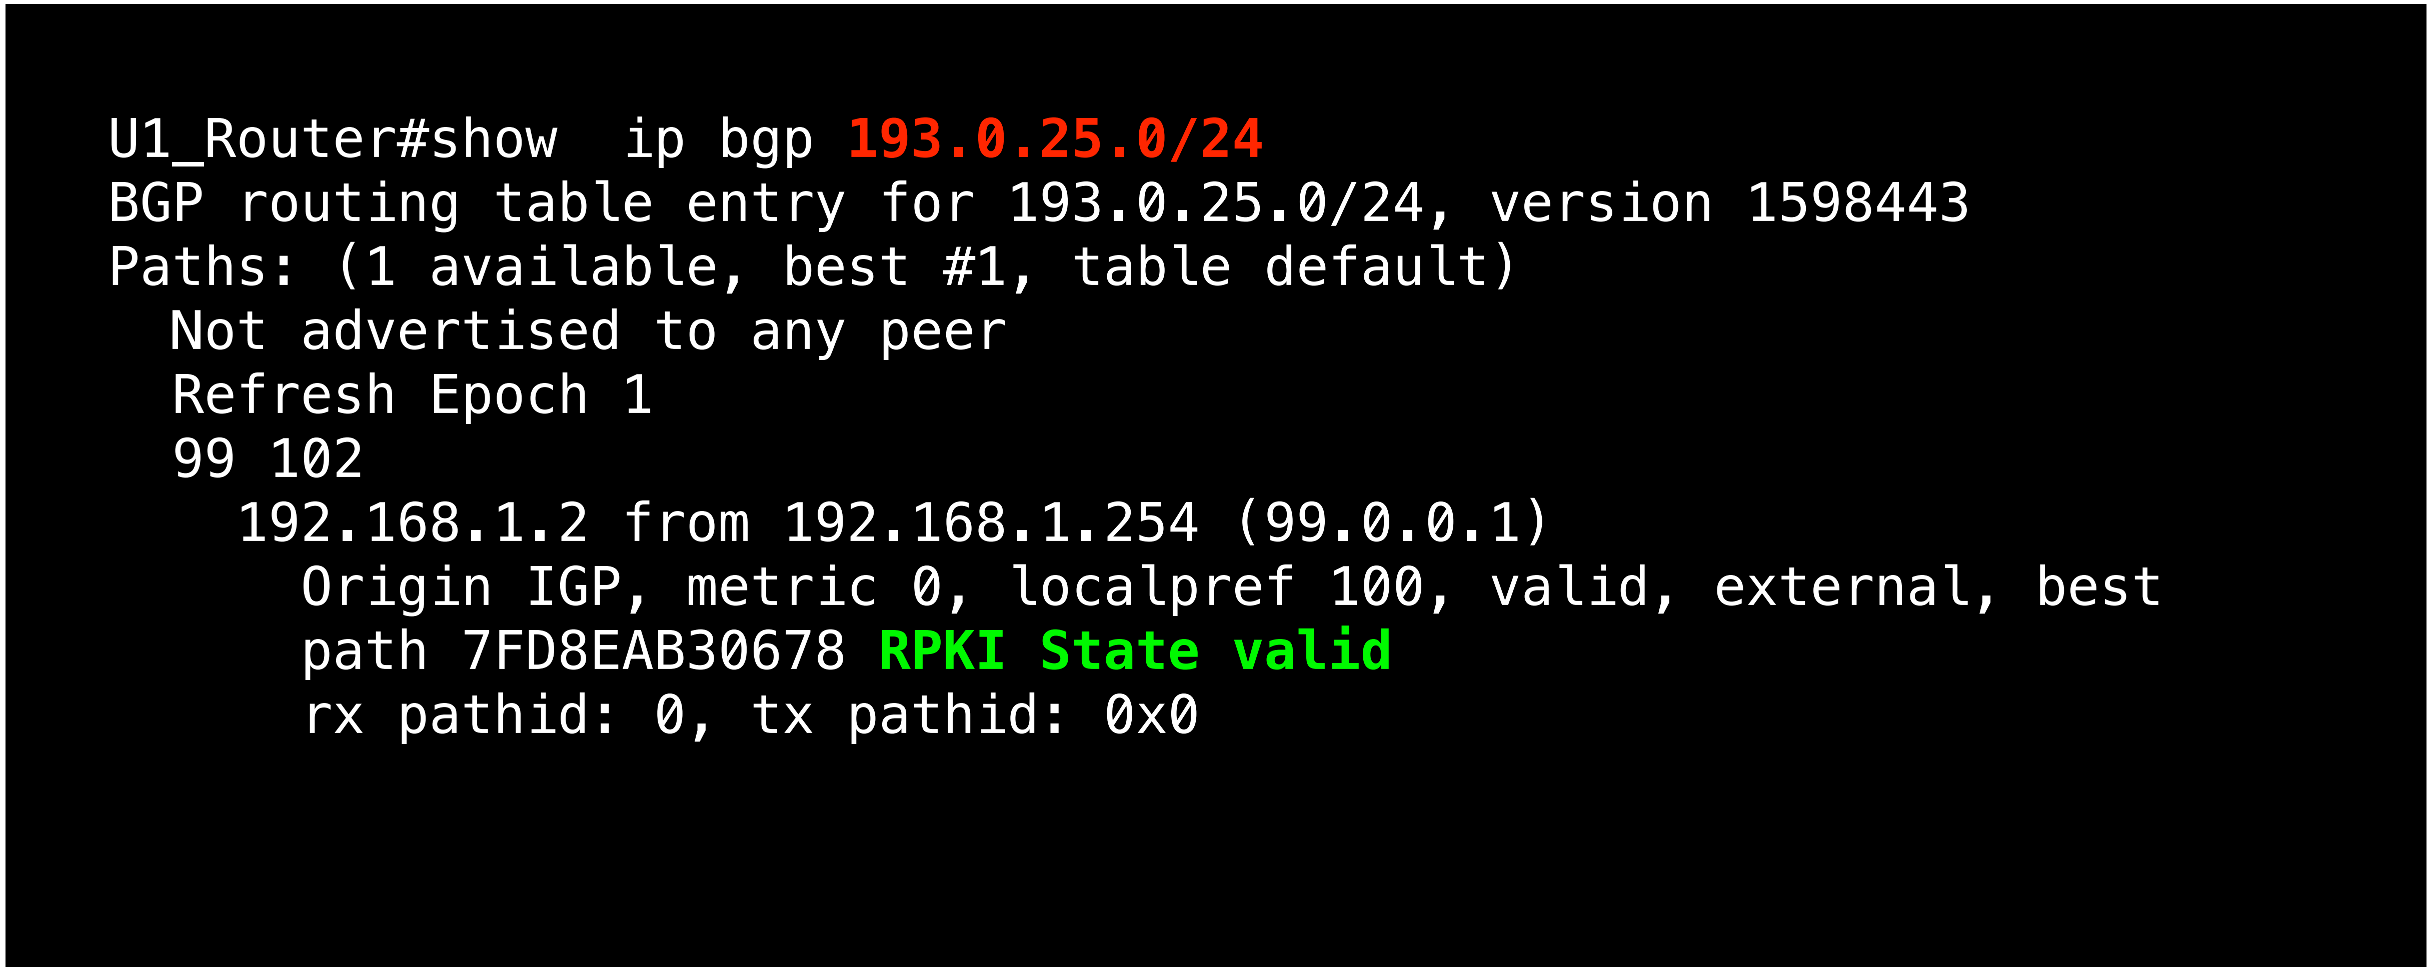

```
U1_Router#show  ip bgp 193.0.25.0/24
BGP routing table entry for 193.0.25.0/24, version 1598443
Paths: (1 available, best #1, table default)
  Not advertised to any peer
  Refresh Epoch 1
  99 102
    192.168.1.2 from 192.168.1.254 (99.0.0.1)
      Origin IGP, metric 0, localpref 100, valid, external, best
      path 7FD8EAB30678 RPKI State valid
      rx pathid: 0, tx pathid: 0x0
```

