RIPE NCC
RIPE NETWORK COORDINATION CENTRE

BGP & Routing Security

Tutorial-RIPE85

RIPE NCC Learning & Development

Cigdem GUR SENOL
e-mail: cgur@ripe.net

Agenda

Vulnerabilities of BGP

How to Secure Internet Routing

Implementing BGP Filters

Routing Security with RPKI

- What is RPKI?
- Registering in the RPKI system (ROAs)
- RPKI Validators
- Validating BGP Announcements

Demo: BGP Origin Validation with RPKI

Vulnerabilities of BGP

BGP has some challenges …

- BGP has some challenges from a routing security perspective:
  - It is only based on trust, no built-in security
  - No verification of the correctness of prefixes or AS paths
- These challenges are discussed in RFC#4272, “BGP Security Vulnerabilities Analysis”

RFC#4272,“BGP Security Vulnerabilities Analysis”

Vulnerabilities of BGP

- Based on RFC, BGP has three fundamental vulnerabilities:

1. No internal mechanism to protect the integrity and source authenticity of BGP messages

2. No mechanism specified to validate the authority of an AS to announce NLRI

3. No mechanism to verify the authenticity of the attributes of a BGP update message

- These vulnerabilities can be exploited either **maliciously** or **accidentally**

Due to these vulnerabilities …

- Many BGP incidents happen every year!
- Attacks can be conducted by exploiting TCP or BGP messages
- Any AS can announce any prefix
  - BGP prefix hijacks due to malicious activity / mis-origination
- Any AS can prepend any ASN to the AS path
  - Path hijacks, MITM

Sometimes, just human errors …

- Typo errors
  - Also known as “fat fingers”
  - May cause mis-origination
- Configuration errors
  - Faulty BGP filter configuration
  - AS path prepending mistake
- Simple mistakes may cause big problems!
  - BGP hijacks or route leaks

But, sometimes they are malicious!

- Attackers can abuse BGP by using its vulnerabilities
- Potential attacks on BGP;
  - TCP/IP Protocol attacks (Spoofing, Session hijacking)
  - Protocol manipulation attacks (MED modification, exploit RFD/MRAI timer)
  - Denial of service attacks via resource exhaustion
  - BGP route manipulation attacks (BGP route hijack, BGP path hijack)

BGP Route Manipulation Attacks

- Attackers can
  - Inject bogus routing information into BGP tables
  - Reroute packets based on its intensions
  - Prevent traffic from reaching to its intended destination
- The goal is to blackhole the traffic, eavesdropping or traffic analysis
- Route manipulation attacks can be classified as
  - BGP Origin Hijacks
  - BGP Path Hijacks

BGP Origin Hijack

- The hijacking AS
  - Abuses mutual trust between ASes
  - Originates a prefix **that it is not authorised to originate**
- Traffic is diverted to the hijacker's network
- Difficult to say whether it was an accident or an attack!
- Hijacker may announce
  - the exact same prefix
  - a more specific prefix

Announcing the same prefix

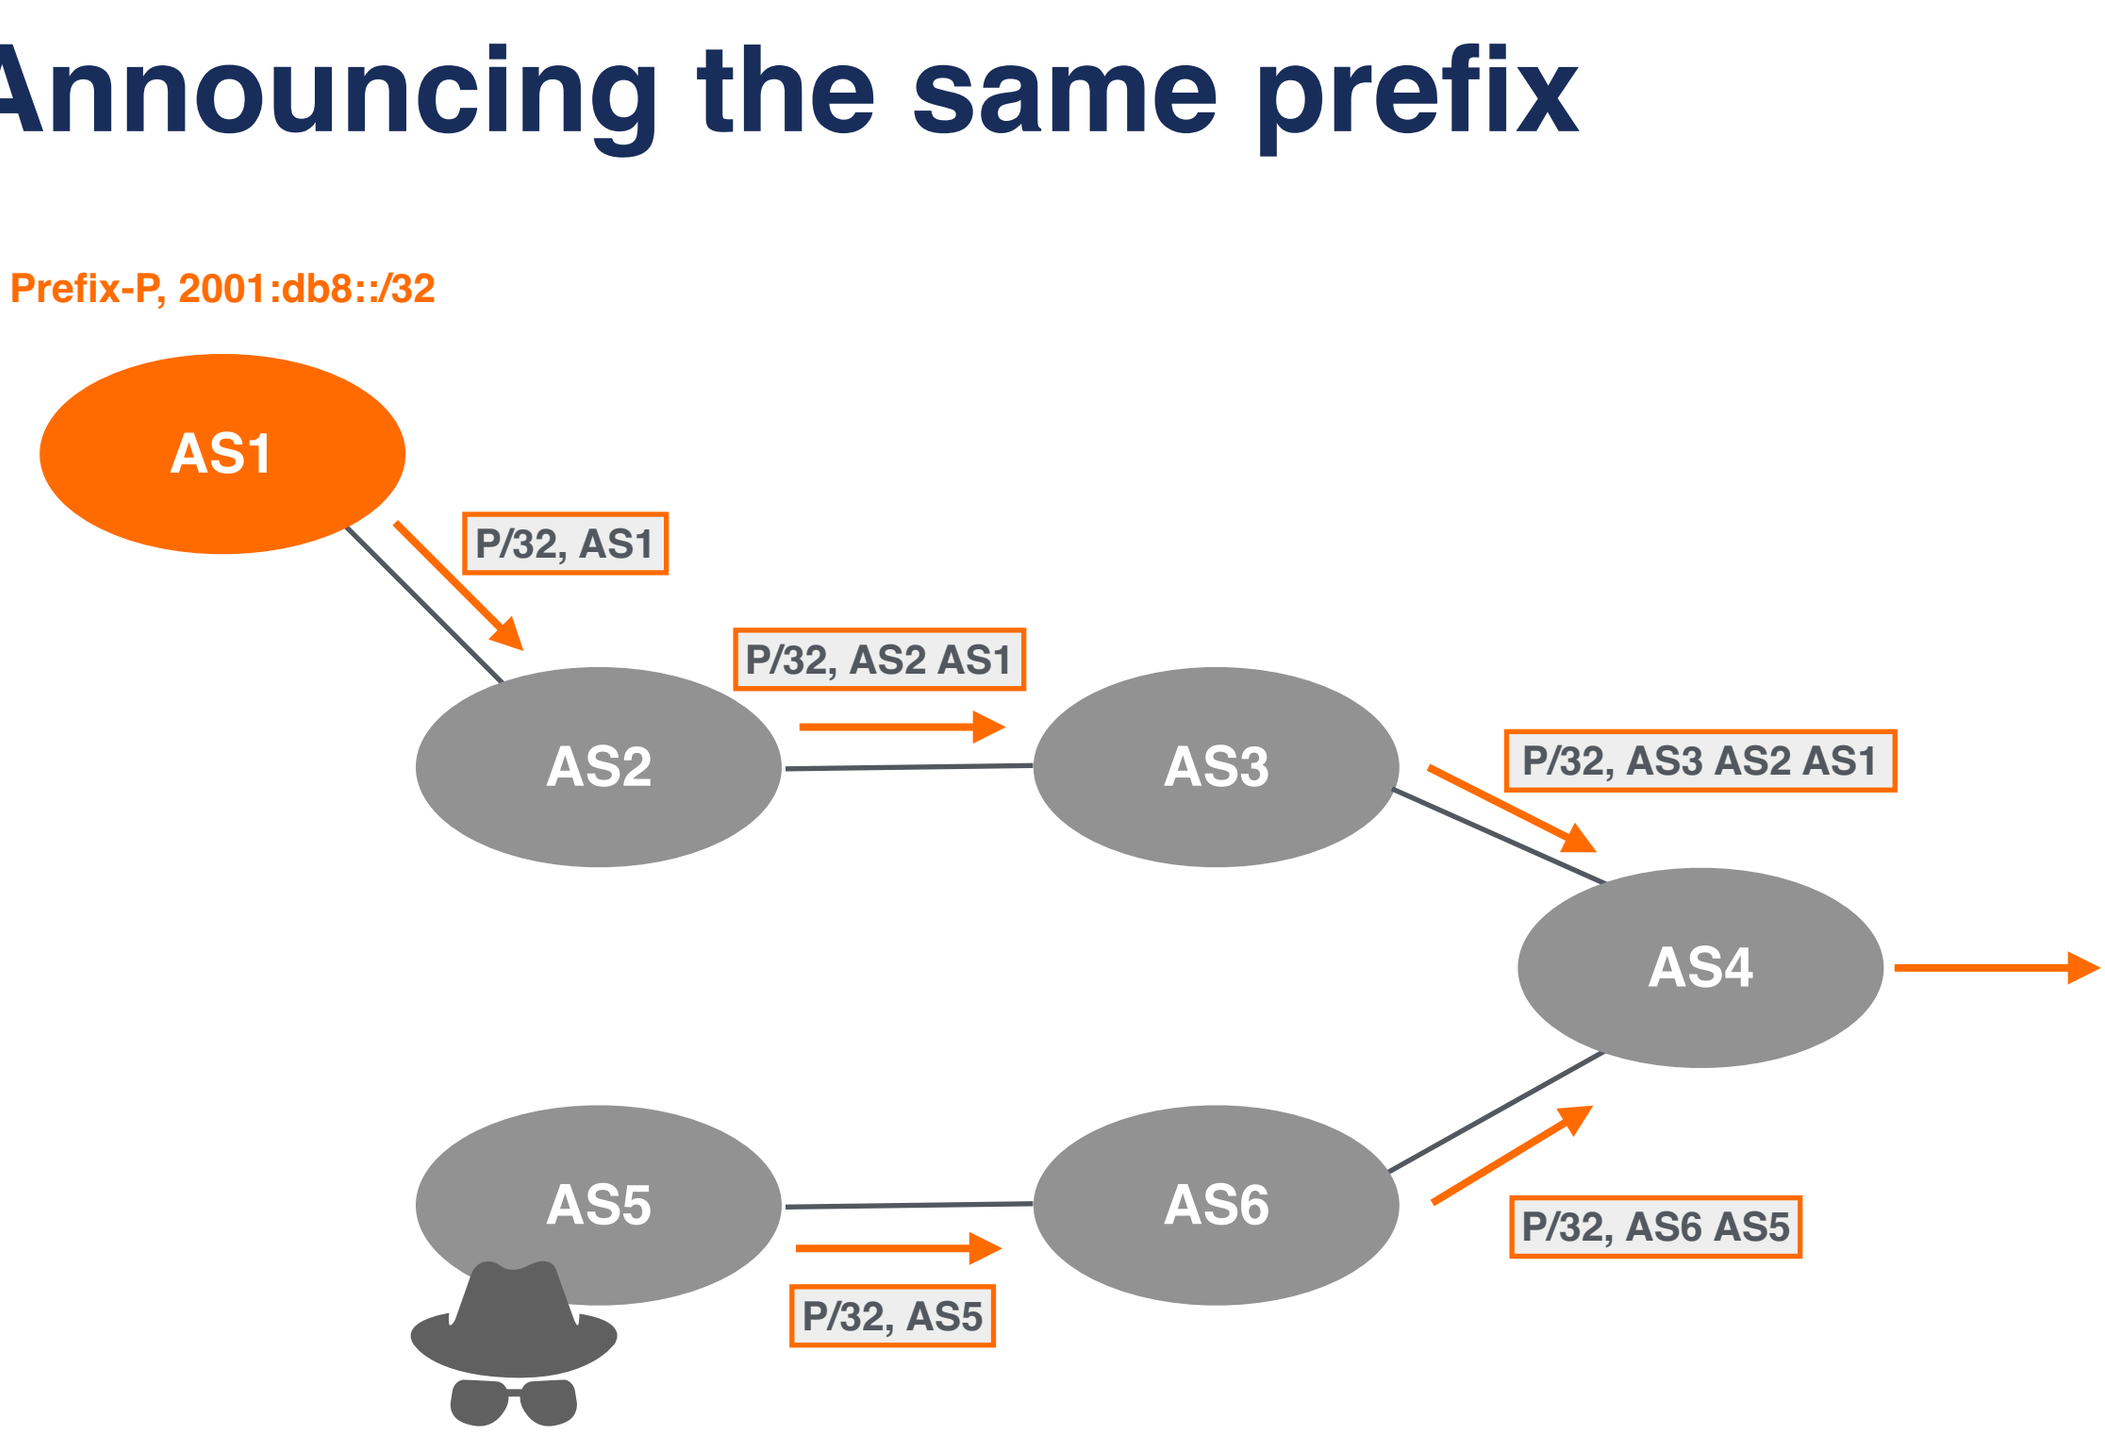

This is a **local hijack!**
Only some networks are affected based on BGP path selection process.

Announcing a more specific prefix

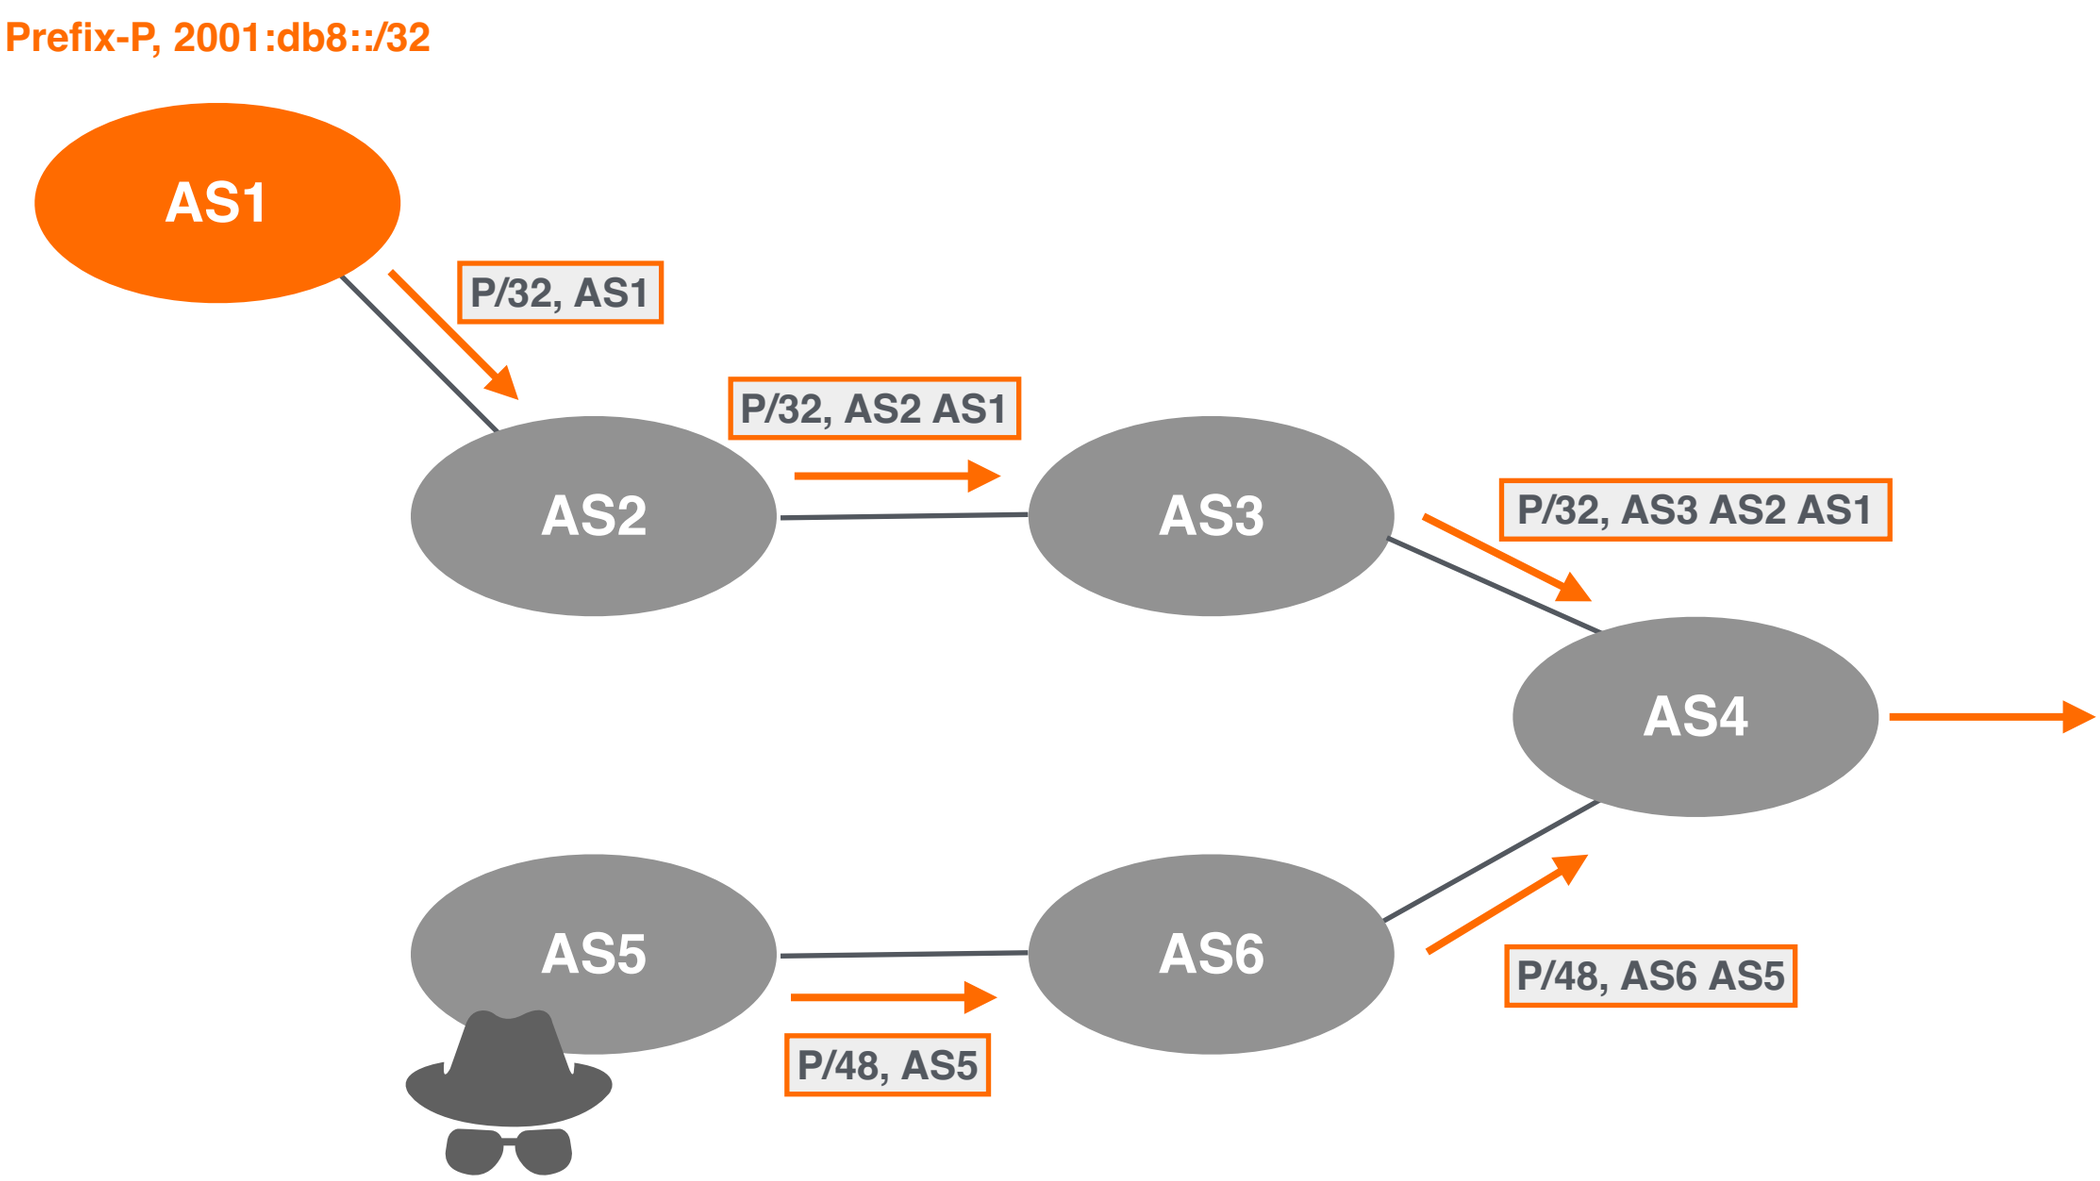

This is a **global hijack!**
All traffic for more specific will be forwarded to the attacker's network network.

BGP Path Hijack

- The attacker can
  - **send a fake path** with correct or different origin and its ASN in the middle
  - or **modify an existing path**
- It can modify the path of the BGP updates by
  - inserting false AS numbers into the AS path
  - or removing some ASes from AS path

Fake path with correct origin

- The origin of the path does not change!
- The attacker creates a forged AS link between two ASes
- The attacker reroutes the traffic to itself

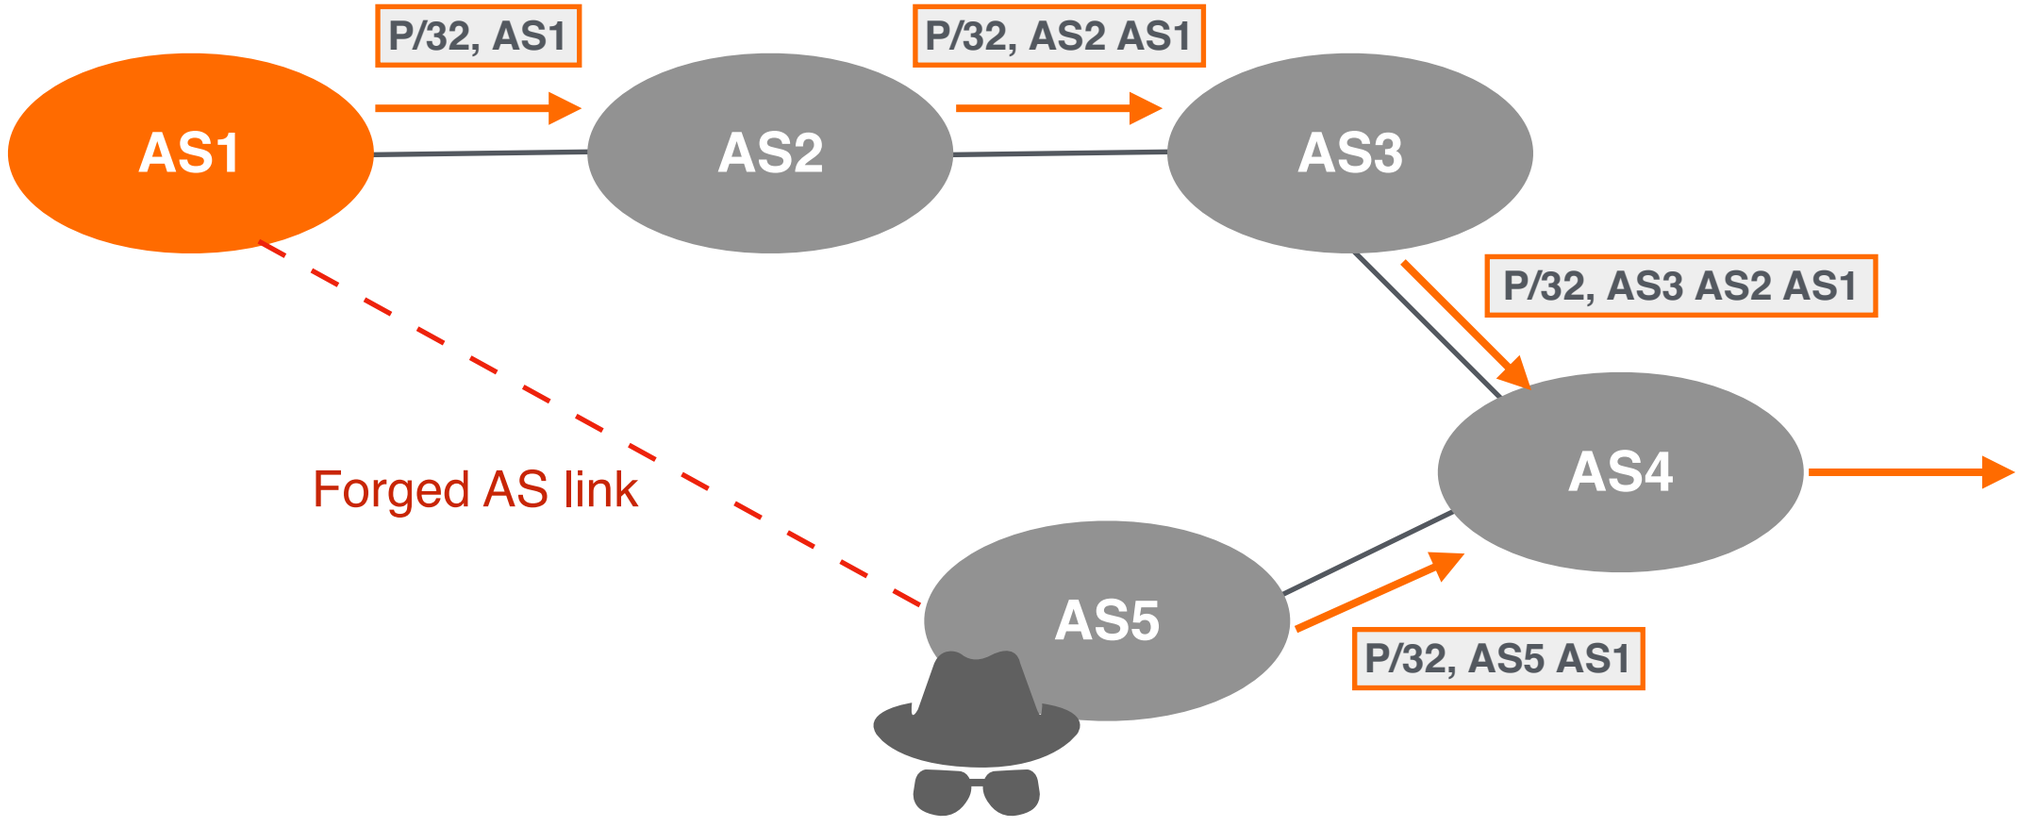

Fake path with different origin

- Looks like an origin hijack
  - But in reality, the origin AS is not the cause of the problem!
- Again, the attacker reroutes the traffic to itself

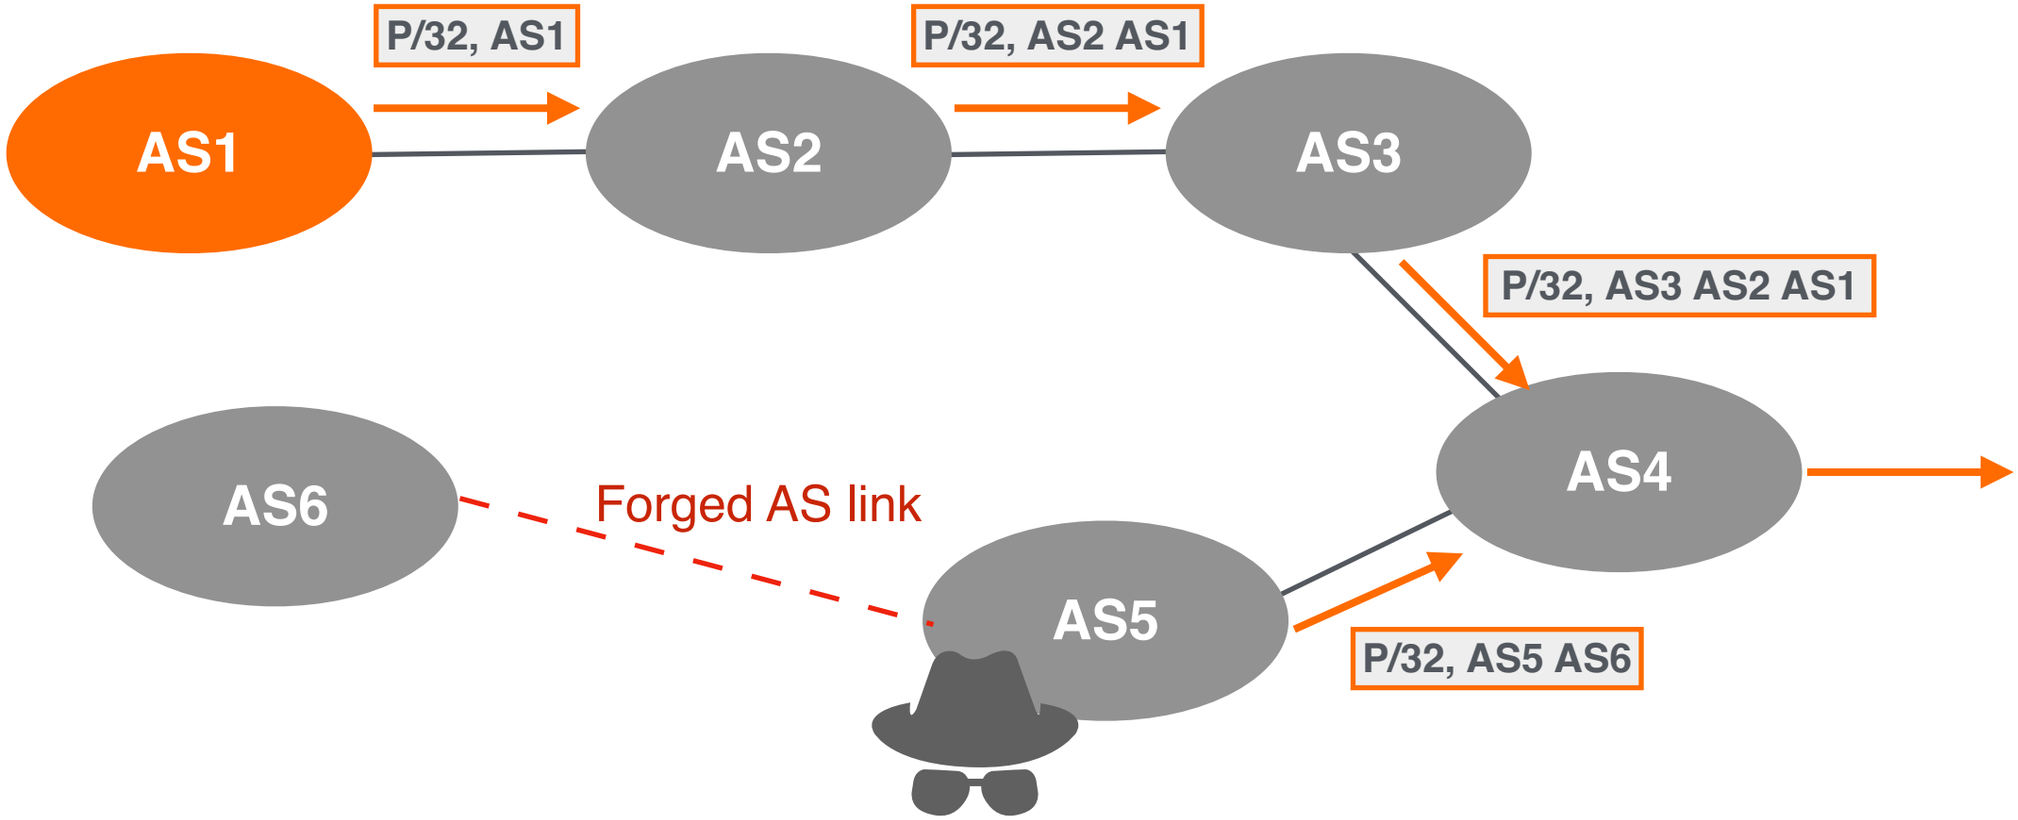

Modifying an existing path

- Neighbours of the attacker receive a false path
- The attacker can do either of these two things:
  - Analyse the traffic and then route to AS1
  - Drop the traffic to AS1

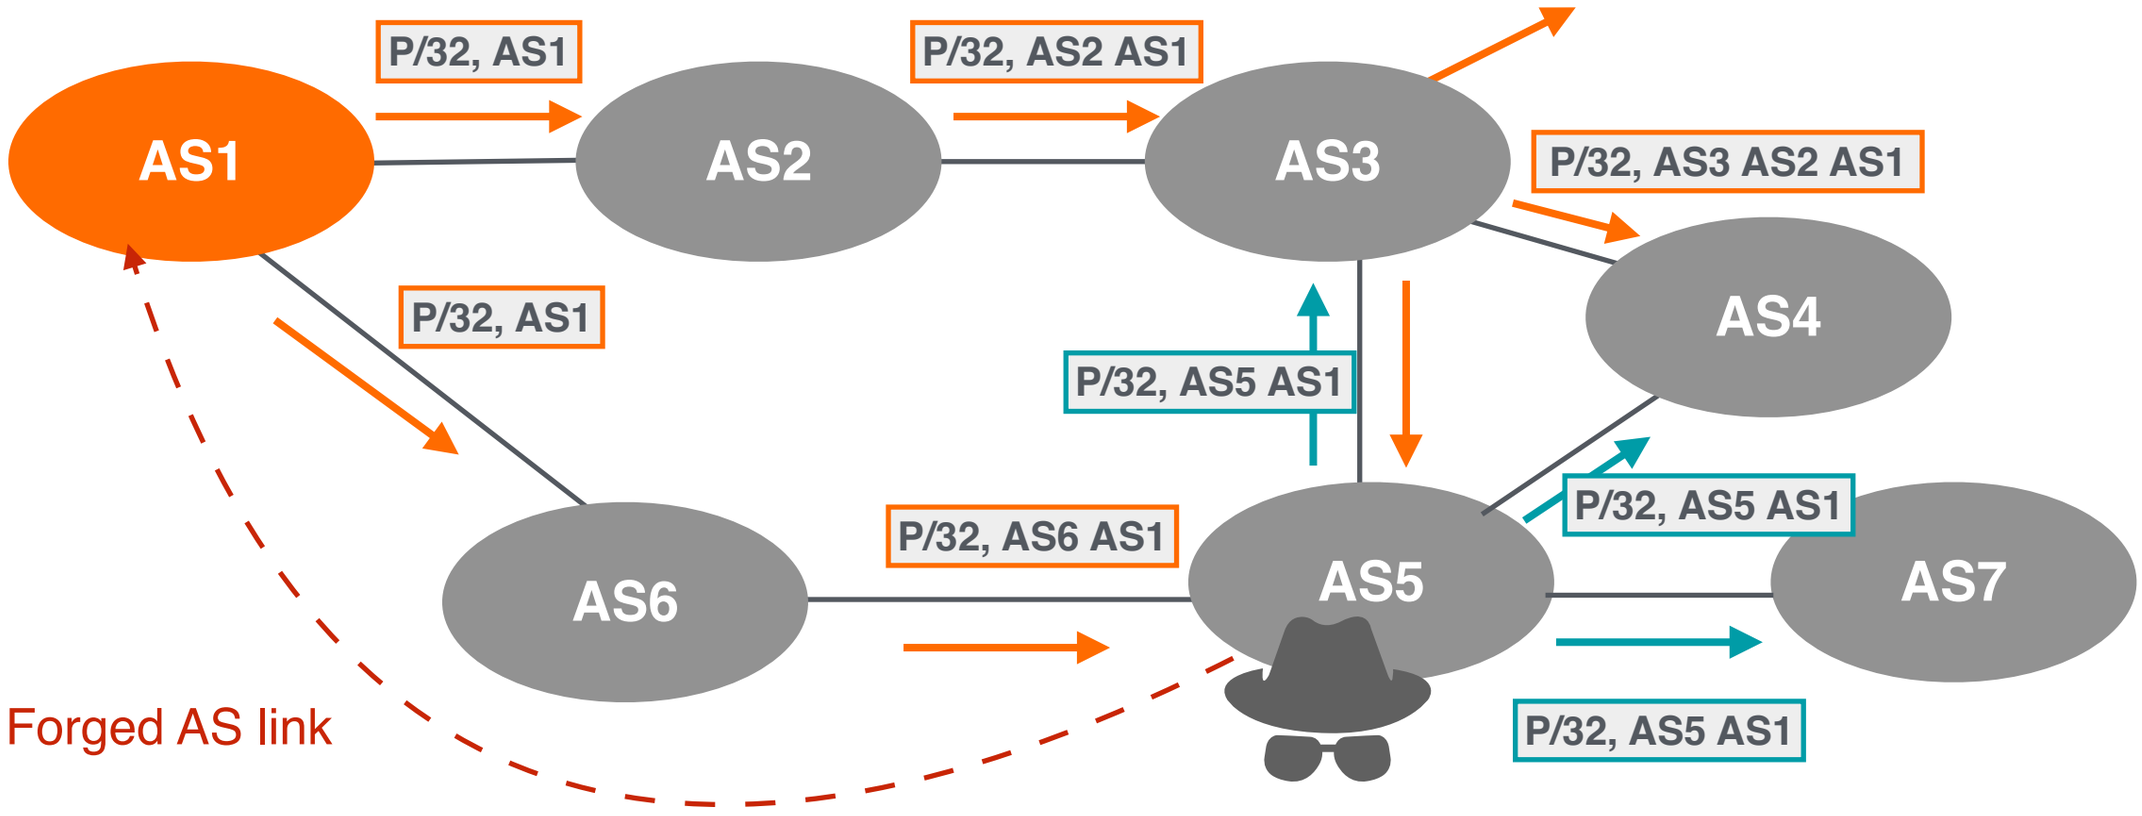

Man-in-the-middle Attack (MITM)

- “Kapela-Pilosov” attack, which was introduced in Defcon , 2008
- Origin is correct and AS path is not altered!
- Goal of this attack is to intercept and log/alter the traffic

Man-in-the-middle Attack (MITM)

1 The attacker identifies a usable path to the victim AS

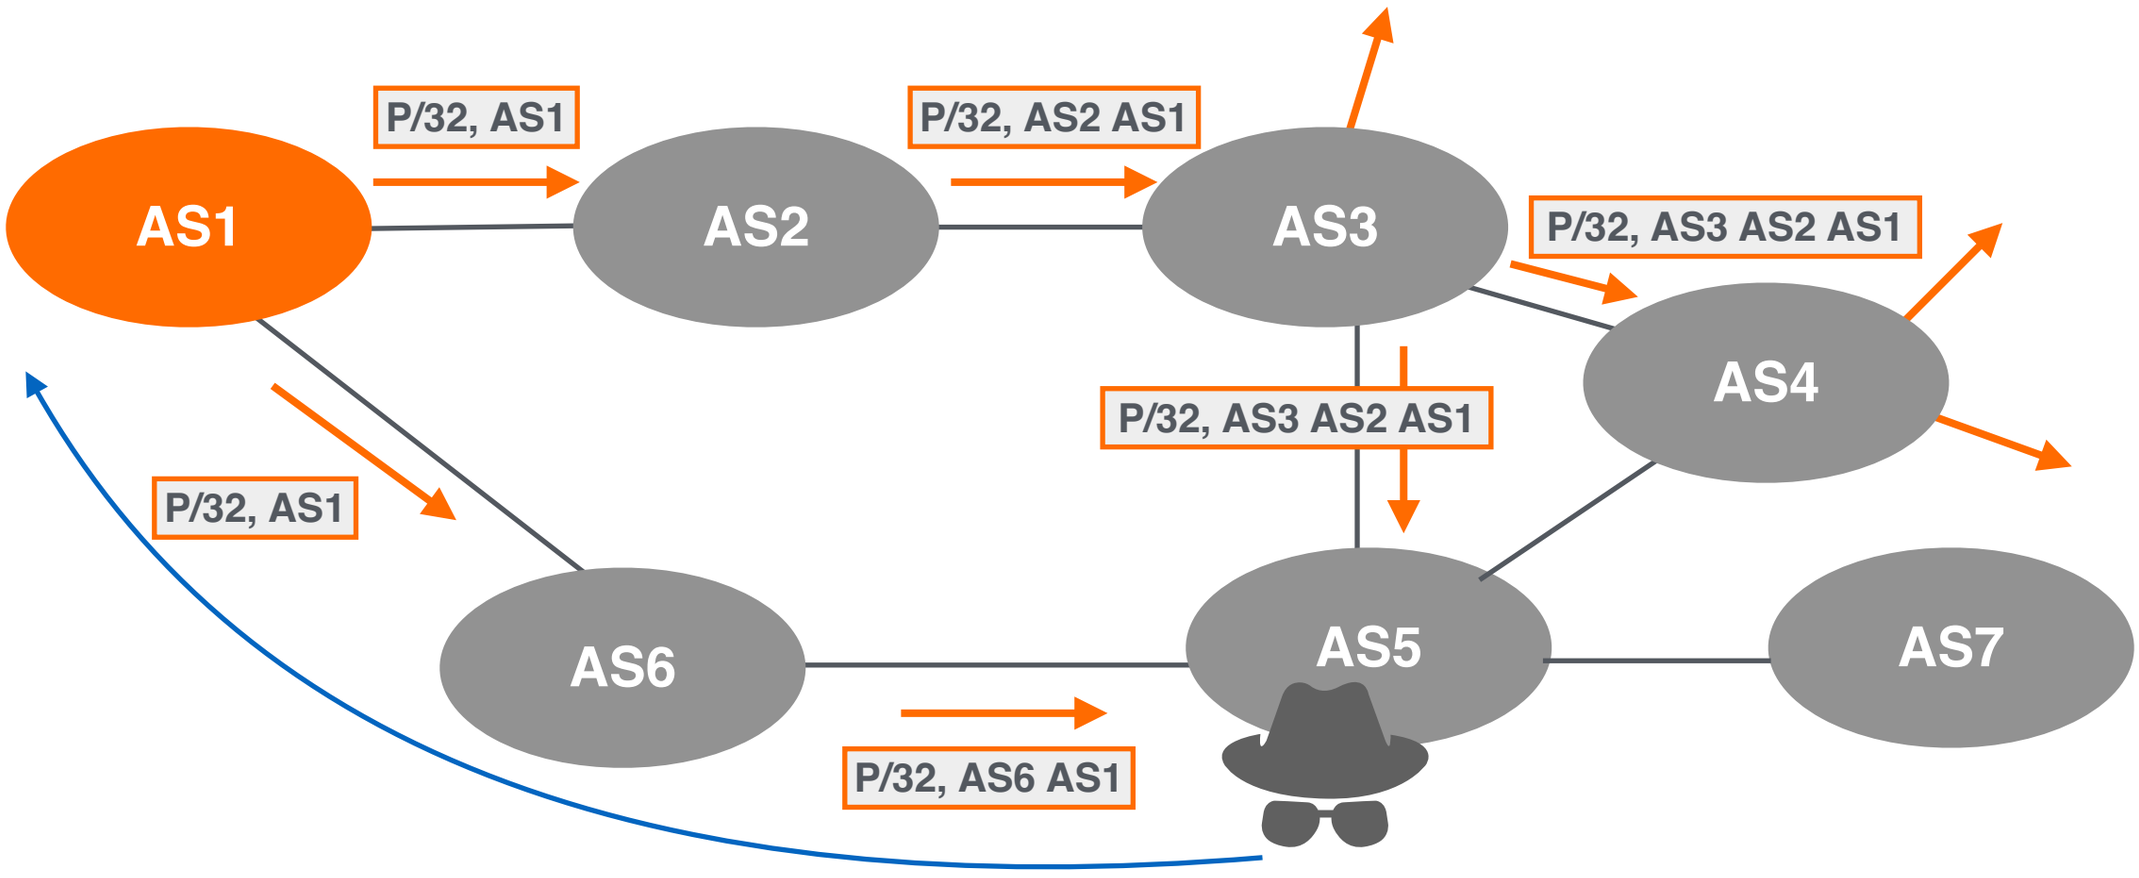

Man-in-the-middle Attack (MITM)

