RIPE NCC

RIPE NETWORK COORDINATION CENTRE

Securing BGP with RPKI

Ondřej Caletka

RIPE NCC Learning & Development

Agenda

BGP and Internet Routing

- Is BGP secure?

Routing Security with RPKI

- What is RPKI?
- Building Blocks of RPKI
- BGP Origin Validation (BGP OV)

BGP and Internet Routing

Is BGP secure?

How does it work?

How does it work?

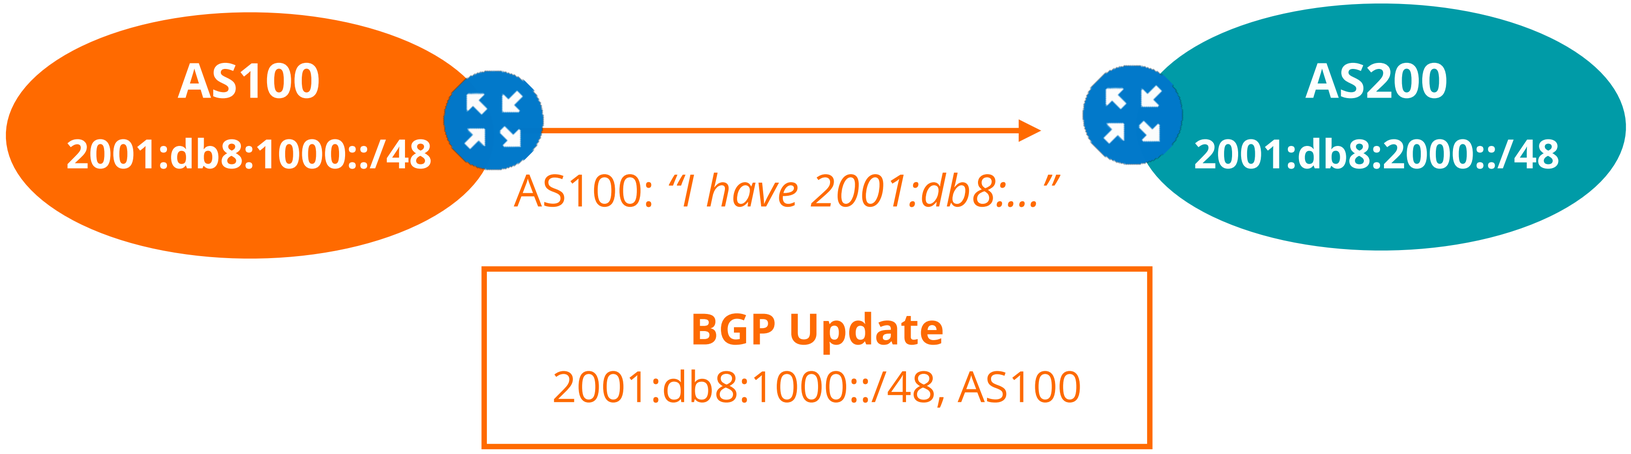

How does it work?

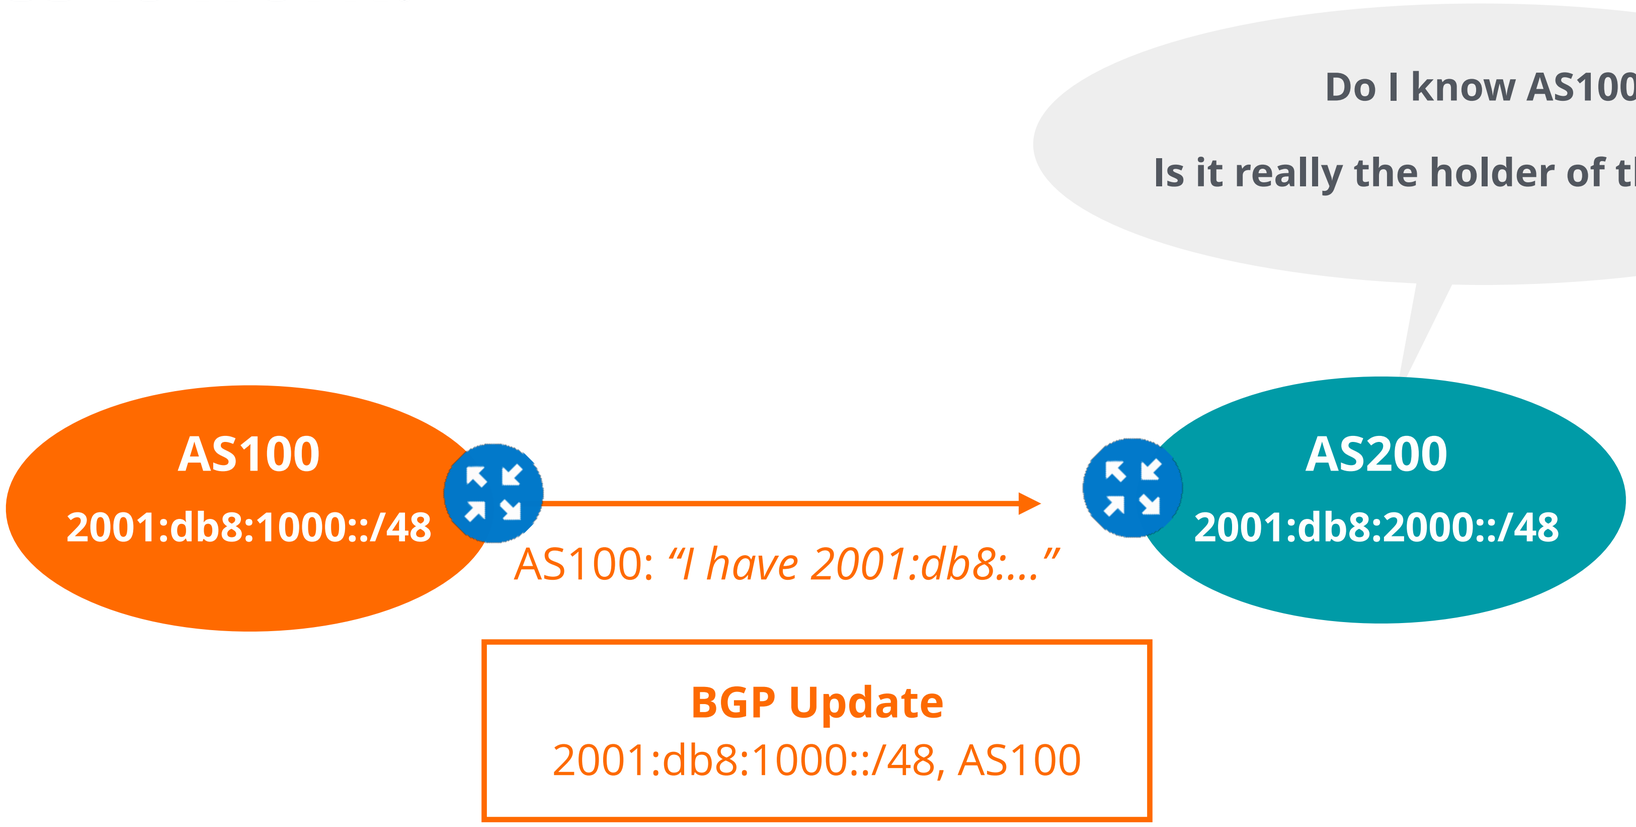

How does it work?

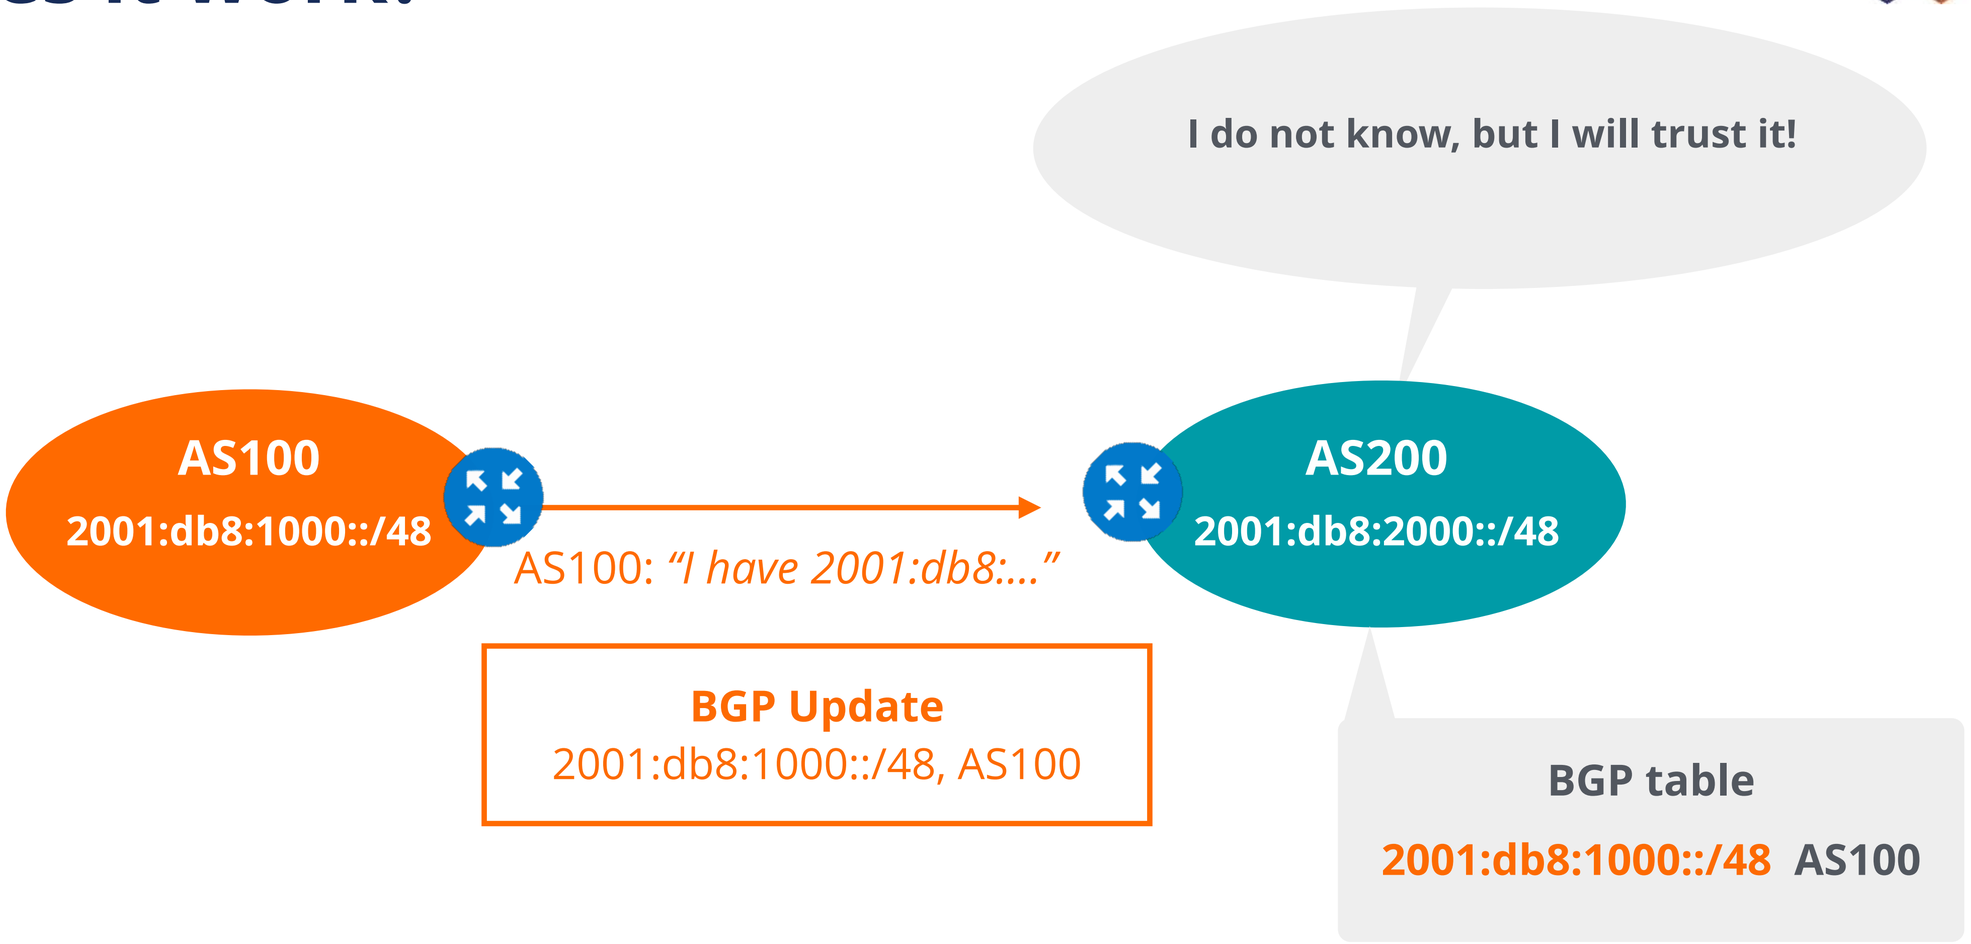

How does it work?

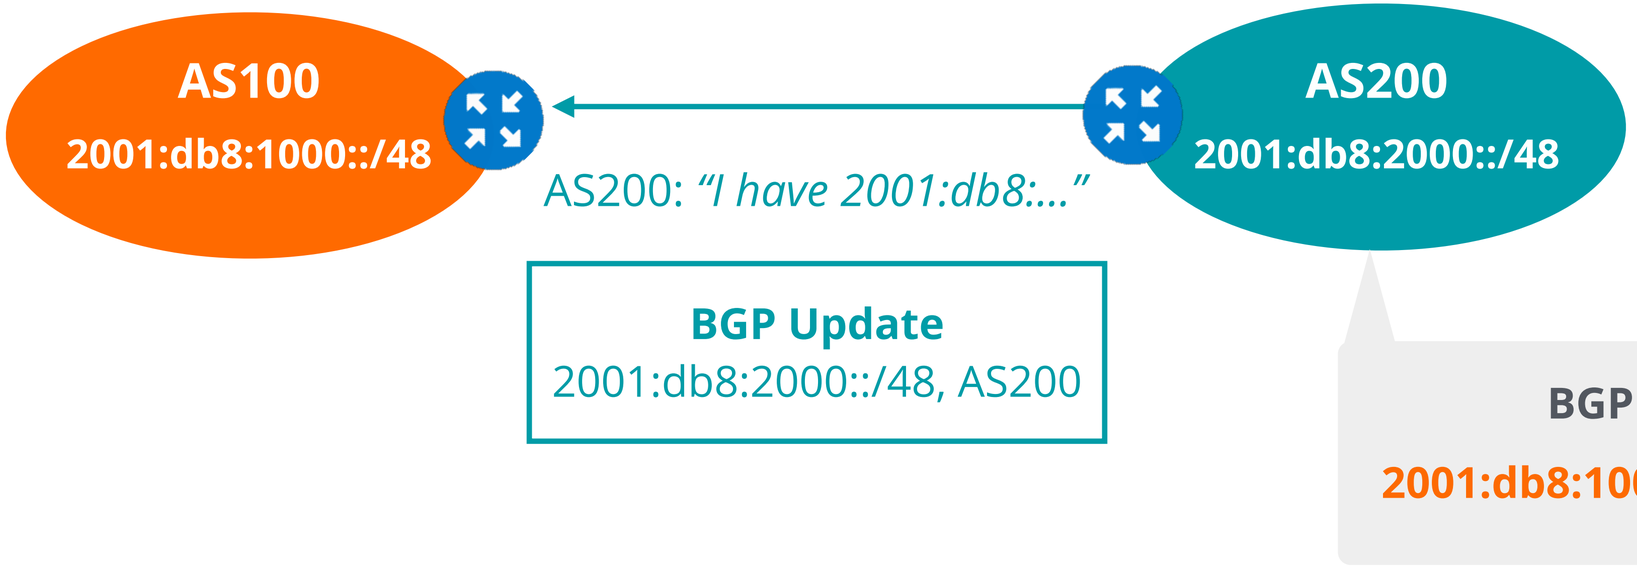

How does it work?

