RIPE NCC
RIPE NETWORK COORDINATION CENTRE

Introduction to RPKI

Tutorial

RIPE NCC Learning & Development

Çiğdem Gür Şenol | ENOG 18 | 7 June 2021

Agenda

Routing on the Internet

Routing Security

How does RPKI work?

Routing on the Internet

How does Internet Routing work?

As you know, we’re using the **BGP** protocol on the Internet.
There are **no alternatives!**

BGP

- The routing protocol of the Internet!
- Runs between independently operated networks (ASes)
- Connects the entire Internet
- Goal of BGP is to exchange routing and reachability info

How Does it Work?

How Does it Work?

How Does it Work?

How Does it Work?

How Does it Work?

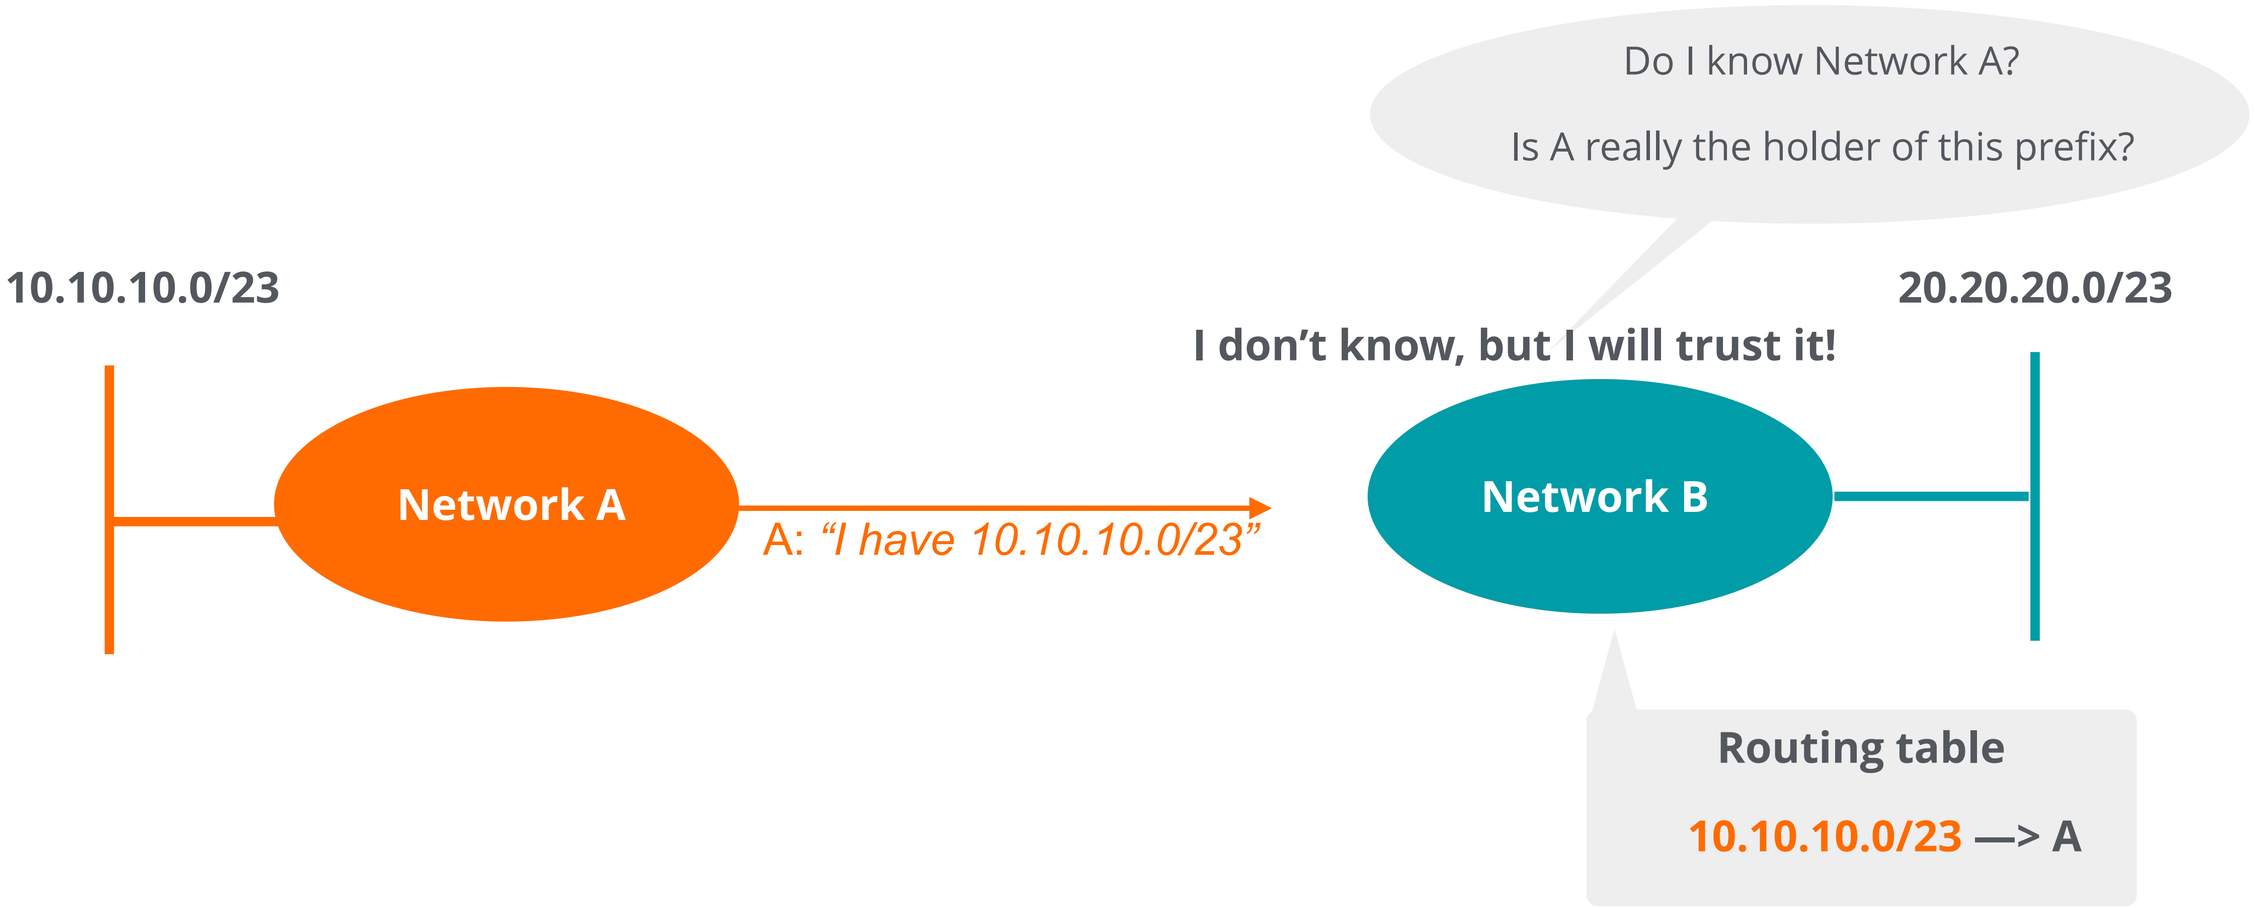

How Does it Work?

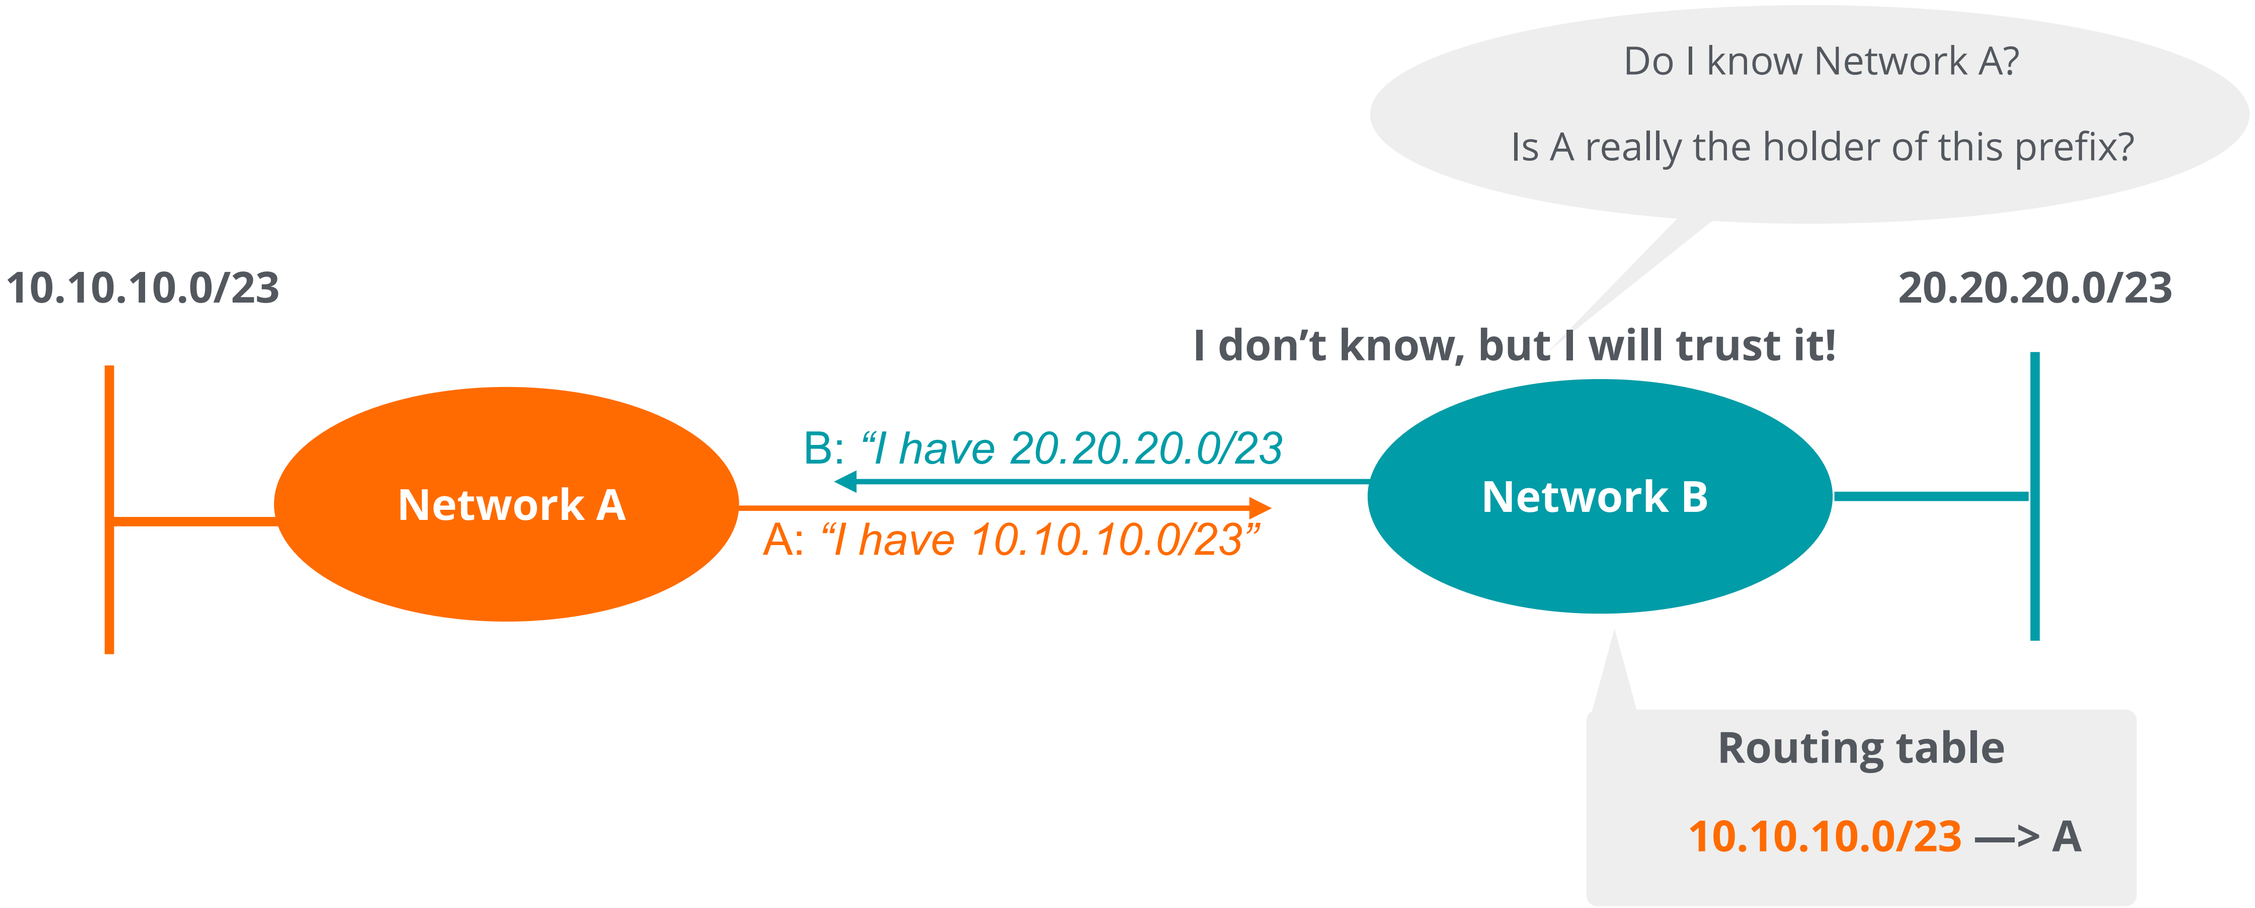

How Does it Work?

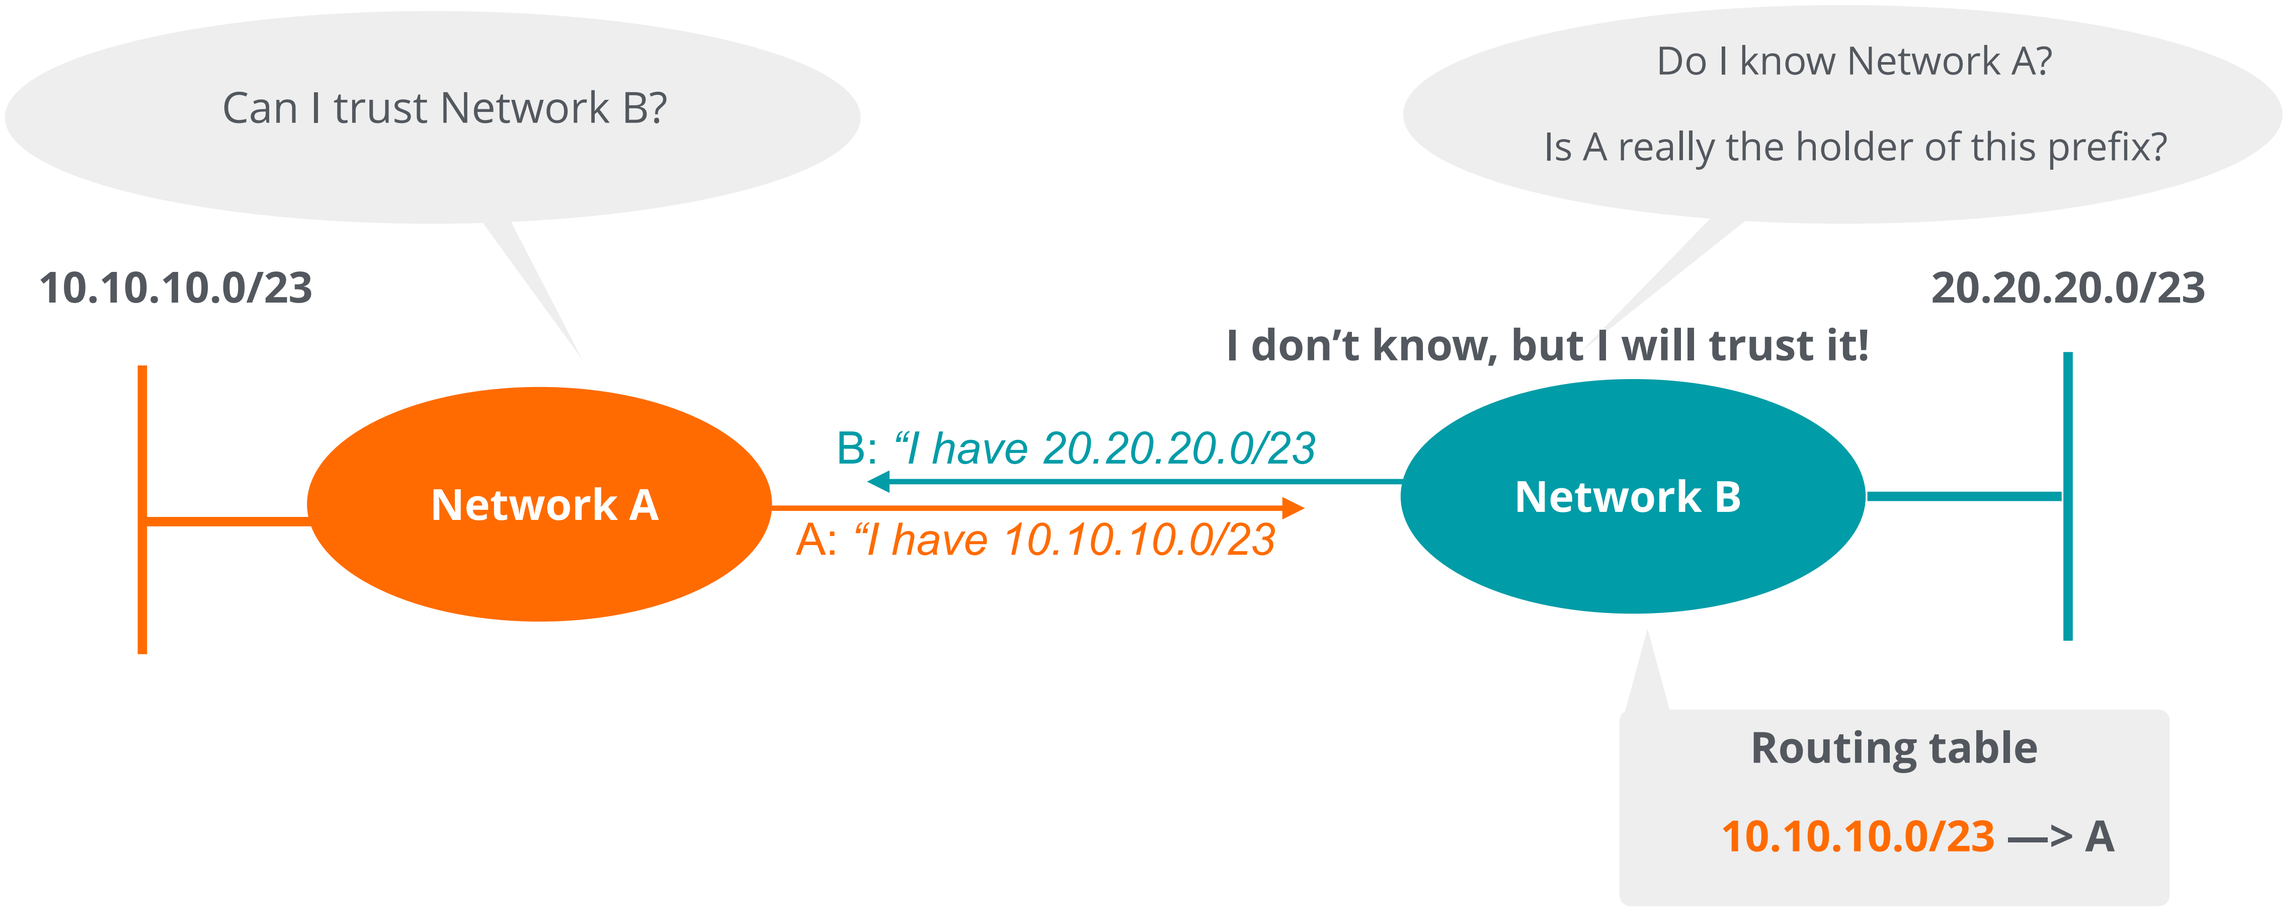

How Does it Work?

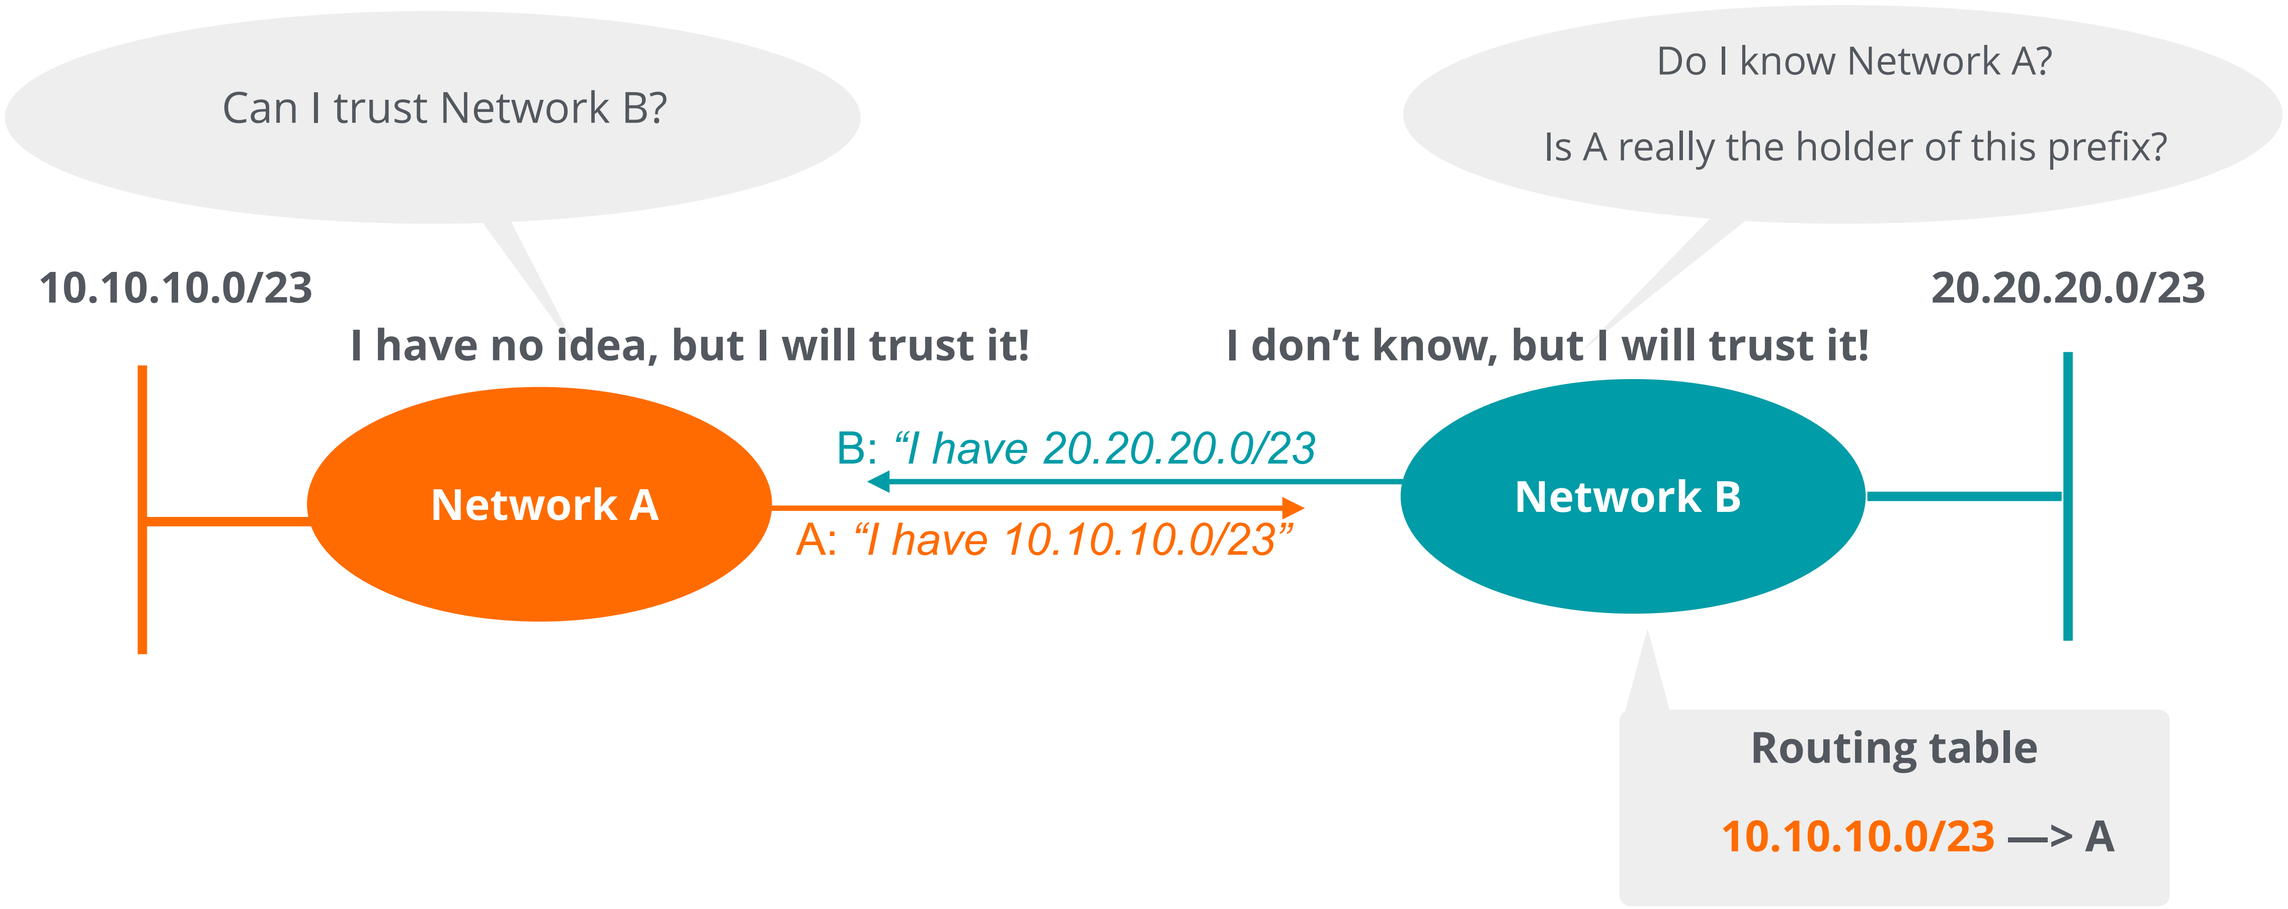

How Does it Work?

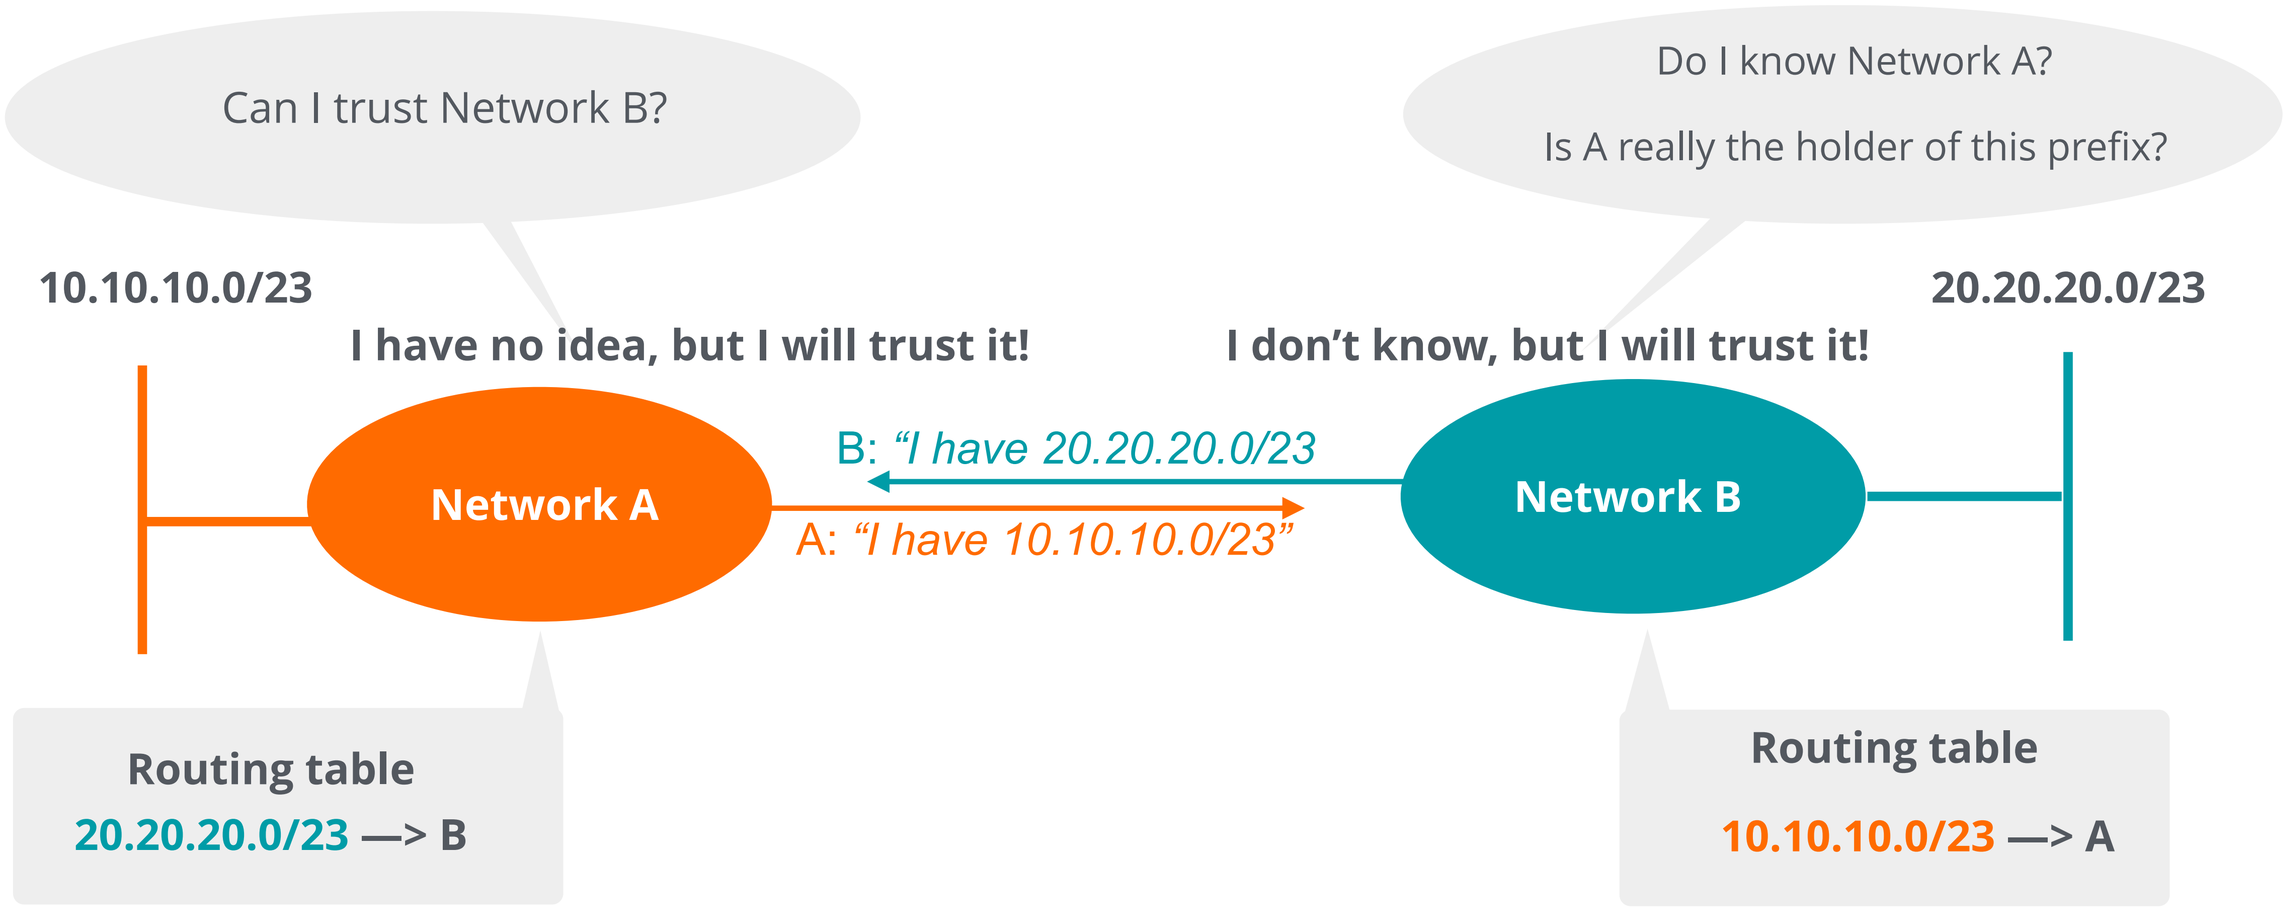

BGP is based on **trust!** There is **no built-in security!**

- Any AS can announce any prefix
- Anyone can prepend any ASN to any path they want
- BGP packets are transmitted without any encryption and authentication mechanism
- No single authoritative source of who should be doing what

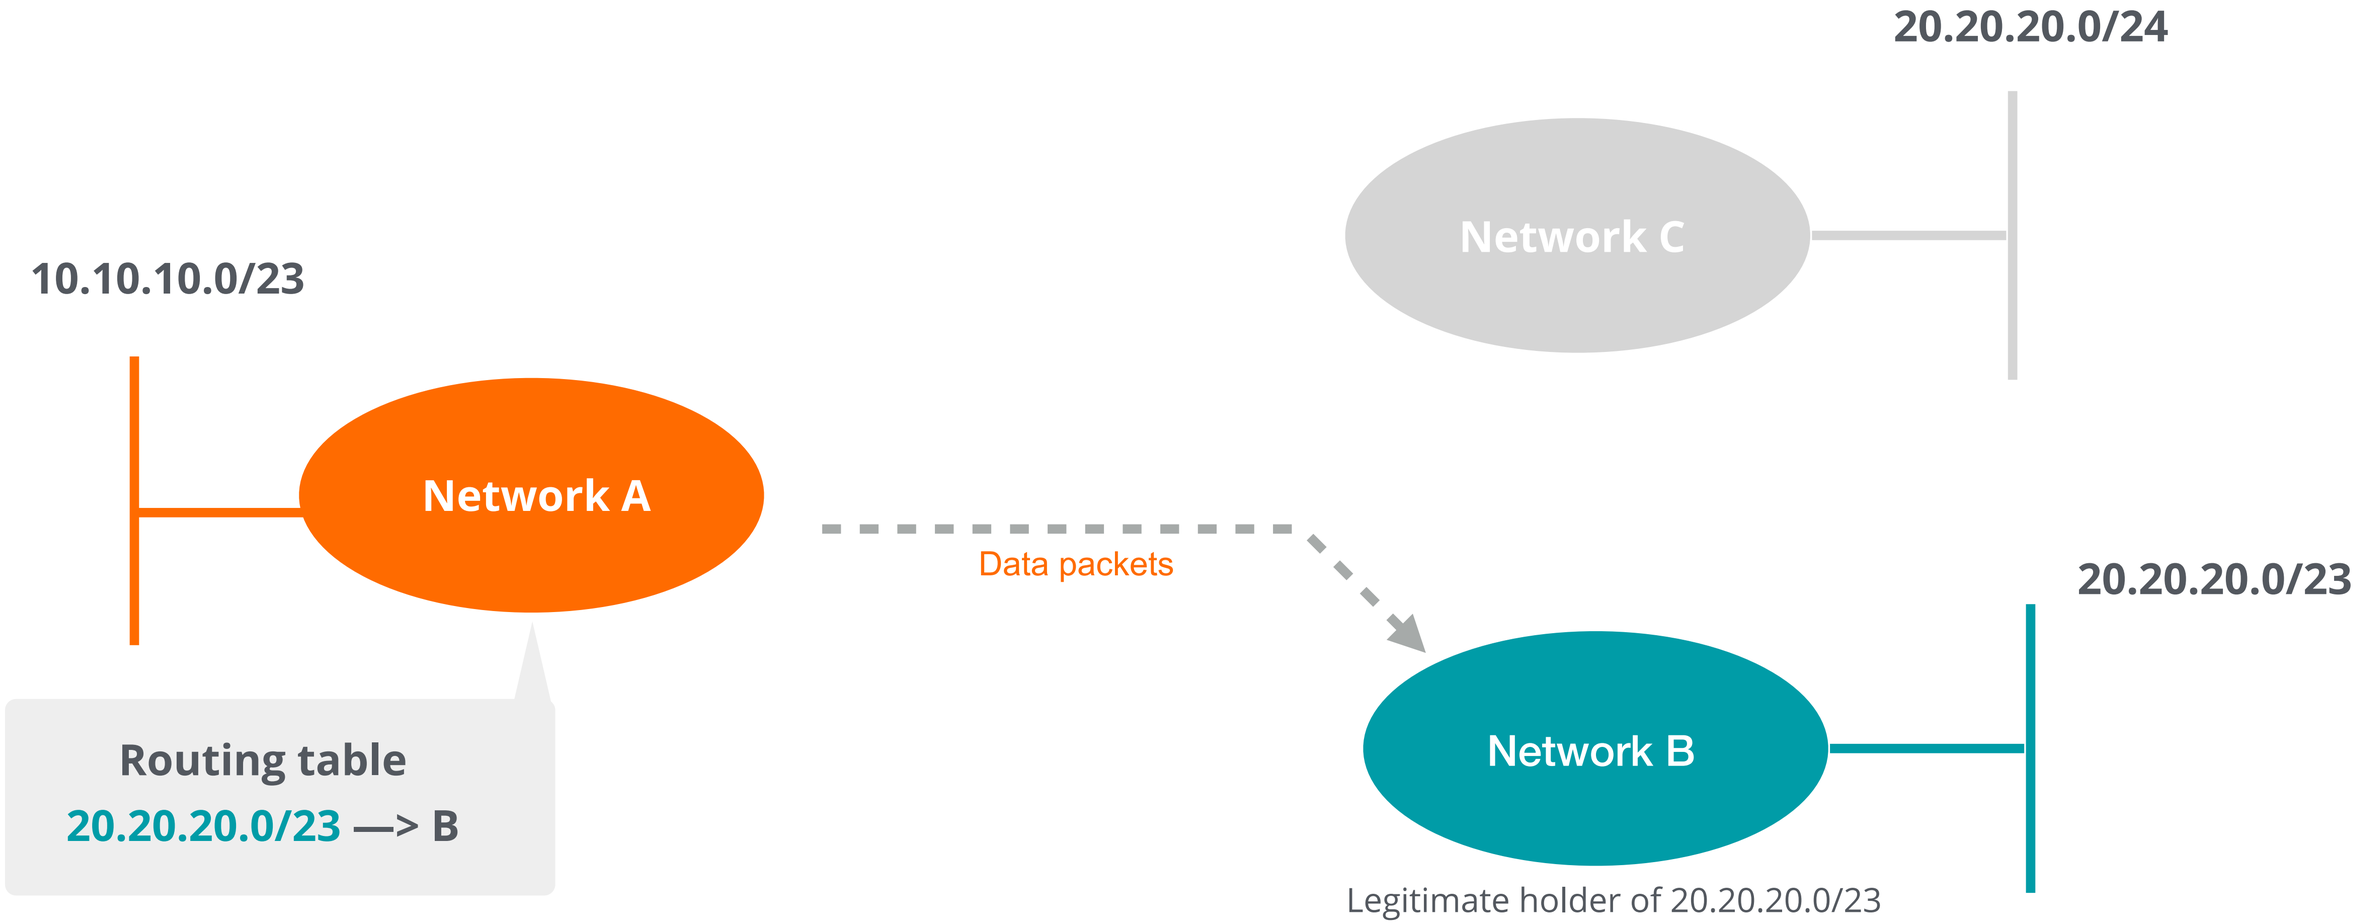
20.20.20.0/24
Network C
10.10.10.0/23
Network A
Data packets
20.20.20.0/23
Routing table
20.20.20.0/23 —> B
Network B
Legitimate holder of 20.20.20.0/23

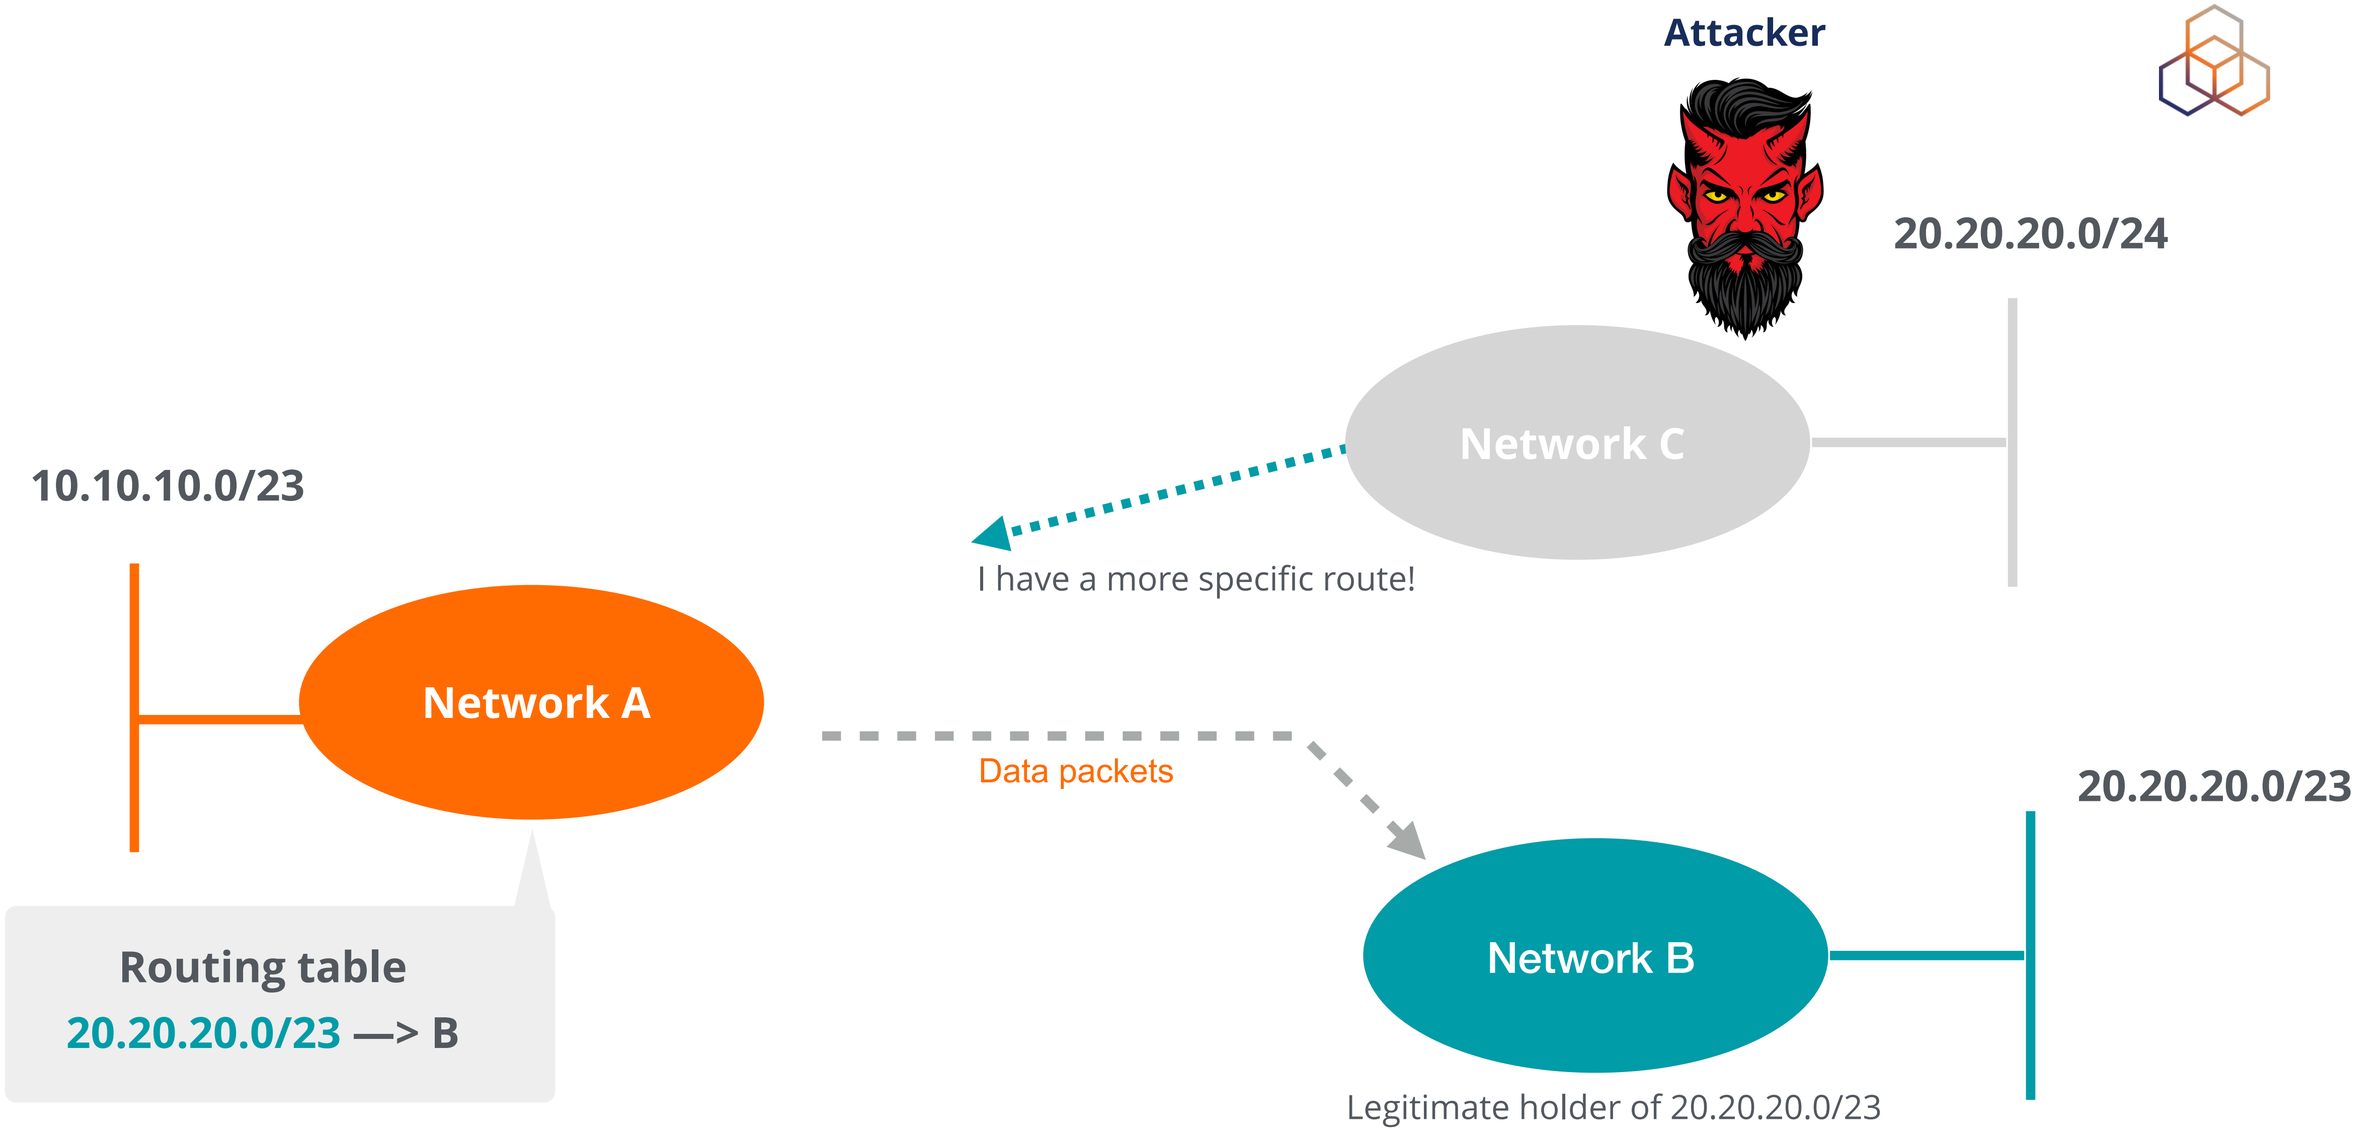
Attacker
20.20.20.0/24
Network C
10.10.10.0/23
I have a more specific route!
Network A
Data packets
20.20.20.0/23
Routing table
20.20.20.0/23 —> B
Network B
Legitimate holder of 20.20.20.0/23

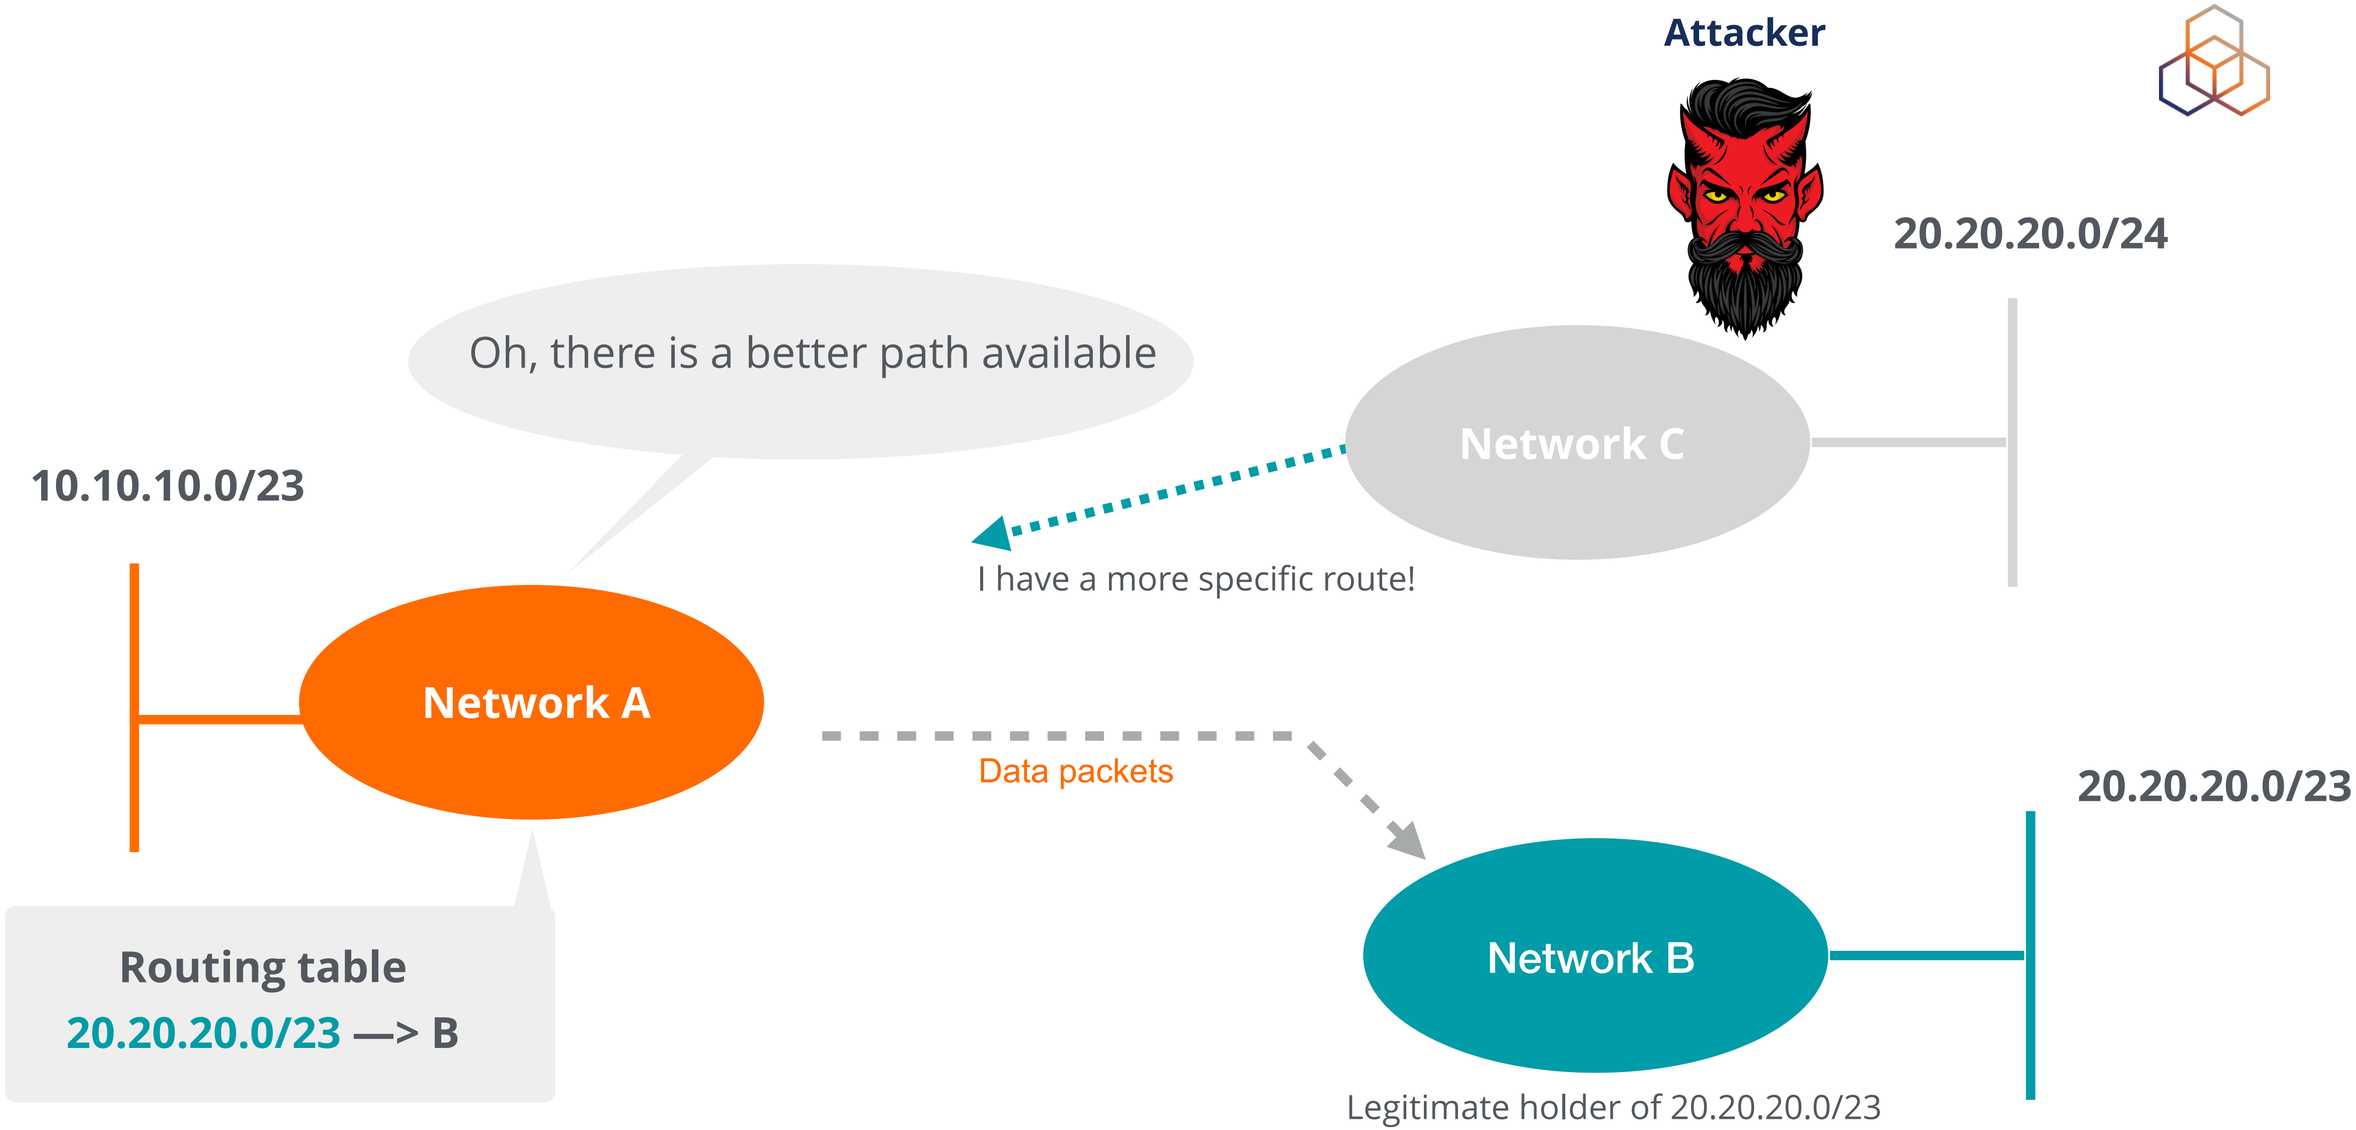
Attacker
20.20.20.0/24
Oh, there is a better path available
Network C
10.10.10.0/23
I have a more specific route!
Network A
Data packets
20.20.20.0/23
Routing table
20.20.20.0/23 —> B
Network B
Legitimate holder of 20.20.20.0/23