2 Replaces the prefix in a received update with a more-specific prefix

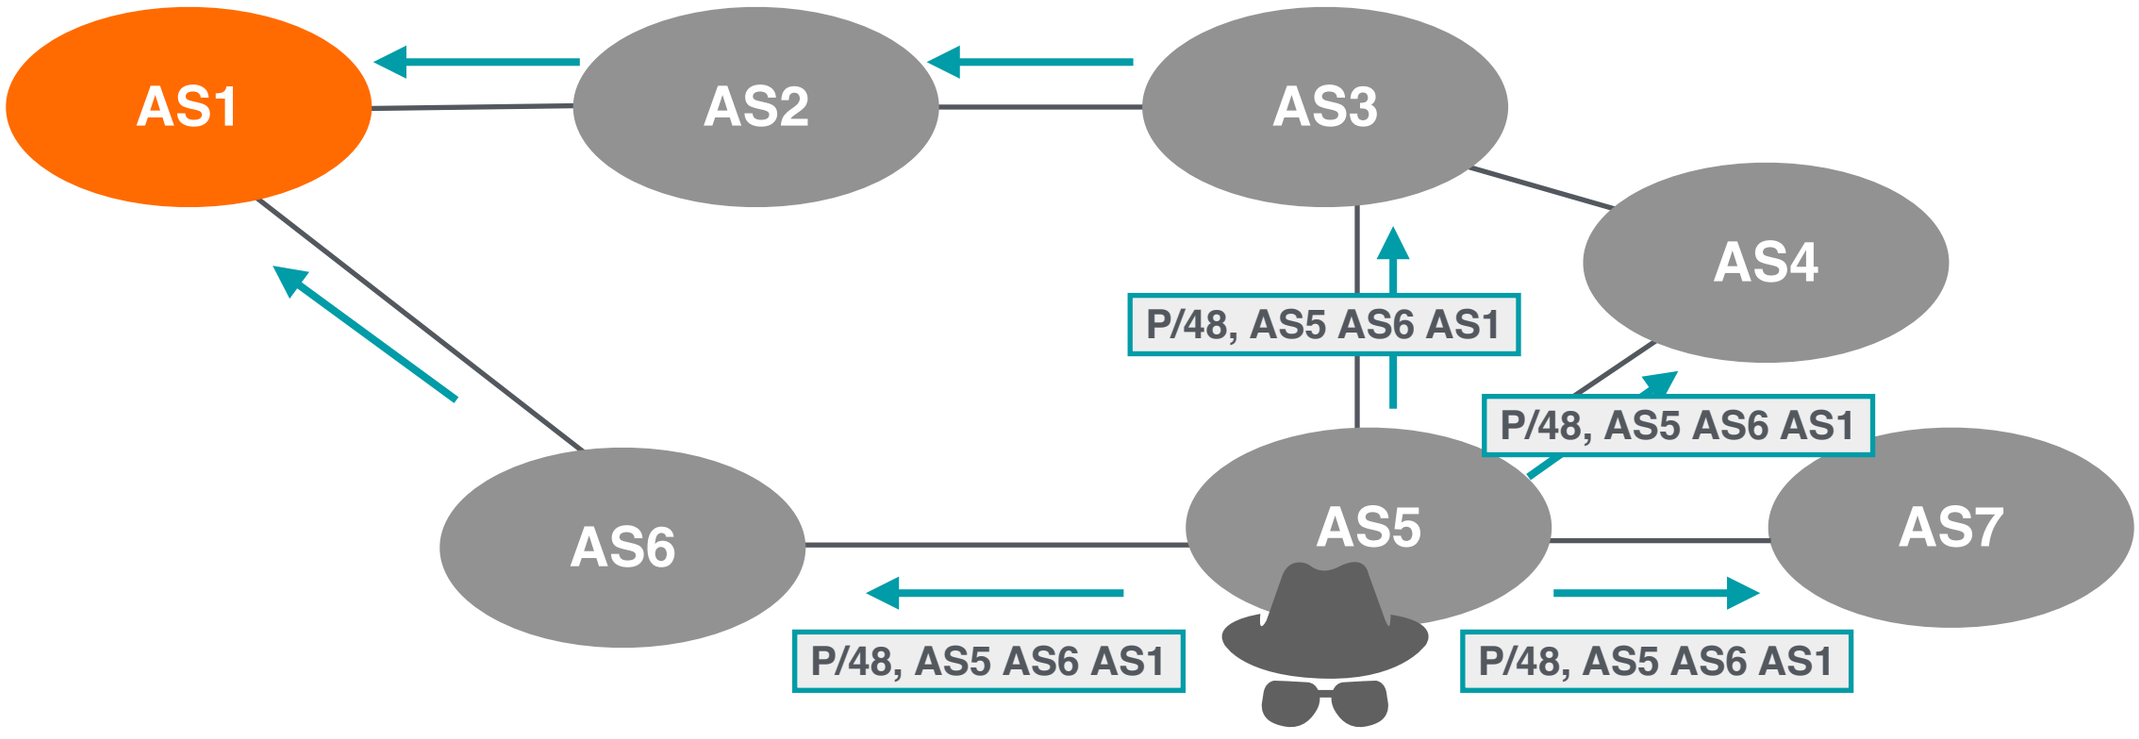

Man-in-the-middle Attack (MITM)

2 Replaces the prefix in a received update with the more-specific prefix

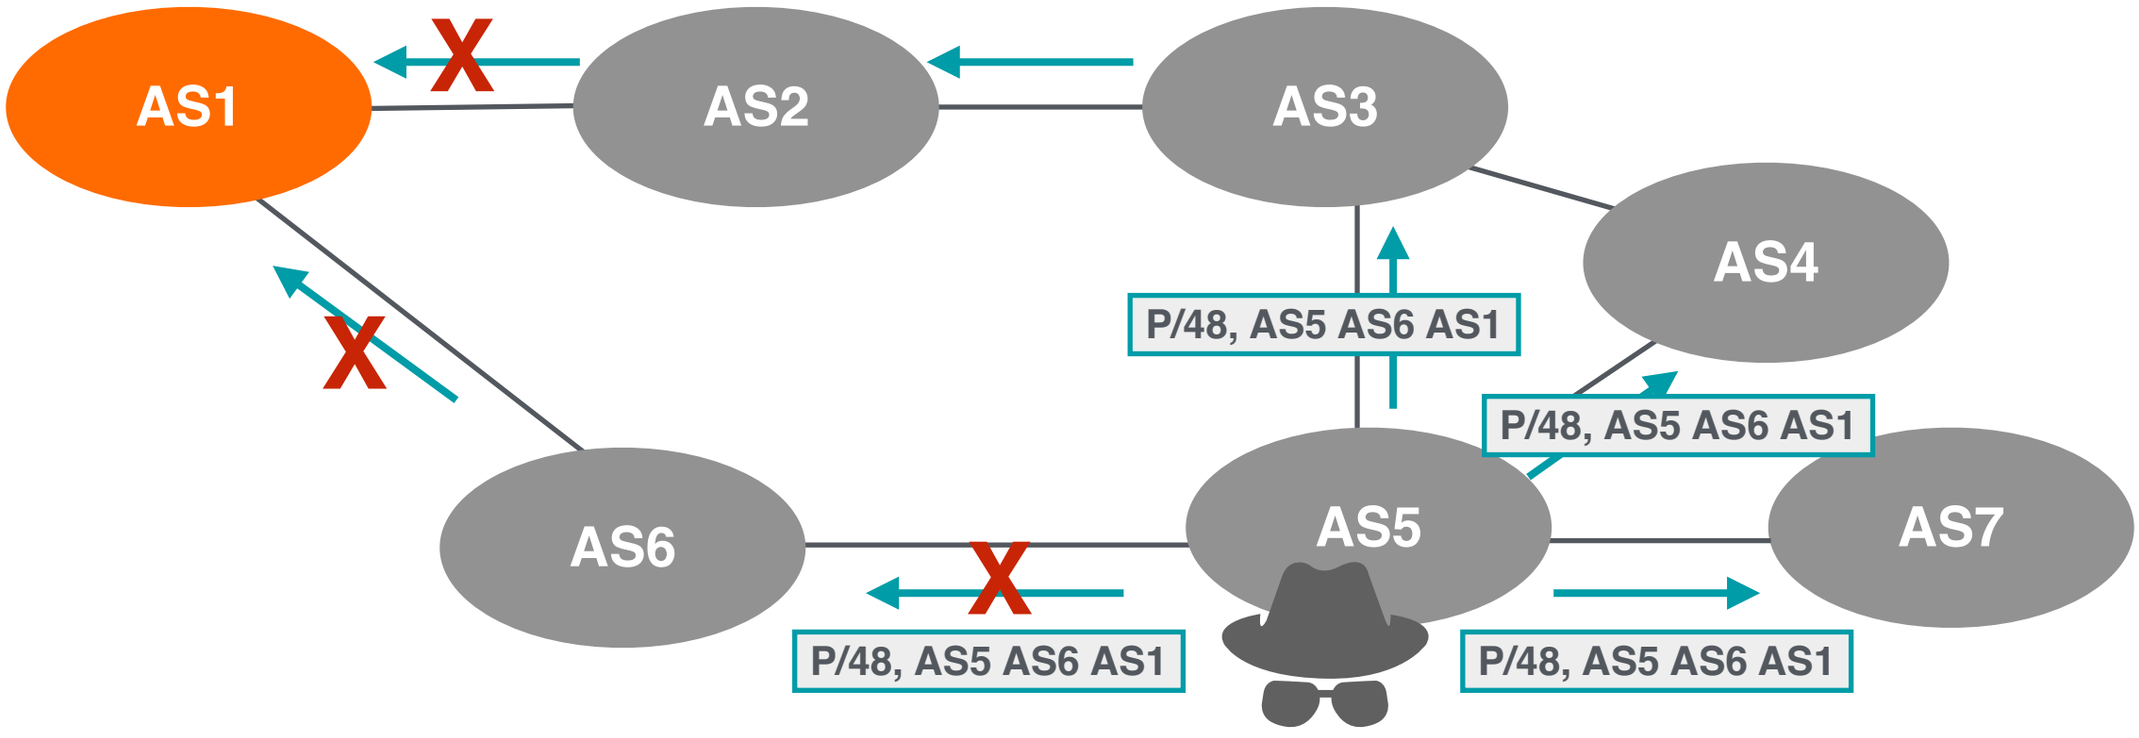

Only ASes on the return path discard this BGP announcement
(AS path loop avoidance)

Man-in-the-middle Attack (MITM)

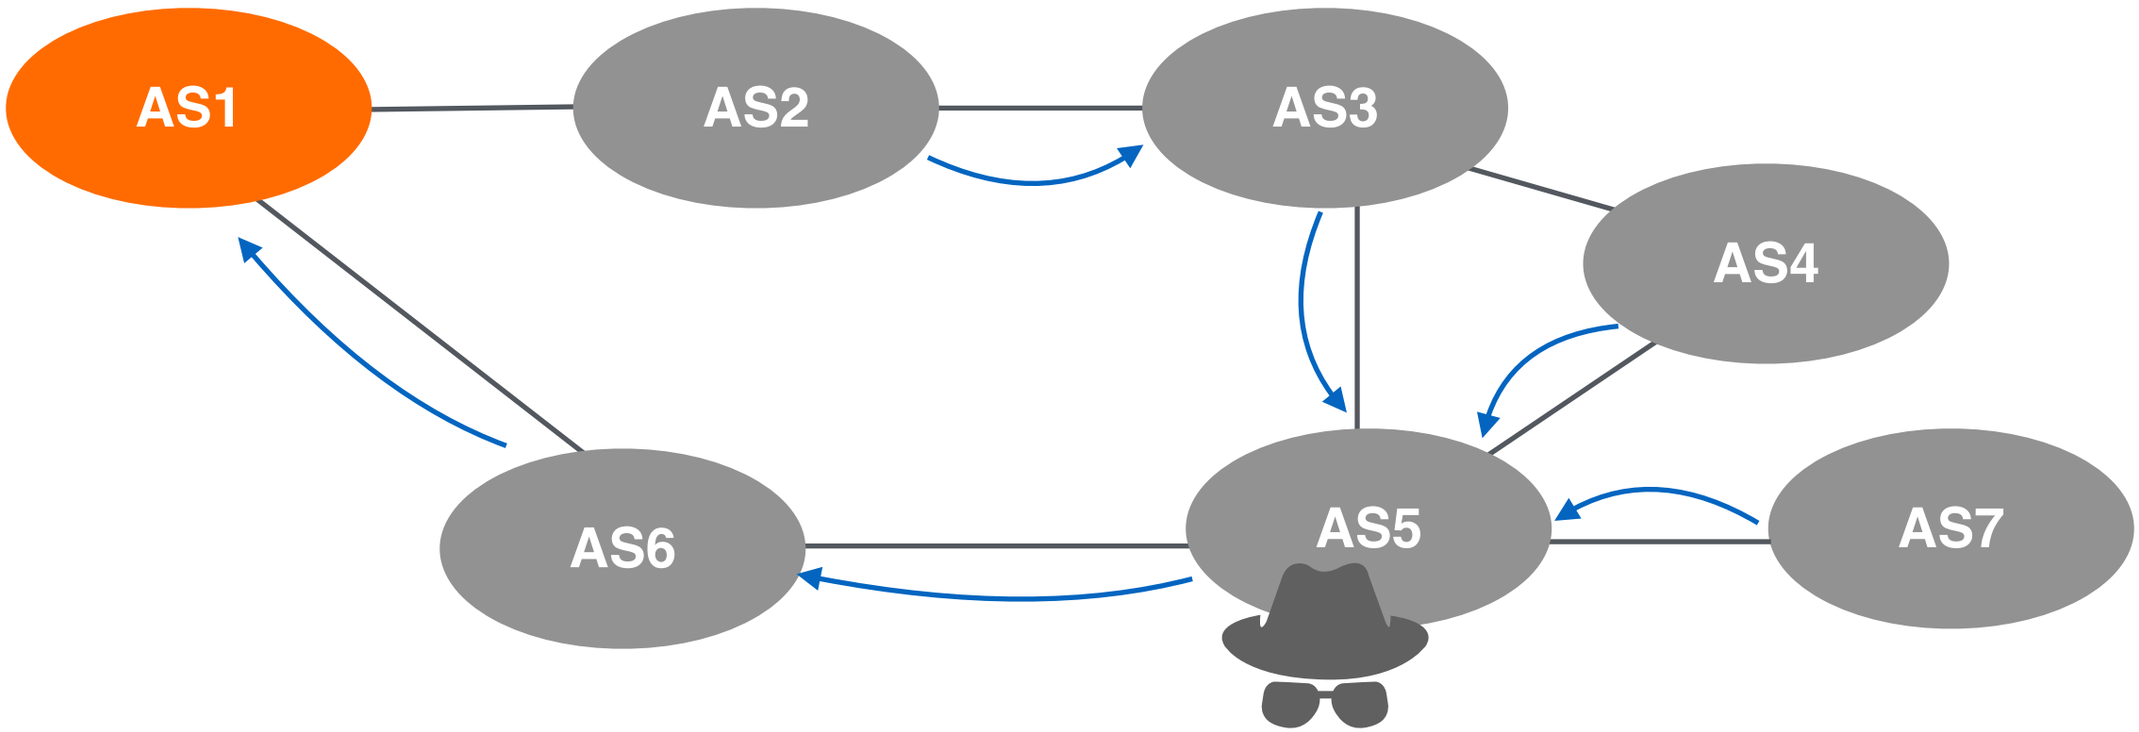

The more specific prefix wins! Traffic follows the incorrect path!

To summarise …

- BGP is vulnerable to mistakes and attacks
- Attackers could:
  - Inject bogus routing information into the BGP table
  - Hijack a BGP session and break peer-to-peer connections,
  - Initiate a DoS attack and exhaust victim's resources
  - Manipulate BGP and reroute packets
  - Intercept and eavesdrop
  - Blackhole the entire network

Questions

How to Secure Internet Routing

In order to secure routing, …

- Announce the right prefixes to the right peers
- Have proper filters in place to eliminate route leaks
- Validate the routing information you receive to mitigate hijacks
- Take all the measures you can to protect your network
- Remember that a single protection mechanism is not enough!
  - Apply several mechanisms together

What measures?

- RFC#7454 documents major countermeasures for BGP Operations and Security
- According to RFC and best practices, you should
  - Protect your BGP speaker (control plane and data plane filters)
  - Protect your BGP Sessions (MD5, TCP-AO)
  - Register your routing information in IRR system
  - Implement Route Filtering
  - Implement RPKI and validate the origin of received BGP routes

Protect your BGP Speakers

- Allow only BGP neighbours to send packets to TCP 179
  - Implement Control Plane Policing (CoPP)
  - Use data plane filters (ACLs) (If CoPP is not supported)
- Limit accepted BGP traffic
- uRPF to mitigate DoS/DDoS attacks

Protect your BGP Sessions

- Authenticate your BGP sessions and ensure the integrity of BGP messages
- MD5
  - Legacy solution, still widely deployed in many networks
  - Not a strong authentication mechanism, obsoleted by TCP-AO
- TCP-AO
  - Supports multiple stronger authentication algorithms
  - Provides better key management and agility
  - Supported by some vendors

Register your routing in IRR

- Create route, route6 objects in IRR database
- Update your routing registry information regularly
- Create filters based on IRR data
  - Automation relies on the IRR being complete
  - Check your output before using it
- Help others by documenting your policy

Mitigate BGP hijacks and route leaks

- In order to mitigate BGP hijacks and prevent route leaks
  - Implementing BGP filters is essential!
  - Create authorised statements in RPKI system for your prefixes
  - Validate the origin of received BGP routes

Questions

Implementing BGP Filters

Filtering is …

- The most basic protection mechanism to prevent malicious or accidental BGP incidents
- The technique used to control prefixes on the BGP peering
  - Which prefixes will you accept into your network?
  - Which prefixes will you advertise to your peers?

Why filter BGP prefixes?

- Remember some of the recent BGP incidents
  - Youtube (2008),
  - AWS route leak (2016),
  - Google prefix leak (2018),
  - Akamai, Amazon, Alibaba (2020) …
- Cause is poor filtering by upstream providers/peers
- They could have been prevented by implementing filters

BGP Filters (BGP Policies)

- Used to filter prefixes exchanged between BGP peers
- Should be applied on each eBGP peer
  - For received and advertised routes
- Filters can match on
  - IP prefixes
  - AS paths
  - Or any other BGP attributes such as BGP communities

BGP Filters (BGP Policies)

- Implemented on both ingress and egress
- Inbound policy
  - For Incoming (received) routes
  - Detects configuration mistakes and attacks
- Outbound policy
  - For outgoing (advertised) routes
  - Limits propagation of routing information

Implementing BGP Filters

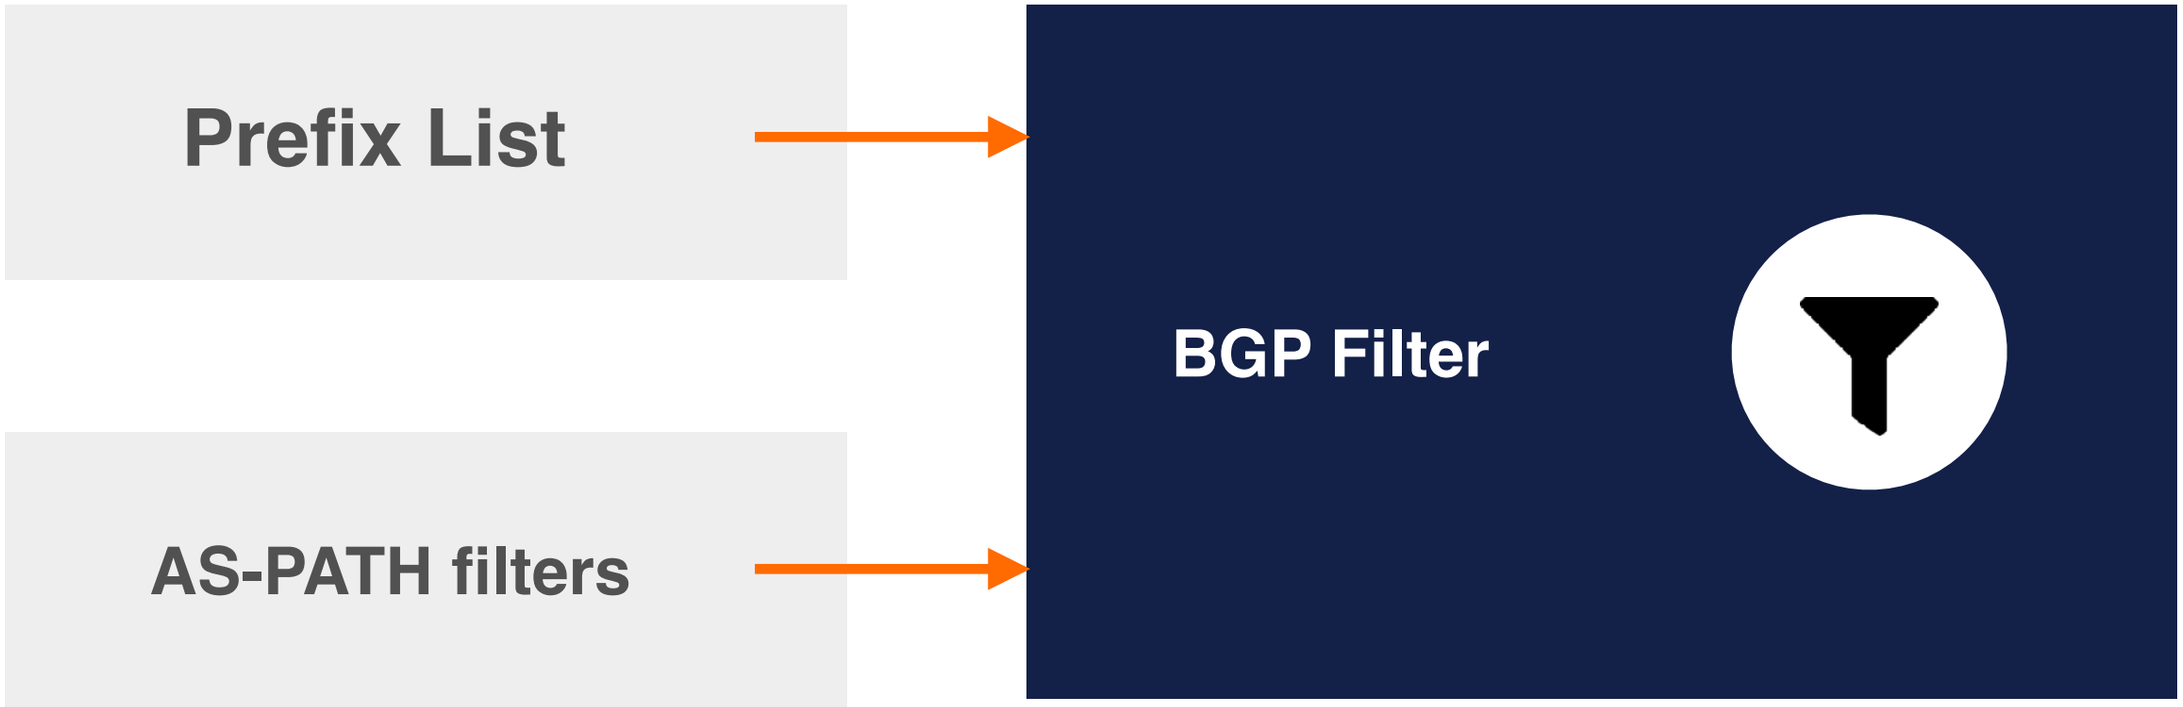

Prefix-list

- Lists of routes you want to **accept** or **announce**
- You can create them **manually** or **automatically** with data from IRRs
- It can be done using scripts or some tools
  - such as Filtergen (Level3), bgpq3, peval …
- Easy to use, but not highly scalable

AS-Path Filtering

- Filter routes based on AS-PATH
- Widely used and highly scalable
- Applied the same way as prefix-list filters

```
router bgp 65564
  network 10.0.0.0 mask 255.255.255.0
  neighbor 172.16.1.1 remote-as 65563
  neighbor 172.16.1.1 filter-list 1 out
  neighbor 172.16.1.1 filter-list 2 in

ip as-path access-list 1 permit 65564
ip as-path access-list 2 permit 65563
```

Which prefixes should be filtered?

- RFC#7454, “BGP Operations and Security”, lists the prefixes to be filtered
  - Special-purpose prefixes (IPv4/IPv6) (Martians)
  - Unallocated prefixes
  - Prefixes that are too specific
  - Prefixes owned by an AS
  - IXP LAN prefixes
  - The default route (0.0.0.0/0, ::/0)

Prefix filtering recommendations

- In full routing networks, some policies should be applied
  - On each BGP peer
  - For both received and advertised routes (Inbound&outbound)
- Recommendations vary based on type of BGP peering relationships
  - Public/Private peering
    - Transit provider
    - Customer

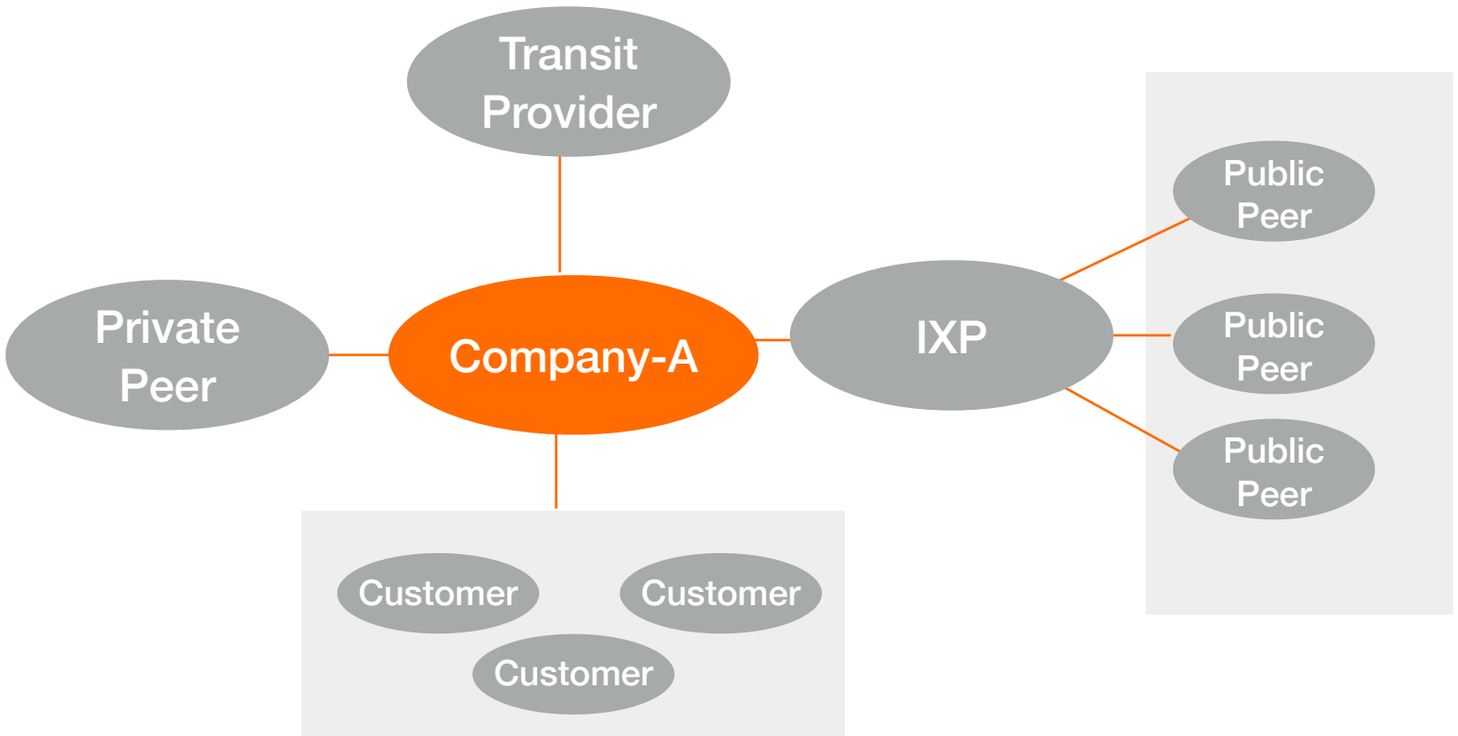

Filters with Peers (inbound)

- Filters with Public/Private Peers
- On **inbound**, strict or loose filtering could be implemented
- Strict filtering:
  - Makes sure advertisements conform to what is declared in IRRs
  - Impact should be checked before applying the policy
- Loose filtering:
  - Filters the routes based on RFC#7454 recommendations

Prefixes that are not globally routable	Prefixes belonging to the local AS
Prefixes not allocated by IANA (IPv6 only)	IXP LAN prefixes
Routes that are too specific	The default route

Filters with Peers (outbound)

- Only **locally originated and customers' prefixes** should be sent
  - If possible, list the prefixes to be advertised, and deny the rest!
- Additional filters could be added to filter the followings
  - Prefixes that are not globally routable
  - Routes that are too specific
  - IXP LAN prefixes
  - The default route
  - Prefixes learnt from other peers or transit providers

Filters with Transit (inbound)

- If FRT is desired,
  - RFC#7454 recommendations are the same with public/private peers
  - except the default route
- If upstream provider is supposed to announce the default route only
  - accept only the default route

Filters with Transit (outbound)

- The same outbound filters should be applied as those for public/ private peers
- Make sure that only authorised prefixes are sent
  - Locally originated and customers' prefixes
- Filter the prefixes learnt from other peers or other transit providers

Filters with Customers (inbound)

- If all customer prefixes are known,
  - Accept customer prefixes only and discard the rest!
- What if you do not have this information? Filter the followings:
  - Special purpose prefixes
  - Unallocated prefixes
  - Prefixes that are too specific
  - Prefixes belonging to the local AS
  - The default route

Filters with Customers (outbound)

- According to RFC#7454, it may vary depending on customers preferences
- If customer asks for default route
  - send only default
- For other cases, filter the following prefixes:
  - Prefixes that are not globally routable
  - Too specific routes
  - The default route (?)

Data sources

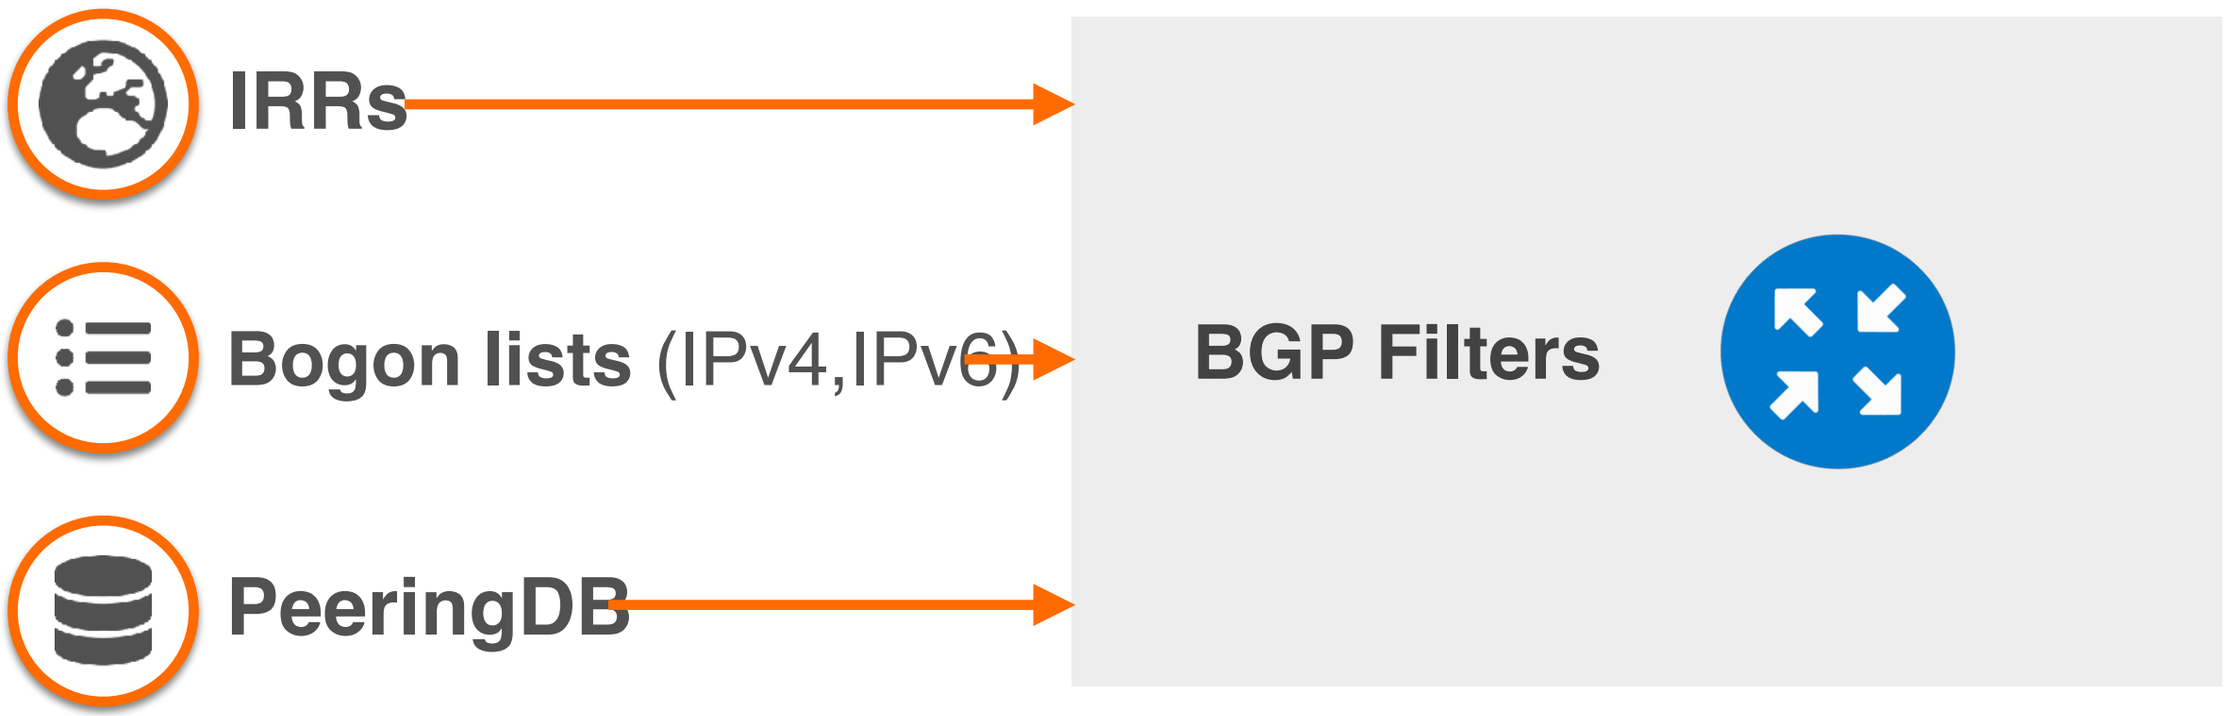

Bogon lists

- **Bogons** are prefixes that should never appear in the Internet routing table!
  - Martians (RFC#1918 Private addresses + Reserved space)
  - IANA unallocated space
- **Full Bogons** should be filtered as well
  - Bogons + RIR unallocated/assigned
- The bogon and full bogon lists are not static
- Team Cymru provides lists of bogons and full bogons

https://www.team-cymru.com/bogon-reference-http

ASN Bogons

ASNs	Reserved?
0	Reserved - RFC7607
23456	AS_TRANS - RFC6793
64496-64511 and 65536-65551	Reserved for use in docs and code - RFC5398
64512-65534 and 4200000000-4294967294	Reserved for Private Use - RFC6996
65535 and 4294967295	Last 16 and 32 bit ASNs - RFC 7300
65552-131071	Reserved - IANA

Questions

Routing Security with RPKI

What is RPKI?

What is RPKI?

- RPKI is …
  - **Resource certification** (X.509 PKI certificates)
  - A security framework
- It is used to make Internet routing more secure and reliable

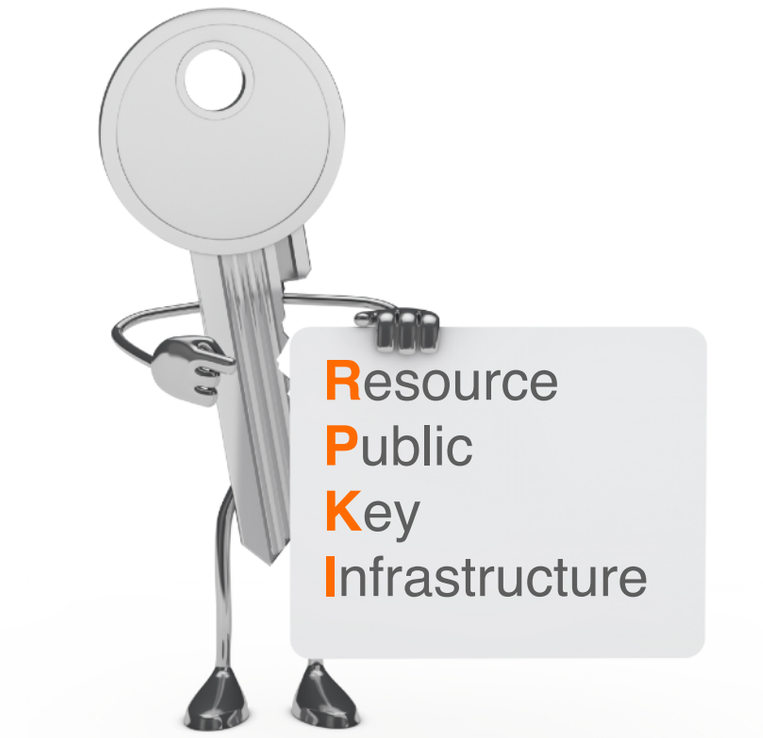

How RPKI enables Routing Security

- Verifies the association between resource holders and their resources.
  - Proves holdership through a public key and certificate infrastructure
- Used to validate the origin of BGP announcements
  - Is the originating ASN authorised to originate a particular prefix?
- Stepping stone to “Path Validation”

Implementing RPKI helps to prevent…

- BGP Origin Hijacks
  - Caused by malicious activities
- Mis-origination
  - Due to typos/fat fingers
- Route leaks
  - Caused by configuration mistakes

How does it work?