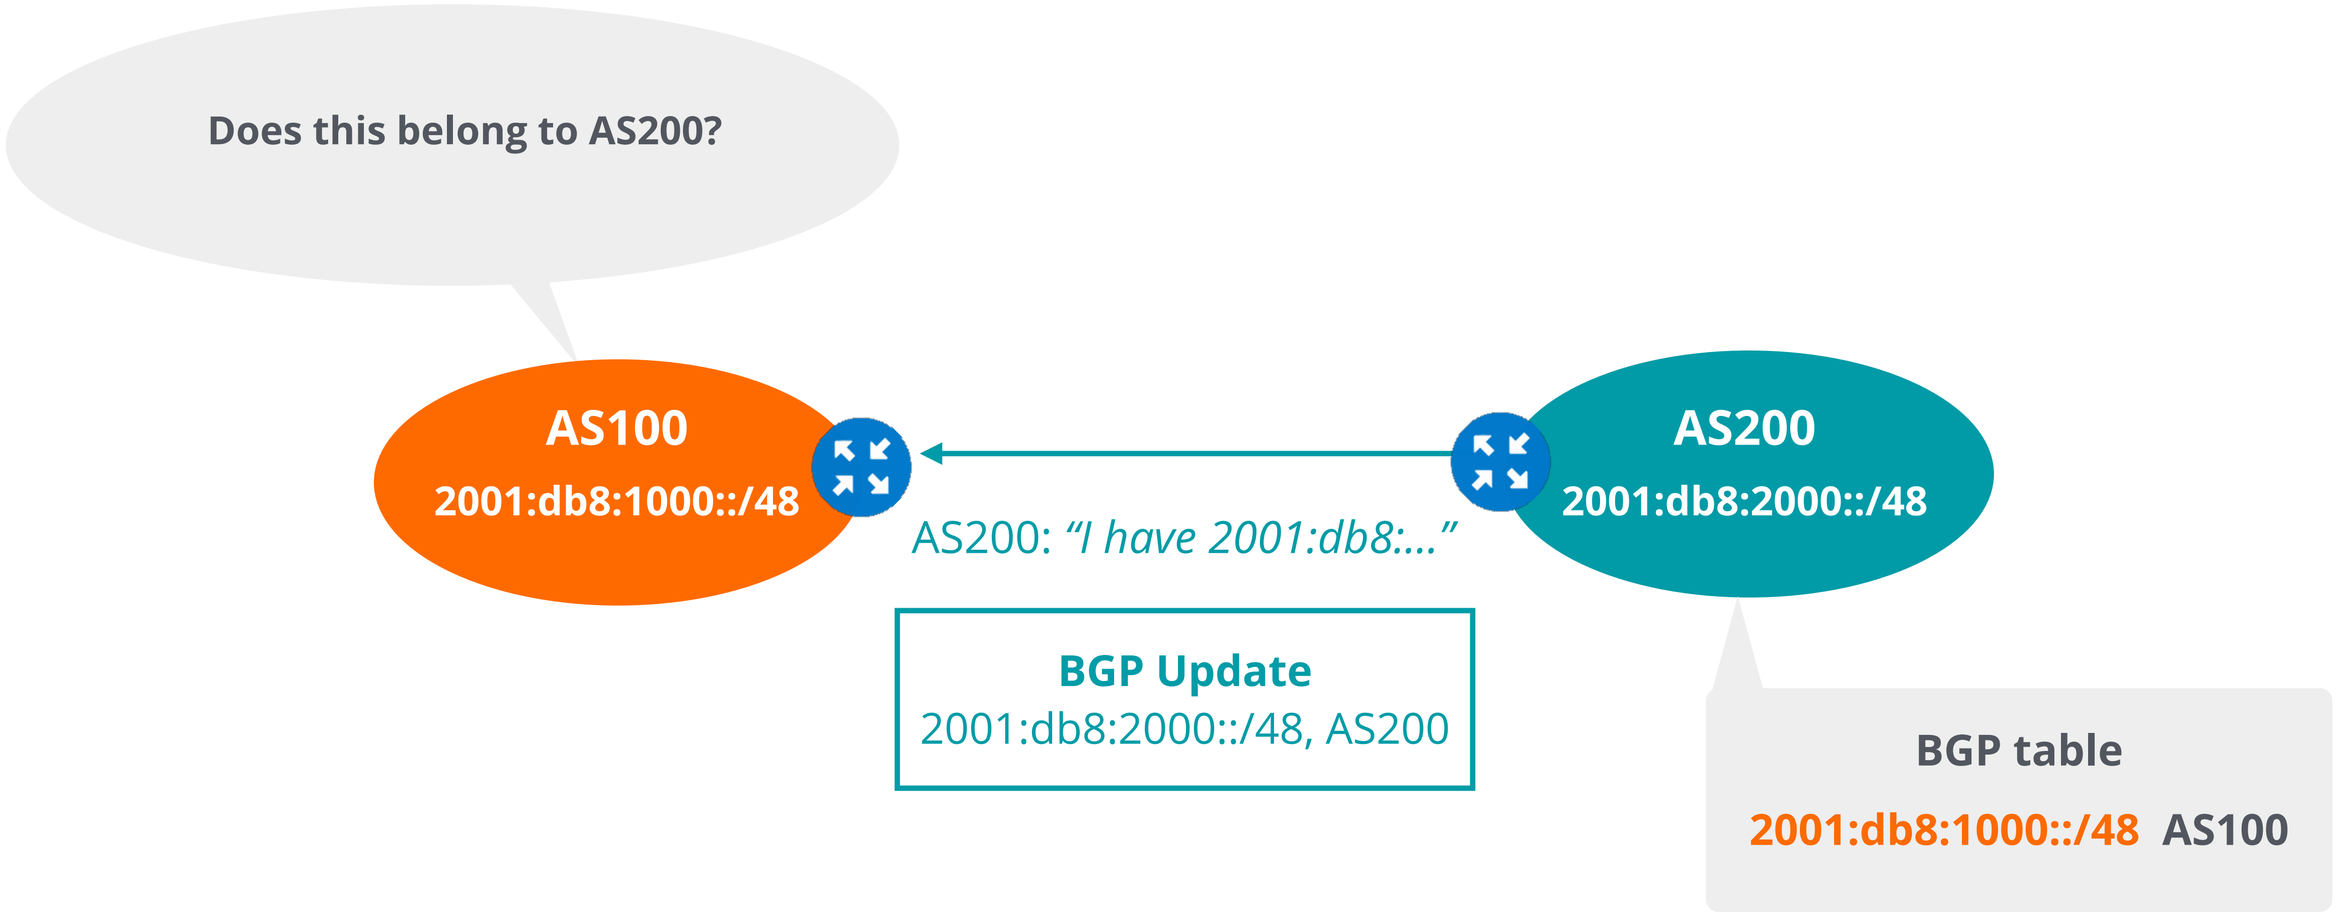

How does it work?

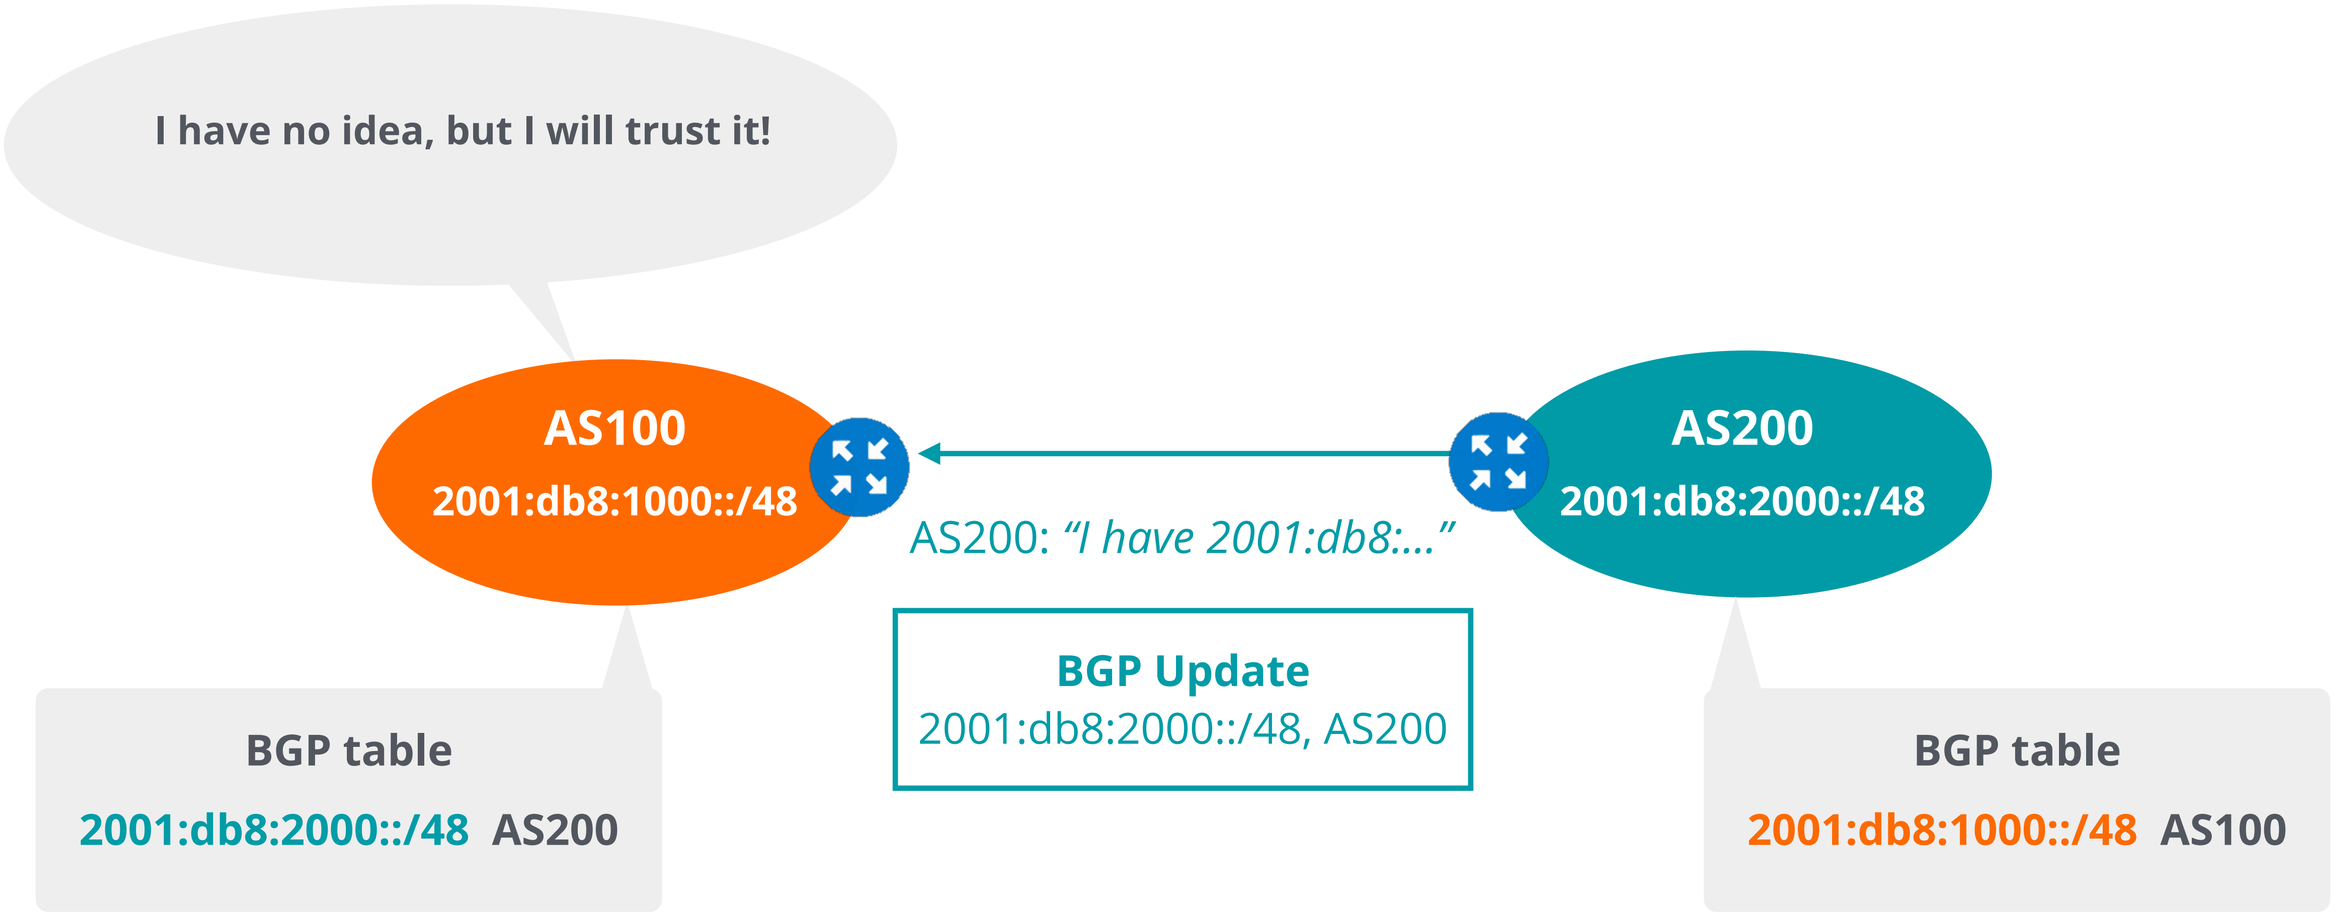

BGP assumes that everybody is telling the truth!

But what if someone lies?

A hijacker may impersonate the legitimate holder!