RPKI Invalid

Prefix belongs to AS103!

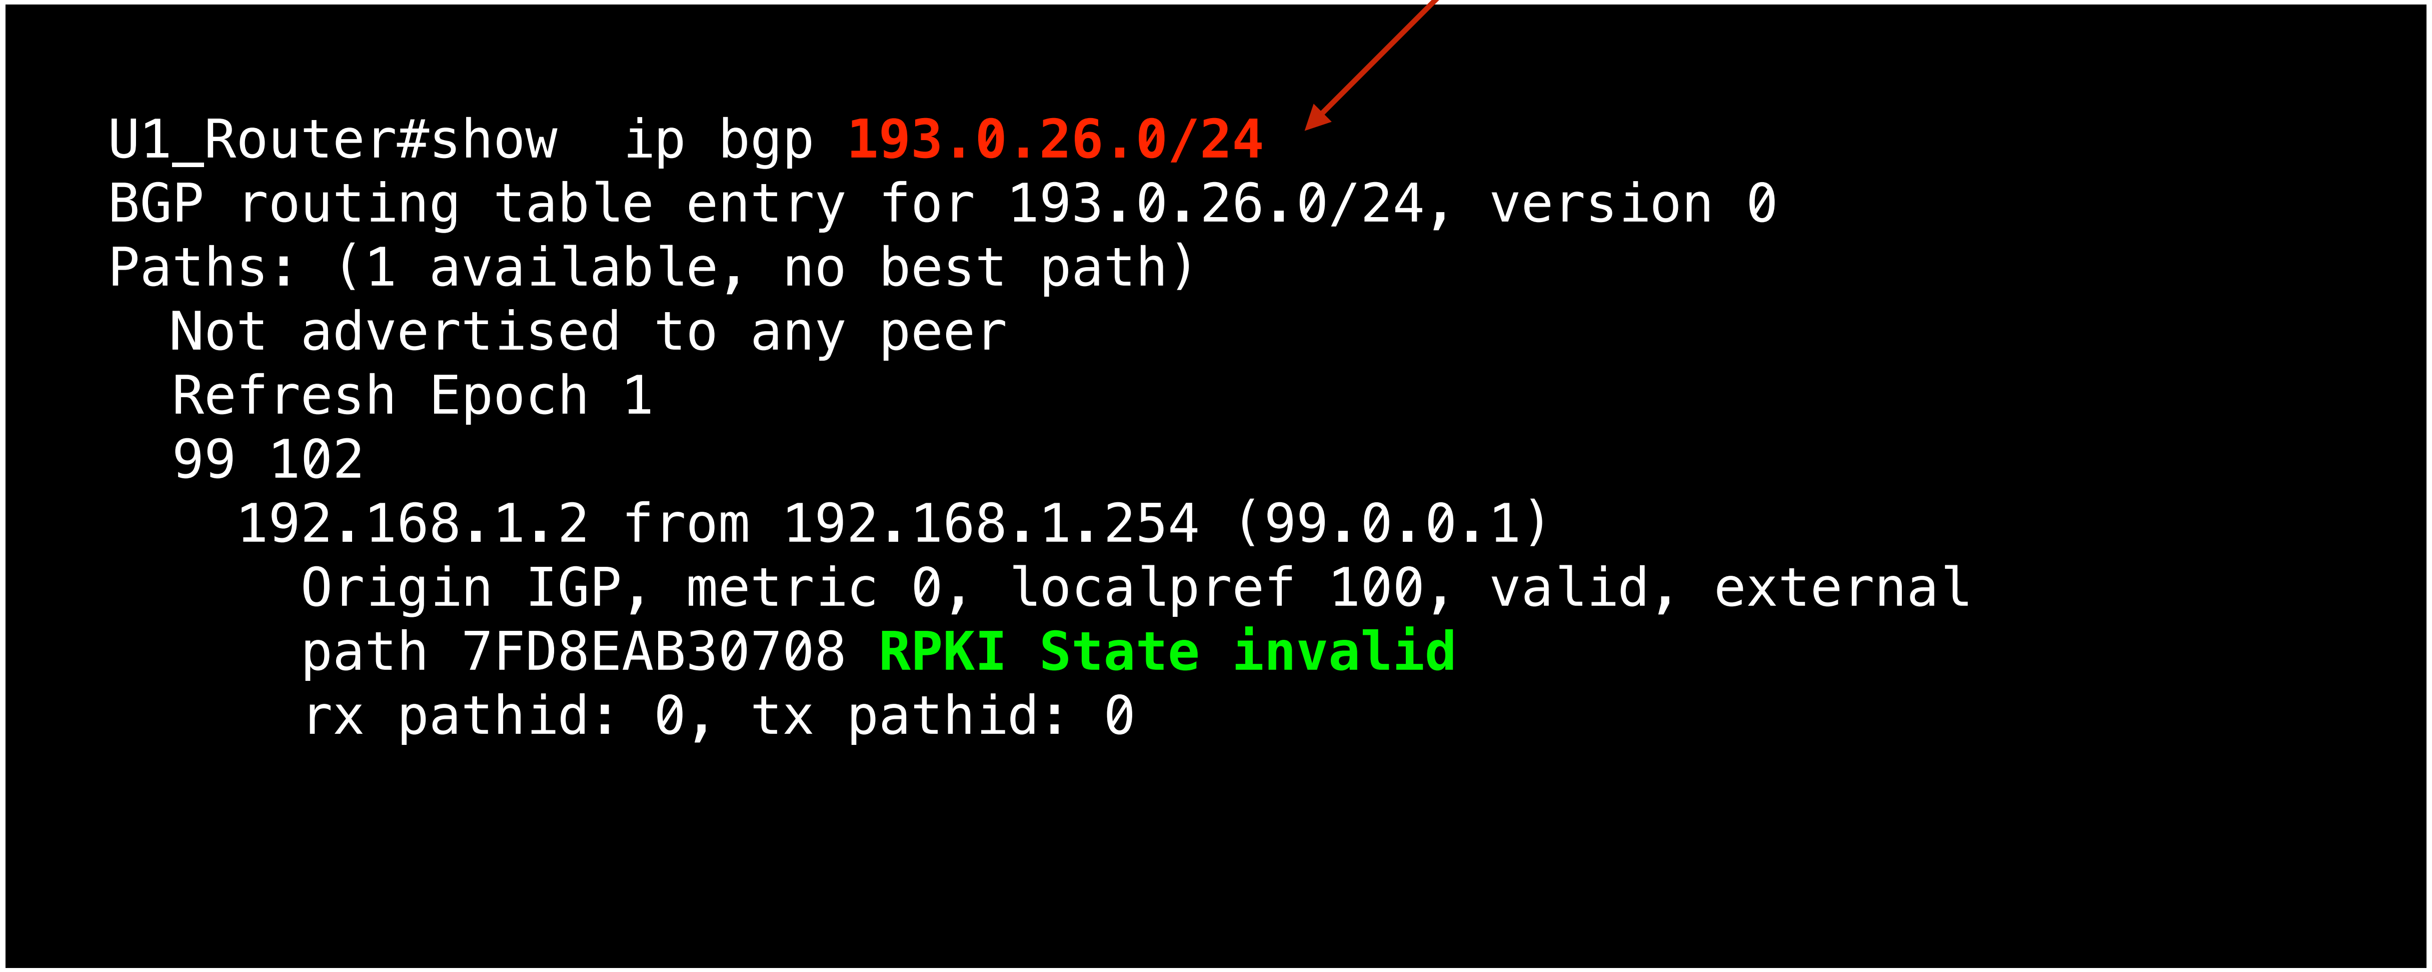

```
U1_Router#show  ip bgp 193.0.26.0/24
BGP routing table entry for 193.0.26.0/24, version 0
Paths: (1 available, no best path)
  Not advertised to any peer
  Refresh Epoch 1
  99 102
    192.168.1.2 from 192.168.1.254 (99.0.0.1)
      Origin IGP, metric 0, localpref 100, valid, external
      path 7FD8EAB30708 RPKI State invalid
      rx pathid: 0, tx pathid: 0
```

Prefix Without a ROA

No ROA for this one!