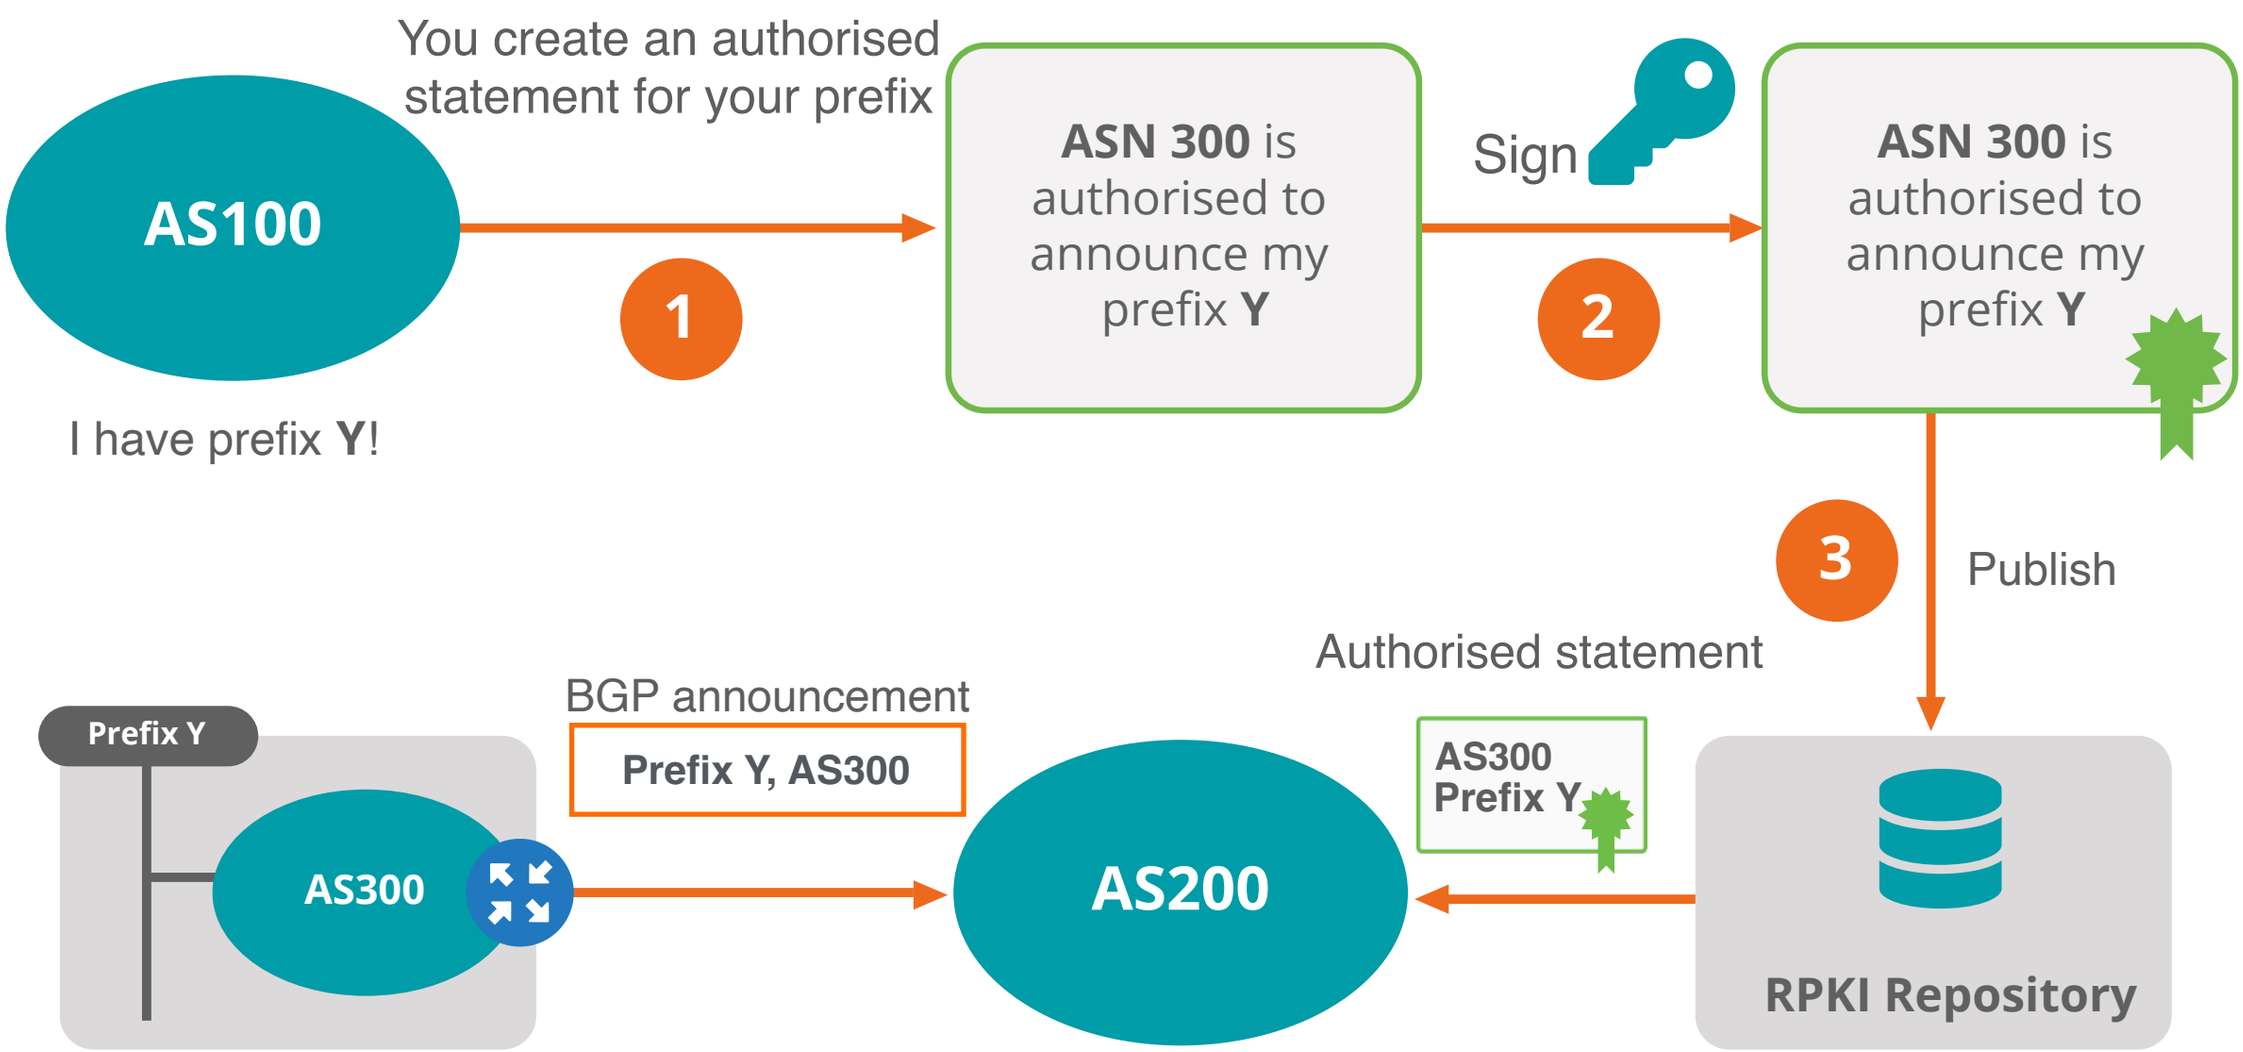

4 Others use those statements to make better routing decisions!

Trust in RPKI

- RPKI relies on the five RIRs as Trust Anchors
- Certificate structure follows the RIR hierarchy
- RIRs issue certificates to resource holders

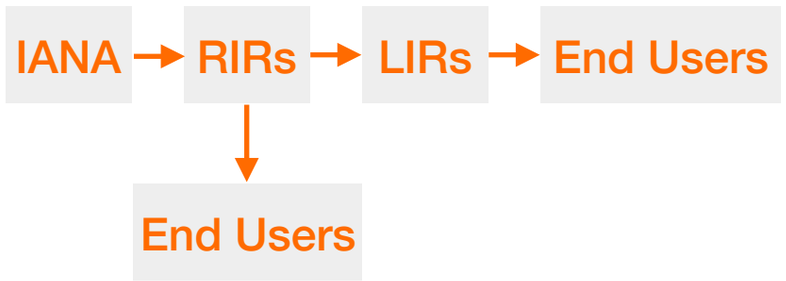

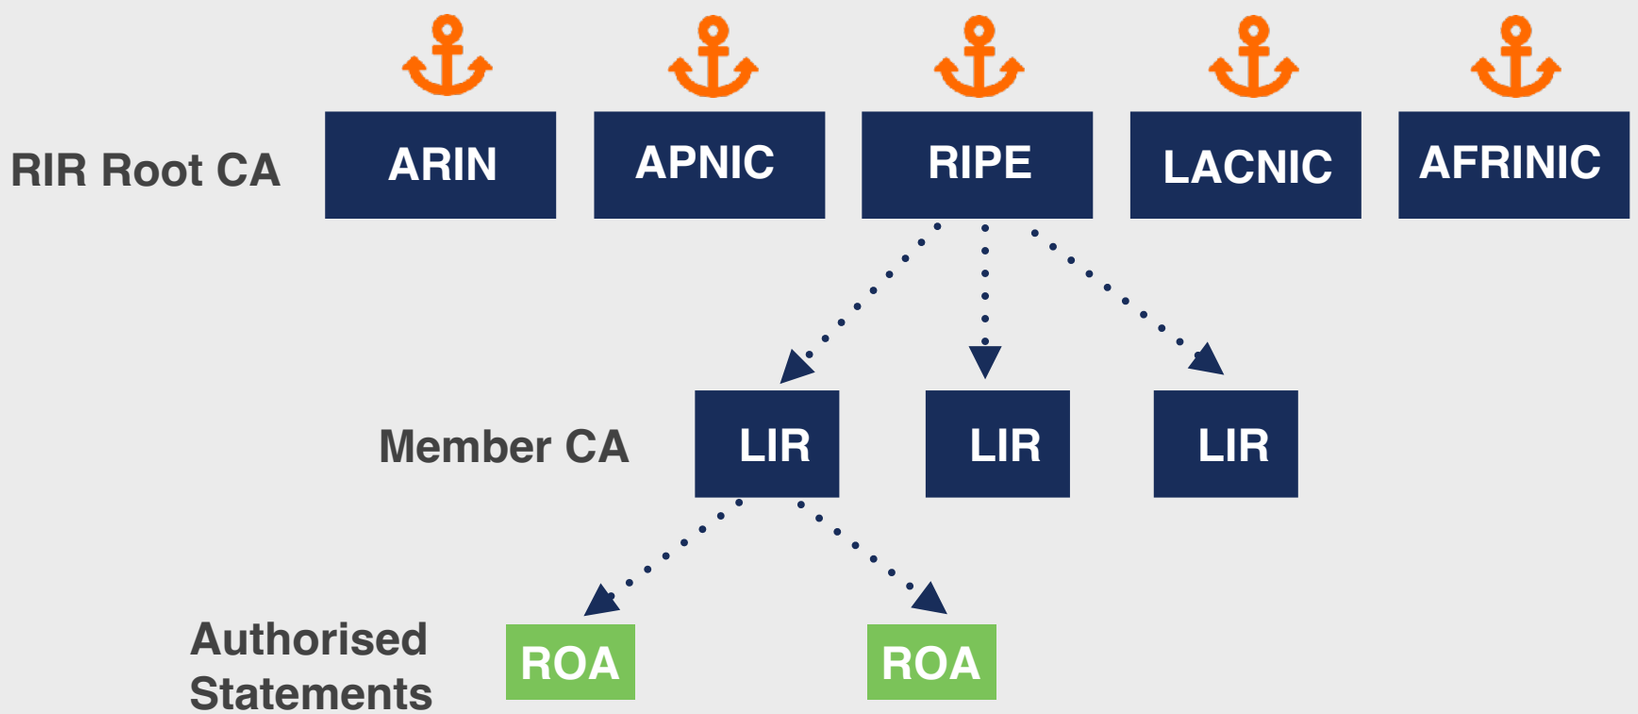

Trust in RPKI

- **Root certificate**
  - **Self-signed**
  - RIRs use root certificate to sign LIRs' certificates

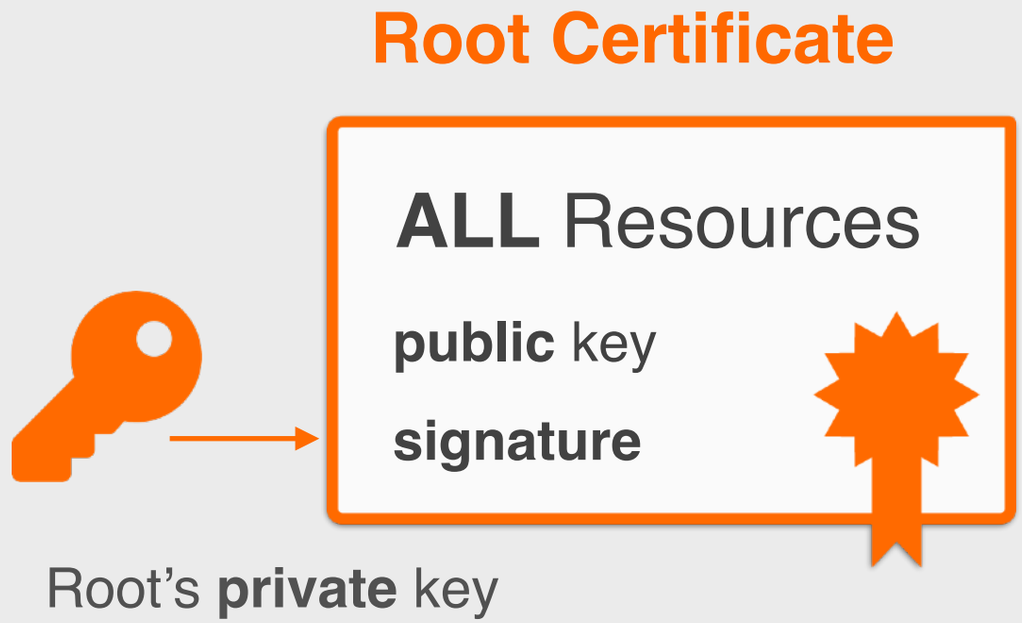

Trust in RPKI

- **Root certificate**
  - **Self-signed**
  - RIRs use root certificate to sign LIRs' certificates
- **LIR certificate**
  - Resource certificate for member allocations
  - Binds LIR's resources to LIR's public key
  - Proves legitimate holdership for the LIR's resources

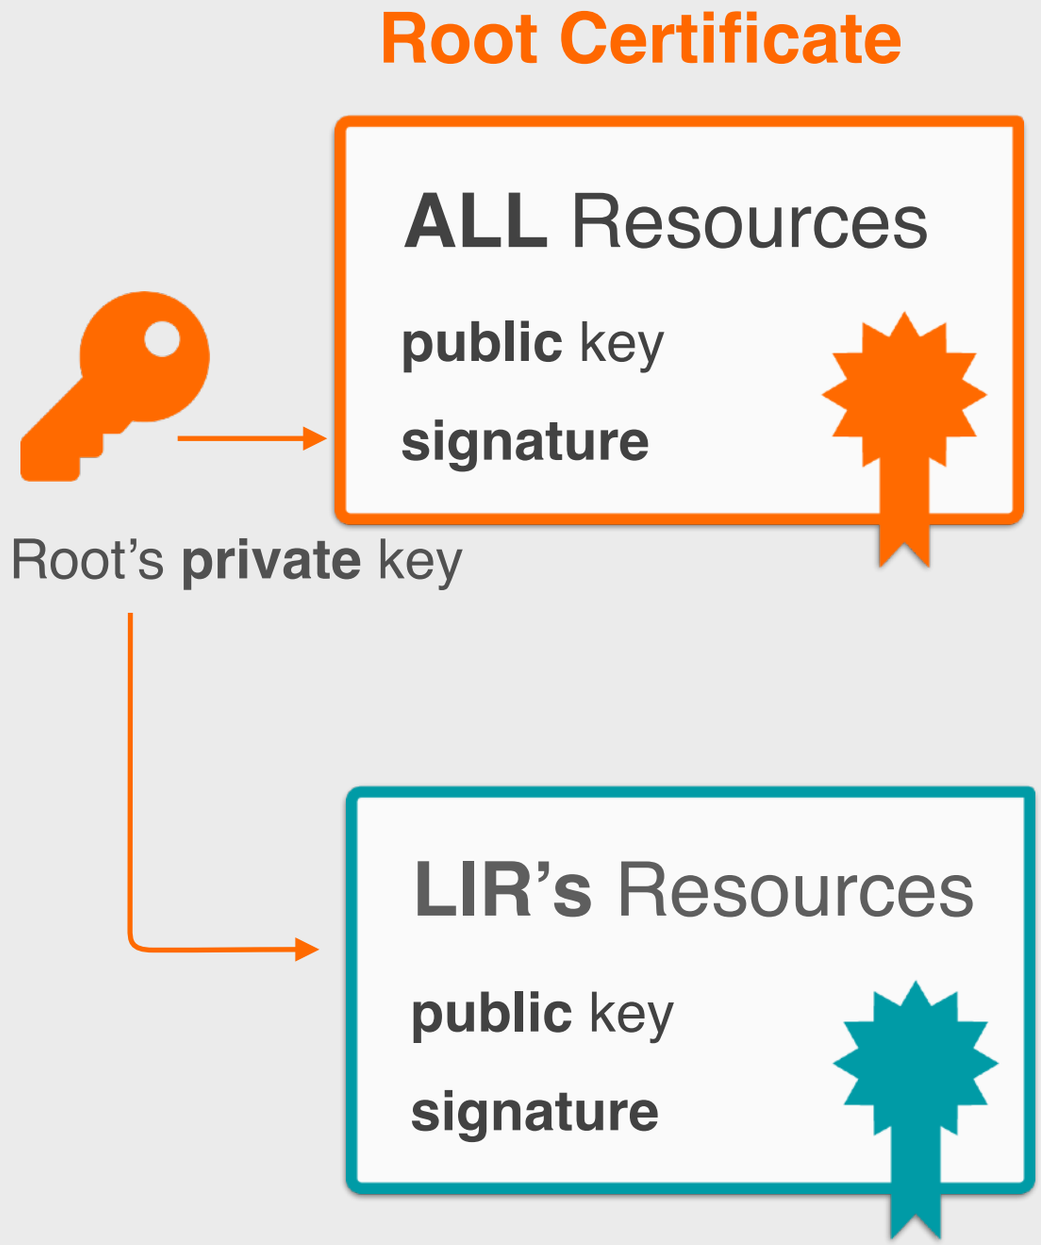

Trust in RPKI

- **Authorised statements**
  - Known as a ROA (Route Origin Authorisation)
  - Cryptographically signed object
  - Signed by LIR’s private key

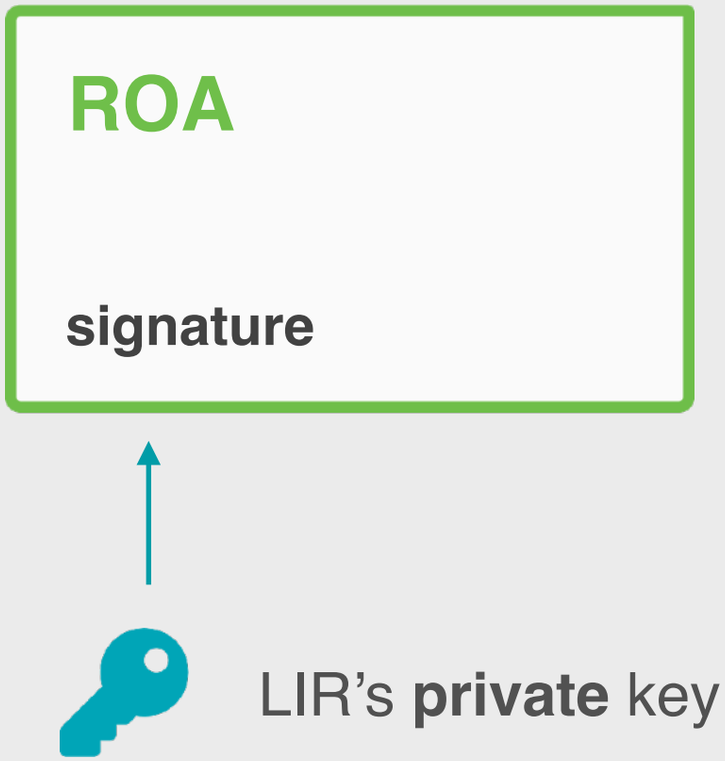

RPKI Chain of Trust

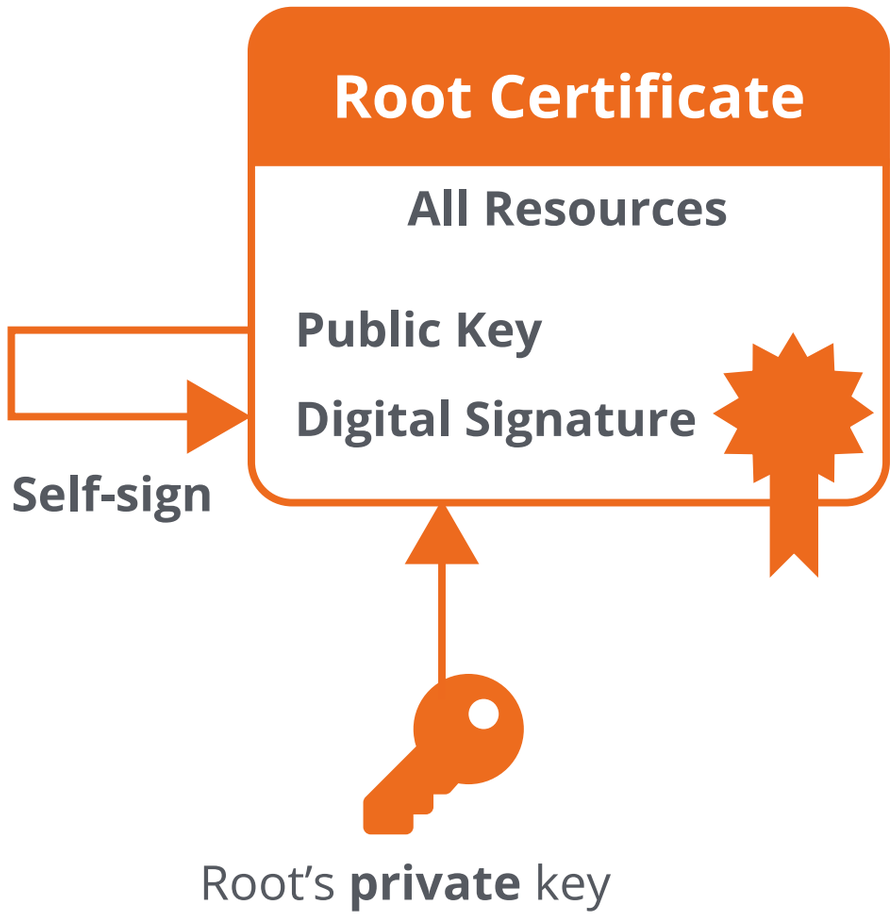

RPKI Chain of Trust

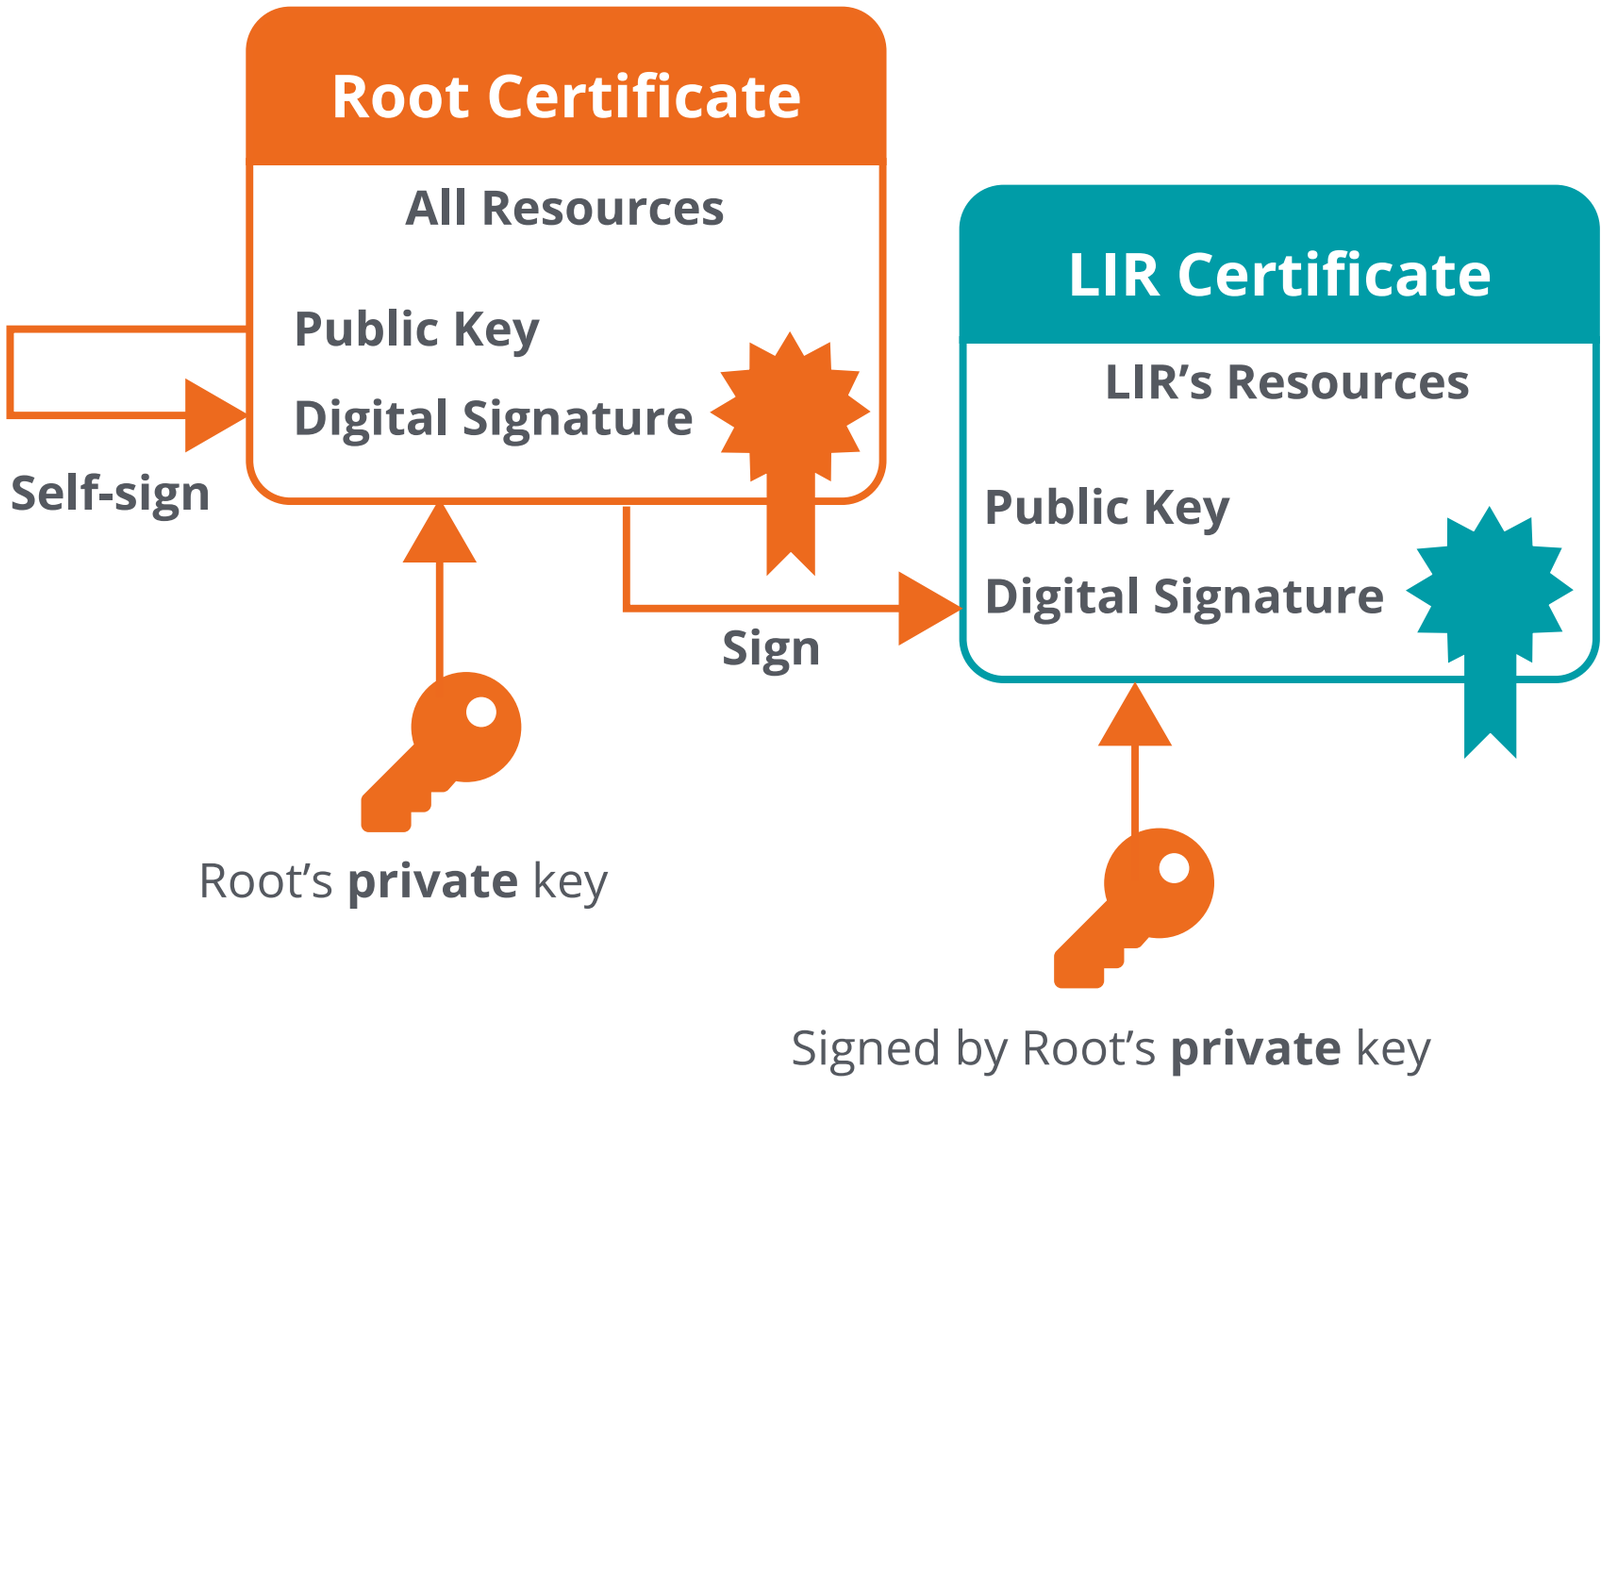

RPKI Chain of Trust

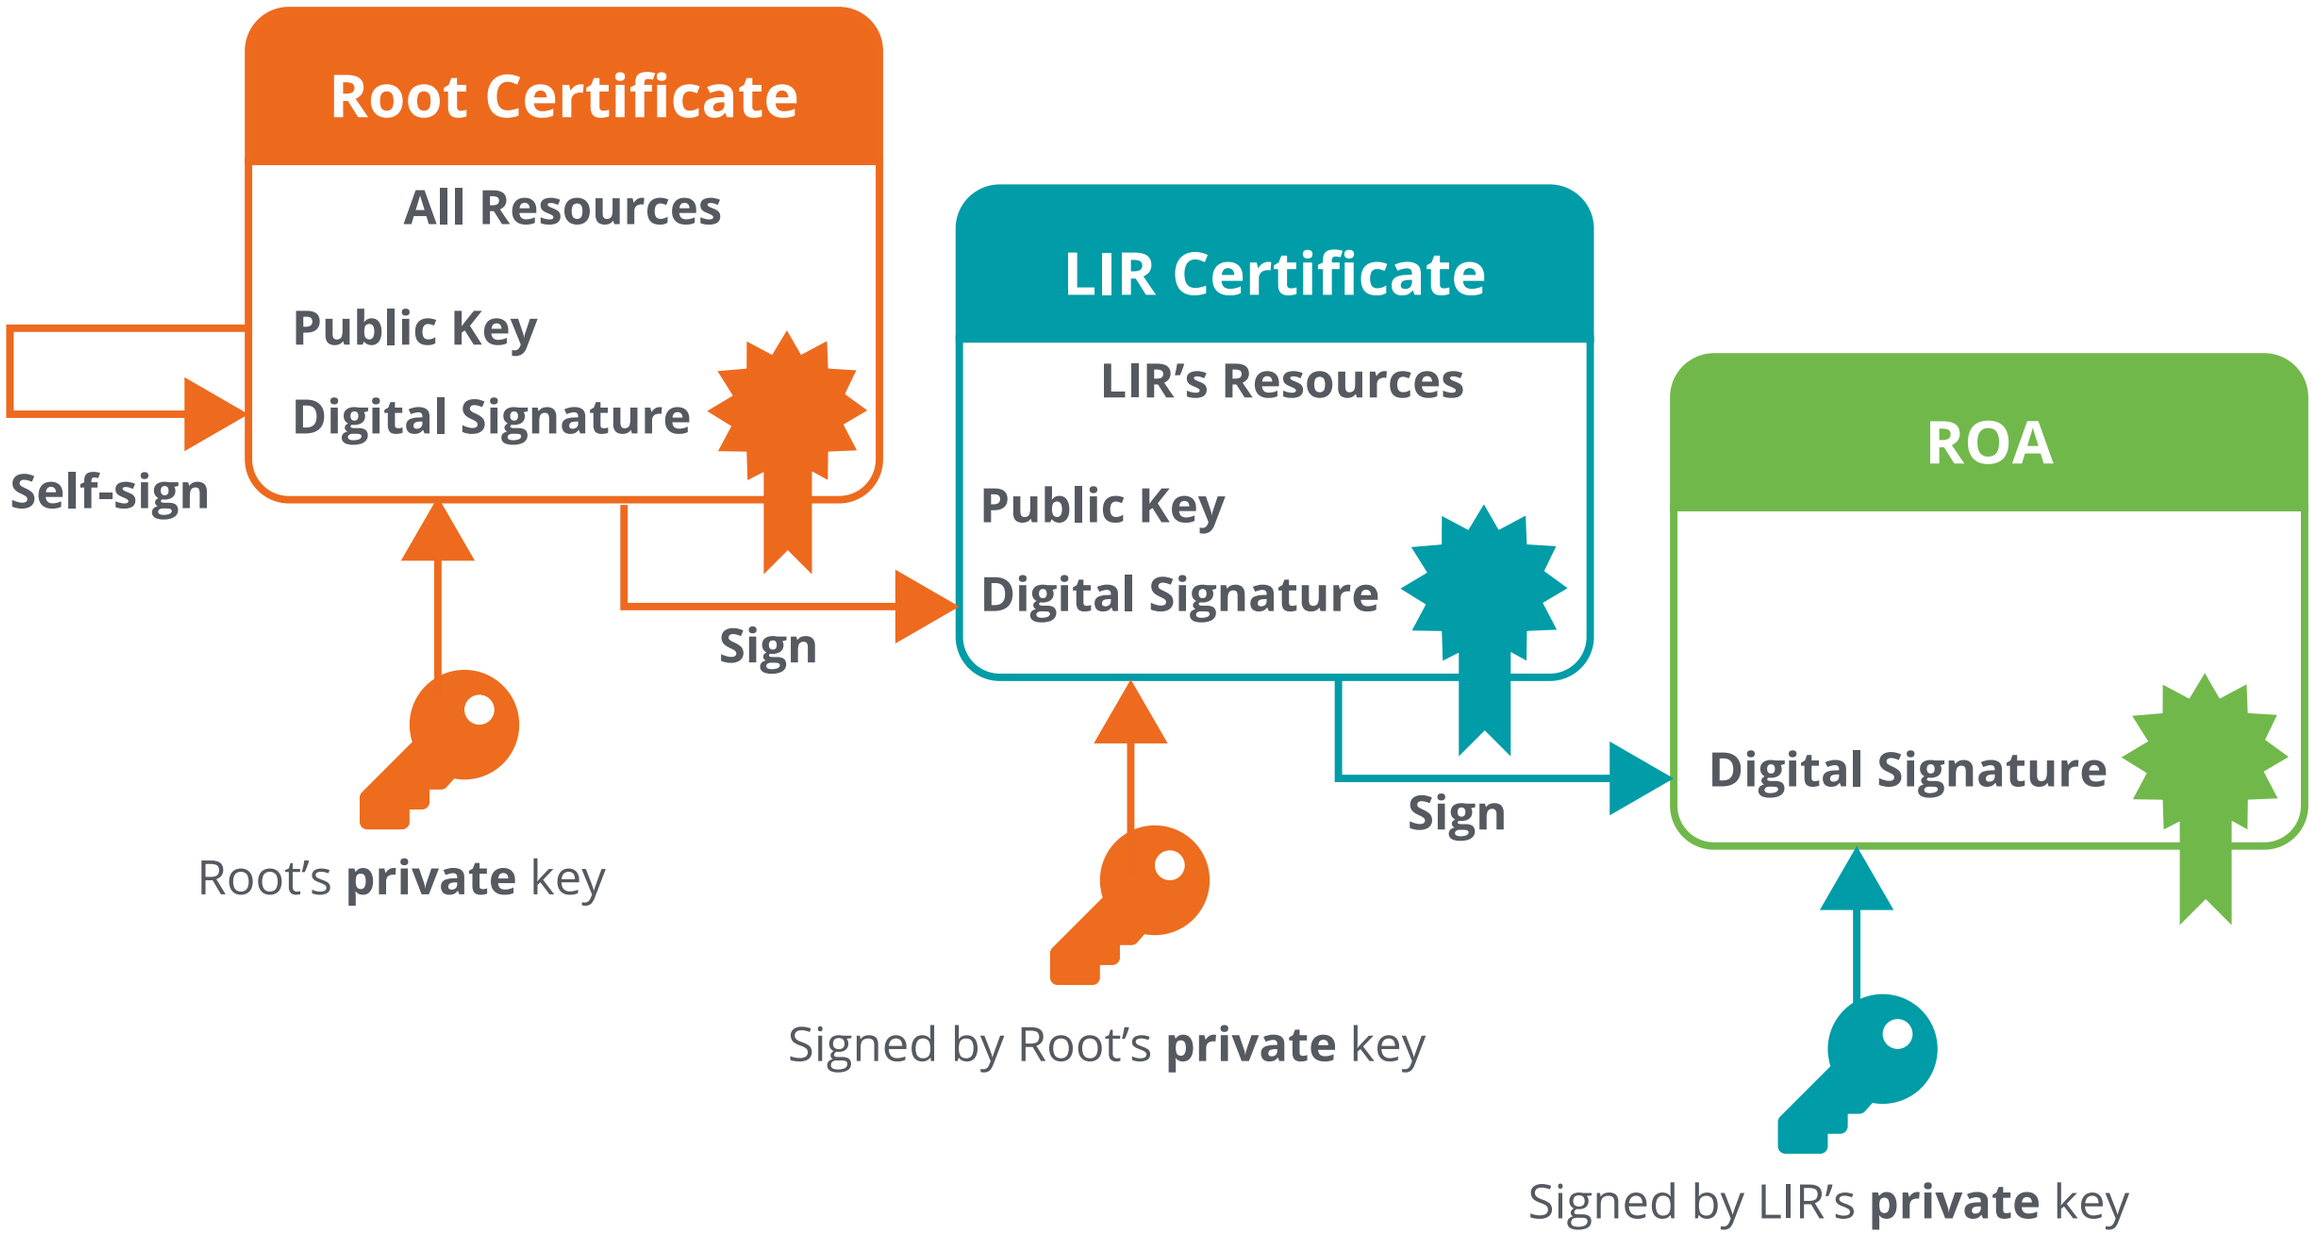

Elements of RPKI

- RPKI system consists of two parts…

SIGNING

Create ROAs for your prefixes in the RPKI system

+

VALIDATION

Verify the information provided by others

Elements of RPKI

- RPKI system consists of two parts…

SIGNING

Create ROAs for your prefixes in the RPKI system

+

VALIDATION

Verify the information provided by others

Routing Security with RPKI

Registering in the RPKI system (ROA)

What are ROAs?

- An **authorised statement** created by the resource holder
- It states that a certain prefix can be originated by a certain AS
- LIRs can create ROAs for their resources
- Multiple ROAs can exist for the same prefix
- ROAs can overlap

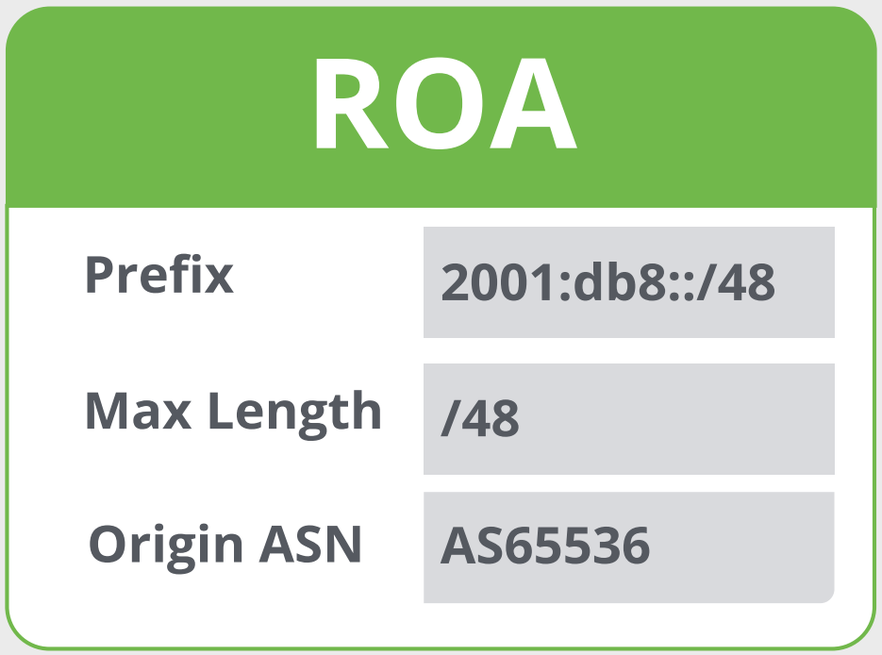

What is in a ROA?