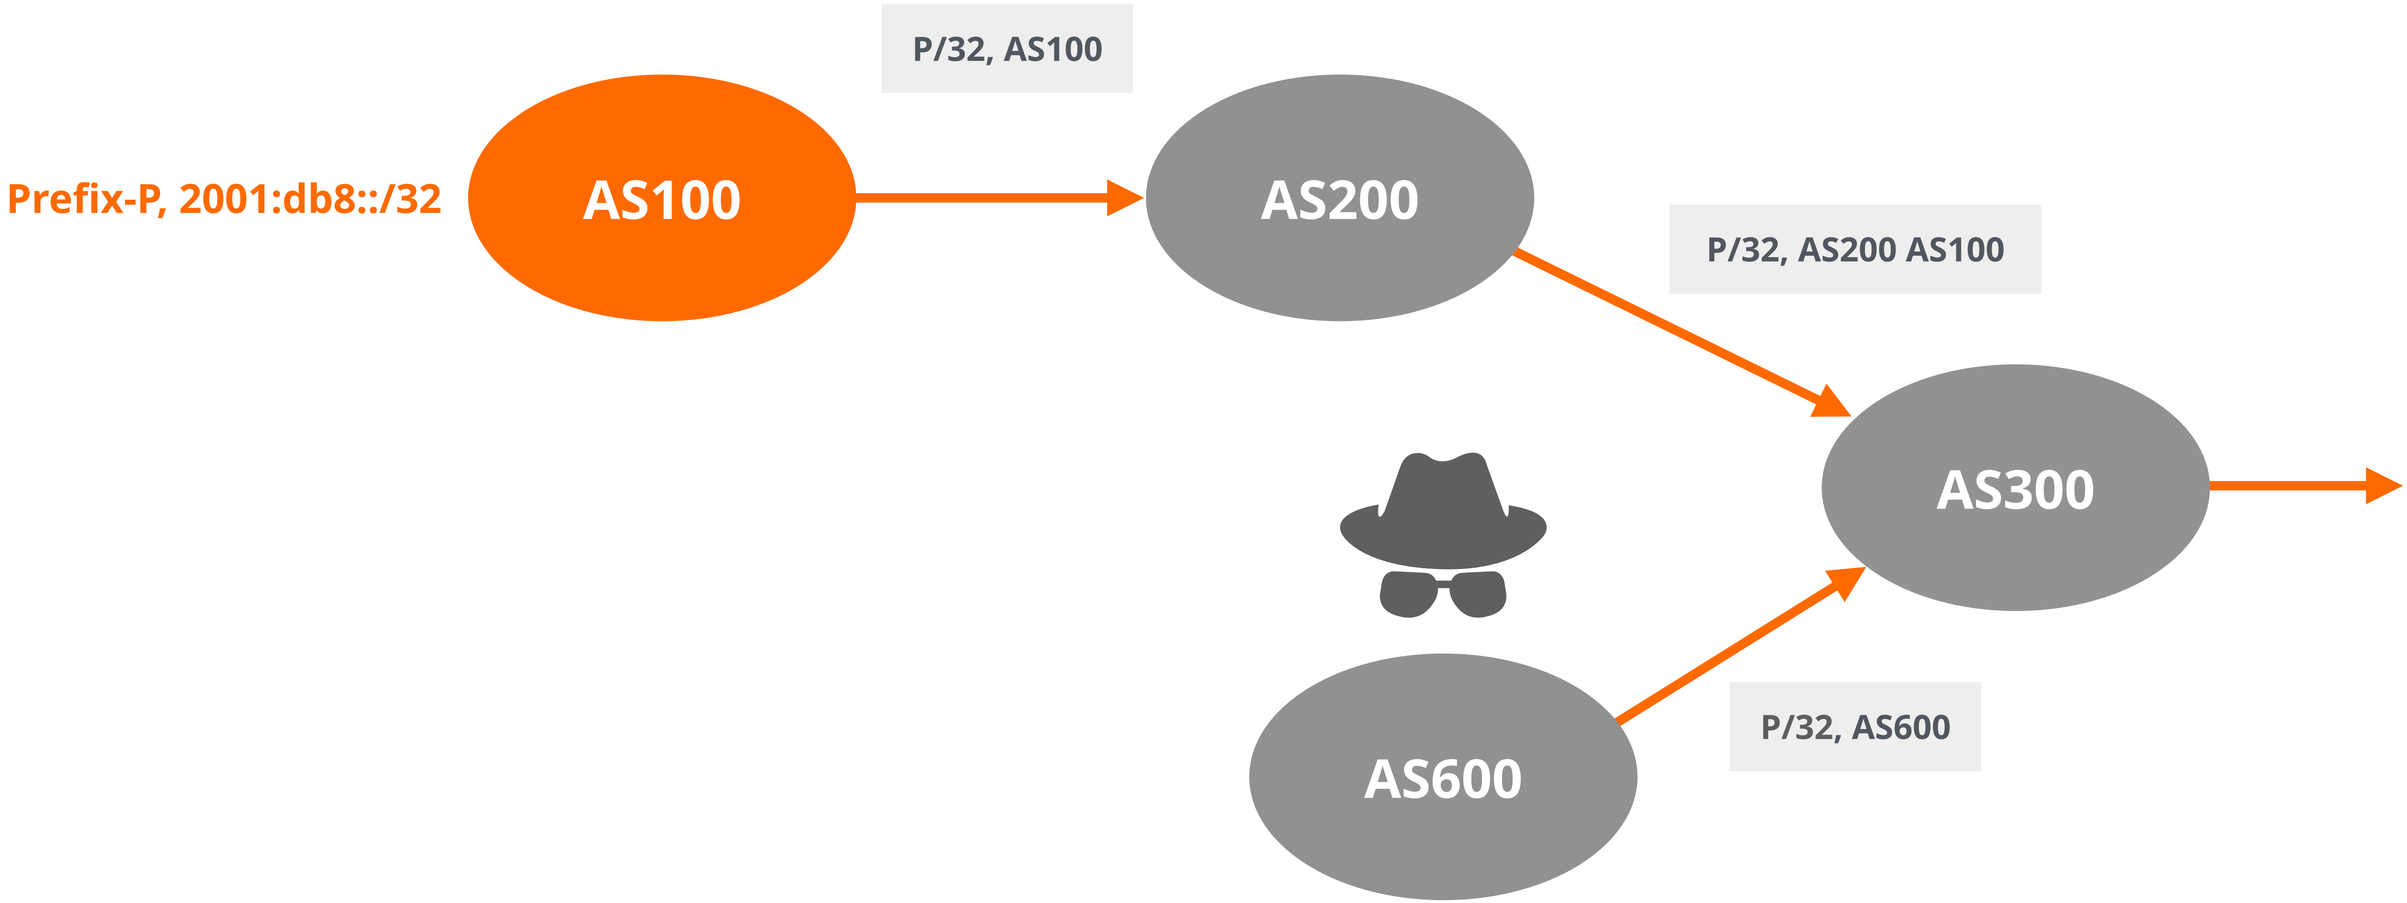

… and may announce the exact same prefix!

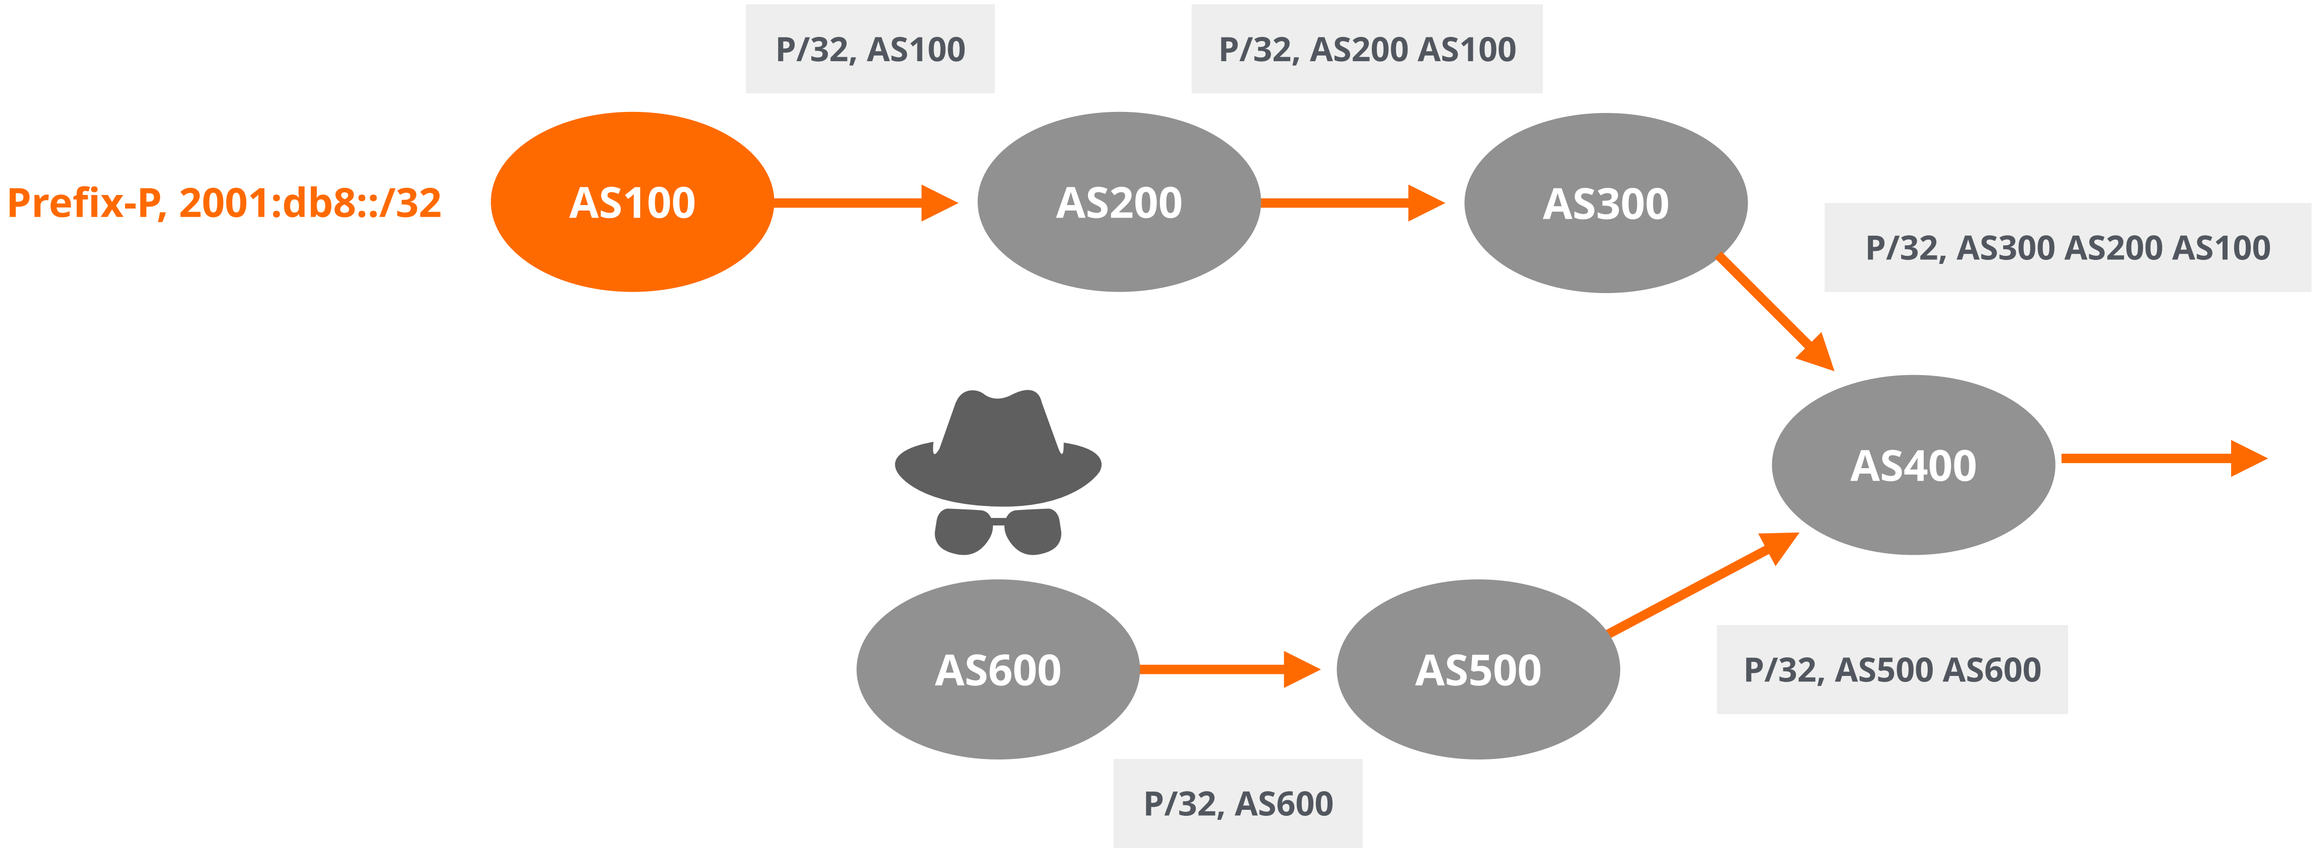

This is a **local hijack!** Only some networks are affected based on BGP path selection.

… or may announce a more specific prefix!