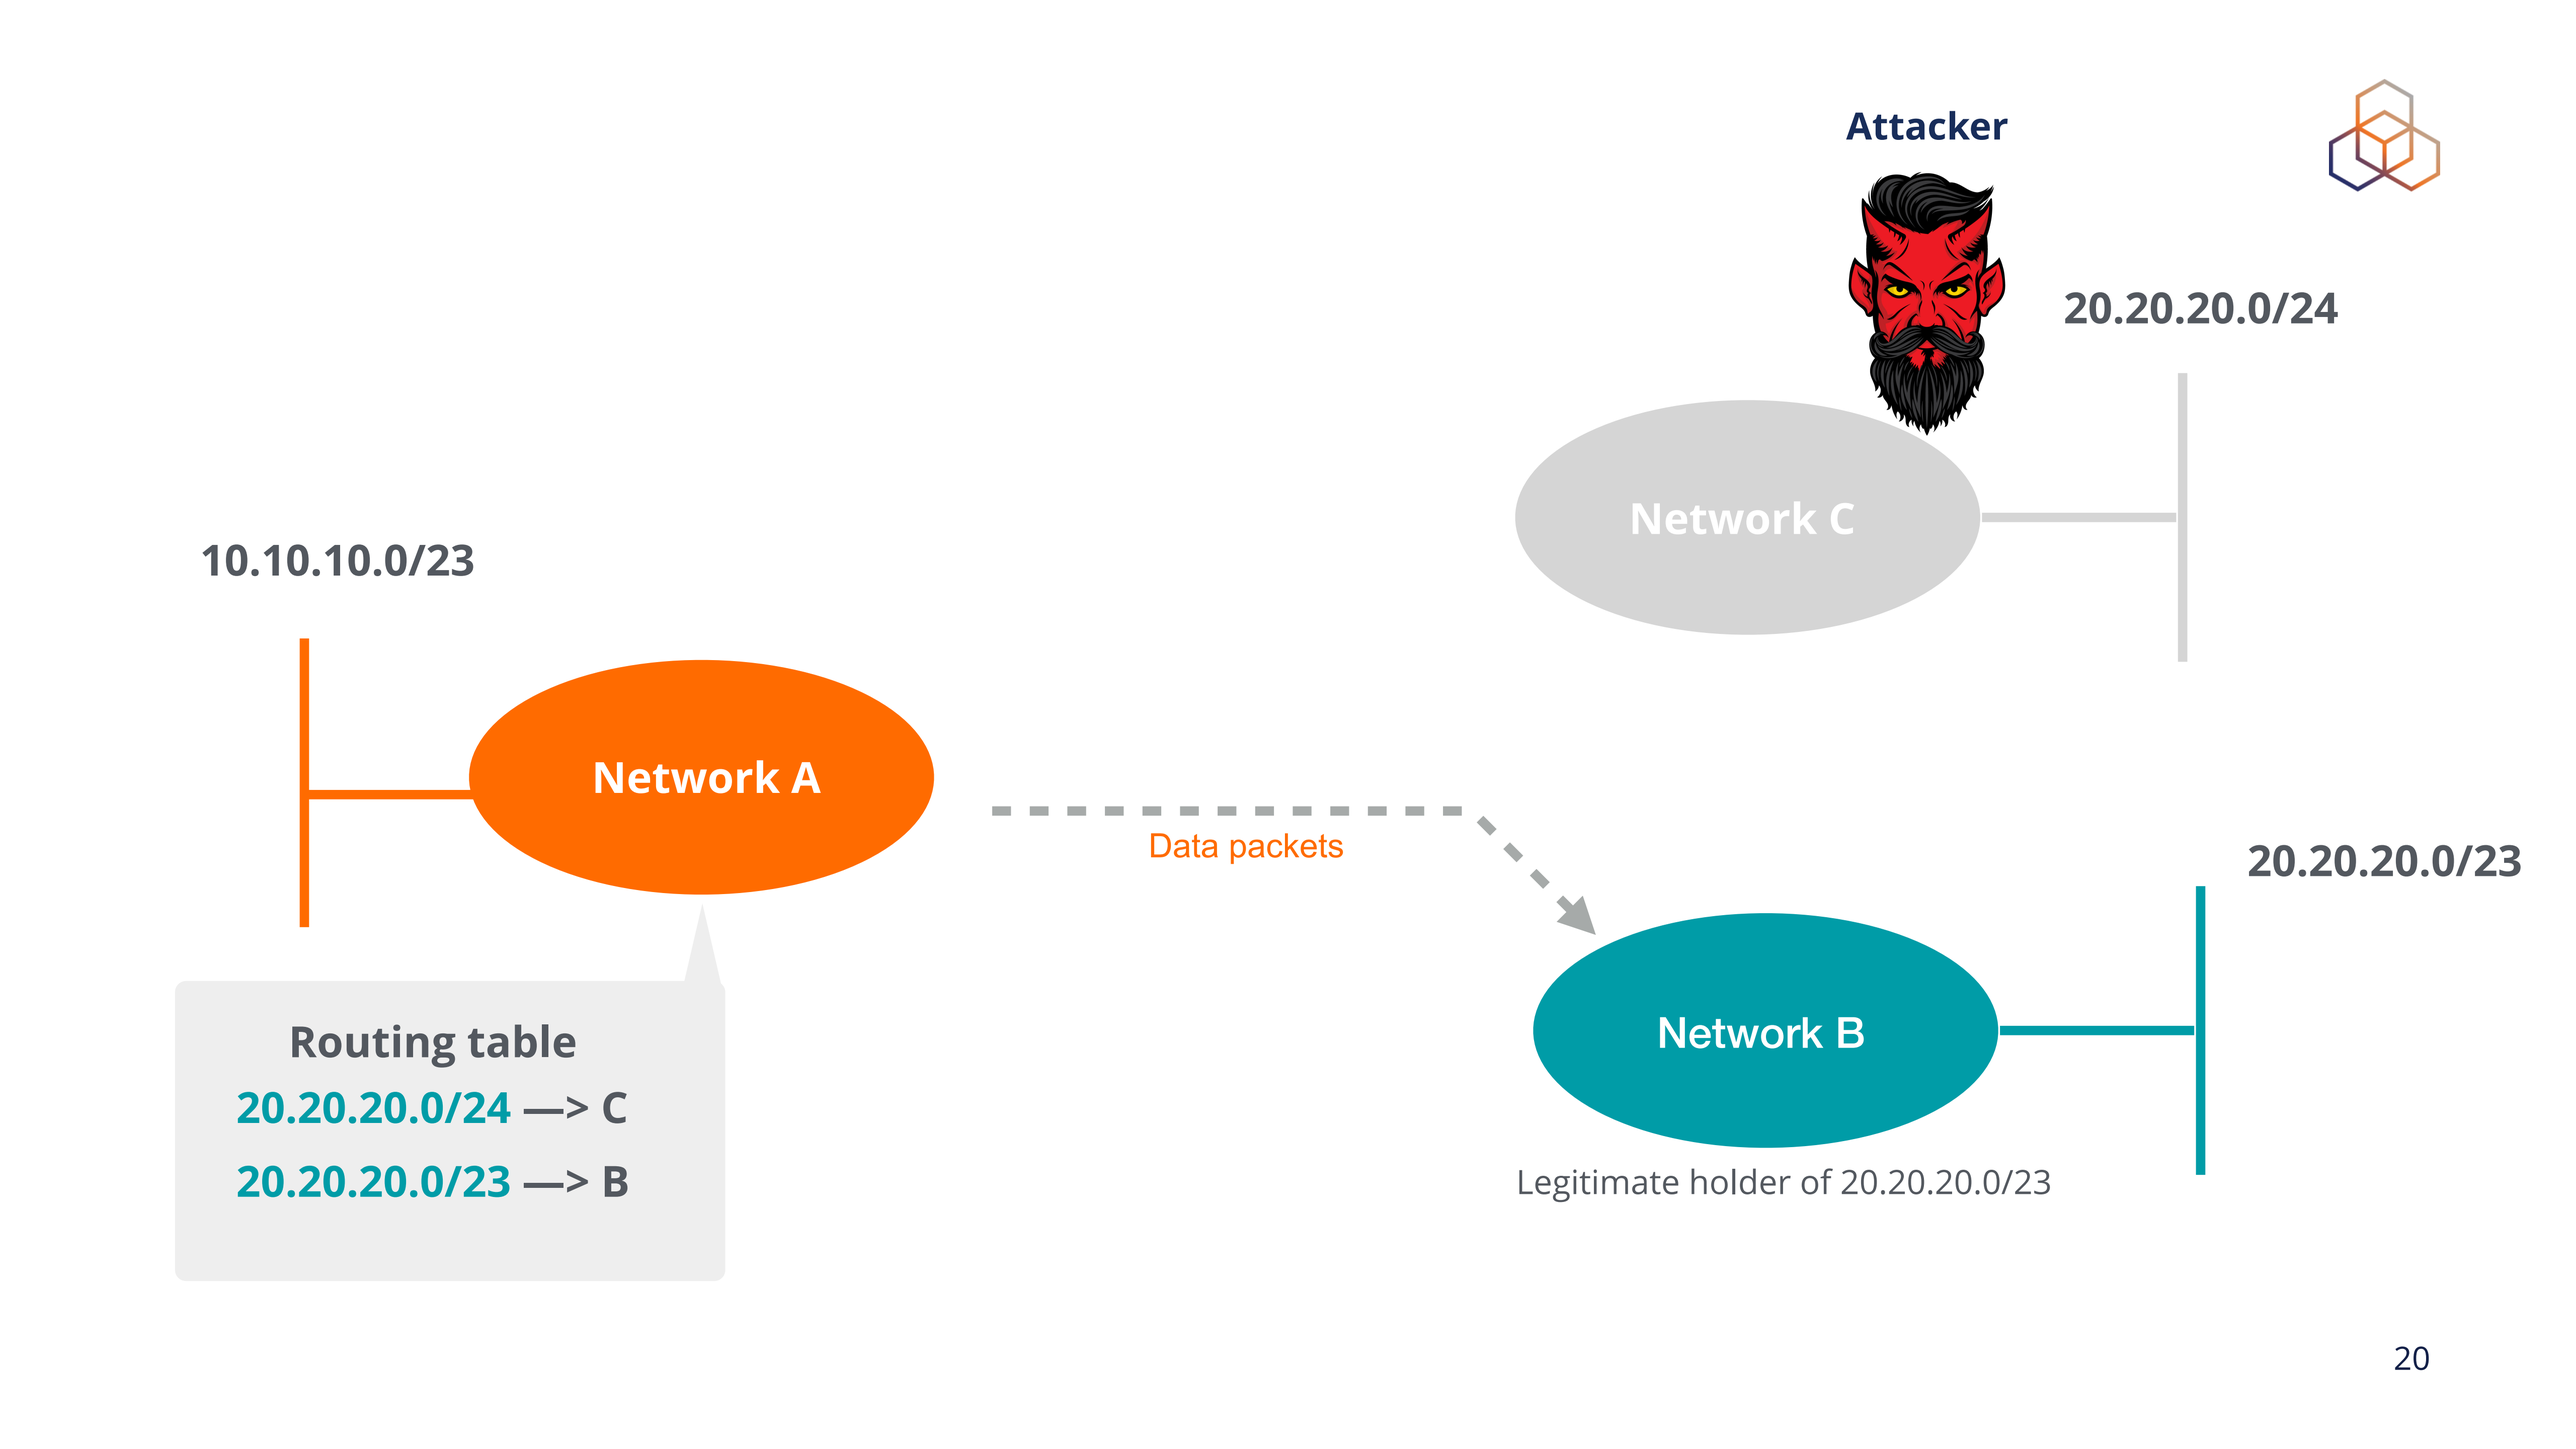
Attacker
20.20.20.0/24
Network C
10.10.10.0/23
Network A
Data packets
20.20.20.0/23
Network B
Legitimate holder of 20.20.20.0/23
Routing table
20.20.20.0/24 —> C
20.20.20.0/23 —> B

All traffic to 20.20.20.0/24 is diverted to the attacker's network!

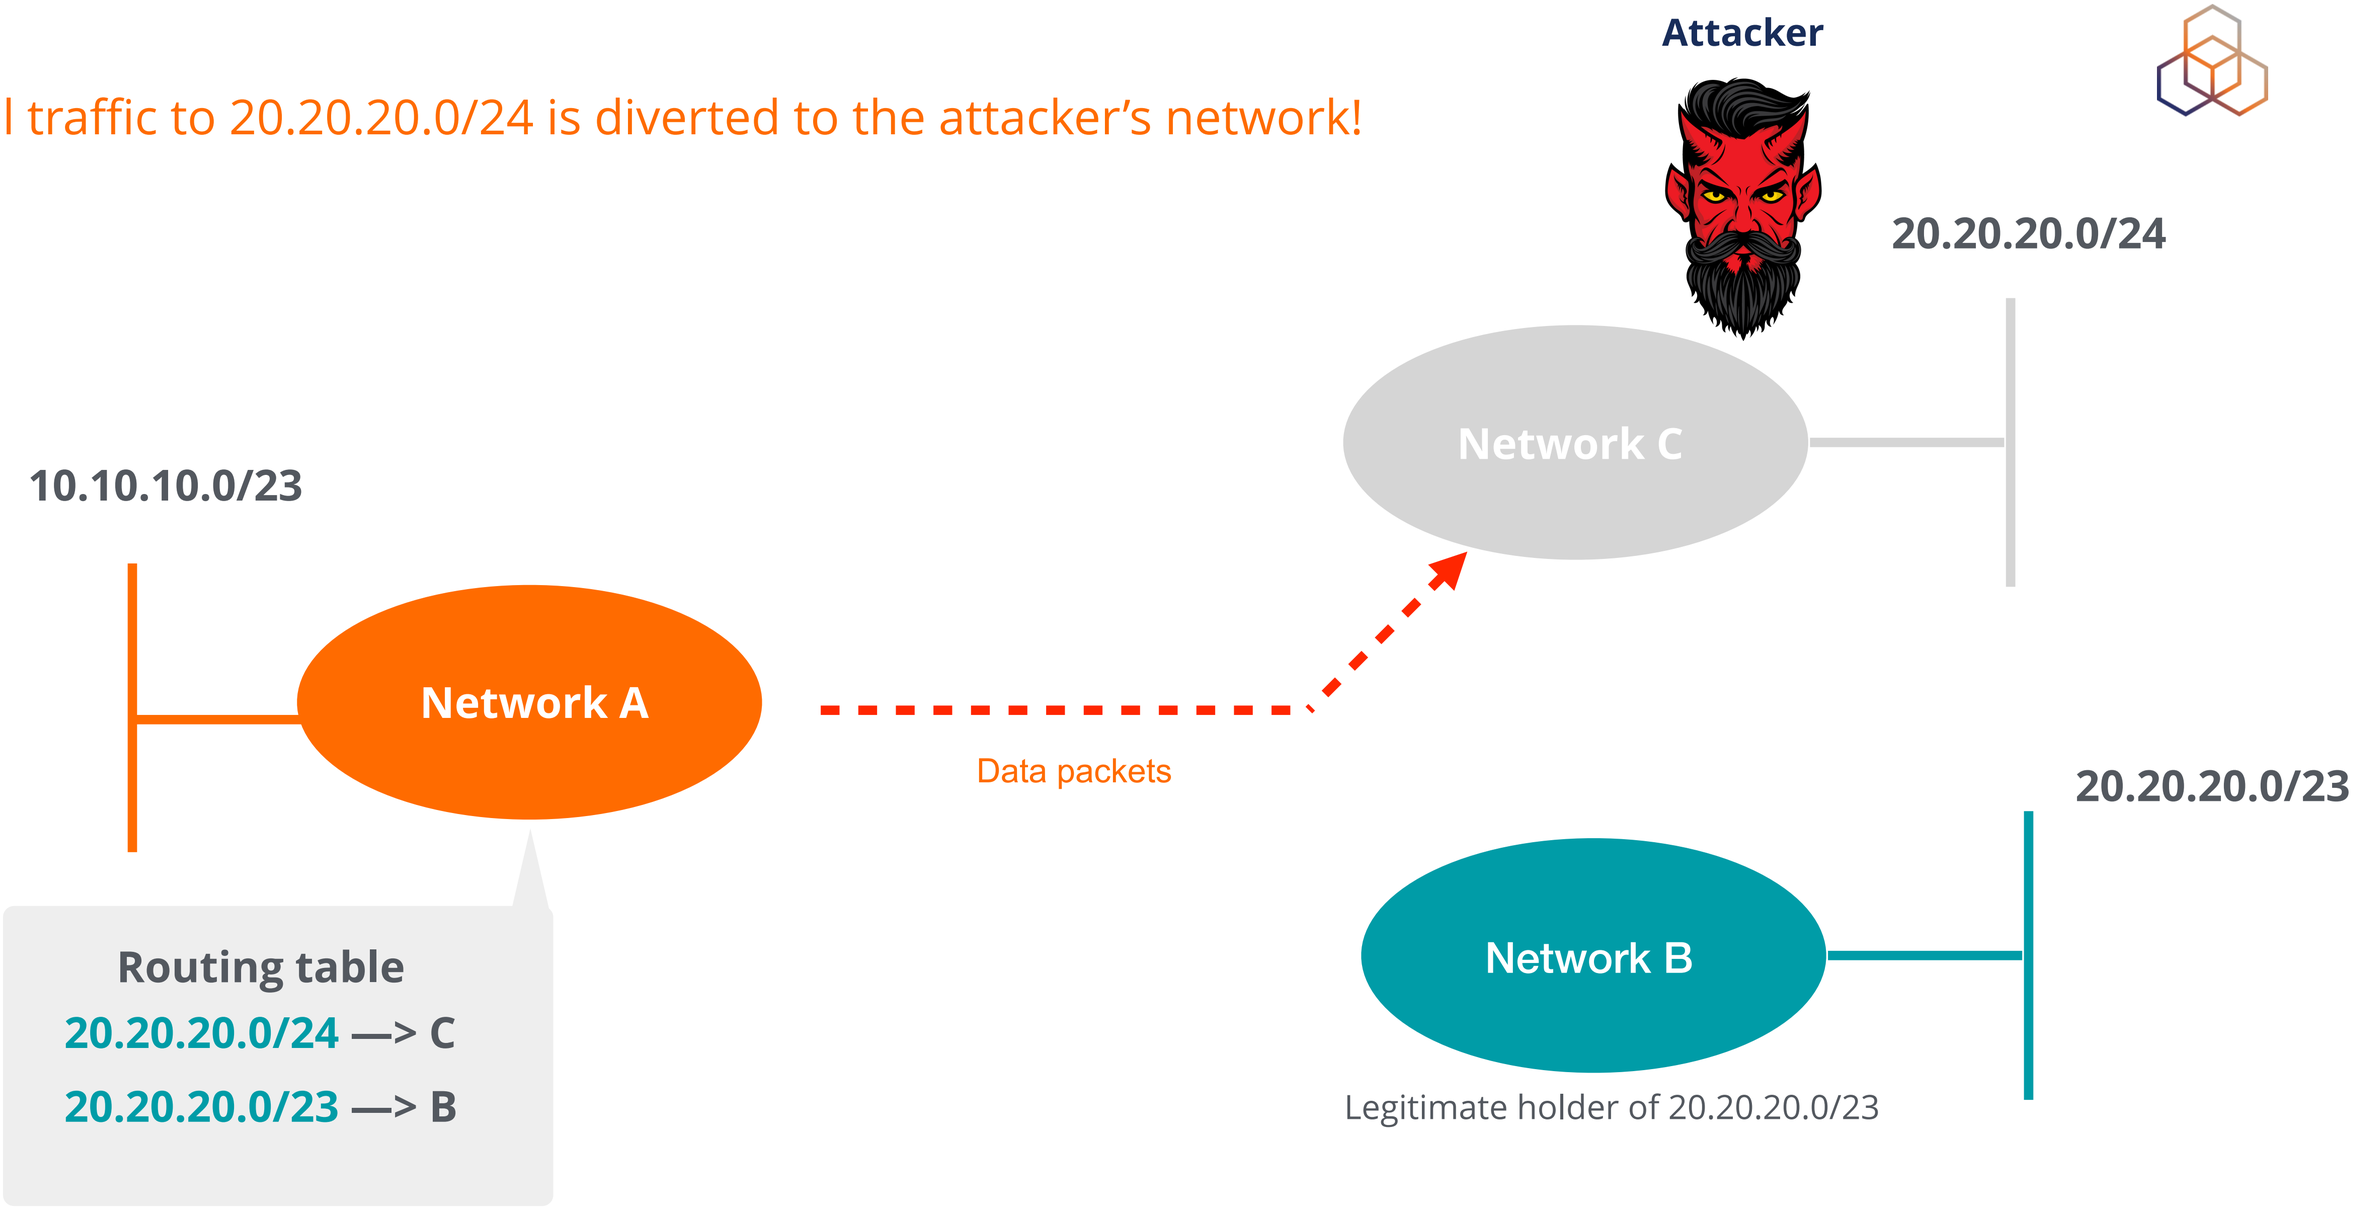

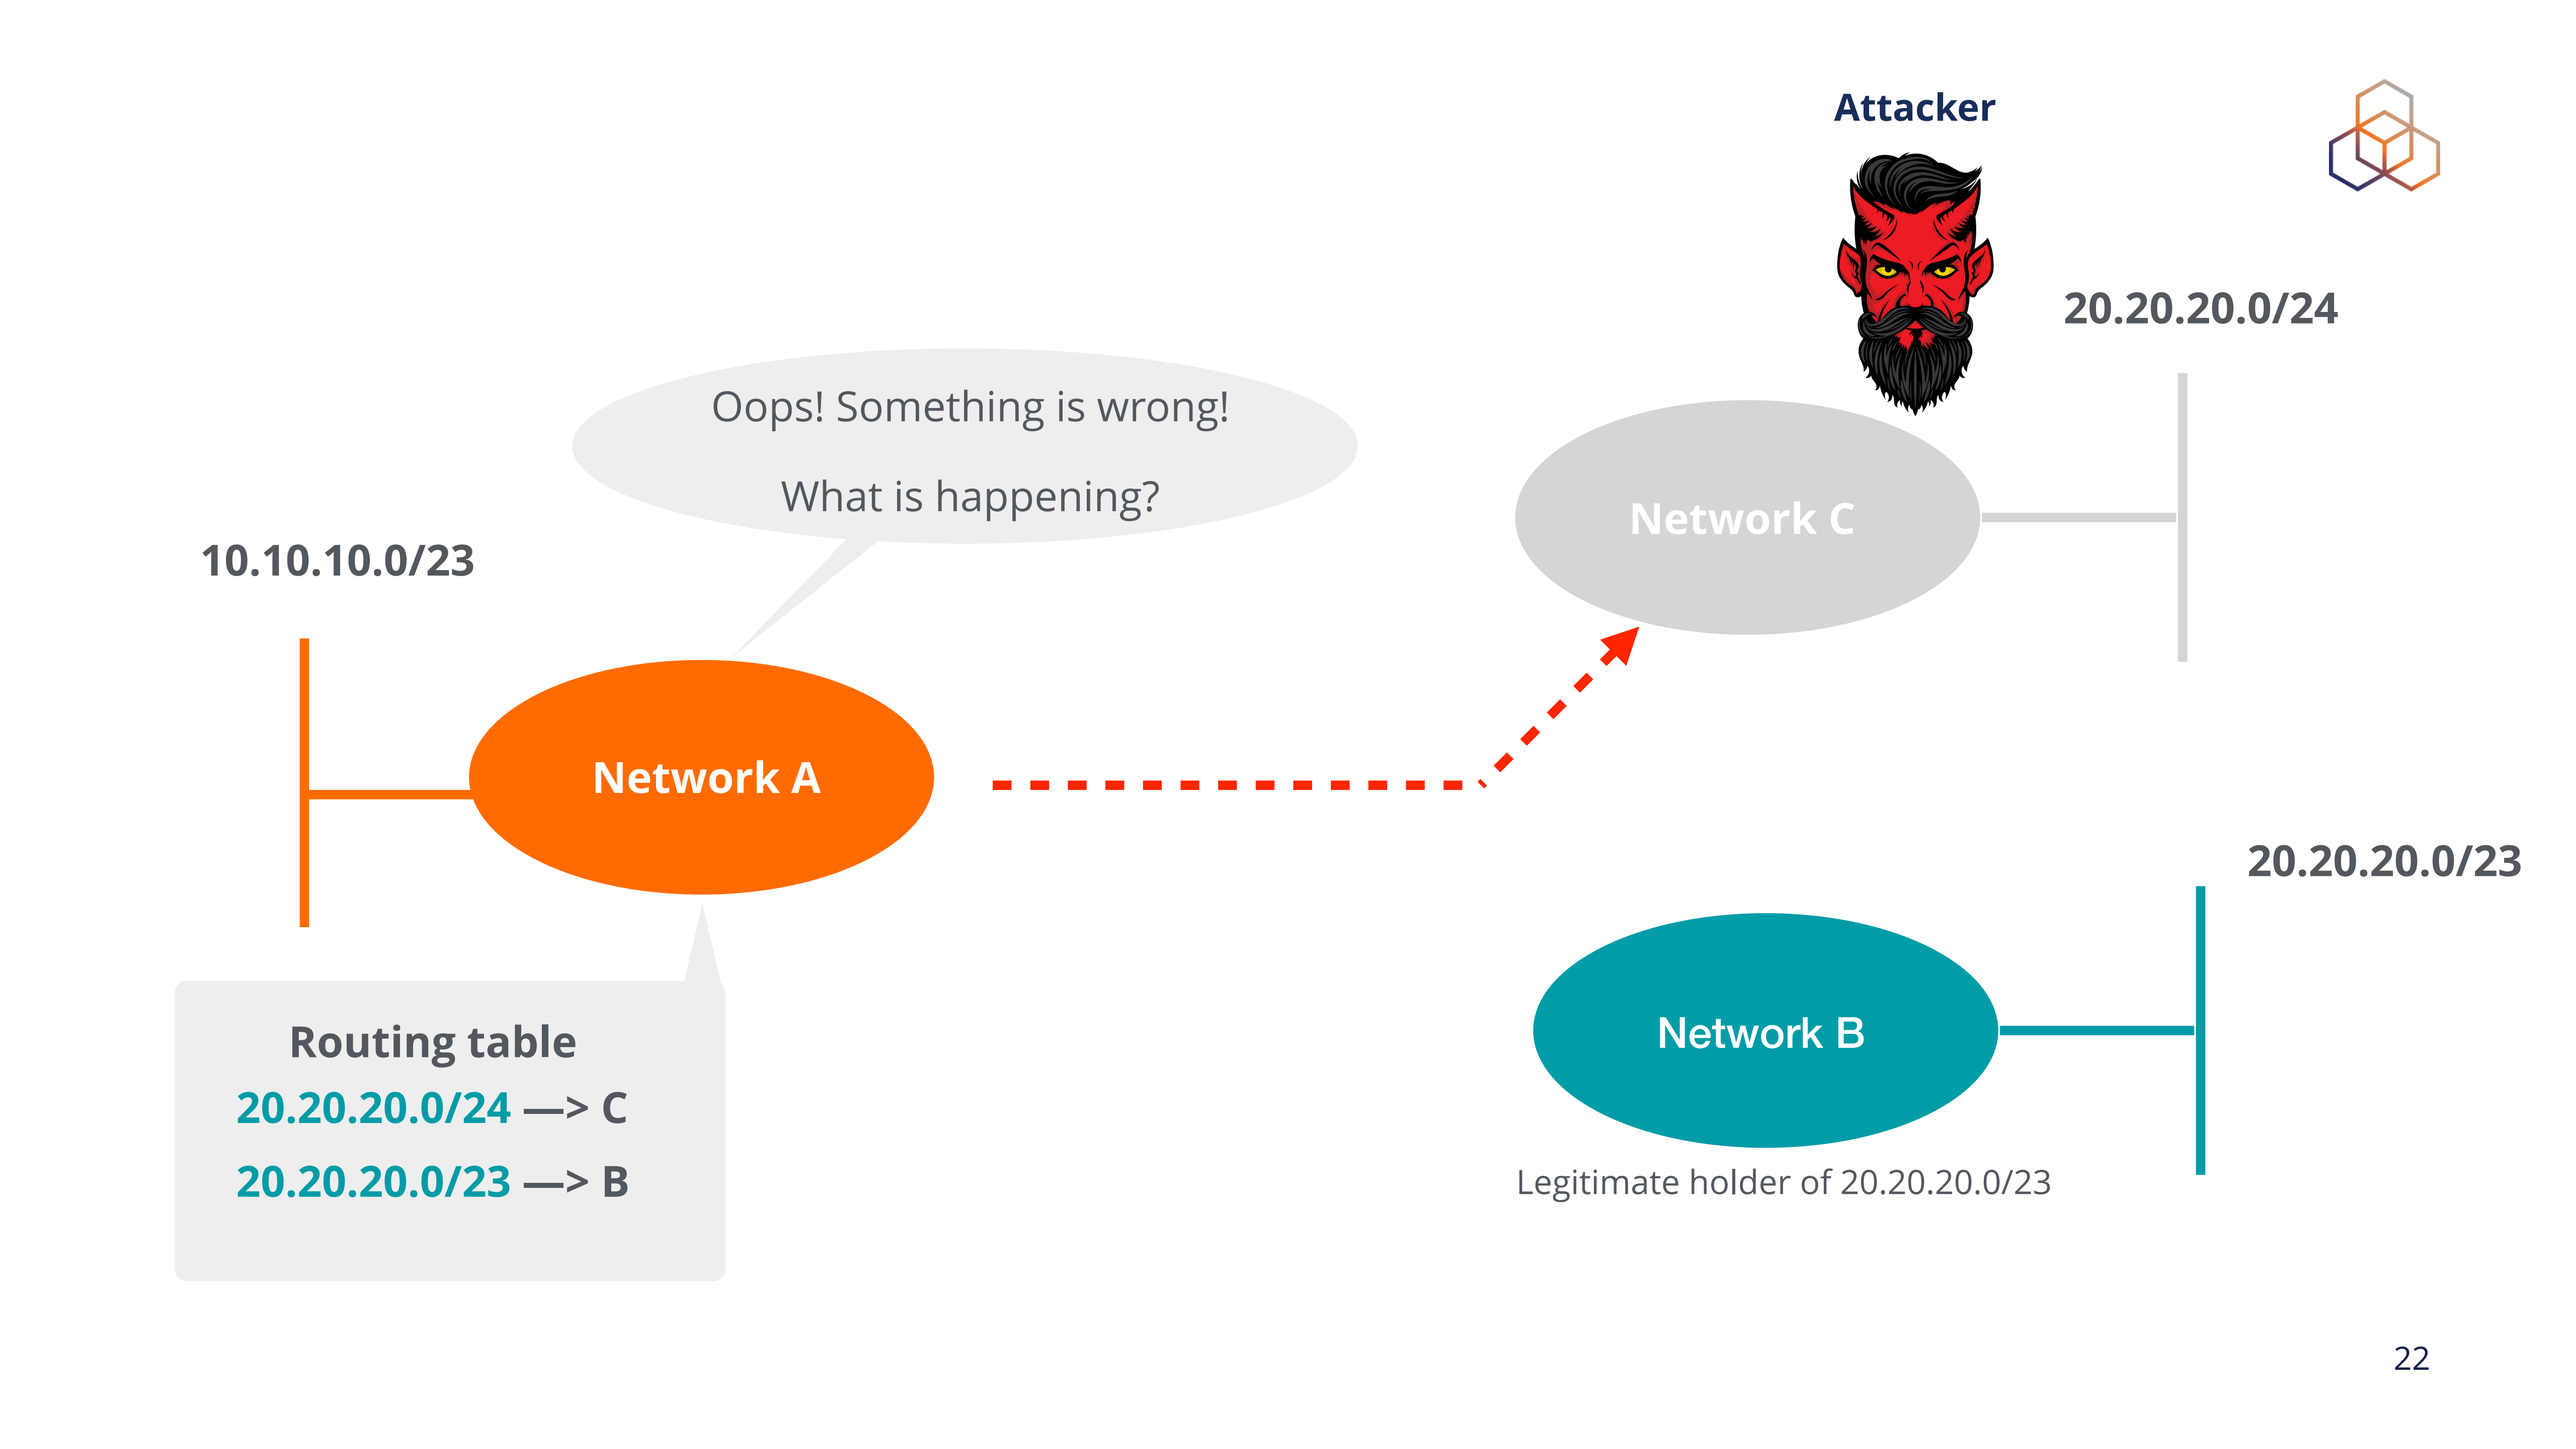
Attacker
20.20.20.0/24
Oops! Something is wrong!
What is happening?
10.10.10.0/23
Network C
Network A
20.20.20.0/23
Routing table
20.20.20.0/24 —> C
20.20.20.0/23 —> B
Network B
Legitimate holder of 20.20.20.0/23

Traffic is blackholed!

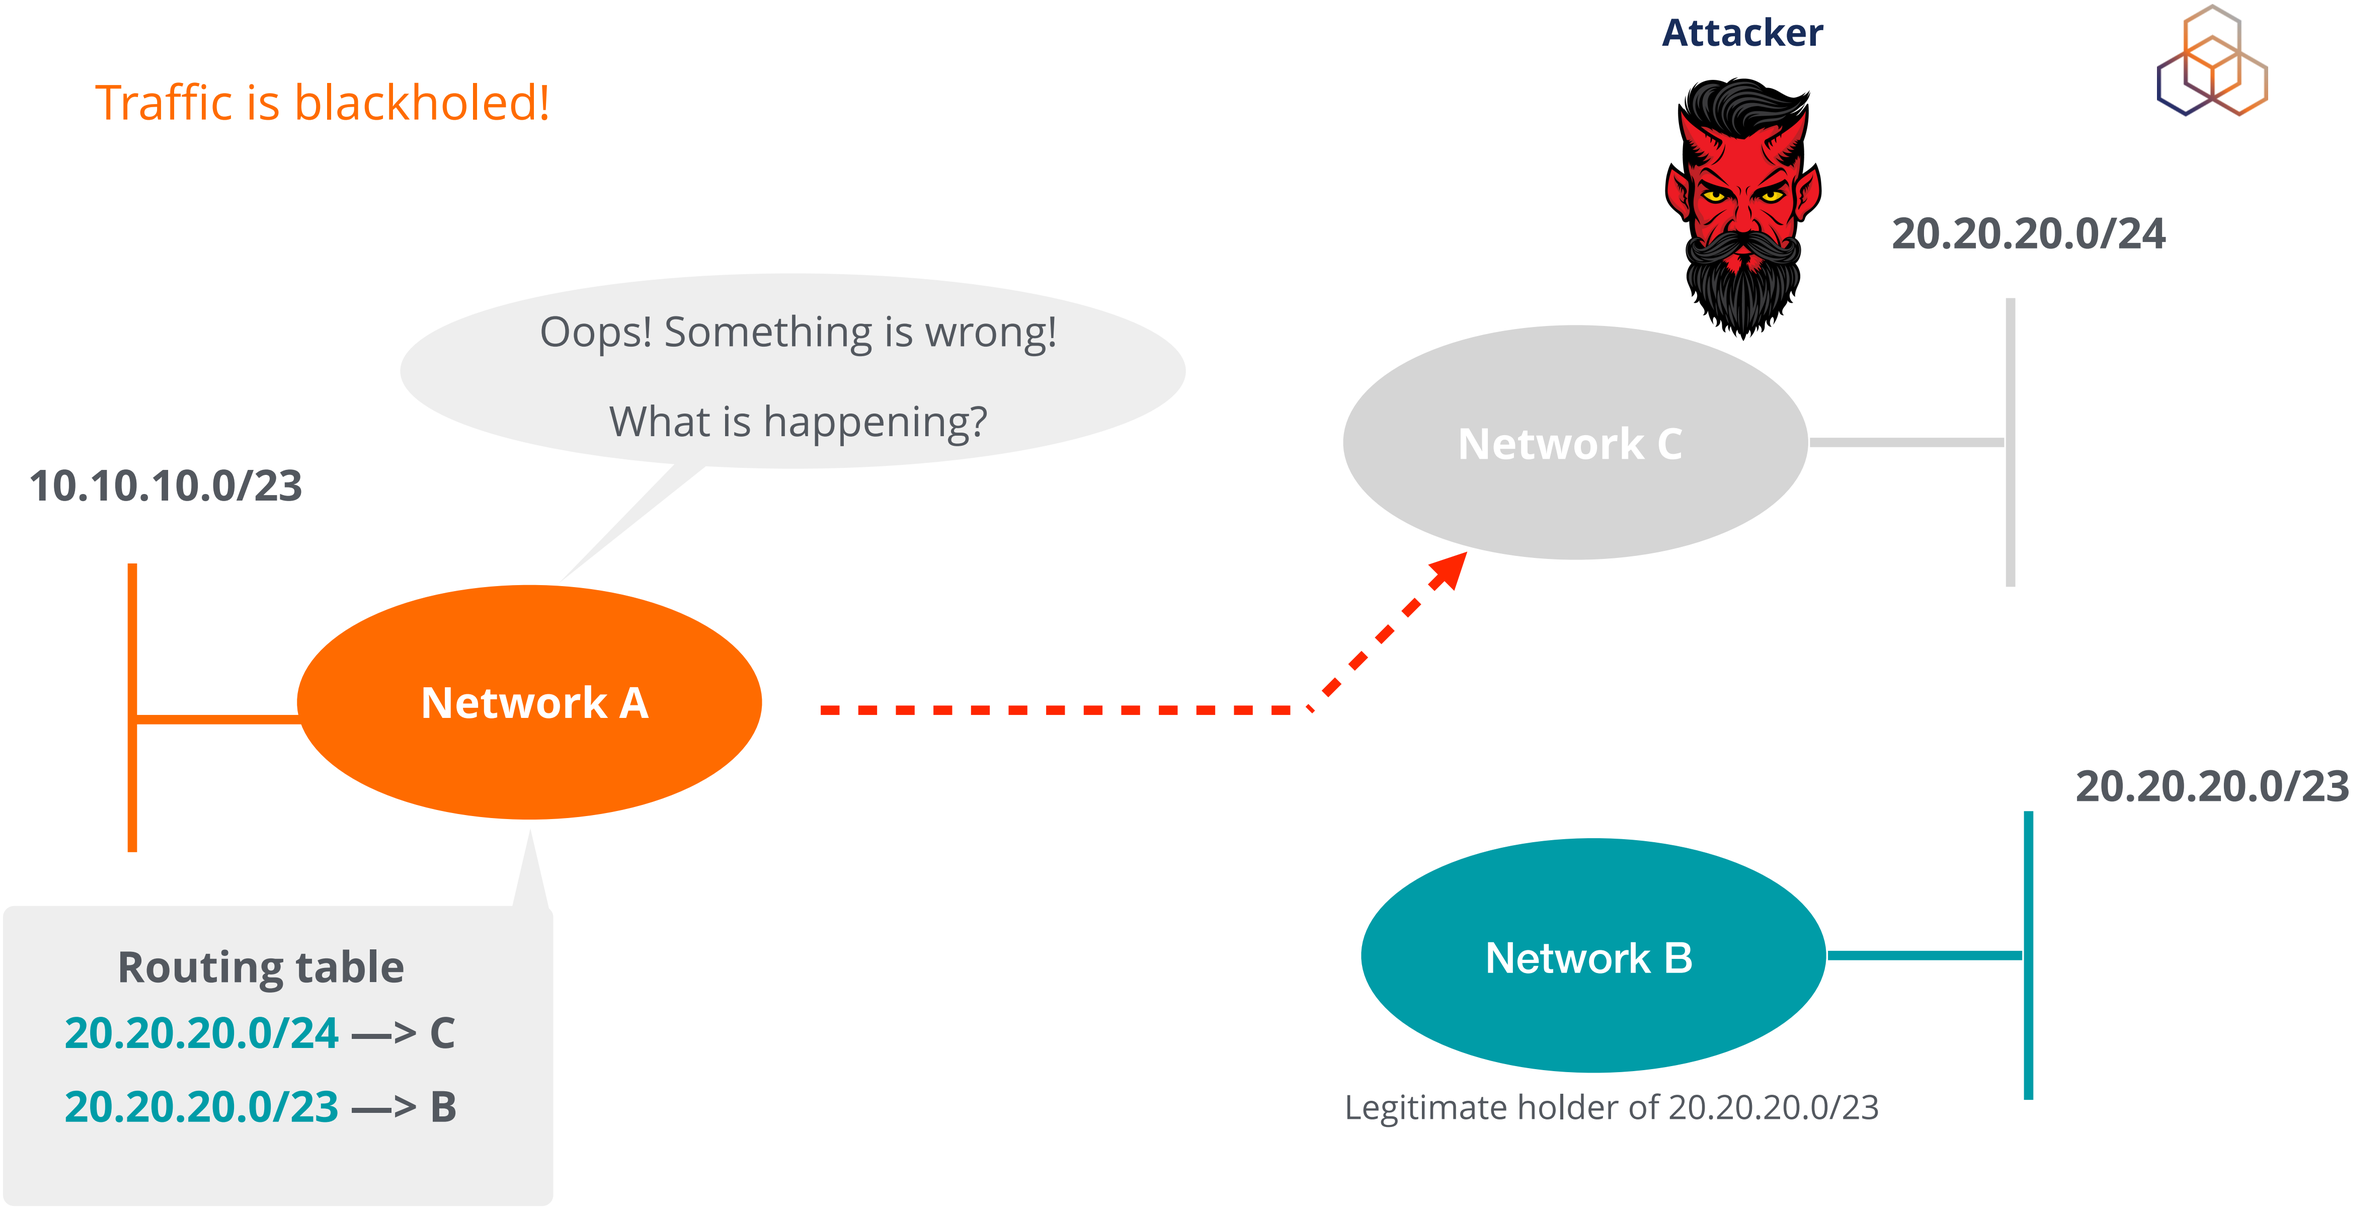

BGP is vulnerable to attacks!

Sometimes it happens accidentally!

- Typing errors
  - Also known as “fat fingers”
  - 2 and 3 are really close on our keyboards…
- Unintentional route leaks
  - Routing policy violations

Accidental or intentional…
Internet Routing Infrastructure is **affected!**

So, how can we secure Internet routing?

Routing Security

How to secure Internet routing?

In order to secure routing…

- We need to find a way
  - to verify whether the prefix is originated by the rightful holder

- But how?
  - By filtering incorrect routing info
  - Registering your routing info in Internet Routing Registries (IRRs)
  - Creating filters based on IRRs
  - Implementing RPKI

How to secure routing with IRR?

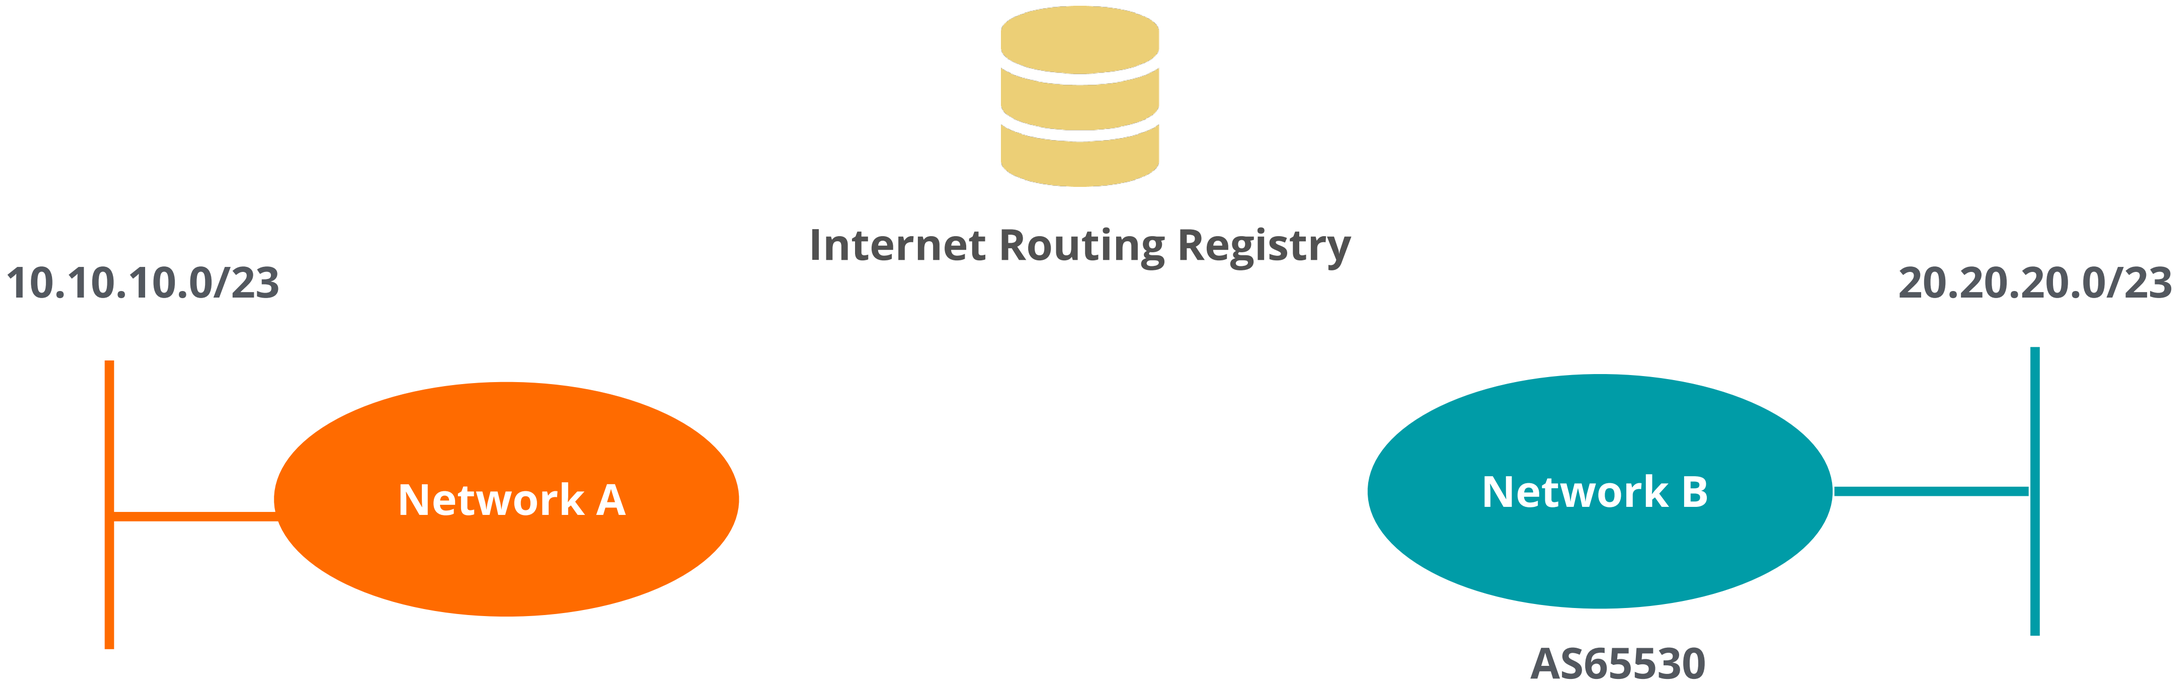

How to secure routing with IRR?

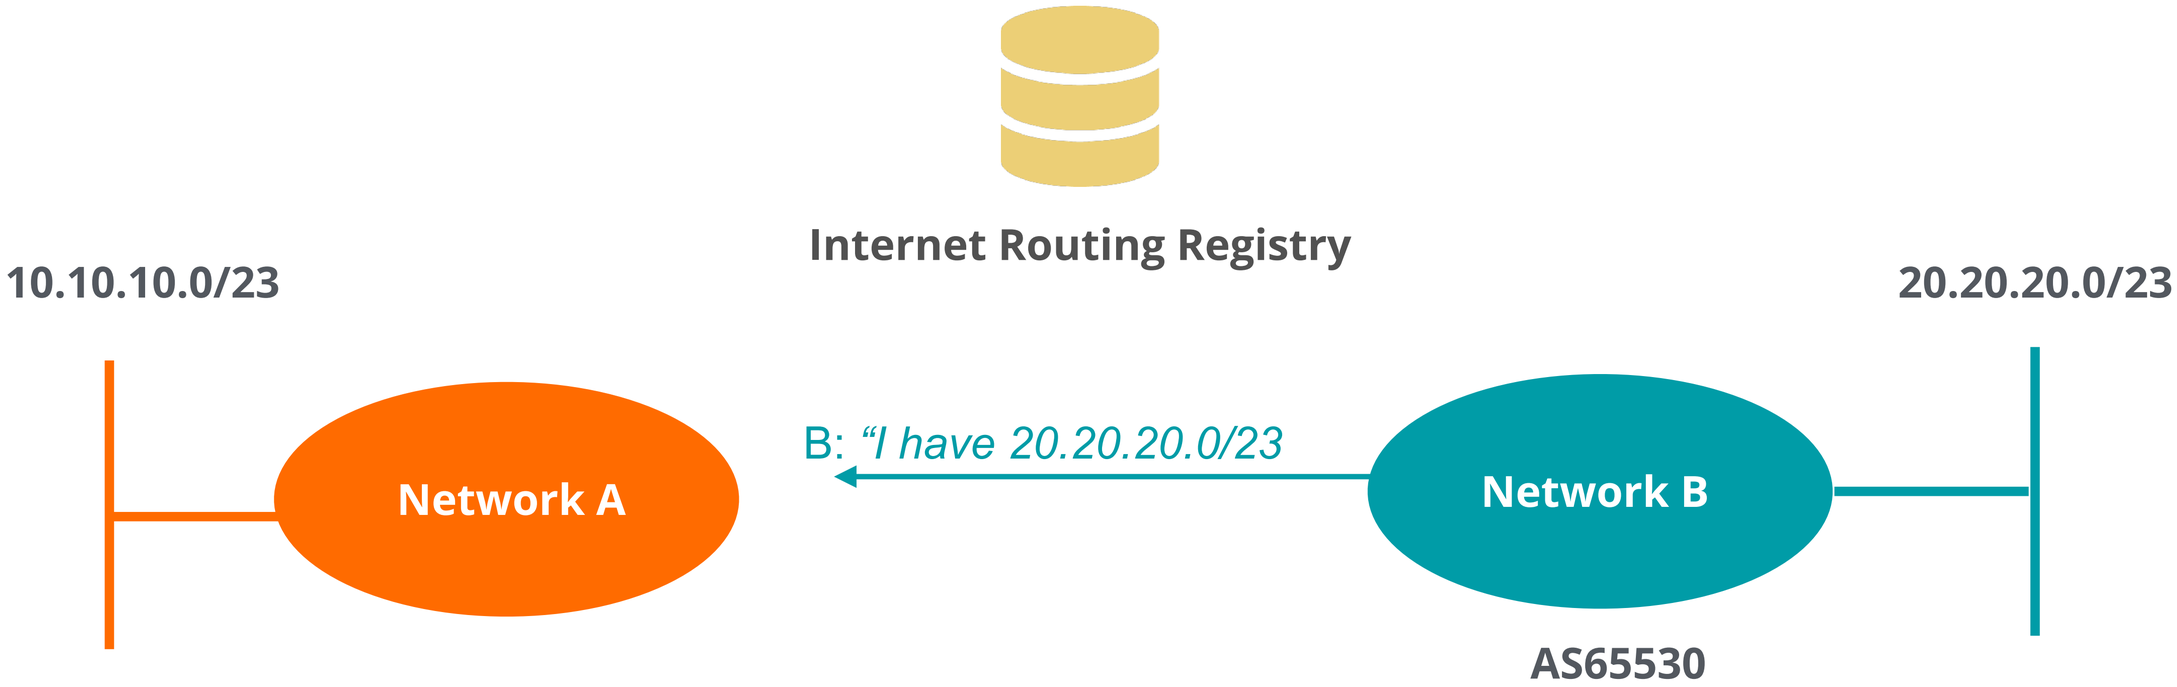

How to secure routing with IRR?

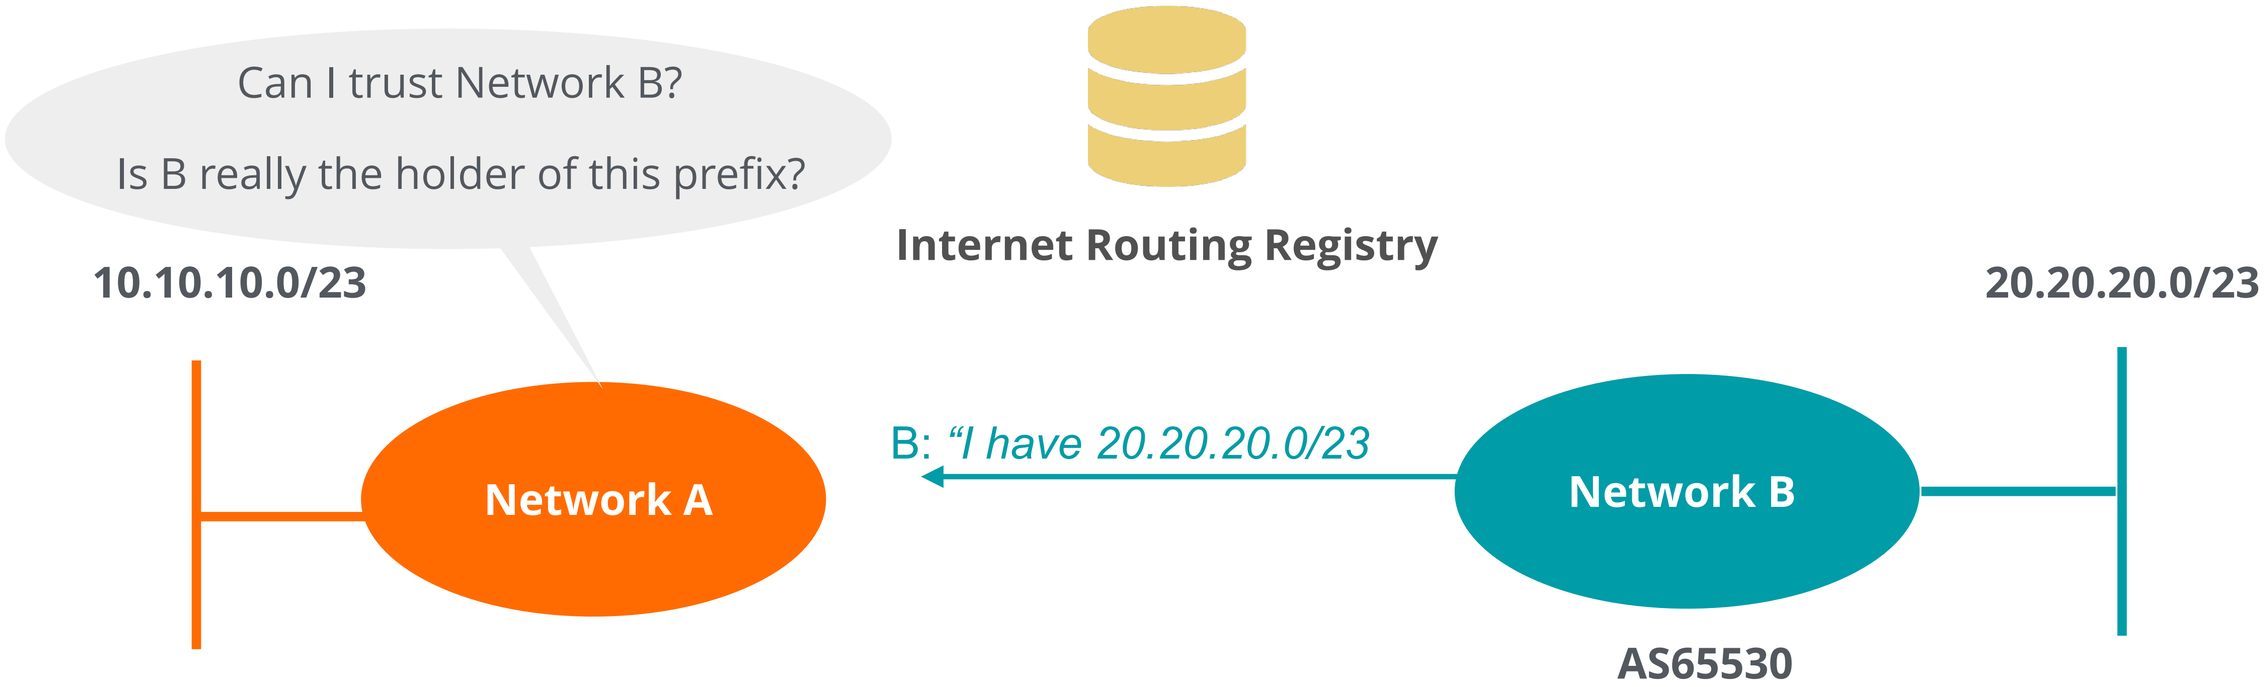

How to secure routing with IRR?