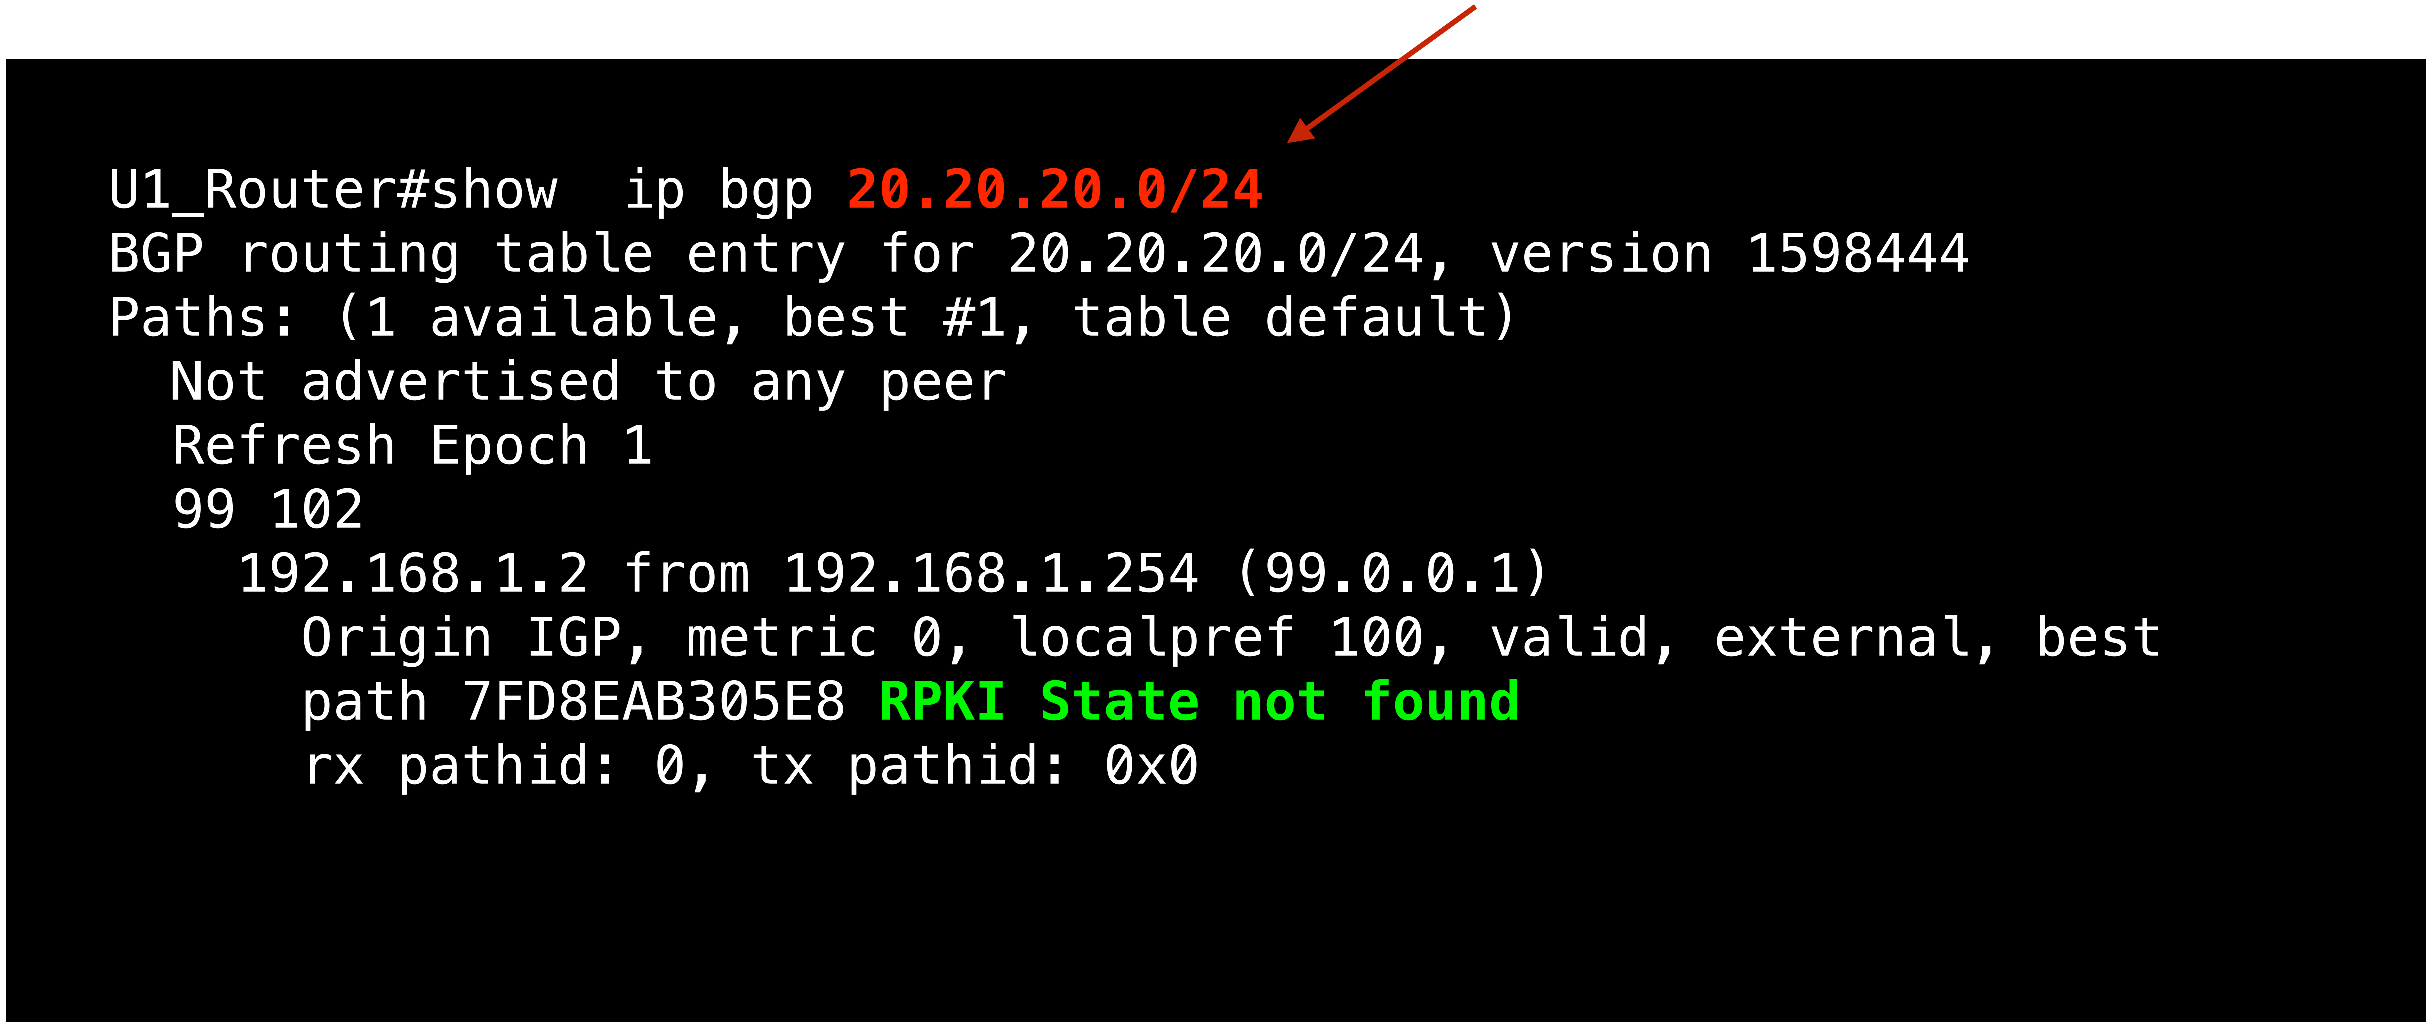

```
U1_Router#show  ip bgp 20.20.20.0/24
BGP routing table entry for 20.20.20.0/24, version 1598444
Paths: (1 available, best #1, table default)
  Not advertised to any peer
  Refresh Epoch 1
  99 102
    192.168.1.2 from 192.168.1.254 (99.0.0.1)
      Origin IGP, metric 0, localpref 100, valid, external, best
      path 7FD8EAB305E8 RPKI State not found
      rx pathid: 0, tx pathid: 0x0
```