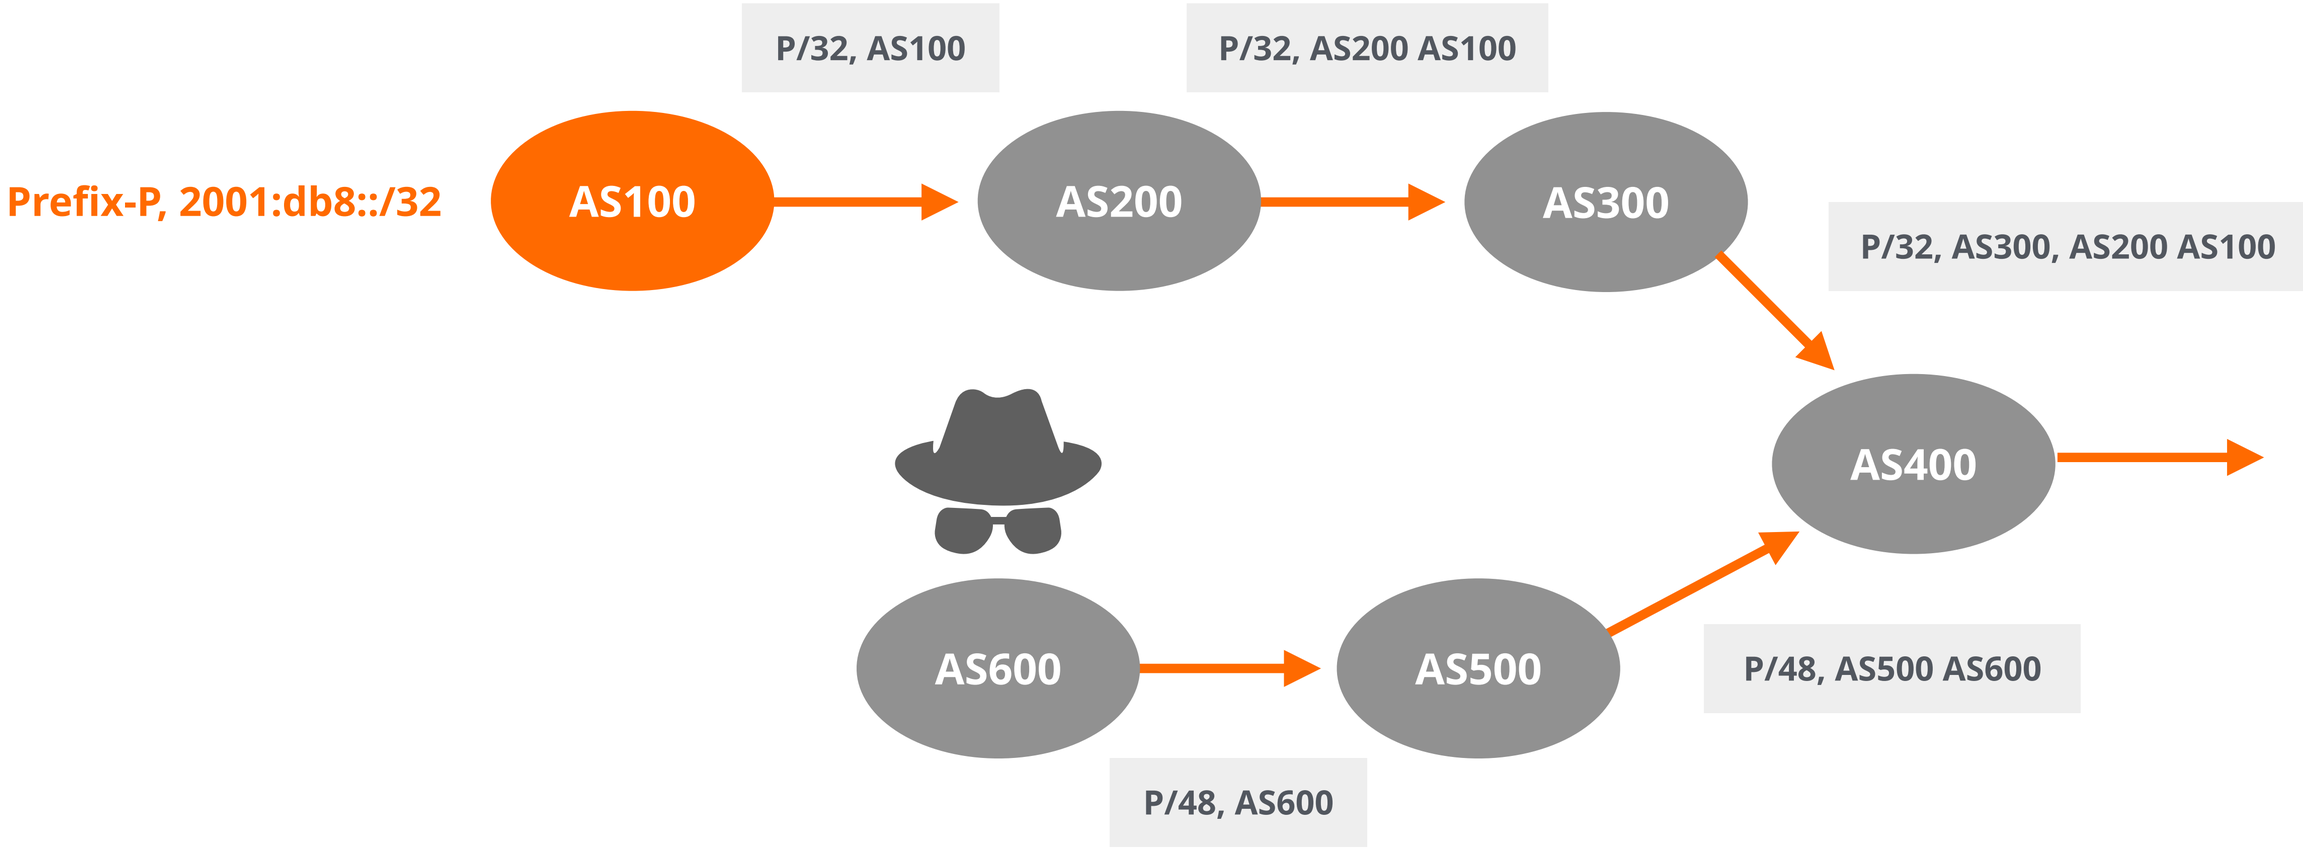

This is a **global hijack!** All traffic for prefix P will be forwarded to the hijacker's network.

It is also possible to hijack the AS path!

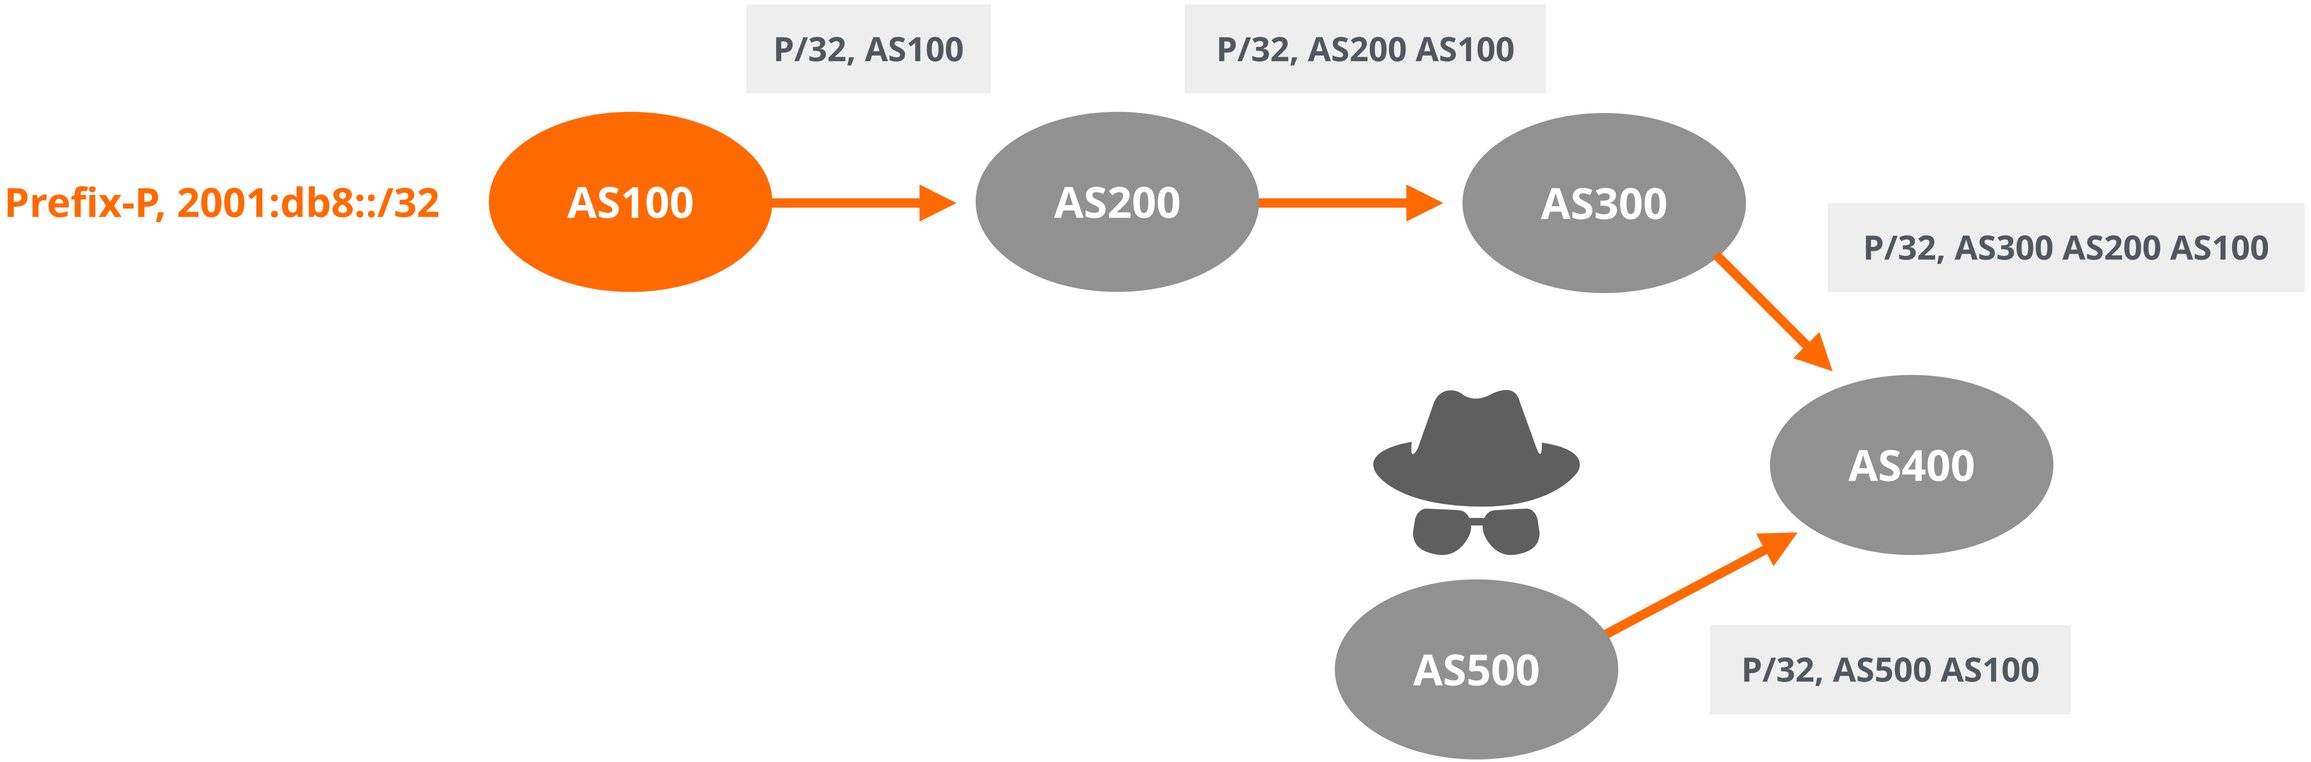

The attacker claims that it has a shorter path to prefix P and hijacks the BGP path!

It happens…

- Because there is no built-in security in BGP!
  - Any AS can announce any prefix
  - Anyone can prepend any ASN to the BGP path
  - BGP announcements are accepted without validation

- Incorrect routing information can be propagated all over the Internet

Sometimes they are just human errors…

- Typo errors
  - Also known as “fat fingers”
  - May cause mis-origination
- Configuration errors

Faulty BGP filter configuration

- Causes routing policy violations
- Unintentional route leaks

AS path prepending mistake

- May cause origin change
- Or forged AS path

A few notable incidents from recent years

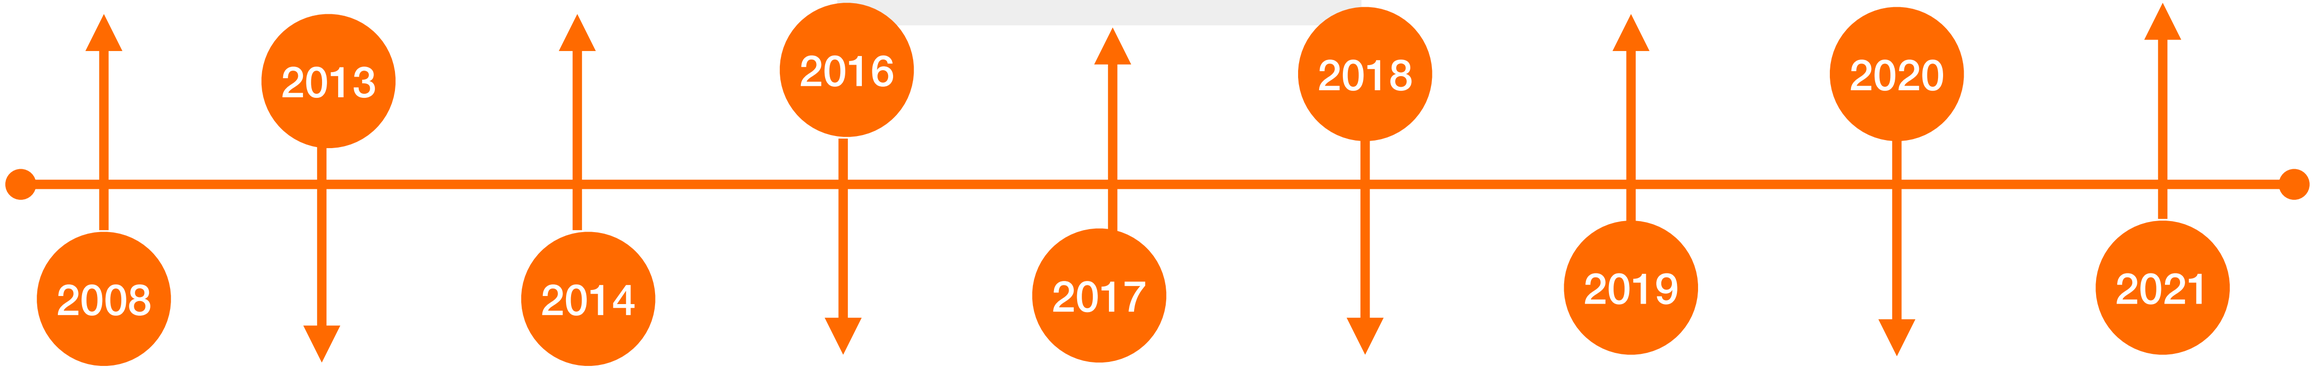

Pakistan Telecom hijacked YouTube

STE hijacked many prefixes including US DoD, YouTube, Level 3, Akamai...

Google, Apple, FB, Microsoft, Riot Games prefixes were hijacked

Taiwan Network Information Center Public DNS in Taiwan was hijacked

AS55410 (Vodafone Idea) hijacked almost 40k prefixes, belonging to 4k+ ASNs in 164 countries.