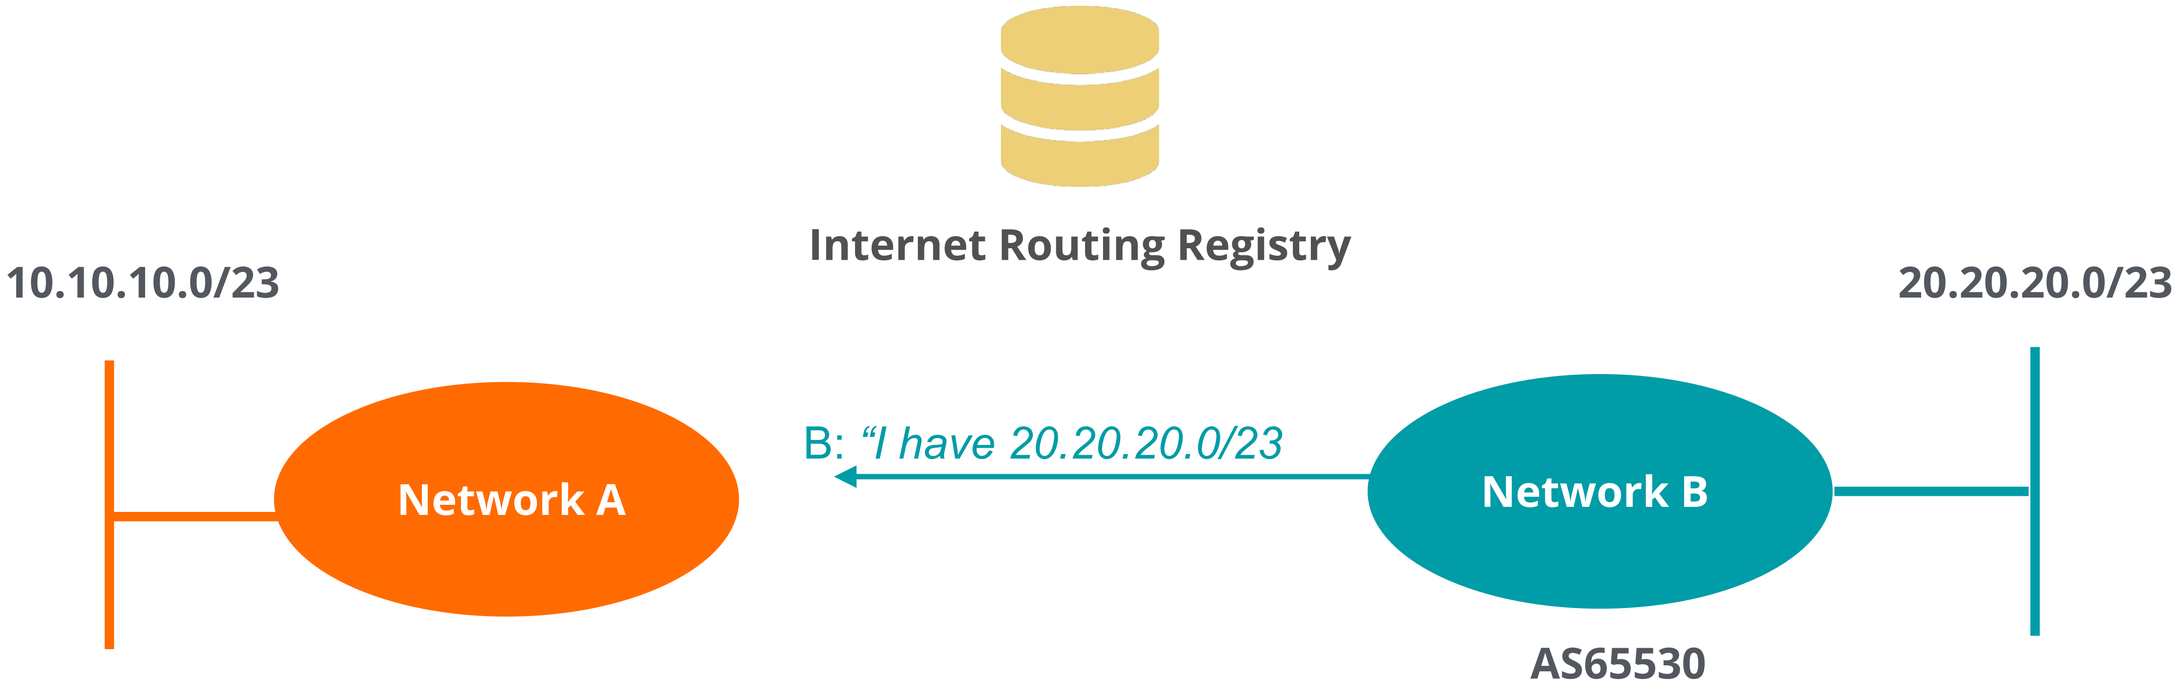

How to secure routing with IRR?

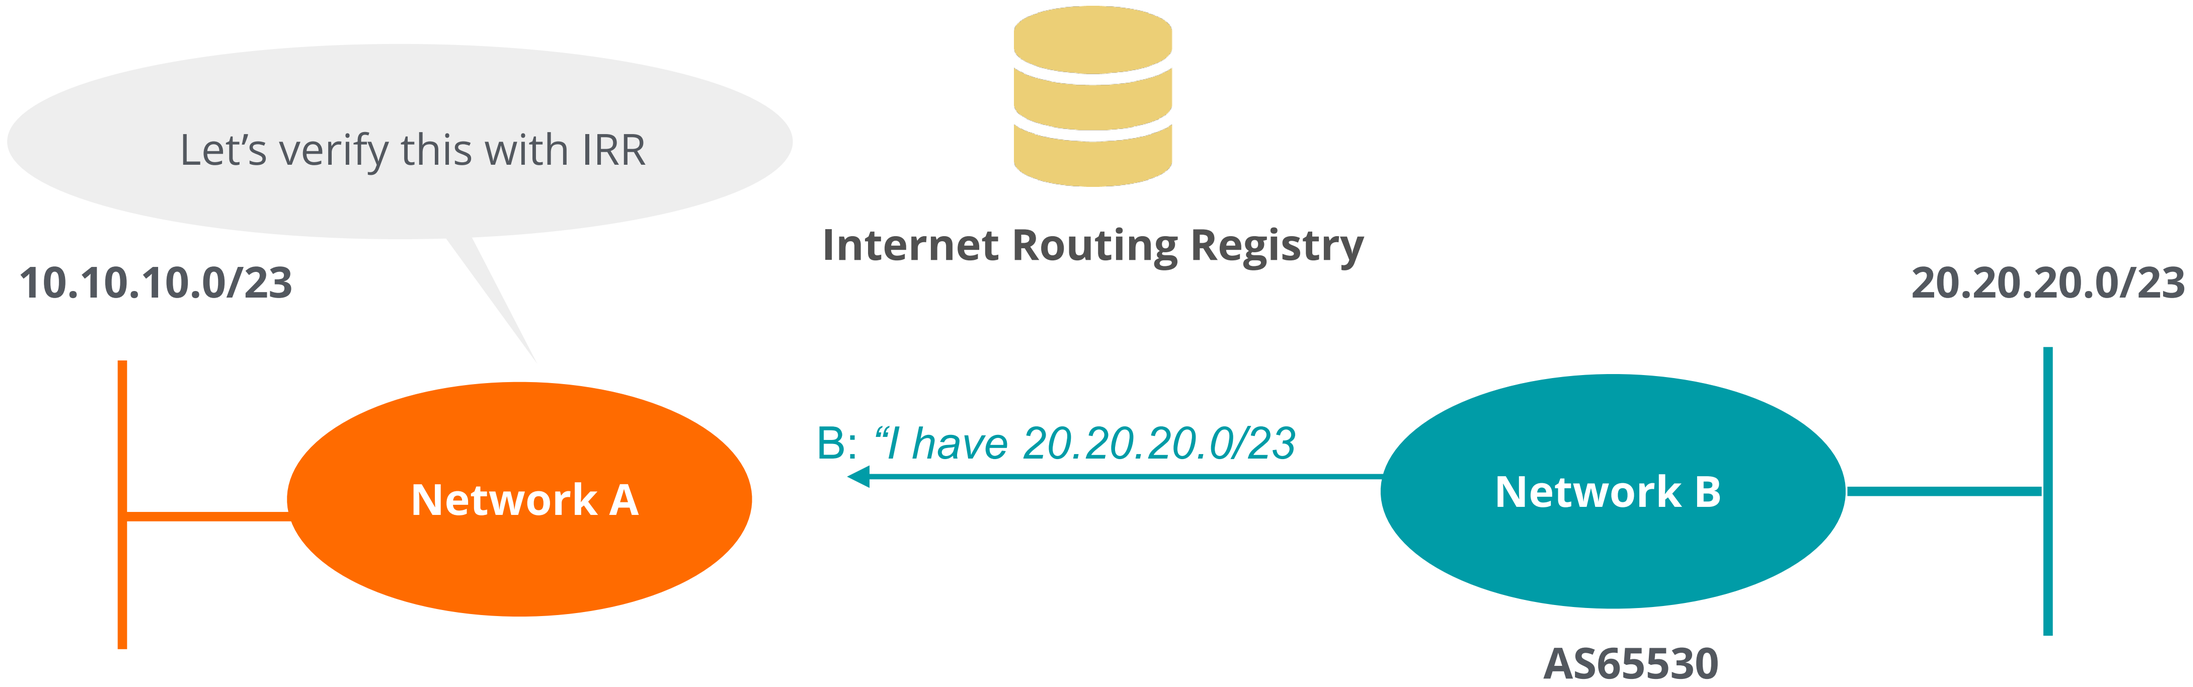

How to secure routing with IRR?

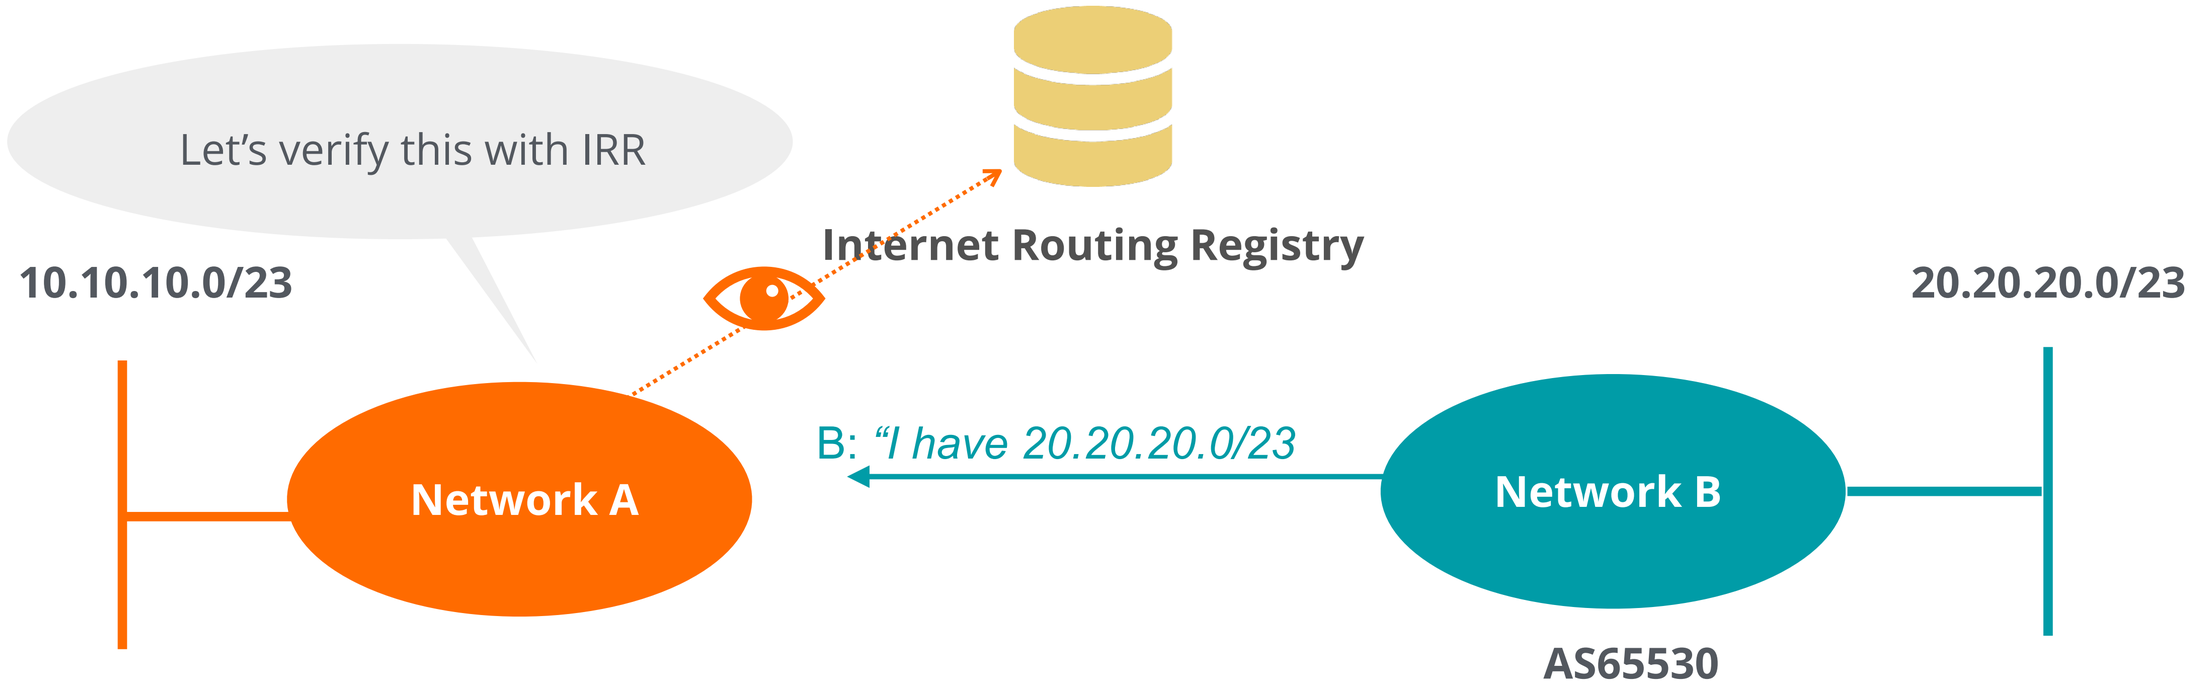

How to secure routing with IRR?

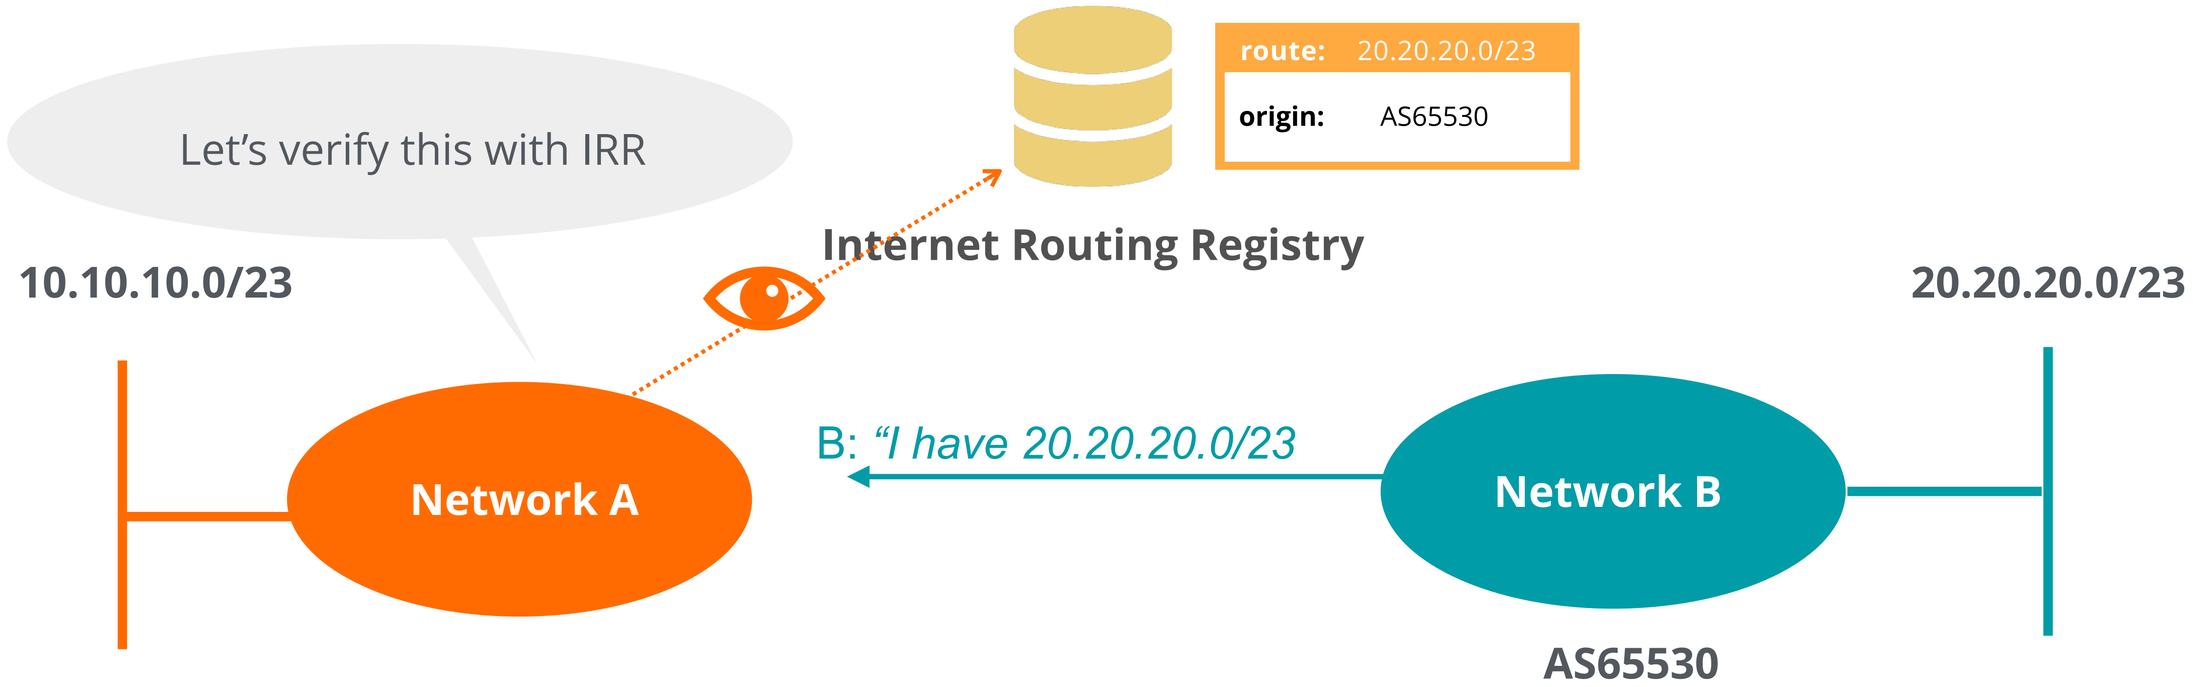

How to secure routing with IRR?

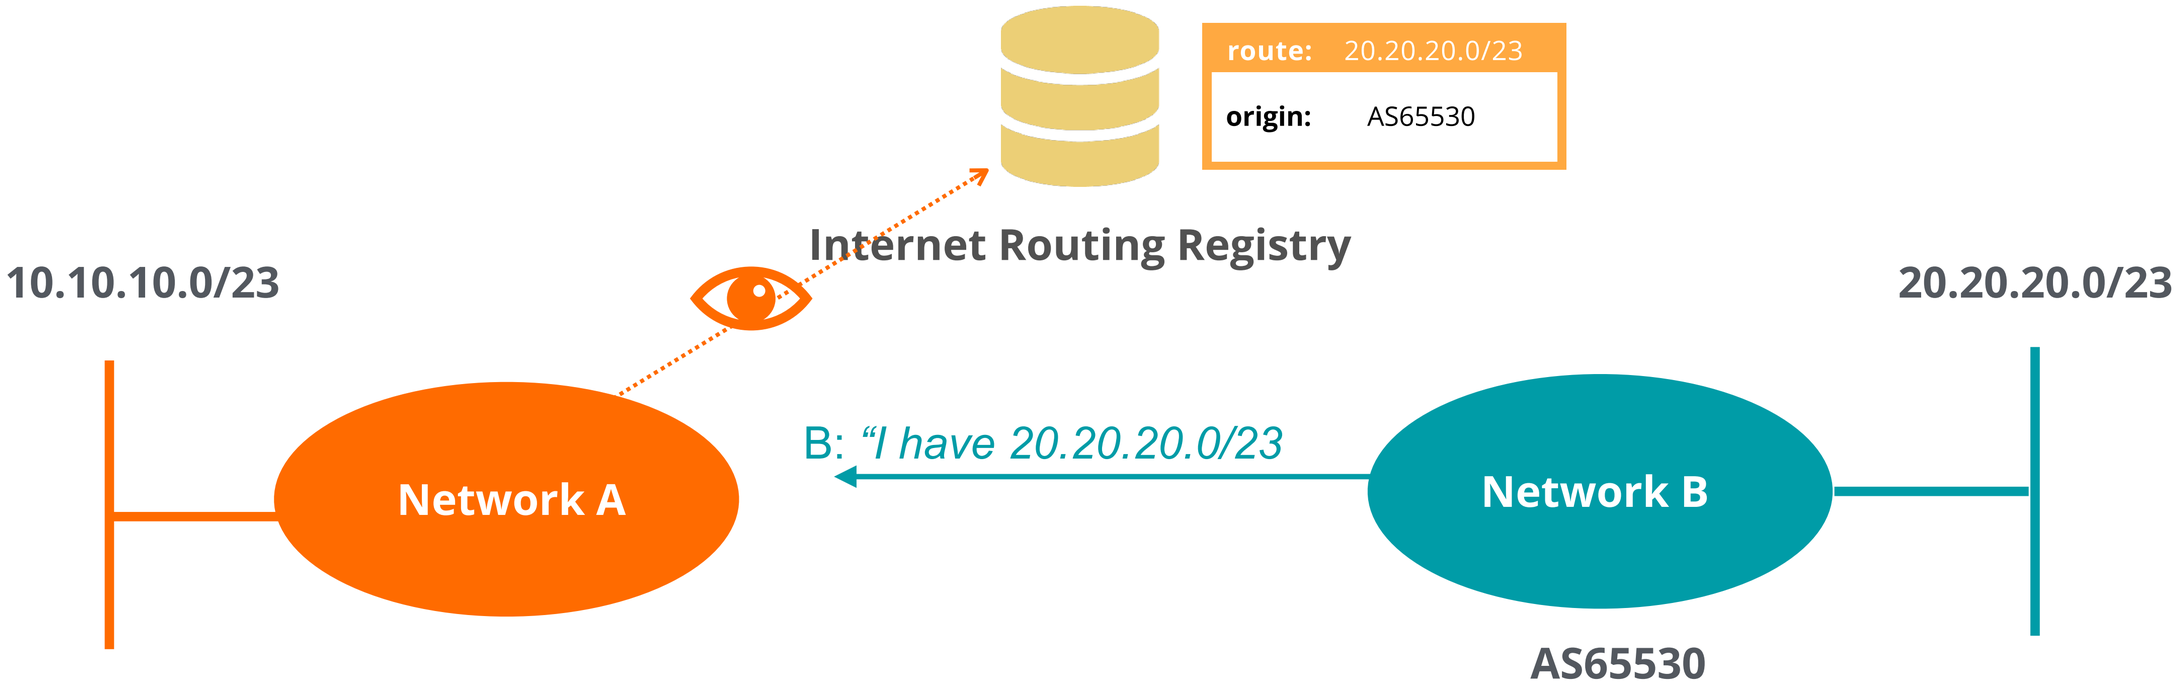

How to secure routing with IRR?

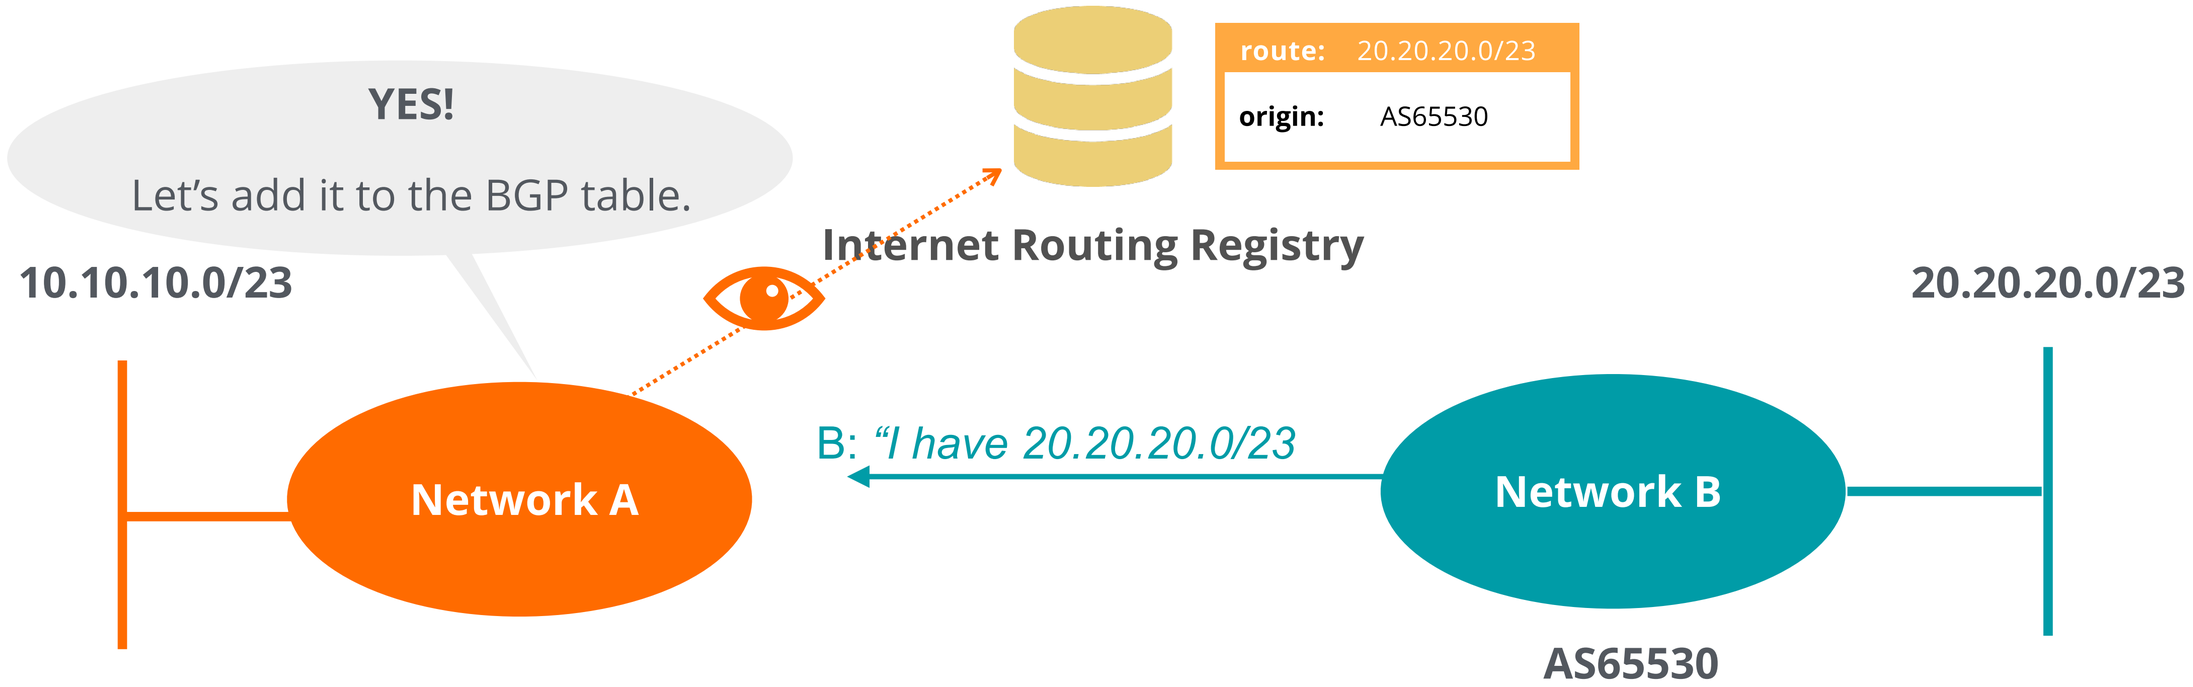

How to secure routing with IRR?

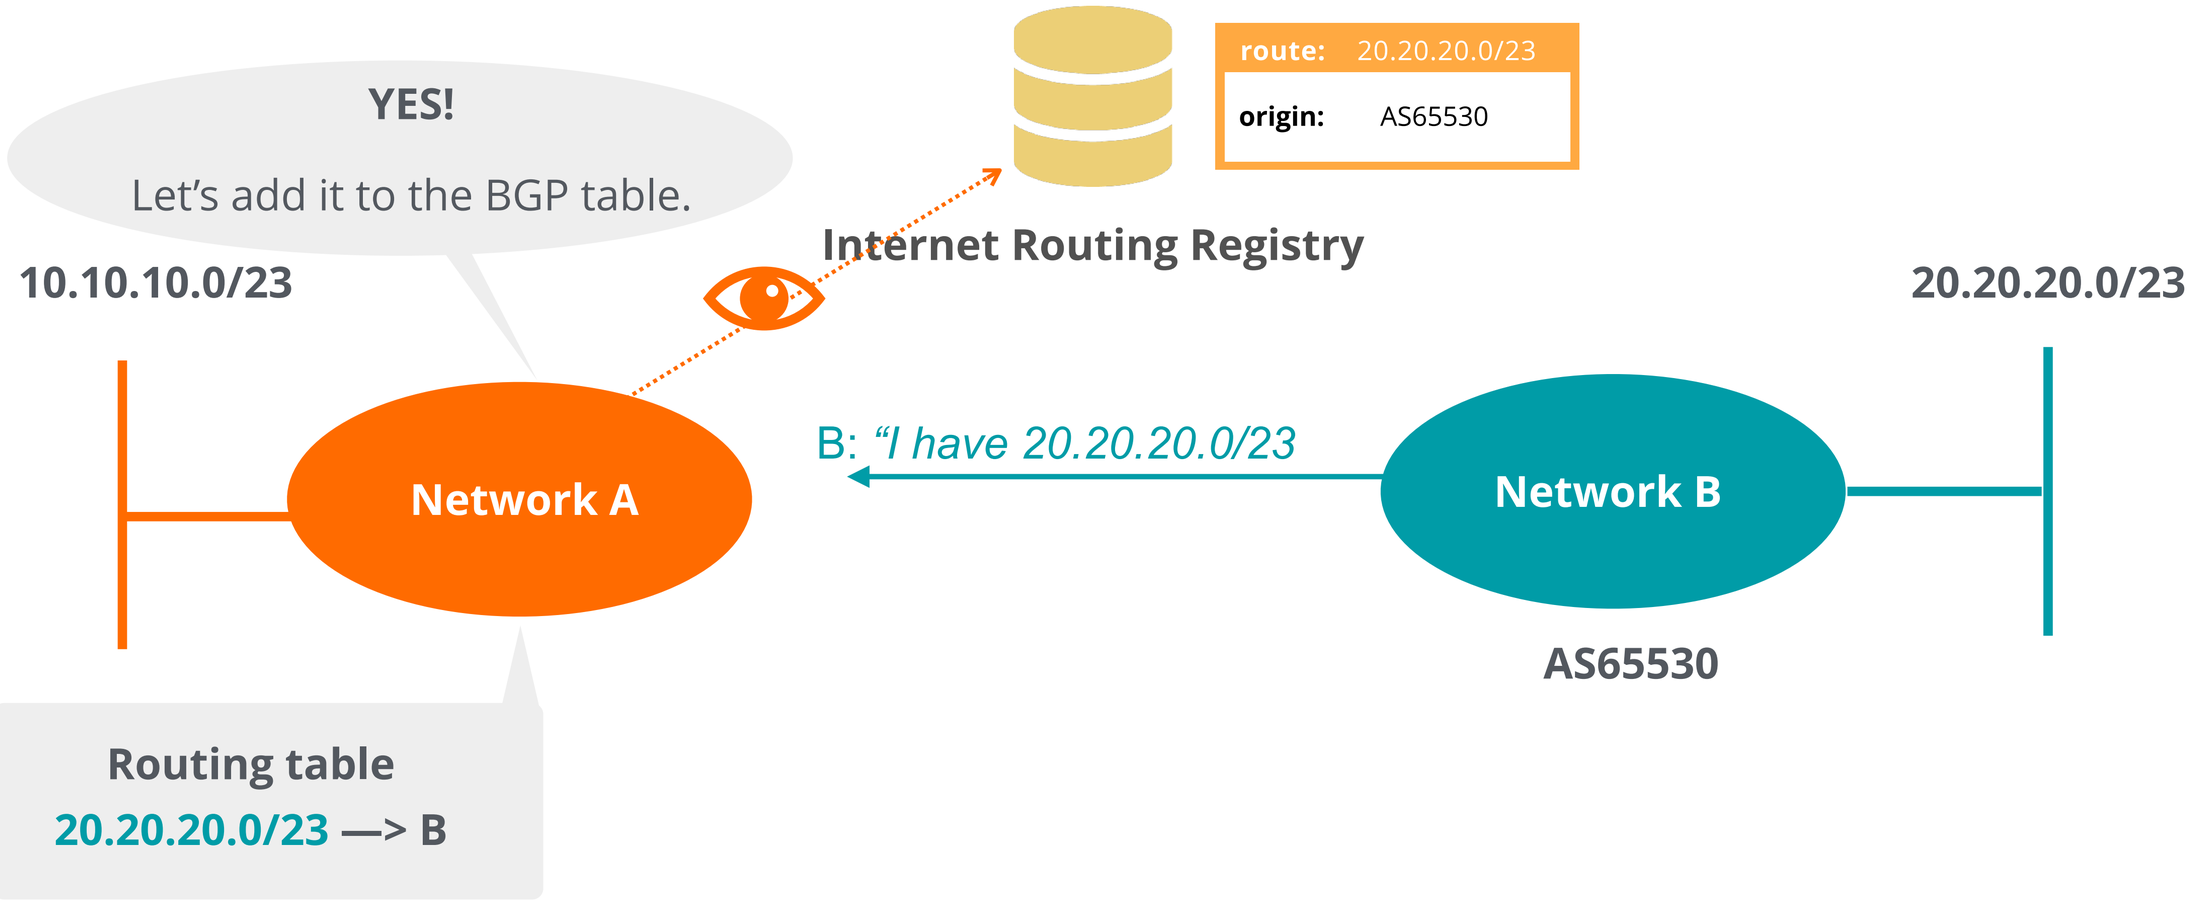

Many upstream providers perform **IRR filtering**.
Create BGP filters based on **route(6)** objects.

If there is an IRR system available to check route origin,
why do we need RPKI?

Because …

There are some issues with the IRR system

- It is not a globally deployed system, just distributed databases
- No central authority, so no authentication for data, anyone can inject anything
- No verification of who holds IPs/ASNs
- Not everybody registers their routing information
- Not updated properly

As a result...

It is not so accurate

Data is incomplete

Not well maintained

That’s why Internet community came up with the **RPKI** solution!

How does RPKI work?

What is RPKI?

What is RPKI?

- Resource Public Key Infrastructure (RPKI)
- Security framework
- Developed by the IETF
- Method to validate the **“origin of BGP announcements”**

Resource Public Key Infrastructure

Ties IP addresses and ASNs to public keys

Follows the RIR hierarchy

Authorised statements from resource holders

- “ASN X is authorised to announce my Prefix Y”
- Signed, holder of Y

RPKI Chain of Trust

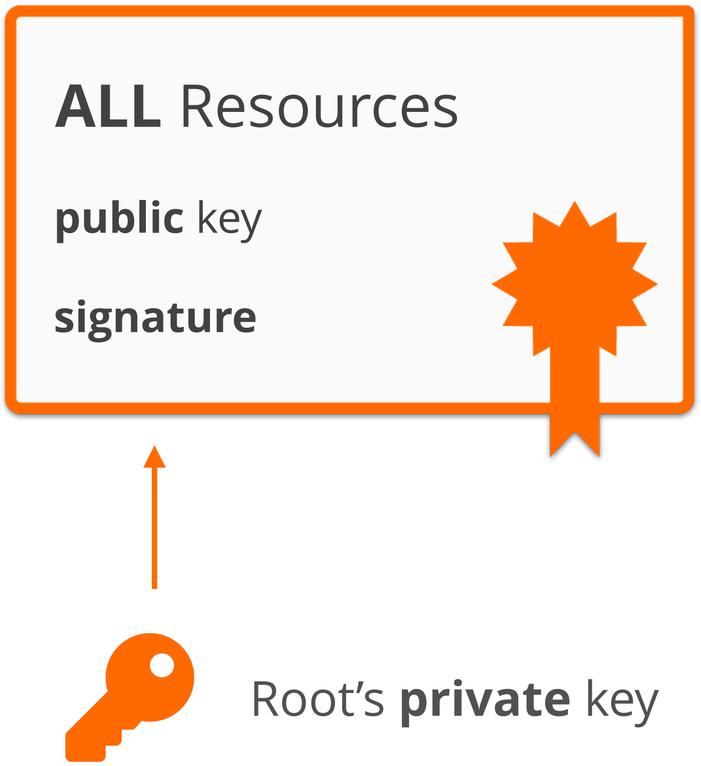

RIPE NCC Root Certificate

Self-signed

Root's **private** key

RPKI Chain of Trust

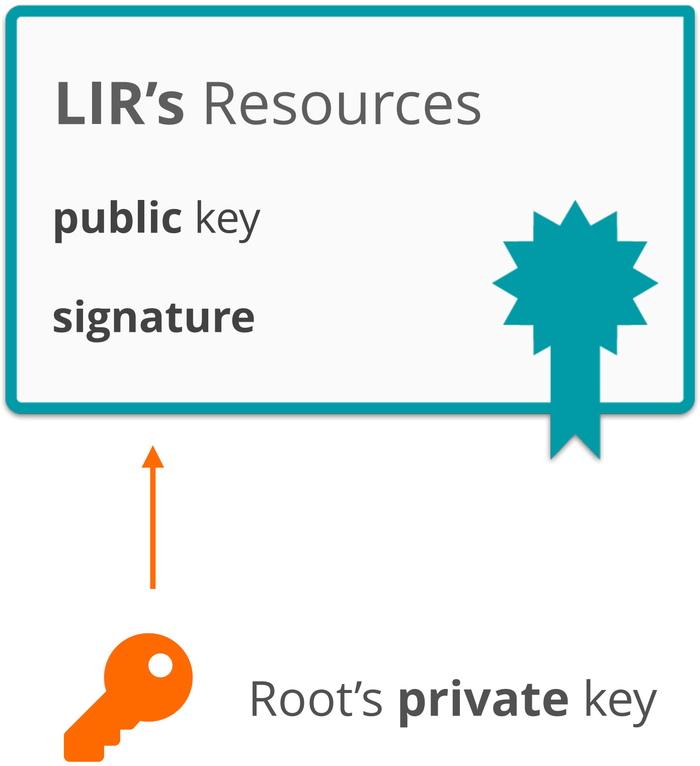

LIR Certificate

Signed by the Root private key

RPKI Chain of Trust

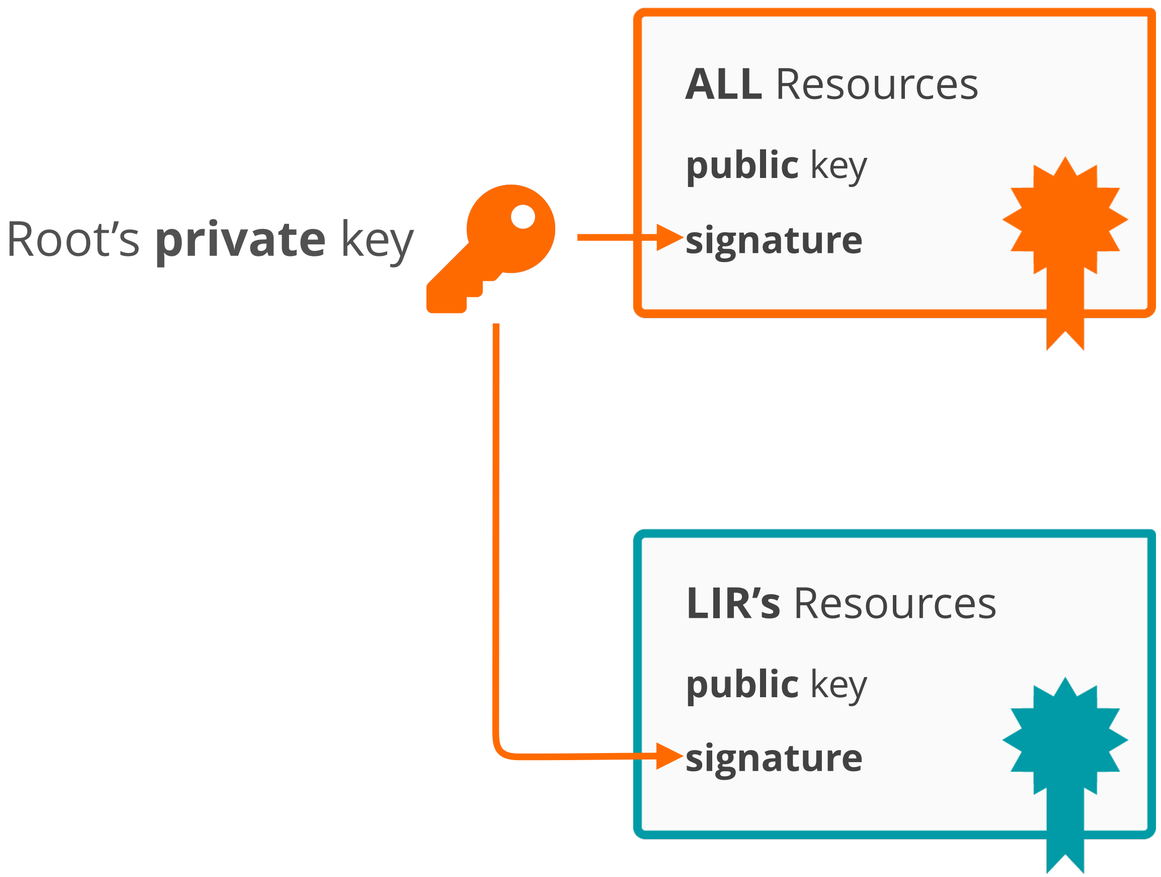

RPKI has two elements

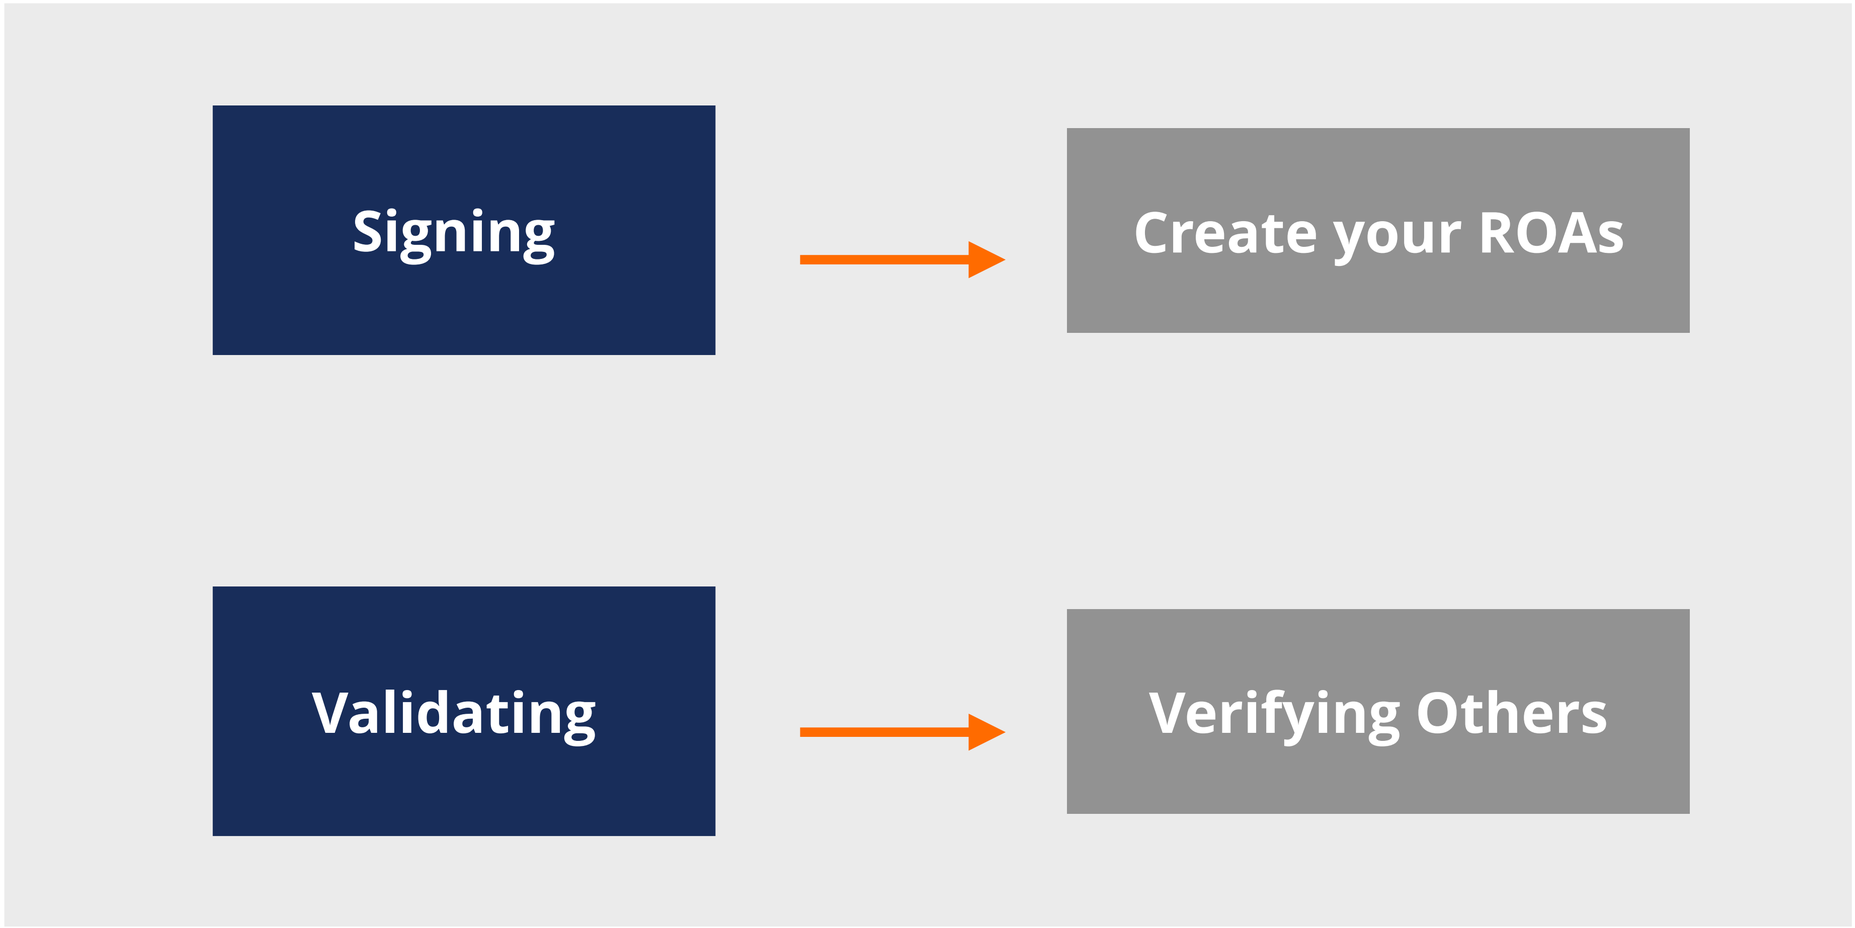

RPKI has two elements

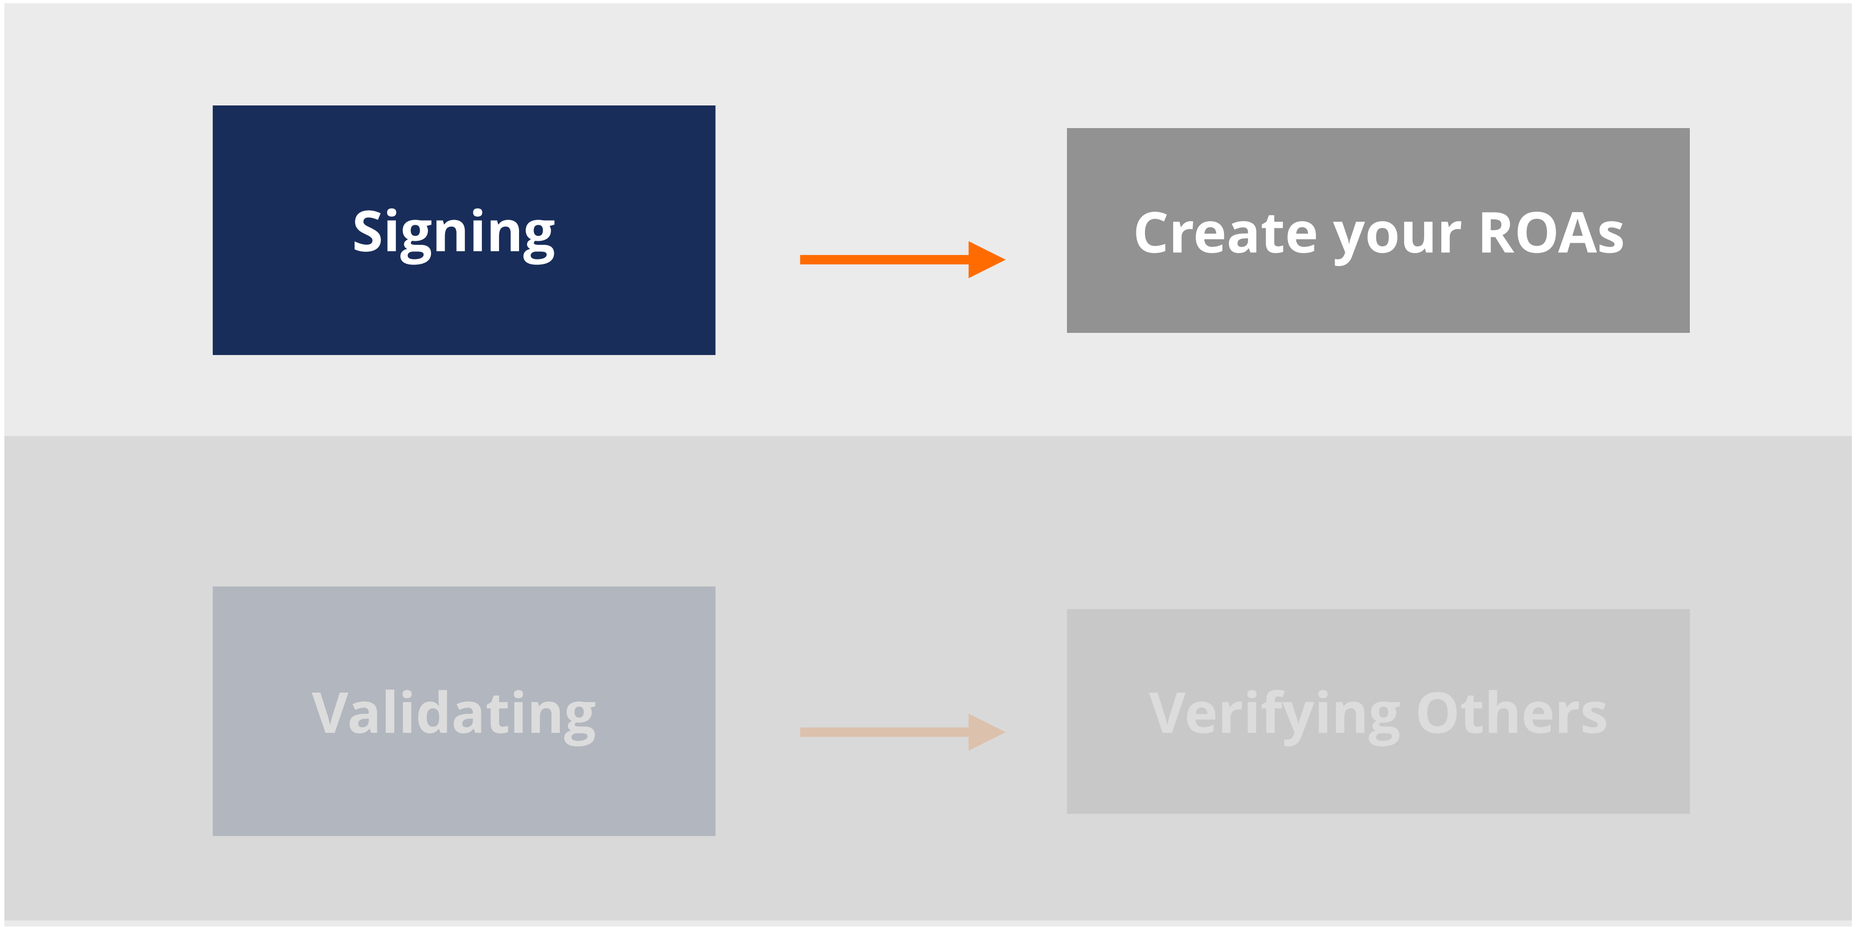

How does RPKI work?

Creating ROAs

ROA (Route Origin Authorization)

- Digitally signed object
- An authorised statement created by the resource holder
- Contains a list of address prefixes and an AS number
- Multiple ROAs can exist for the same prefix

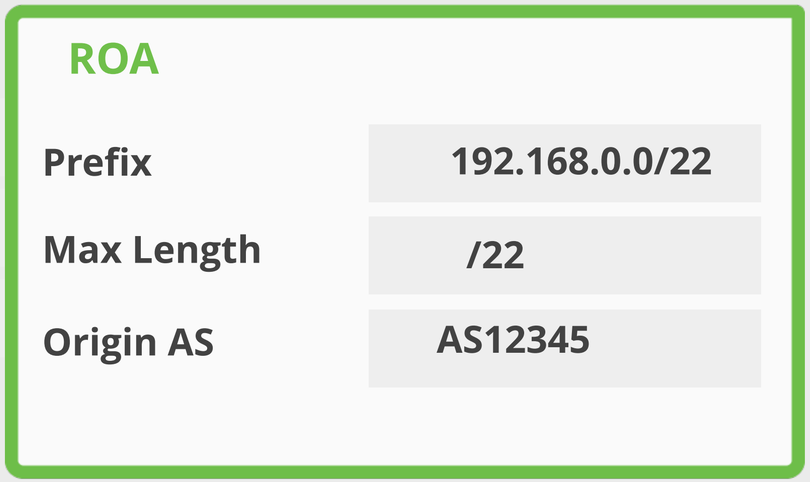

What is in a ROA?

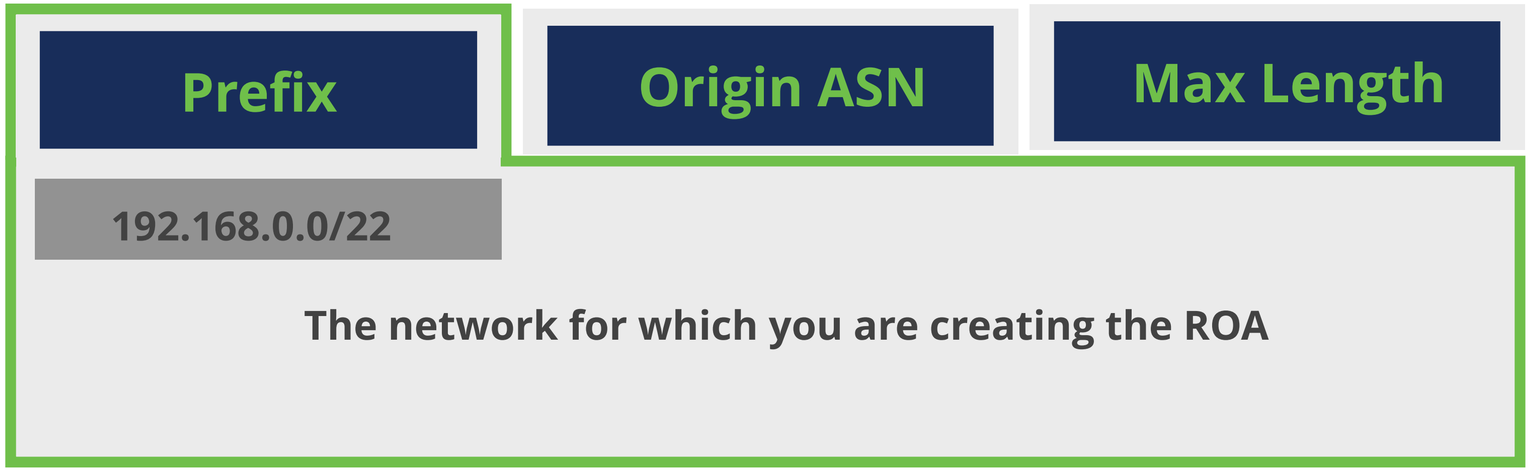

What is in a ROA?