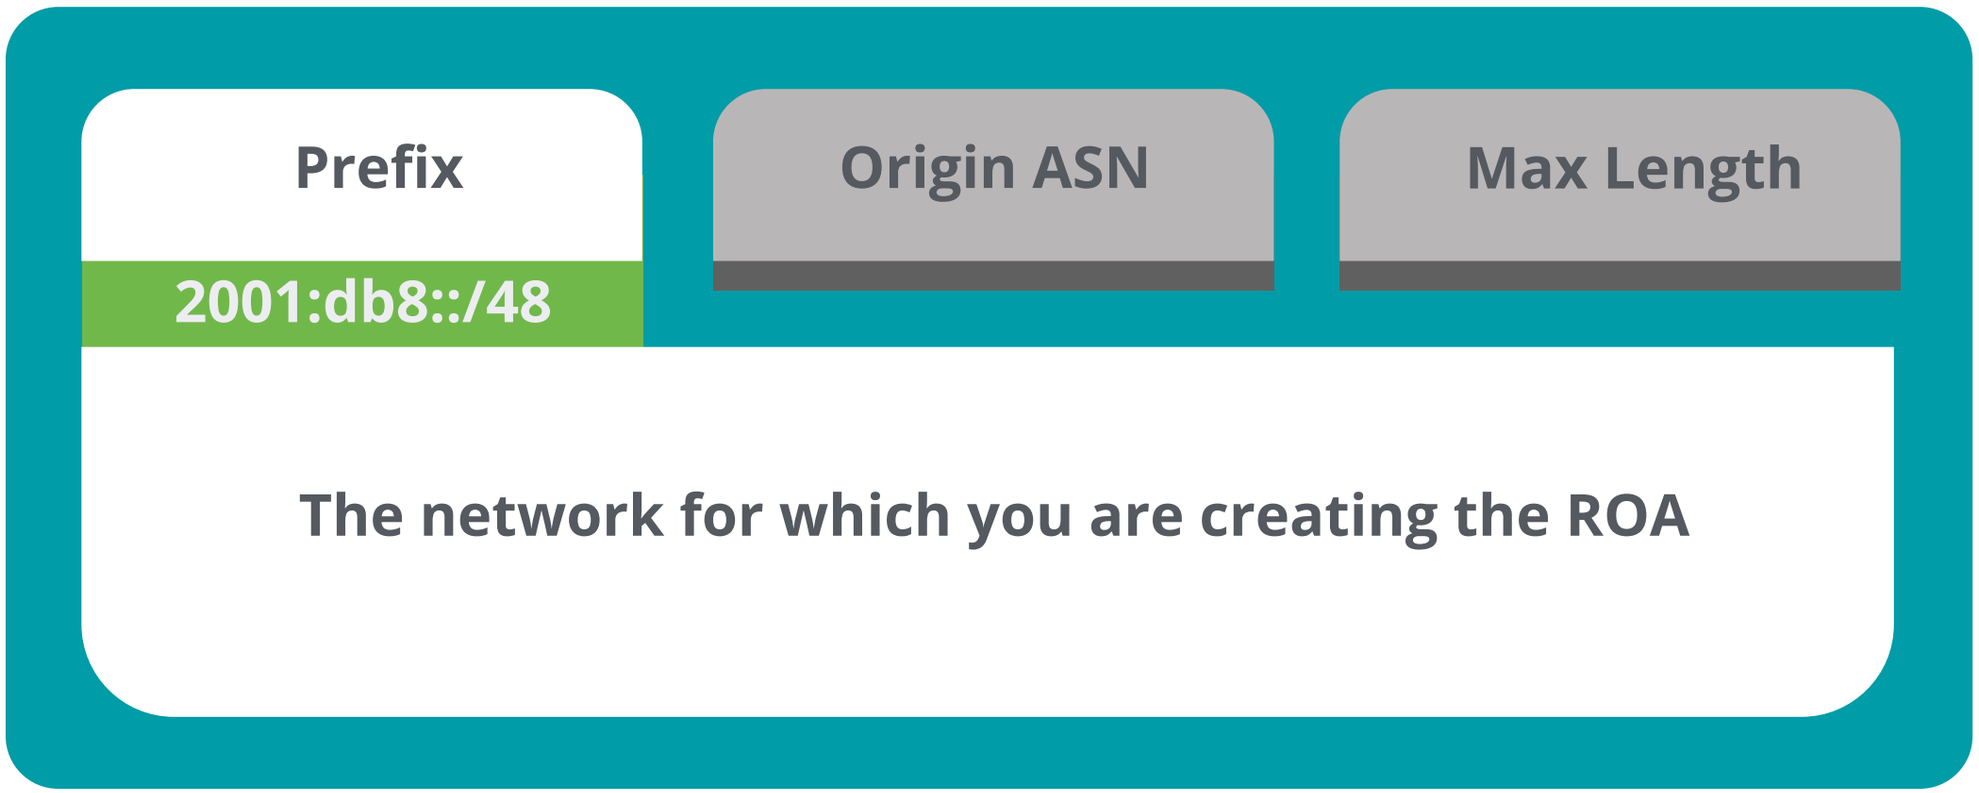

What is in a ROA?

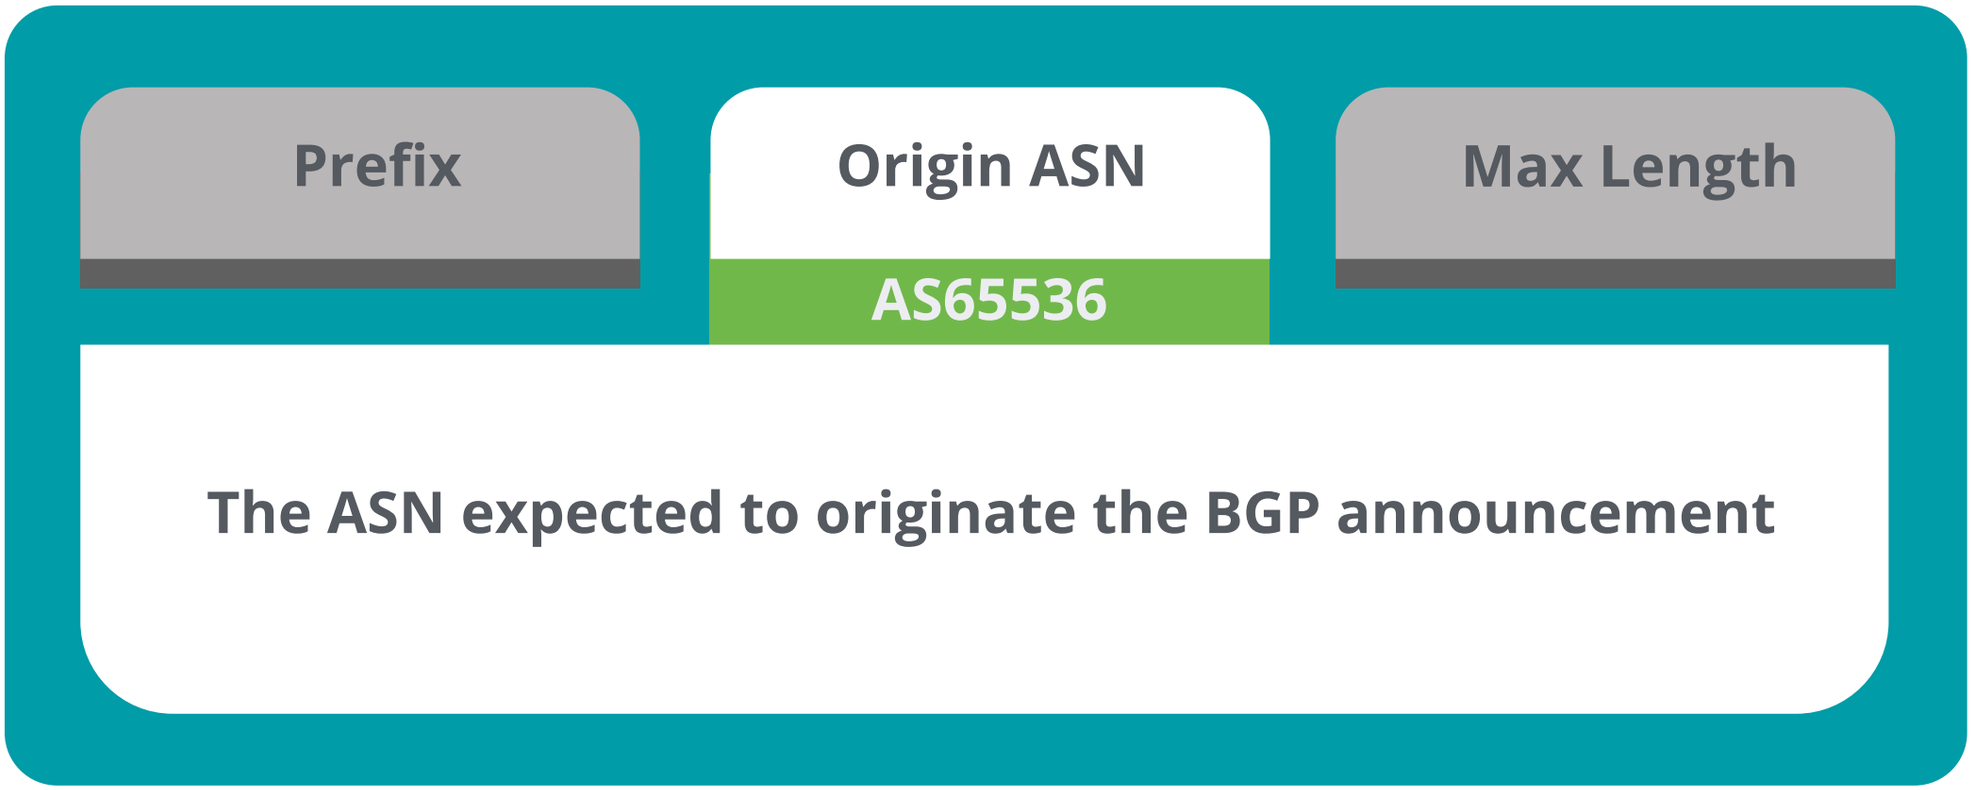

What is in a ROA?

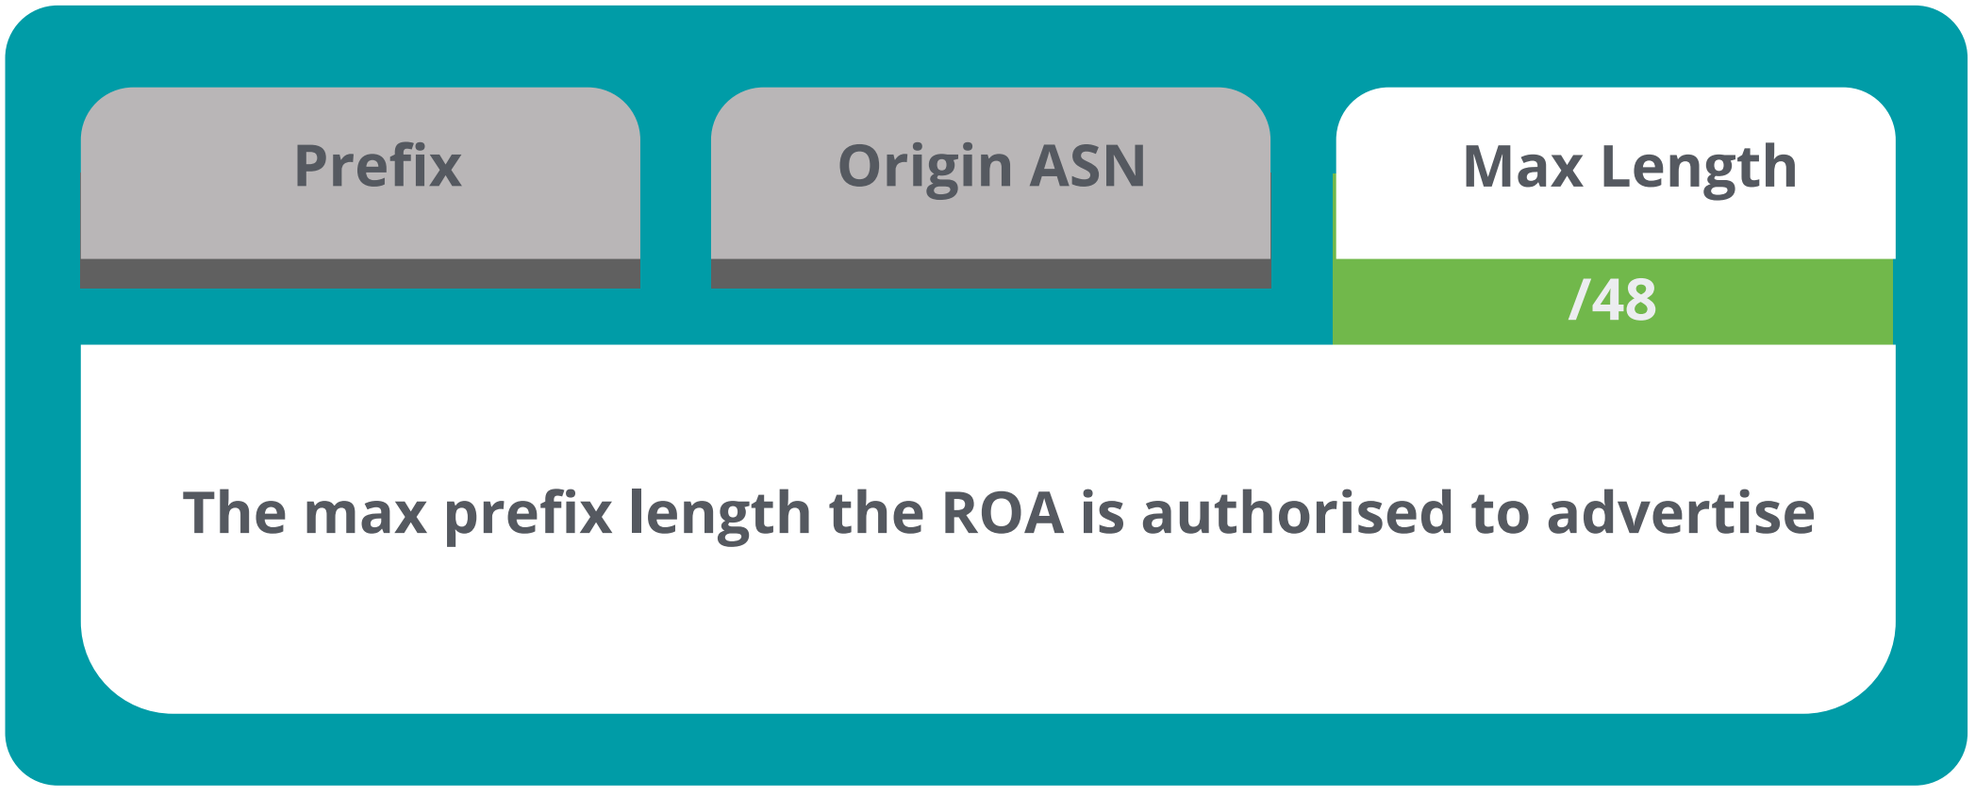

Max-Length

RIPE NCC (AS3333) has an IP address allocation

193.0.0.0/21

Max-Length

RIPE NCC (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 ASN	AS3333

Max-Length

RIPE NCC (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 ASN	AS3333

Max-Length

RIPE NCC (AS3333) has an IP address allocation

AS3333 creates this ROA

According to ROA;

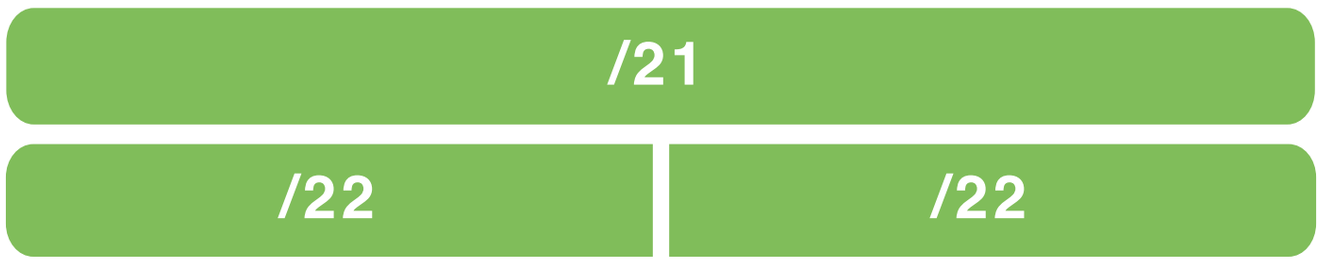

193.0.0.0/21

ROA

Prefix	193.0.0.0/21
Max Length	/22
Origin ASN	AS3333

Max-Length

RIPE NCC (AS3333) has an IP address allocation

AS3333 creates this ROA

According to ROA;

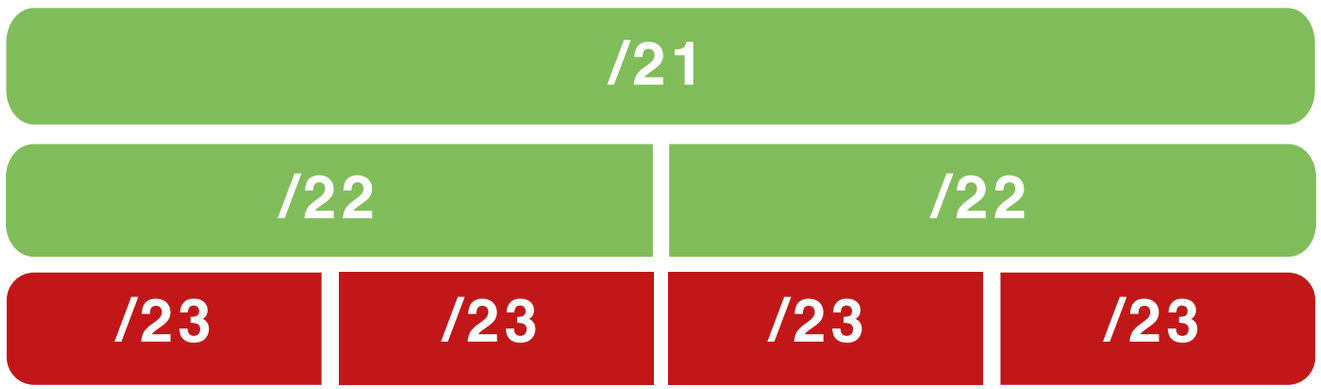

193.0.0.0/21

ROA

Prefix	193.0.0.0/21
Max Length	/22
Origin ASN	AS3333

Max-Length

RIPE NCC (AS3333) has an IP address allocation

AS3333 creates this ROA

According to ROA;

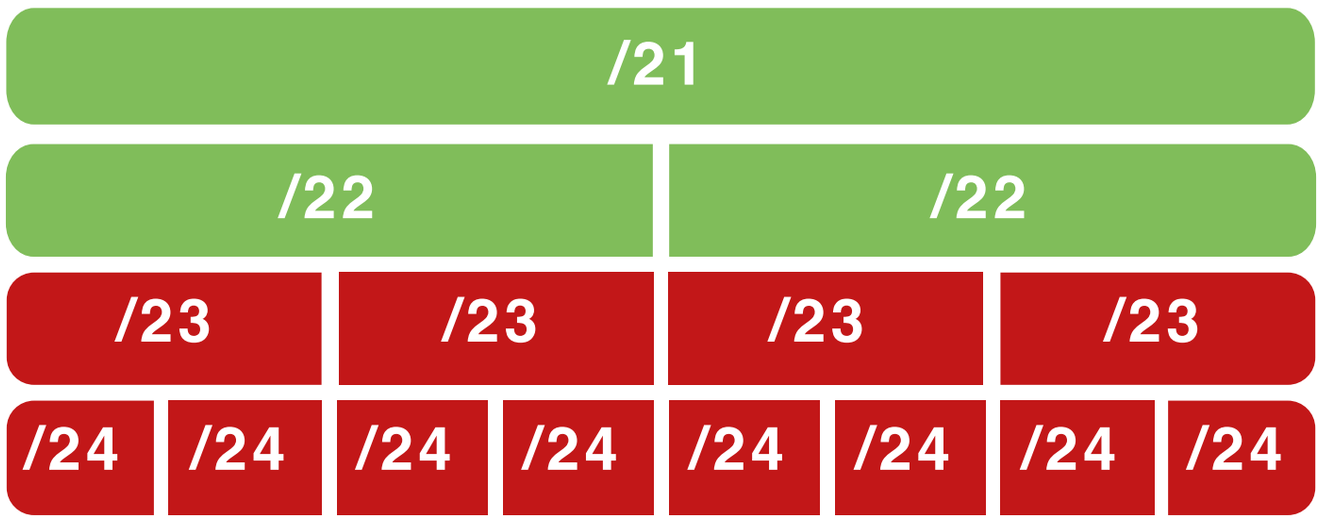

193.0.0.0/21

ROA

Prefix	193.0.0.0/21
Max Length	/22
Origin ASN	AS3333

Max-Length

RIPE NCC (AS3333) has an IP address allocation

AS3333 creates this ROA

According to the ROA;

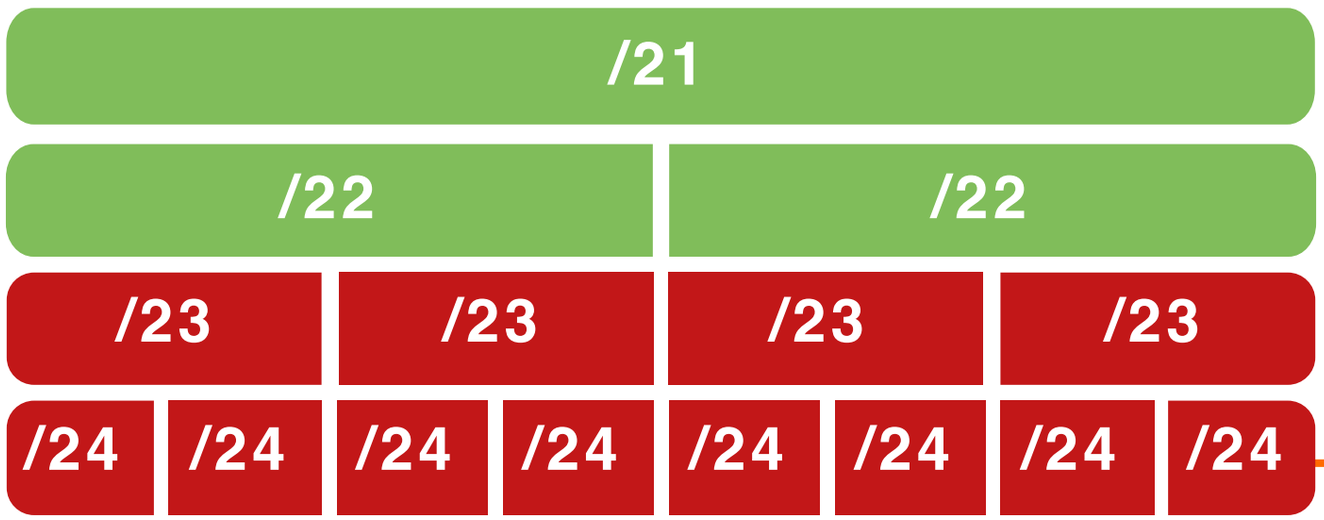

193.0.0.0/21

ROA

Prefix	193.0.0.0/21
Max Length	/22
Origin ASN	AS3333

Any more specific announcements are not authorised by the ROA.

How can you create a ROA?

- Login to the LIR Portal (my.ripe.net)
- Go to the RPKI Dashboard
- Choose which RPKI model to use

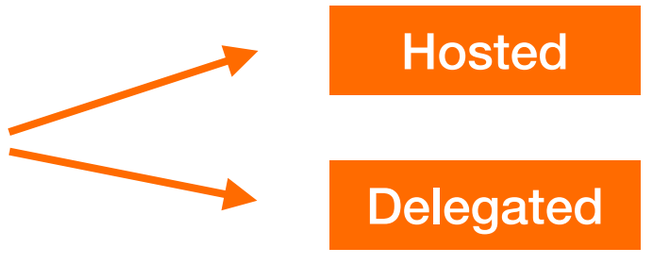

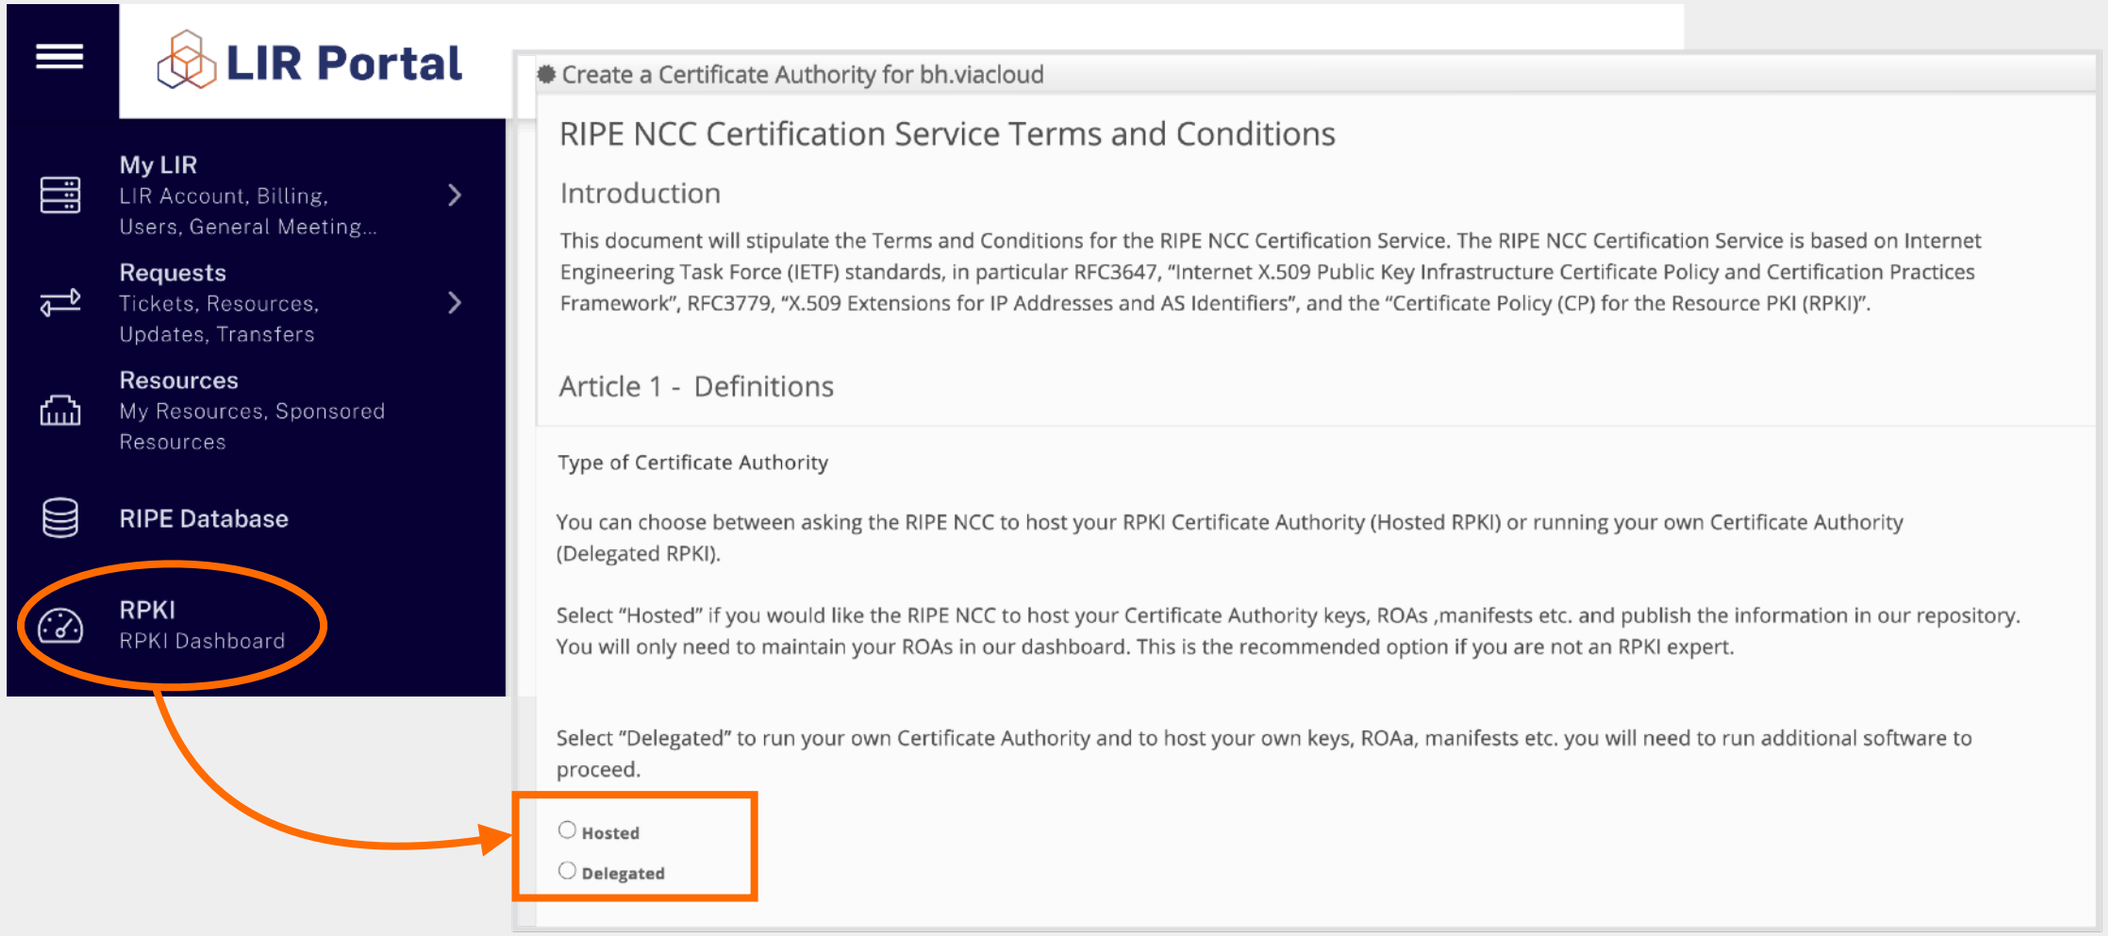

Hosted RPKI

- ROAs are created and published using the **RIR's member portal**
- The RIR hosts a CA for LIRs and signs all ROAs
- Automated signing and key rollovers
- Allows LIRs to focus on creating and publishing ROAs

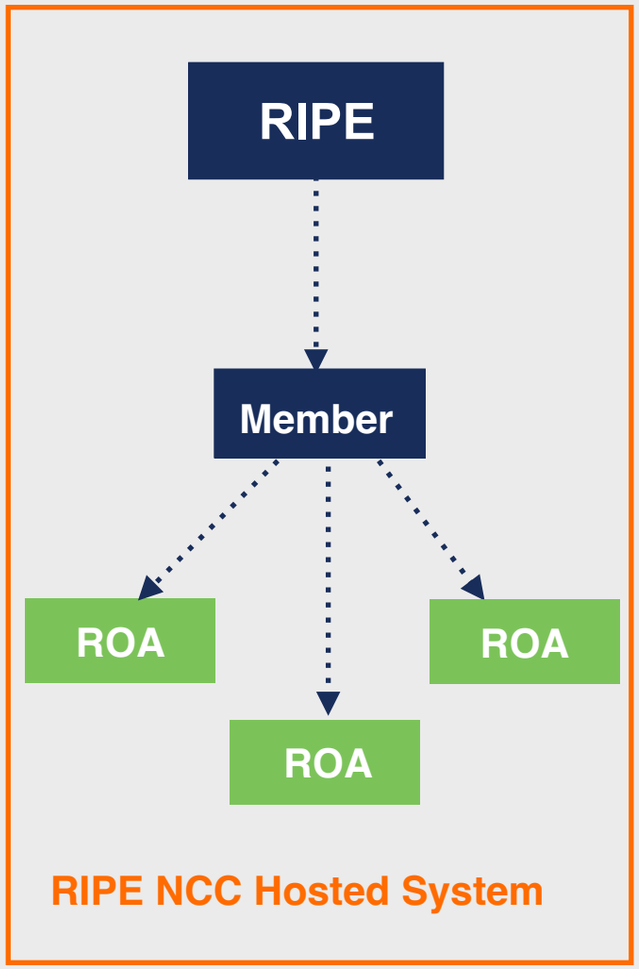

Delegated RPKI

- Each LIR manages its part of the RPKI system
  - Runs its own CA as a child of the RIR
  - Manages keys/key rollovers
  - Creates ROAs in its own platform
  - Signs and publishes ROAs