Ukranian Telco (Vega) hijacked BT prefixes

AWS Route Leak

Amazon Route 53 Hijack myetherwallet.com cryptocurrency hijack

Google Prefix Leak

Telstra hijacked 500 prefixes

>8k BGP prefixes hijacked, affected companies such as Amazon, Akamai, Alibaba

April 2021: BGP hijack by Vodafone Idea, AS55410

- What happened?
  - 34,000+ prefixes hijacked!
  - Impacted major network operators, cloud and CDN providers
  - 13 times more traffic than usual
- Why did it happen?
  - Caused by wrong advertisement
  - Lack of good filtering by upstream providers

April 2020: Akamai, Amazon and Alibaba

- What happened?
  - 8k+ routes hijacked by Rostelecom (AS12389)
  - 200+ CDNs and cloud providers impacted
  - Not known how much data leaked
- Why did it happen?
  - Malicious activity
  - Lack of good filtering by upstream providers/peers

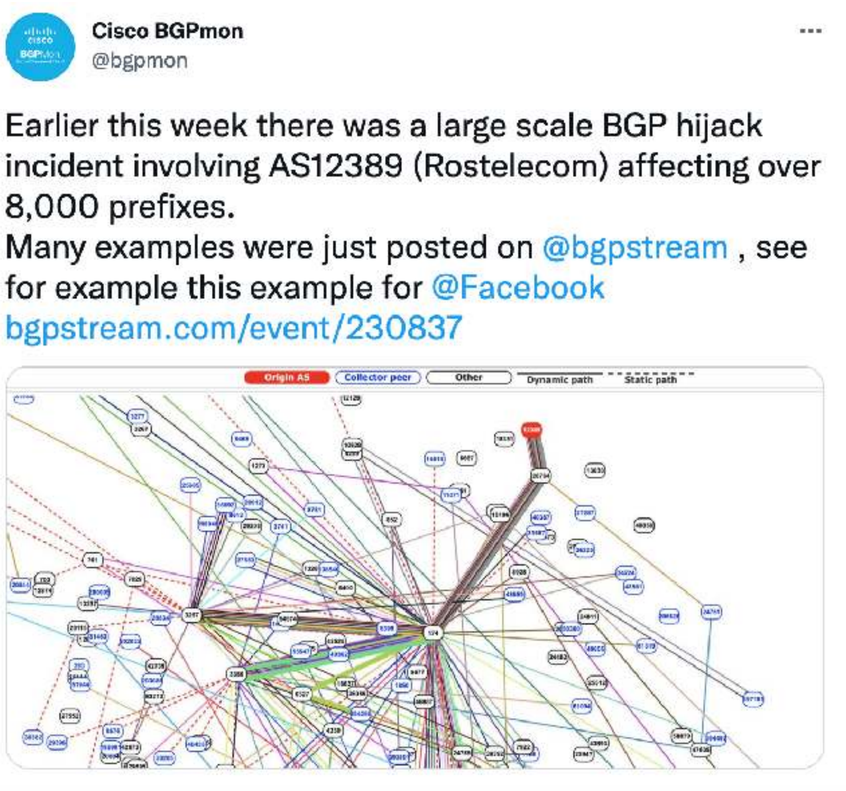

April 2018: Amazon - MyEtherWallet

- BGP hijack of Amazon DNS
- How did it happen?
- Why?
  - Attack to steal cryptocurrency

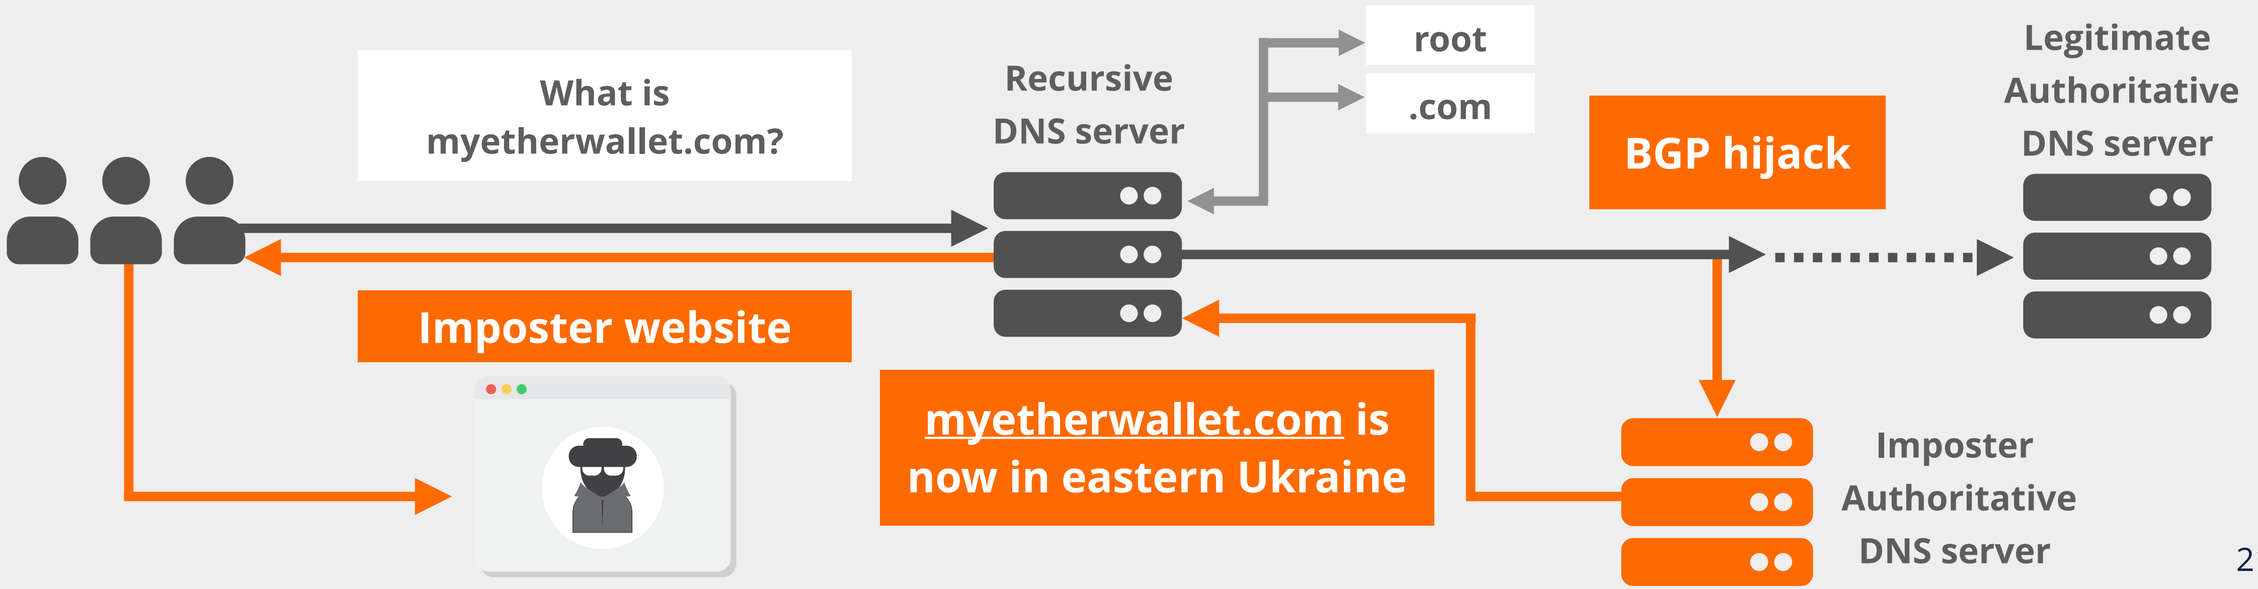

November 2018: Google prefix leak

- MainOne leaked Google routes to CT
- CT propagated them to several transit ISPs
- Google services (G Suite and Google Search) affected by the leak
- Due to misconfigured filters

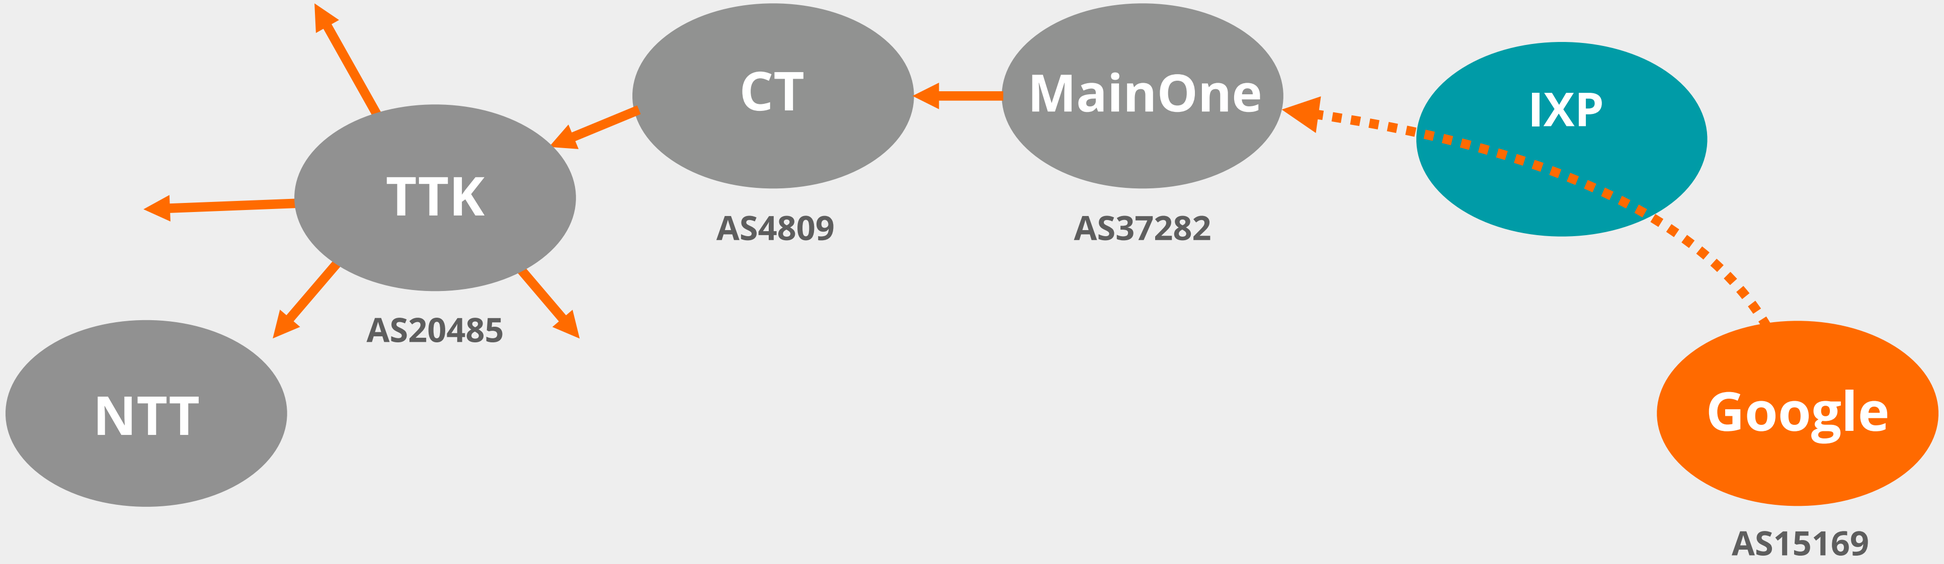

April 2016: AWS route leak

- Private AS originated Amazon's prefixes, but more specific
- Innofield leaked these routes to its upstream
- No big impact because most ISPs didn't accept the bogus route
- Caused by misconfigured route optimiser

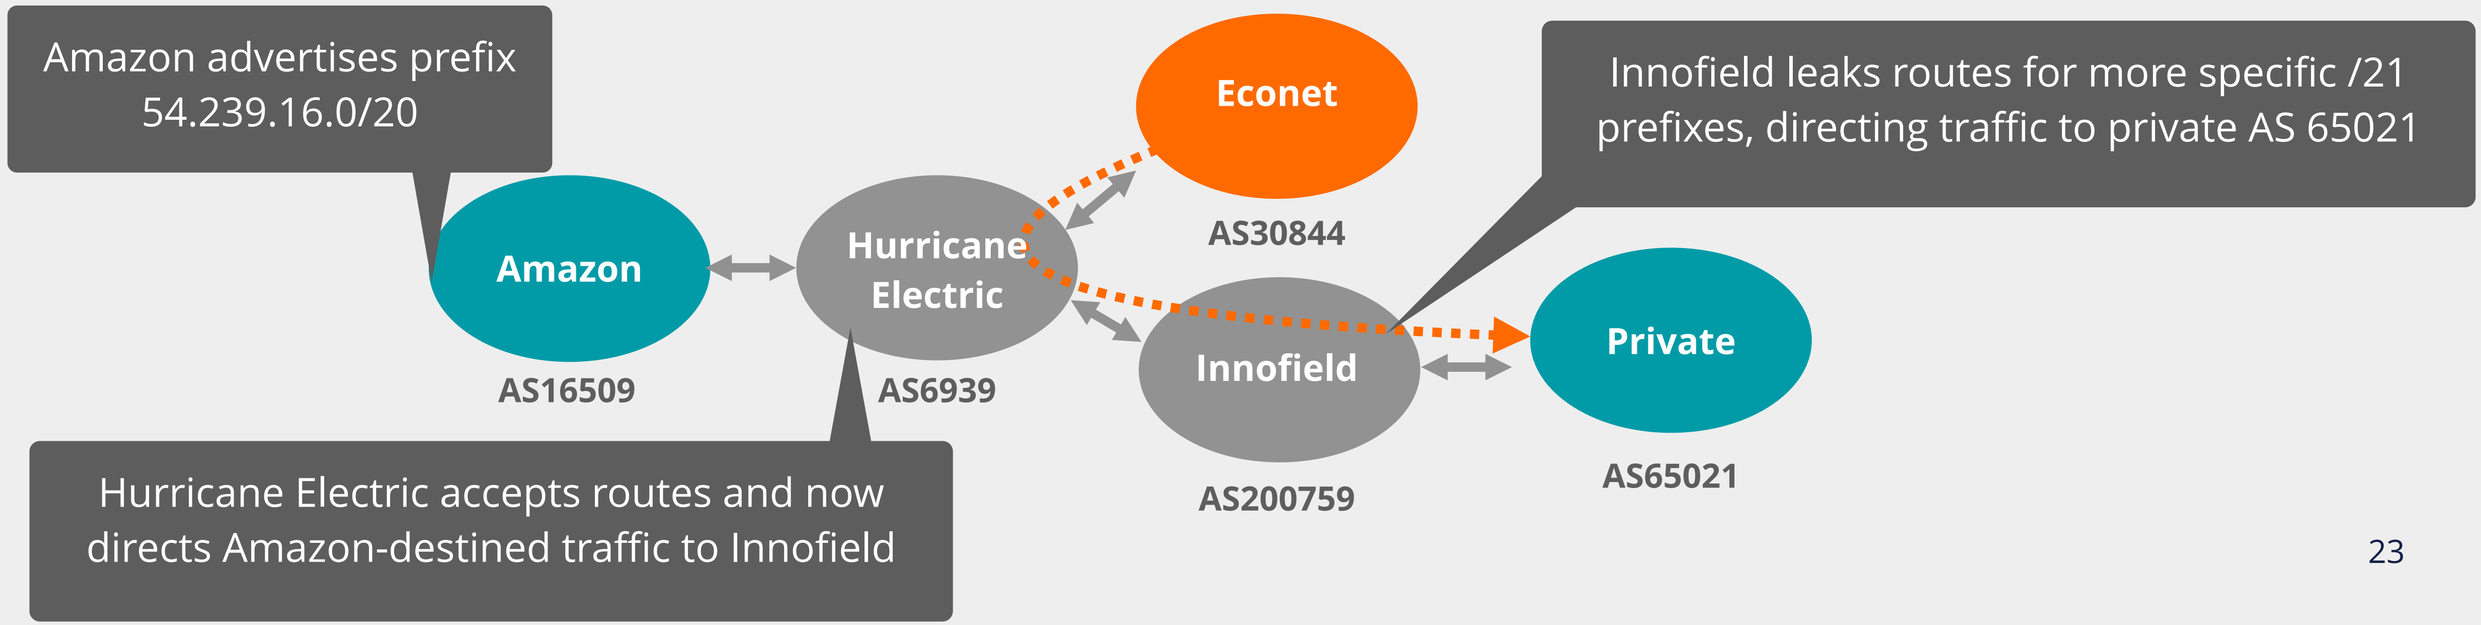

In order to secure routing…

- We need to **verify the routing information**
  - Has the announced prefix been **originated by the legitimate holder**?
  - Has someone **tampered with the AS path** of the BGP update?
- Prevent propagation of **incorrect routing information**

But how?