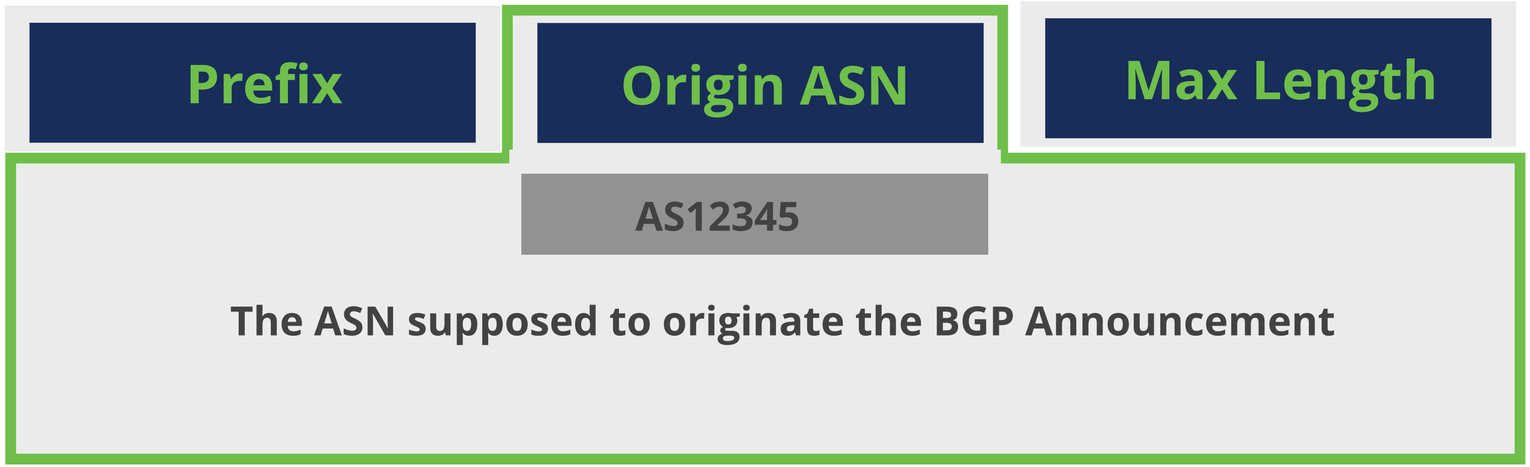

What is in a ROA?

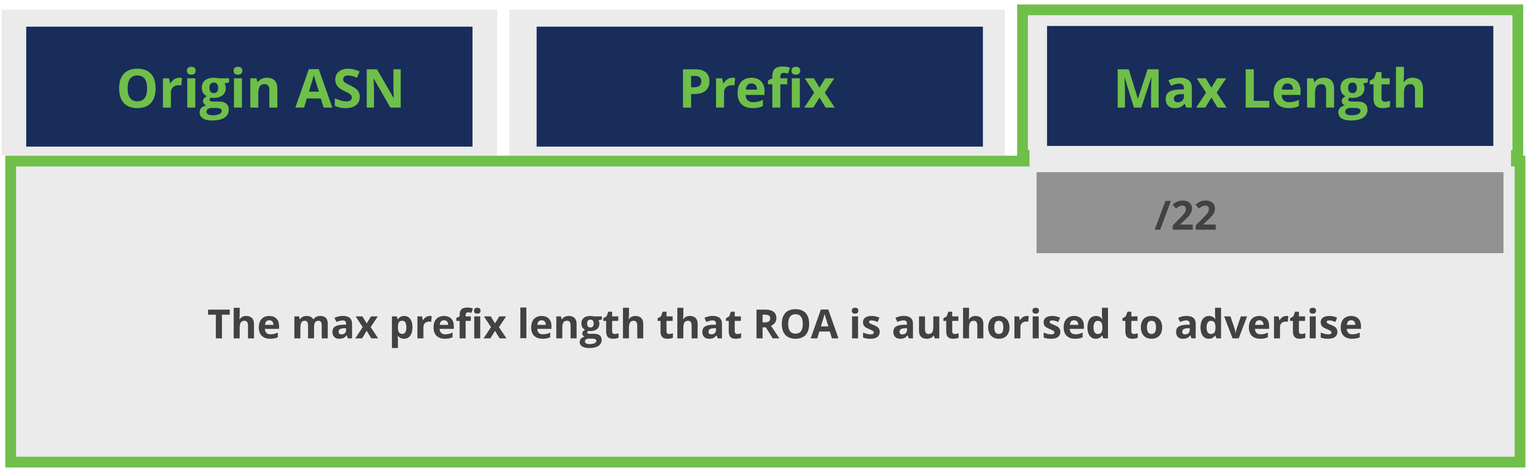

Max-Length

AS3333 has an IP address allocation

193.0.0.0/21

Max-Length

AS3333 has an IP address allocation

AS3333 creates this ROA

193.0.0.0/21

ROA

Prefix	**193.0.0.0/21**
Max Length	**/22**
Origin AS	**AS3333**

Max-Length

AS3333 has an IP address allocation

AS3333 creates this ROA

According to ROA;

193.0.0.0/21

ROA

Prefix	**193.0.0.0/21**
Max Length	**/22**
Origin AS	**AS3333**

Max-Length

AS3333 has an IP address allocation

AS3333 creates this ROA

According to ROA;

/21

193.0.0.0/21

ROA

Prefix	**193.0.0.0/21**
Max Length	**/22**
Origin AS	**AS3333**

Max-Length

AS3333 has an IP address allocation

AS3333 creates this ROA

According to ROA;

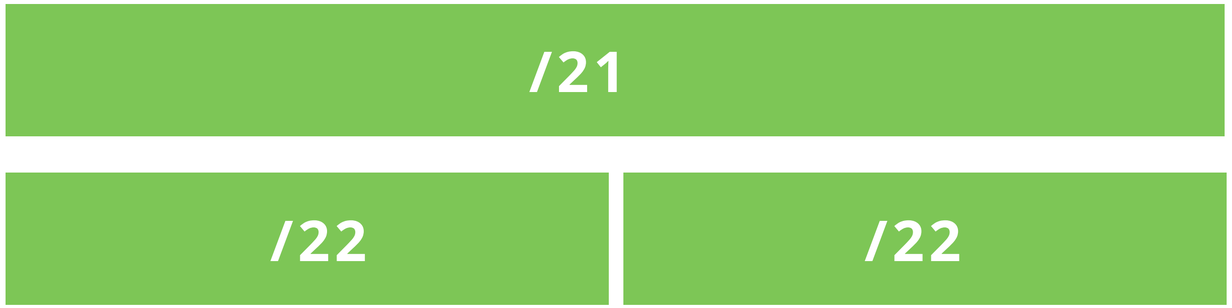

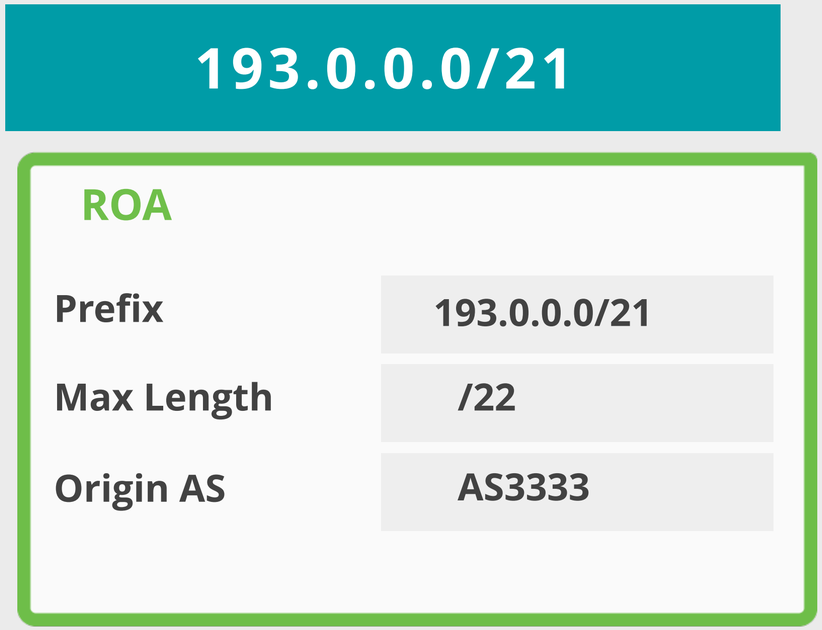

Max-Length

AS3333 has an IP address allocation

AS3333 creates this ROA

According to ROA;

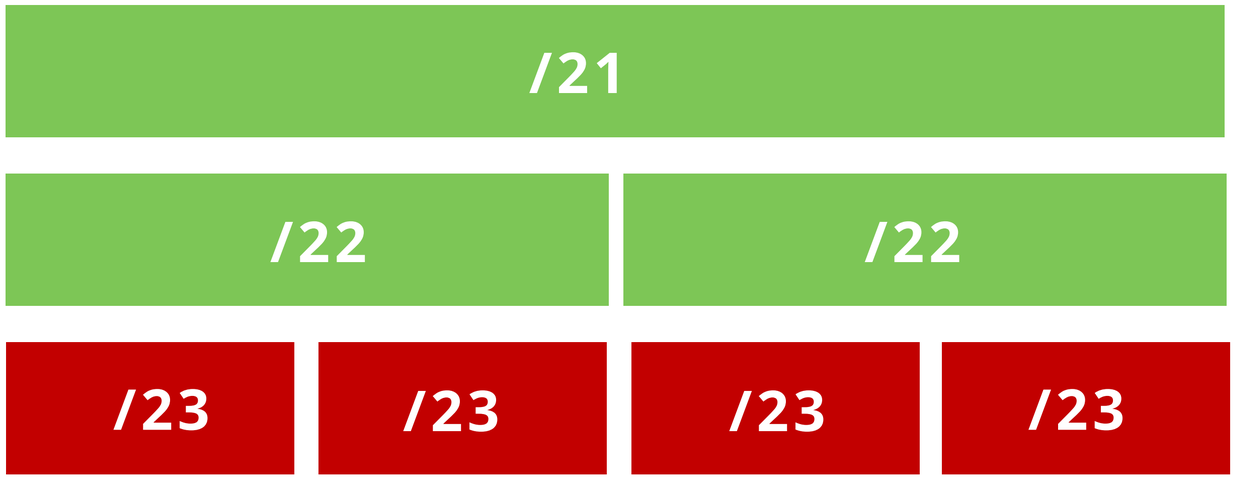

193.0.0.0/21

ROA

Prefix	**193.0.0.0/21**
Max Length	**/22**
Origin AS	**AS3333**

Max-Length

AS3333 has an IP address allocation

AS3333 creates this ROA

According to ROA;

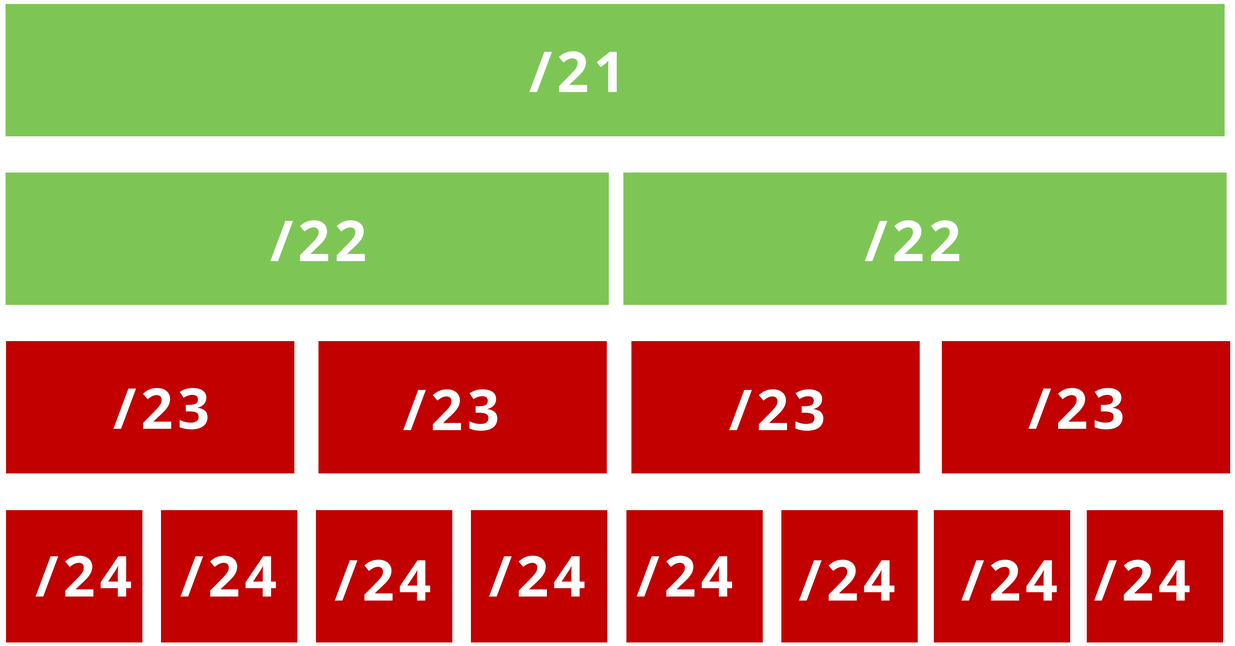

193.0.0.0/21

ROA

Prefix	**193.0.0.0/21**
Max Length	**/22**
Origin AS	**AS3333**

Max-Length

AS3333 has an IP address allocation

AS3333 creates this ROA

According to ROA;

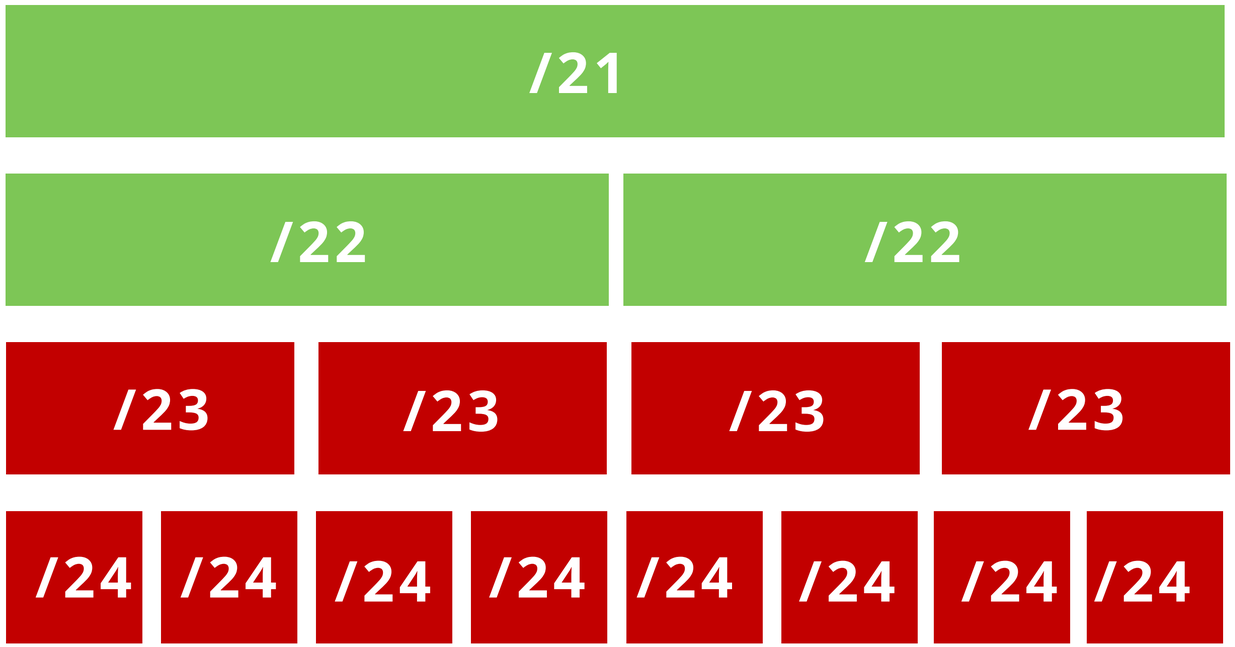

193.0.0.0/21

ROA

Prefix	**193.0.0.0/21**
Max Length	**/22**
Origin AS	**AS3333**

Any more specific announcements are unauthorised by the ROA.

How should we use max-length?

You created a single ROA authorising the entire /22

How should we use max-length?

You created a single ROA authorising the entire /22

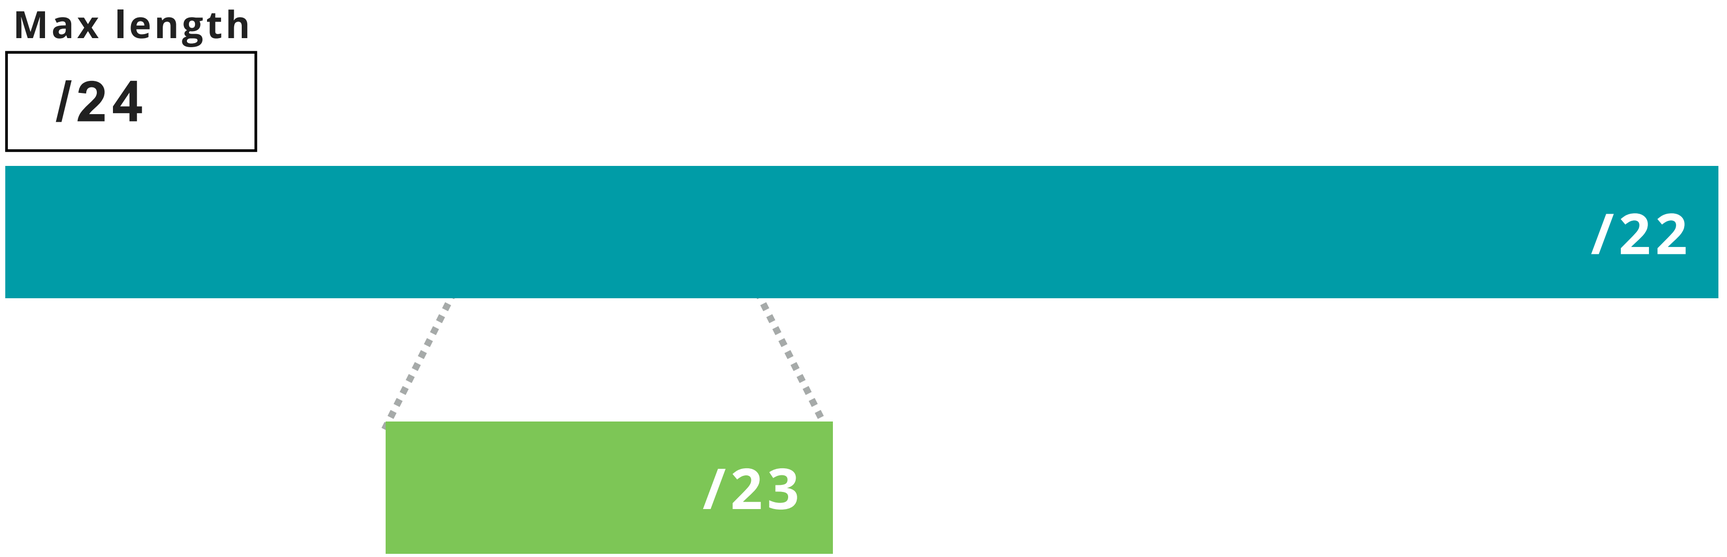

How should we use max-length?

You created a single ROA authorising the entire /22

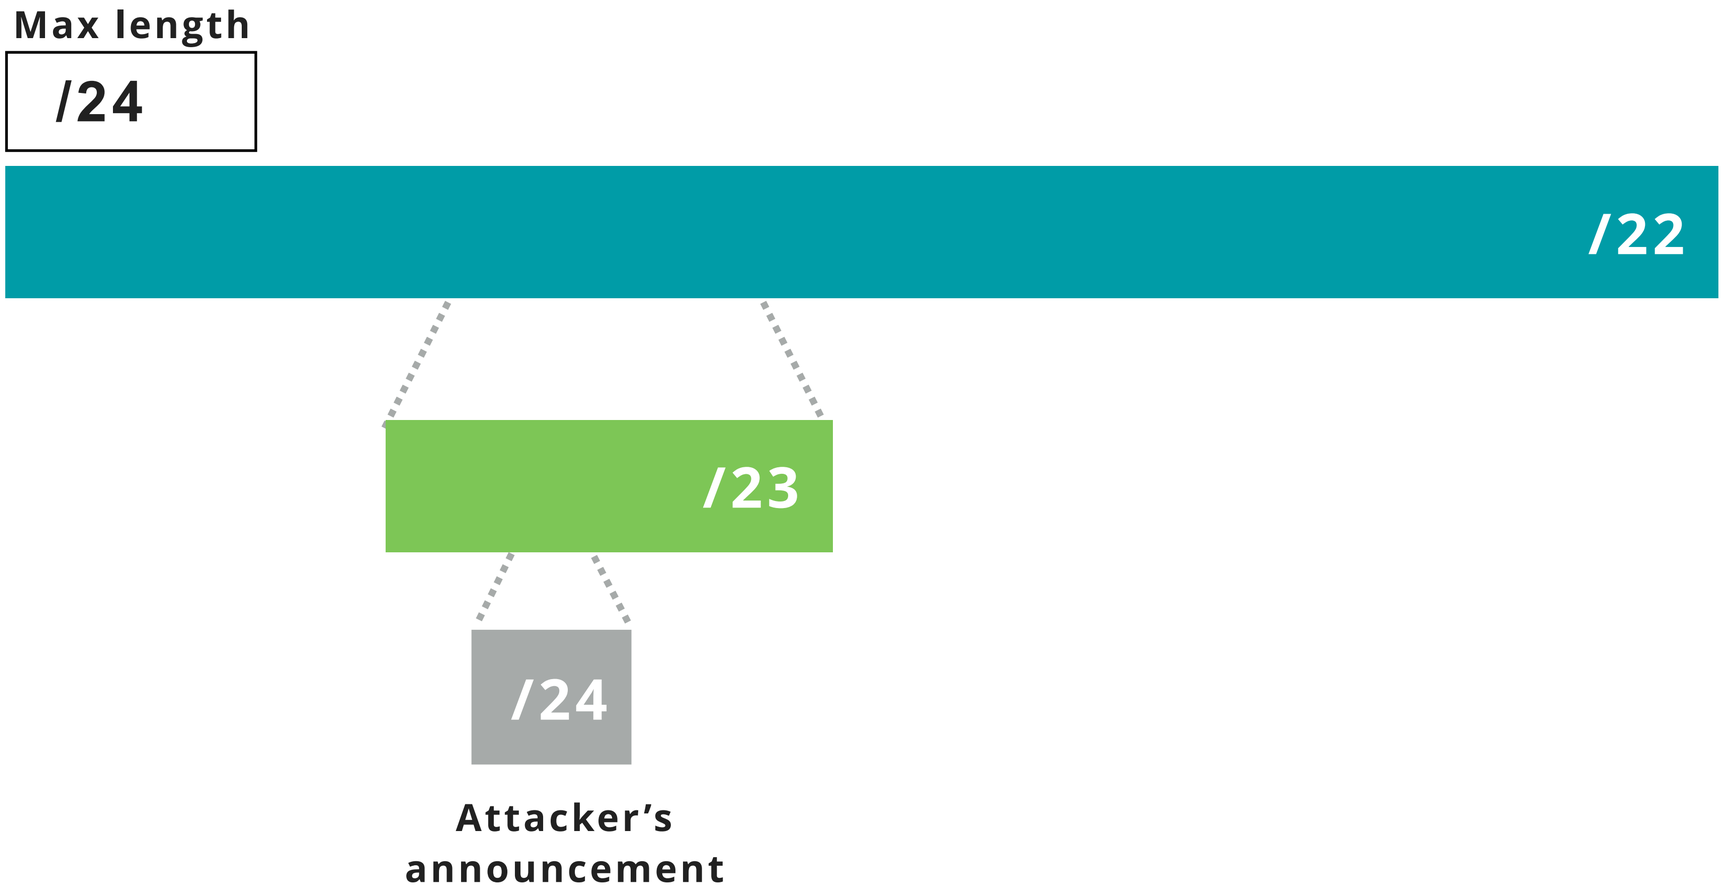

How should we use max-length?

You created a single ROA authorising the entire /22

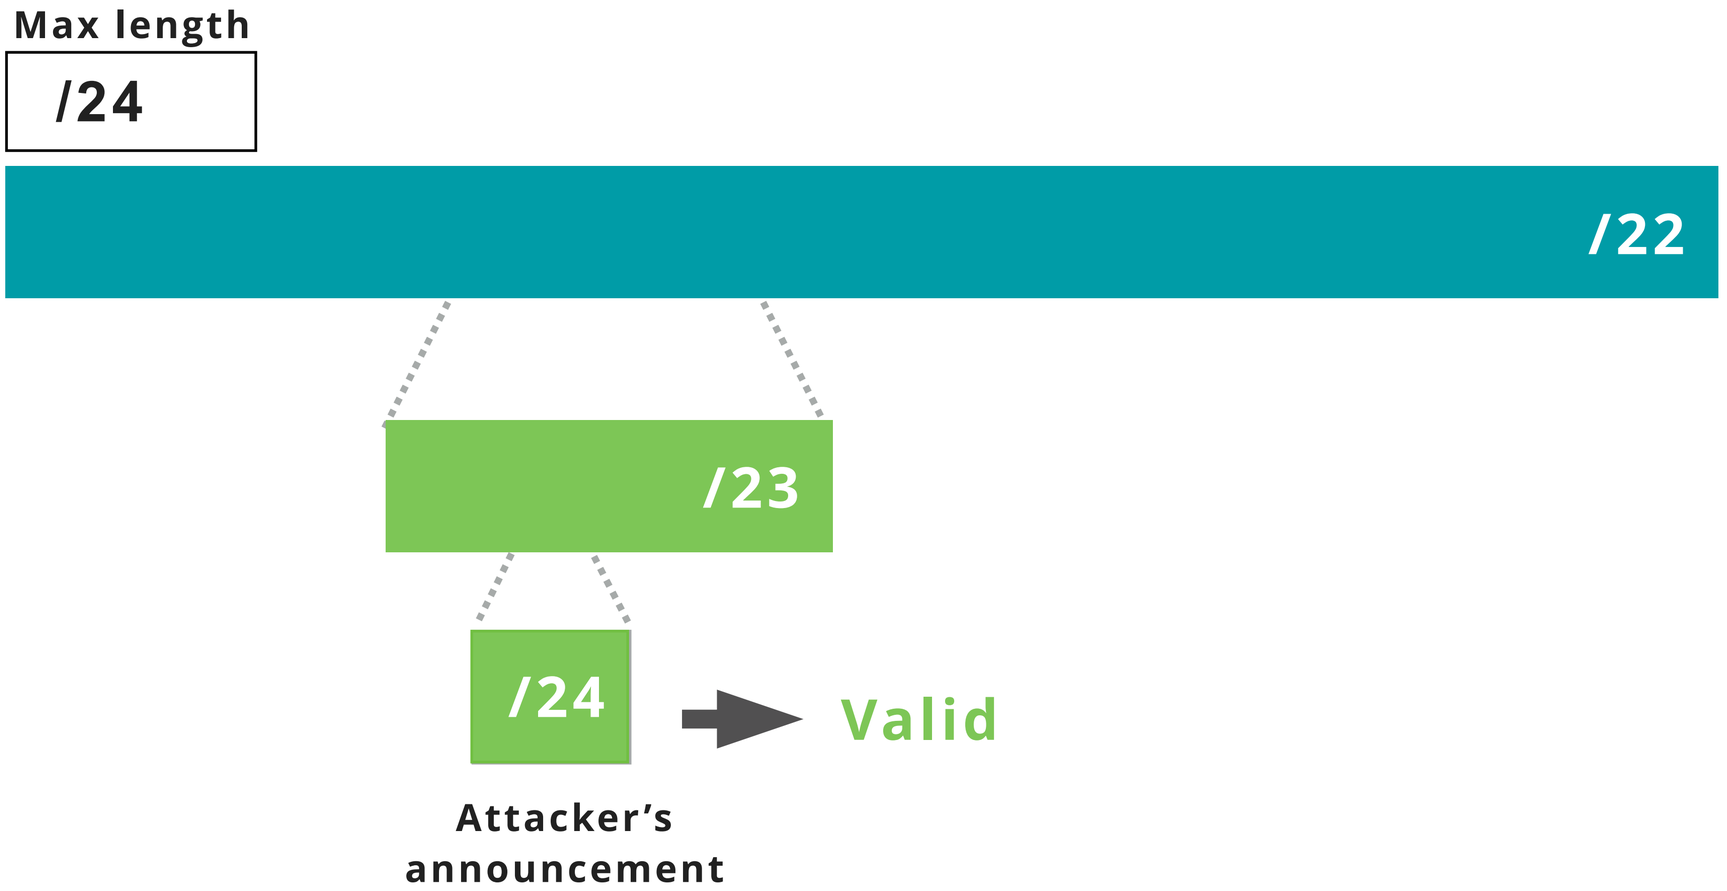

How should we use max-length?

You created ROAs only for your BGP announcements

How should we use max-length?

You created ROAs only for your BGP announcements

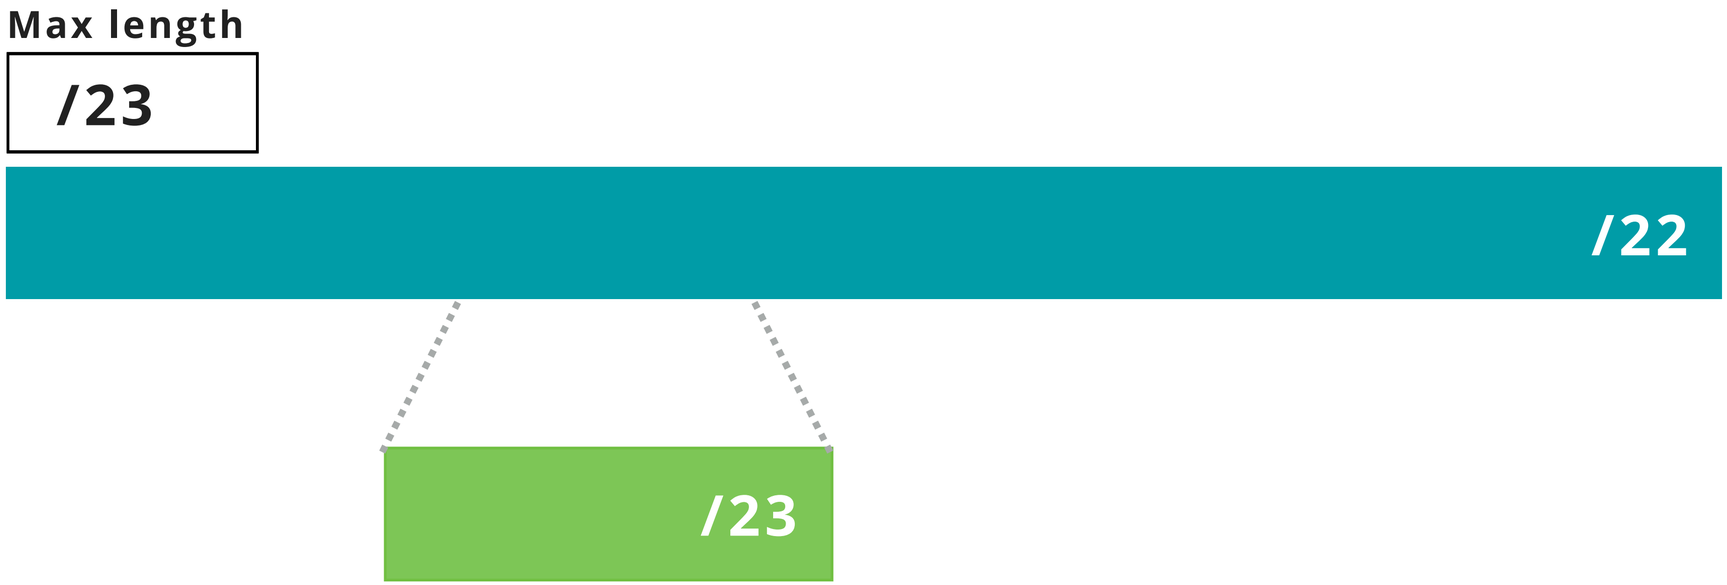

How should we use max-length?

You created ROAs only for your BGP announcements

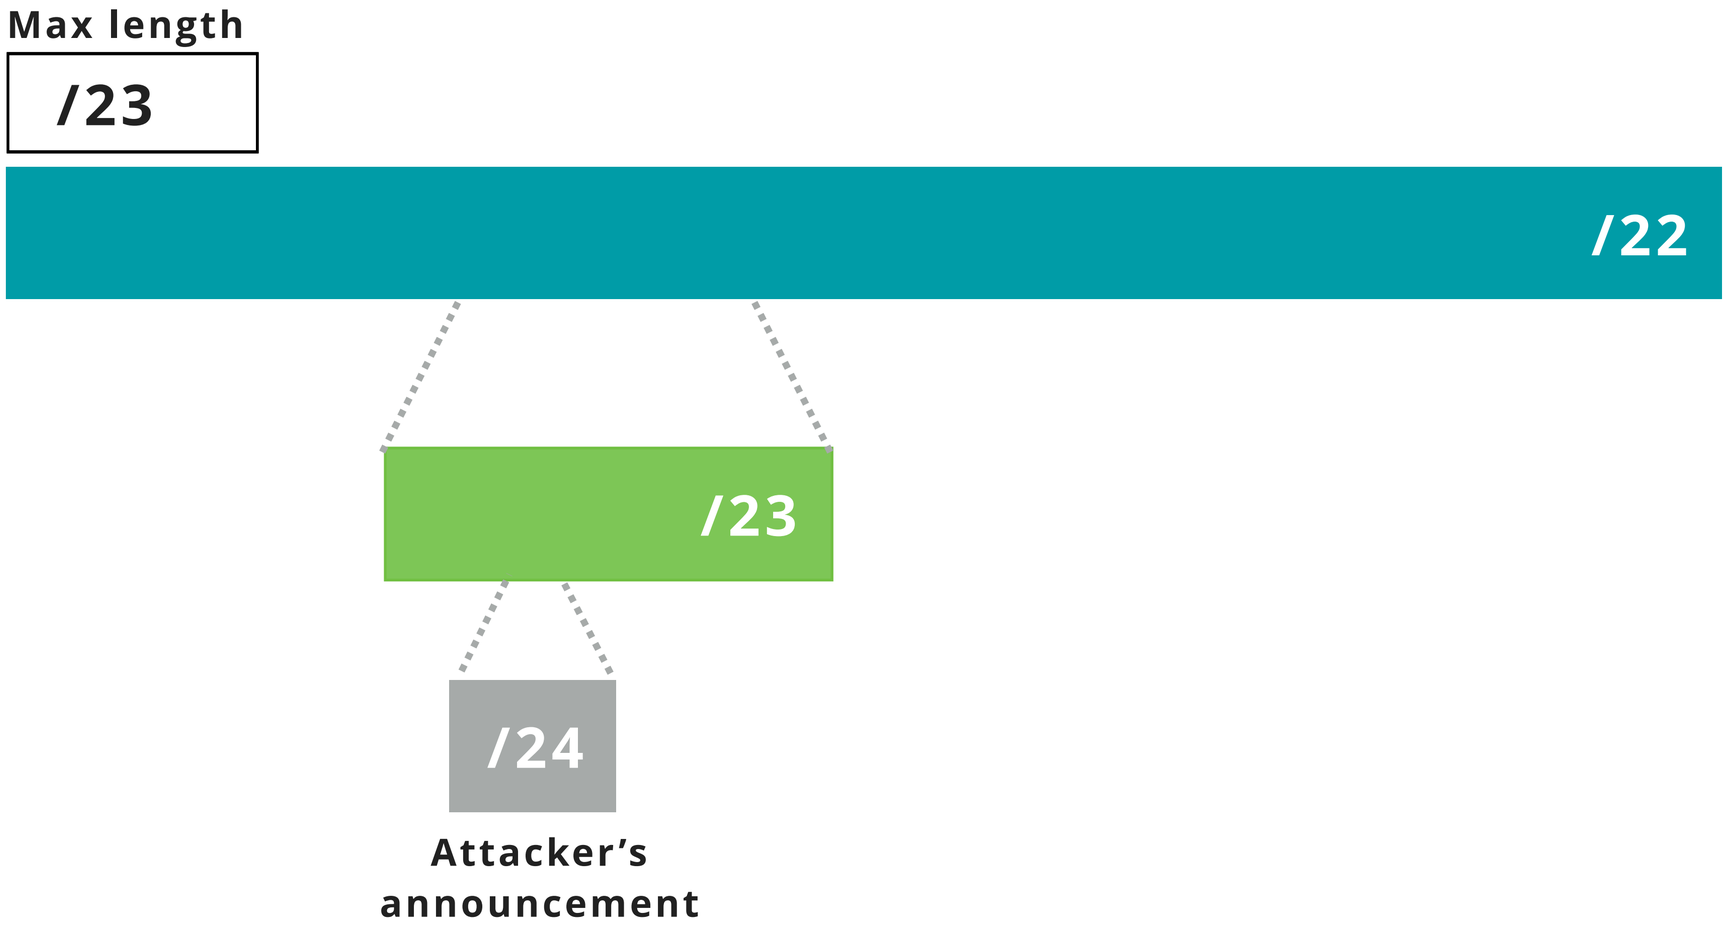

How should we use max-length?

You created ROAs only for your BGP announcements

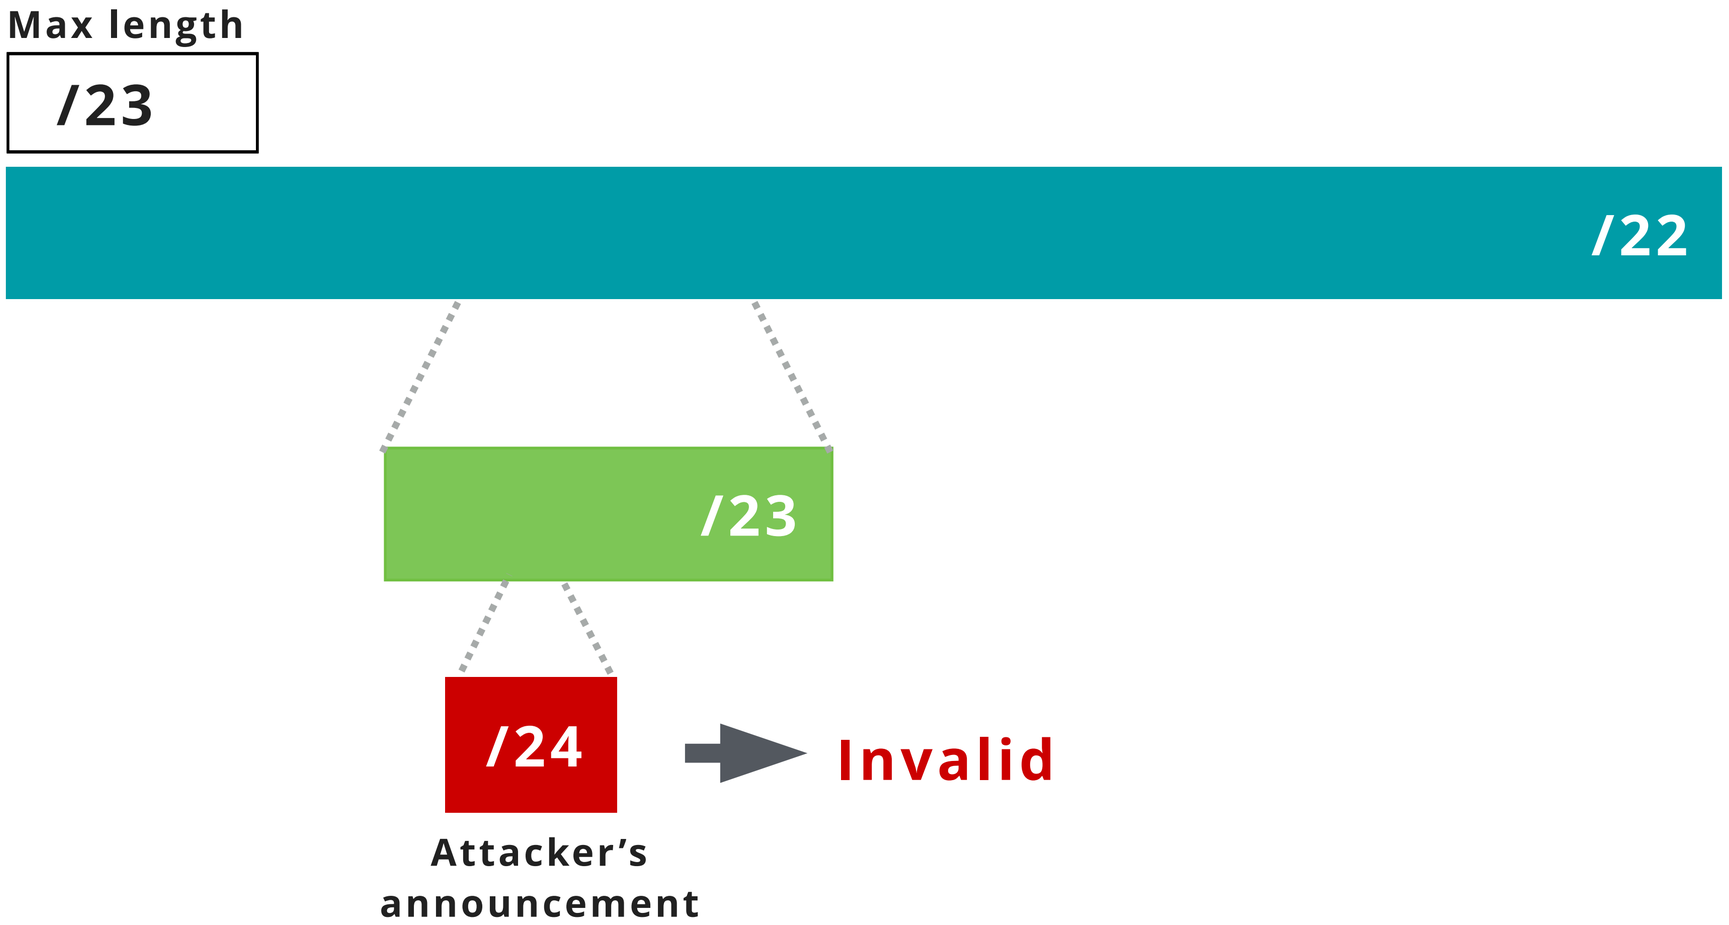

How should we use max-length?

You created ROAs only for your BGP announcements

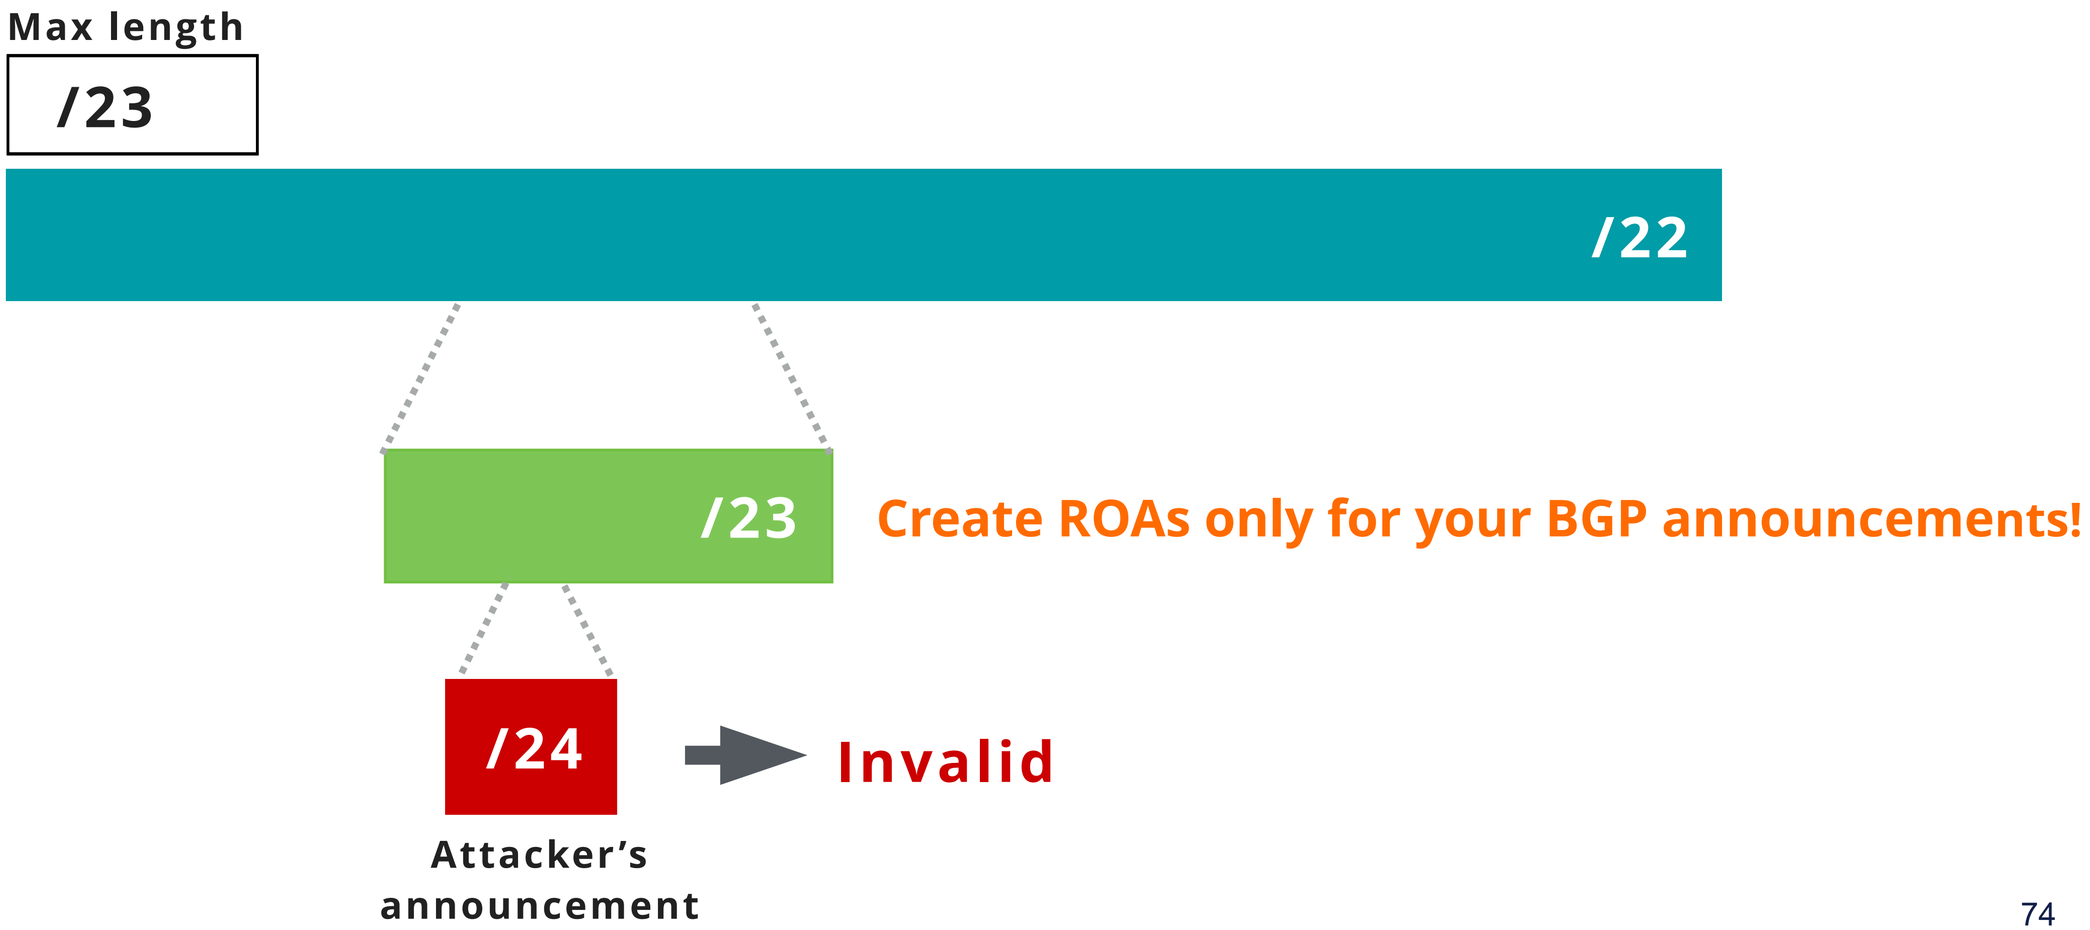

Take the poll!

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

ROA

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

1 min.

RPKI Certificate Structure

- RPKI relies on just five Trust Anchors
- 5 RIRs run a root CA with a Trust Anchor
- RIRs can verify who is creating objects in RPKI system

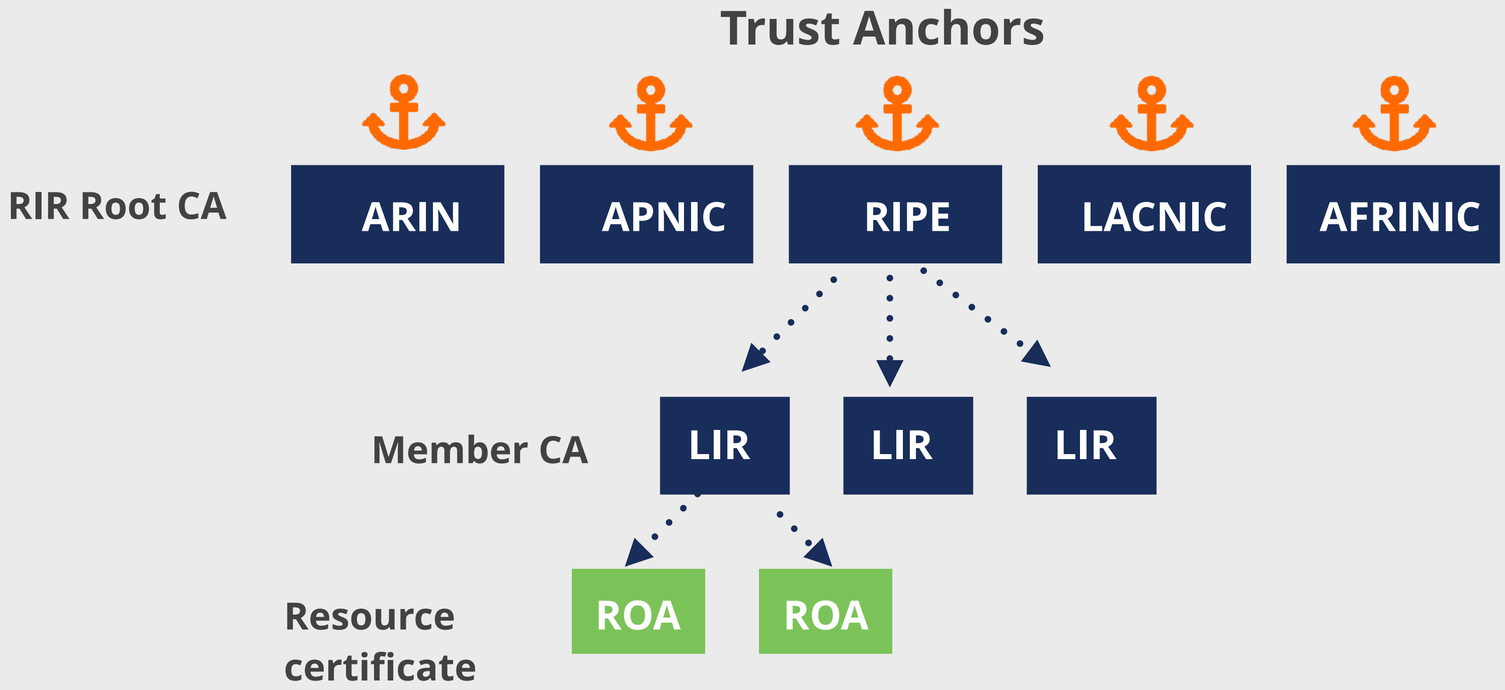

RPKI Chain of Trust

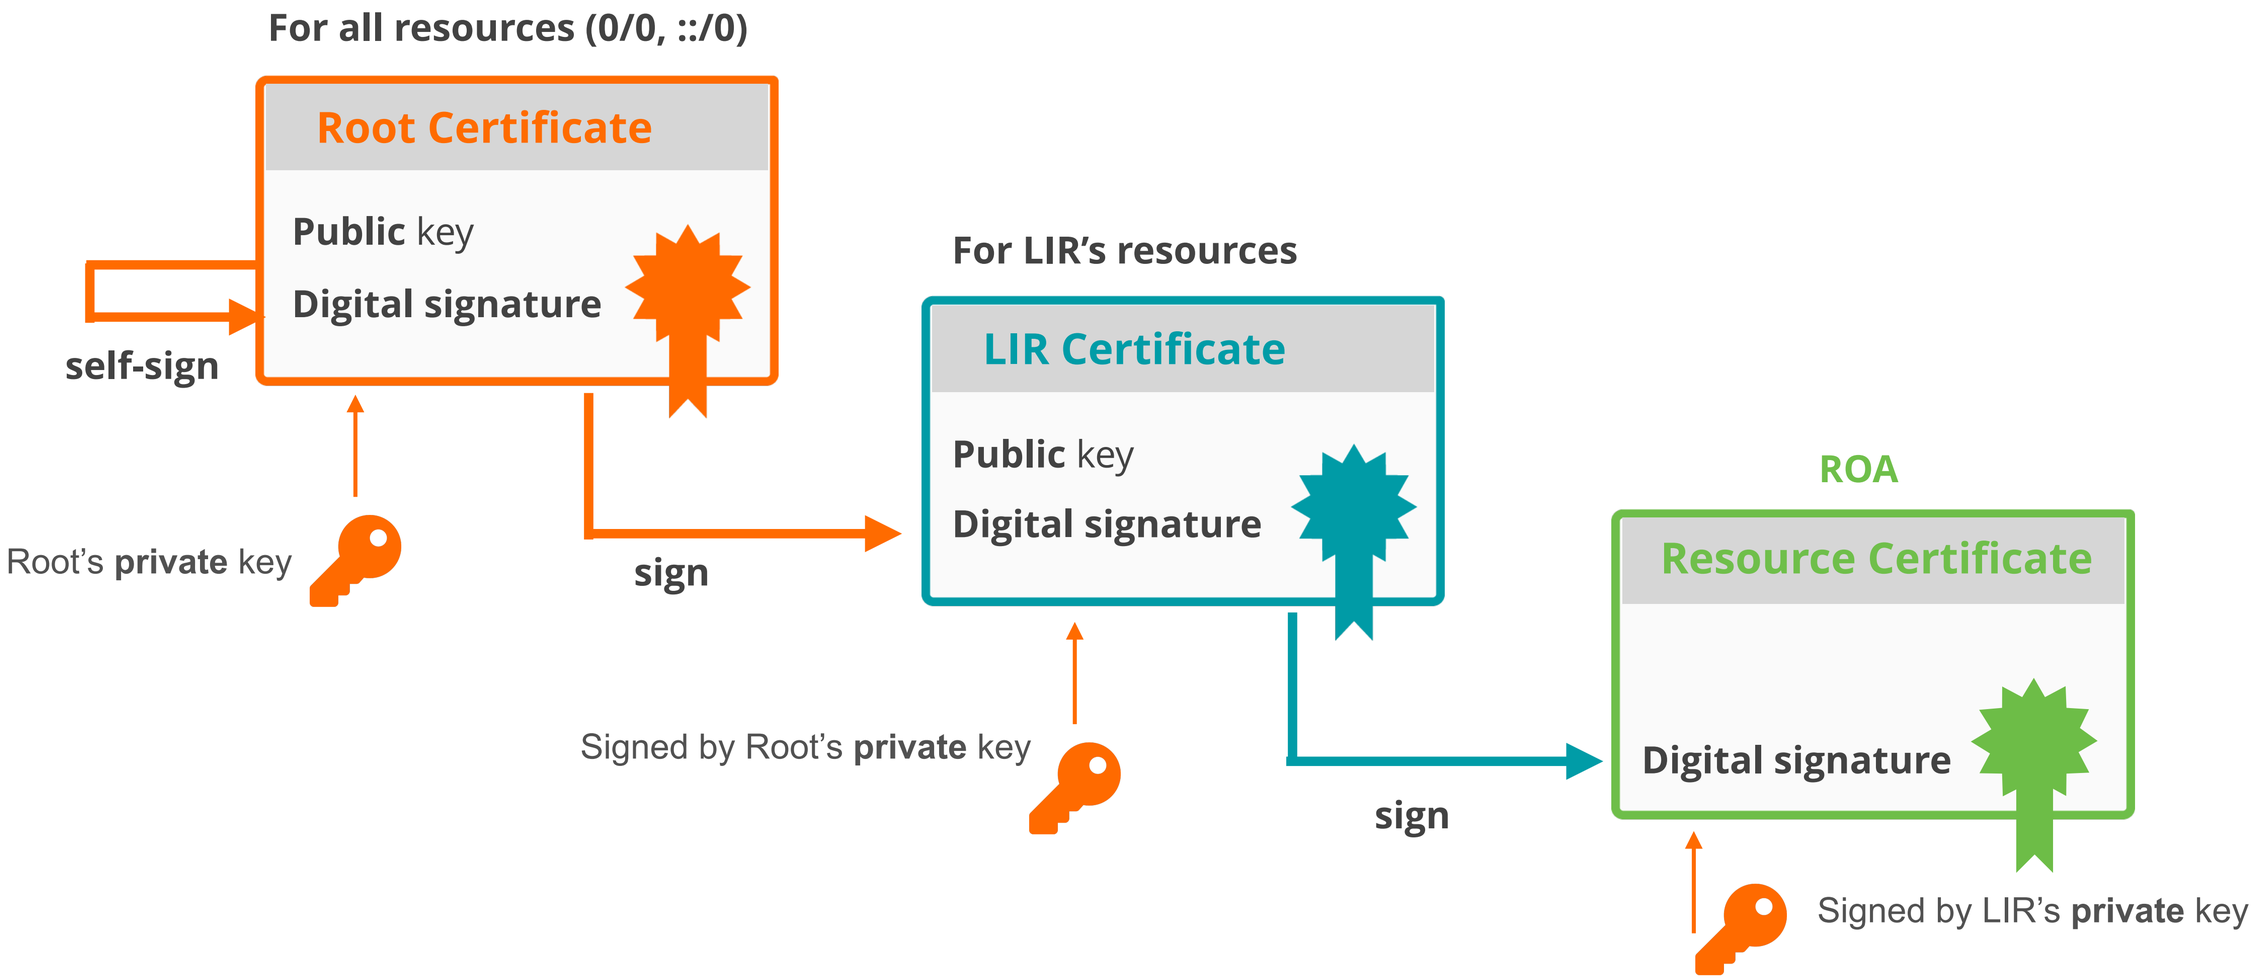

How to implement RPKI?