Questions ?

Demo!

Discarding BGP Invalids

After Validating...

- You have to make **decisions**
  - Accept or discard the BGP Announcement
  - As temporary measure, you could influence other attributes, such as Local Preference
- You can manage this by using **route-map**

Configure Route Maps

Configure Route-map on the router of **AS101**

```
(config-router)# route-map rpki-accept permit 10
(route-map)# match rpki valid
(route-map)# set local-preference 110
(route-map)# route-map rpki-accept permit 20
(route-map)# match rpki not-found
(route-map)# set local-preference 80
```

Add Route Map to Neighbour

```
(config)# router bgp 101
(config)# address-family ipv4
(config)# neighbor 192.168.1.254 route-map rpki-accept in
```

Reconfigure Your BGP Sessions

```
# clear bgp ipv4 unicast 192.168.1.254
```

And have a bit of patience. The full routing table for both IPv4 and IPv6 needs to be re-evaluated.

Check Your Work

```
# show ip bgp XXX
```

RPKI Valid

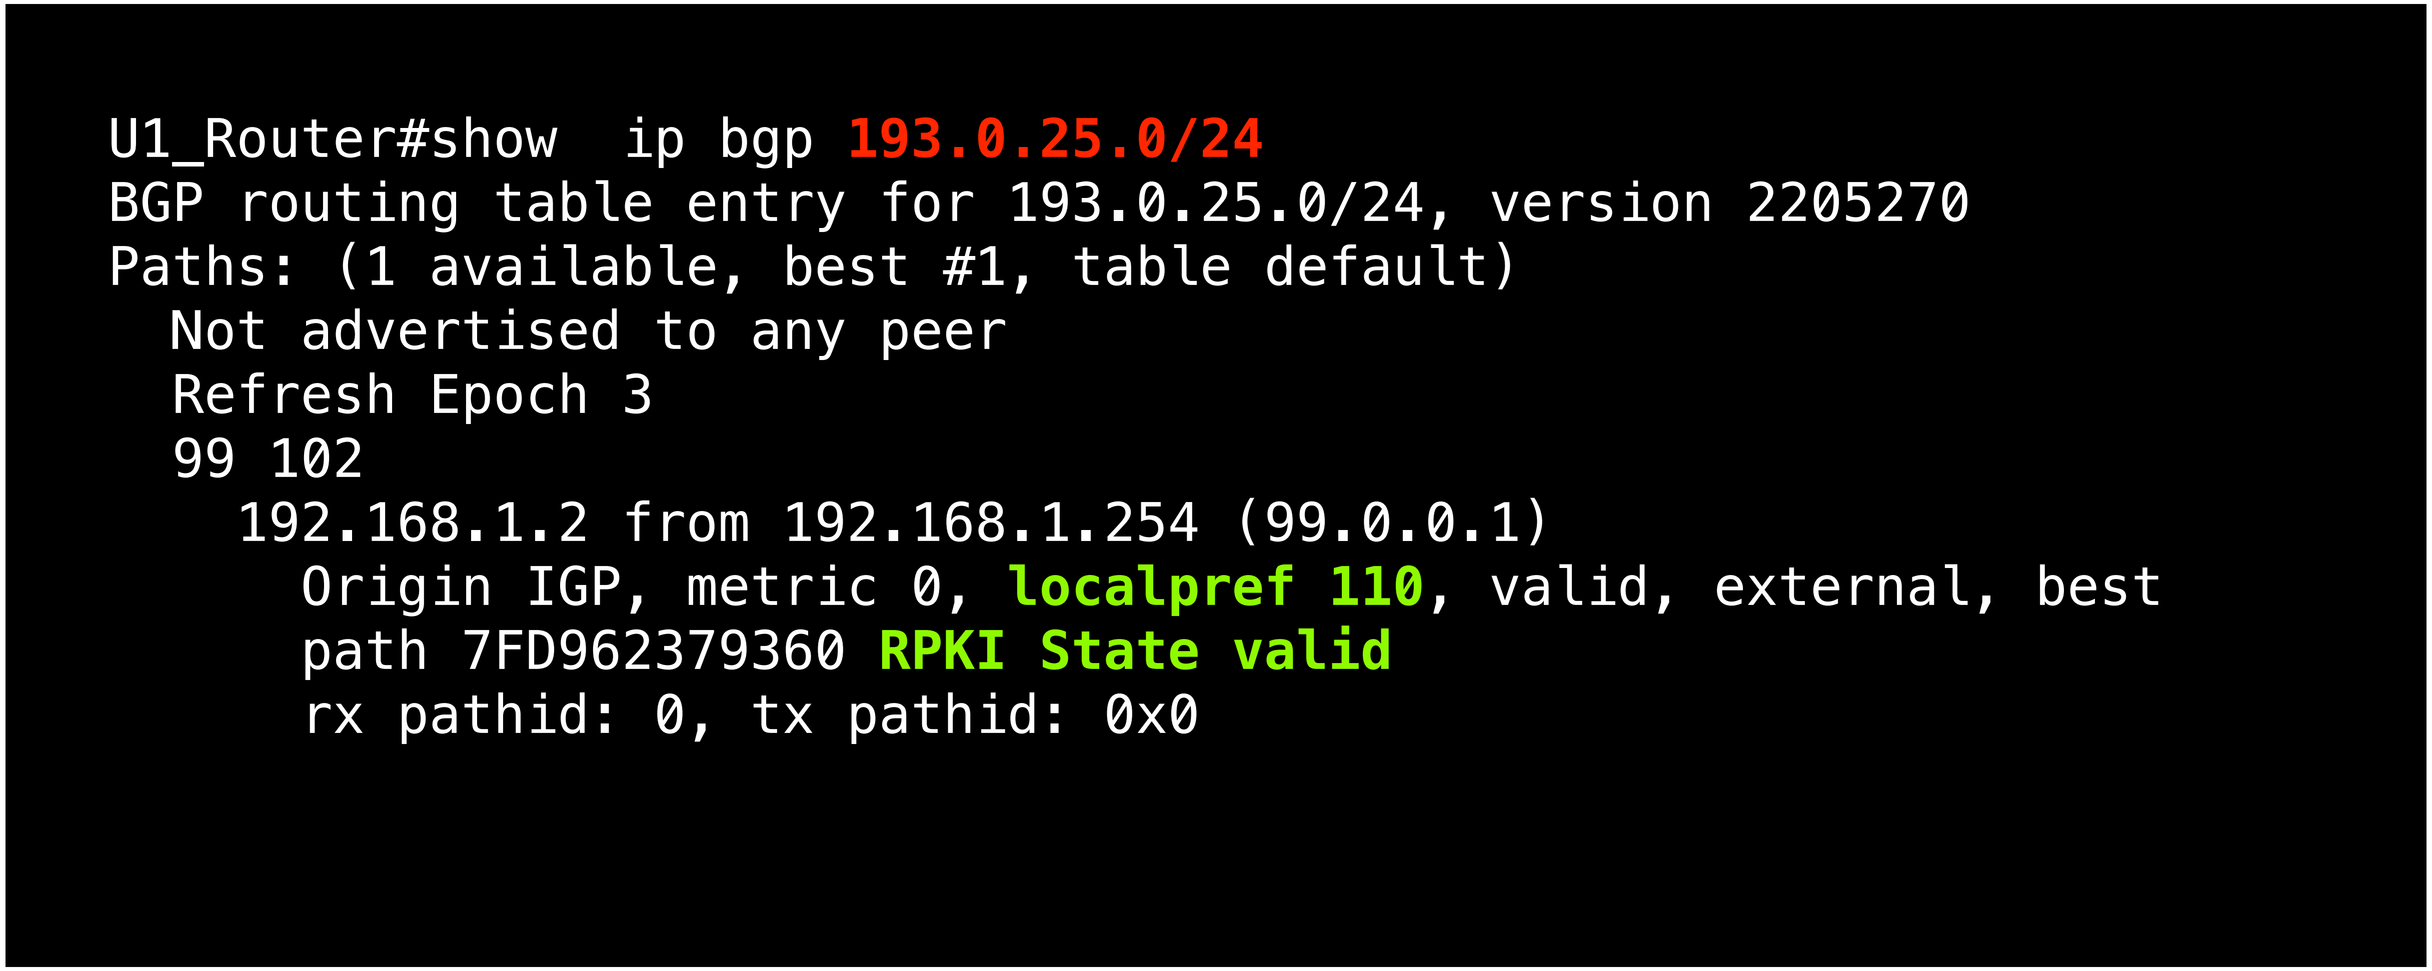

```
U1_Router#show  ip bgp 193.0.25.0/24
BGP routing table entry for 193.0.25.0/24, version 2205270
Paths: (1 available, best #1, table default)
  Not advertised to any peer
  Refresh Epoch 3
  99 102
    192.168.1.2 from 192.168.1.254 (99.0.0.1)
      Origin IGP, metric 0, localpref 110, valid, external, best
      path 7FD962379360 RPKI State valid
      rx pathid: 0, tx pathid: 0x0
```

RPKI Invalid

Prefix belongs to AS103!