1. **Check prefixes** before announcing

2. **Register** your routing information in **IRRs**

3. **Filter** BGP routes from your peers, customers and upstreams

4. Implement BGP filters based on **verifiable information**

Concerns with the IRR system

1

Not globally deployed

Just distributed databases

2

No central authority

Who will verify the accuracy of the data?

3

No verification of holdership

Anyone can input anything

4

Not updated properly

Information is missing, outdated or incorrect

As a result…

IRRs are **not so accurate**

Data in IRRs is **incomplete**

They’re not well-maintained

IRR filters are good **only if the IRR entries are correct!**

Routing Security with RPKI

What is RPKI?

What is RPKI?

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

How does RPKI help with routing security?

- Verifies the association between resource holders and their Internet number 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 is it different than the IRR system?

- RPKI is based on RIRs as Trust Anchors
  - RIRs have control over the accuracy of registered data

- Cryptography is used to verify the holdership
  - Provides data you can trust

How does it work?

How does it work?

- RPKI attaches a digital certificate to IP addresses and AS numbers

- Digital signatures authorise the use of resources
  - Private key to sign, public key to validate

How to provide trust in RPKI?

- It relies on the 5 RIRs as Trust Anchors
- Certificate structure follows the RIR hierarchy
- RIRs issue certificates to resource holders

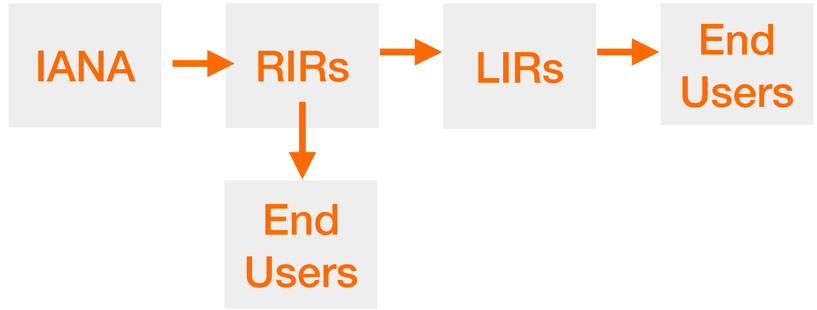

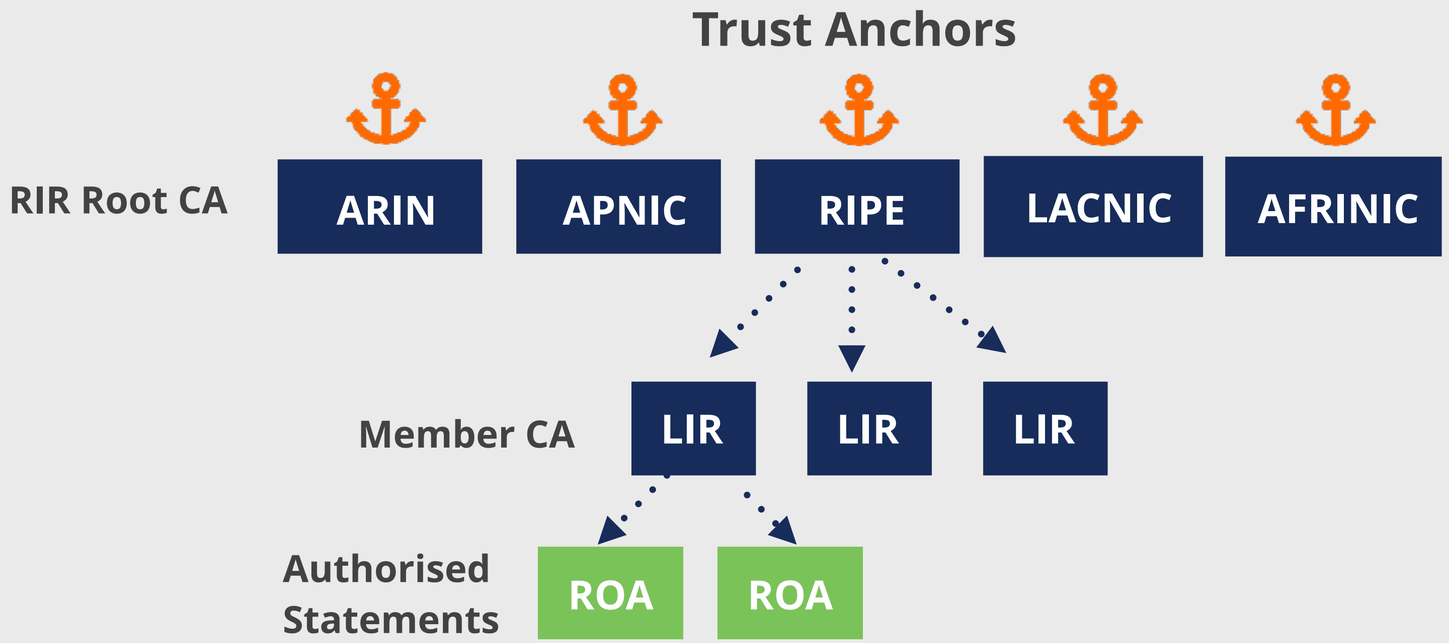

Root Certificate

- RIRs have a **self-signed** root certificate for all resources (0/0 for IPv4, ::/0 for IPv6)
- This signs the resource certificates for all member allocations

LIR Certificate

- Resource certificate for member allocations
- Signed by root’s private key
- Binds LIR’s resources to LIR’s public key
- Proves legitimate holdership for the LIR’s resources

Authorised Statement

- Called as ROA (Route Origin Authorisation)
- Cryptographically signed object
- Signed by LIR’s private key

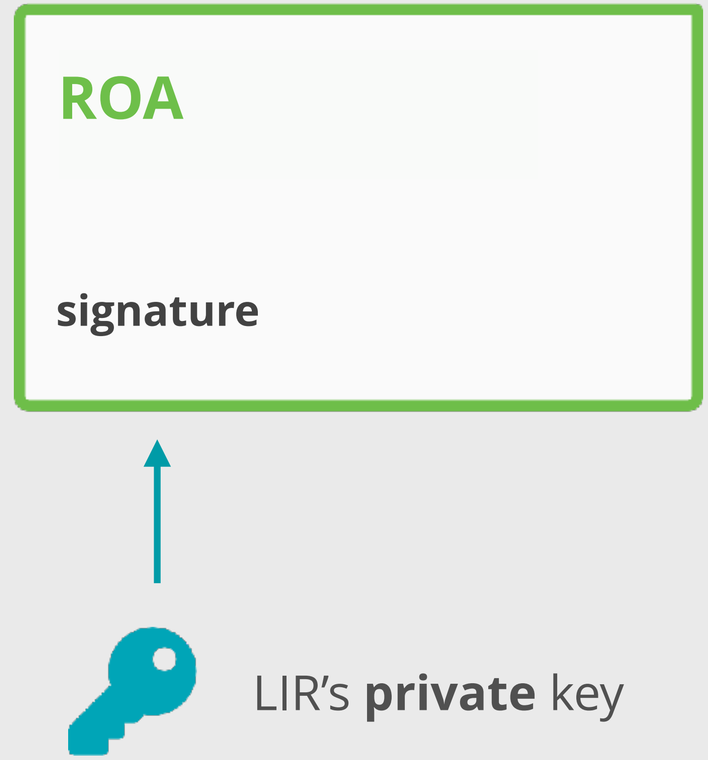

RPKI Chain of Trust

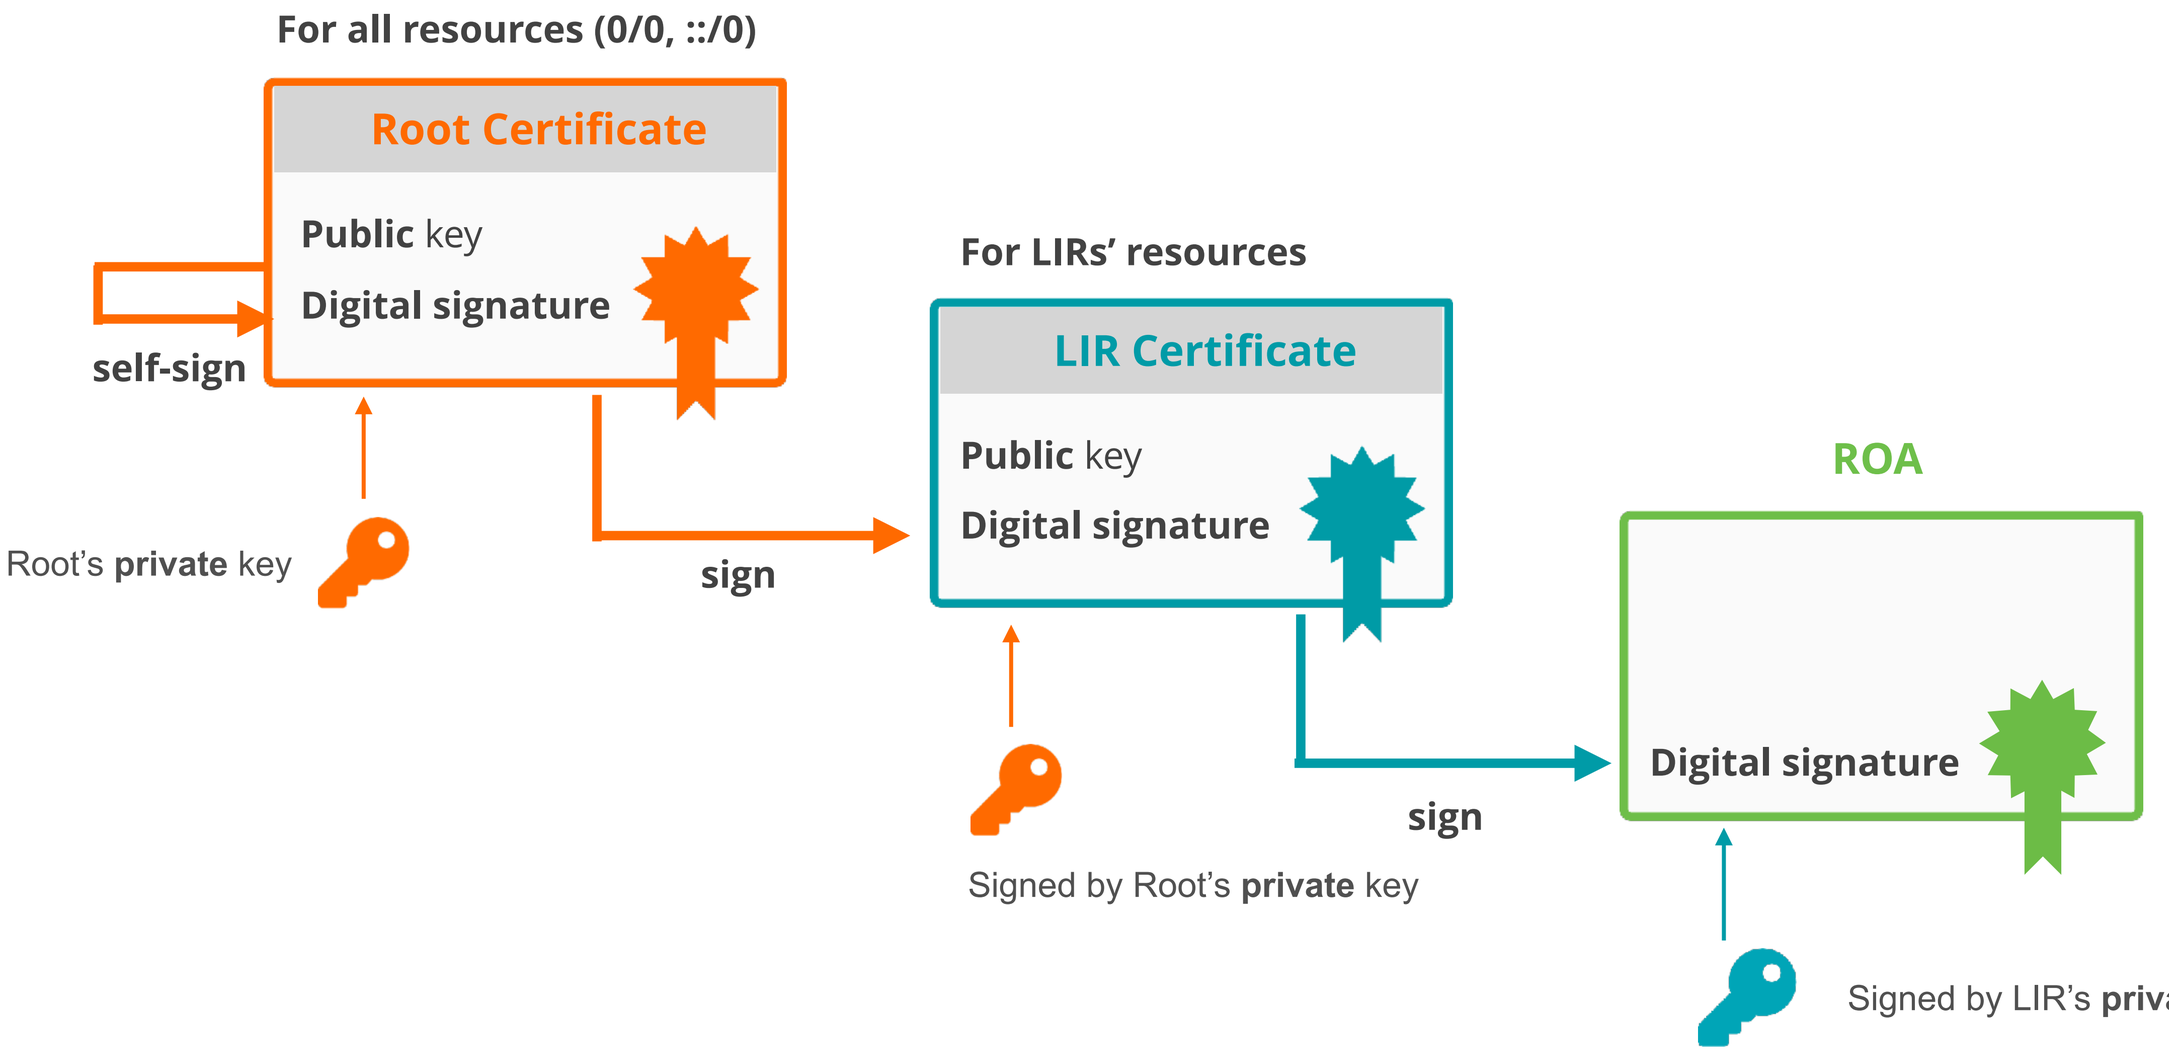

Route Origin Authorisation (ROA)

- Contains a list of address prefixes and an AS number
- LIRs can create a ROA for their resources
- Multiple ROAs can exist for the same prefix
- ROAs can overlap

ROA

Prefix	**2001:db8::/48**
Max Length	**/48**
Origin AS	**AS65536**

Prefix | **Origin ASN** | **Max Length**

2001:db8::/48

The network for which you are creating the ROA

Route Origin Authorisation (ROA)