- First you need to decide which RPKI implementation to use:
  - Hosted RPKI
  - Delegated RPKI

Hosted RPKI

- ROAs are created and published using the **RIR member portal**
- RIR hosts a CA and signs all ROAs
- Automate signing and key rollovers
- Allows you focus on creating and publishing ROAs

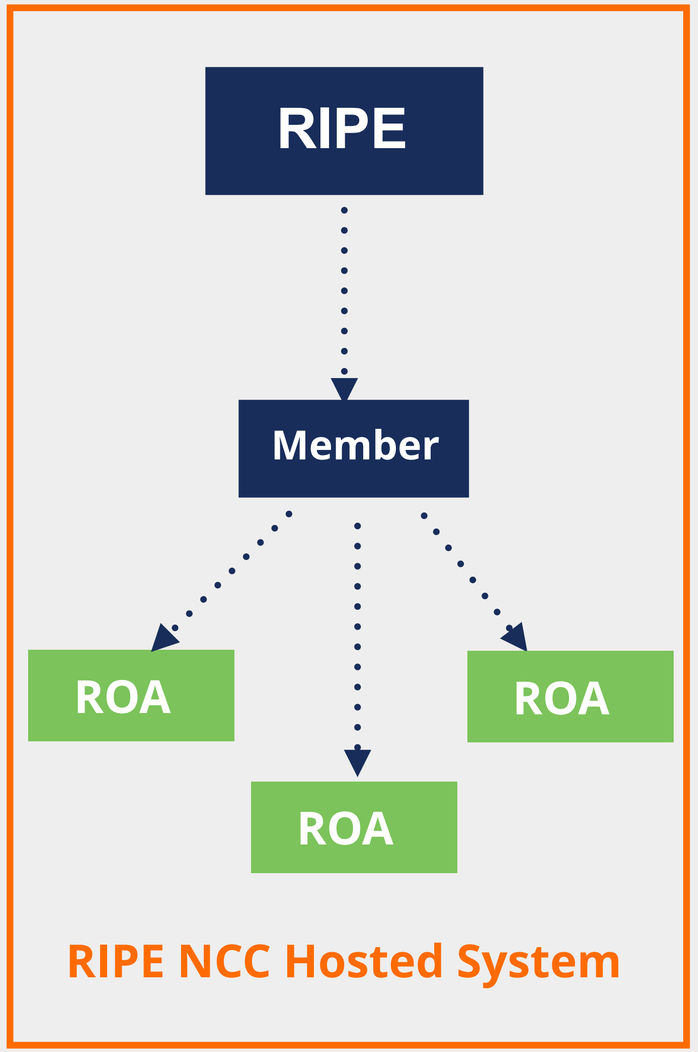

Delegated RPKI

- Run your own Certificate Authority software
  - Dragon Research Labs, RPKI Toolkit
  - NLnet Labs, Krill
- Create ROAs in your own platform
- Manage your own keys/key rollovers
- Setup connection with RIPE NCC CA
- Generate LIR certificate
- Get it signed by parent CA

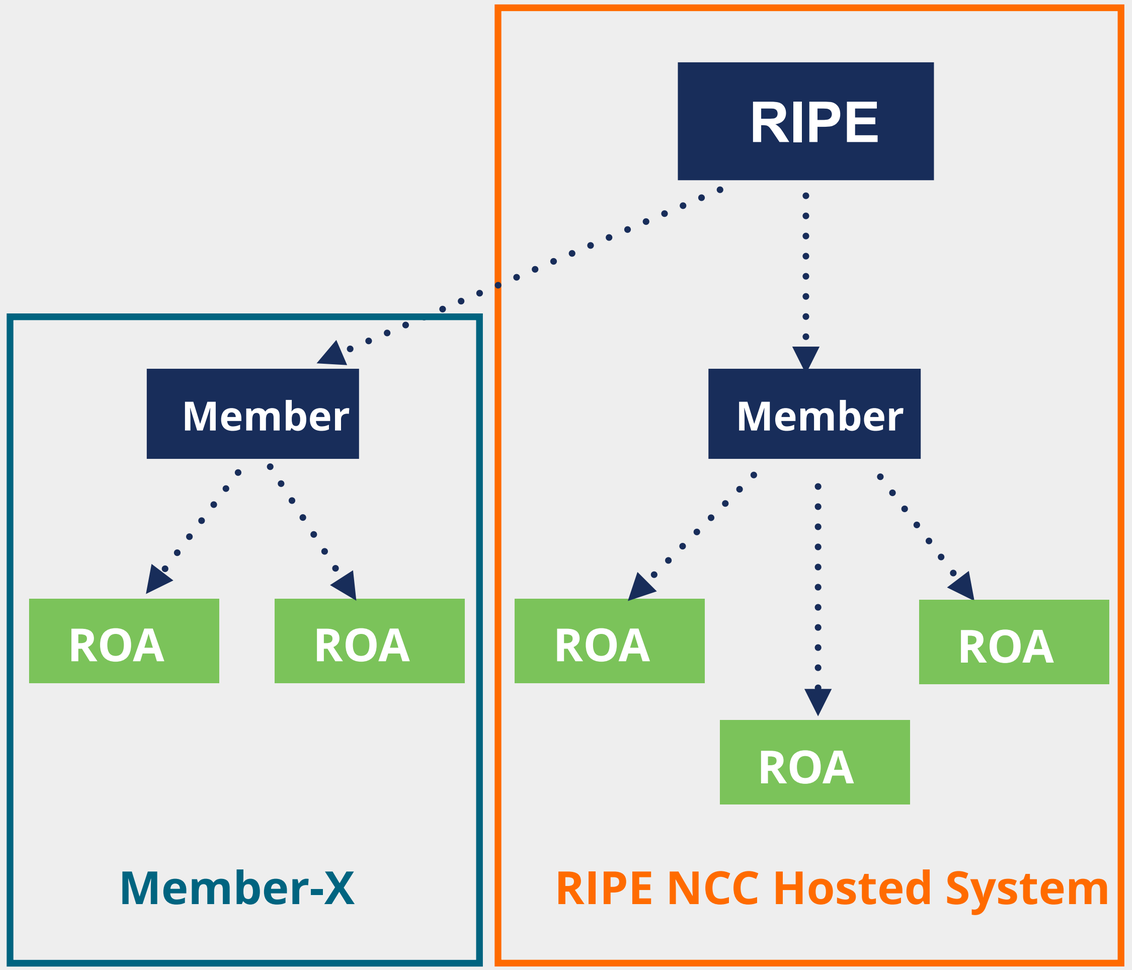

RIPE NCC RPKI Dashboard

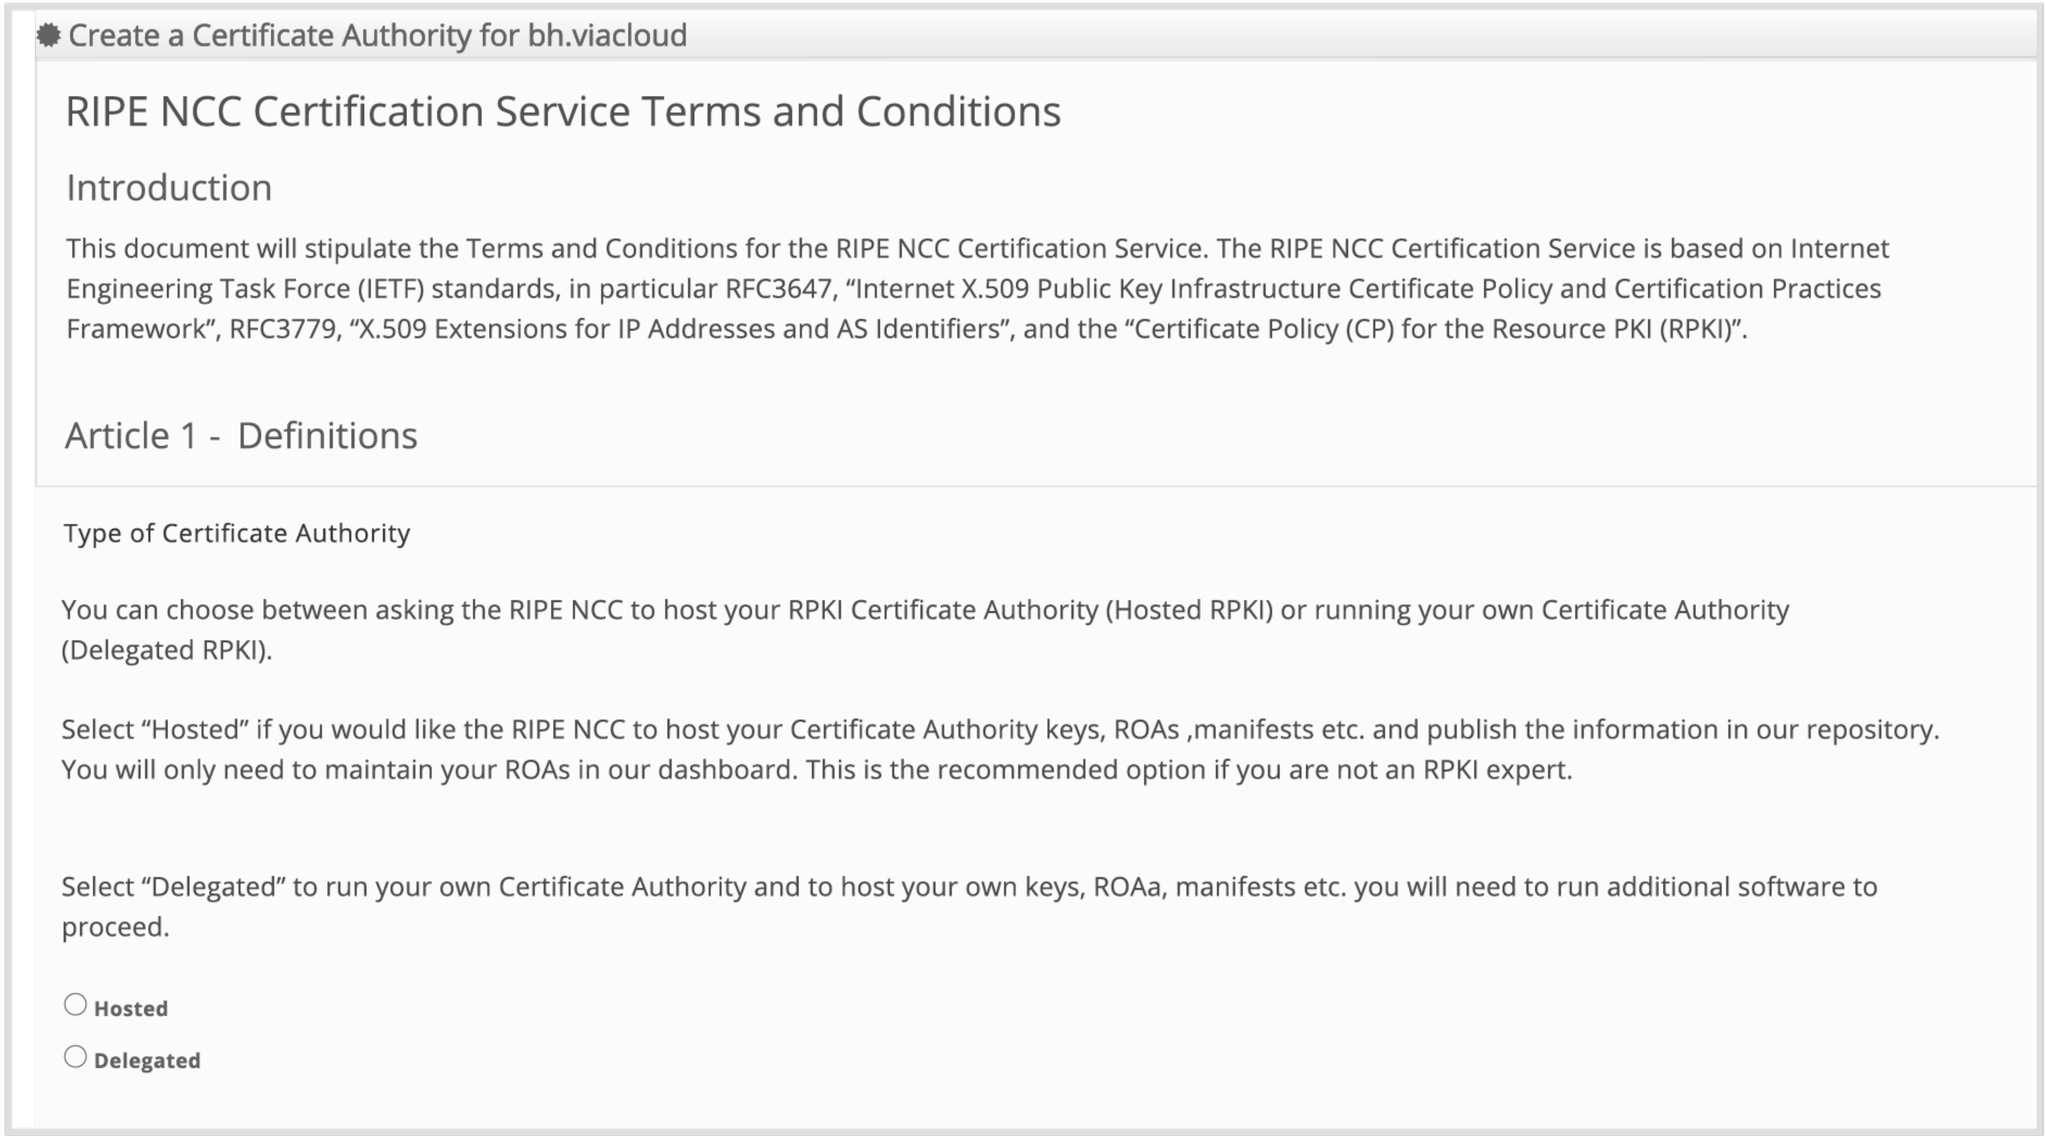

RIPE NCC Hosted Solution

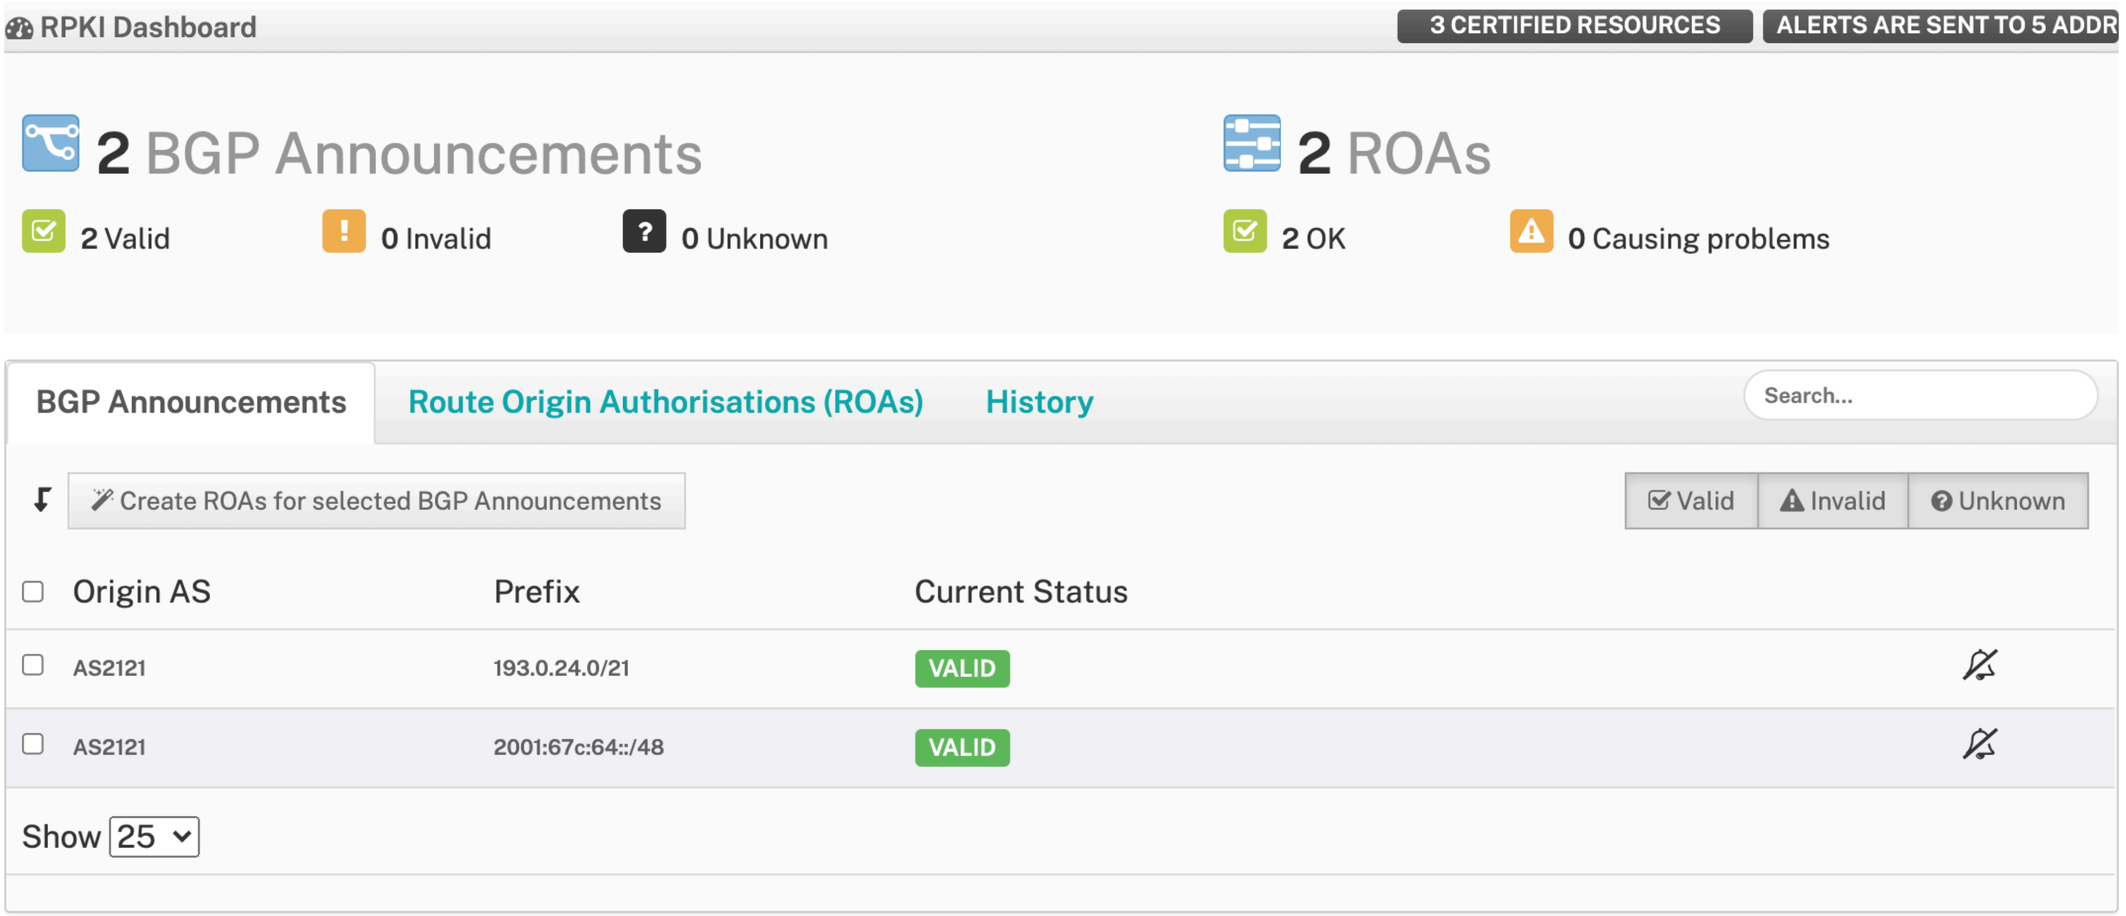

Looking for ROA Certification for PI resources?

Revoke hosted CA

RIPE NCC Hosted Solution

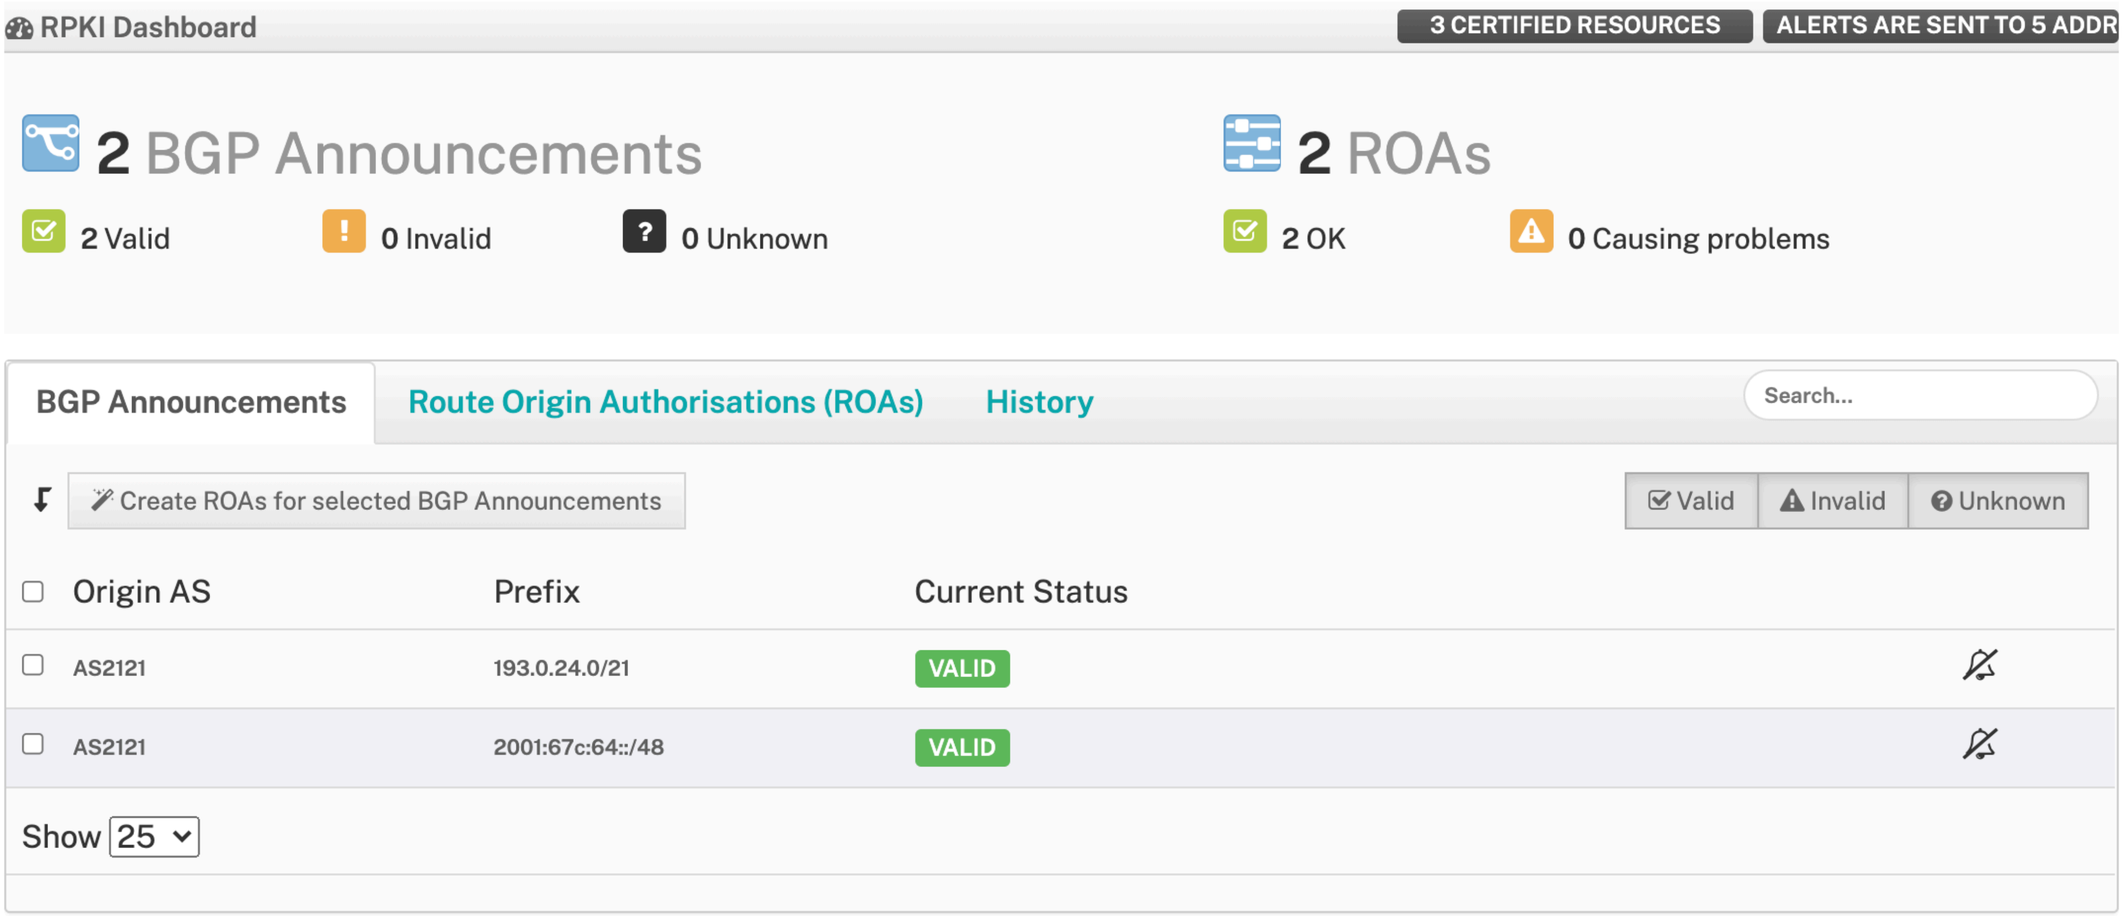

Looking for ROA Certification for PI resources?

Revoke hosted CA

Certifying PI Resources

Requested and managed by PI End User or by Sponsoring LIR

1. Complete the wizard successfully

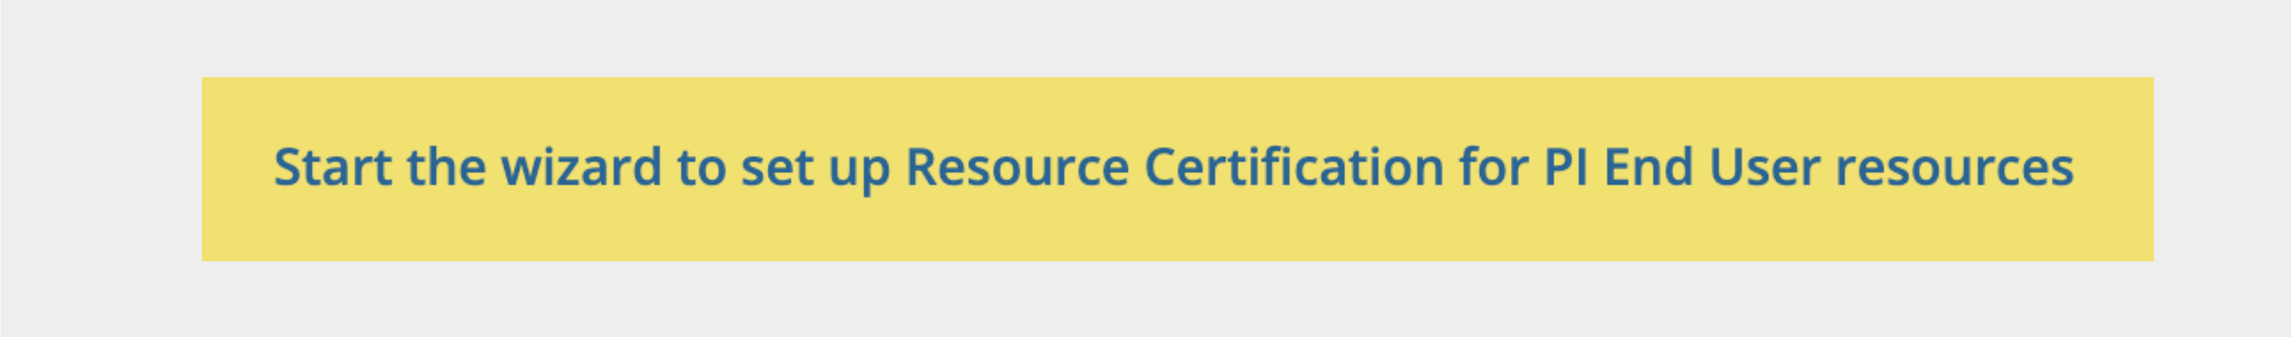

2. Login to https://my.ripe.net and request a certificate
   - Sign in with your RIPE NCC Access account
3. Manage your ROAs

How does RPKI work?

Validation

RPKI has two elements

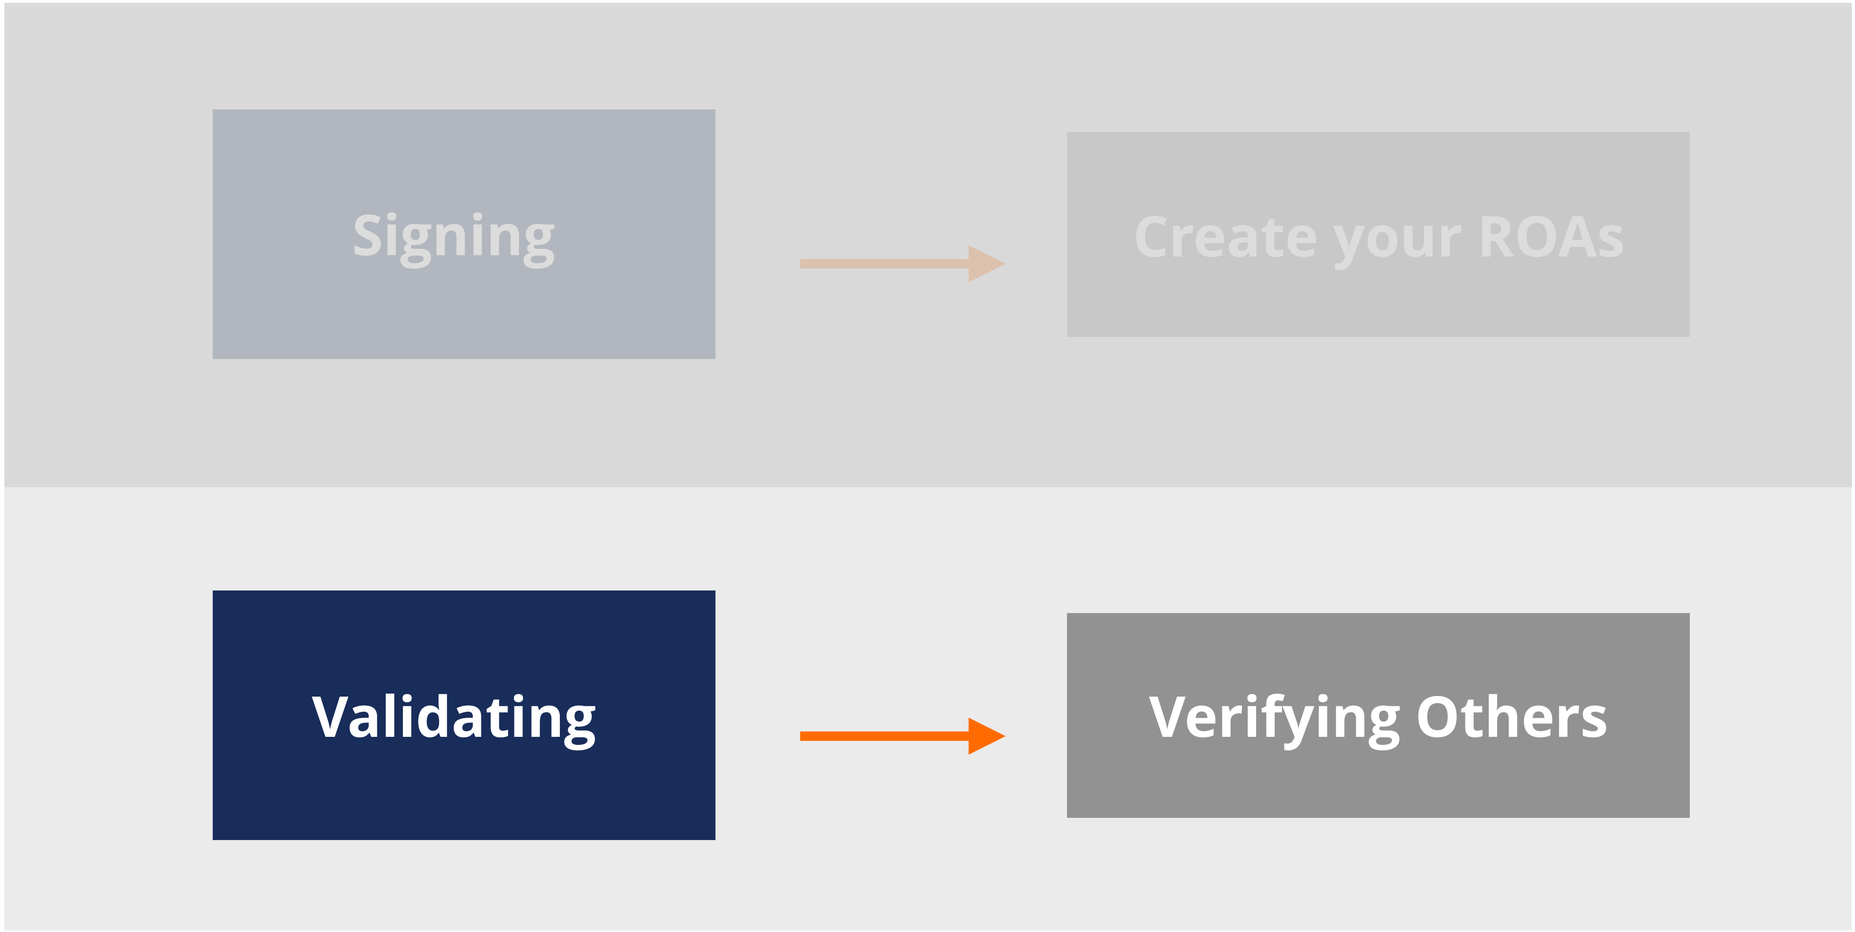

Validation

- You are verifying the info provided by others
- Key part of public key infrastructure
- Sometimes people call it BGP OV
- Yes, it’s the goal but there is more…
- There are two types of validation:
  - ROA Validation
  - BGP Origin Validation (BGP OV)

ROA Validation

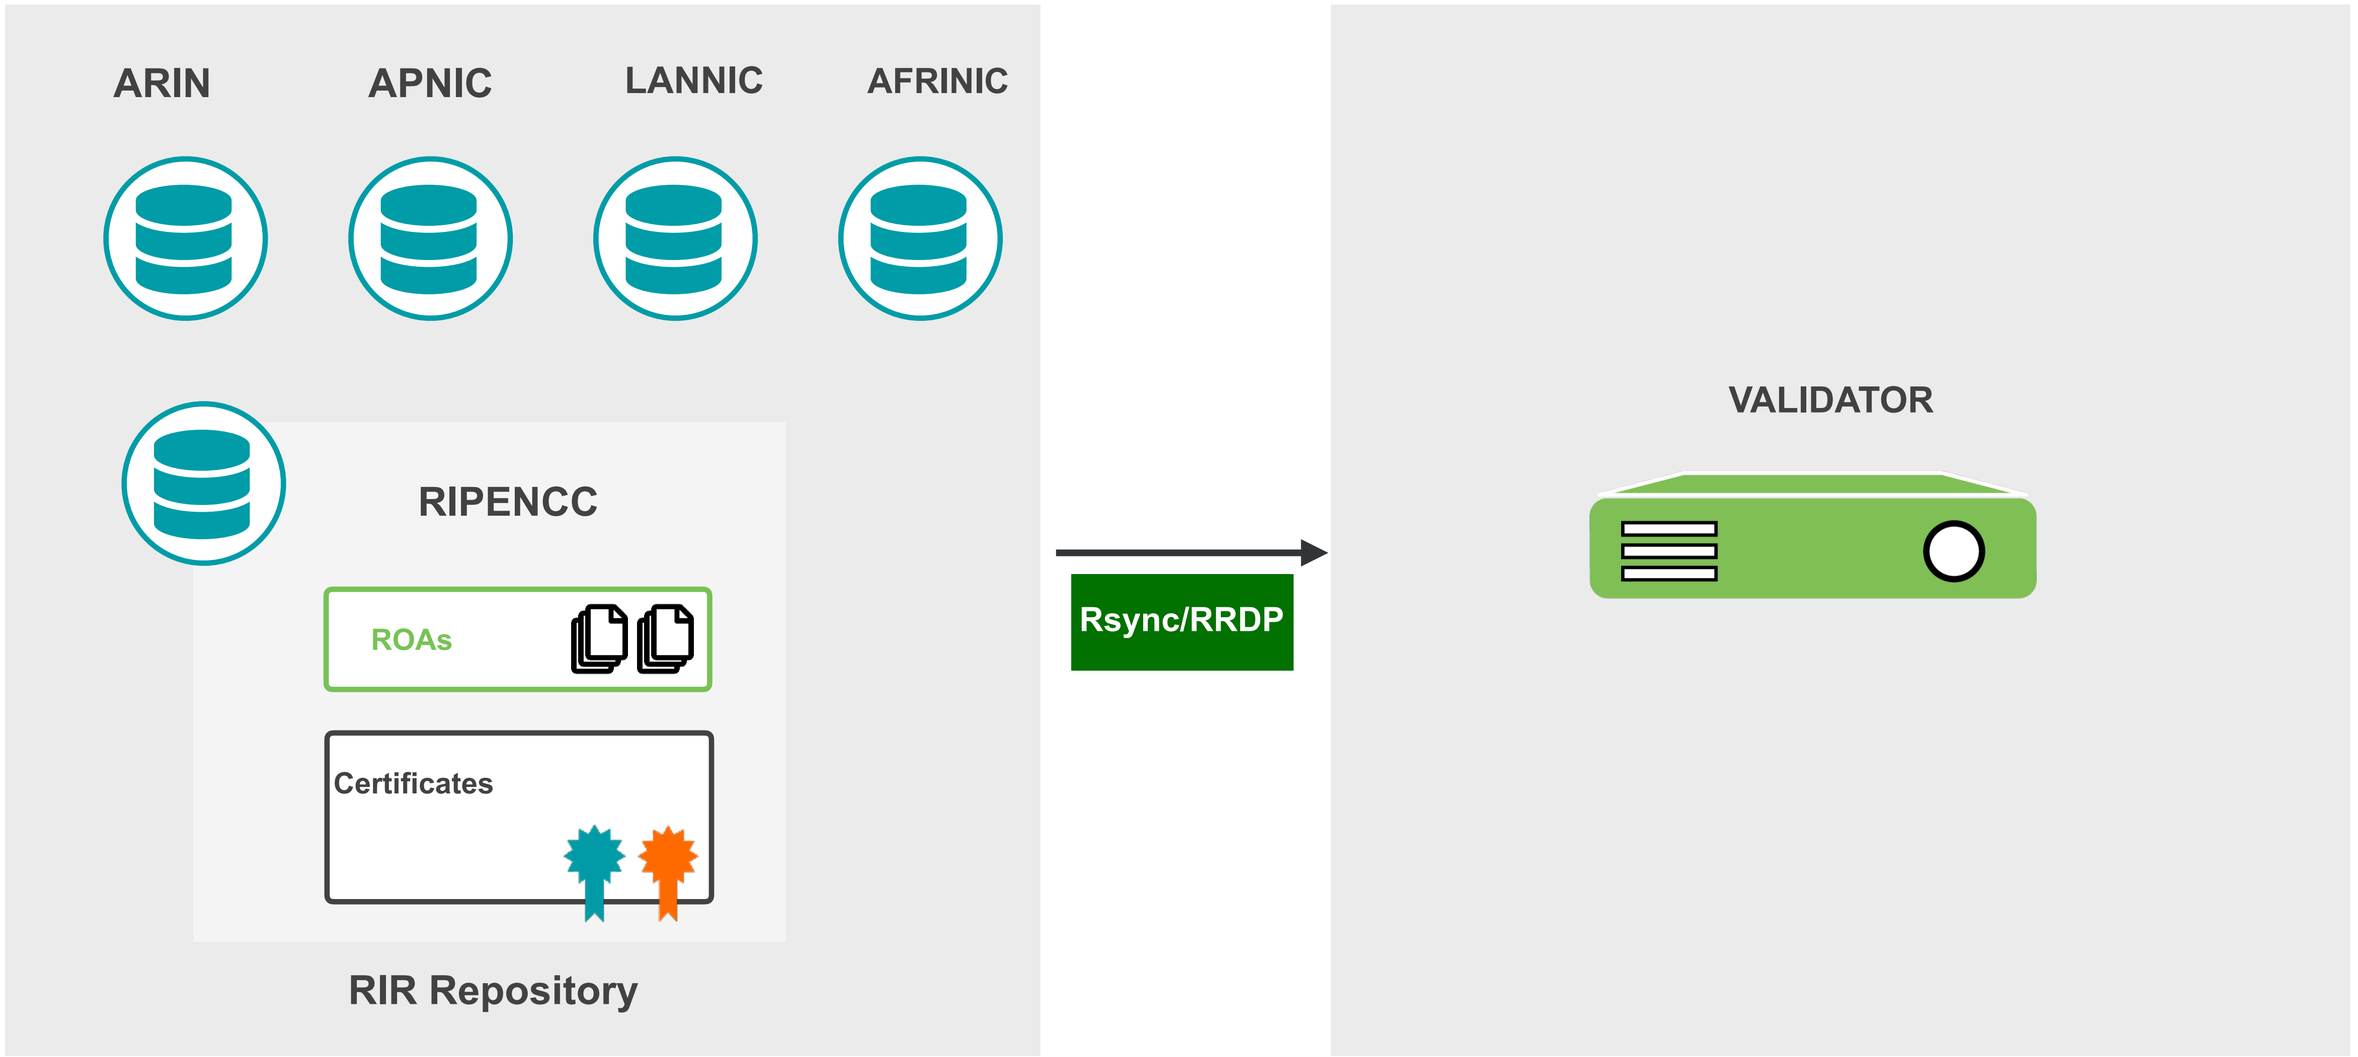

ROA Validation

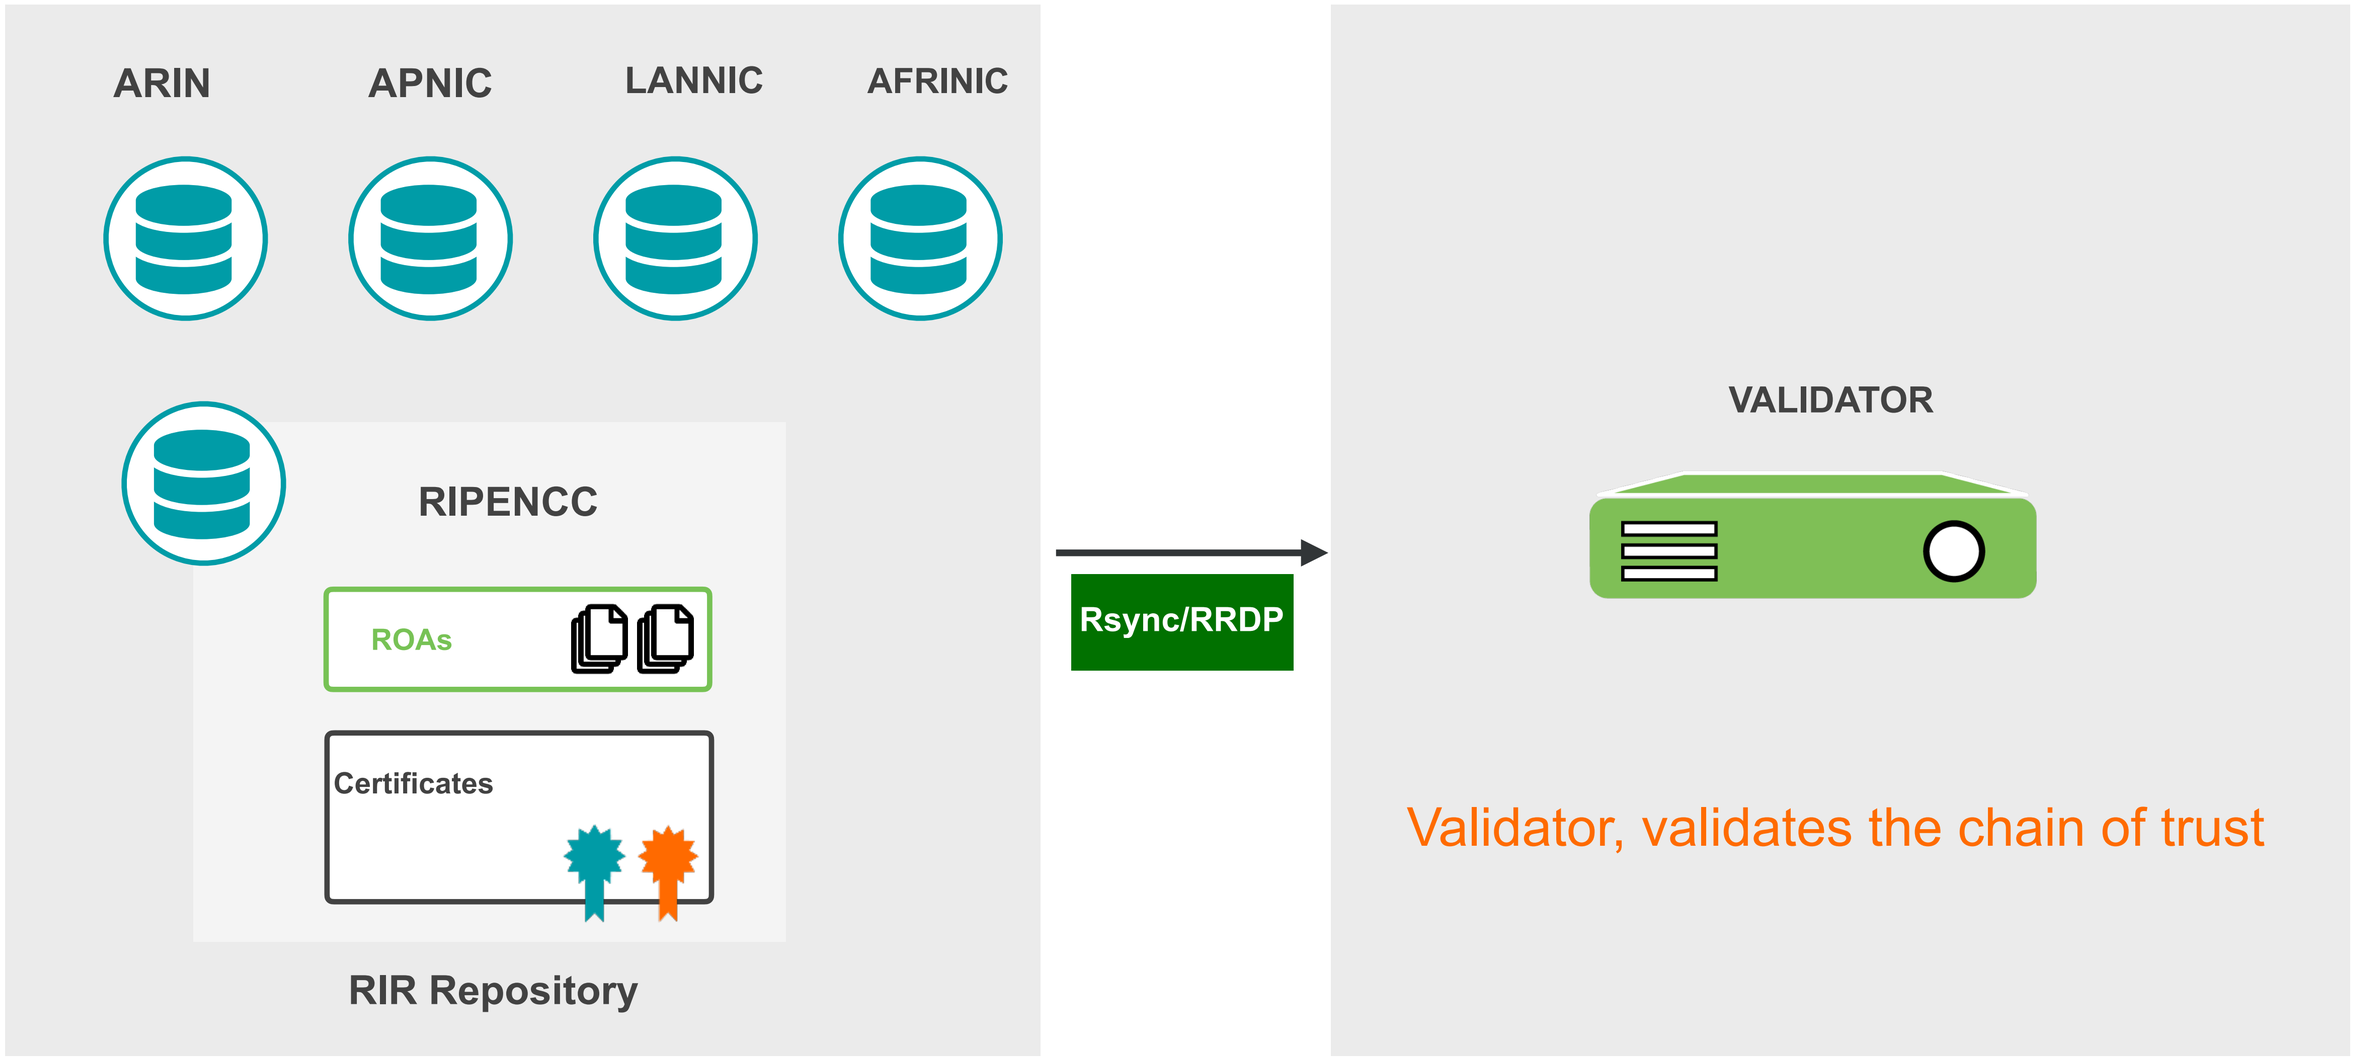

RPKI Validators

- Validator is a software
- Downloads the RPKI repository from the RIRs
- Validates the chain of trust of all ROAs and associated CAs
- Several validator options are available…

RPKI Validator Options

- **RIPE NCC Validator 3.2 (DEPRECATED)**
  - Java based
- **Routinator**
  - Built with Rust, built by NLNetlabs
- **OctoRPKI**
  - Cloudflare’s Relying Party software, written in Go
- **FORT**
  - Open source RPKI validator, Written in C

January 1, 2021 no new features!

July 1, 2021 end of support!