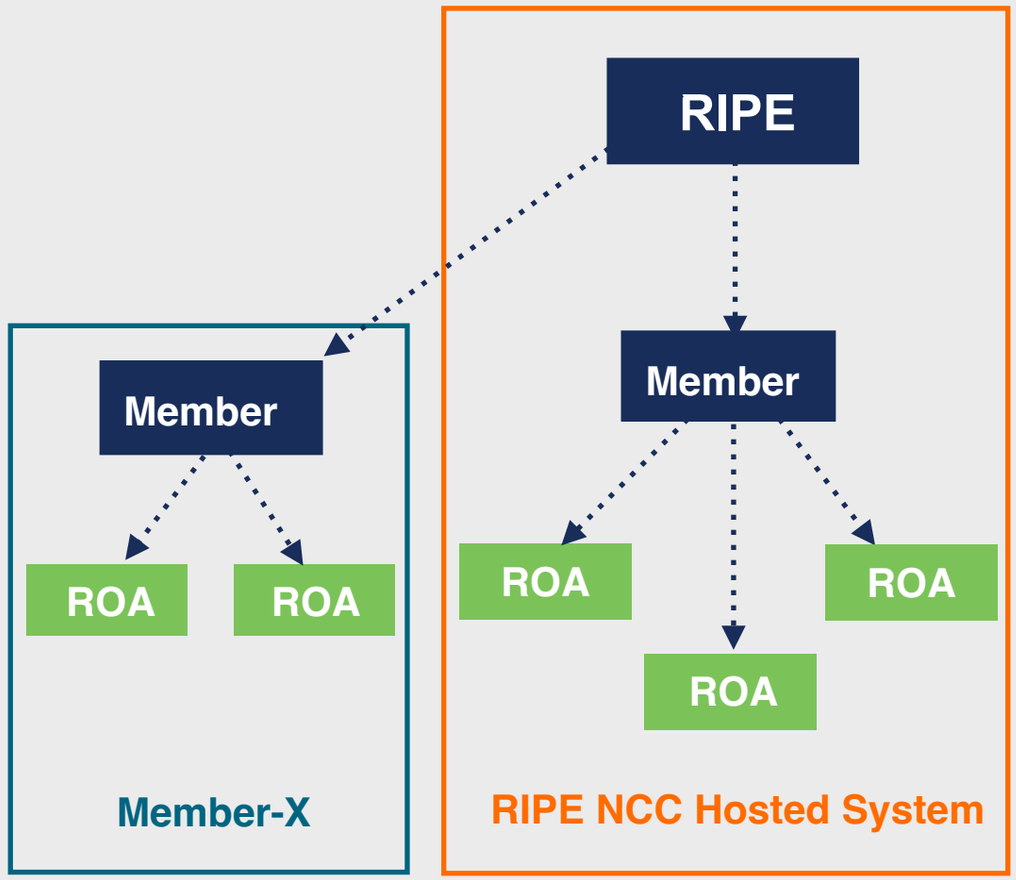

RIPE NCC Hosted Solution

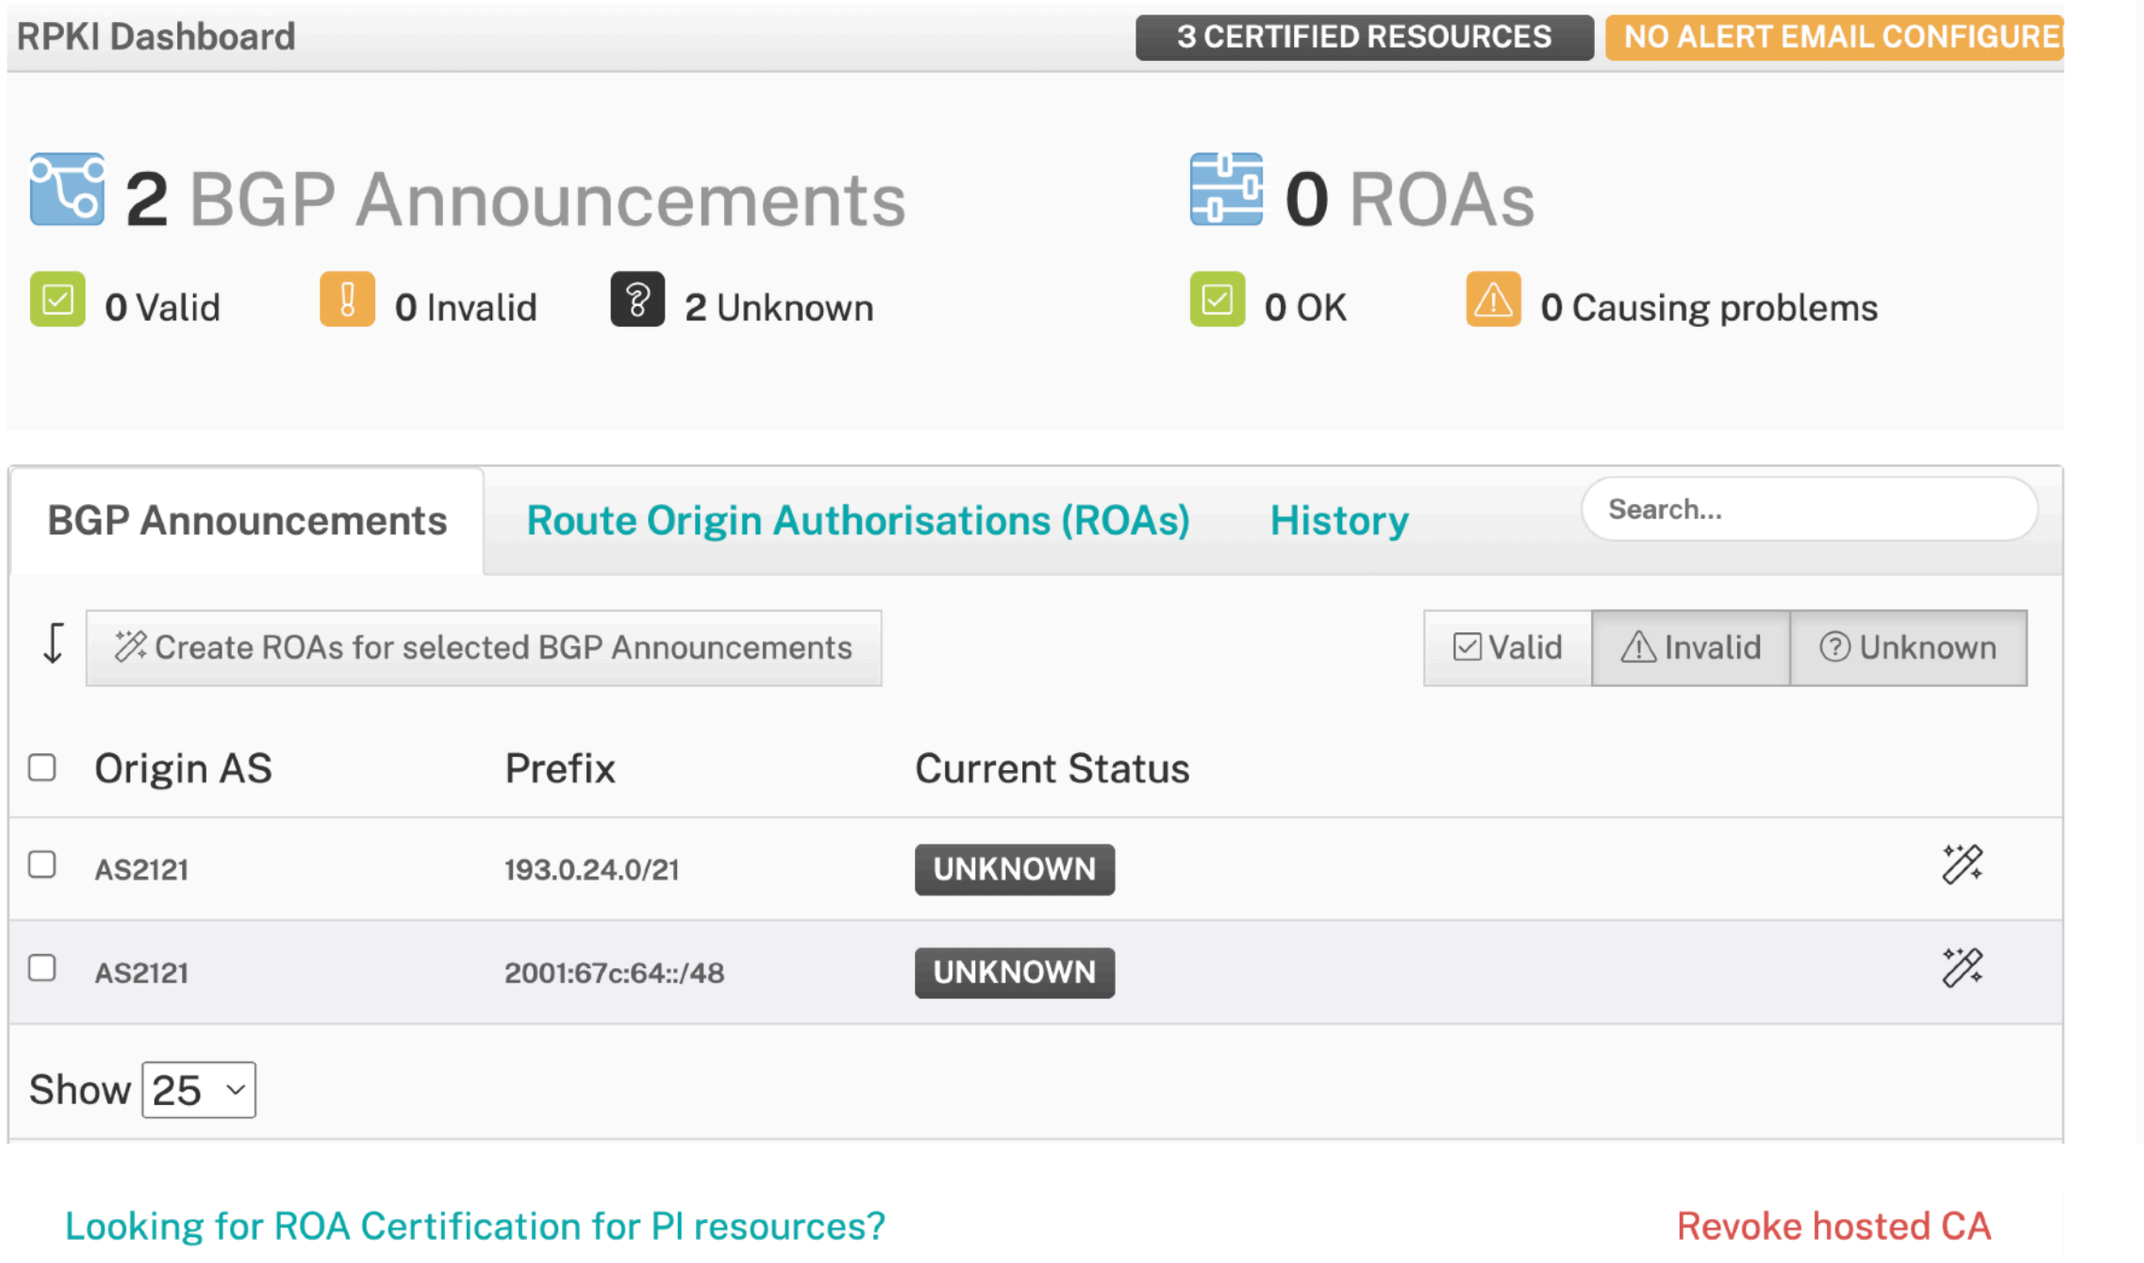

RIPE NCC Hosted Solution

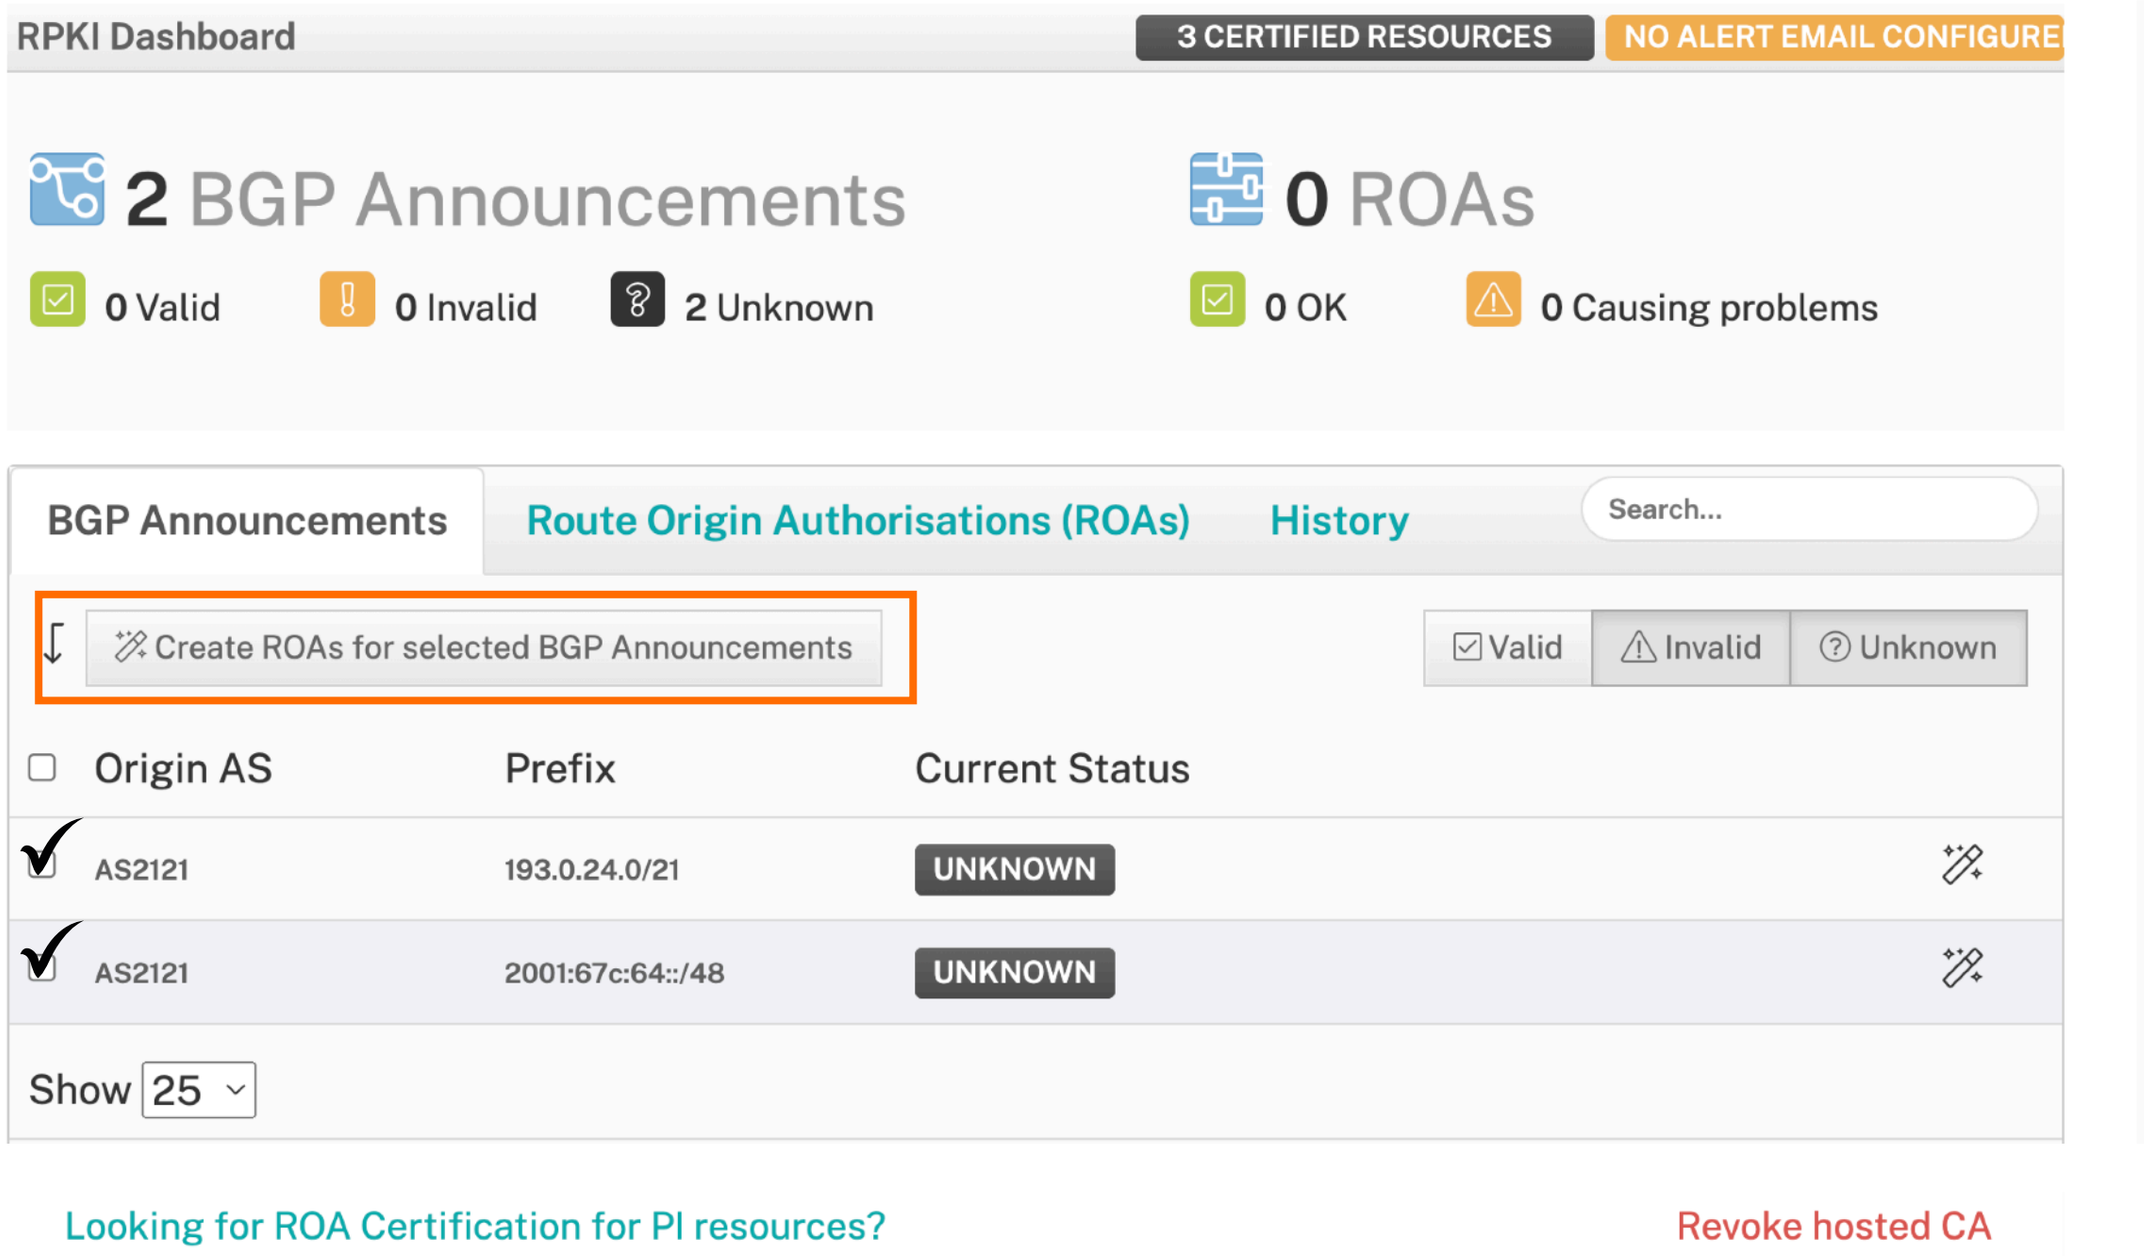

RIPE NCC Hosted Solution

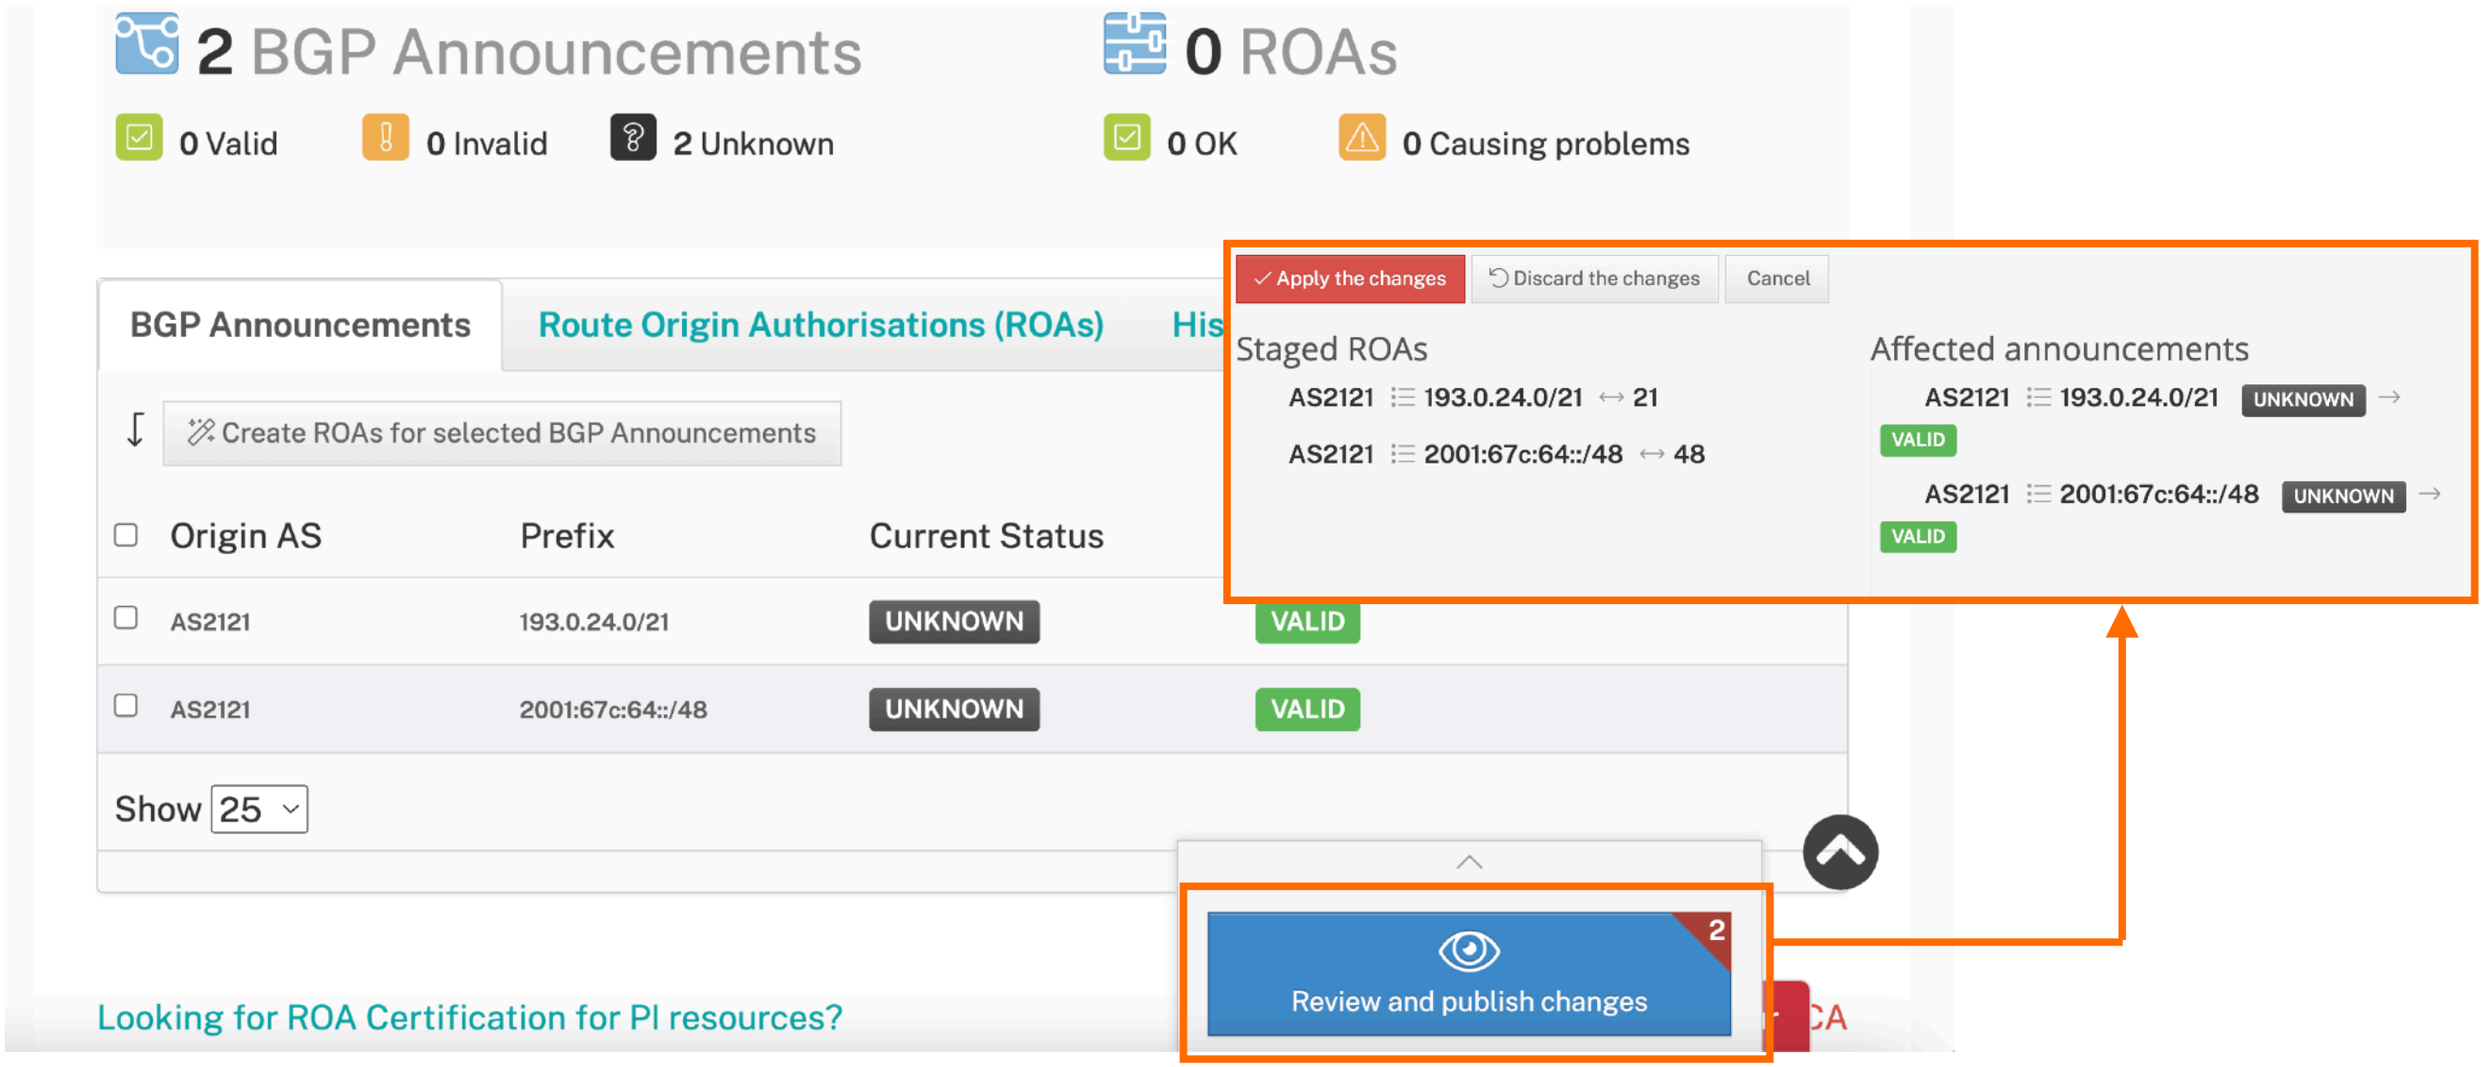

RIPE NCC Hosted Solution

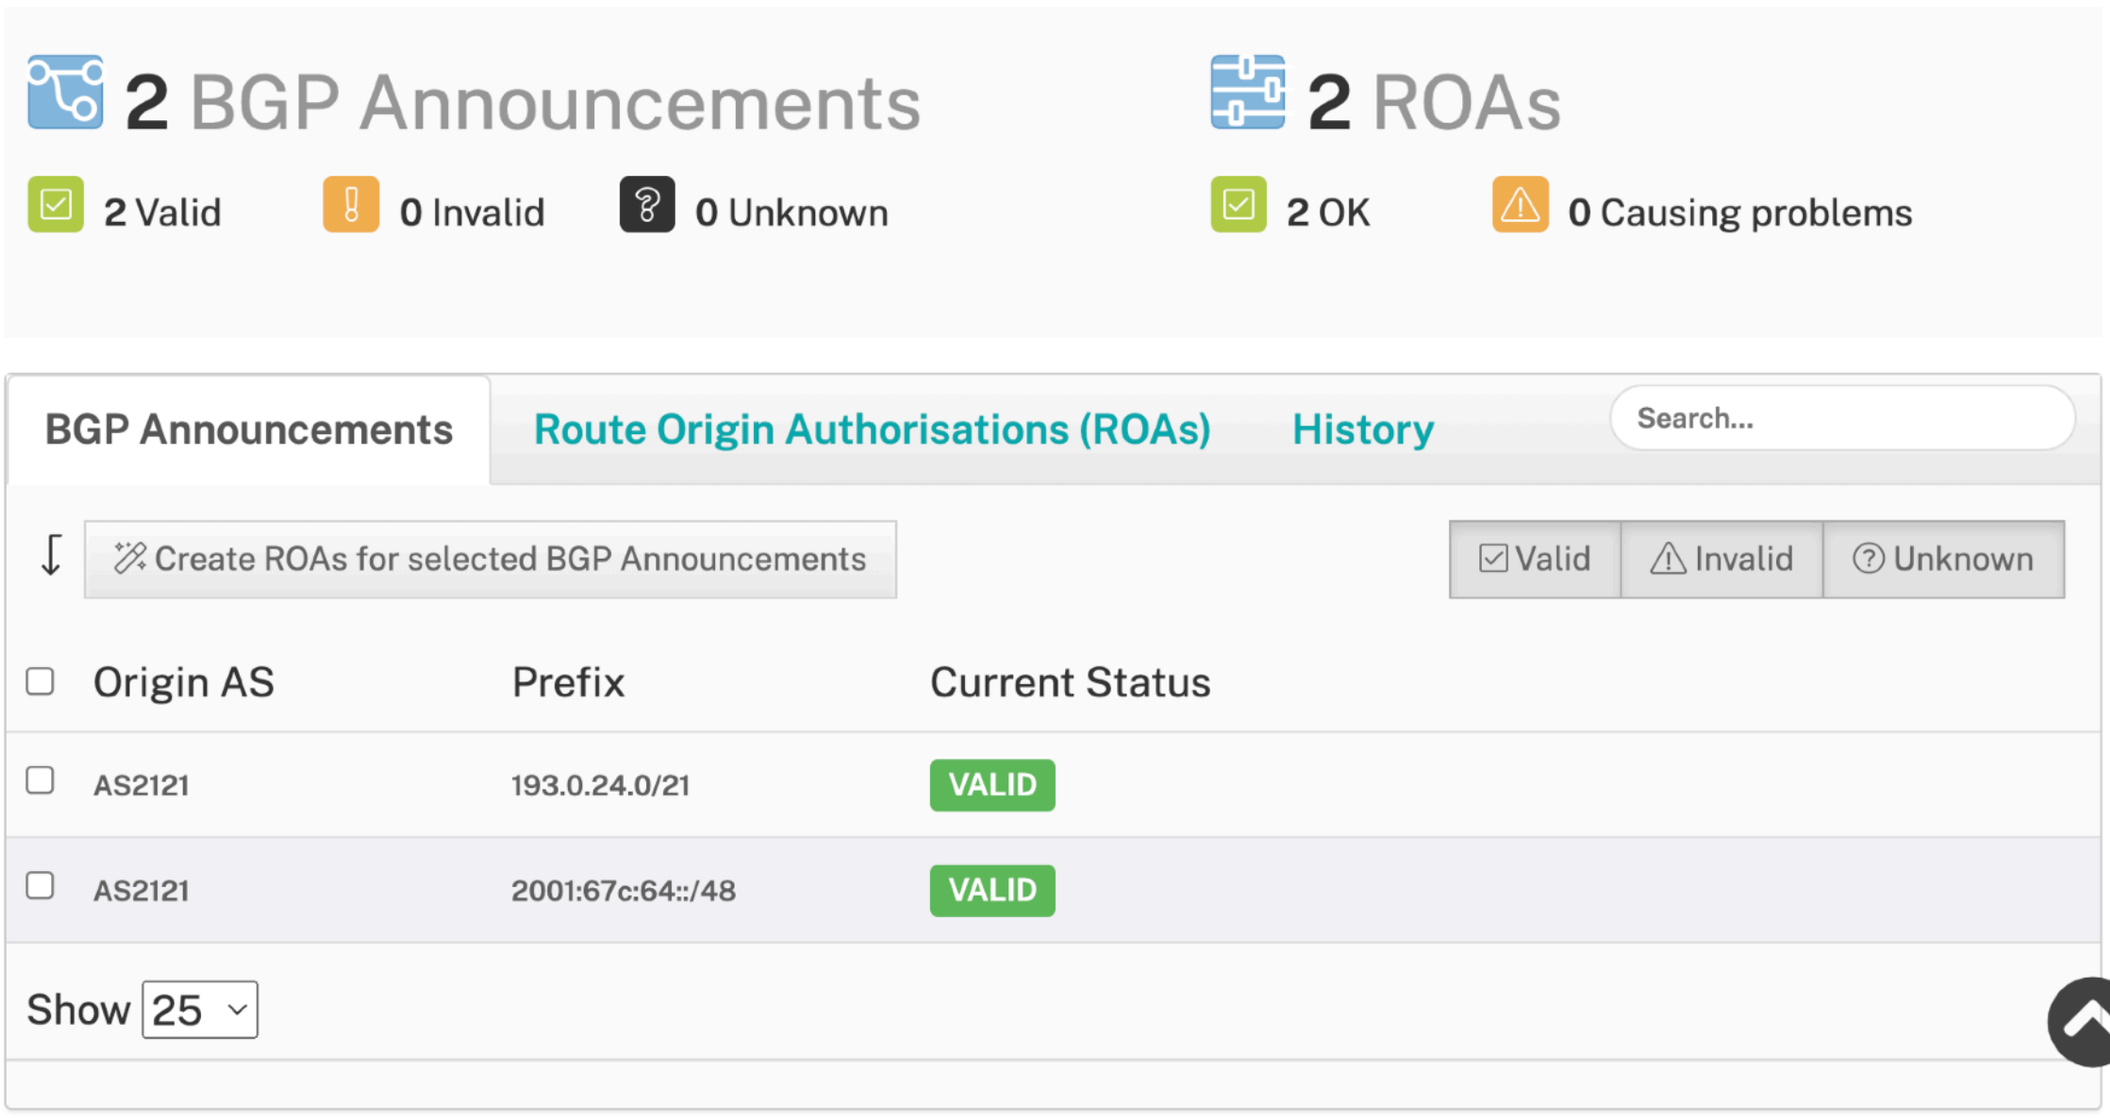

Routing Security with RPKI

RPKI Validators

Elements of RPKI

- RPKI system consists of two parts…

SIGNING

Create ROAs for your prefixes in the RPKI system

+

VALIDATION

Verify the information provided by others

RPKI Validation

- Verifying the information provided by others
  - Proves holdership through a public key and certificate infrastructure
- In order to validate RPKI data, you need to …
  - install a validator software locally in your network
- Goal is to validate the “origin of BGP announcements”
  - Known as BGP Origin Validation (BGP OV) or Route Origin Validation (ROV)

RPKI Validators

- Also known as Relying Party Software
- Connects to RPKI repositories via rsync or RRDP protocol
- Checks the information in TALs to connect to the repositories

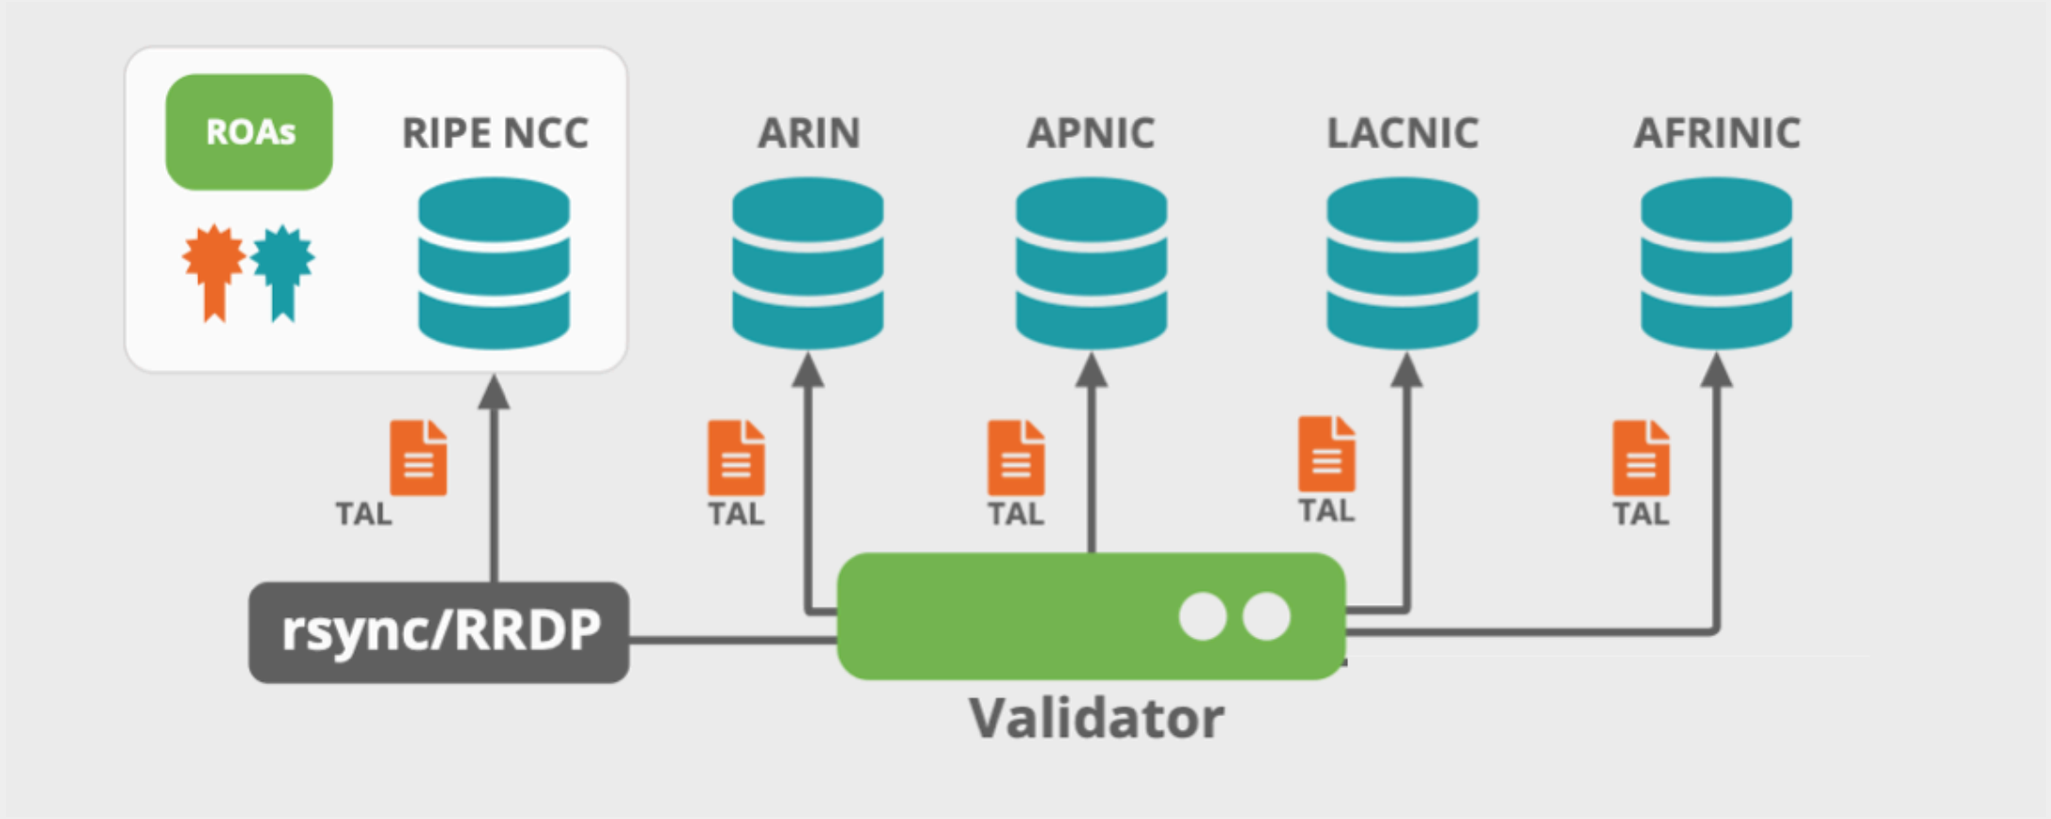

RPKI Validators

- Validator
  - Downloads all ROAs from RPKI repositories (from RIRs and external repos)
  - Validates the chain of trust for all ROAs and associated CAs
  - Creates a local “validated cache” with all the valid ROAs