- Contains a list of address prefixes and an AS number
- LIRs can create a ROA for their resources
- Multiple ROAs can exist for the same prefix
- ROAs can overlap

ROA

Prefix	**2001:db8::/48**
Max Length	**/48**
Origin AS	**AS65536**

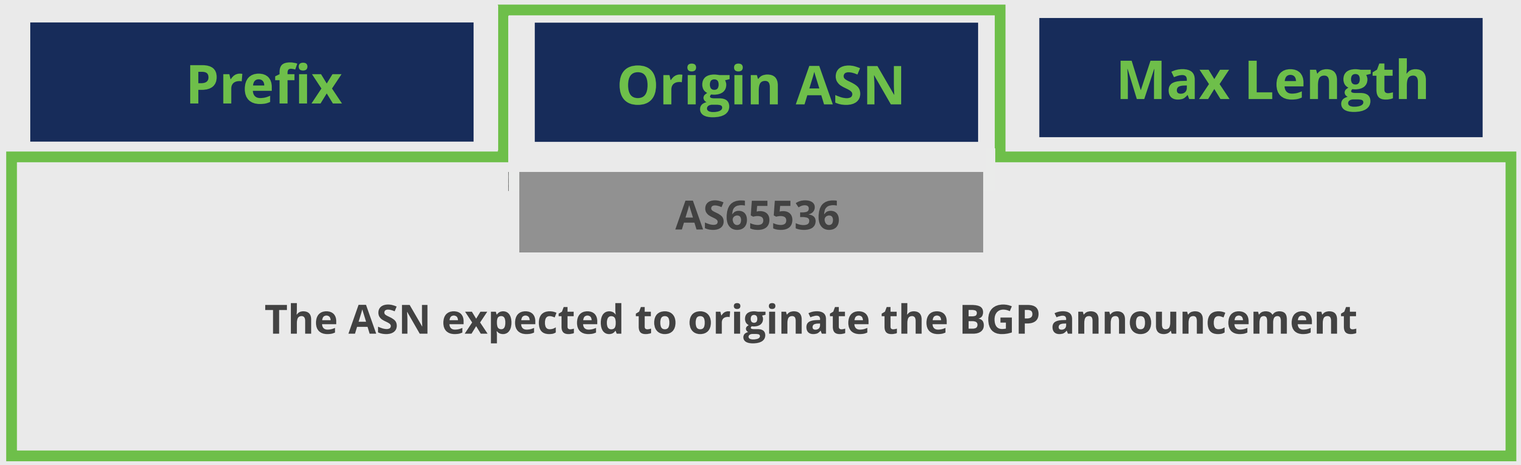

Route Origin Authorisation (ROA)

- Contains a list of address prefixes and an AS number
- LIRs can create a ROA for their resources
- Multiple ROAs can exist for the same prefix
- ROAs can overlap

ROA

Prefix	**2001:db8::/48**
Max Length	**/48**
Origin AS	**AS65536**

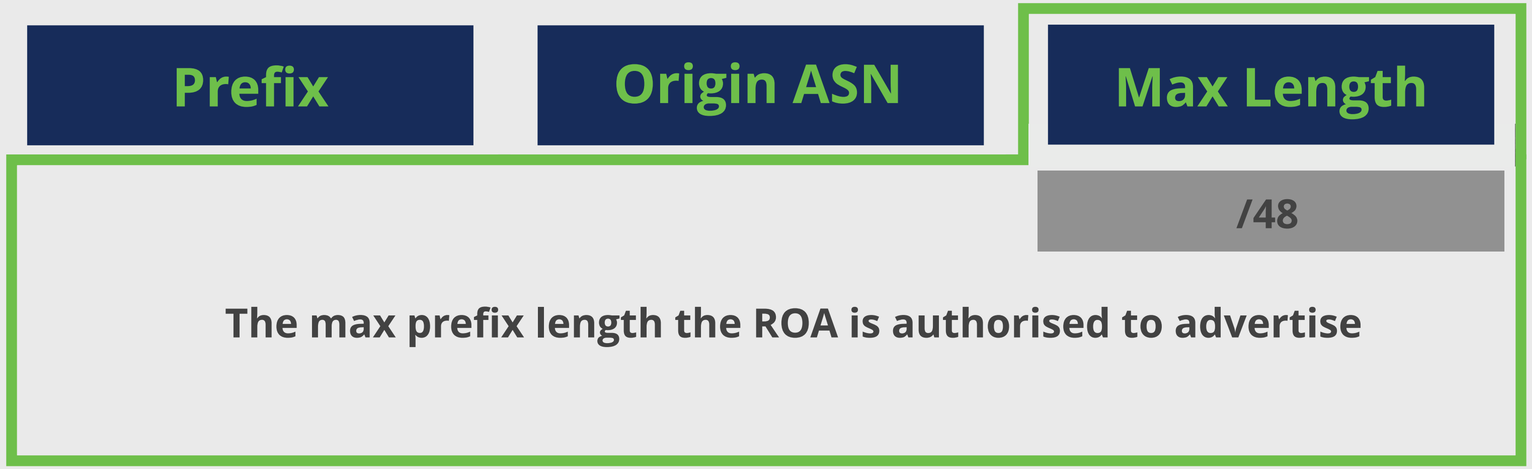

Routing Security with RPKI

Building Blocks of RPKI

Elements of RPKI

- The RPKI system consists of two parts

SIGNING + **VALIDATION**

SIGNING

Create ROAs for your prefixes in the RPKI system

RIPE NCC RPKI Dashboard

BGP Announcements and Route Origin Authorizations (ROAs)

Last BGP import: 3 hours and 24 minutes ago

BGP Announcements: 0 Invalid: 0 Unknown: 0 Valid: 0

ROAs: 2 Causing invalid announcements: 0 Ok: 2

Alert Configuration

Configure alert recipients and notification preferences.

⚠ No recipients are configured.

History

Time (UTC)	User	Summary
01/11/2024, 15:42:20	dmelnik@ripe.net	Updated ROA configuration. Additions: [asn=AS2121, prefix=193.0.24.0/21, maximumLength=21], [asn=AS2121, prefix=2001:67c:64::/48, maximumLength=48]. Deletions: none.

SIGNING

Create ROAs for your prefixes in the RPKI system

RIPE NCC RPKI Dashboard

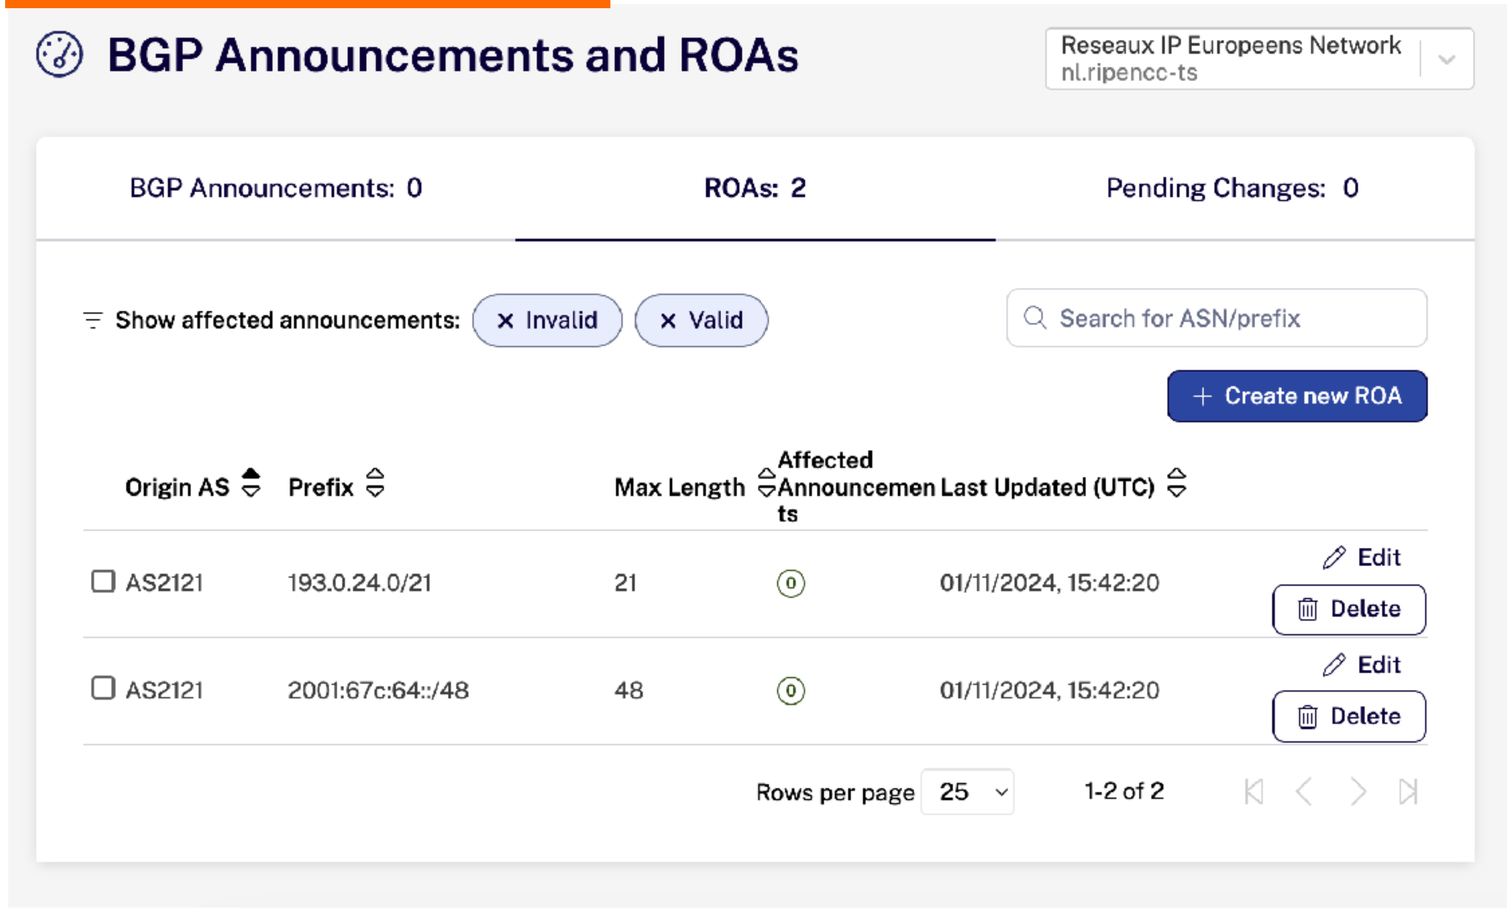

Publication

RIPE NCC RPKI Repository

VALIDATION

Verify information provided by others

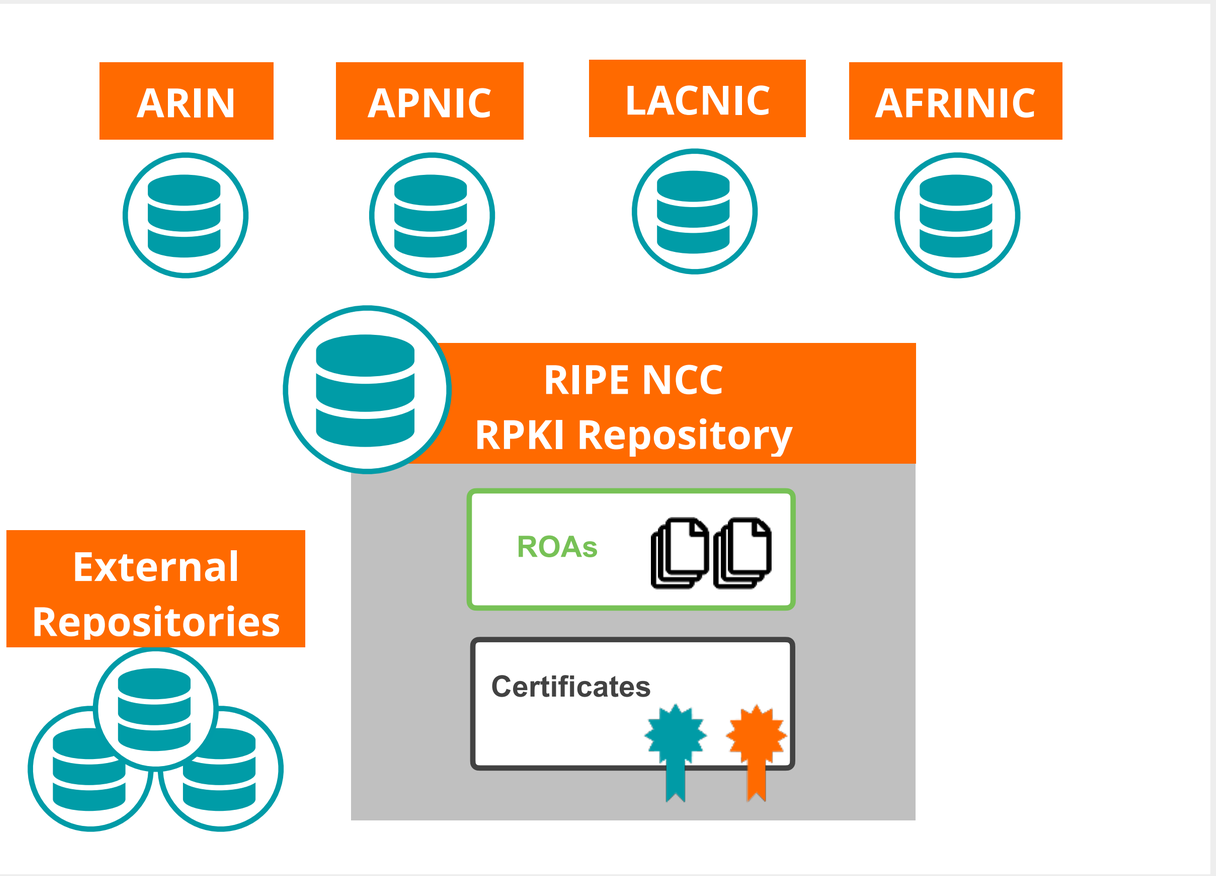

rsync/RRDP

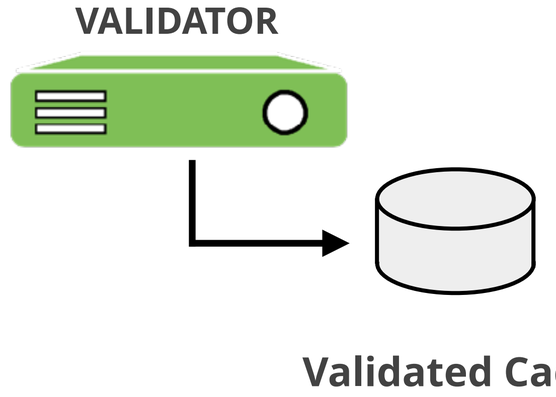

Validation of ROAs

- ROAs are validated by a **validator**, also known as “relying party software”
  - Validates the **chain of trust** and builds a **“validated cache”**
  - Routinator, Fort, rpki-client, etc.