Links for Validators

RPKI Validators:

https://github.com/RIPE-NCC/rpki-validator

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

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

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

For more info…

https://rpki.readthedocs.io

ROA Validation Process

ROA Validation Process

ROA Validation Process

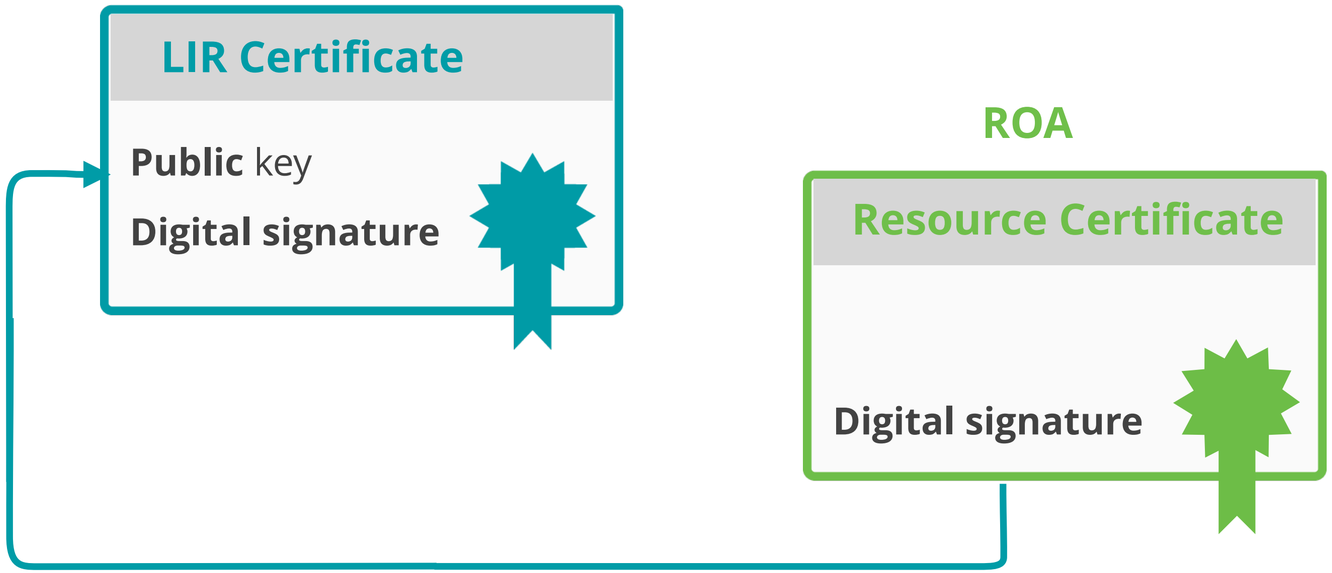

ROA Validation Process

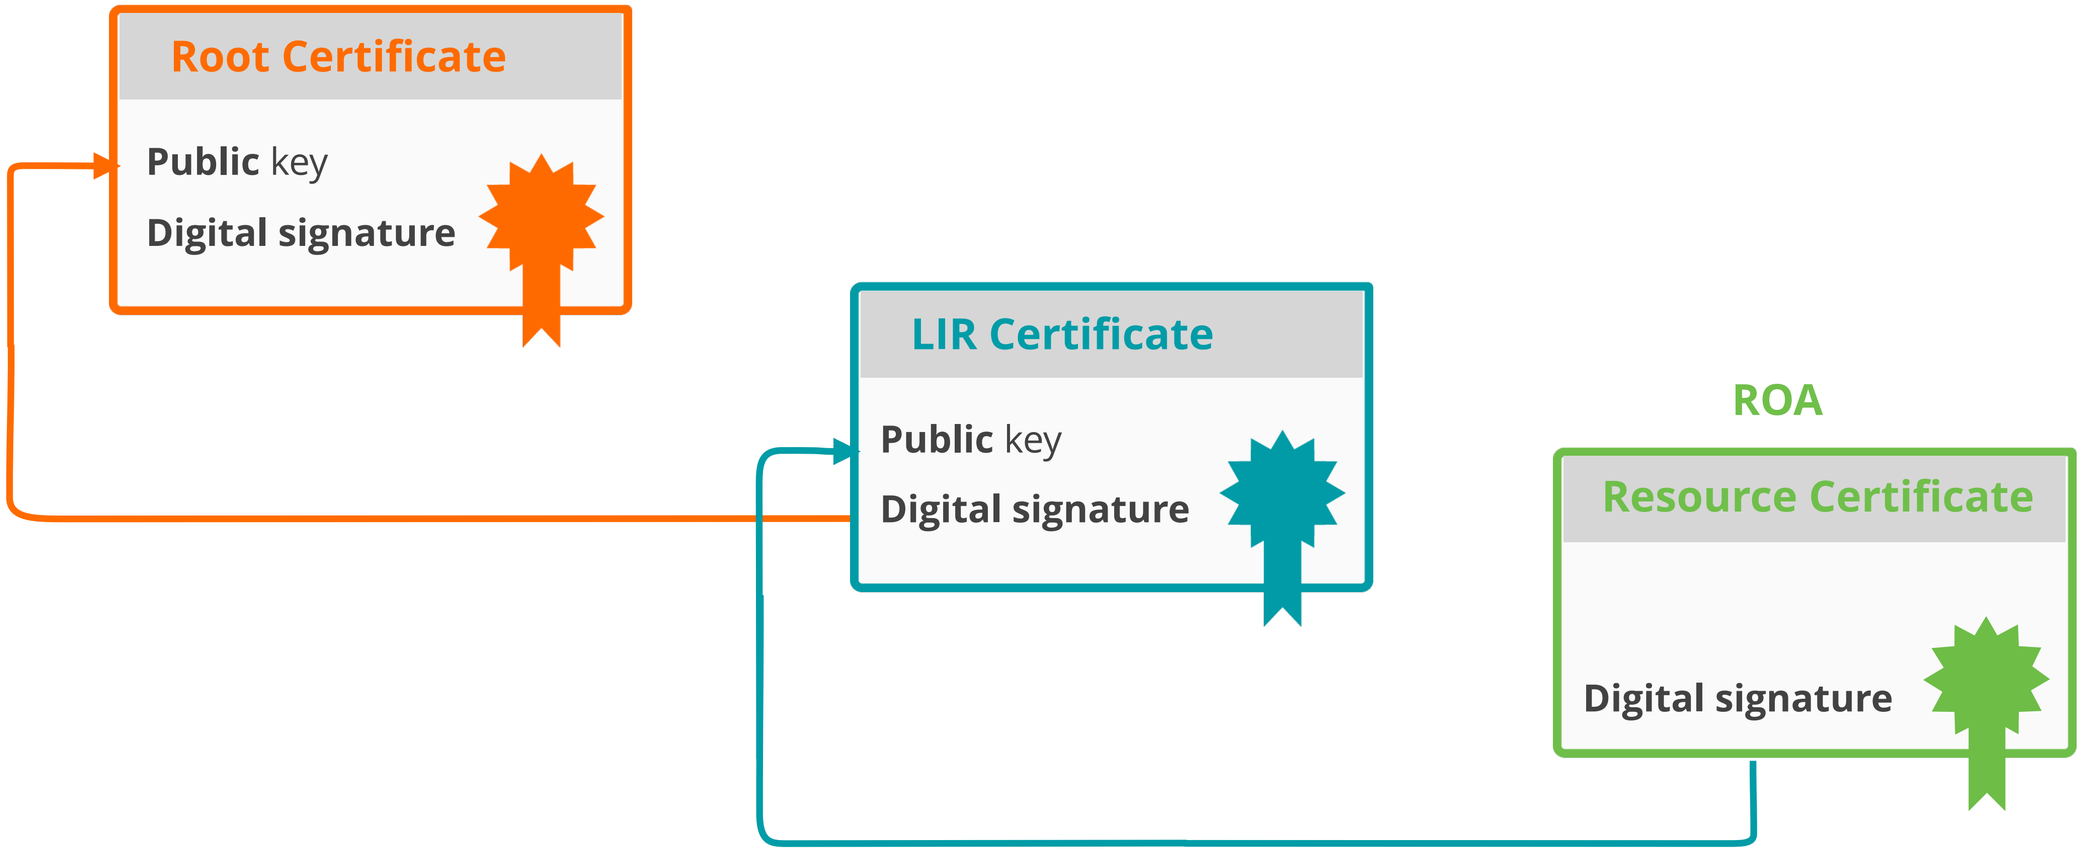

ROA Validation Process

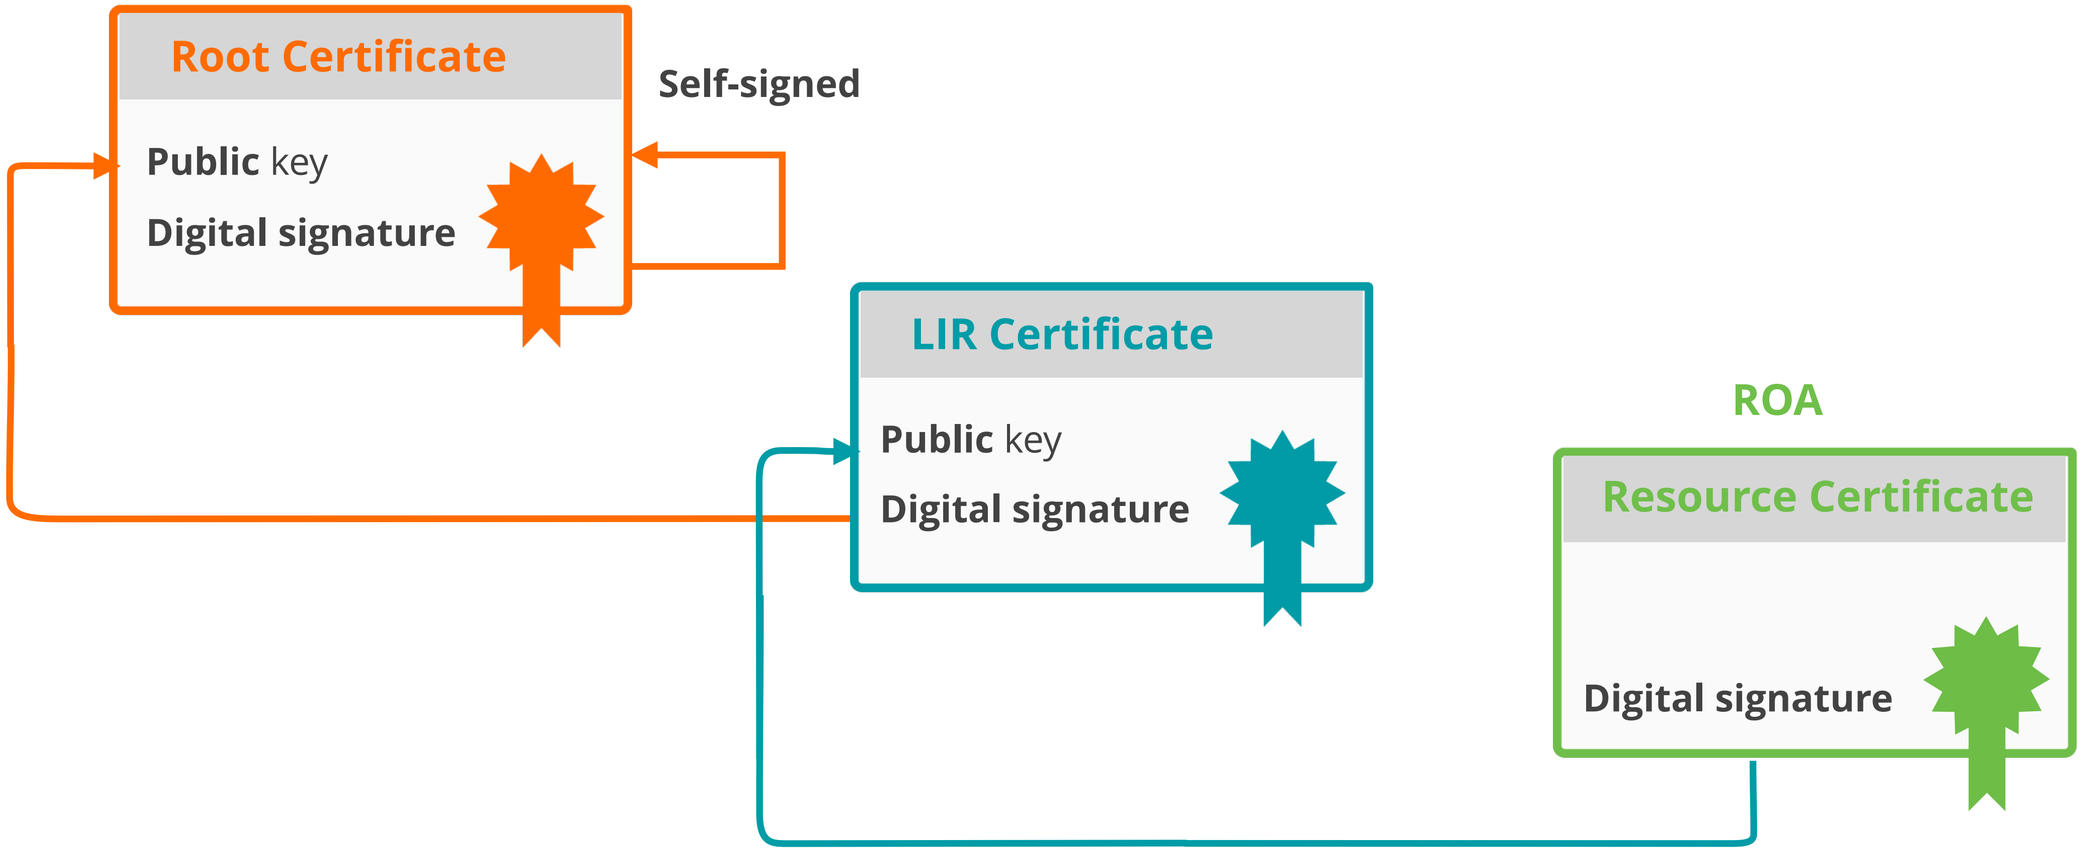

ROA Validation Process

IF chain is complete, it means **ROA is VALID!**

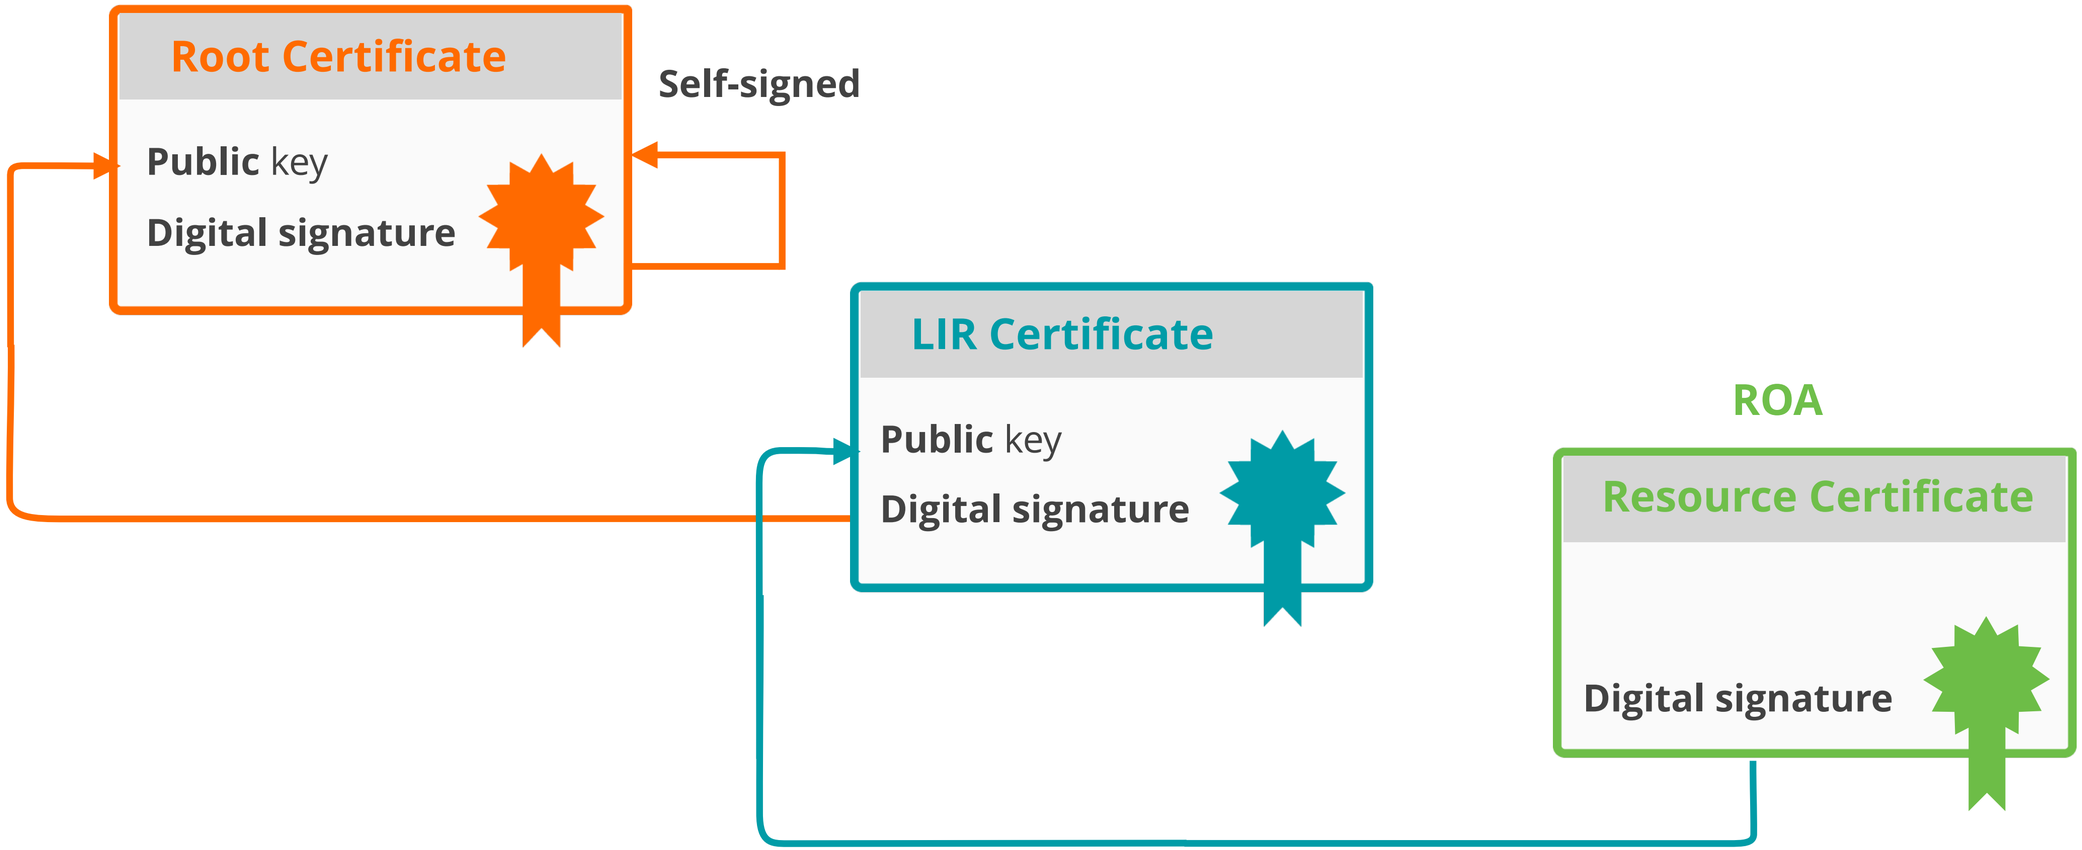

ROA Validation Process

IF chain is complete, it means **ROA is VALID!**

ELSE validation is unsuccessful, **ROA is INVALID!**

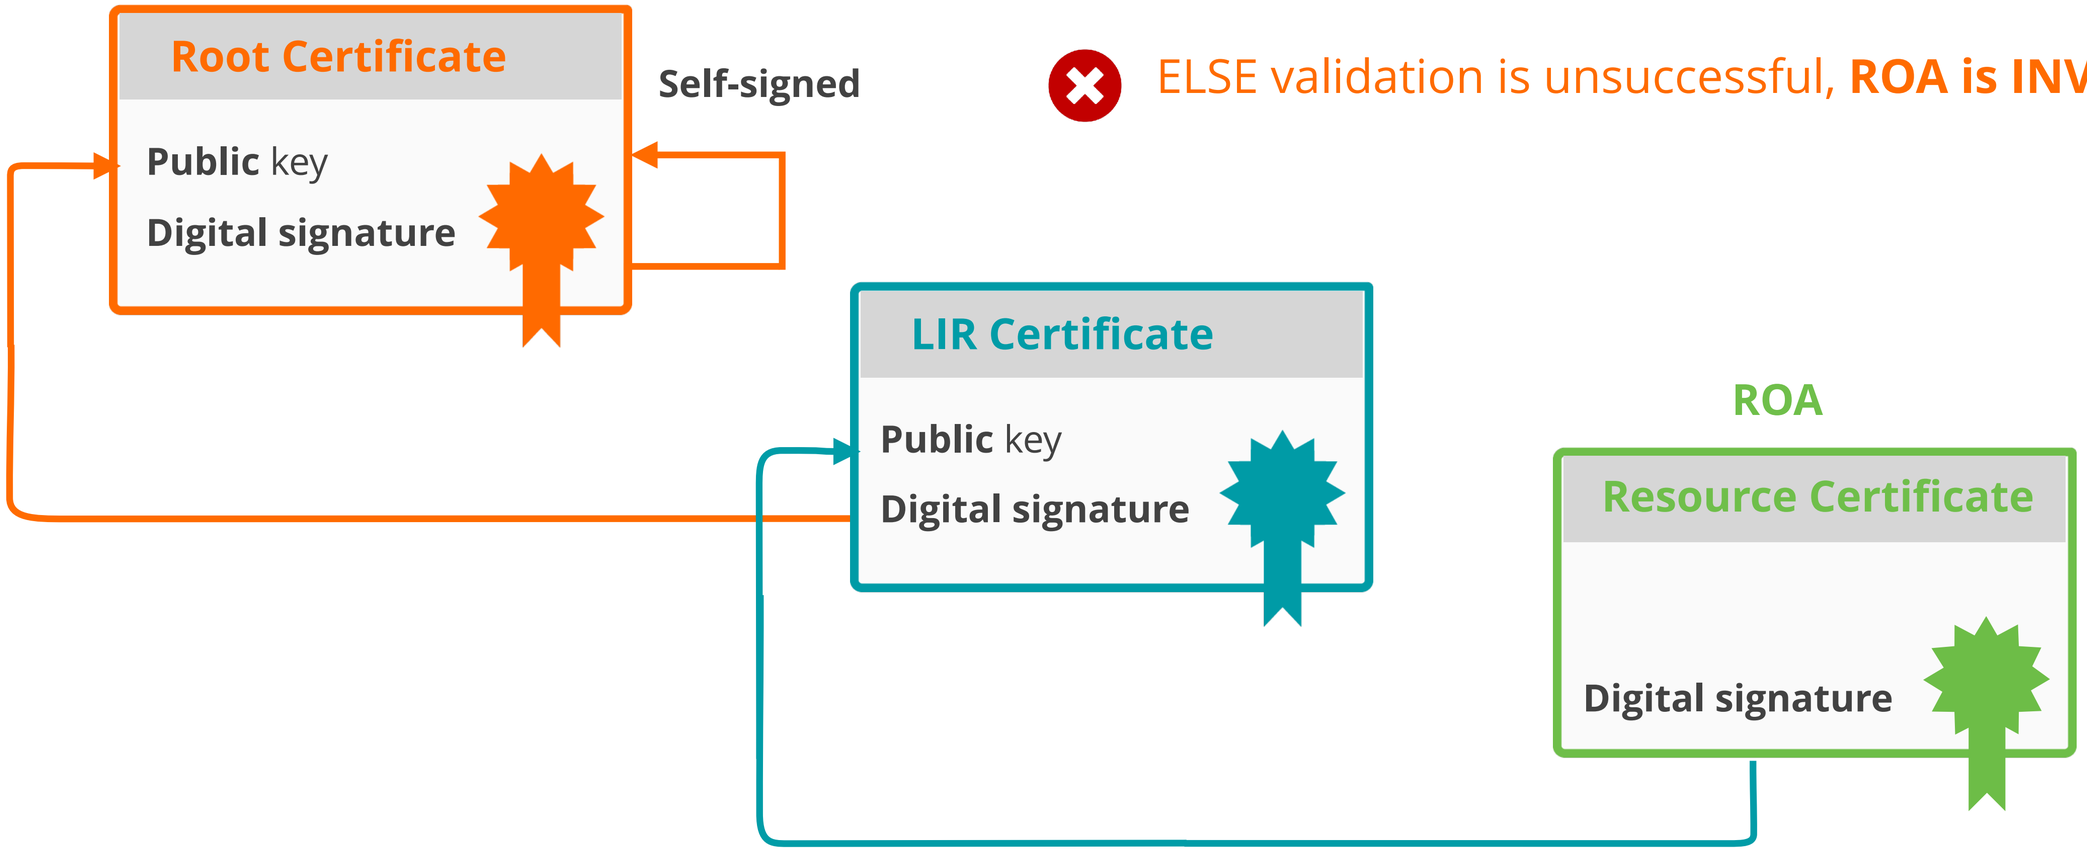

Take the poll!

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

Please choose all the options that apply.

BGP Origin Validation (BGP OV)

- RFC#6811
- BGP Filtering with ROA
- Validating BGP announcements by using RPKI infrastructure