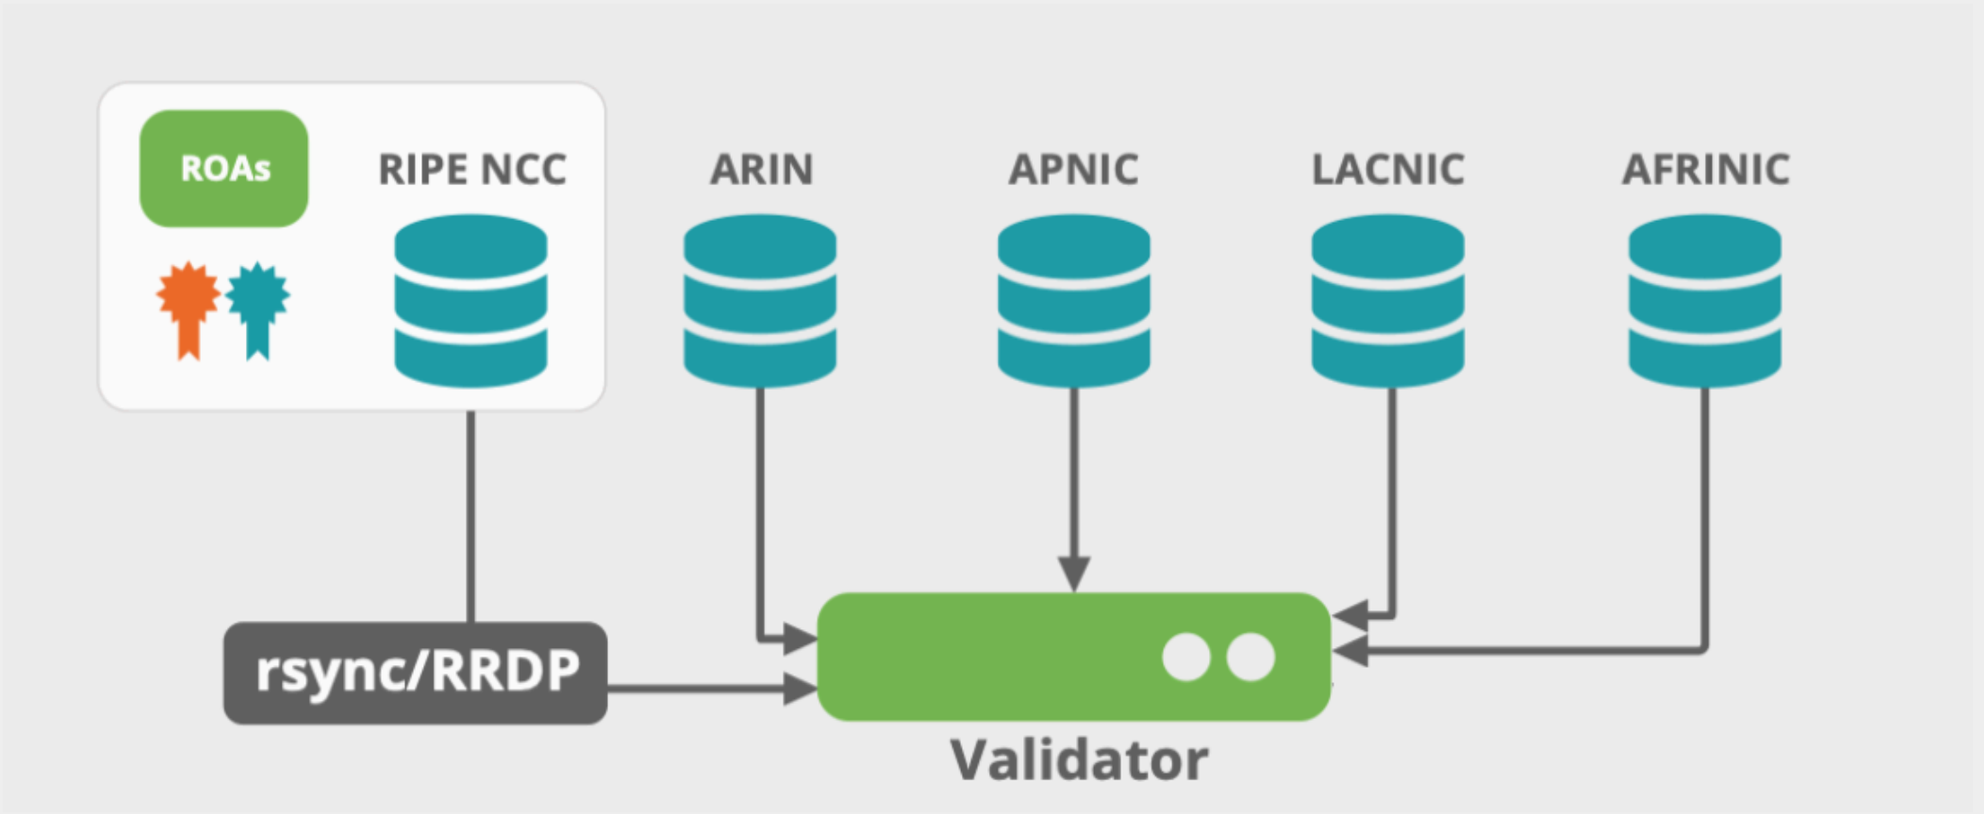

ROA Validation Process

ROA Validation Process

ROA Validation Process

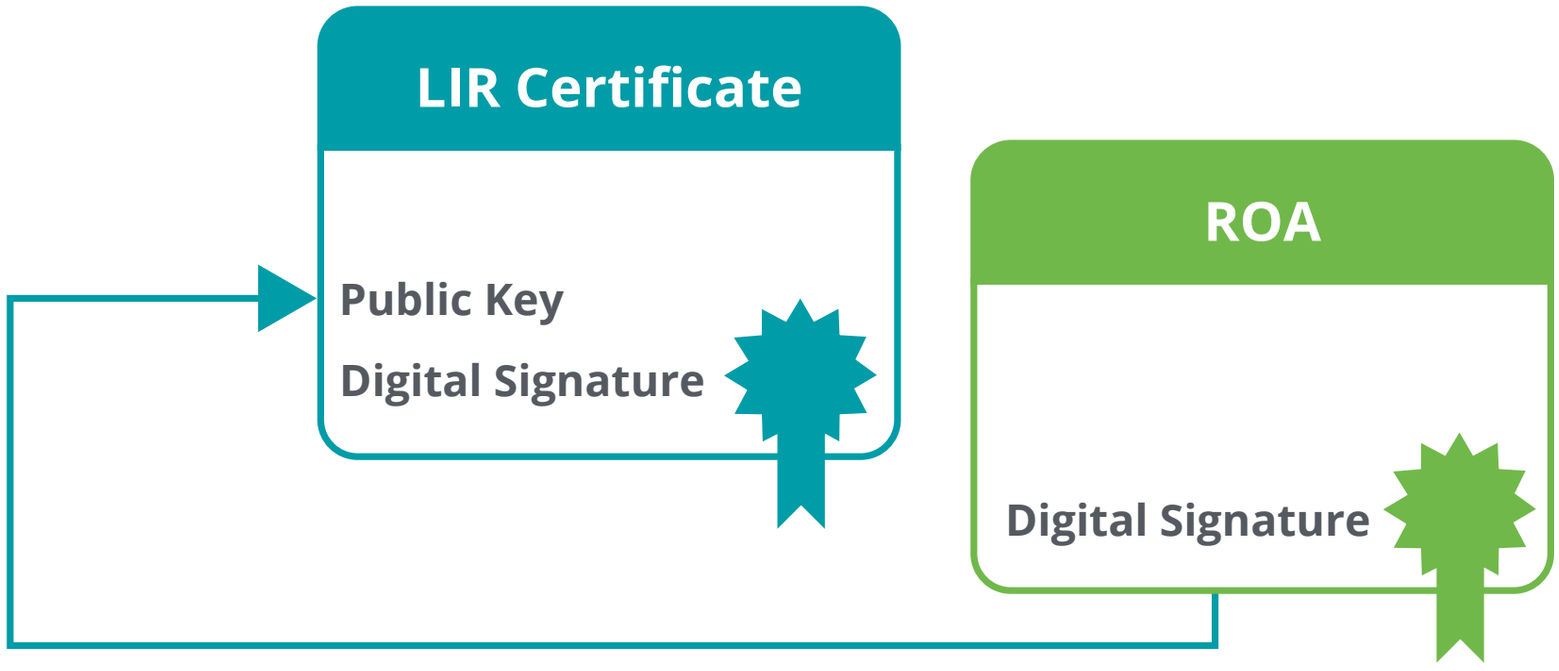

ROA Validation Process

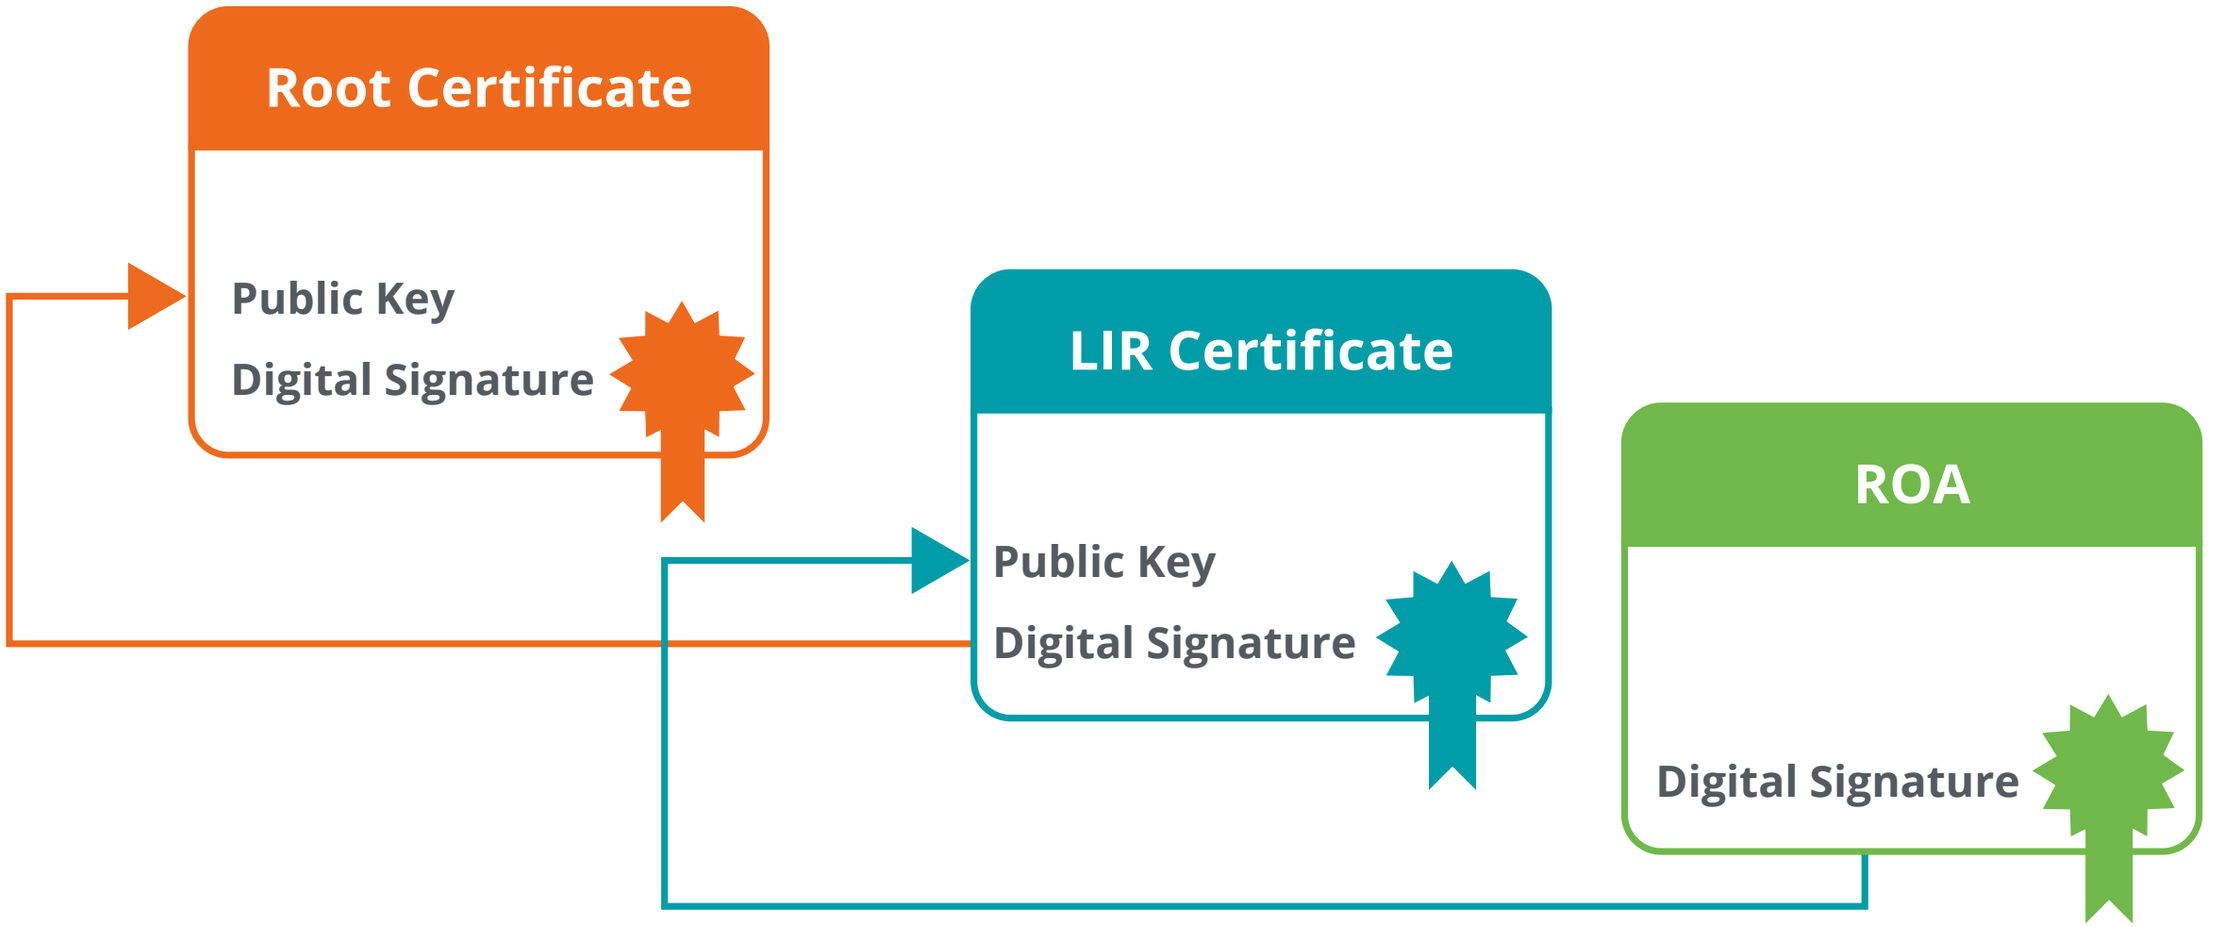

ROA Validation Process

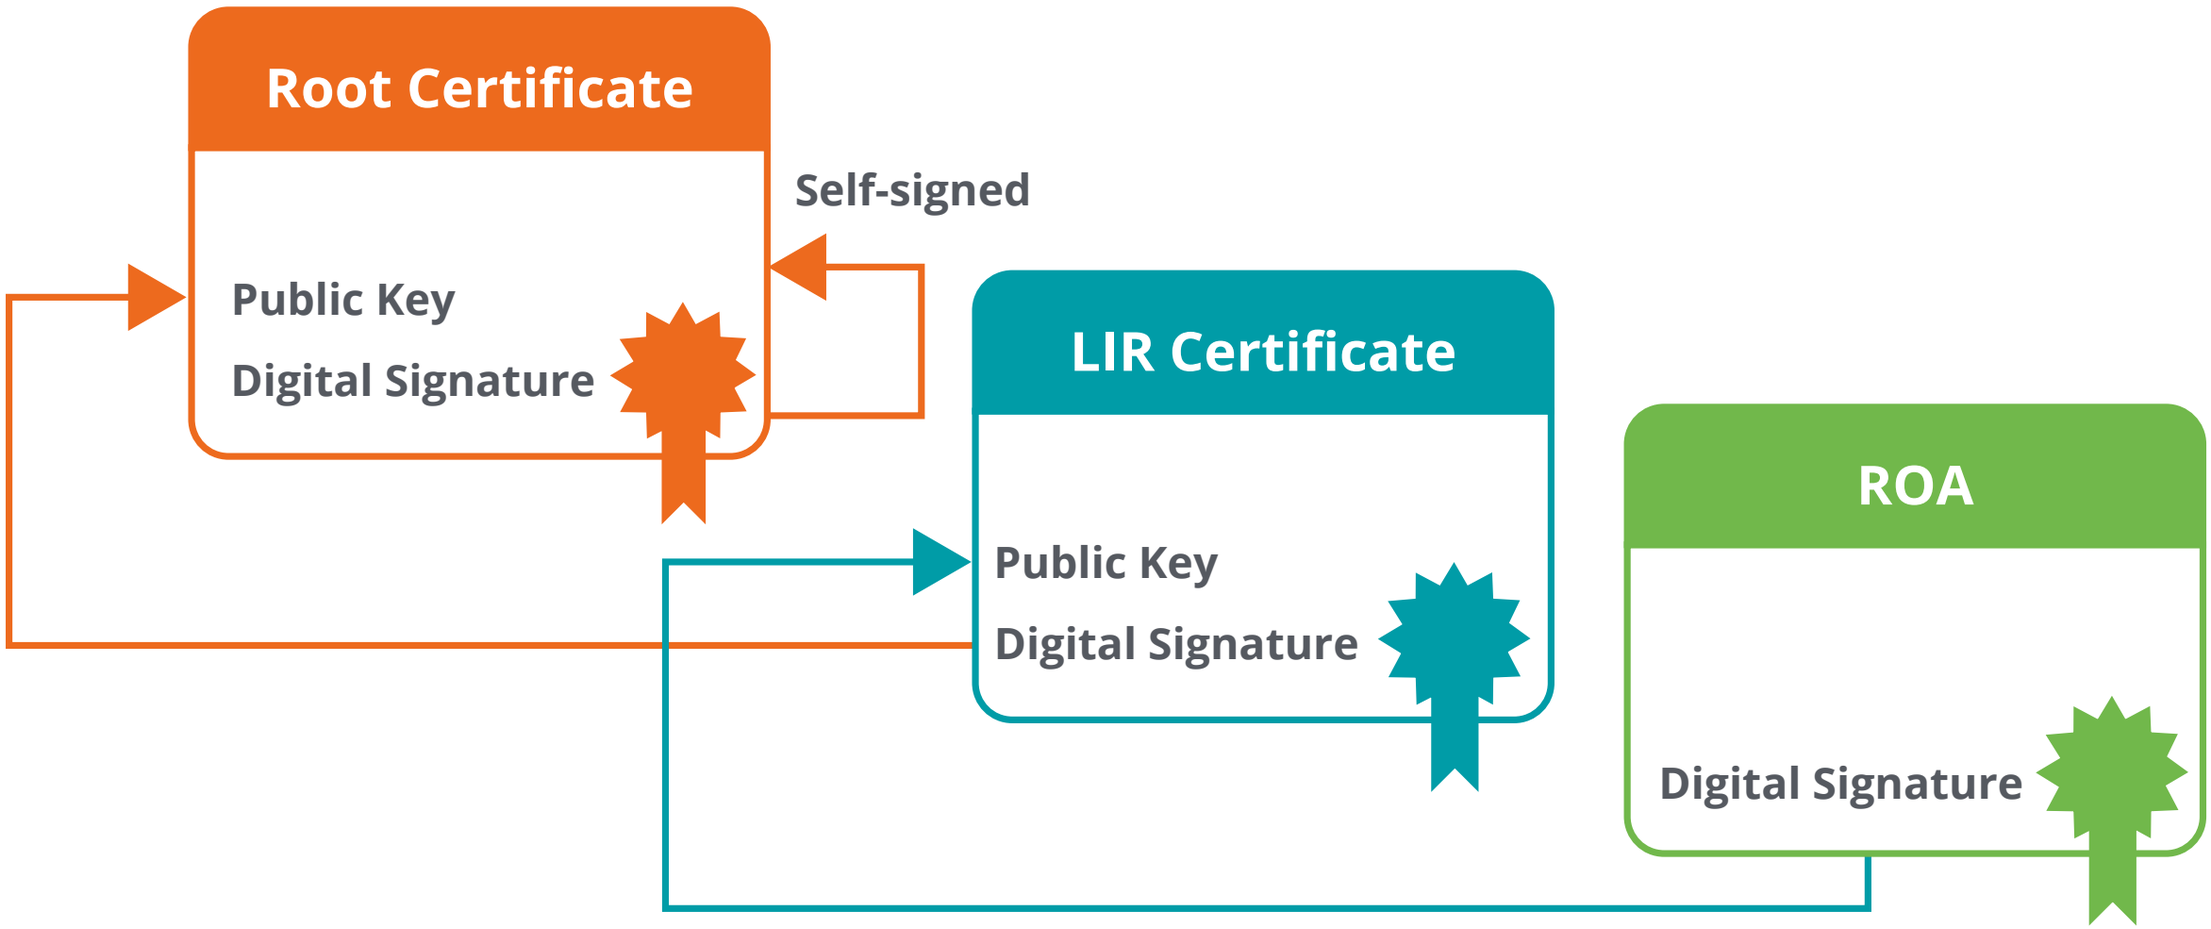

ROA Validation Process

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

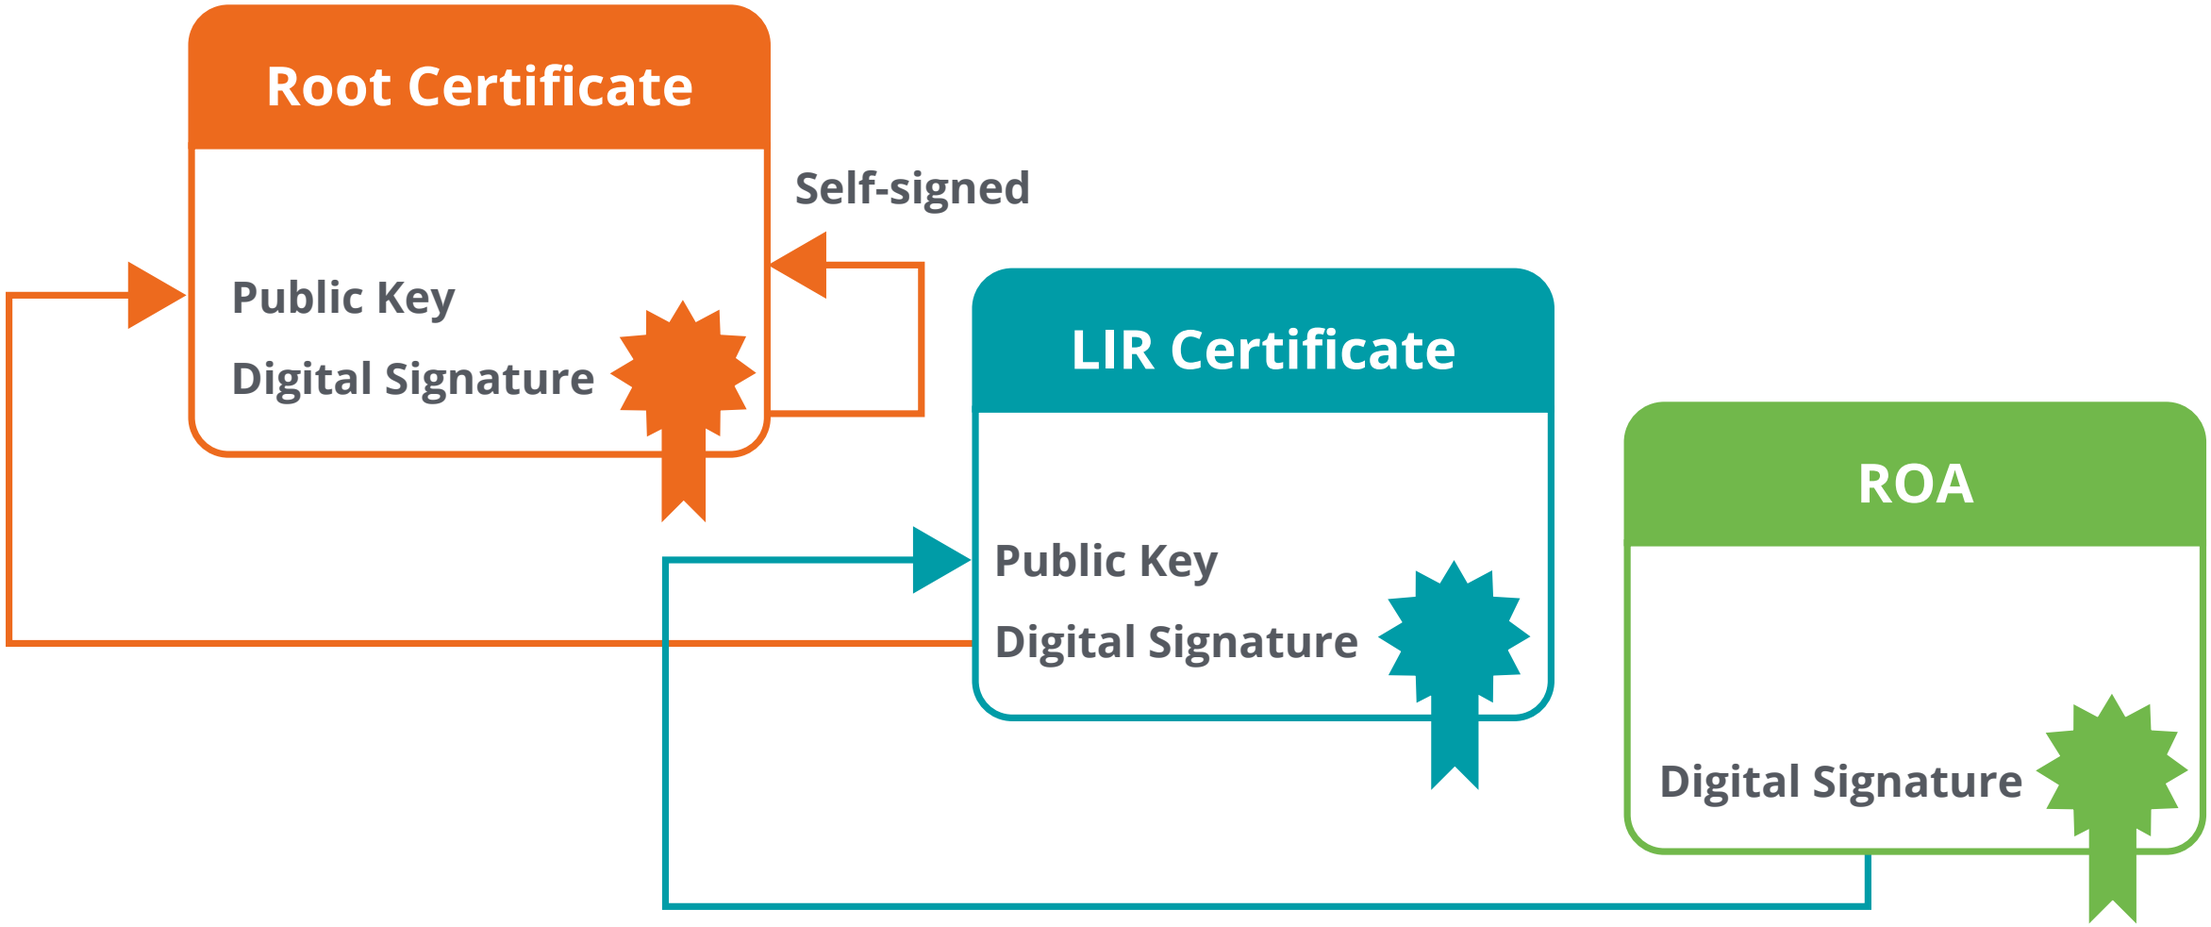

ROA Validation Process

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

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

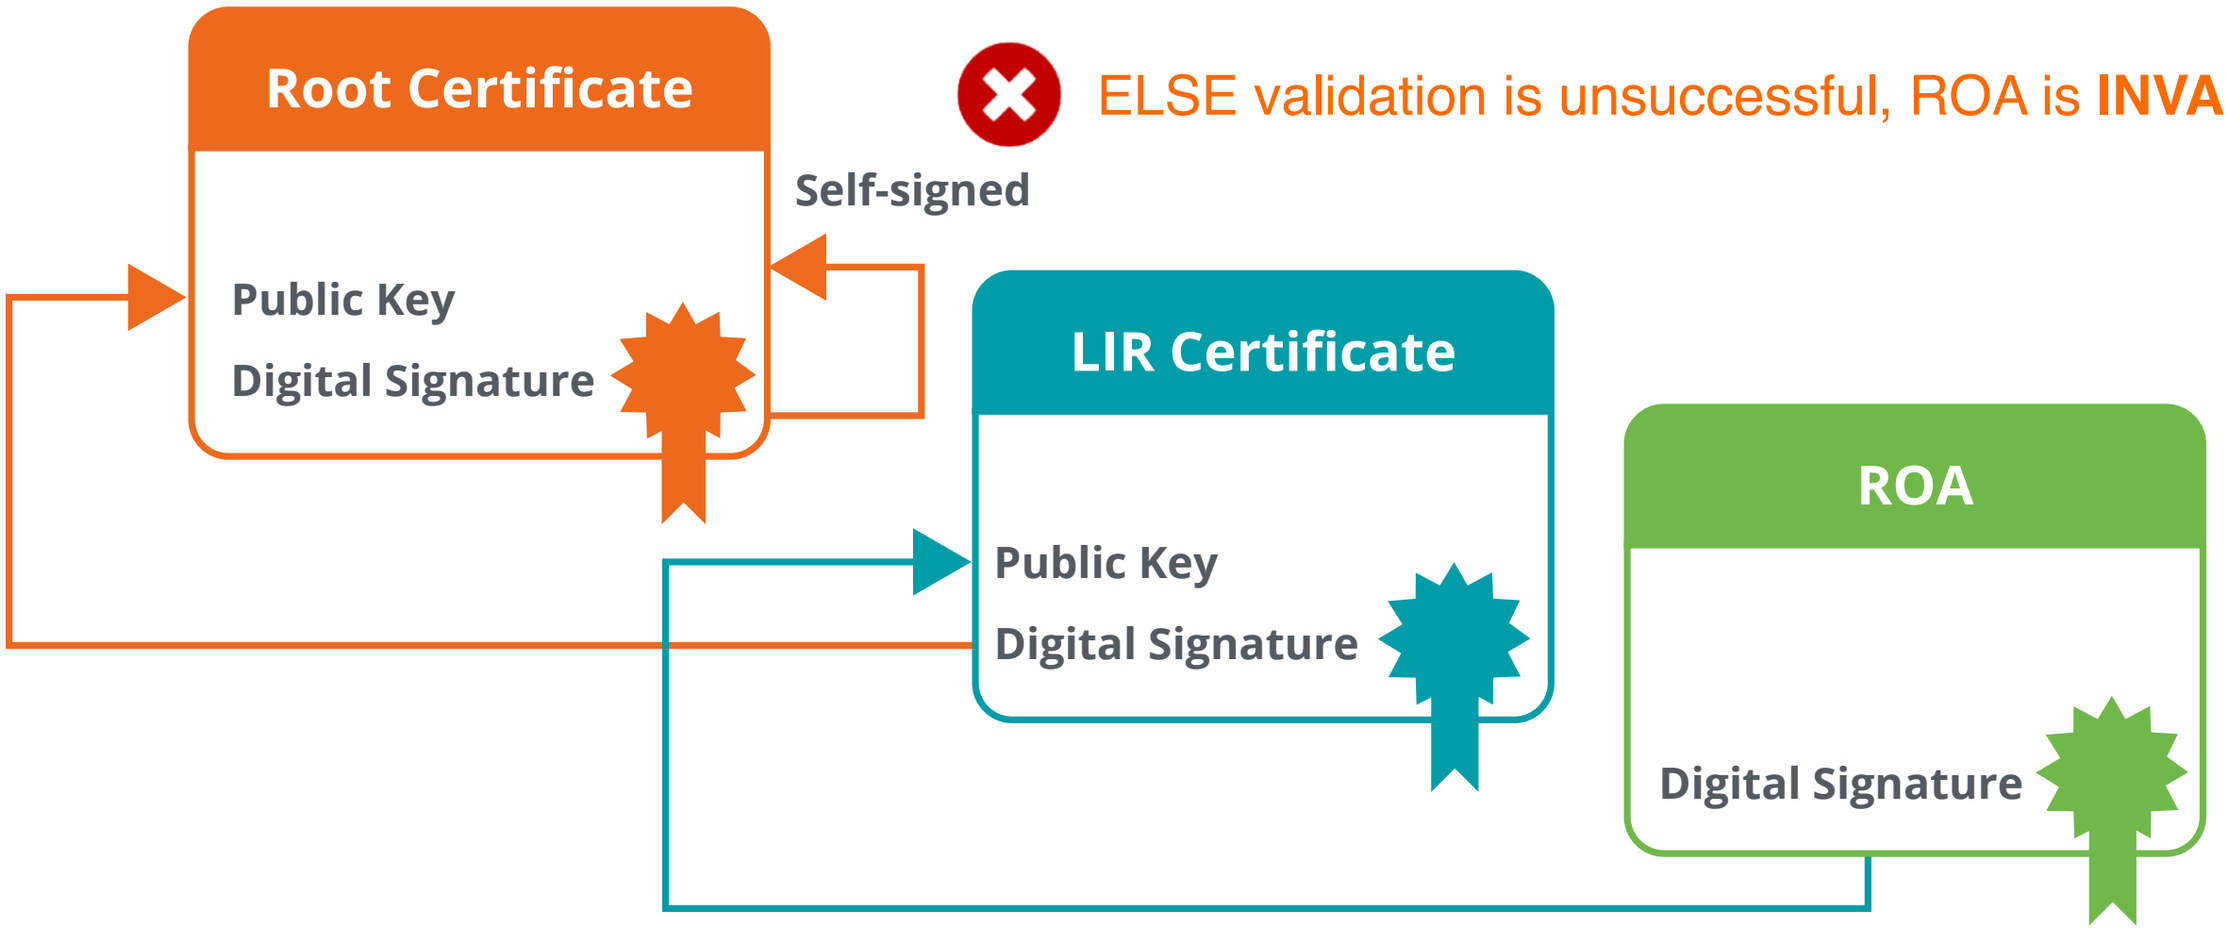

Valid ROAs are sent to the router!

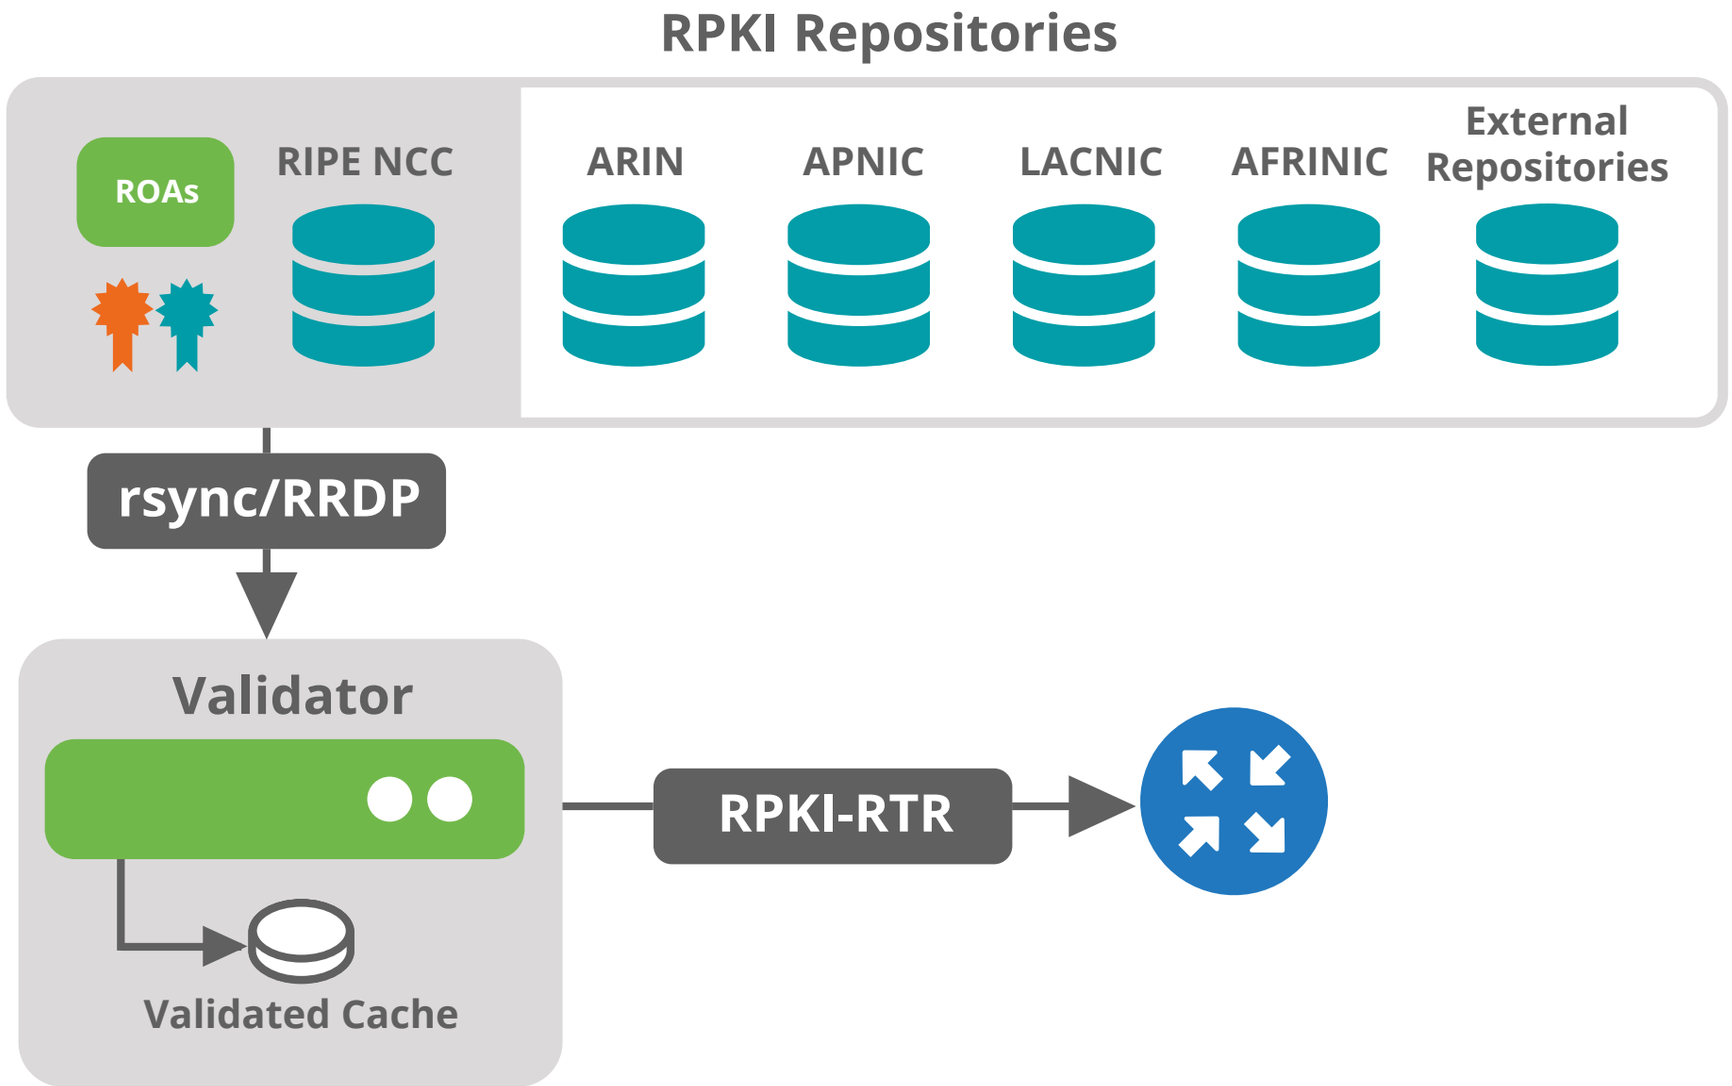

Valid ROAs are sent to the router!

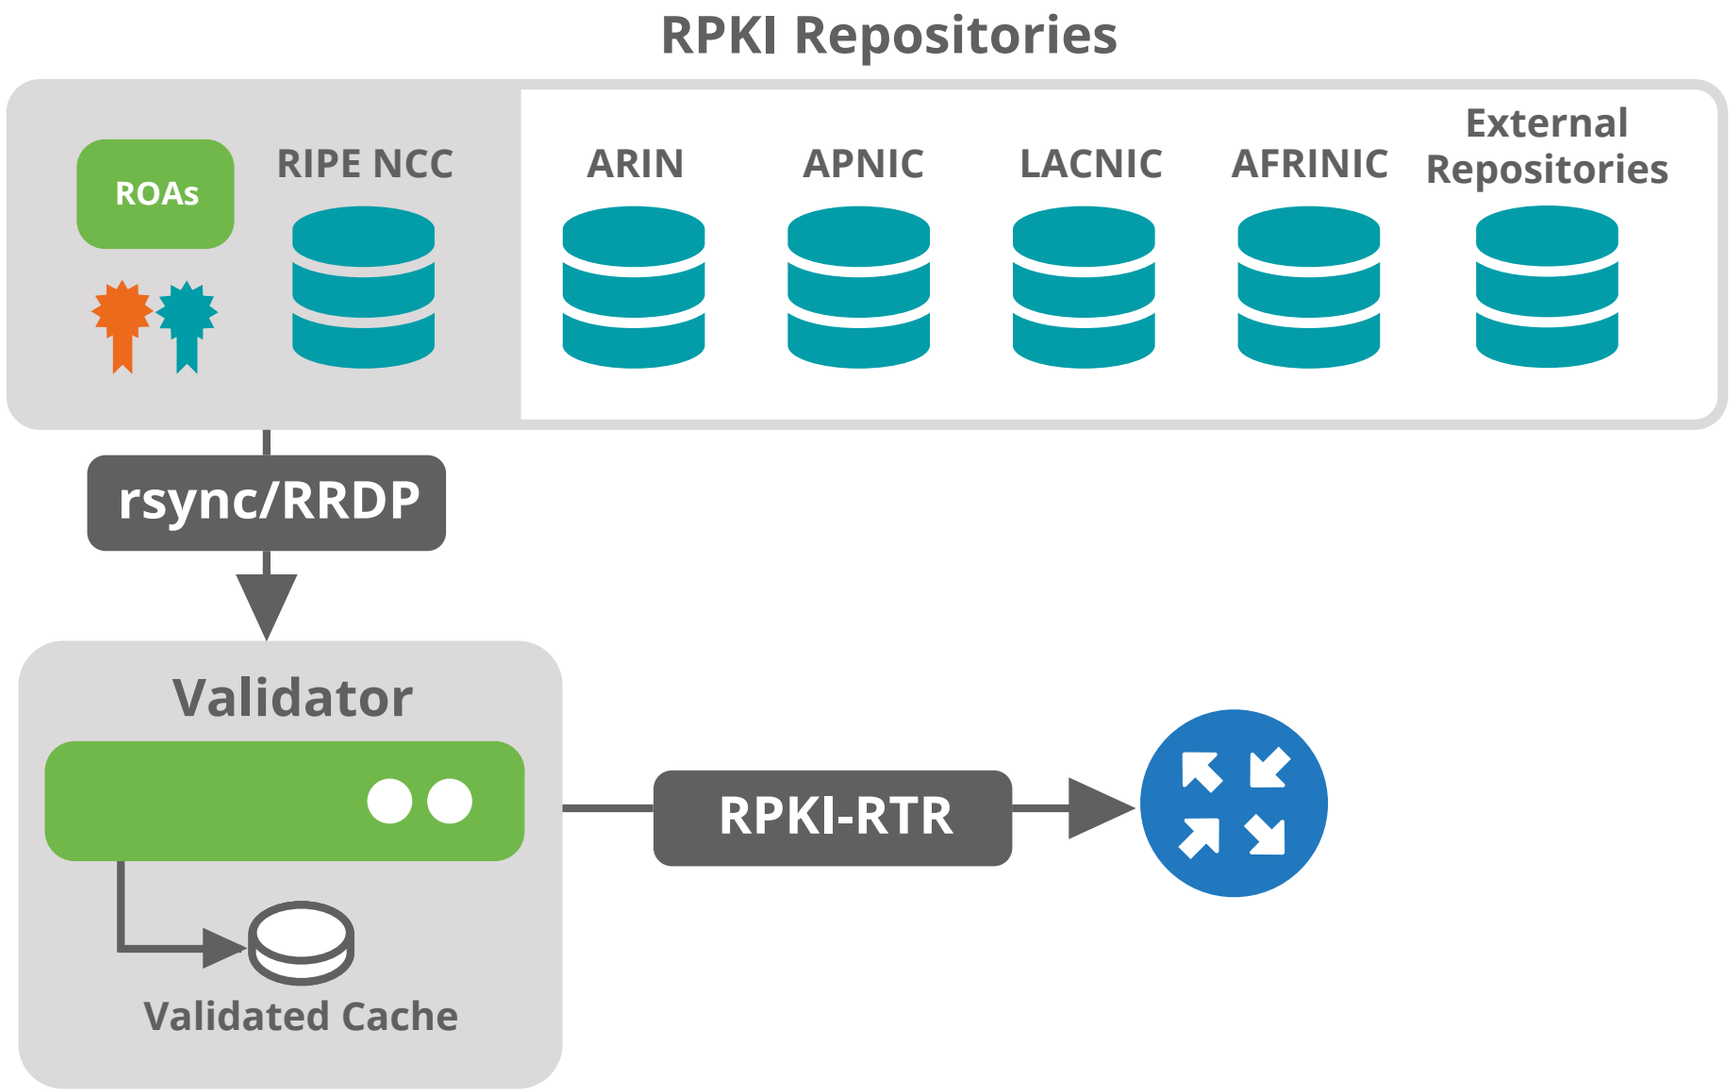

The router uses this information to make better routing decisions!

OR

RPKI Validator Options

- **Routinator**
  - Built by NLNetlabs
- **OctoRPKI**
  - Cloudflare’s relying party software
- **FORT**
  - Open source RPKI validator
- **rpki-client**
  - Integrated in OpenBsd
-

Links for RPKI Validators

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

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

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

https://www.rpki-client.org/

For more info…

https://rpki.readthedocs.io

Routing Security with RPKI

Validating BGP Announcements

BGP Origin Validation (BGP OV)

- RPKI based route filtering, RFC#6811
- BGP announcements are compared against the **valid** ROAs
  - **origin ASN** and **max-length** must match!
- Router decides the validation states of routes: Valid, Invalid and Not Found

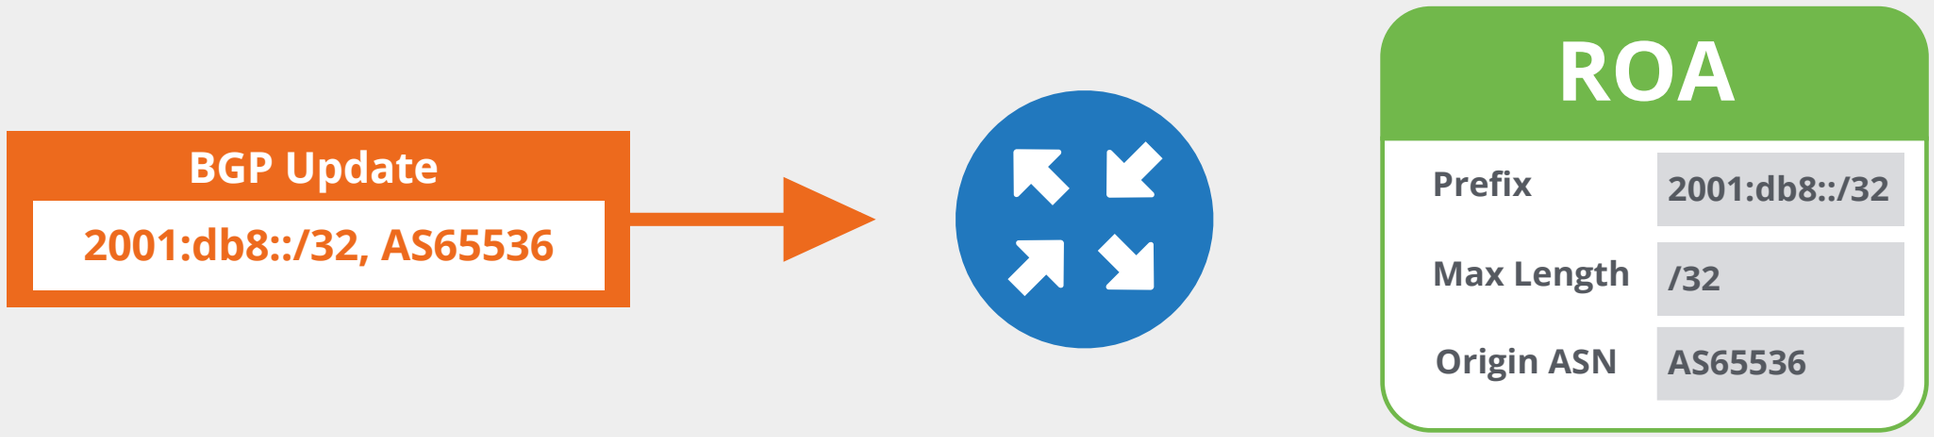

RFC#6811-BGP Prefix Origin Validation https://datatracker.ietf.org/doc/html/rfc6811

How does RPKI validate the origin?