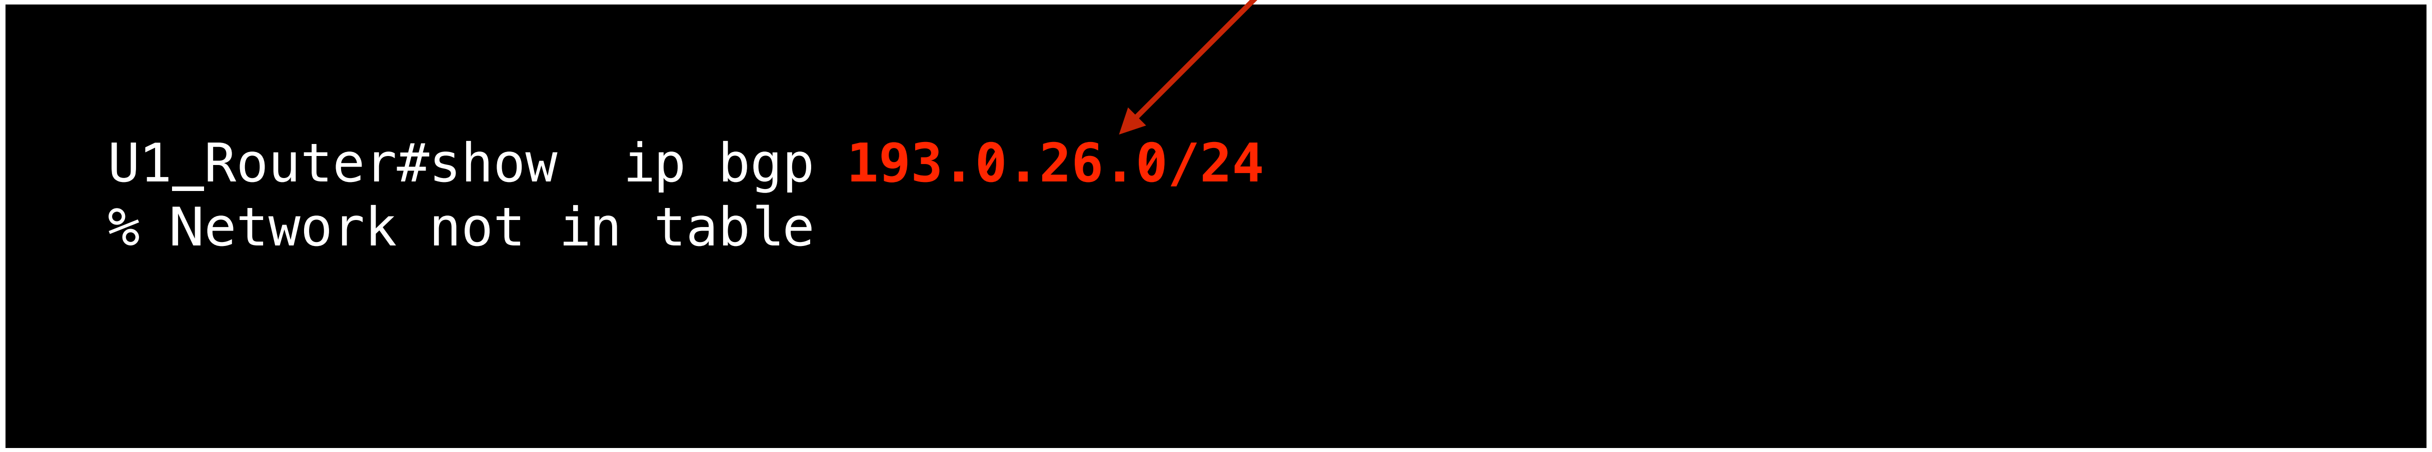

```
U1_Router#show  ip bgp 193.0.26.0/24
% Network not in table
```

Because RPKI state is Invalid!