BGP Origin Validation

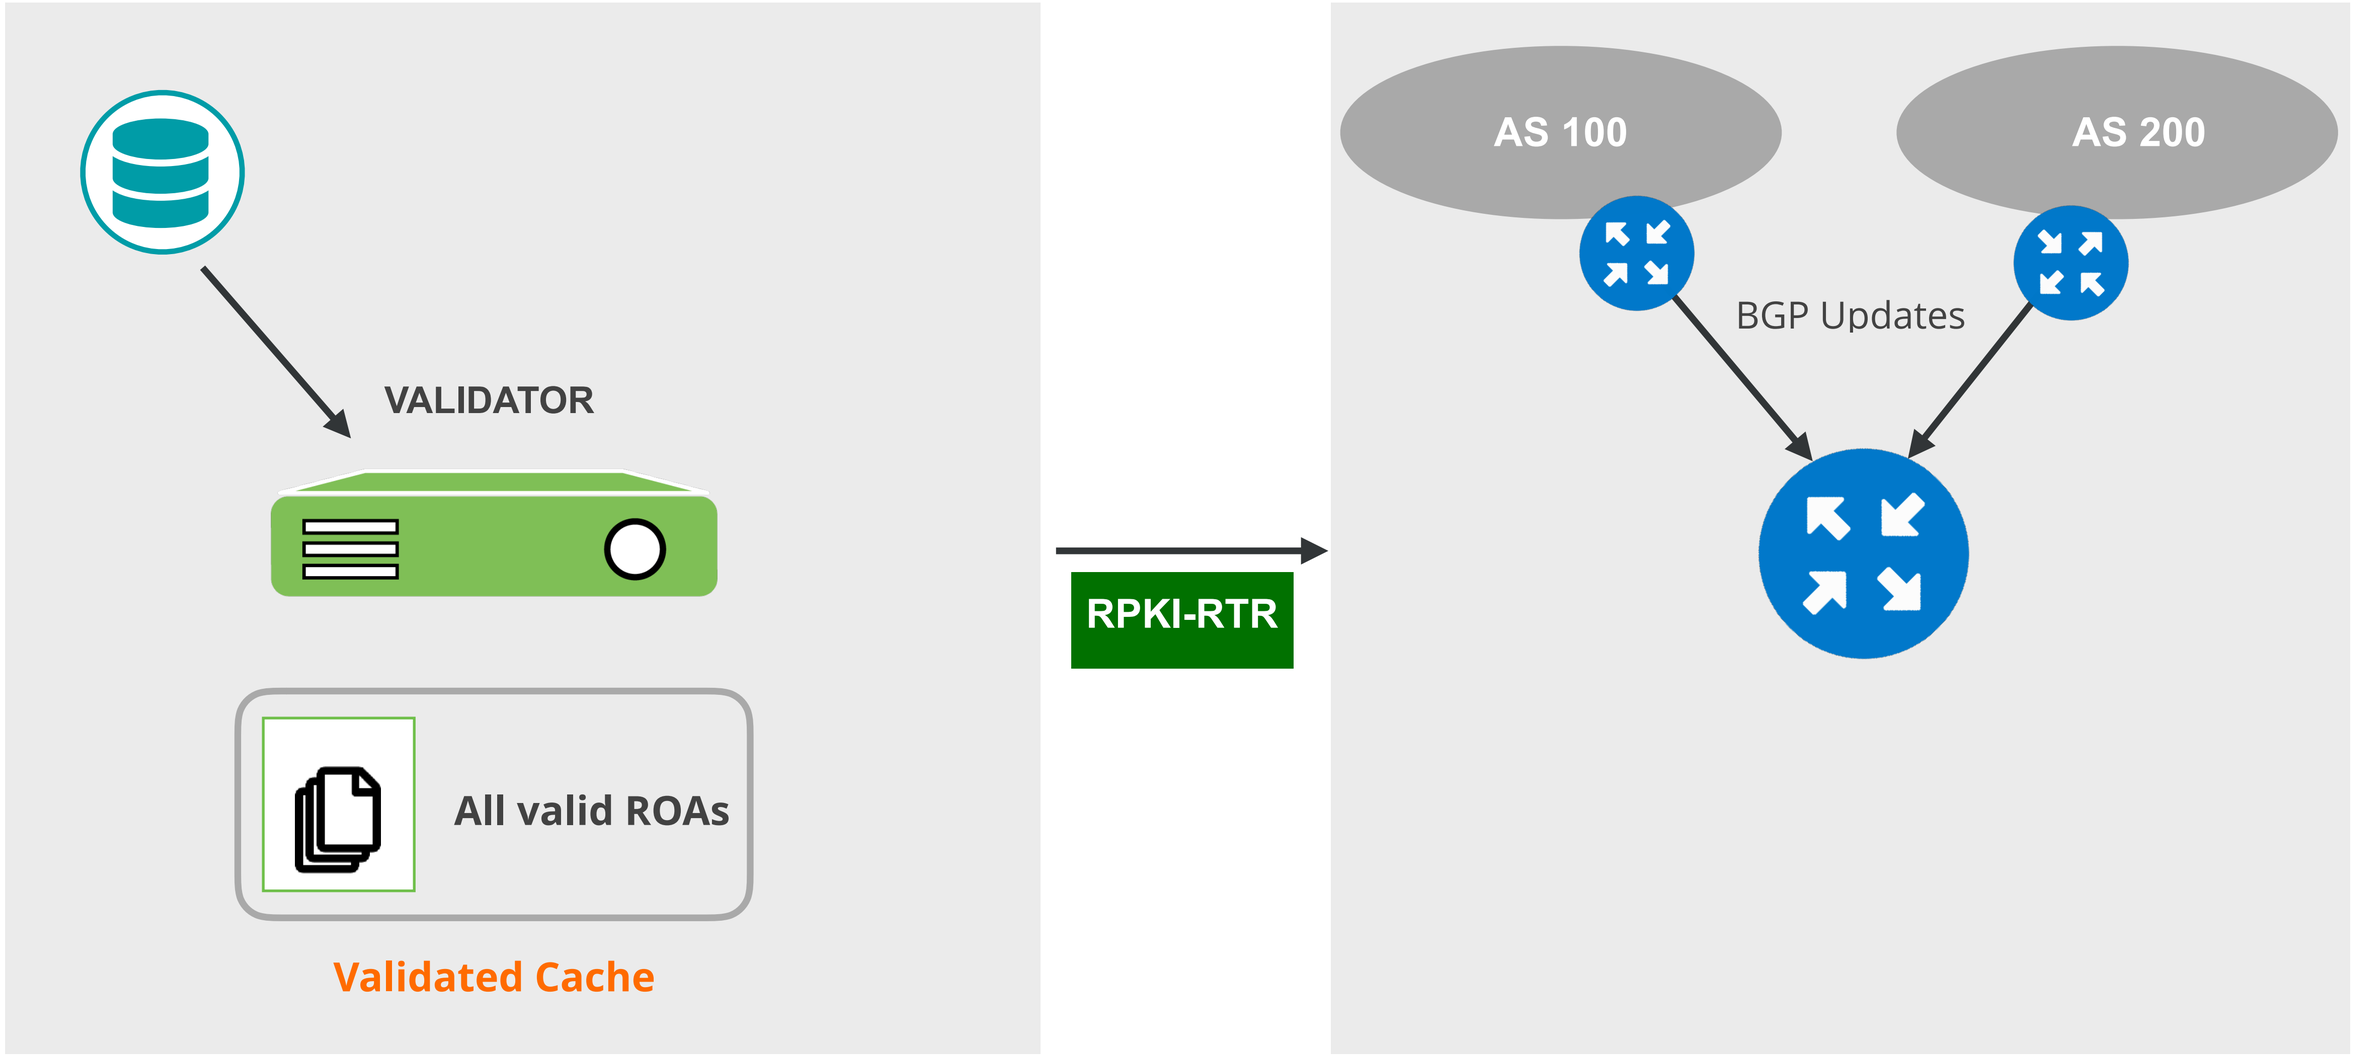

BGP Origin Validation

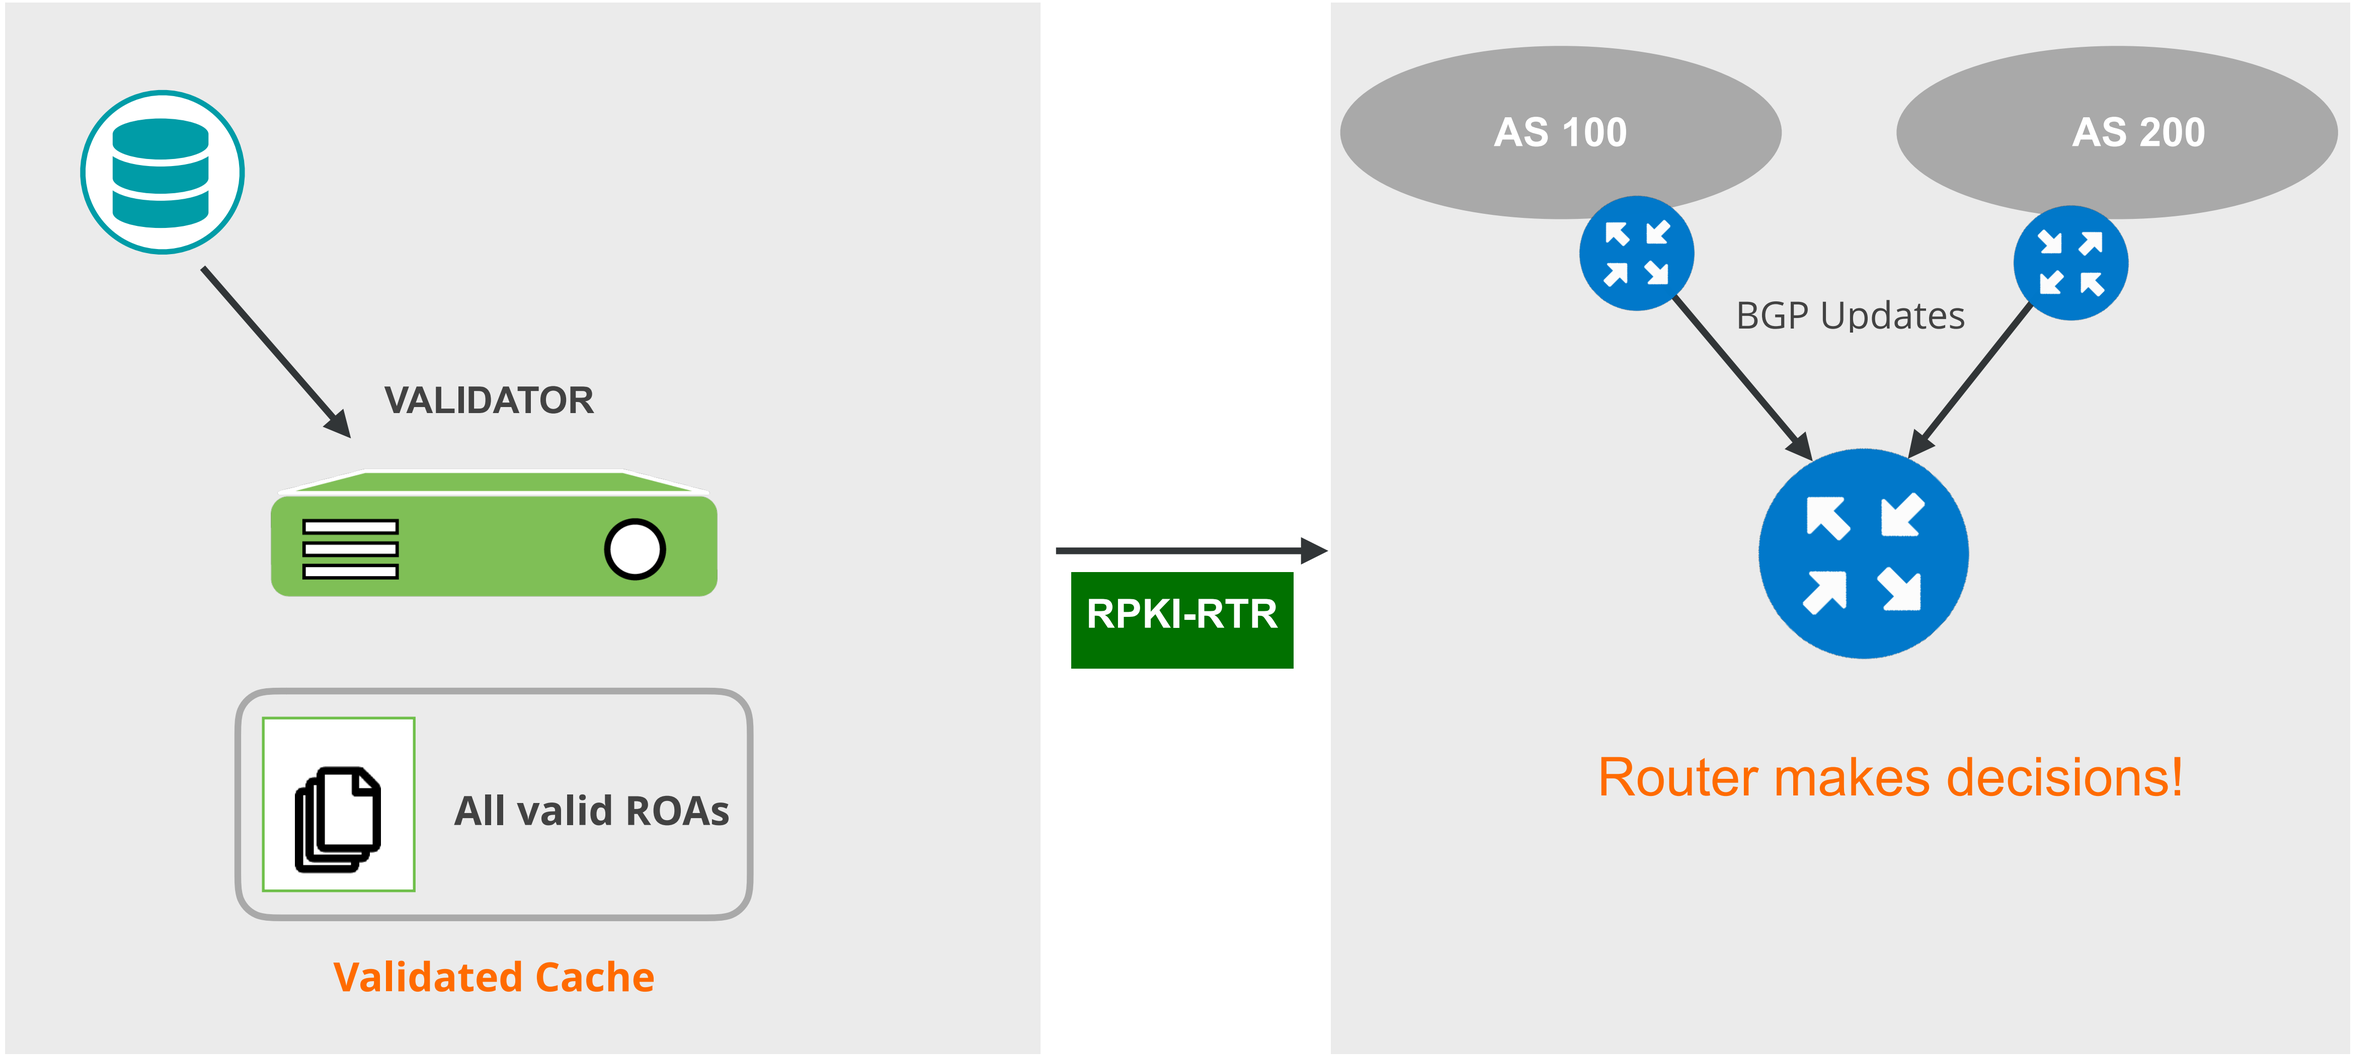

BGP Origin Validation

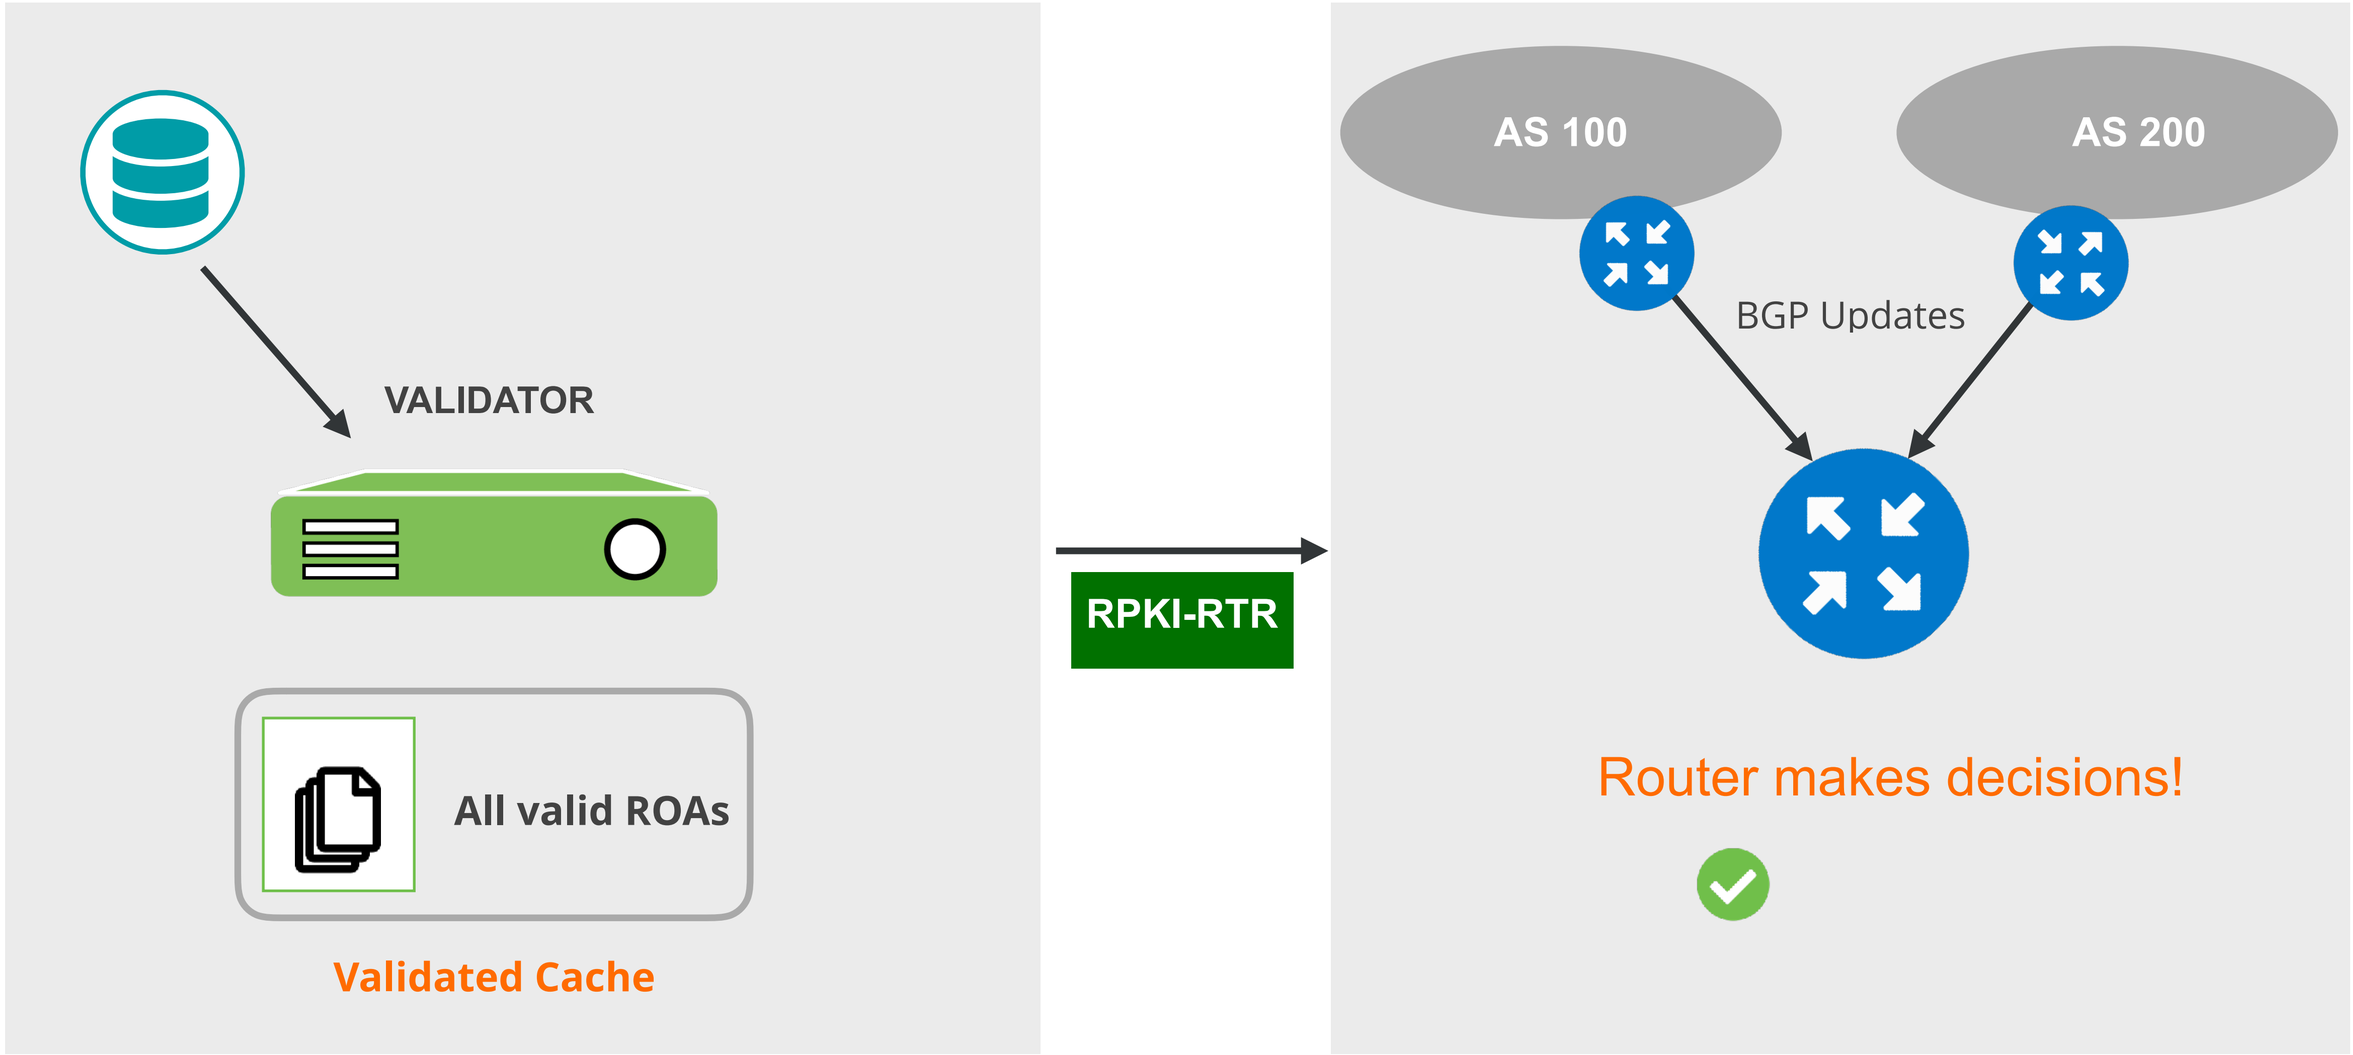

BGP Origin Validation

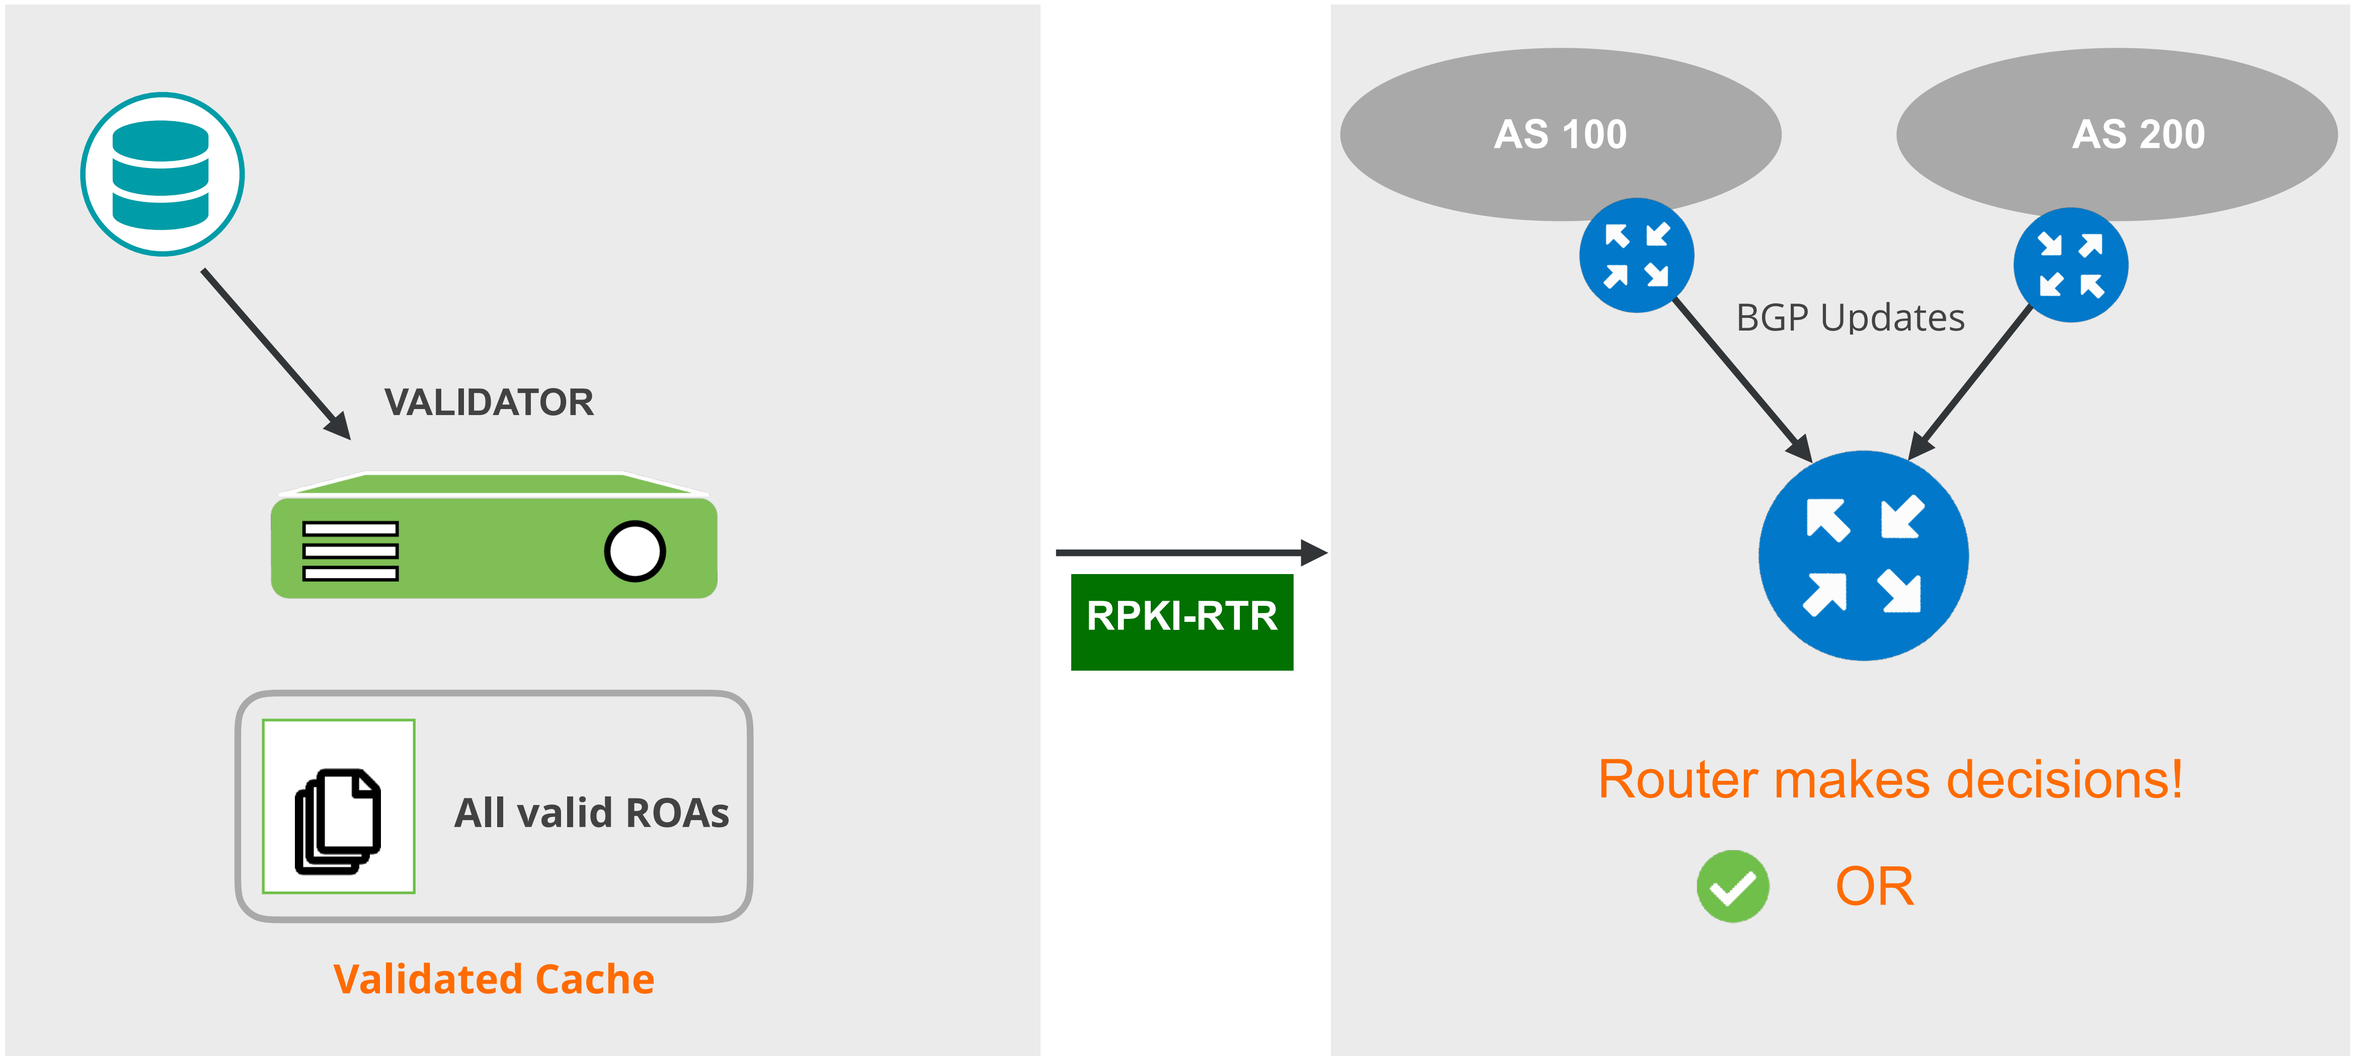

BGP Origin Validation

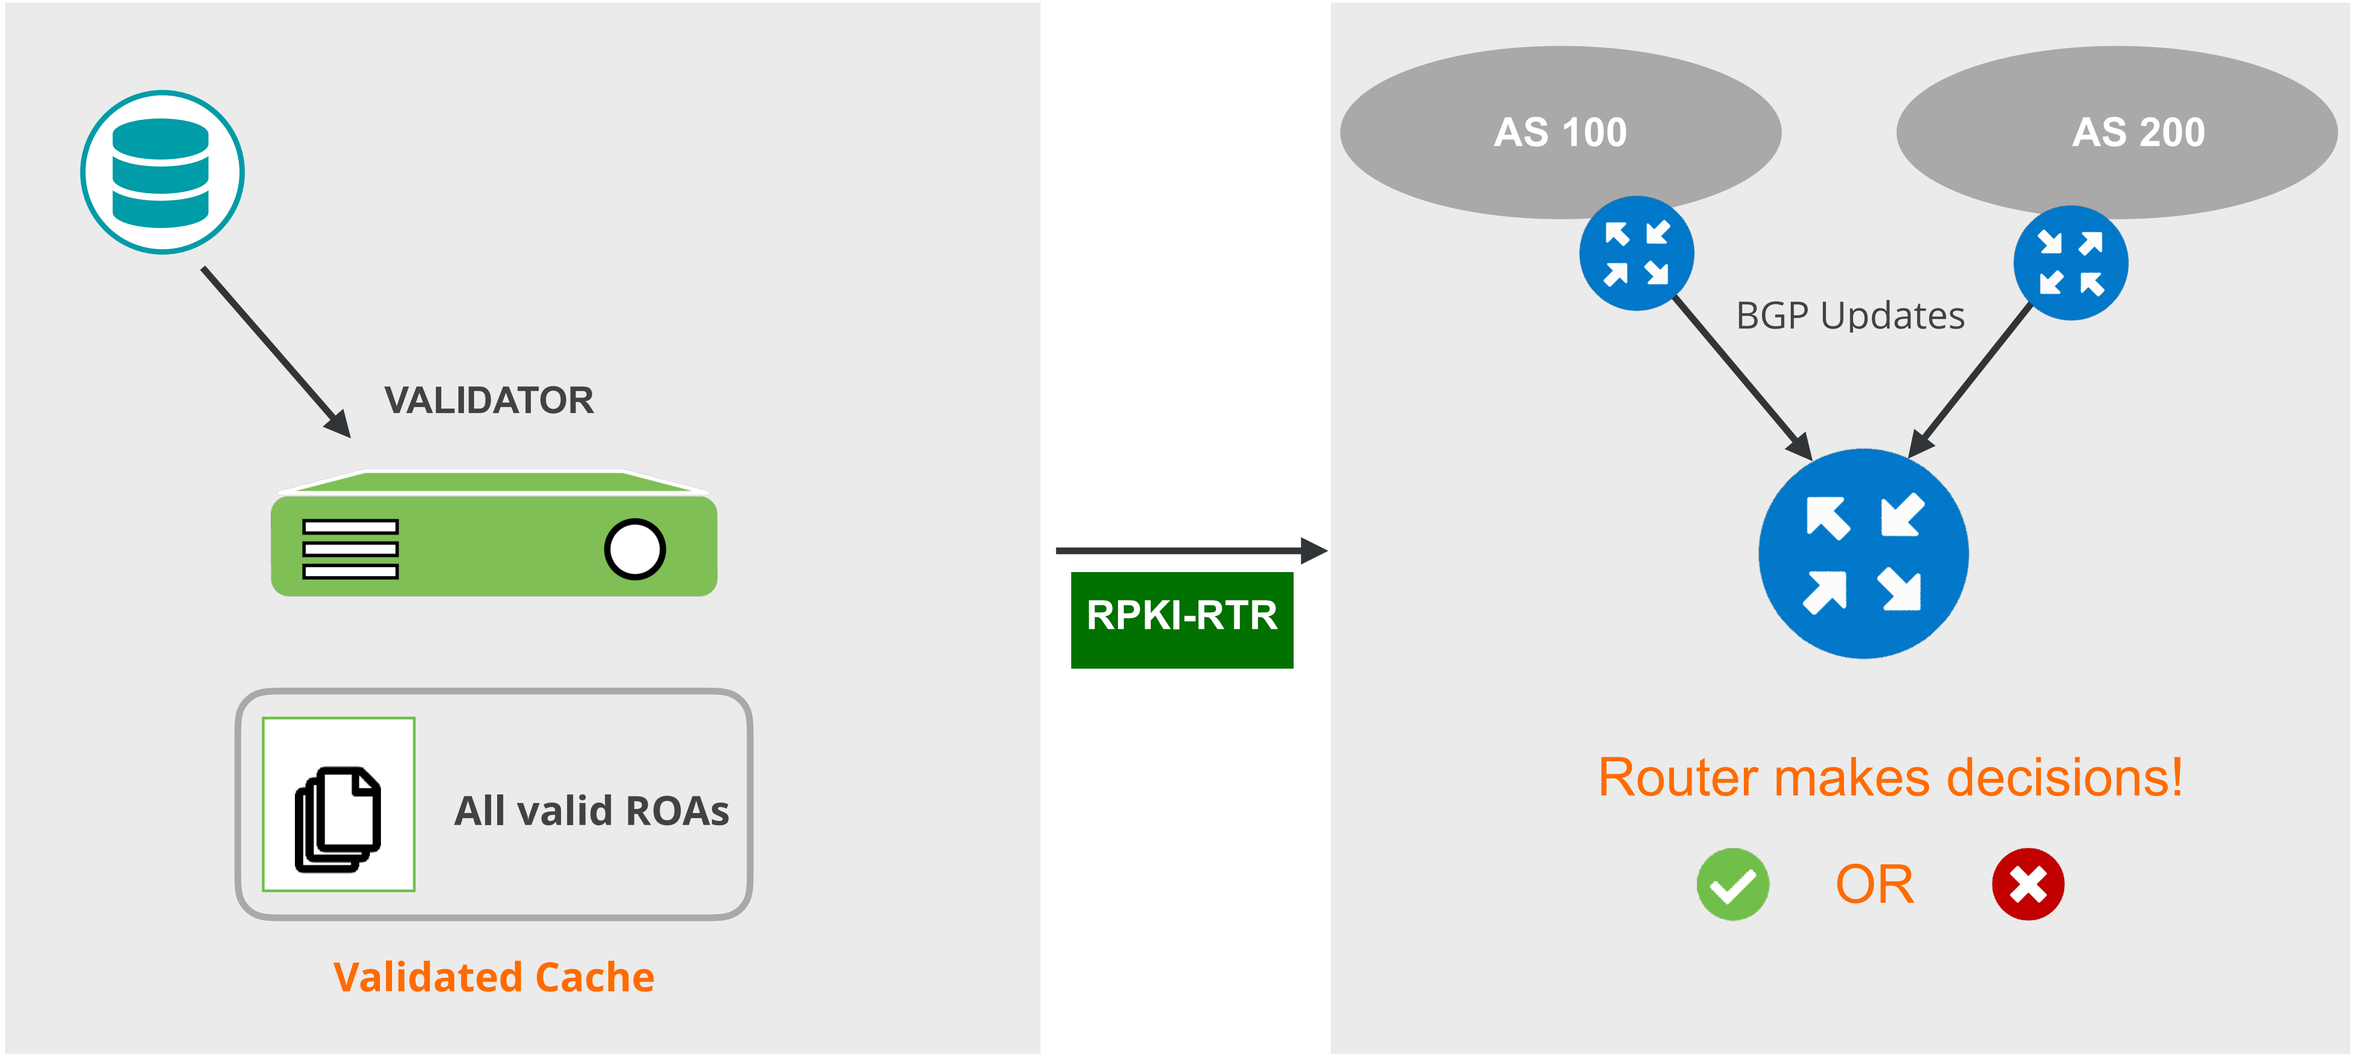

Let’s explain it with examples…

BGP Valid

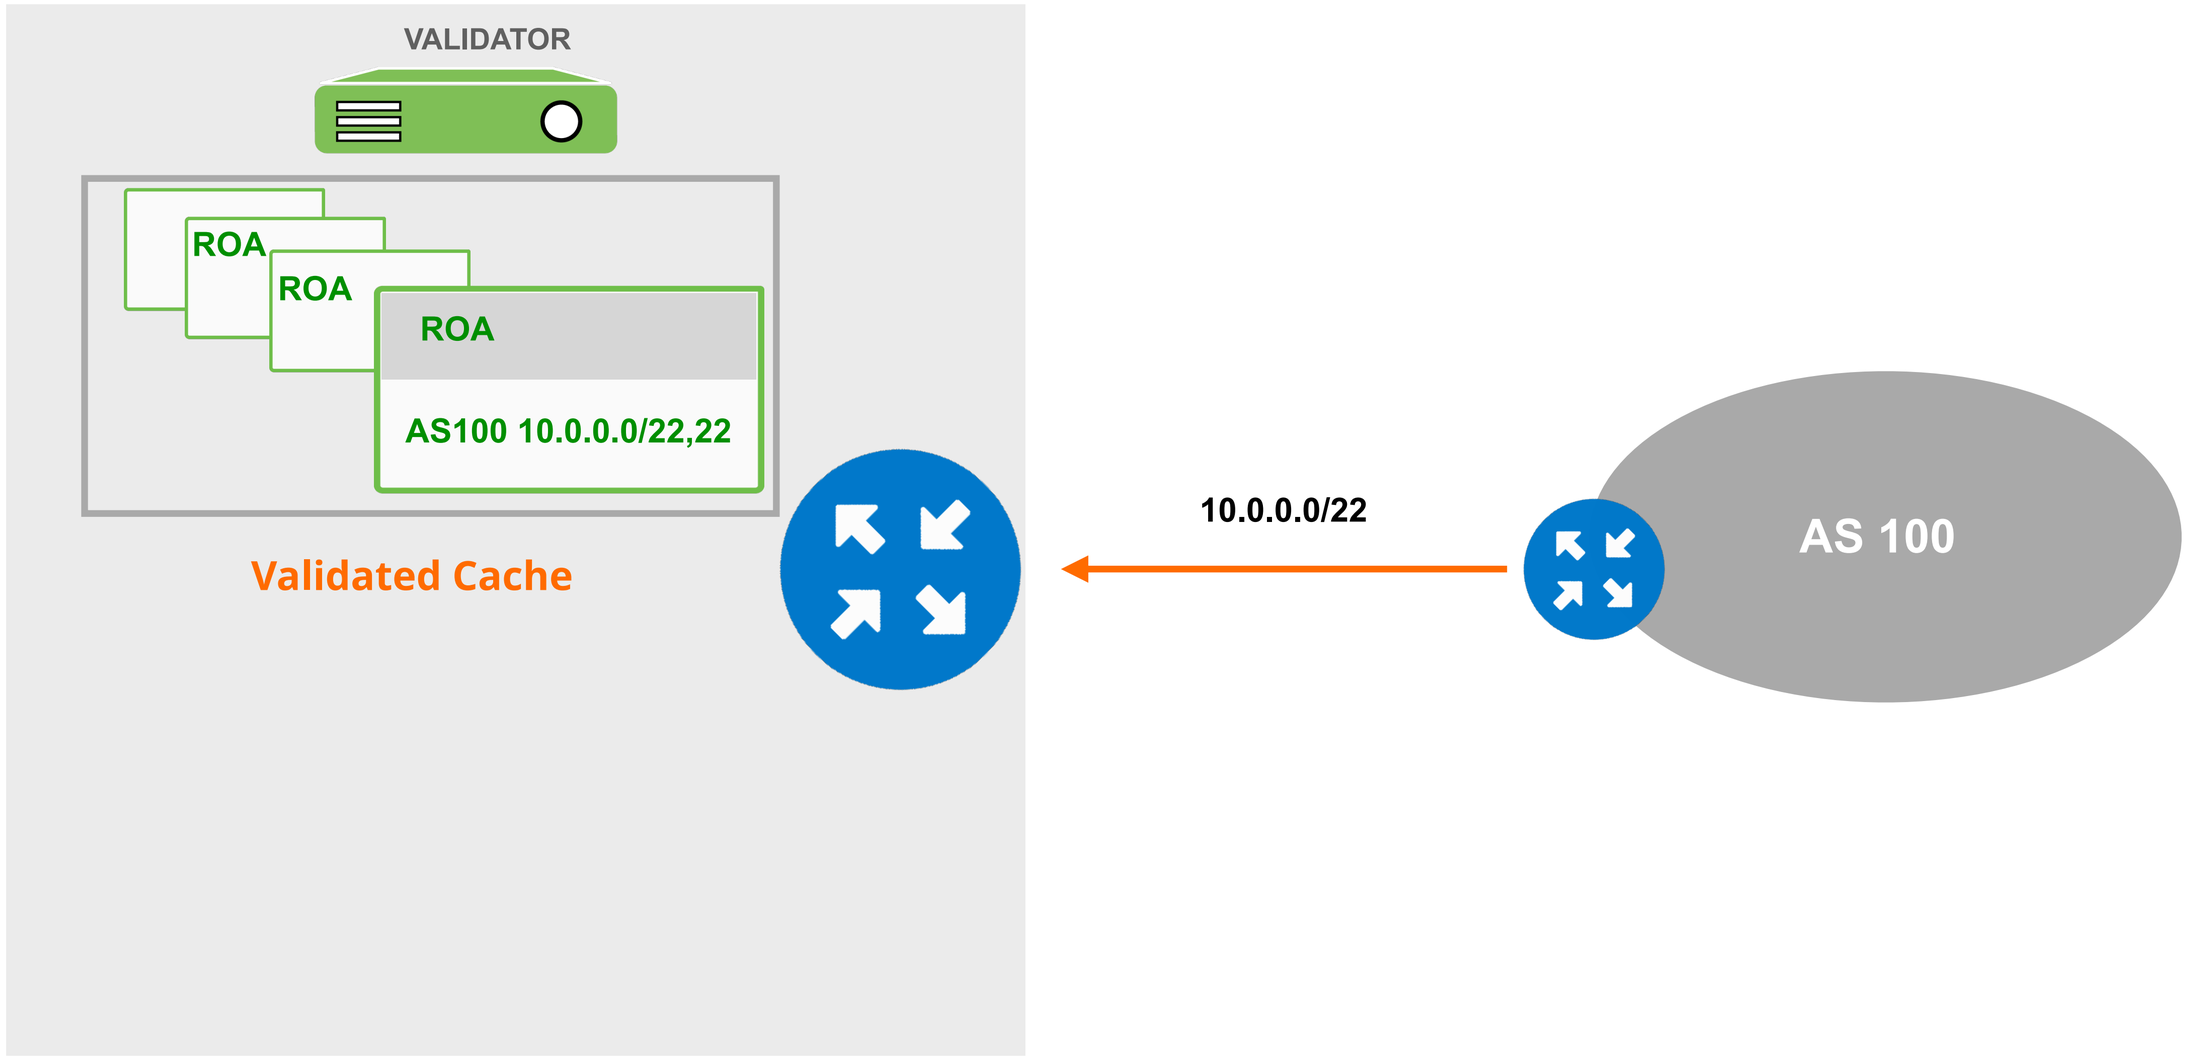
VALIDATOR
ROA
ROA
ROA
AS100 10.0.0.0/22,22
Validated Cache
10.0.0.0/22
AS 100

BGP Valid

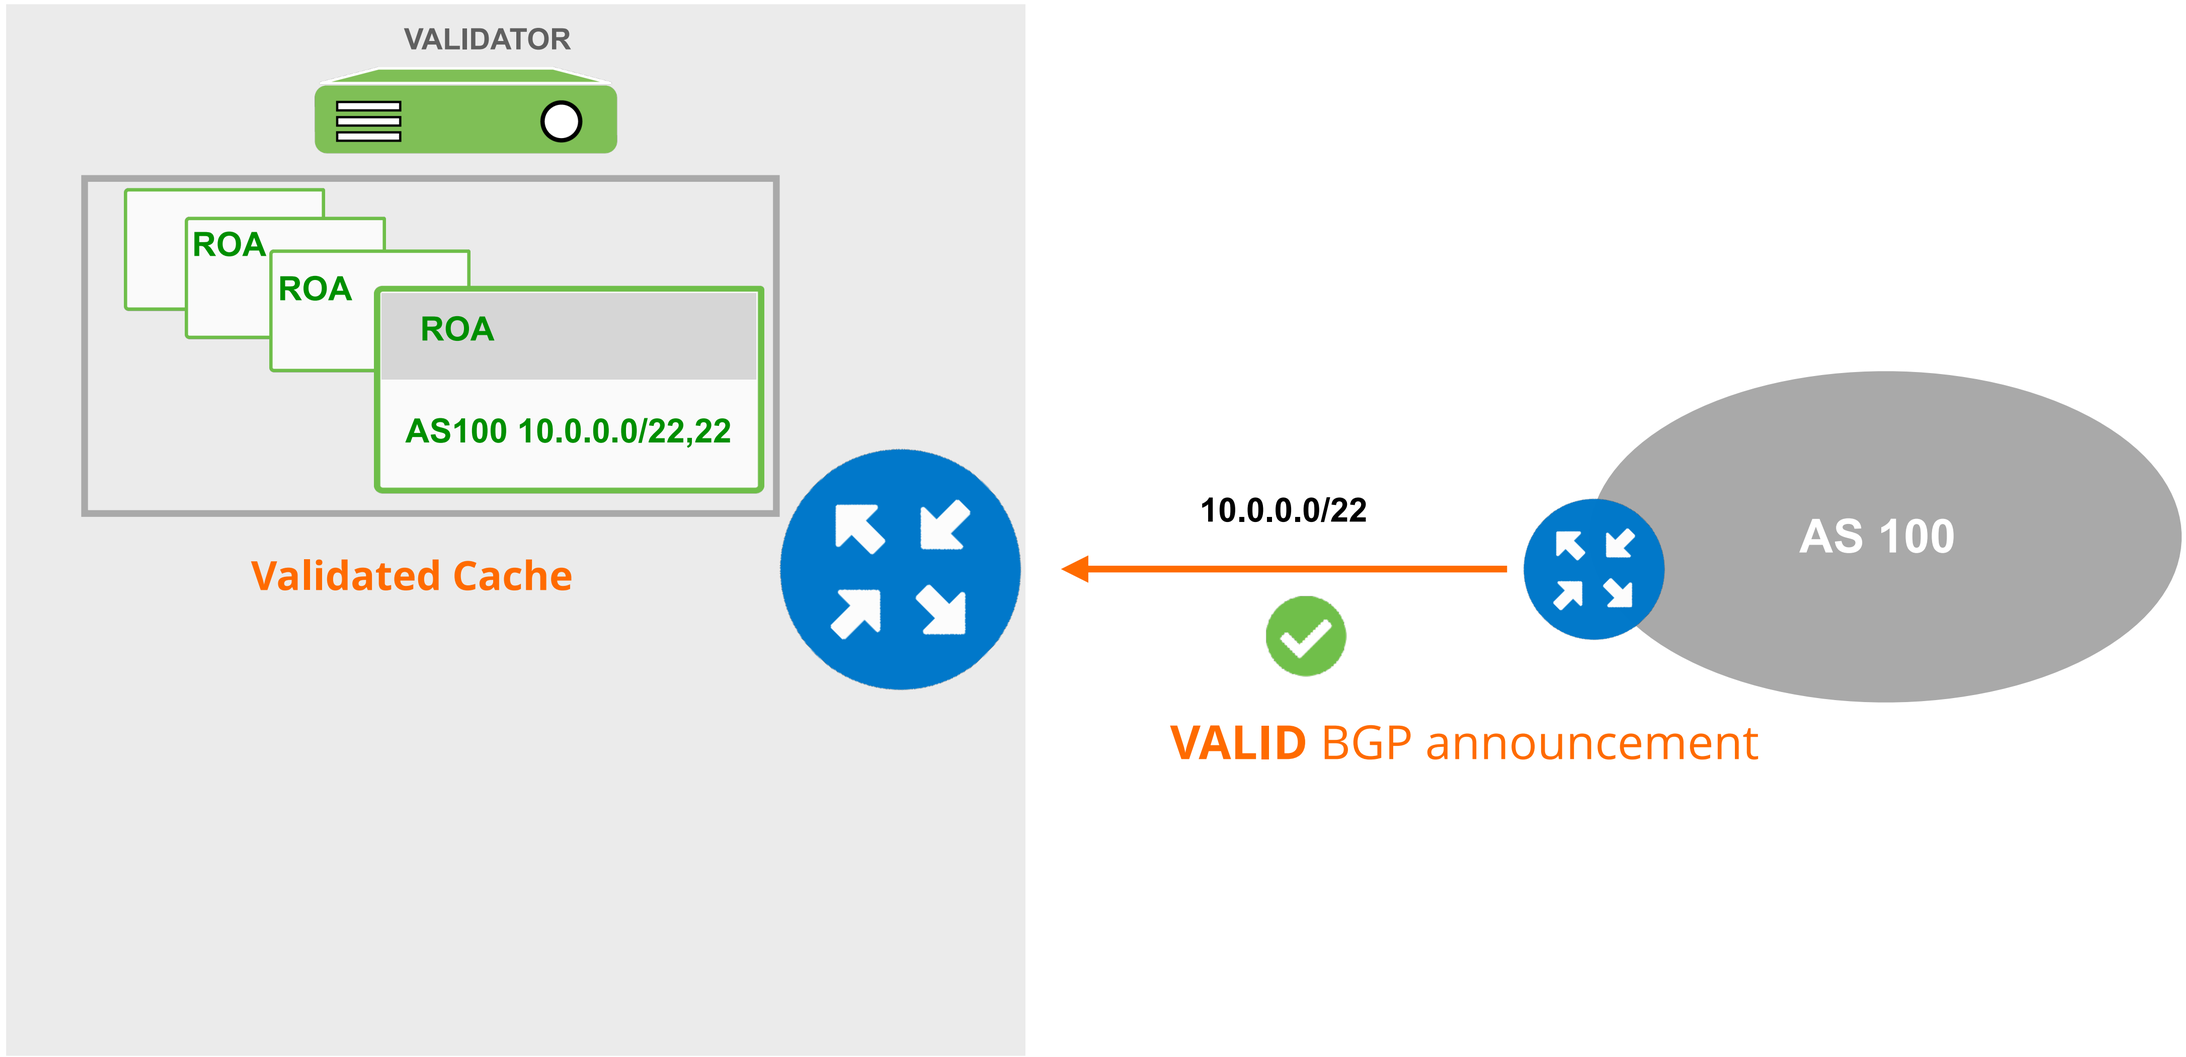
VALIDATOR
ROA
ROA
ROA
AS100 10.0.0.0/22,22
Validated Cache
10.0.0.0/22
AS 100
VALID BGP announcement

BGP Invalid

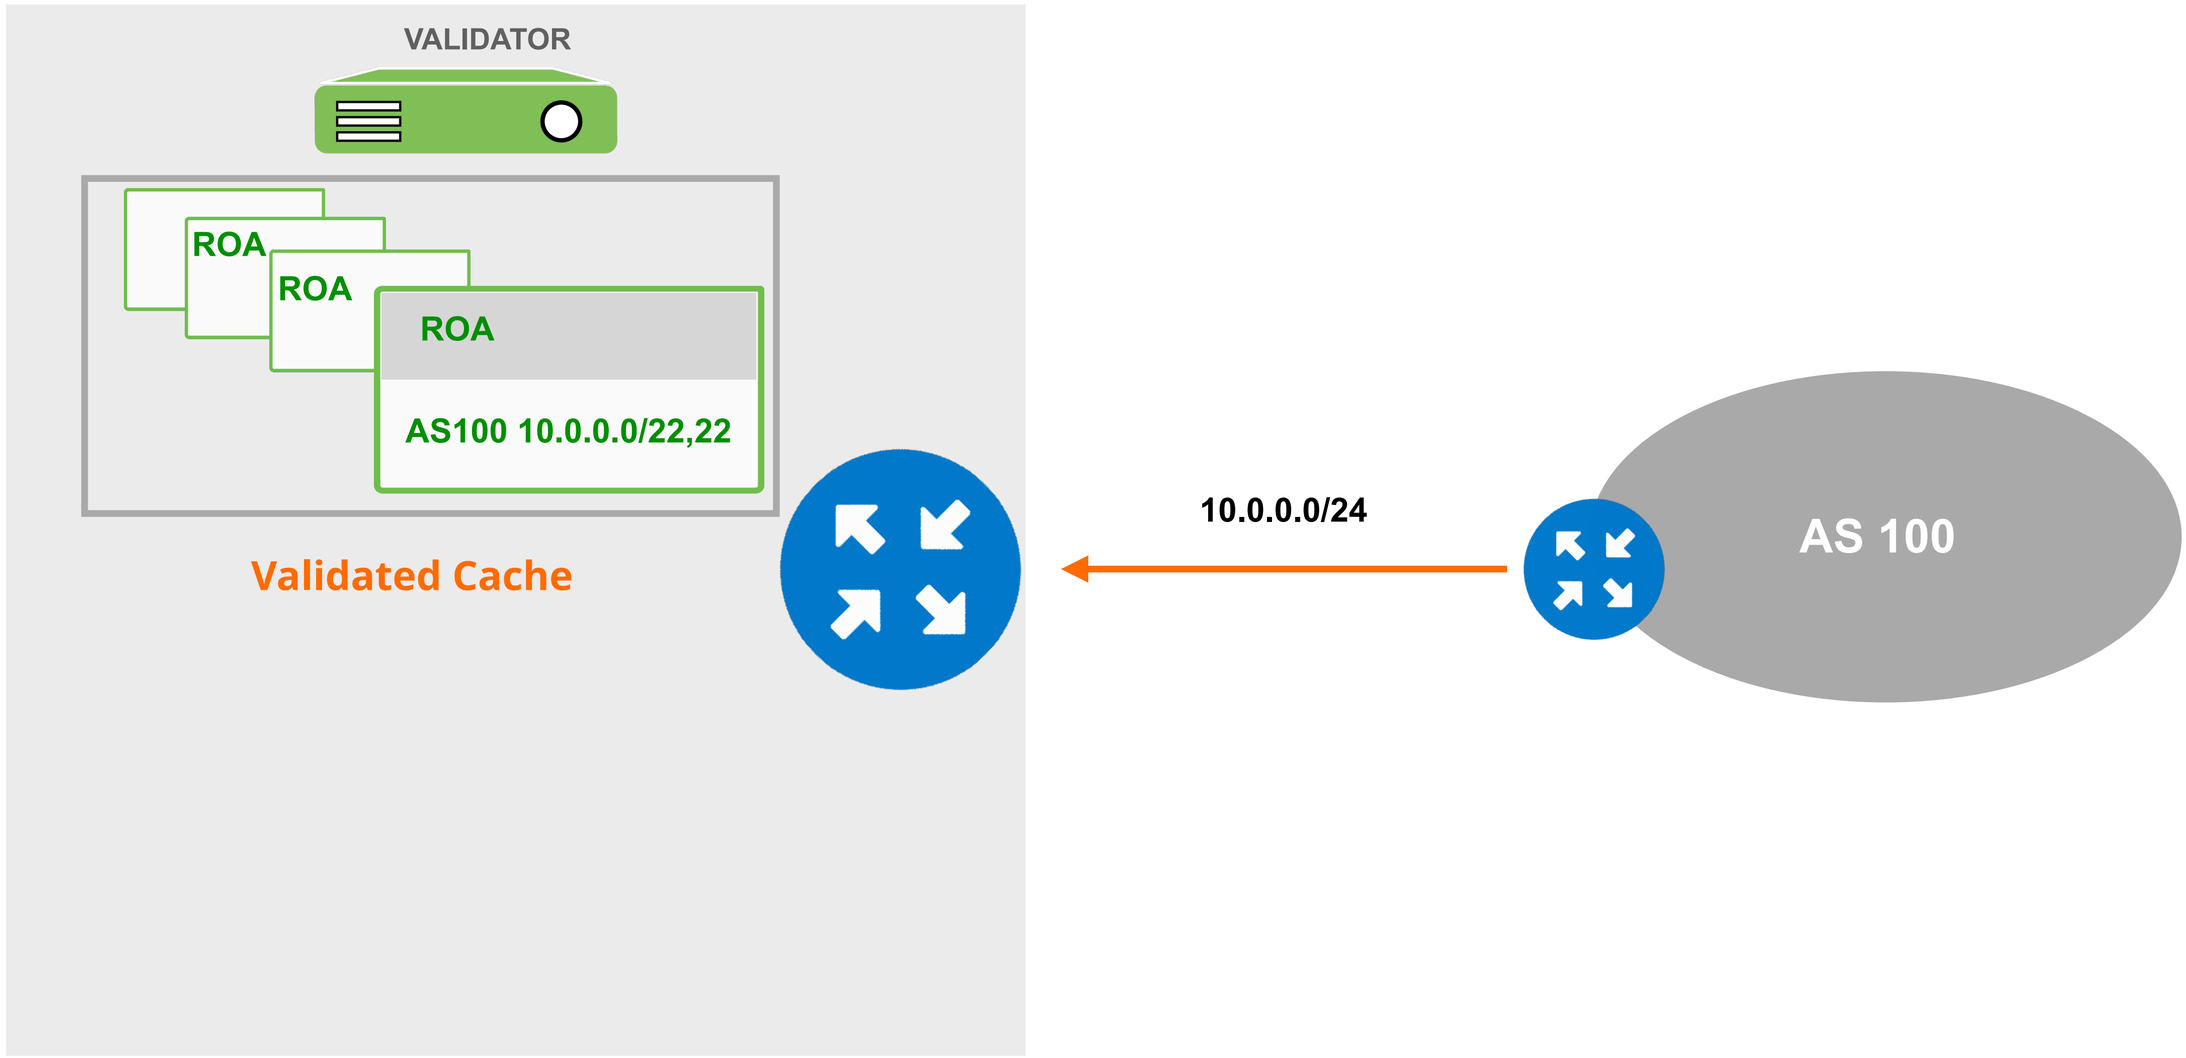
VALIDATOR
ROA
ROA
ROA
AS100 10.0.0.0/22,22
Validated Cache
10.0.0.0/24
AS 100

BGP Invalid

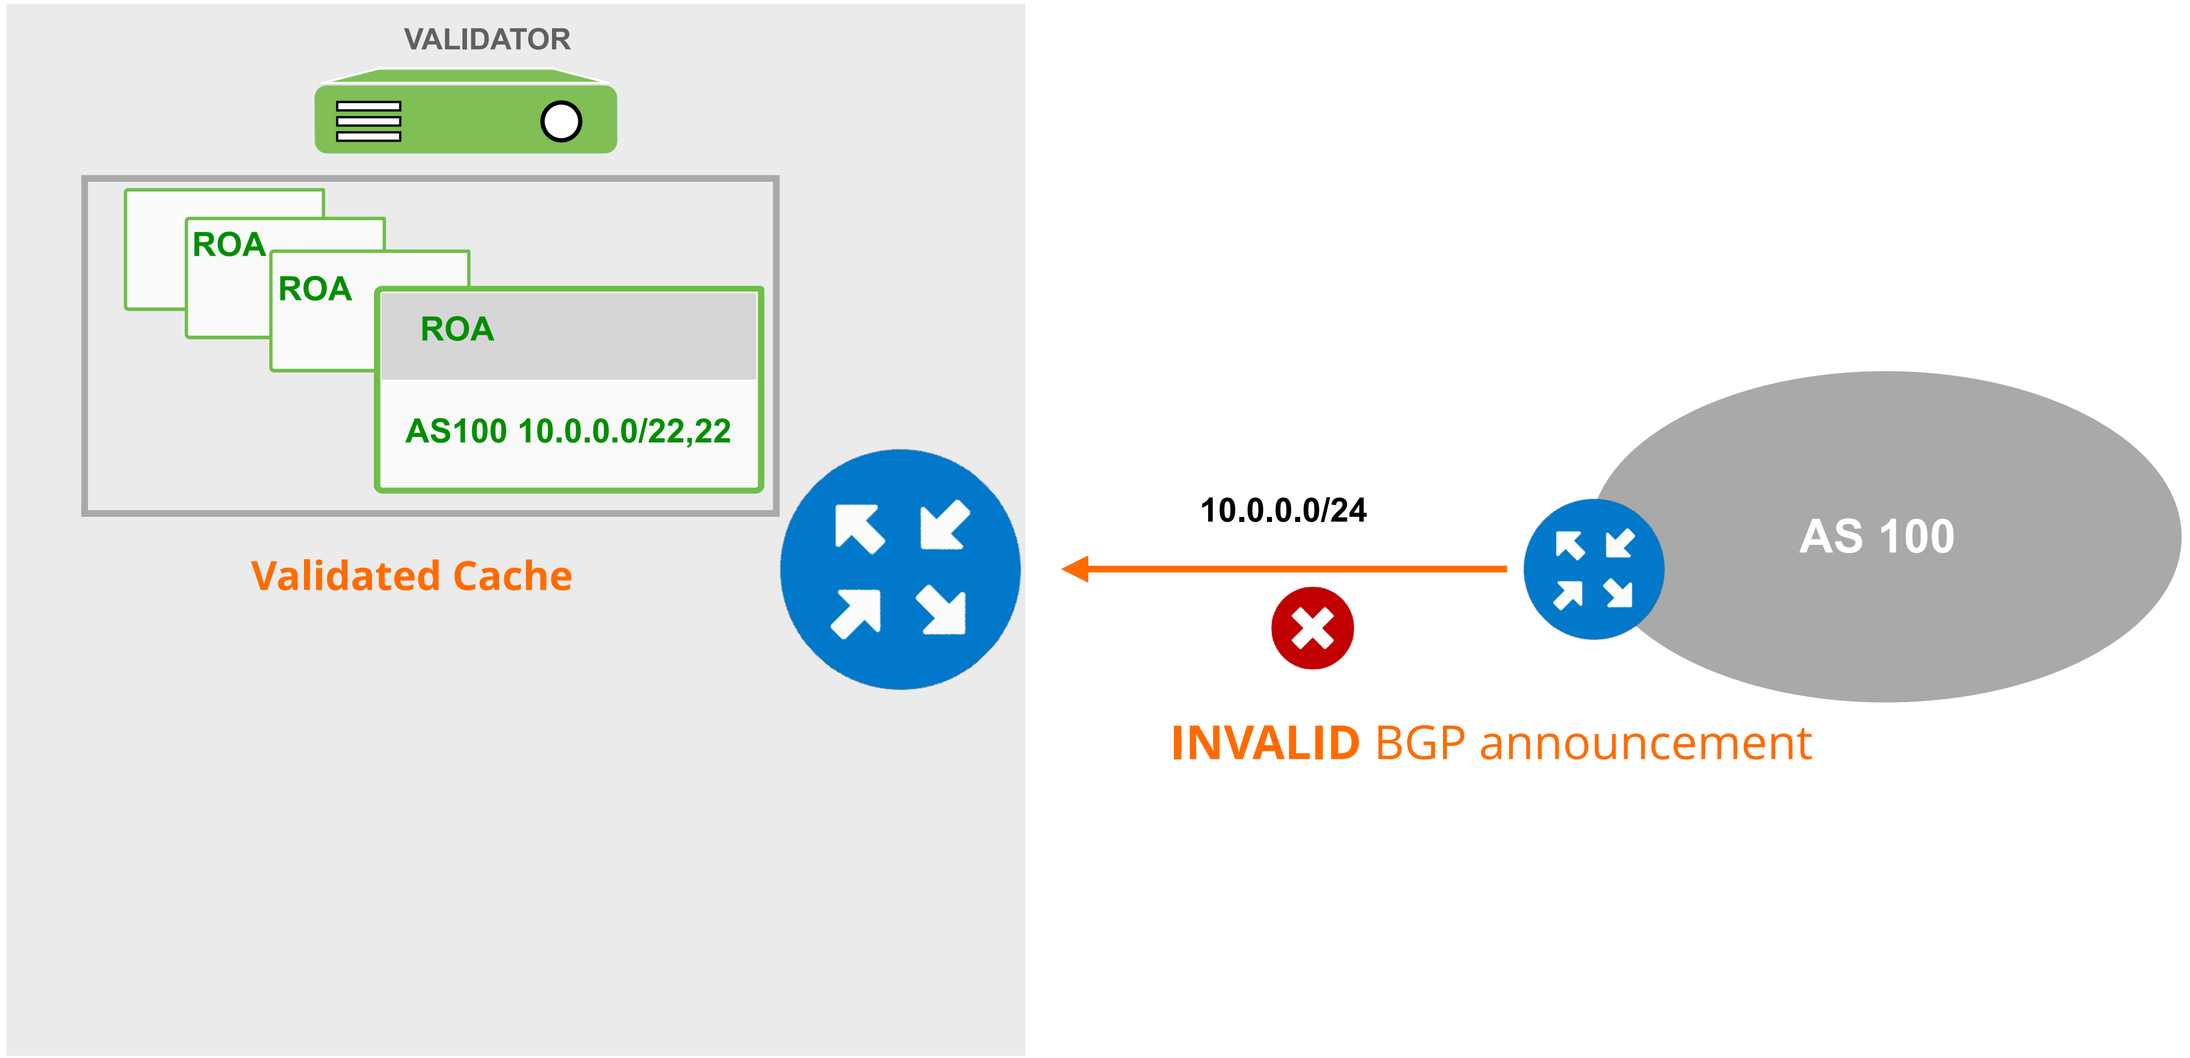
VALIDATOR
ROA
ROA
ROA
AS100 10.0.0.0/22,22
Validated Cache
10.0.0.0/24
AS 100
INVALID BGP announcement

BGP Invalid

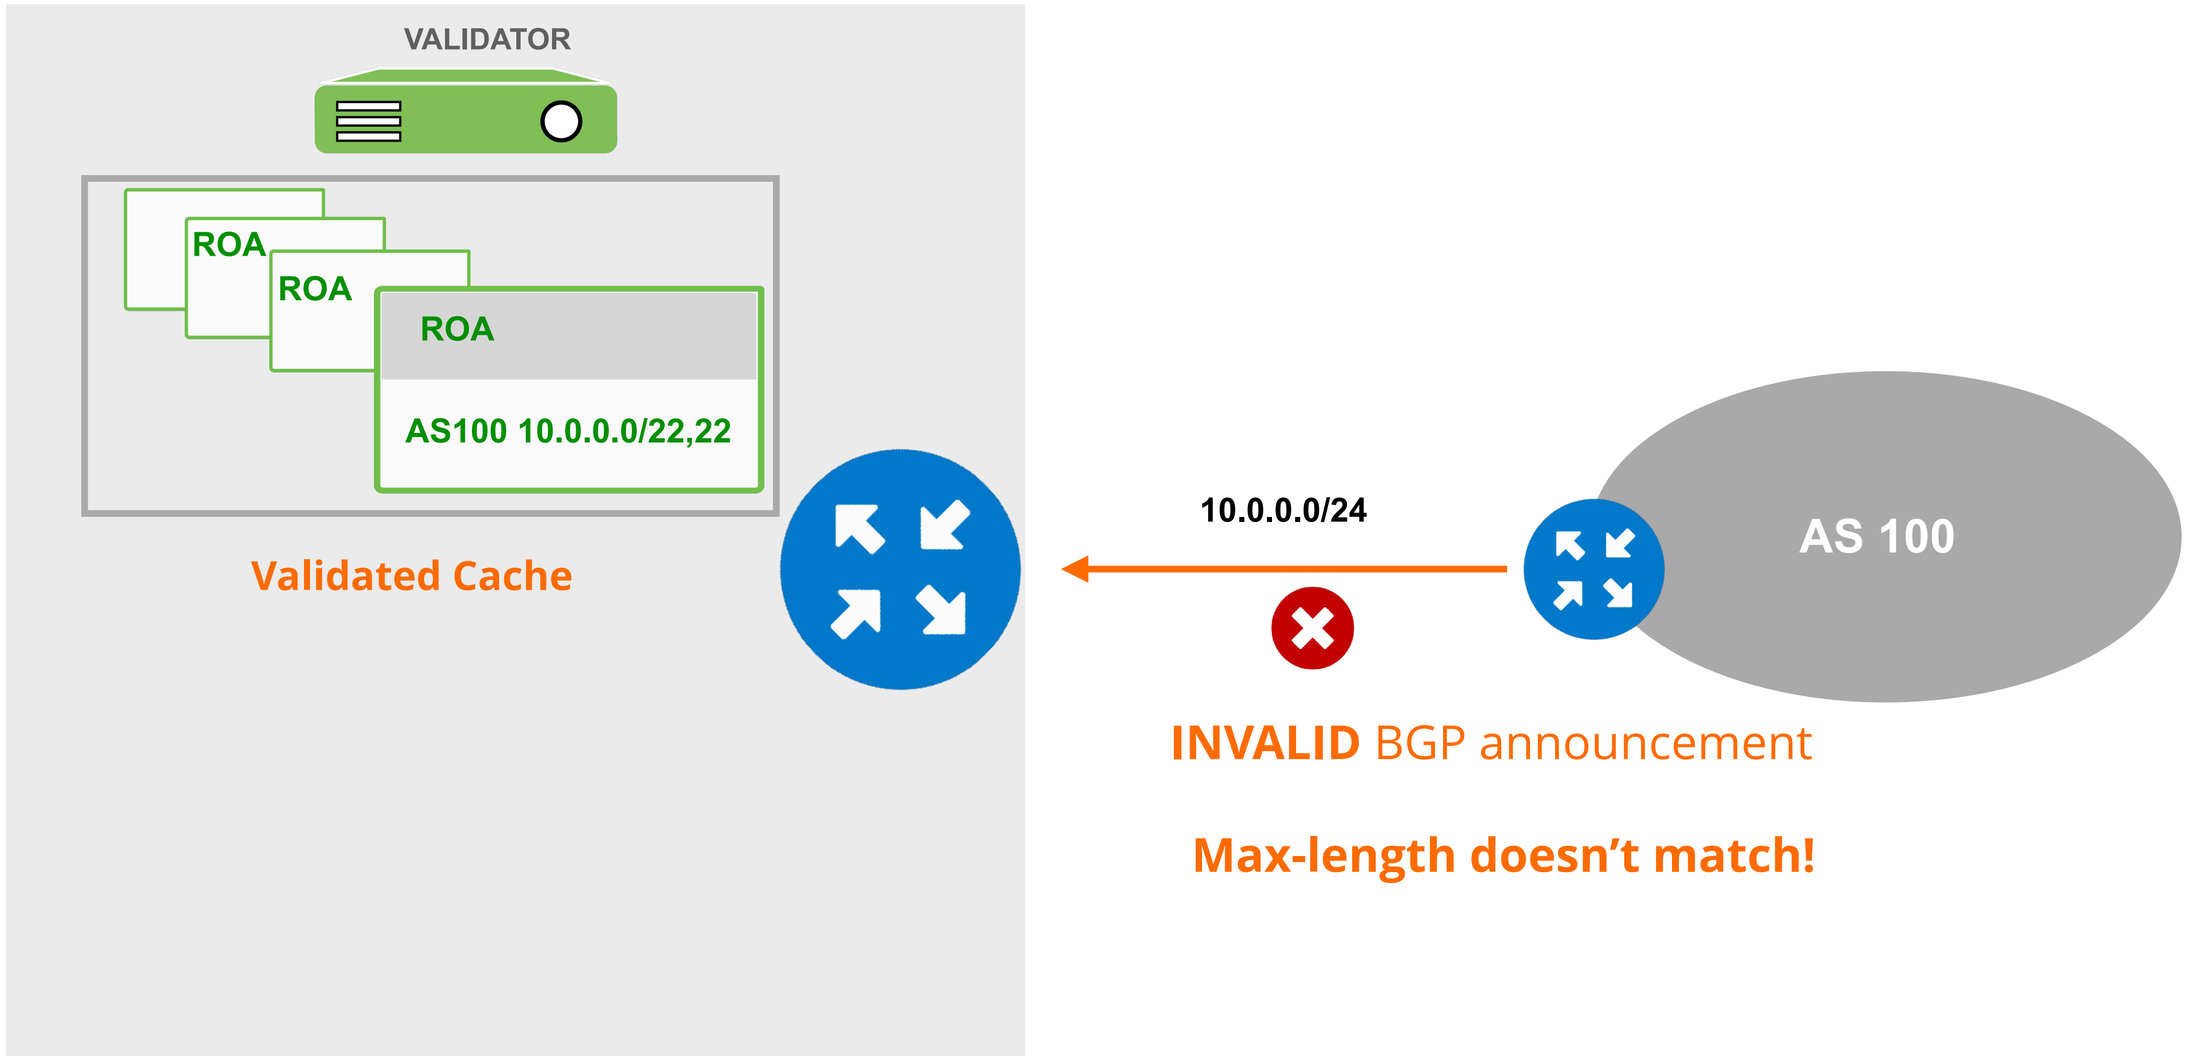
VALIDATOR
ROA
ROA
ROA
AS100 10.0.0.0/22,22
Validated Cache
10.0.0.0/24
AS 100
INVALID BGP announcement
Max-length doesn't match!

BGP Invalid

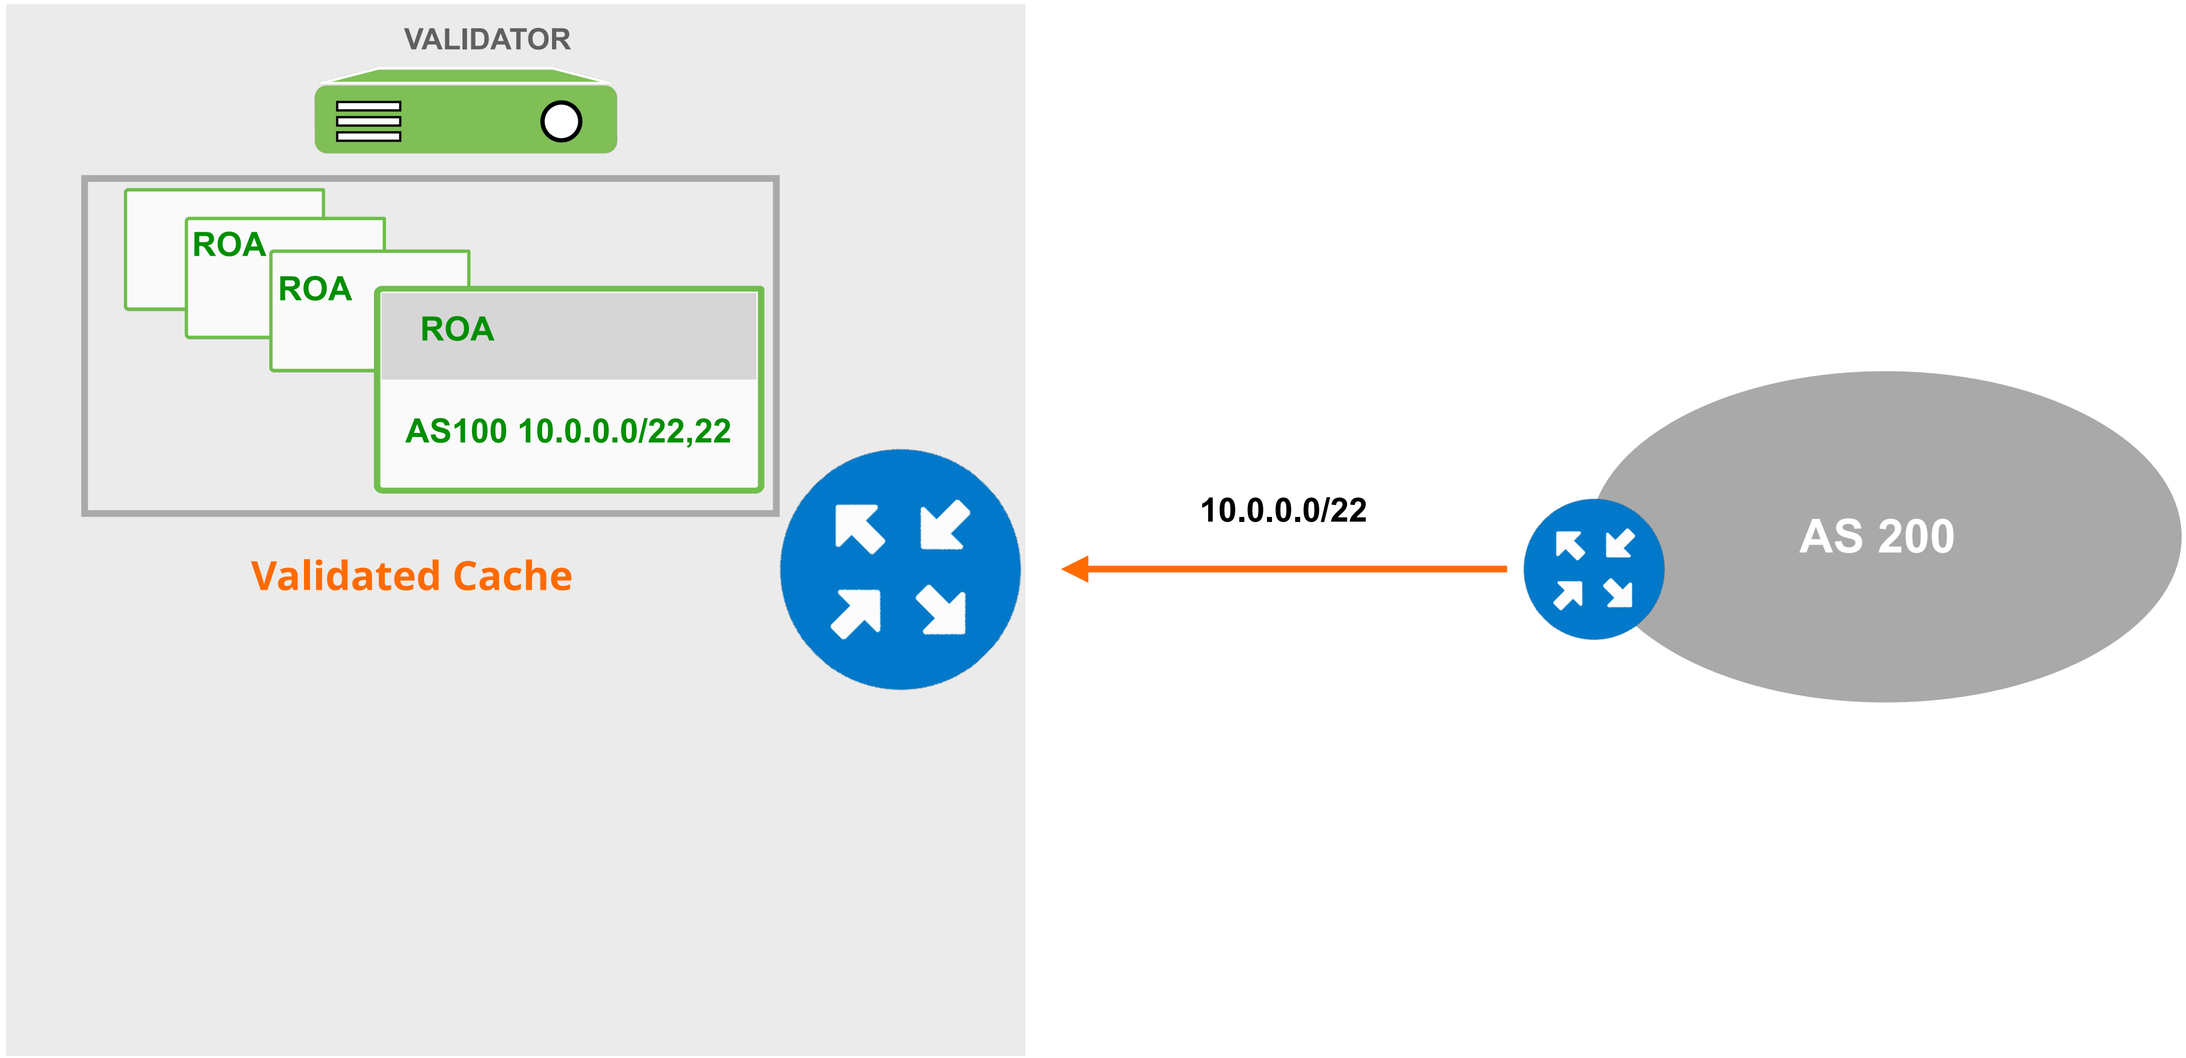
VALIDATOR
ROA
ROA
ROA
AS100 10.0.0.0/22,22
Validated Cache
10.0.0.0/22
AS 200

BGP Invalid

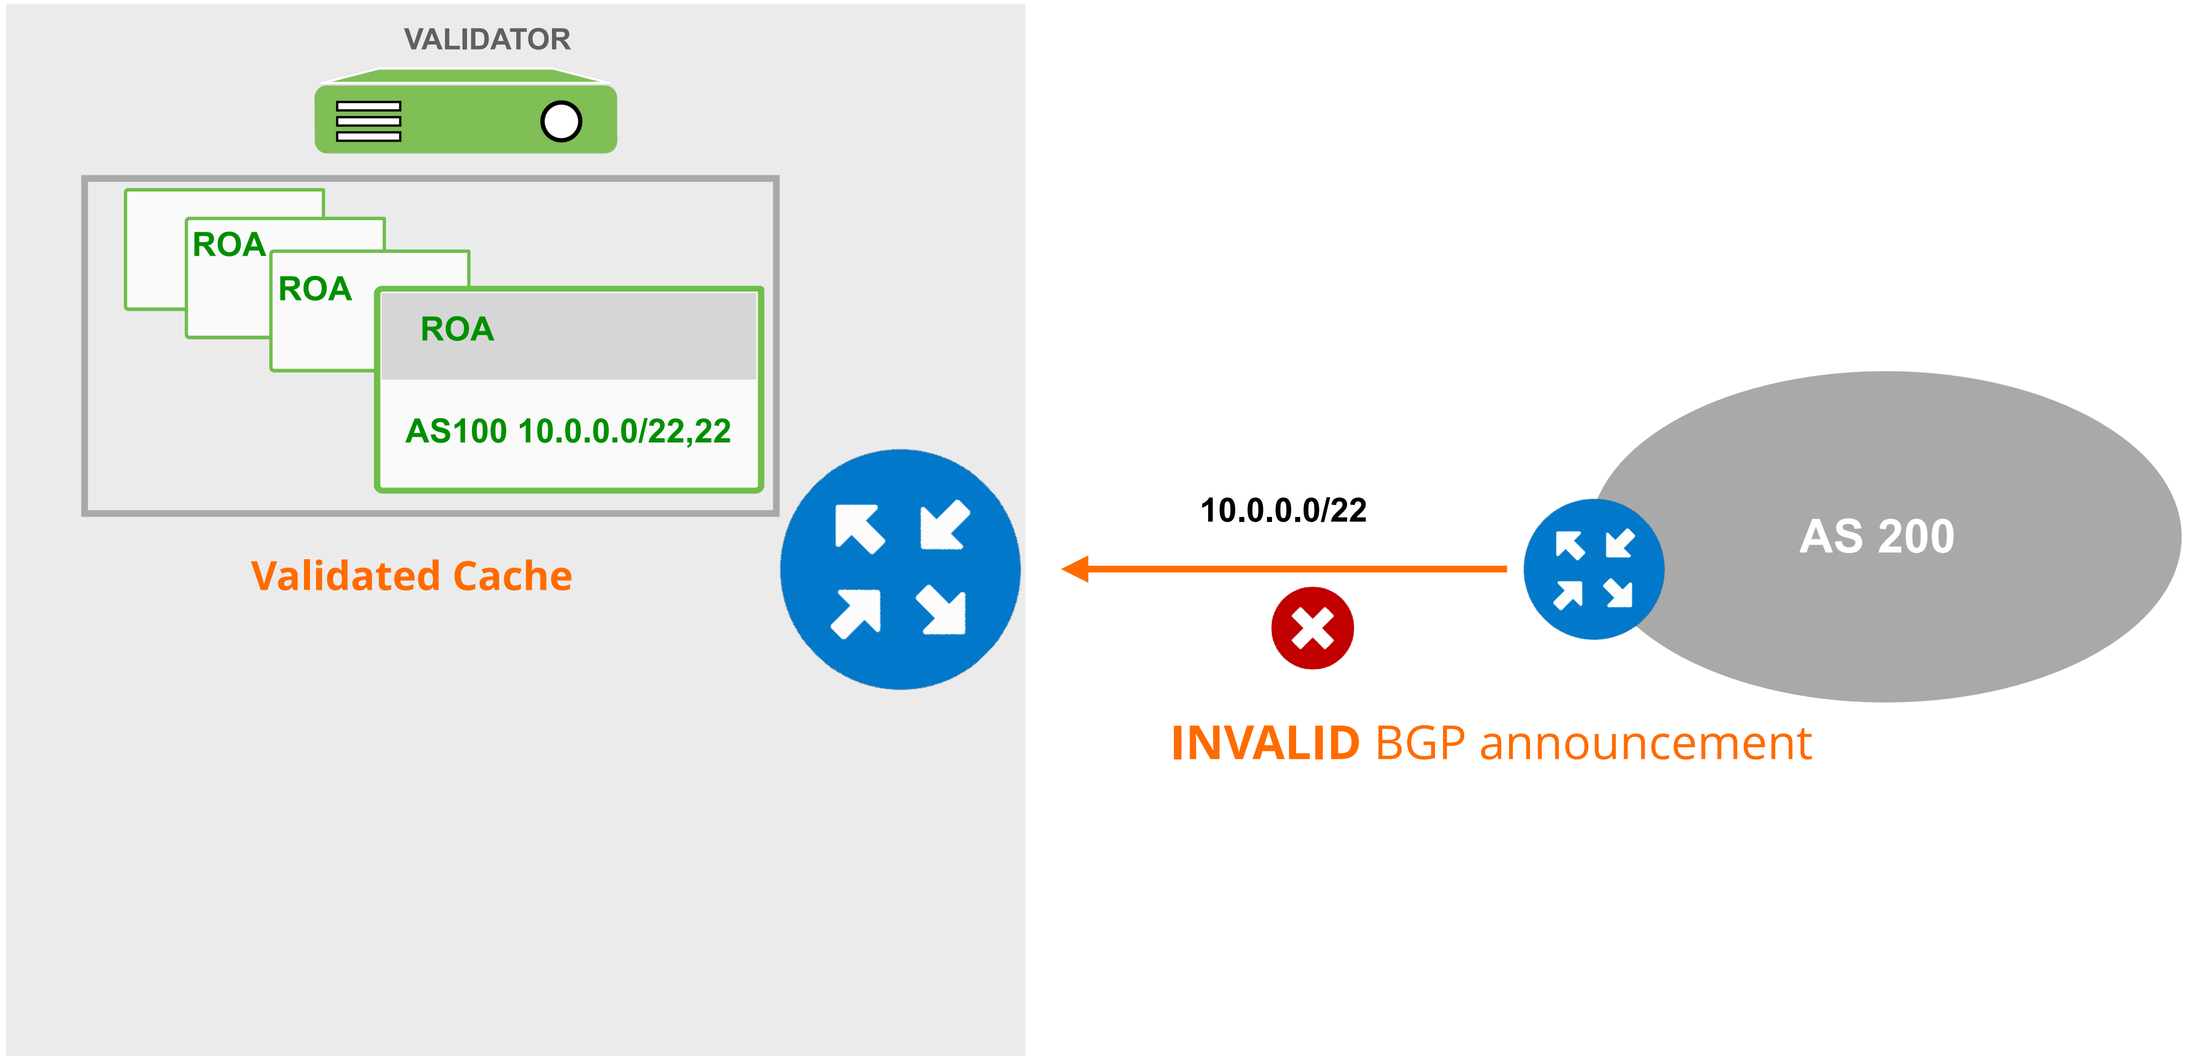
VALIDATOR
ROA
ROA
ROA
AS100 10.0.0.0/22,22
Validated Cache
10.0.0.0/22
AS 200
INVALID BGP announcement

BGP Invalid

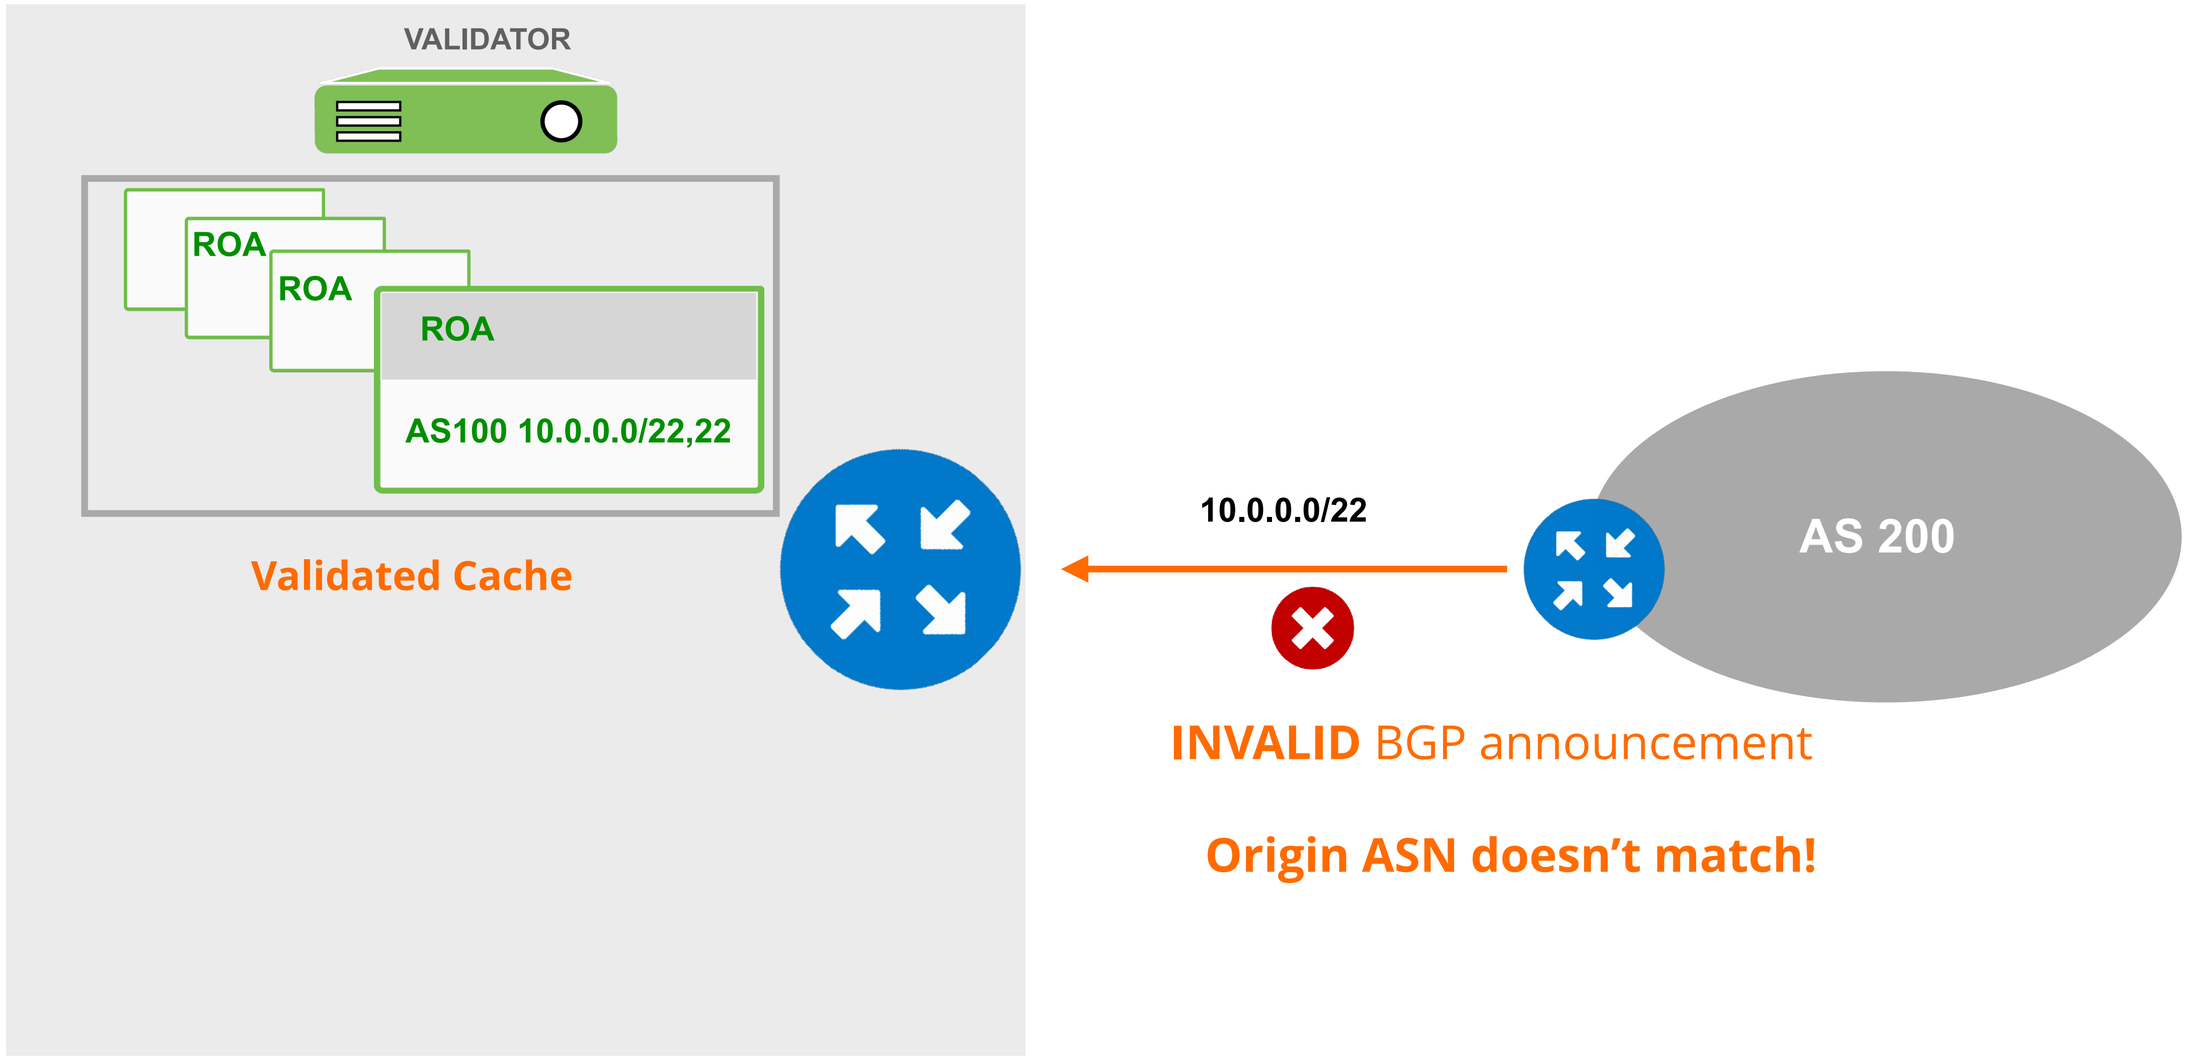
VALIDATOR
ROA
ROA
ROA
AS100 10.0.0.0/22,22
Validated Cache
10.0.0.0/22
AS 200
INVALID BGP announcement
Origin ASN doesn't match!