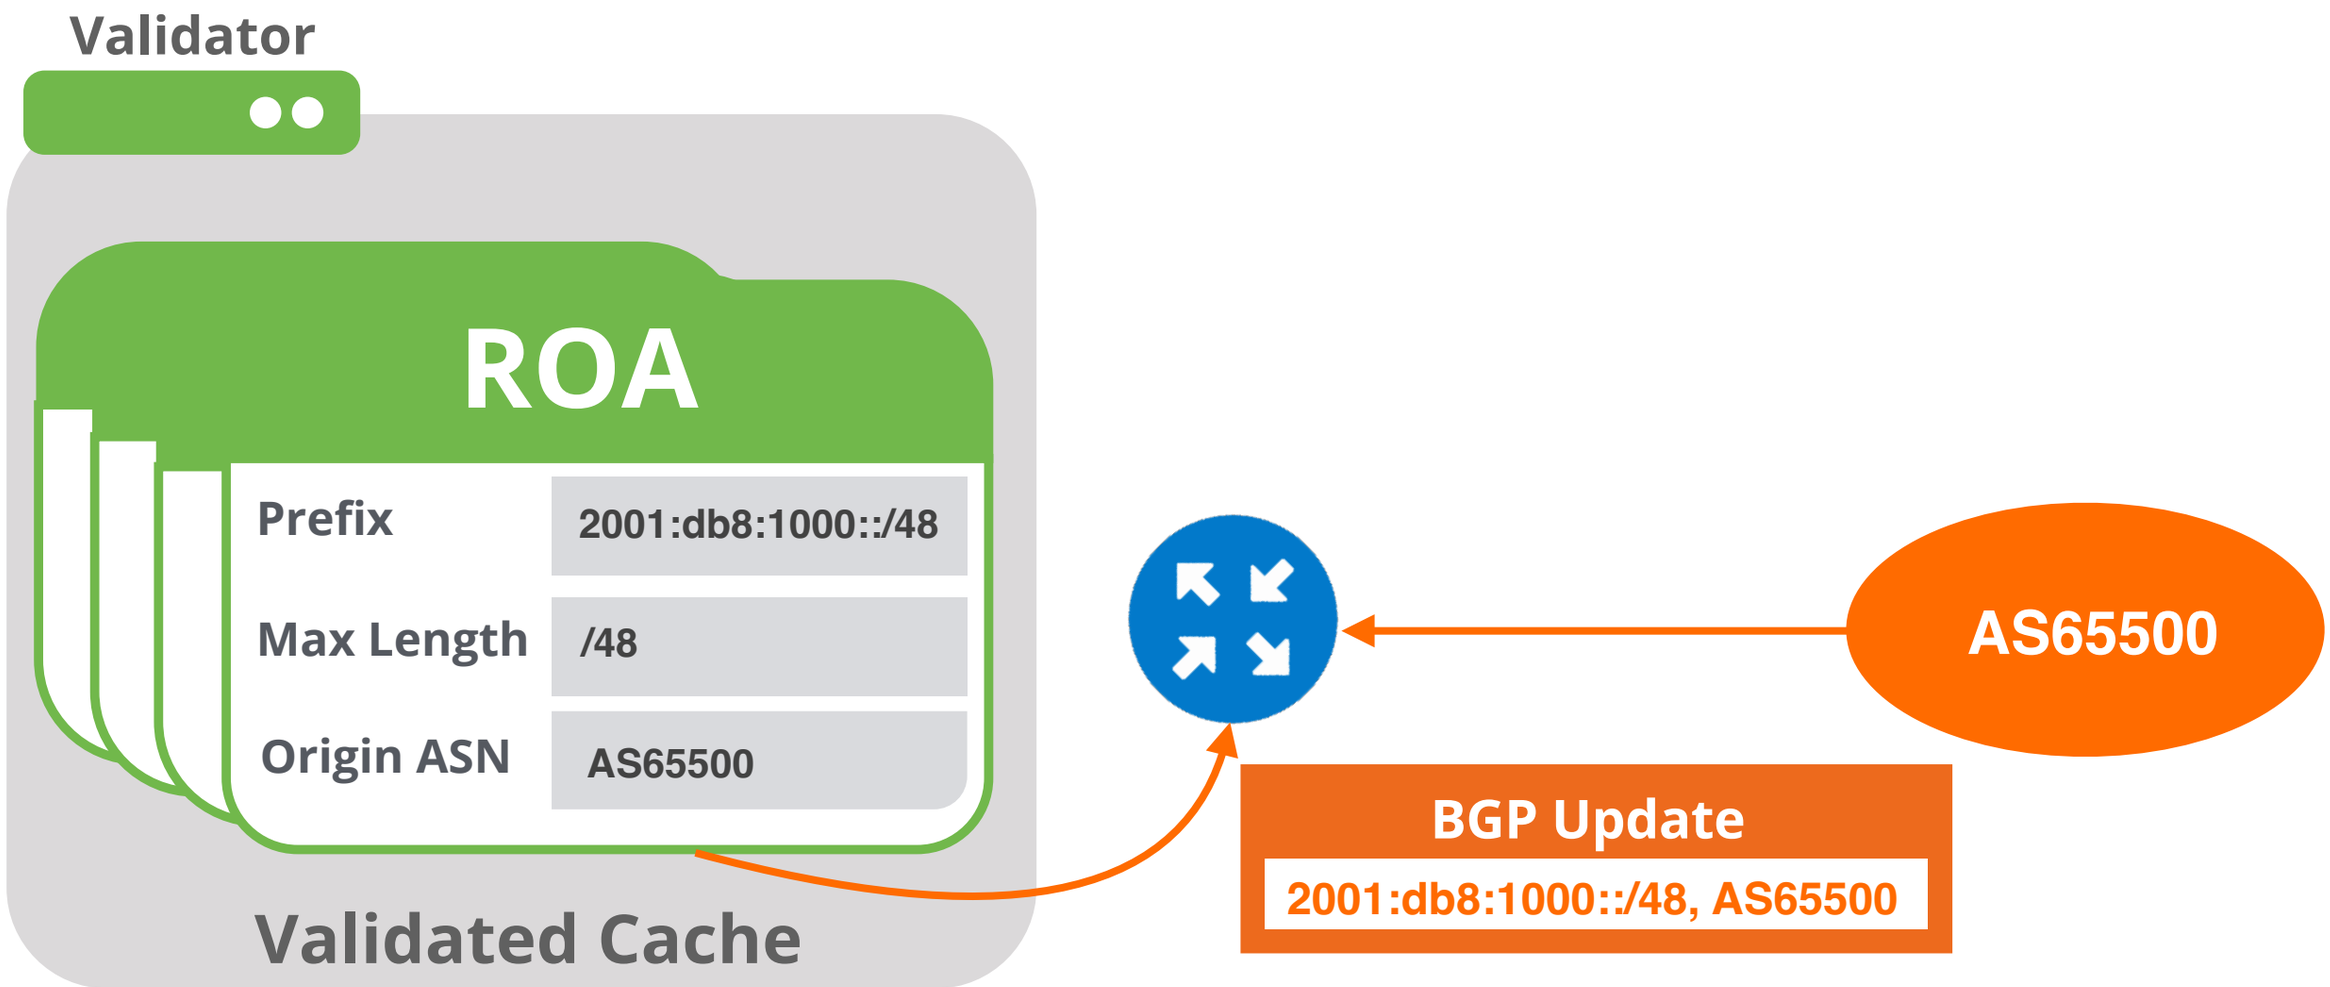

How does RPKI validate the origin?

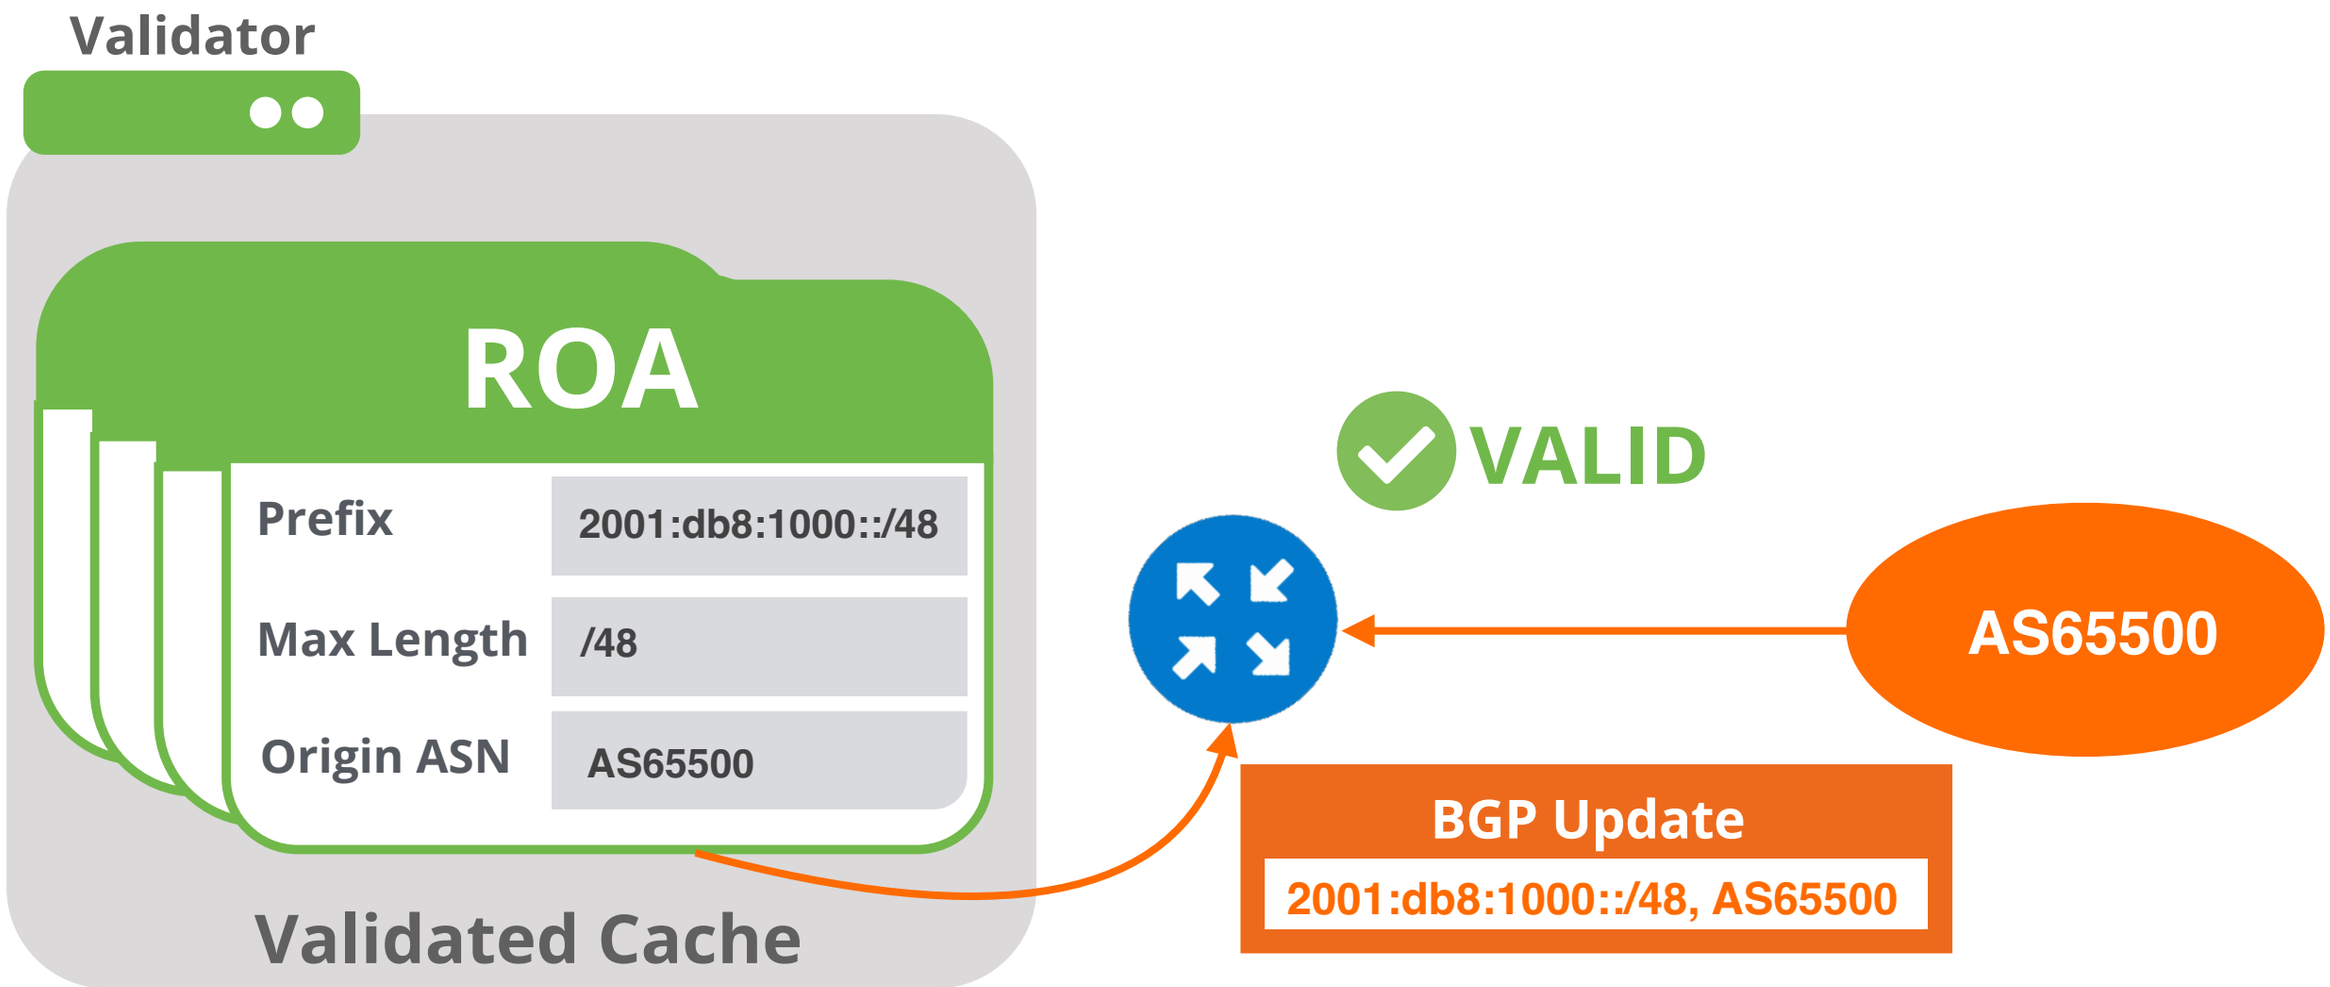

How does RPKI validate the origin?

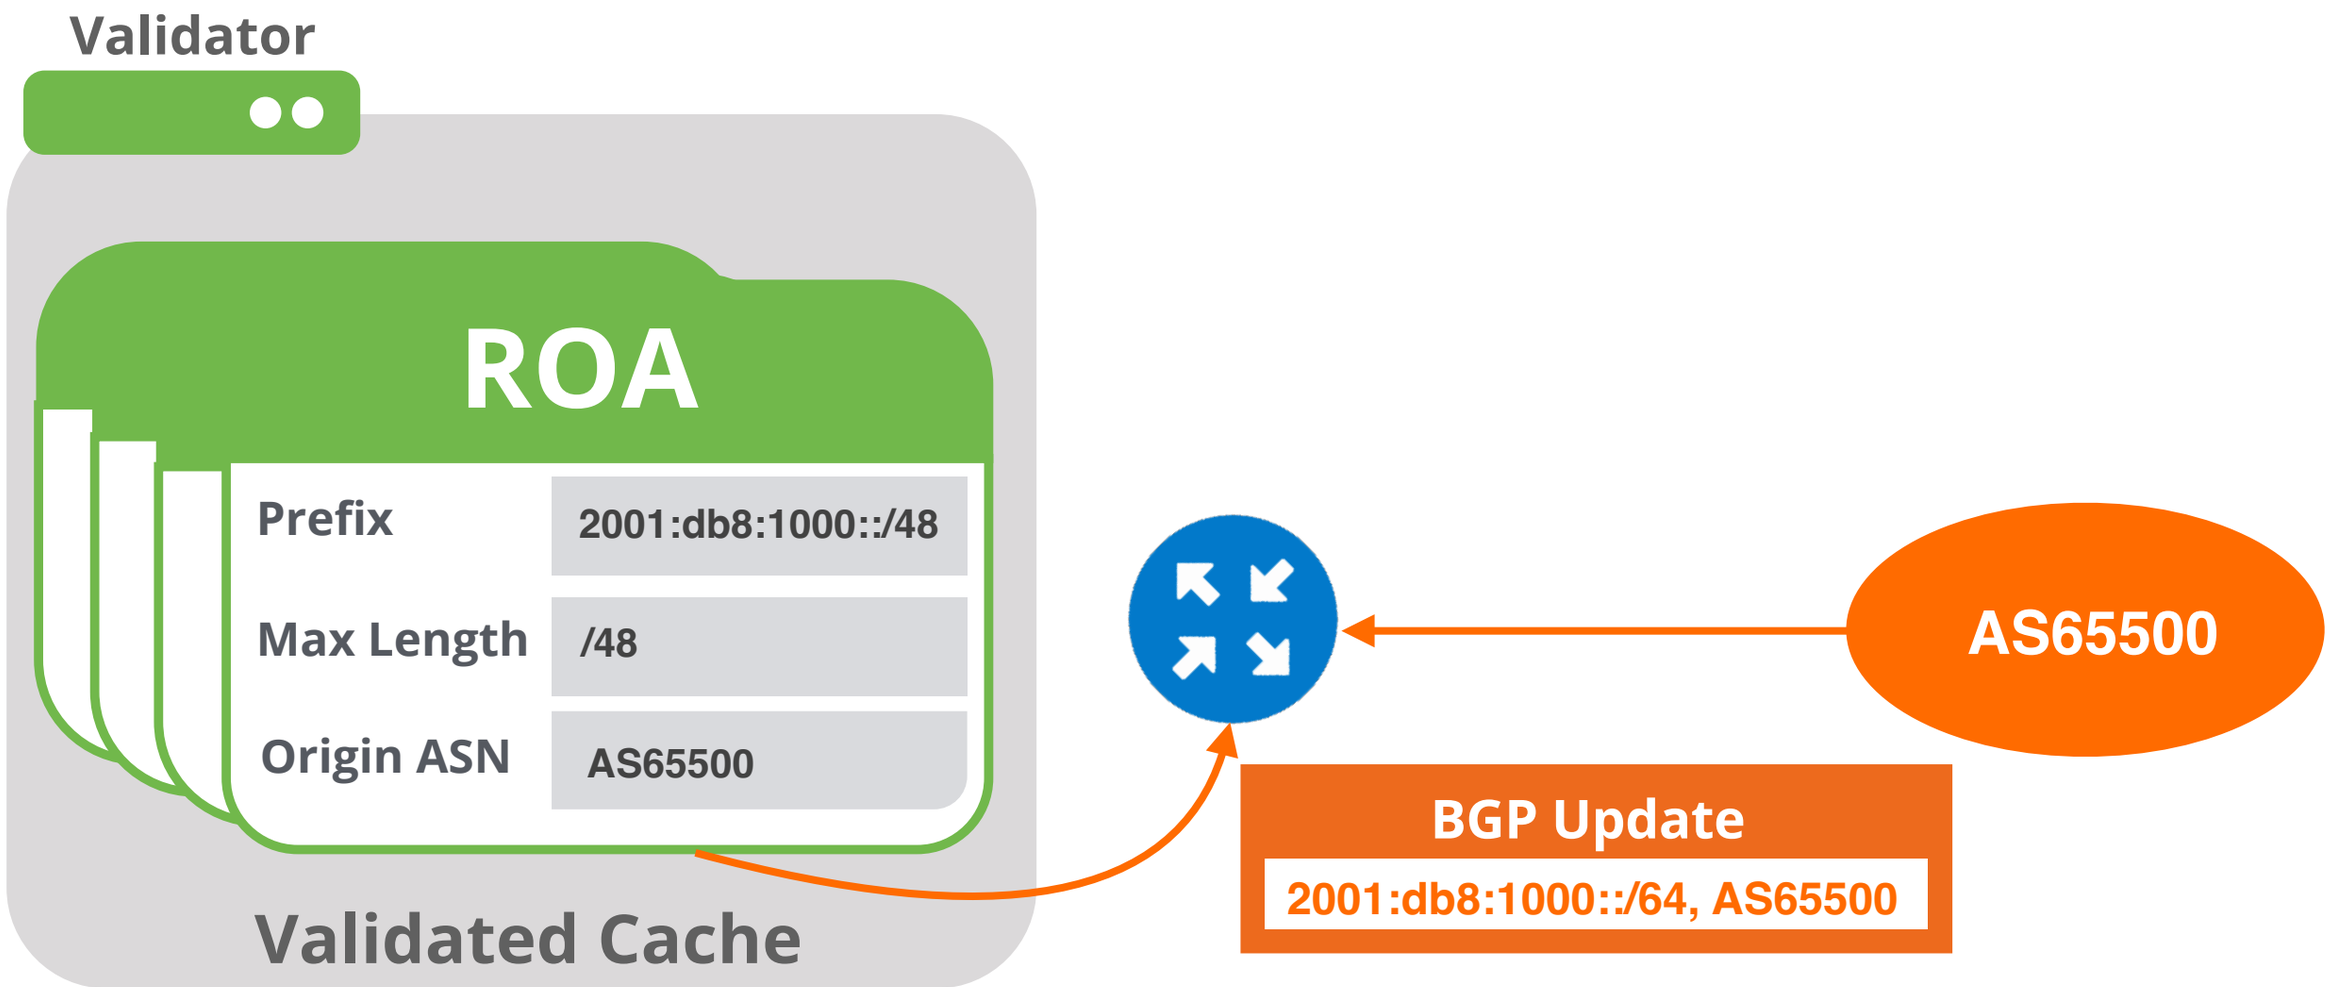

How does RPKI validate the origin?

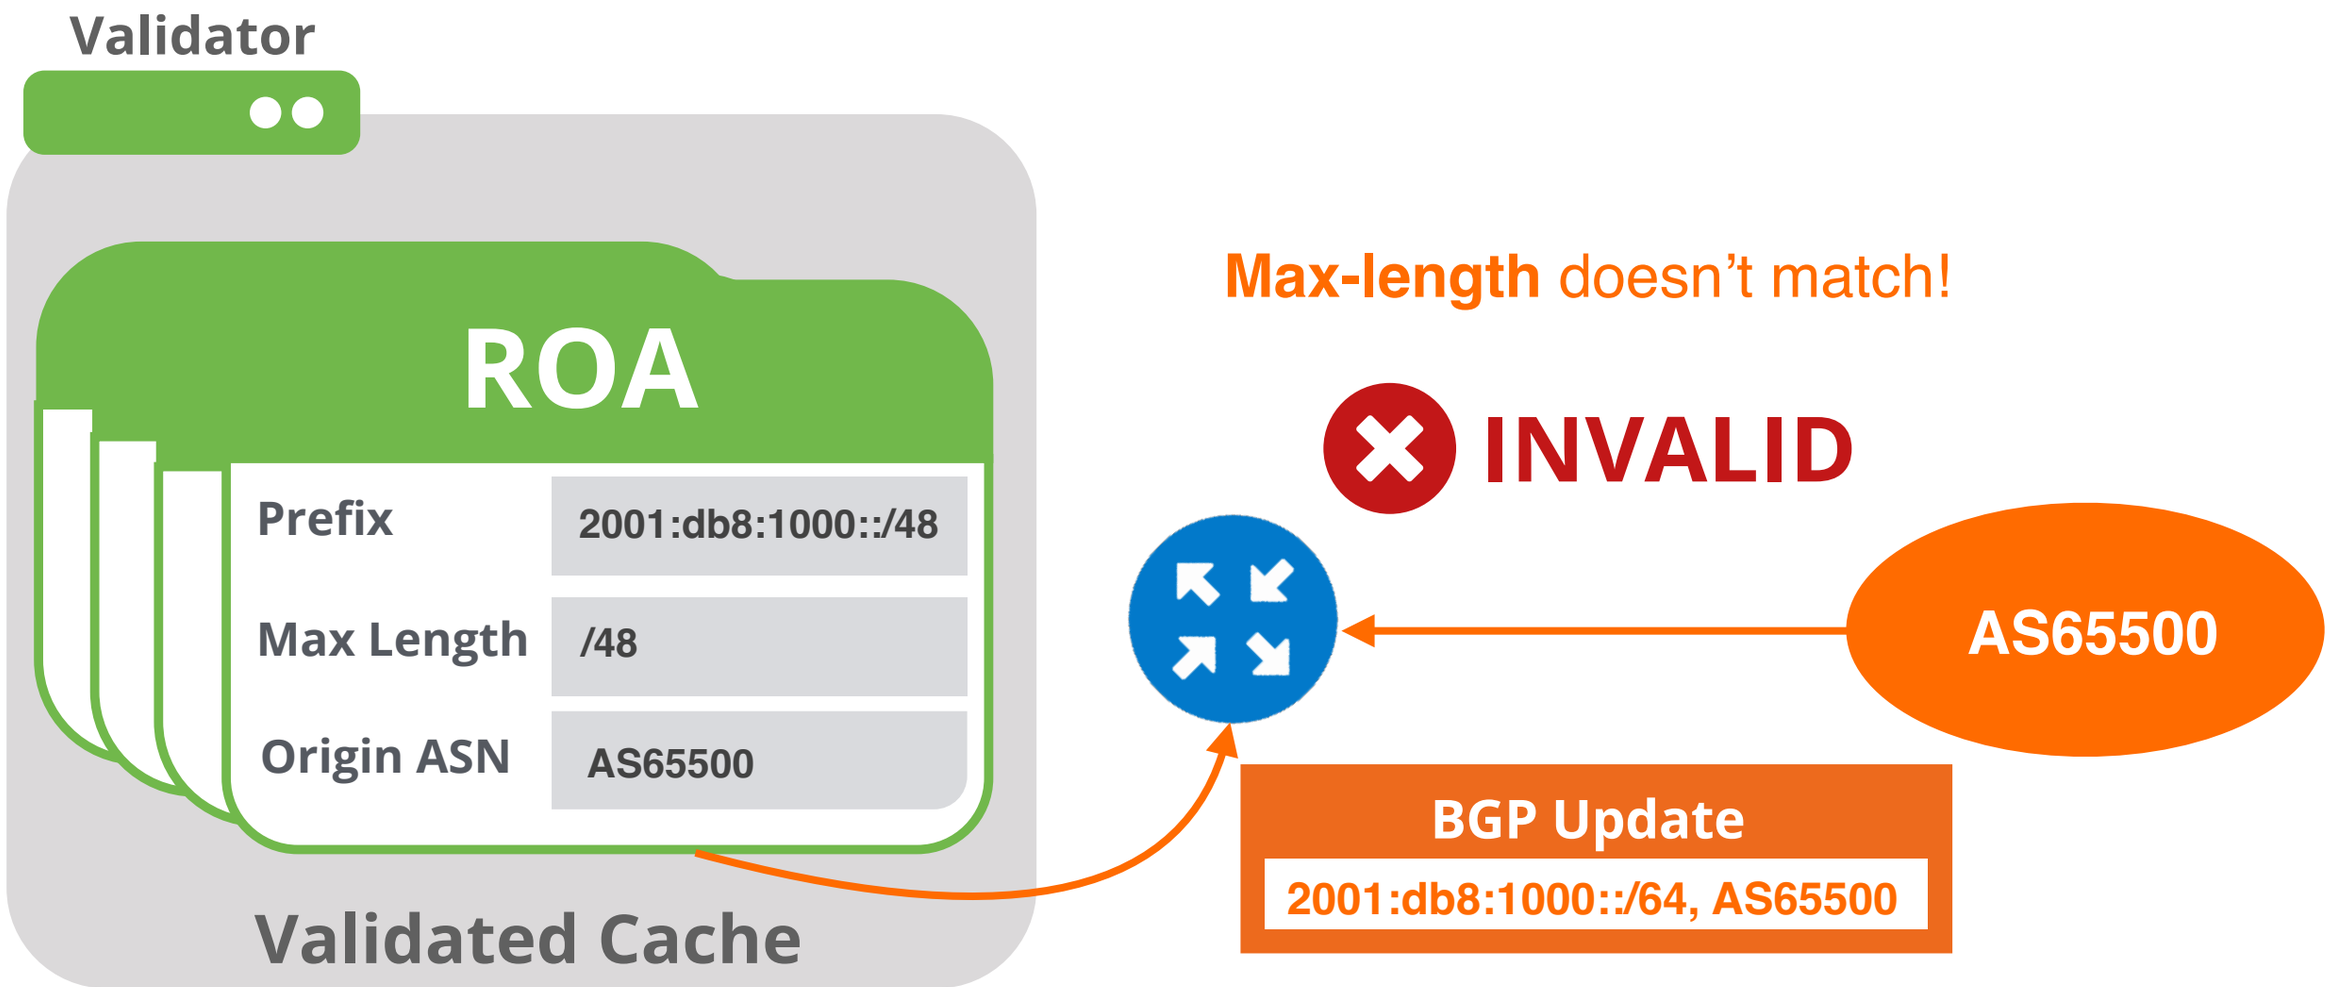

How does RPKI validate the origin?

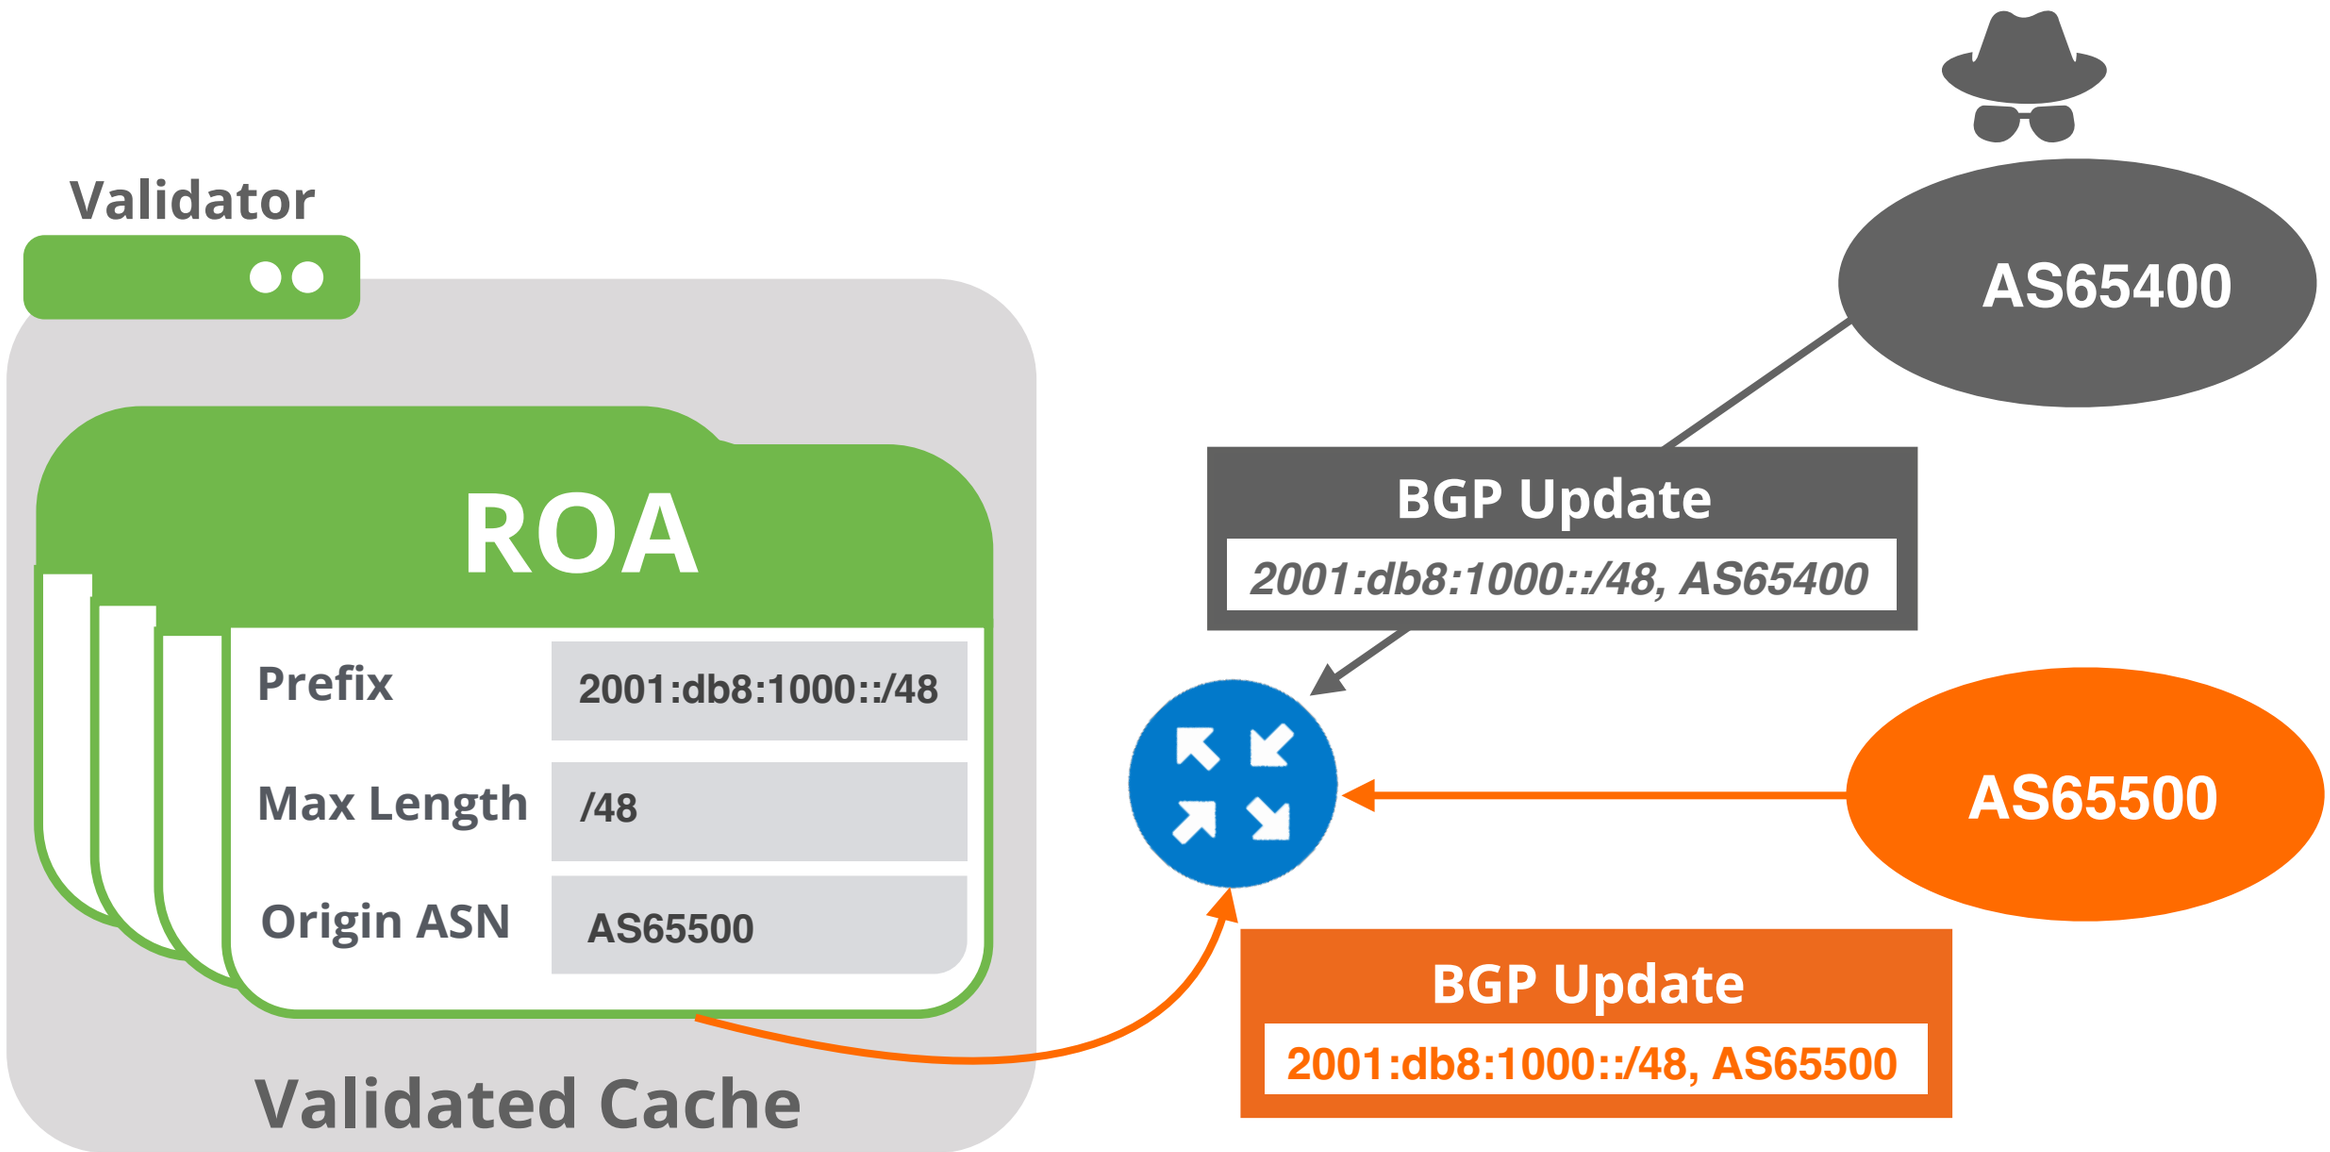

How does RPKI validate the origin?