Prefix Without ROA

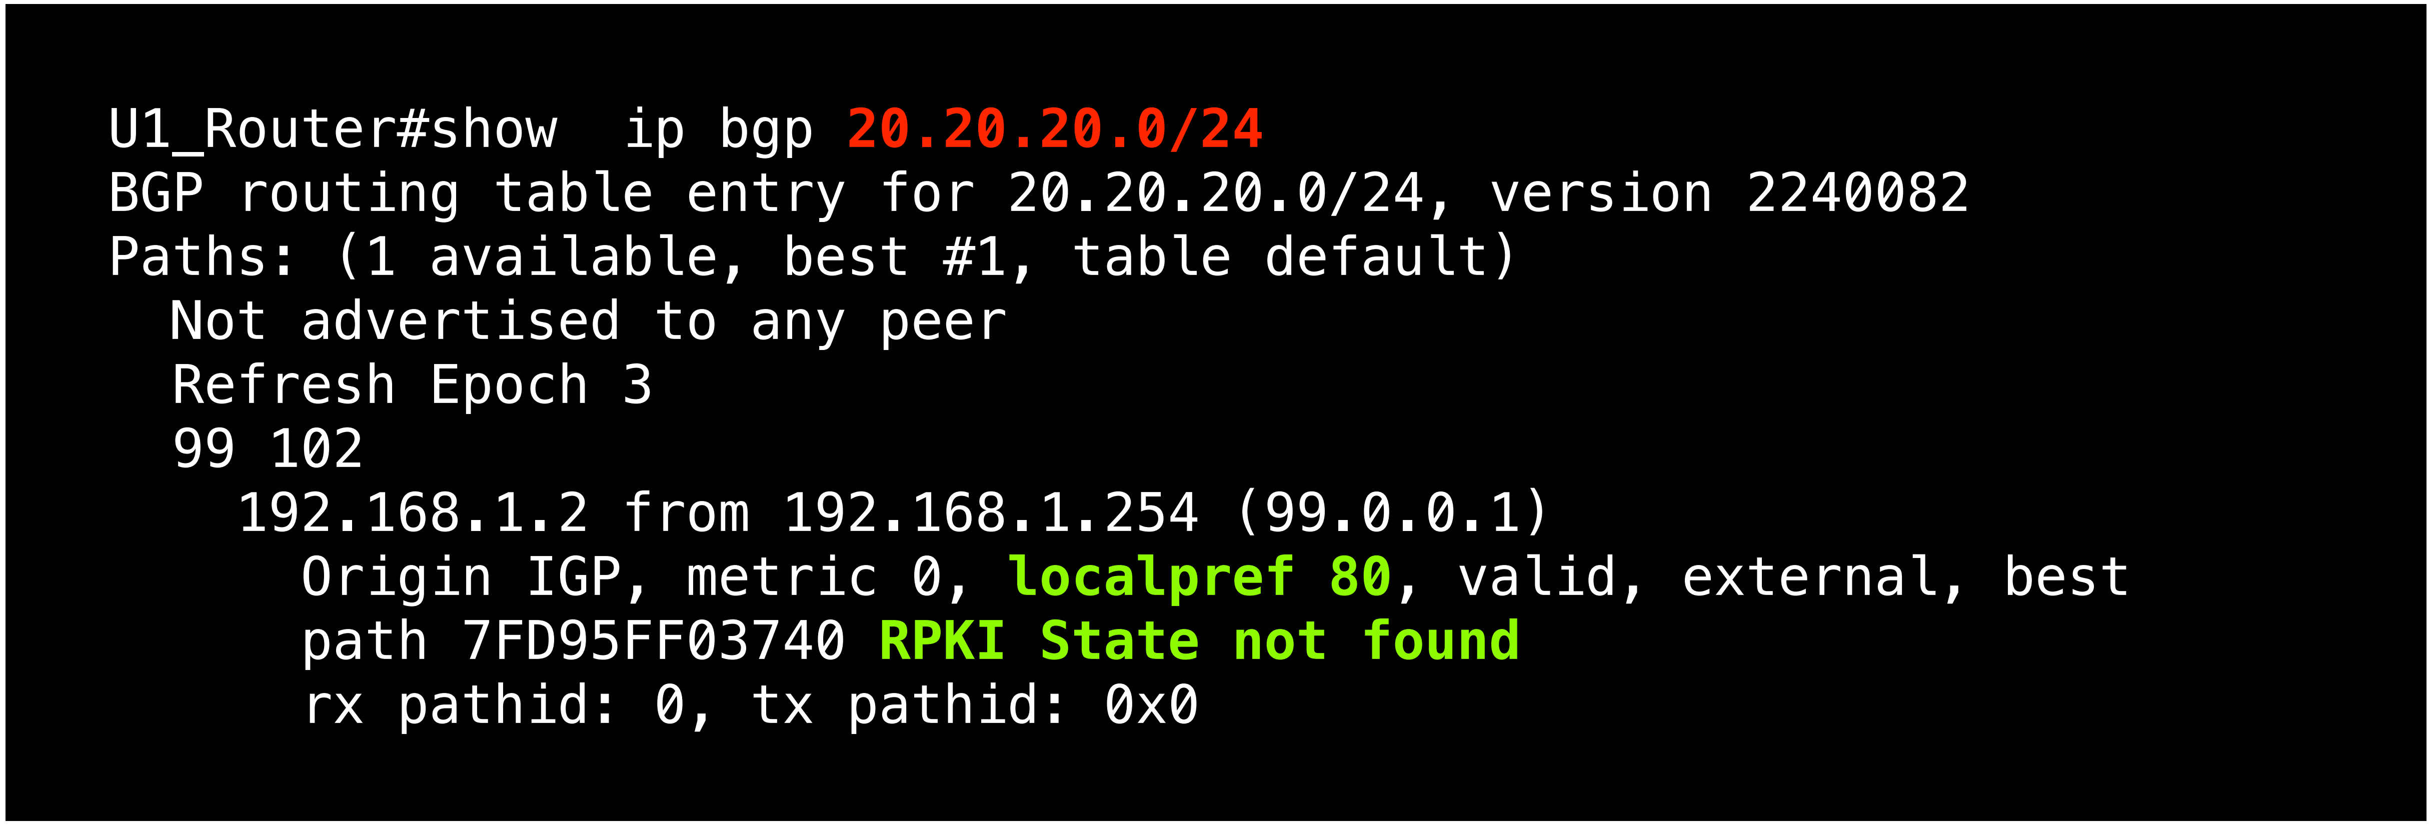

```
U1_Router#show  ip bgp 20.20.20.0/24
BGP routing table entry for 20.20.20.0/24, version 2240082
Paths: (1 available, best #1, table default)
  Not advertised to any peer
  Refresh Epoch 3
  99 102
    192.168.1.2 from 192.168.1.254 (99.0.0.1)
      Origin IGP, metric 0, localpref 80, valid, external, best
      path 7FD95FF03740 RPKI State not found
      rx pathid: 0, tx pathid: 0x0
```

Questions

RIPE NCC
Academy

Learn something new today!

academy.ripe.net

RIPE NCC Certified Professionals

RIPE NCC Certified Professionals
IPv6 Fundamentals
Analyst

RIPE NCC Certified Professionals
RIPE Database
Associate

RIPE NCC Certified Professionals
IPv6 Security
Expert

https://www.ripe.net/certifiedprofessionals

The End!

Край

Y Diwedd

Fí

Finis

النهاية

Соңы

Վերջ

Liðugt

Ende

Finvezh

Кінець

Konec

Kraj

Ënn

Fund

پایان

Lõpp

Beigas

Vége

Son

An Críoch

Крај

הסוף

Fine

Endir

Sfârşit

Fin

Τέλος

Einde

Конец

Slut

Slutt

დასასრული

Pabaiga

Fim

Amaia

Loppu

Tmiem

Koniec