So, it does two checks to validate BGP announcements:

Max-length and **Origin ASN**

BGP Not-Found

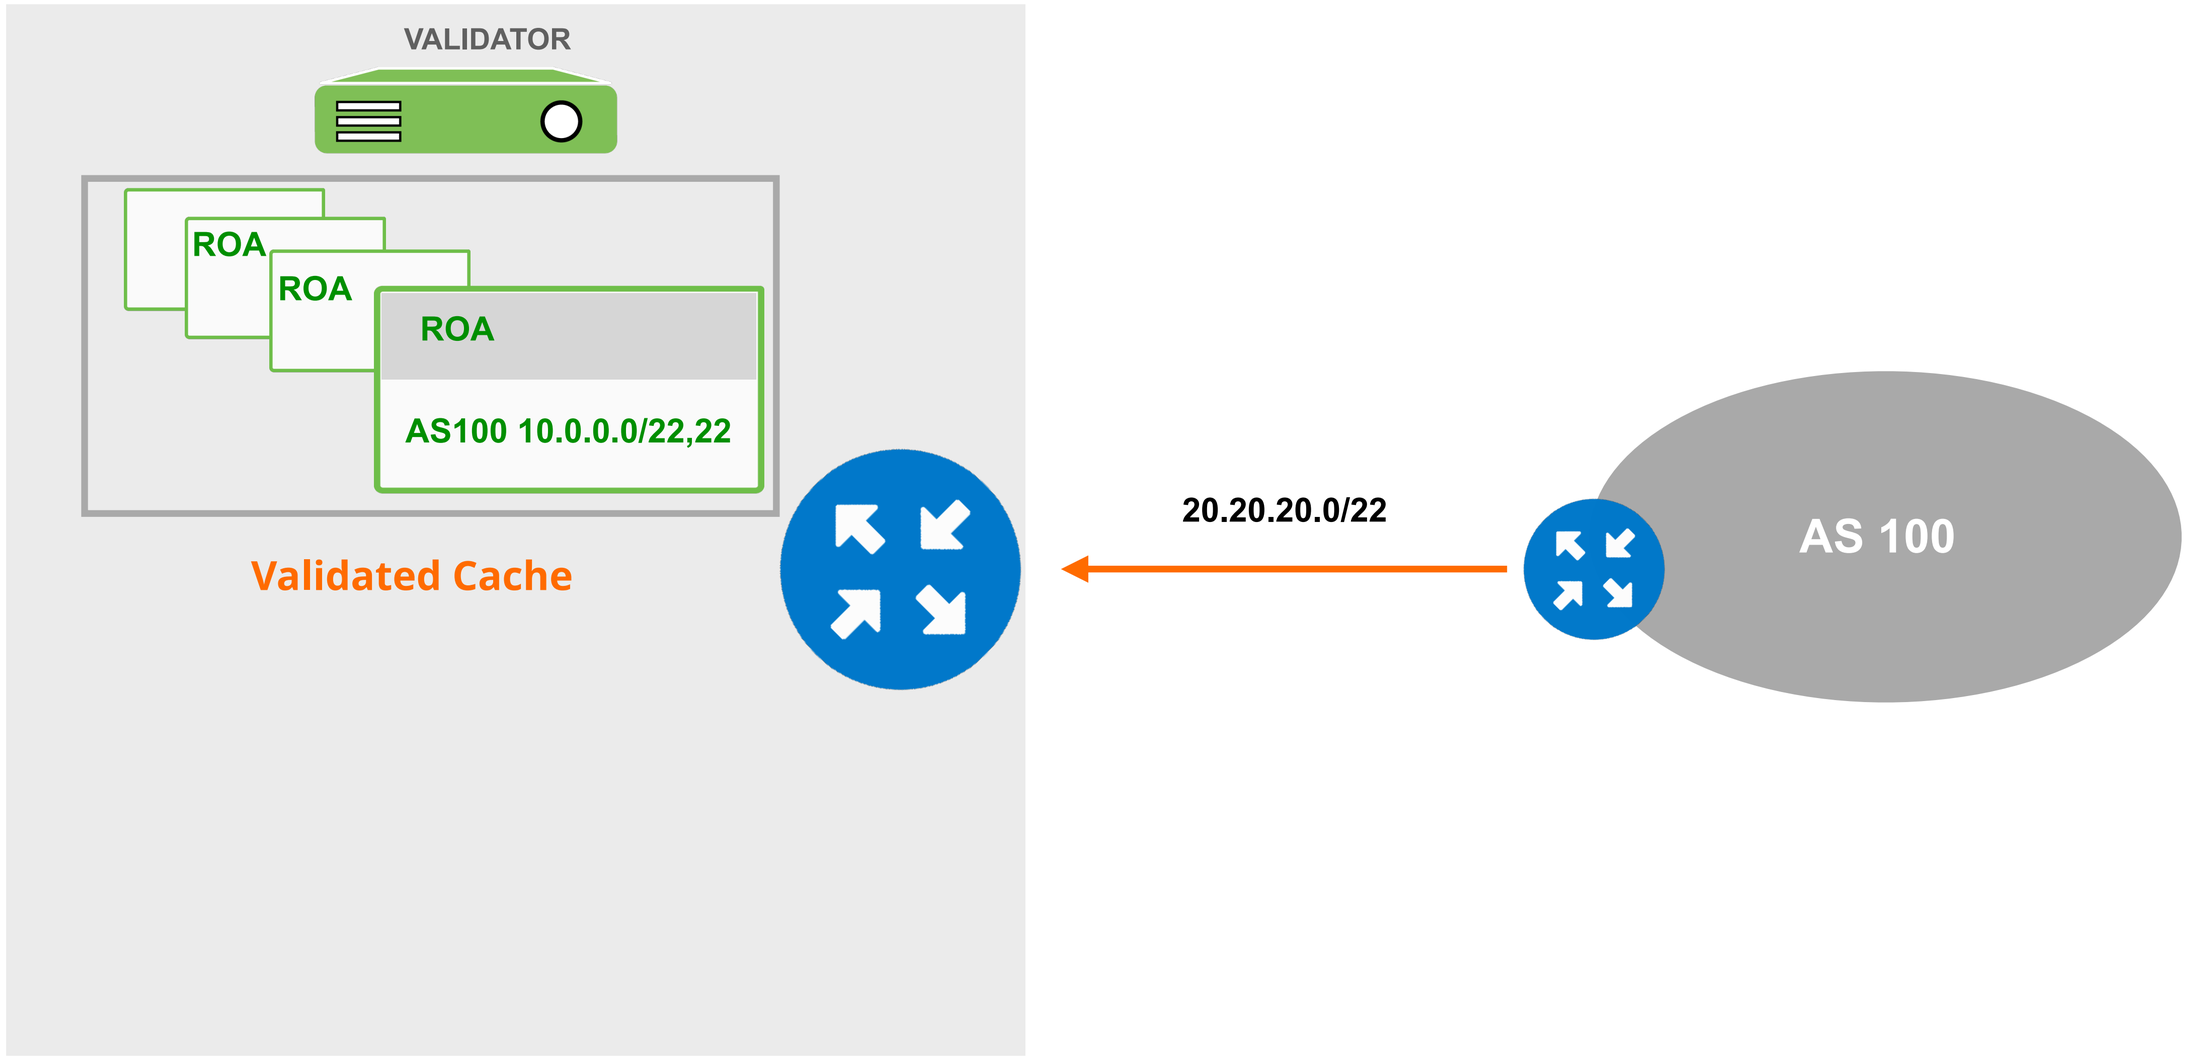

BGP Not-Found

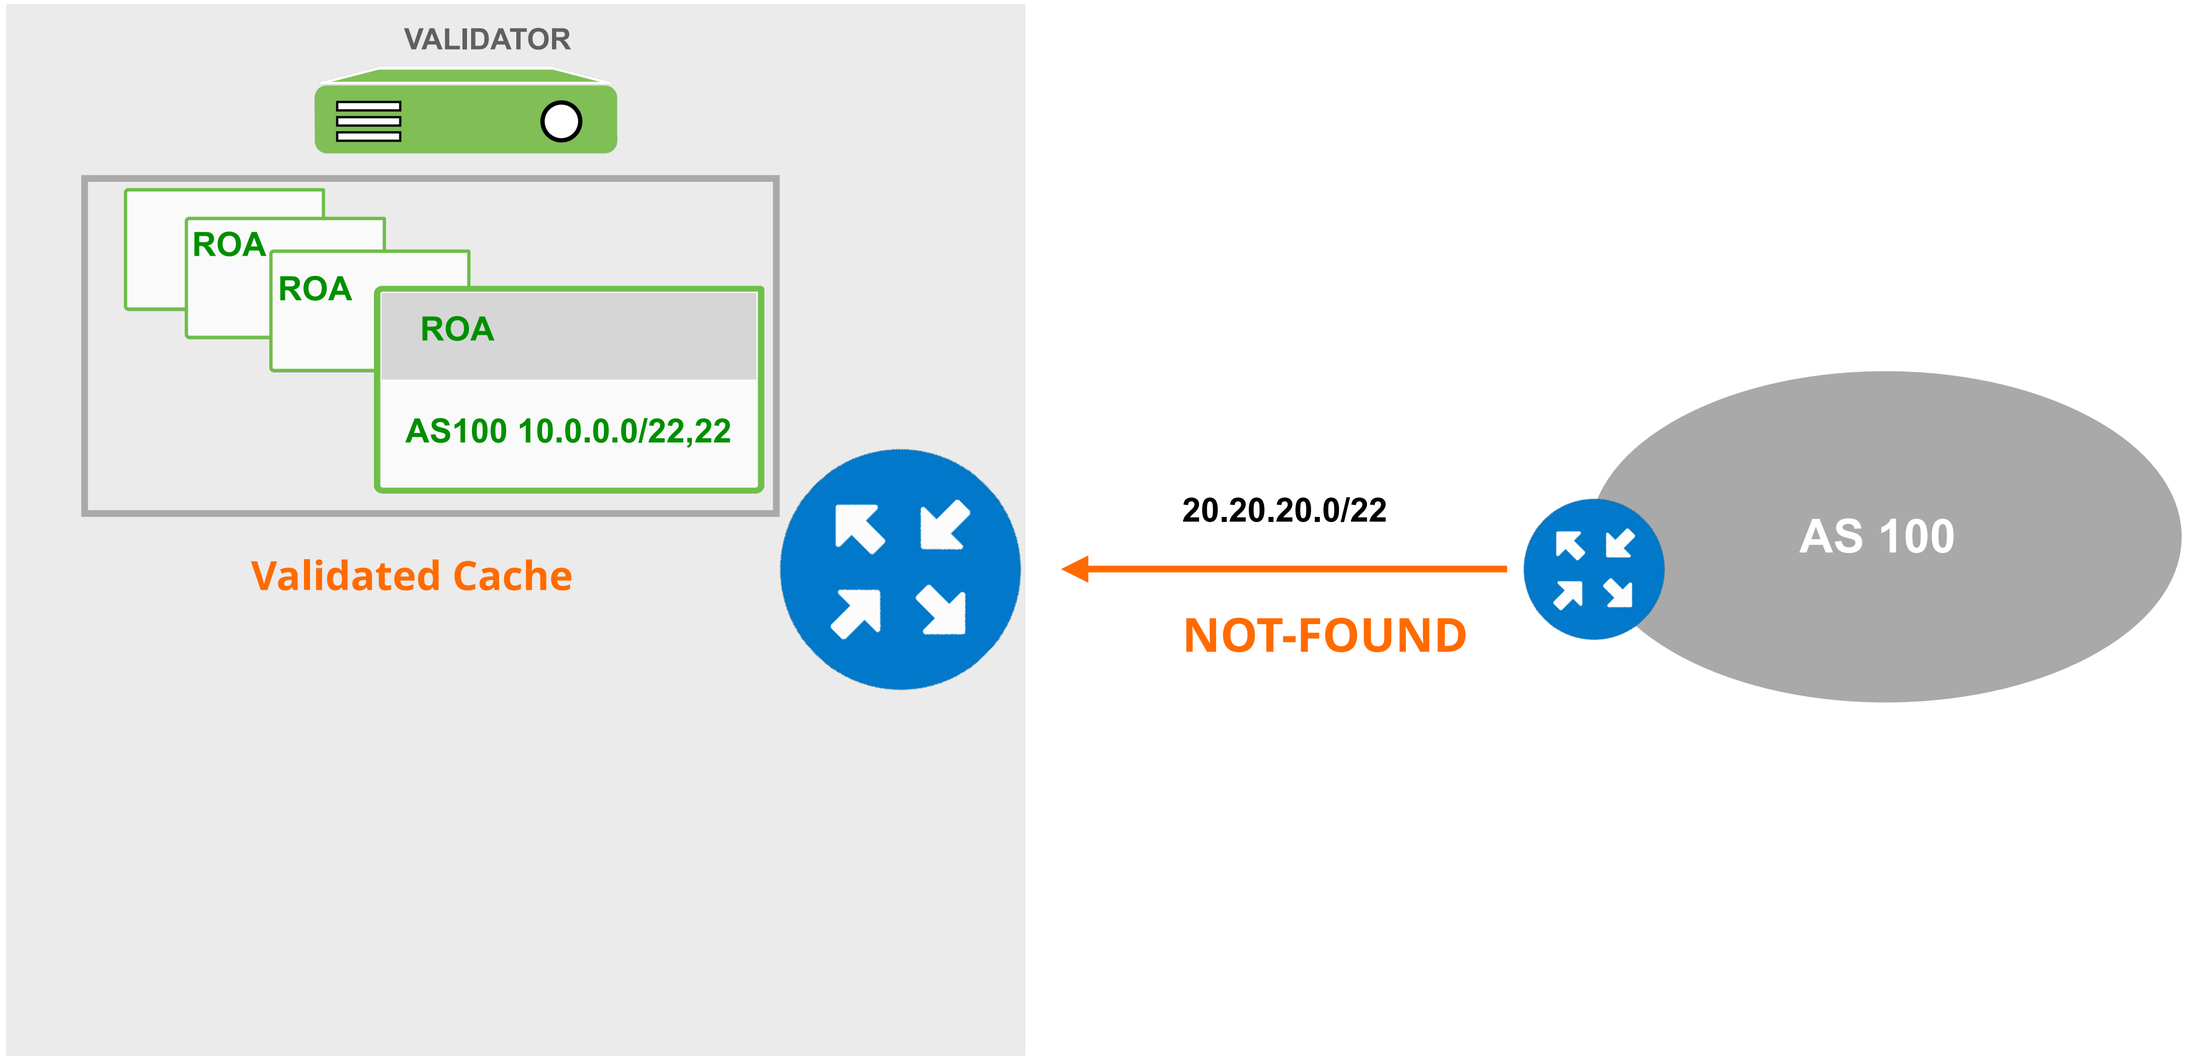

BGP Not-Found

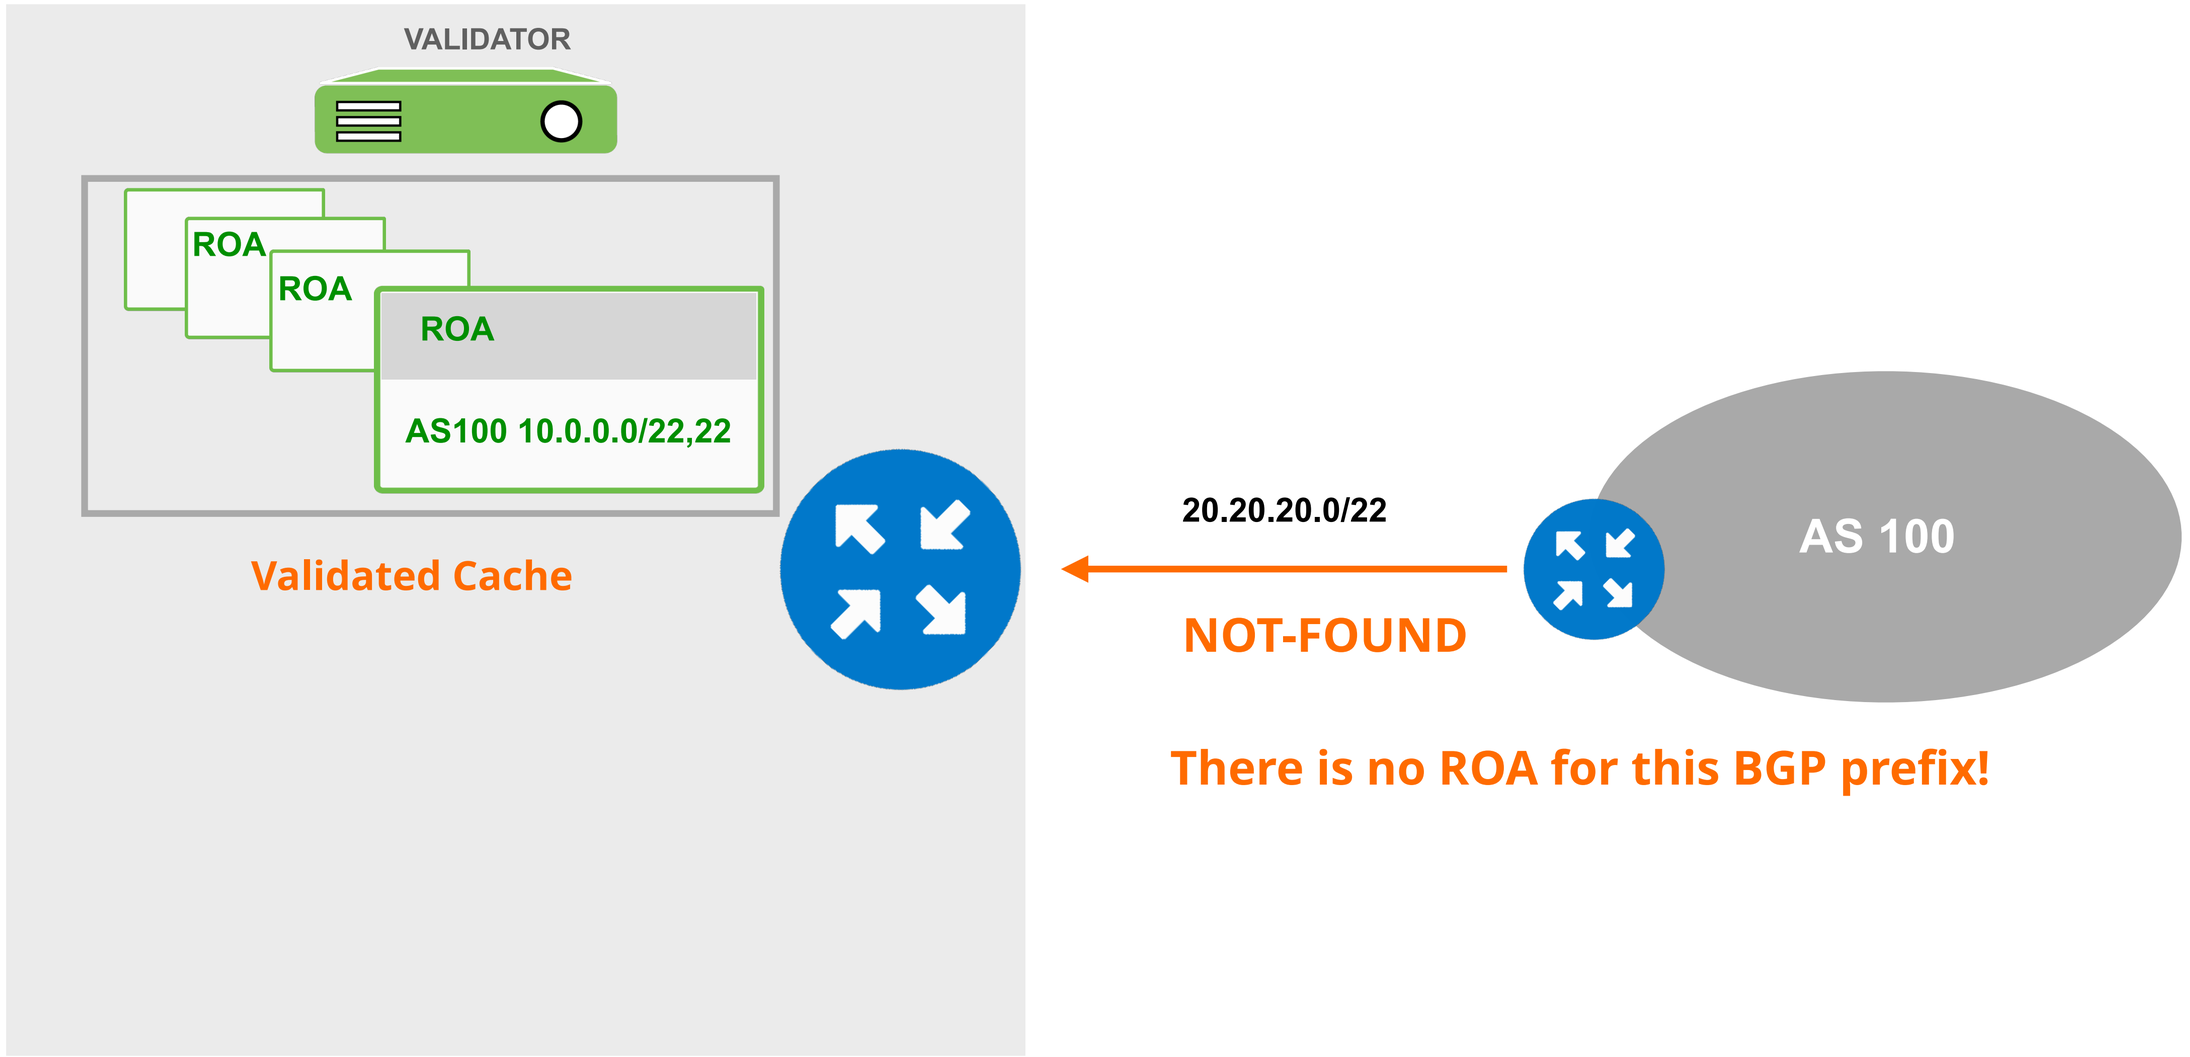

Whitelisting

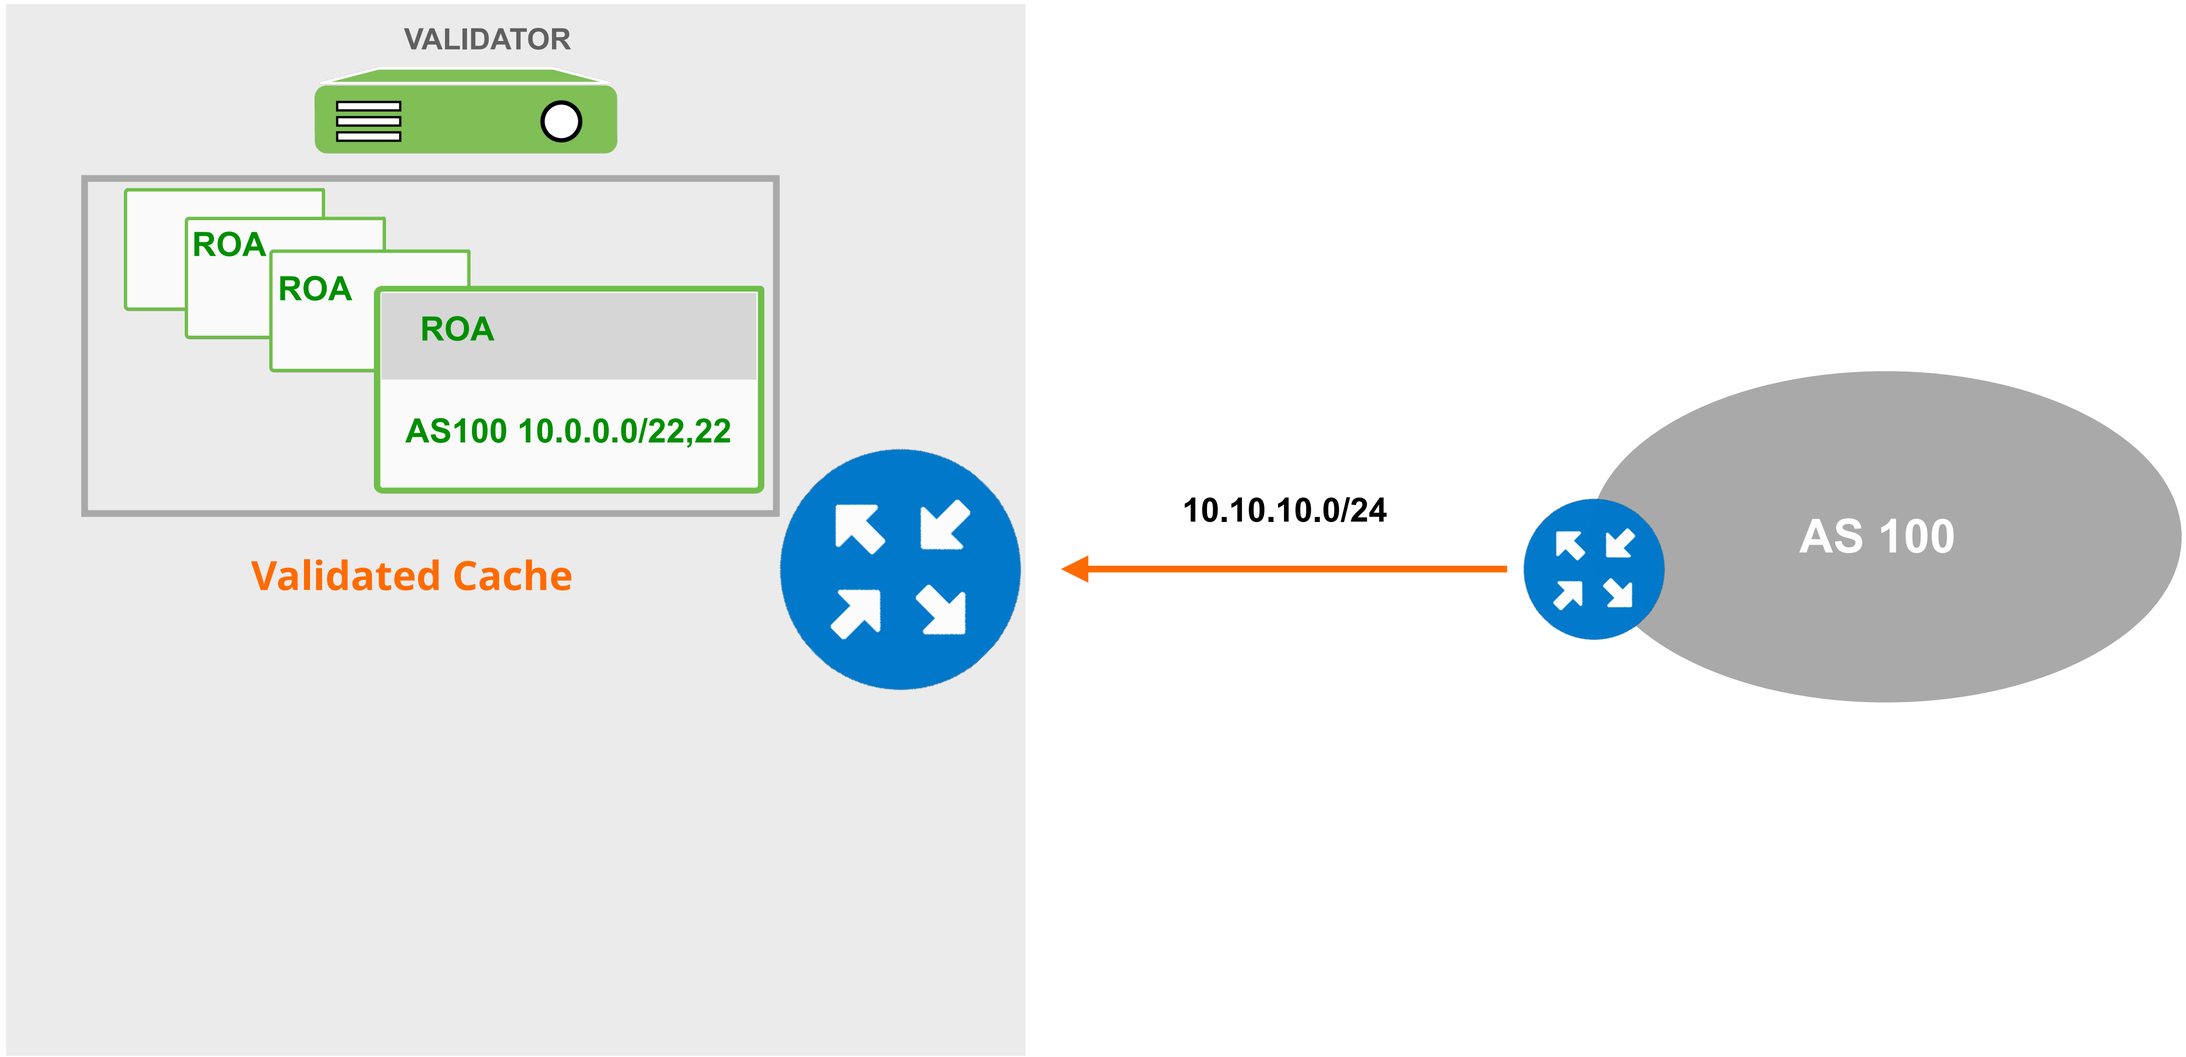

Whitelisting

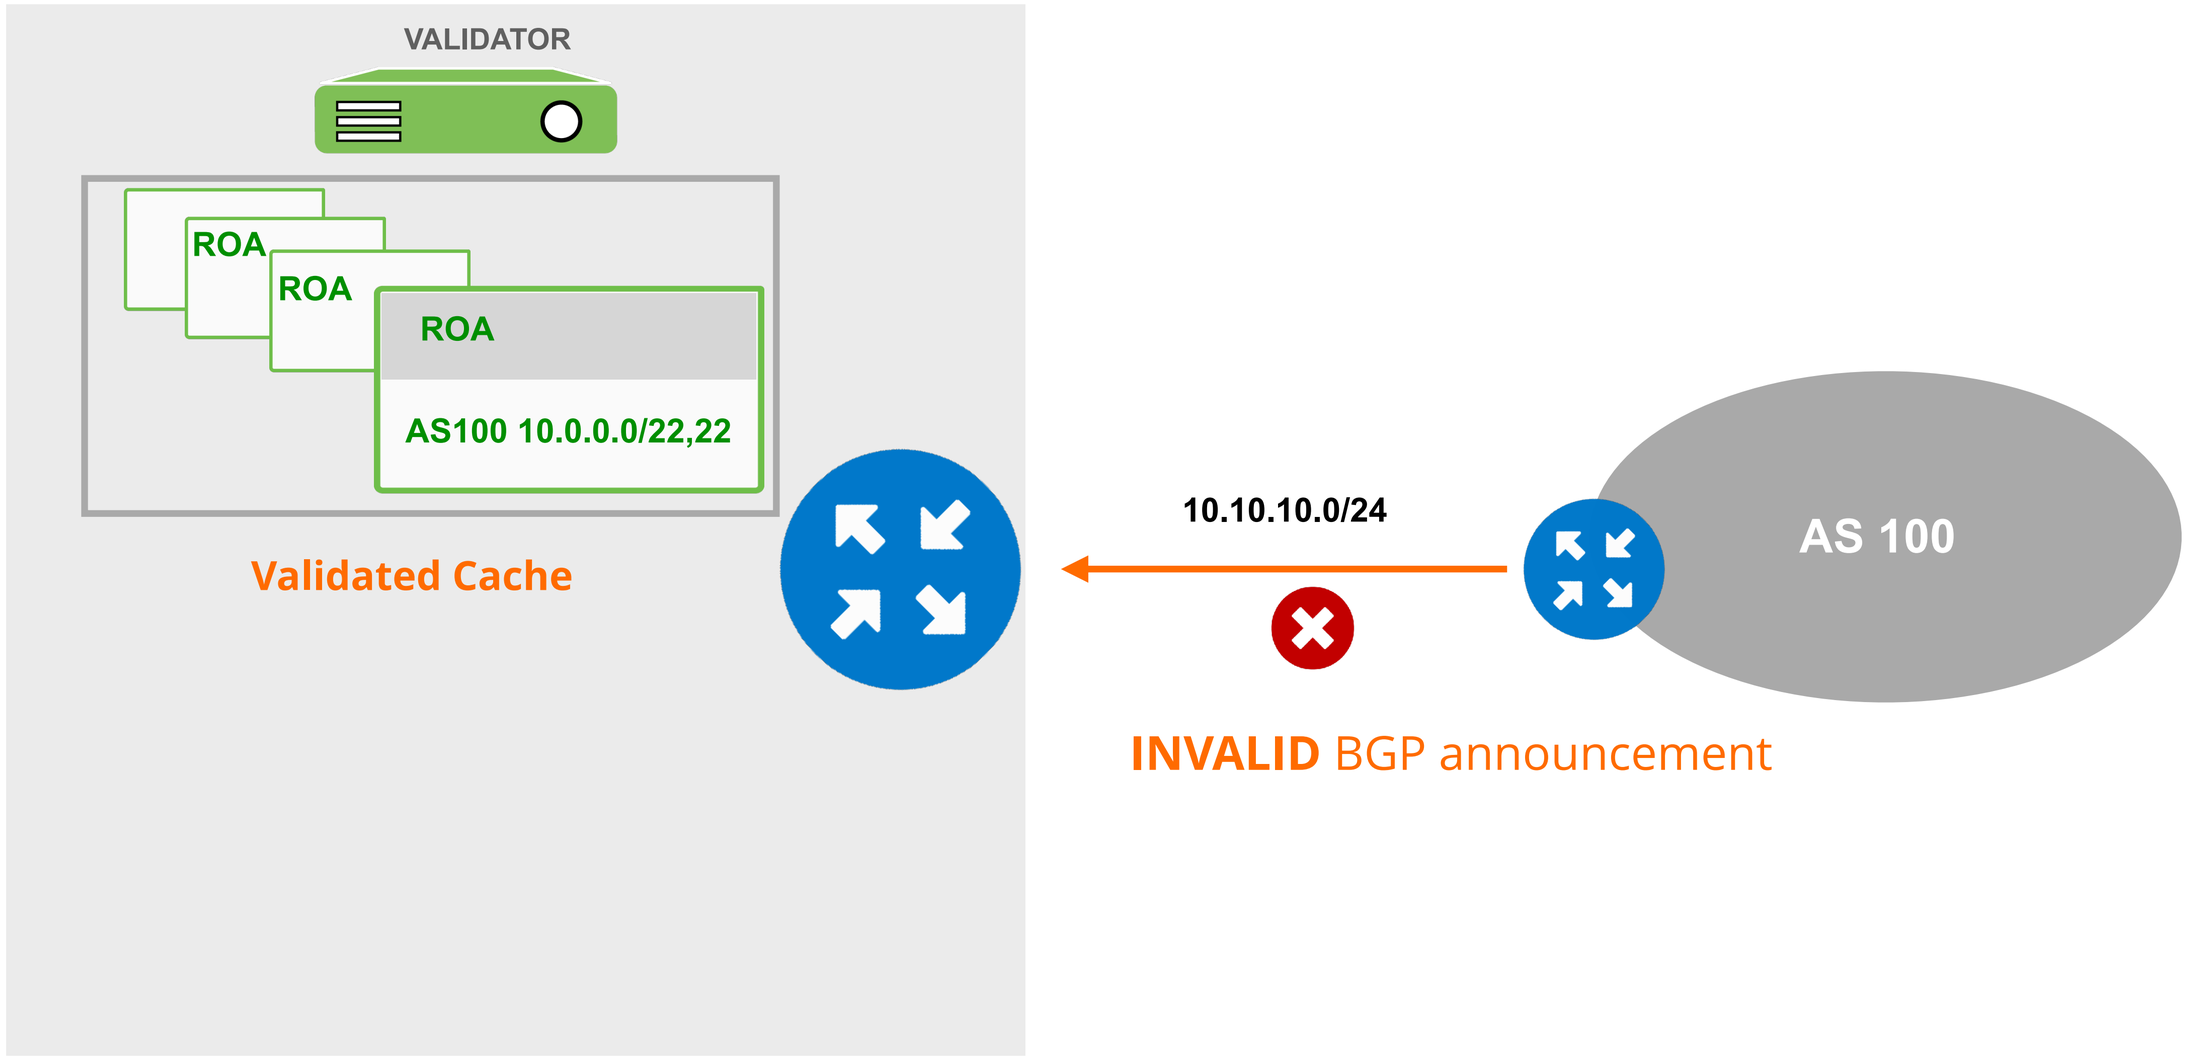

Whitelisting

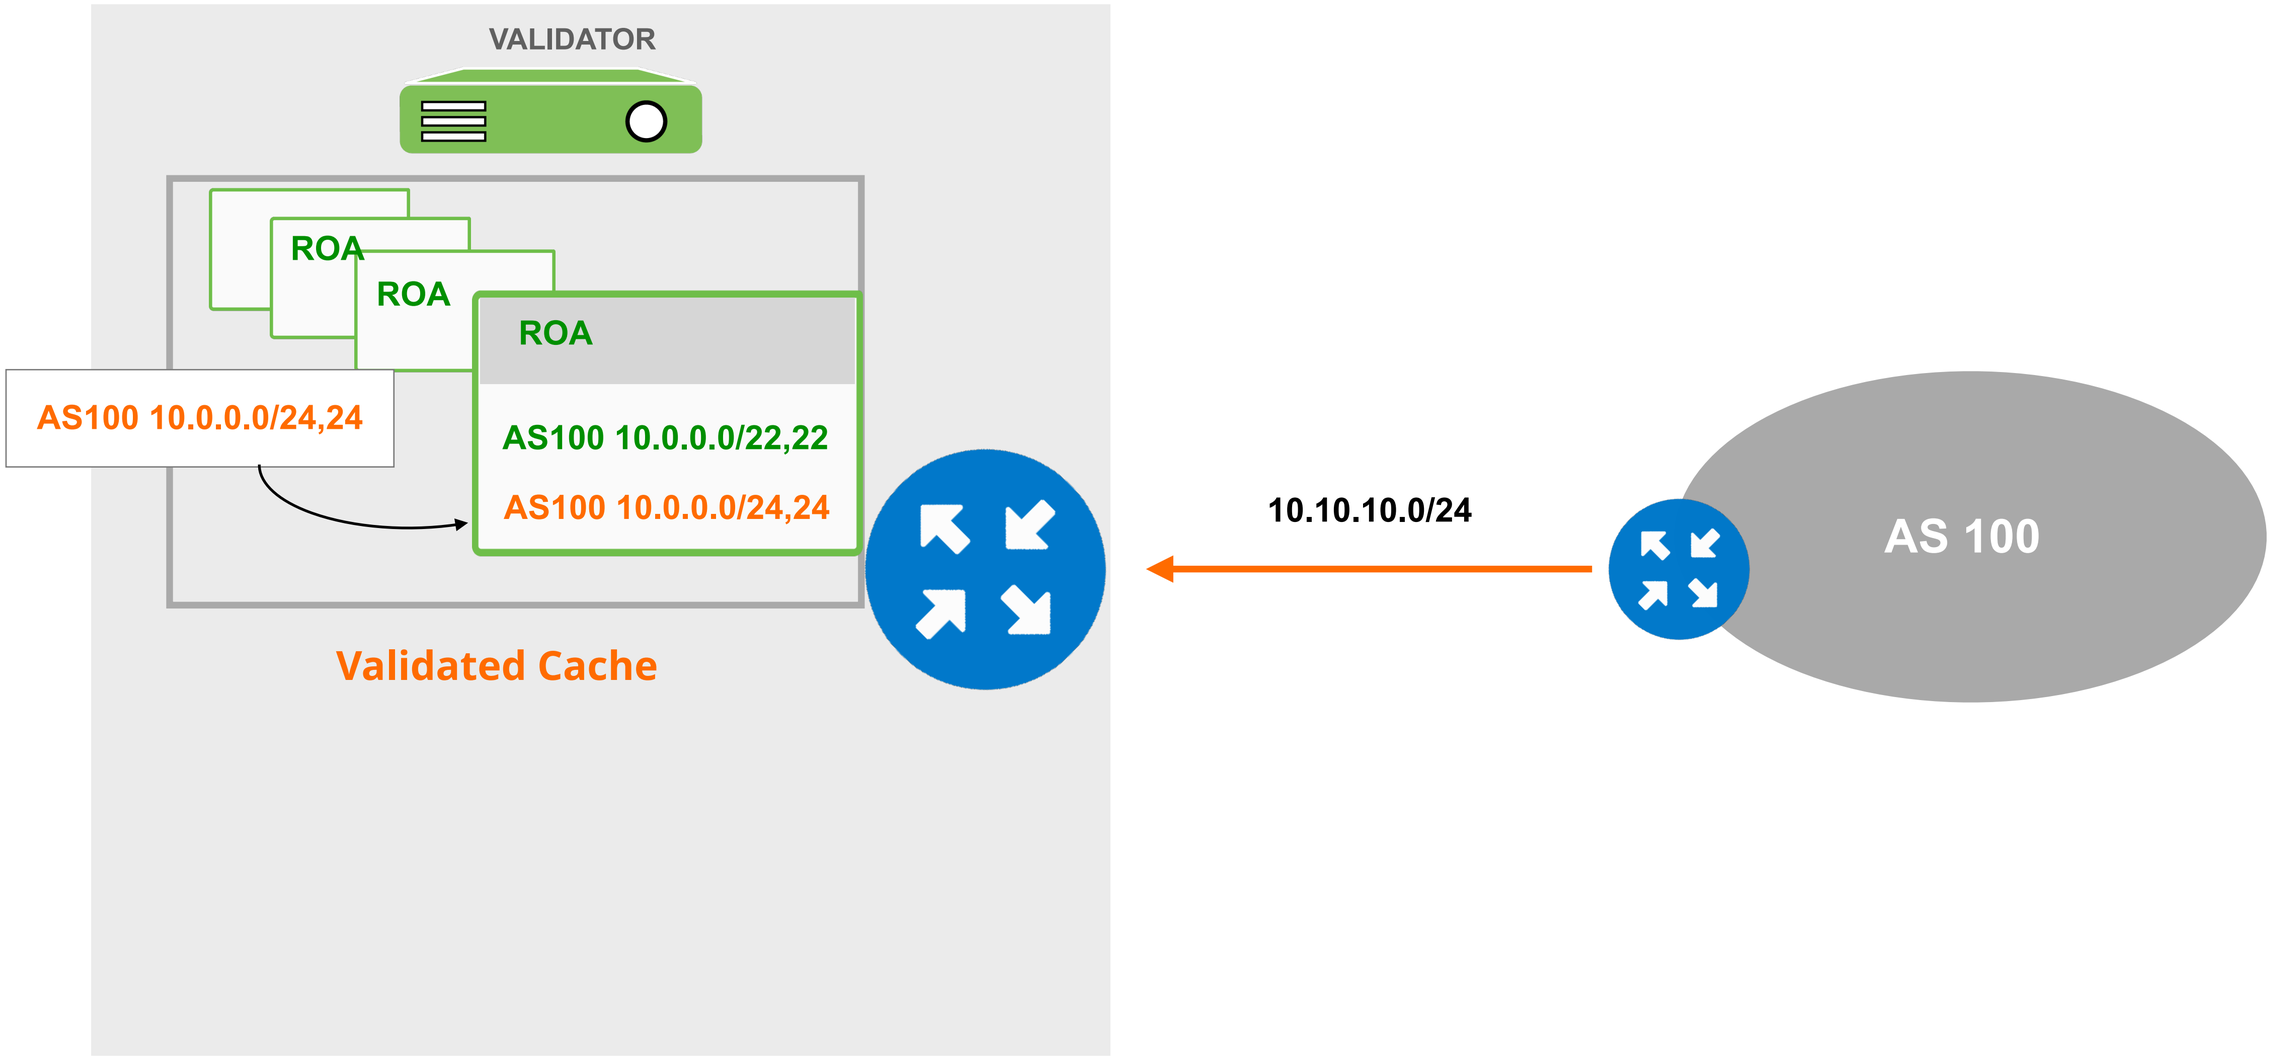

Whitelisting

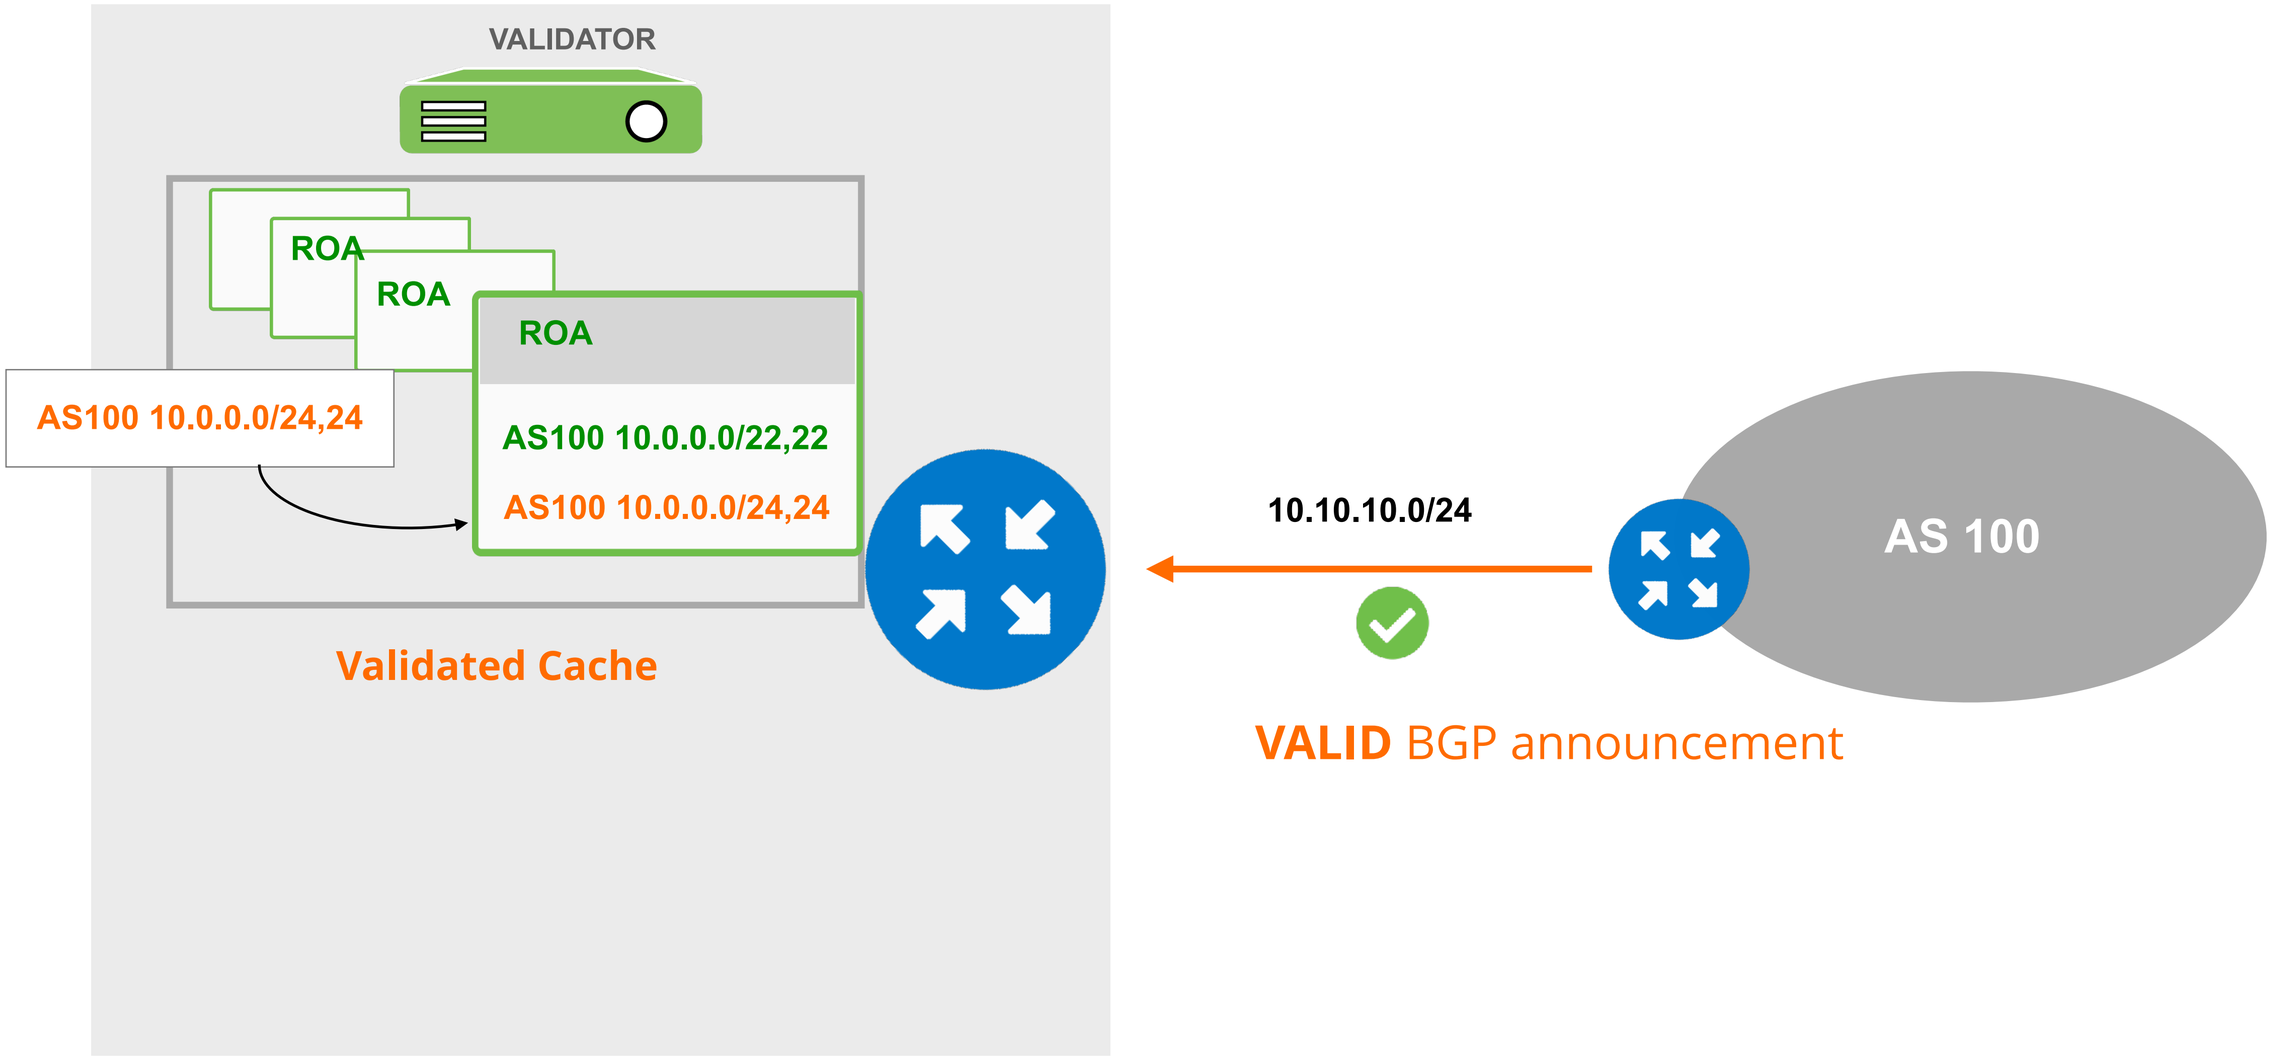

Take the poll!

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

What does this mean?

What to do with INVALIDs?

- For BGP origin validation to achieve its goal
  - Invalids should be dropped!
- As a first step,
  - you can set lower local preference or
  - tag the invalids with a BGP community
- After analysing the effect, you can start dropping INVALIDs

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)

Are you ready to implement RPKI?

Would you like to see how to implement RPKI?

Join our BGP Security webinar series!

BGP Security Webinars

BGP Security: IRR and Filtering

BGP Security: RPKI

RPKI Test Dashboard

https://localcert.ripe.net/#/rpki

- You can create test ROAs for your BGP announcements
- It doesn’t affect your network
- It’s just a test dashboard
- You need to sign in with your RIPE NCC Access Account

Questions