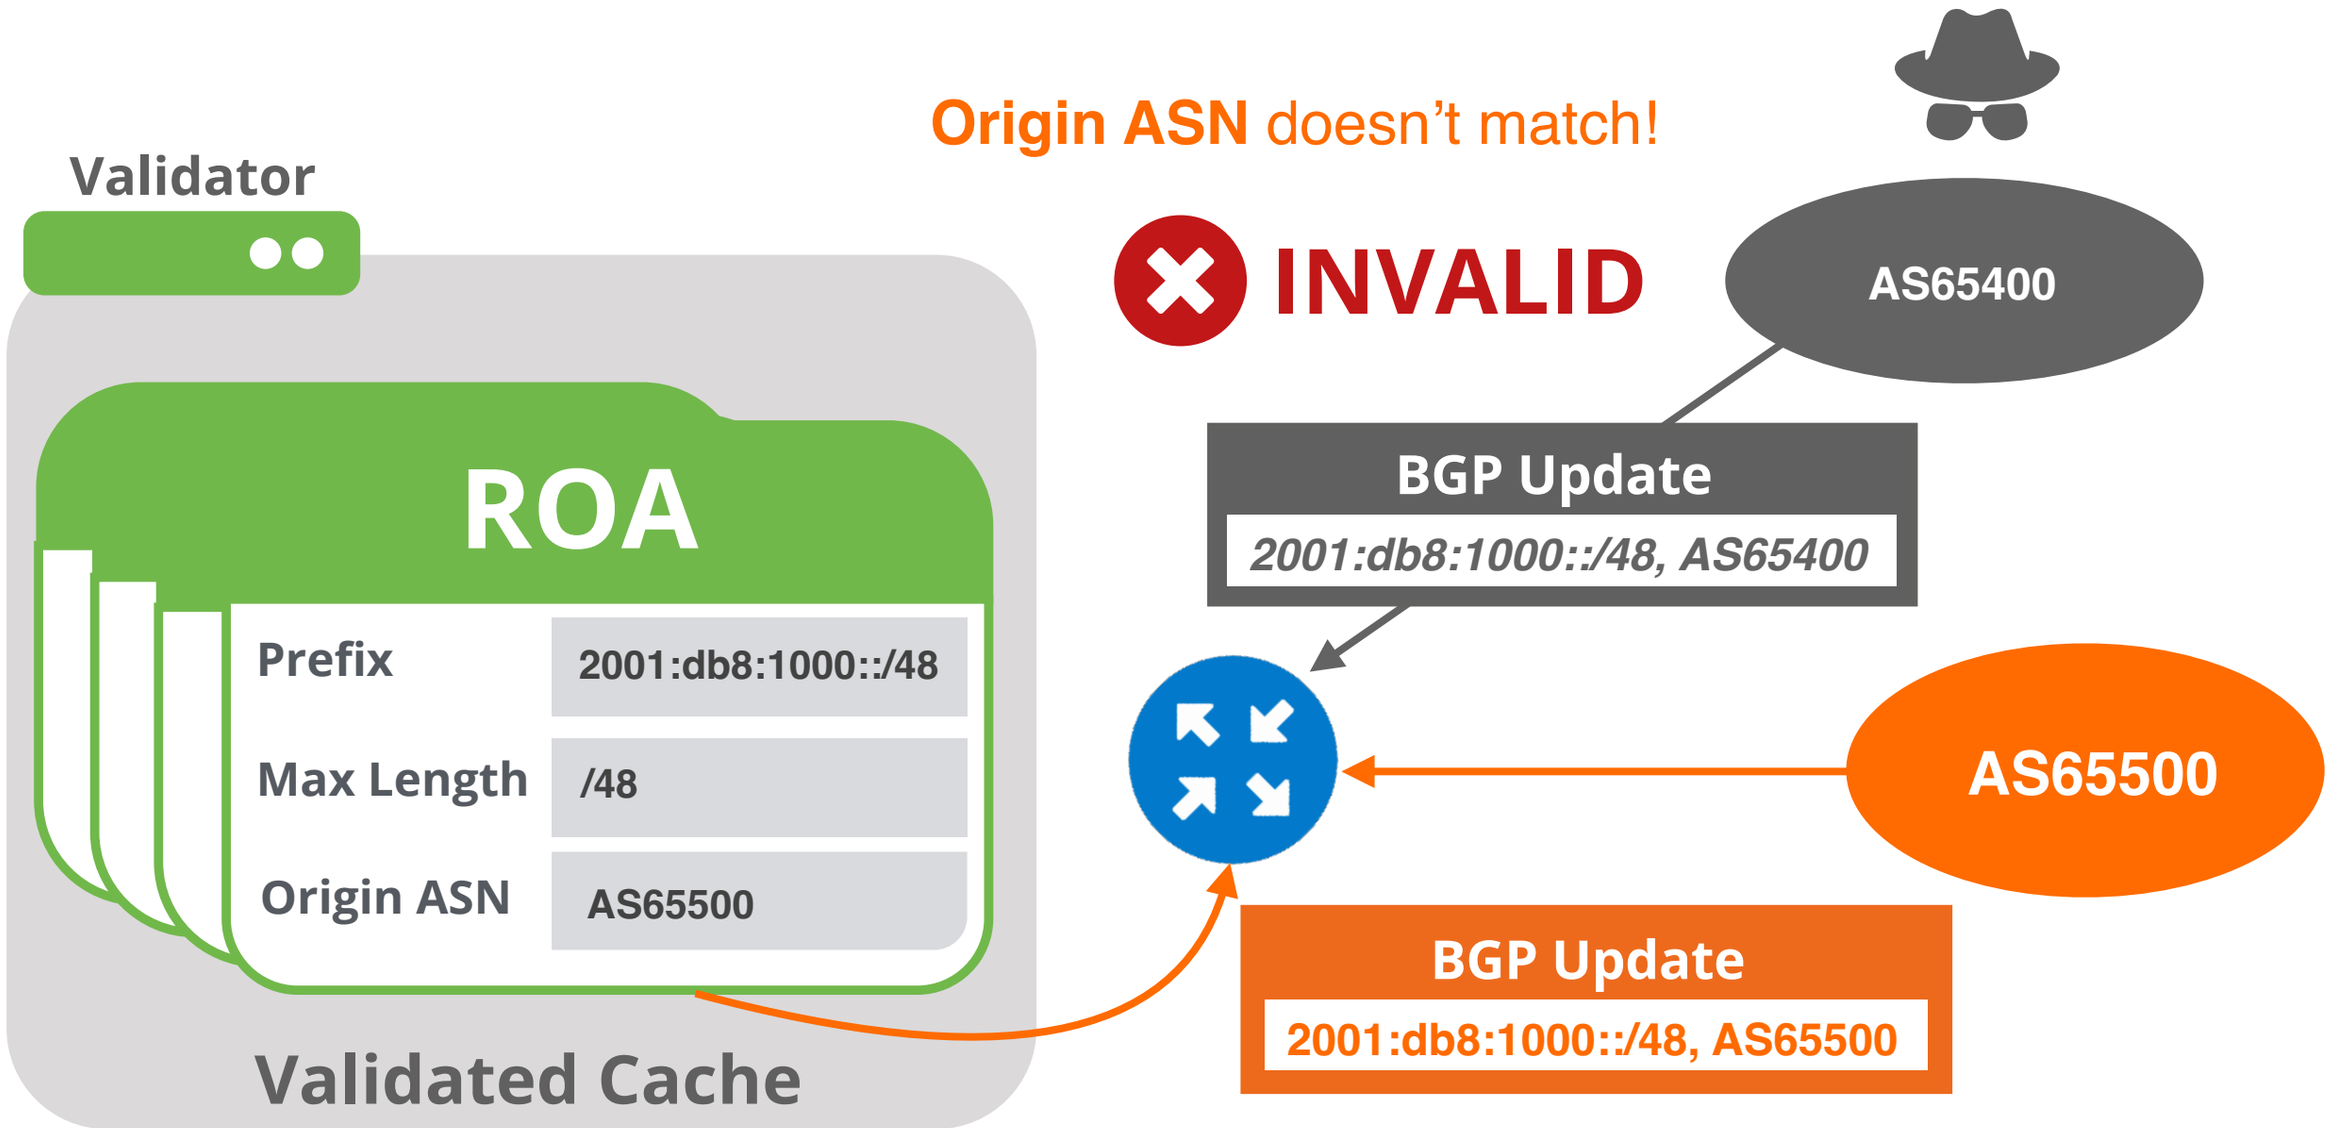

How does RPKI validate the origin?

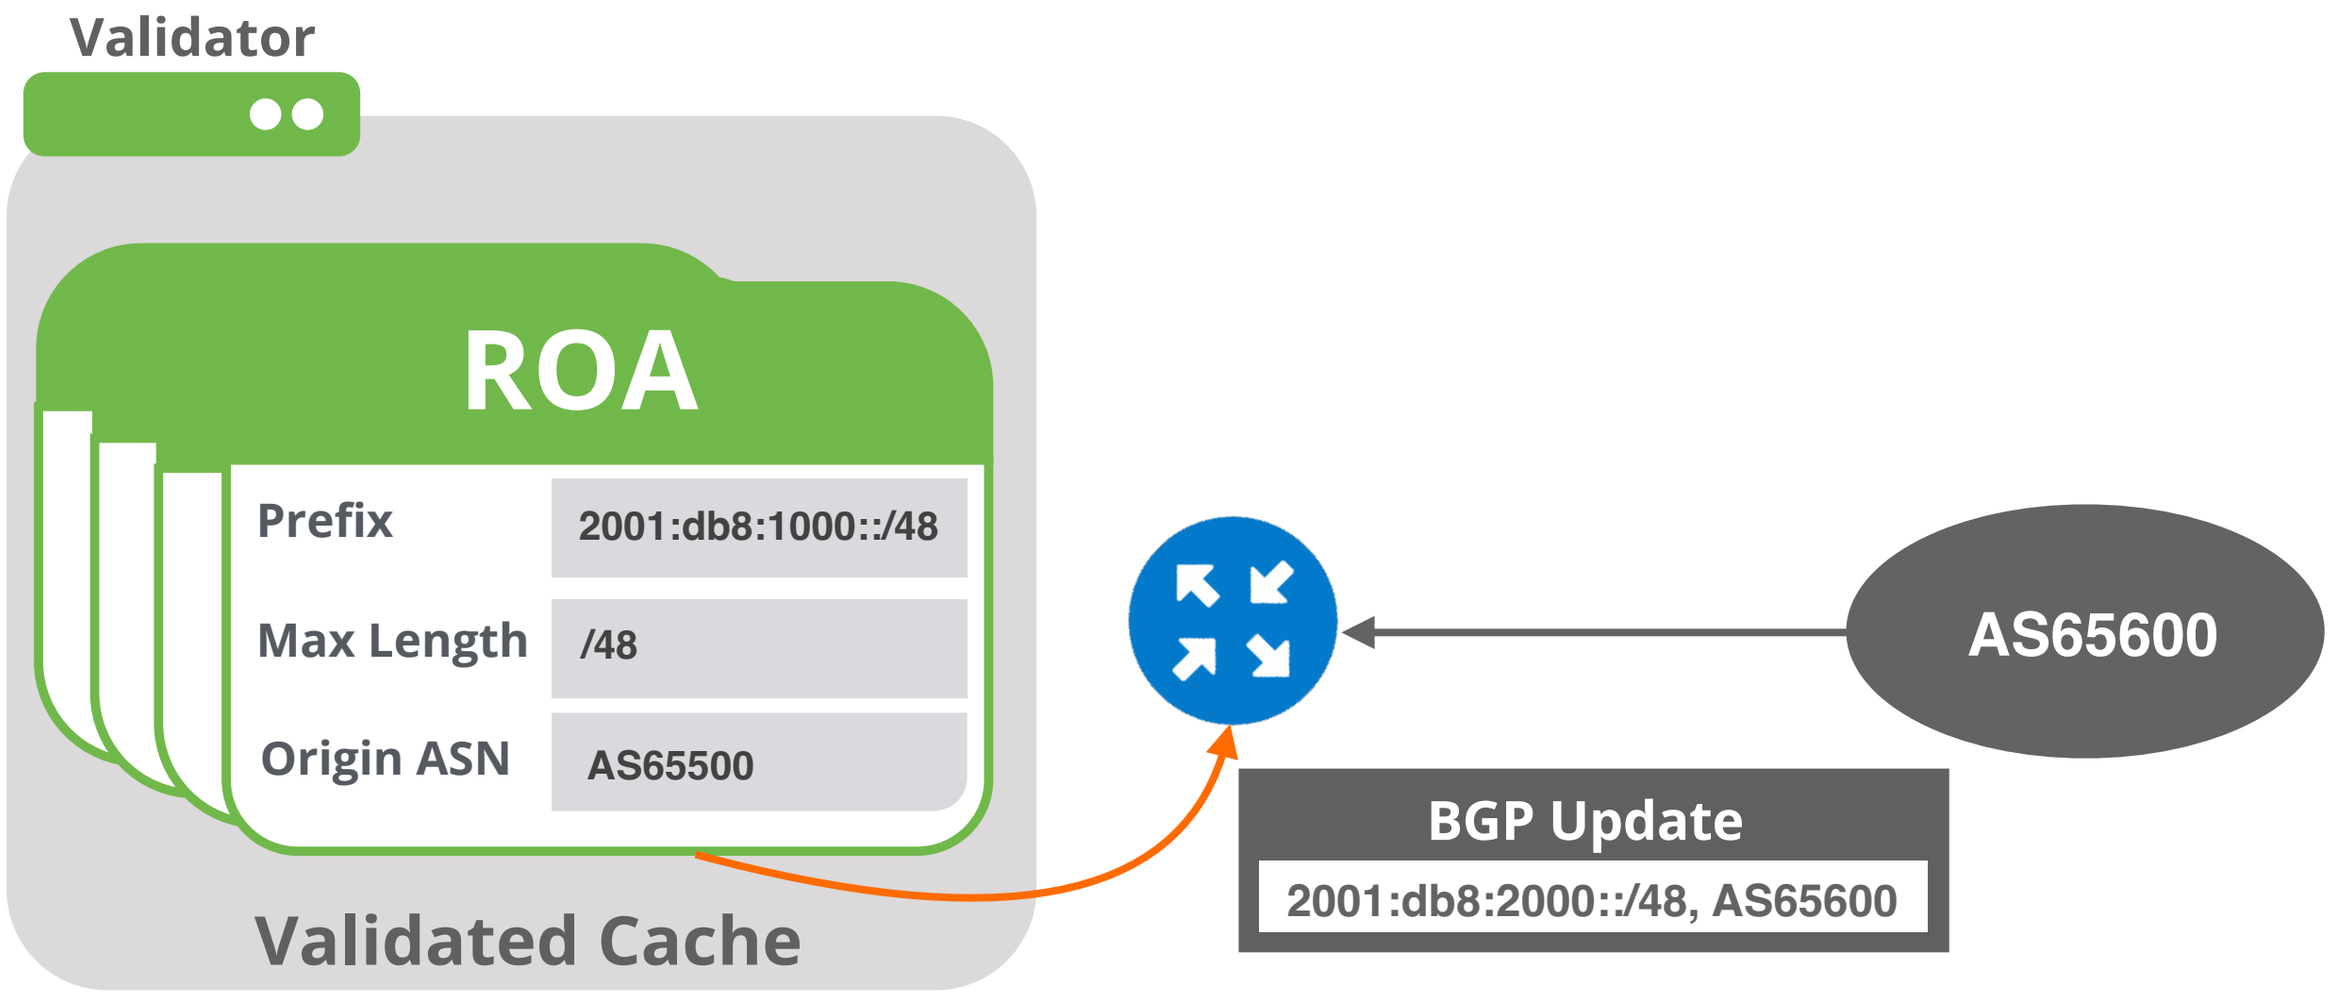

How does RPKI validate the origin?

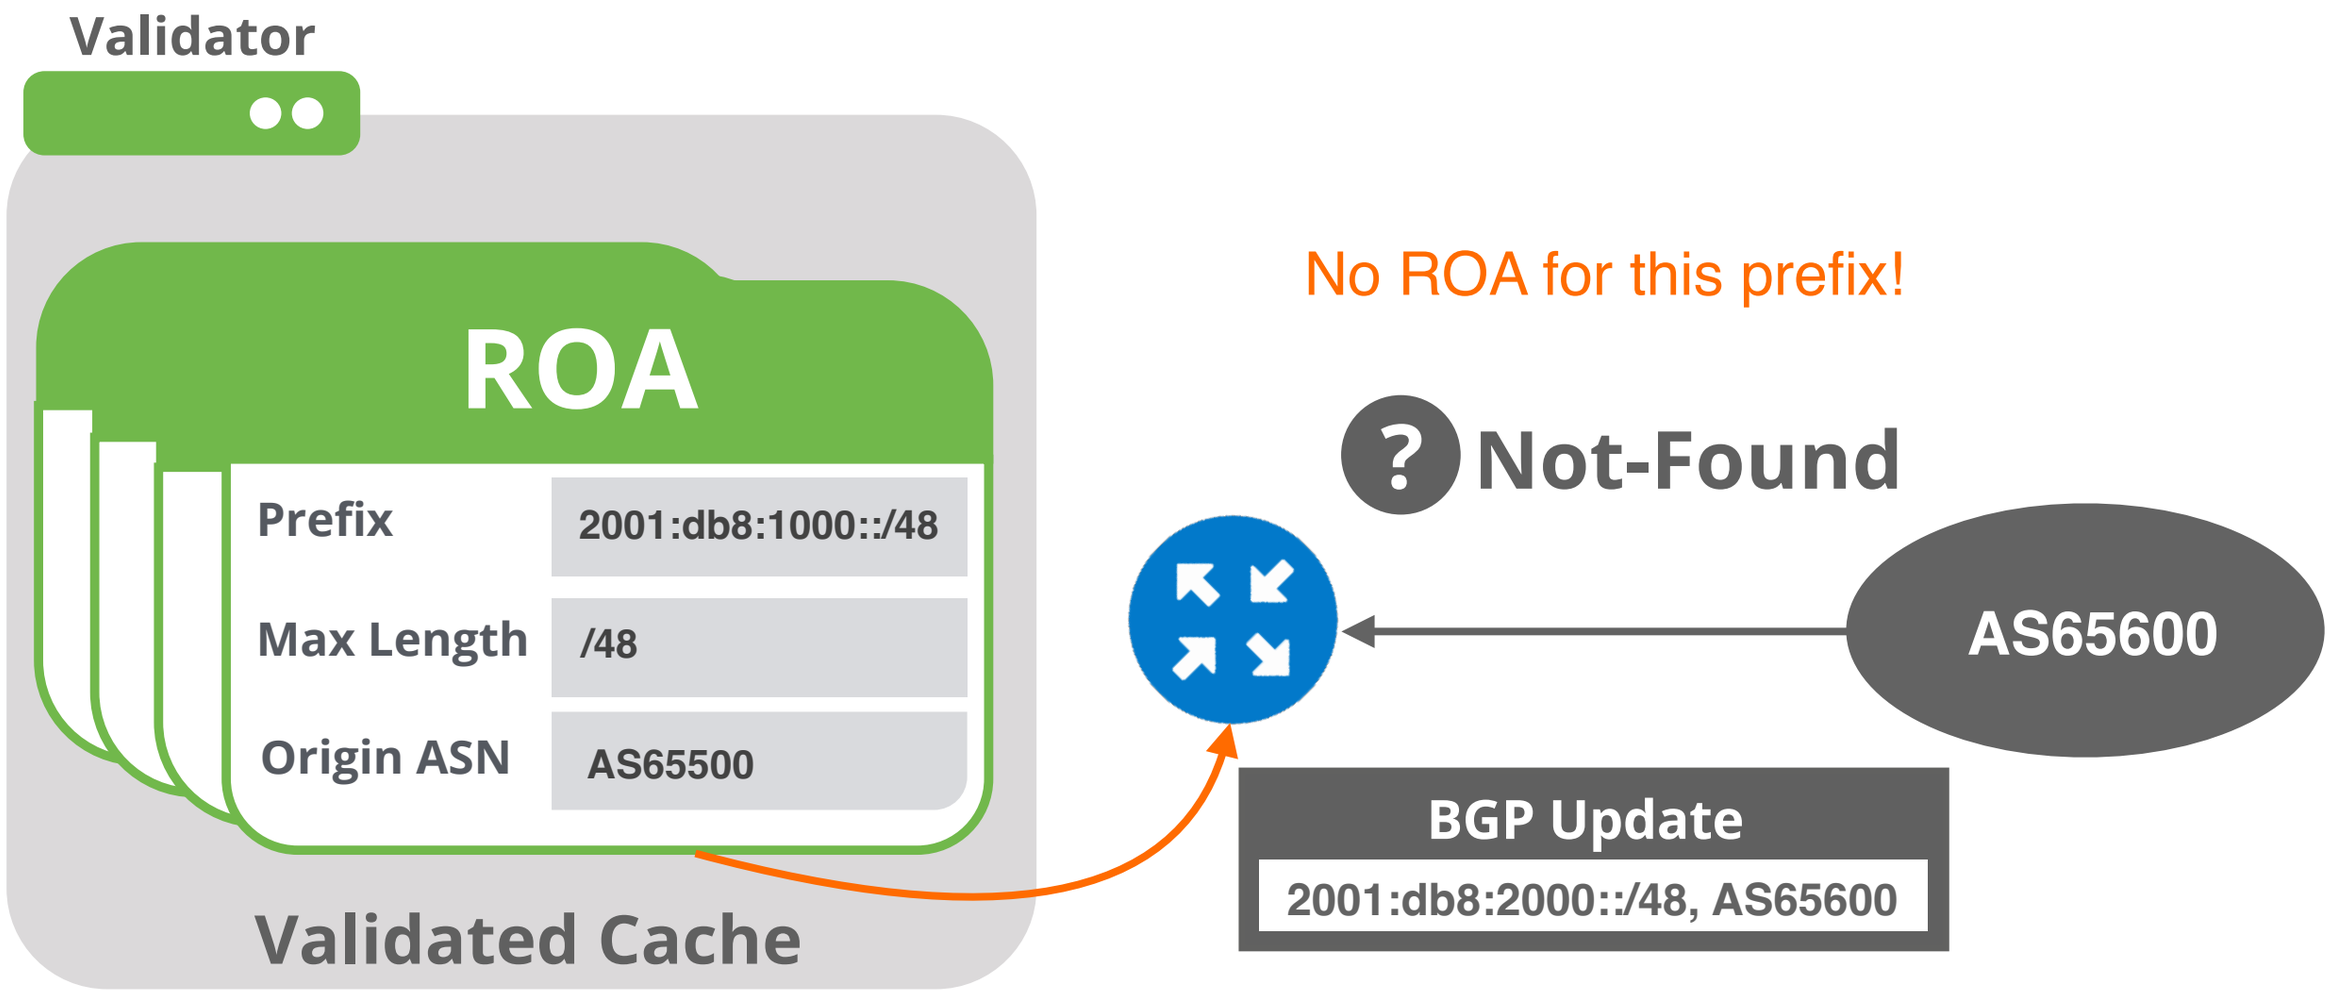

After Validating…

- You have to make a decision : “Accept” or “Discard”

Valid → Accept the prefix (set higher local preference)

Invalid → Discard the prefix

NotFound → Accept the prefix (may set lower local preference)

After Validating…

- You have to make a decision : “Accept” or “Discard”

Valid → Accept the prefix (set higher local preference)

Invalid → Discard the prefix

NotFound → Accept the prefix (may set lower local preference)

Do not consider dropping prefixes with “NotFound” RPKI validation state!

Discarding BGP Invalids

- For BGP origin validation (BGP OV) to achieve its goal…
  - Invalids should be dropped!
- Tag the invalids with a BGP communities
  - or set lower local preference (not a long term solution)
- After analysing the effect, you can start dropping invalids

Discarding BGP Invalids

- Major networks are dropping invalid BGP prefixes!
  - Telia, AT&T, Cloudflare, Netflix, Swisscom, Cogent, …
- April 2021, RIPE NCC (AS3333) started dropping invalids too!
  - Only networks with RPKI Valid or Unknown announcements are allowed
  - This change effects access to RIPE NCC network and the LIR portal
  - K-Root (AS25152) is not part of AS3333

Questions

DEMO

BGP Origin Validation

Goals

- Validate the origin of BGP announcements on your border router
  - Check RPKI validation states, Valid, Invalid, Not Found
- Discard Invalid BGP announcements

Demo Setup

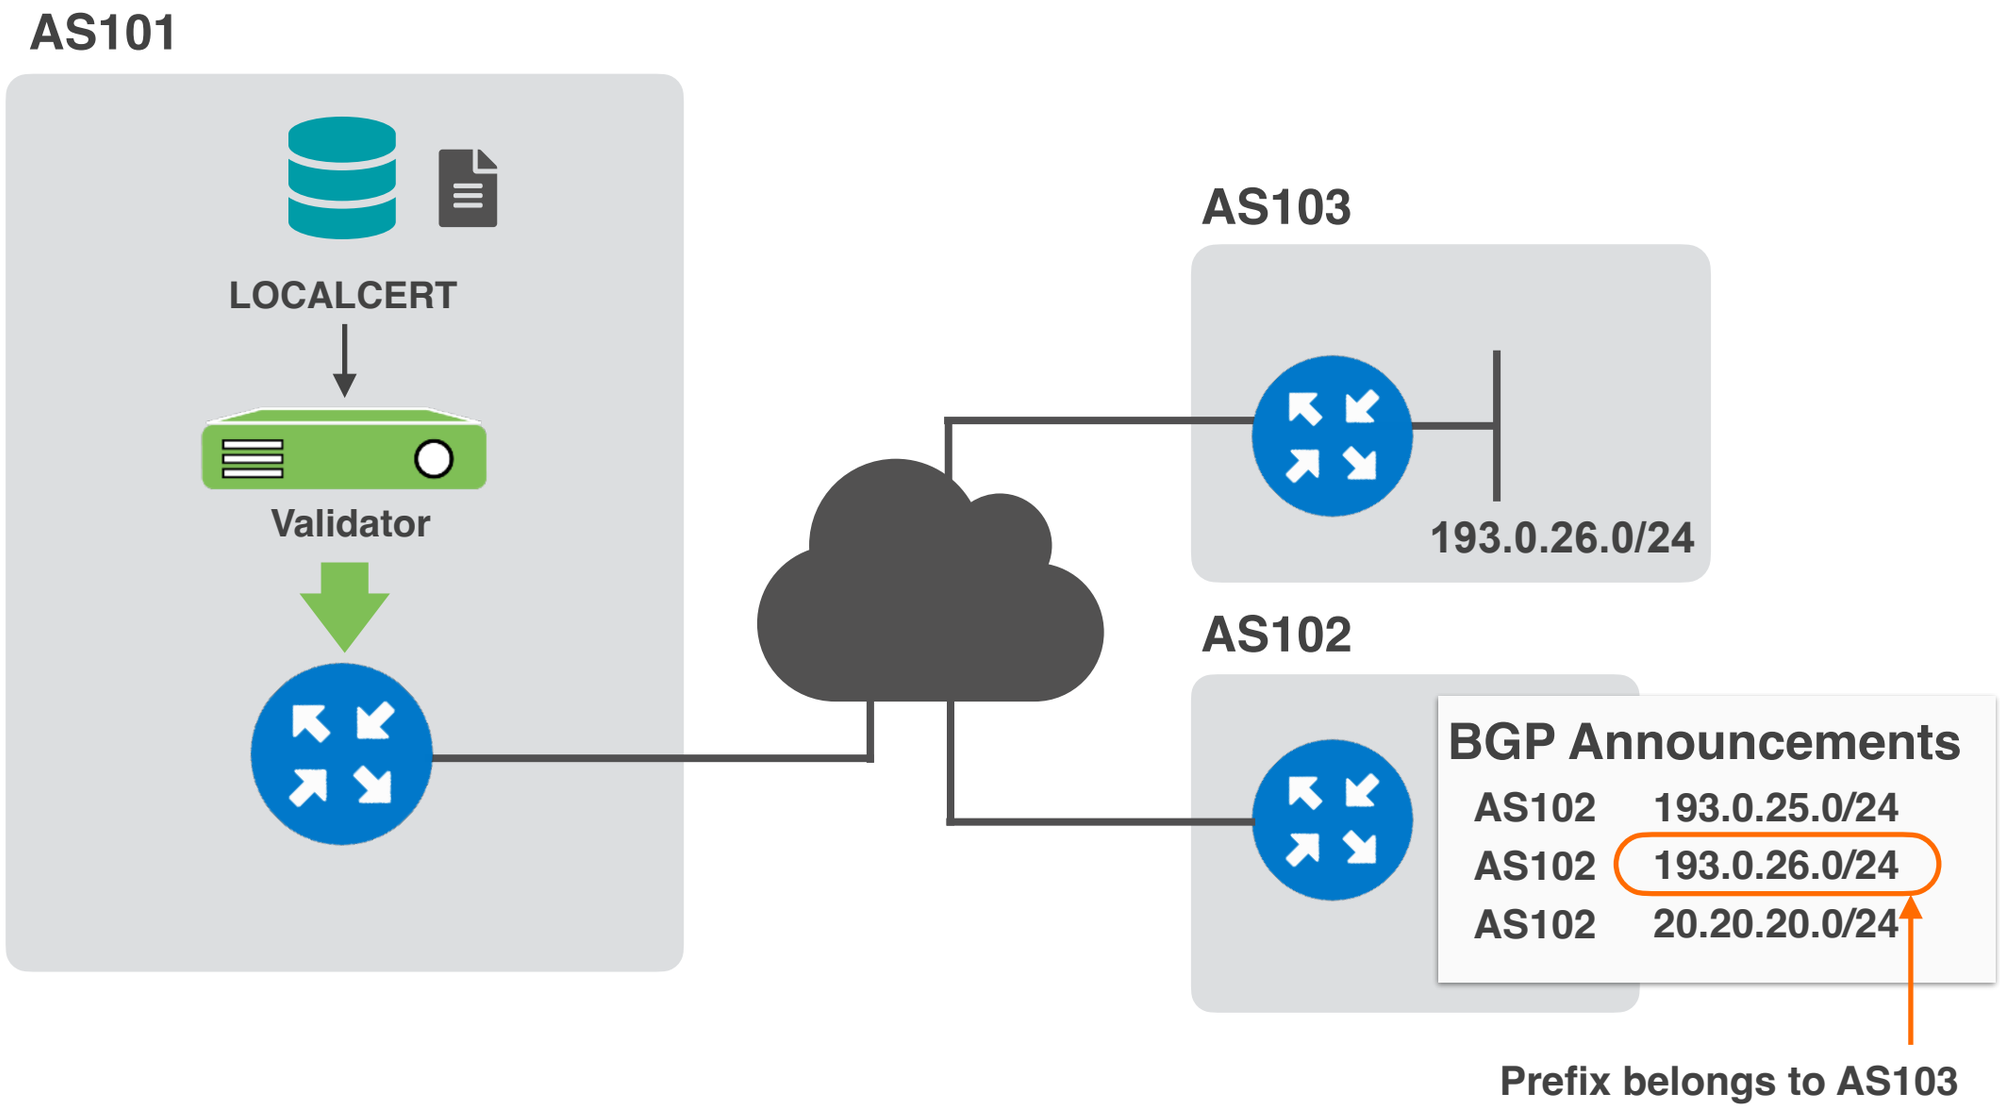
AS101
LOCALCERT
Validator
AS103
193.0.26.0/24
AS102
BGP Announcements
AS102 193.0.25.0/24
AS102 193.0.26.0/24
AS102 20.20.20.0/24
Prefix belongs to AS103

Step-1: Setup validator connection

- Routinator is running on port **3323** and Fort is on **323**
- Configure RPKI-RTR on your router
- Check RPKI prefix table

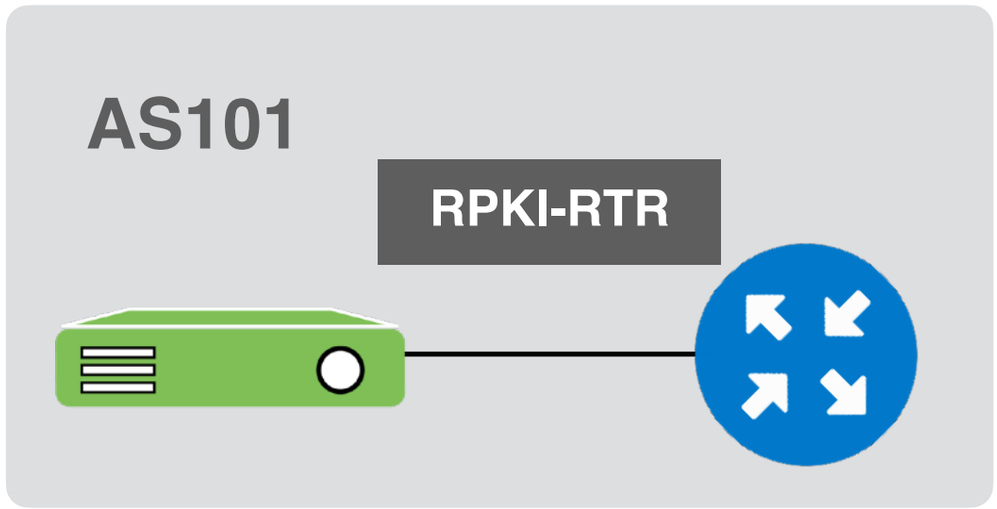

On AS101 router

```
(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
```

```
# show ip bgp rpki table
```

Step-2: Create a BGP hijack

- **AS102** is the hijacker!

It'll originate the prefix of AS103.

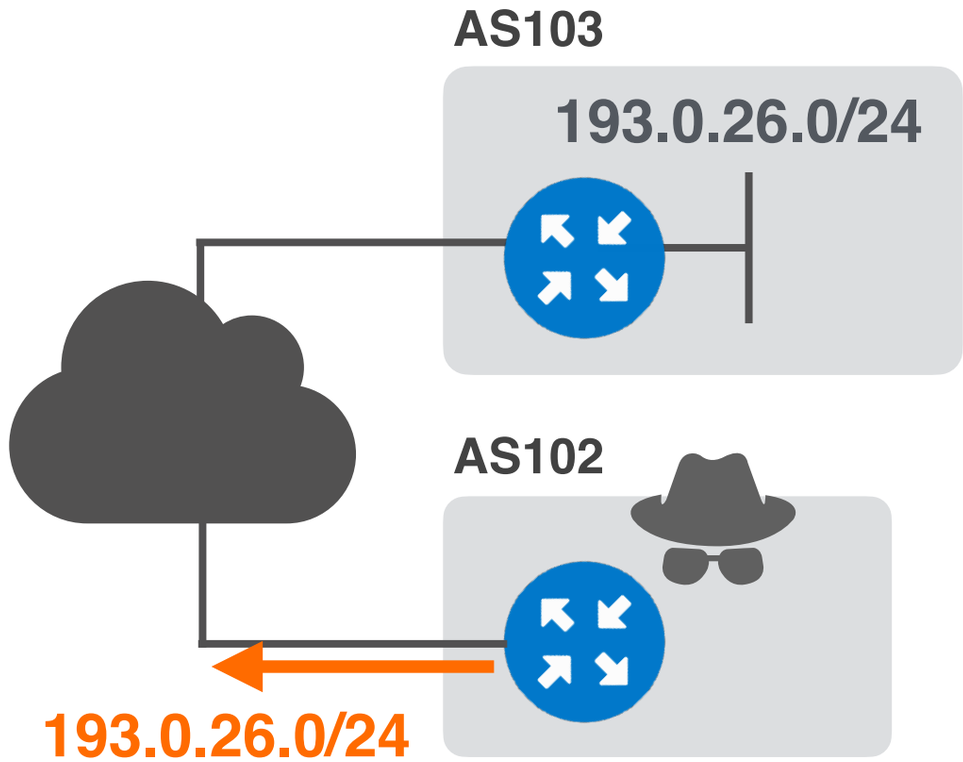

On AS102 router

```
(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!

Step-3: Check validation result

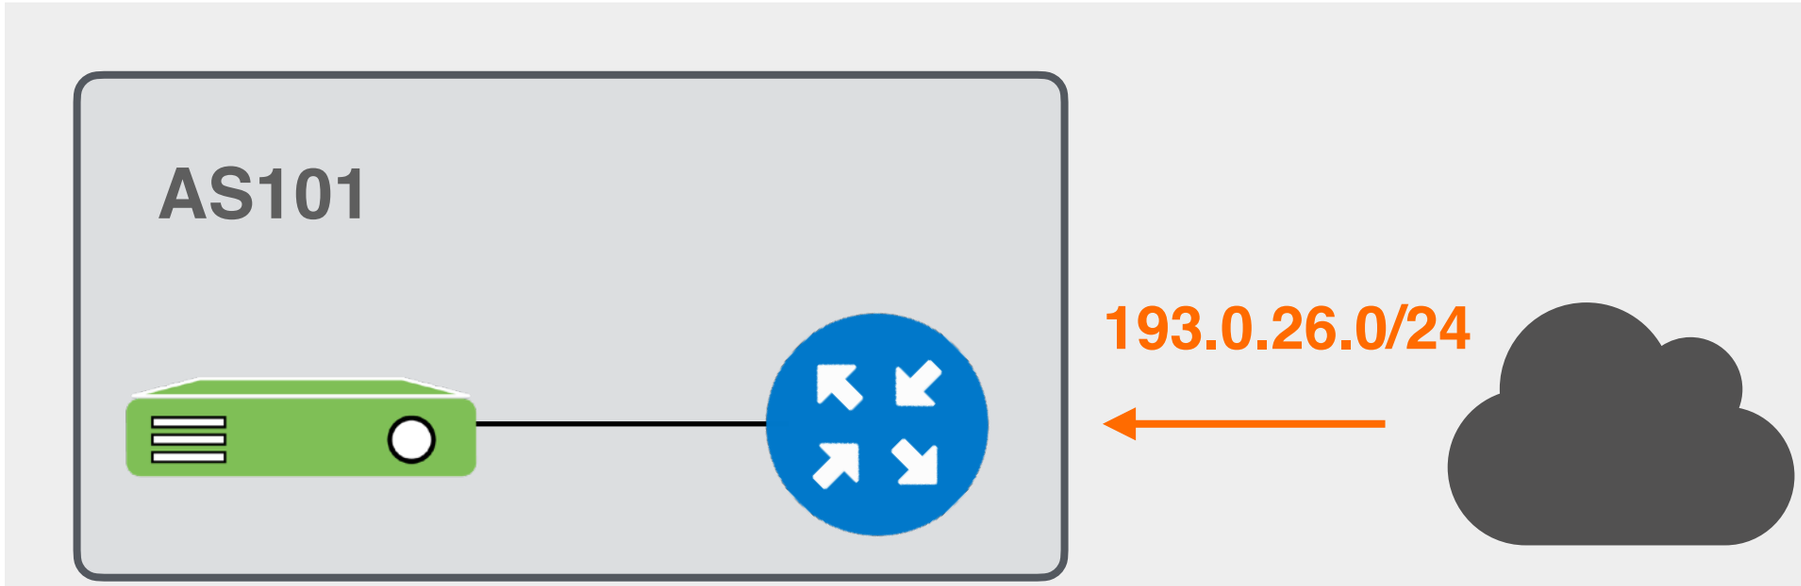

- On your router, check RPKI validation states for the routes in BGP table

On AS101 router

```
# show ip bgp
# show ip bgp ipv6 unicast
```

Step-4: Discard invalids

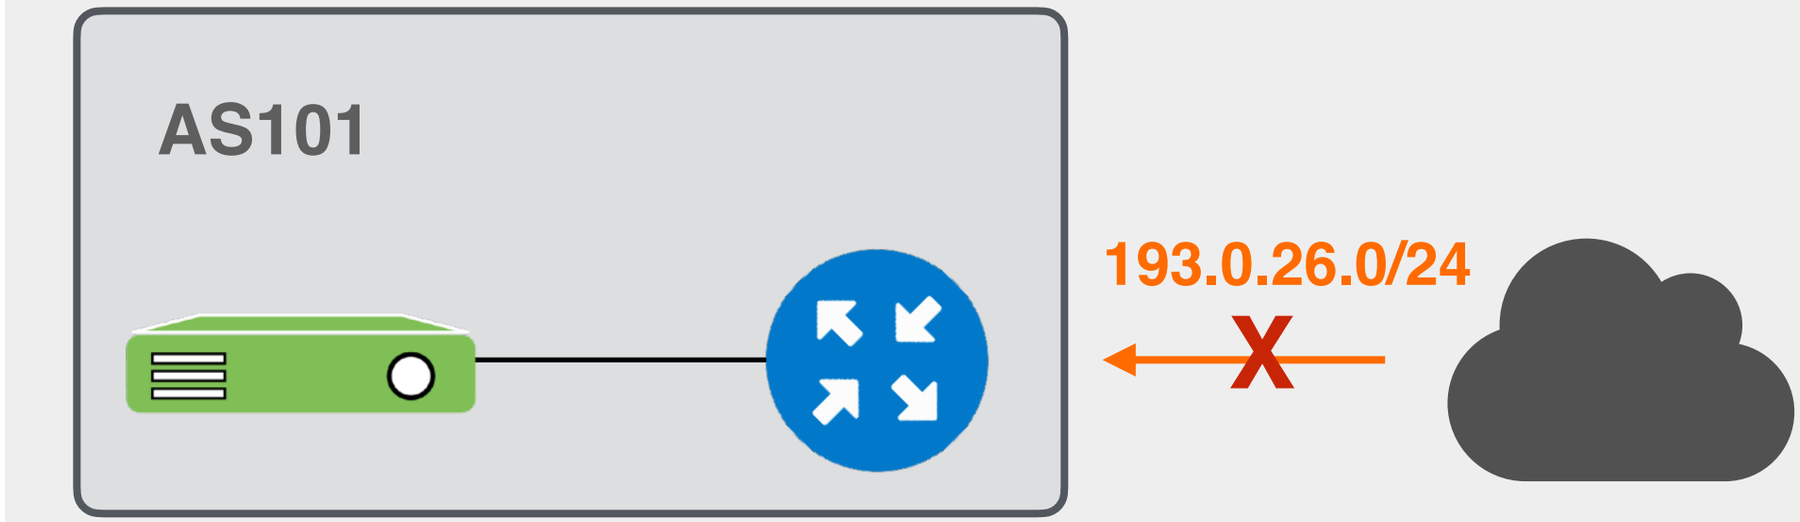

On AS101 router

```
(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
```

Questions

RIPE NCC
Academy

Learn something new today!

academy.ripe.net

RIPE NCC Certified Professionals

Coming soon!

https://getcertified.ripe.net/

We want your feedback!

What did you think about this session? Take our **survey** at:


Ënn
Соңы
An Críoch
پایان
Y Diwedd
Vége
Endir
Finvezh
վերջ
Ende
Koniec
Son
დასასრული
Amaia
הסוף
Кінець
Finis
Lõpp
Tmiem
Sfârşit
Loppu
Liðugt
Kpaj
Slutt
Kraj
النهاية
Конец
Konec
Fund
Fin
Fí
Край
Τέλος
Fine
Einde
Pabaiga
Slut
Fim
Beigas
END