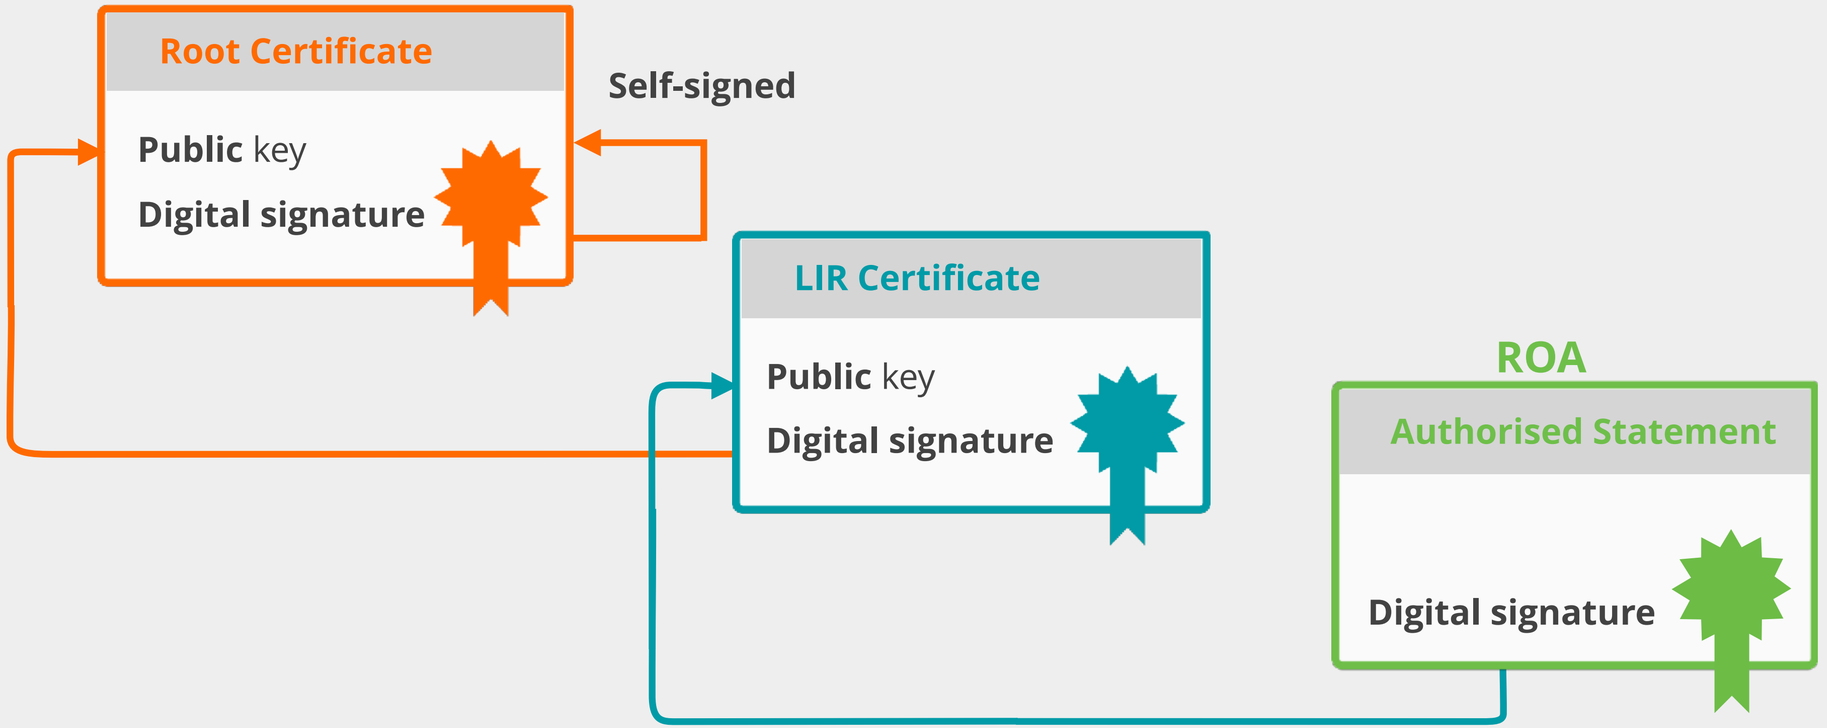

VALIDATION

Verify information provided by others

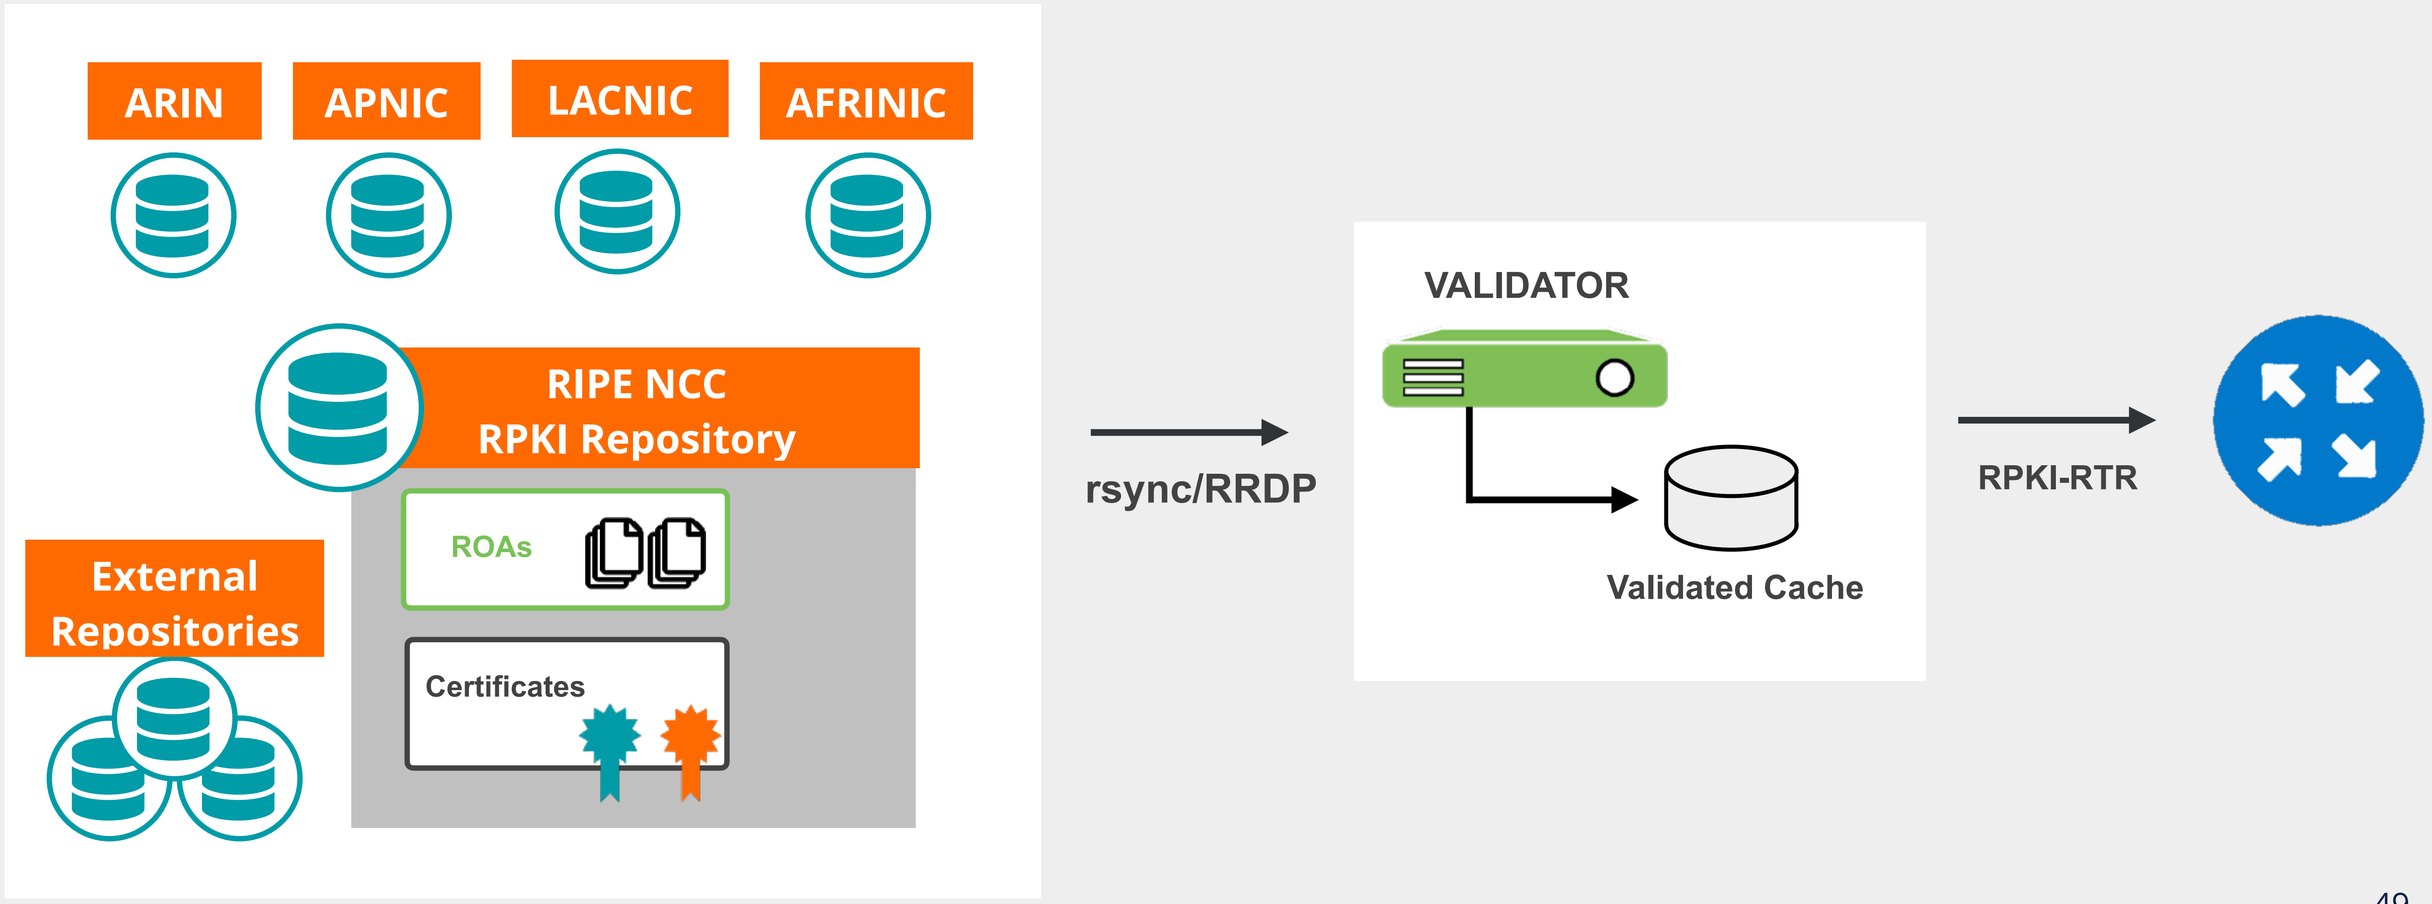

RPKI has two implementations

- **Hosted** RPKI
  - RIRs host CAs for LIRs
  - Automated signing and key rollovers
  - Information published in RIR repository

- **Delegated** RPKI
  - LIR manages full RPKI system
  - Runs its own CA, manages its own keys/key rollovers
  - Creates ROAs in its own platform

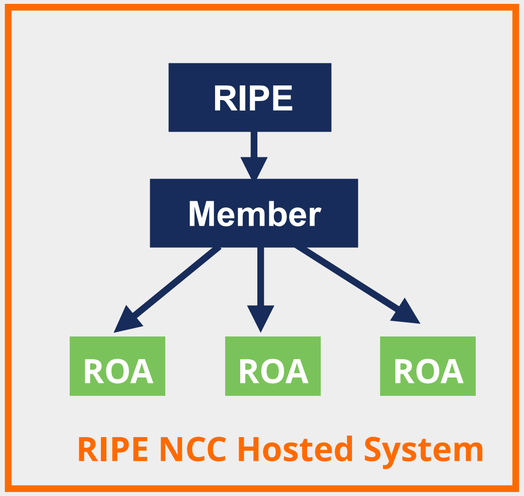

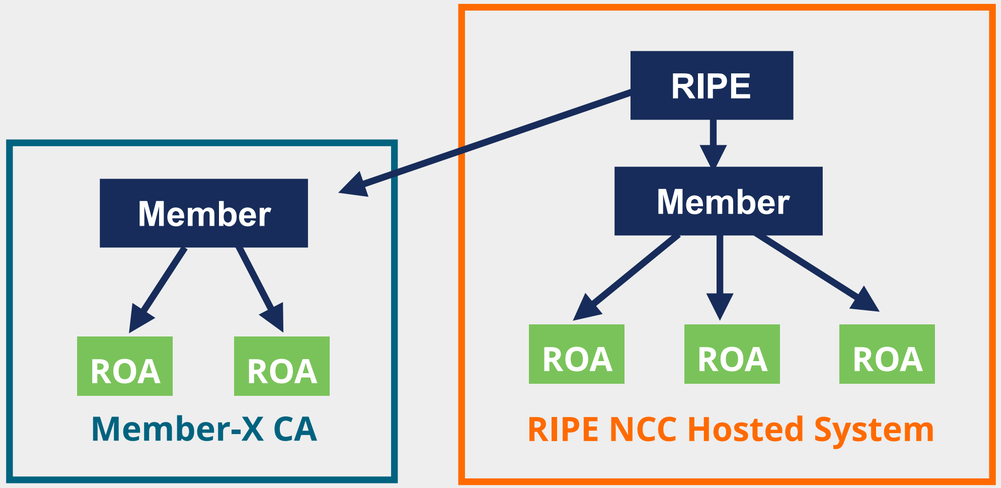

Which RPKI implementation should I choose?

Hosted RPKI

- Easy to implement
  - Request LIR certificate
  - Create your ROAs
- **Recommended option** unless your needs are special
- Everything is managed by RIR
  - Signing, key management, publication, etc.

Delegated RPKI

- Gives **more control**
  - Create ROAs in your own platform and keep in your repository
  - Sign and publish your ROAs
  - Store your keys, manage key rollovers
- Good option if you have **resources from many RIRs**
  - Single system to manage all your ROAs
- Option to delegate to customers

Routing Security with RPKI

BGP Origin Validation (BGP OV)

BGP Origin Validation (BGP OV)

- RPKI-based route filtering
- BGP announcements are compared to the valid ROA
- Origin ASN and Max Length must match!
- Router decides the validation states: Valid, Invalid and Not Found

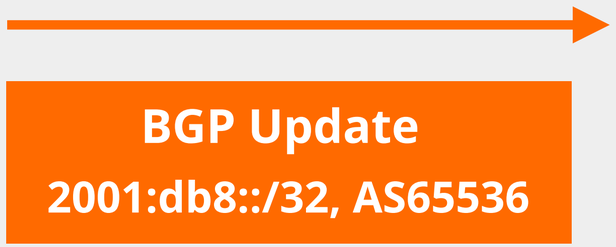

ROA	
Prefix	2001:db8::/32
Max Length	/32
Origin AS	AS65536

How does RPKI validate the origin of BGP routes?

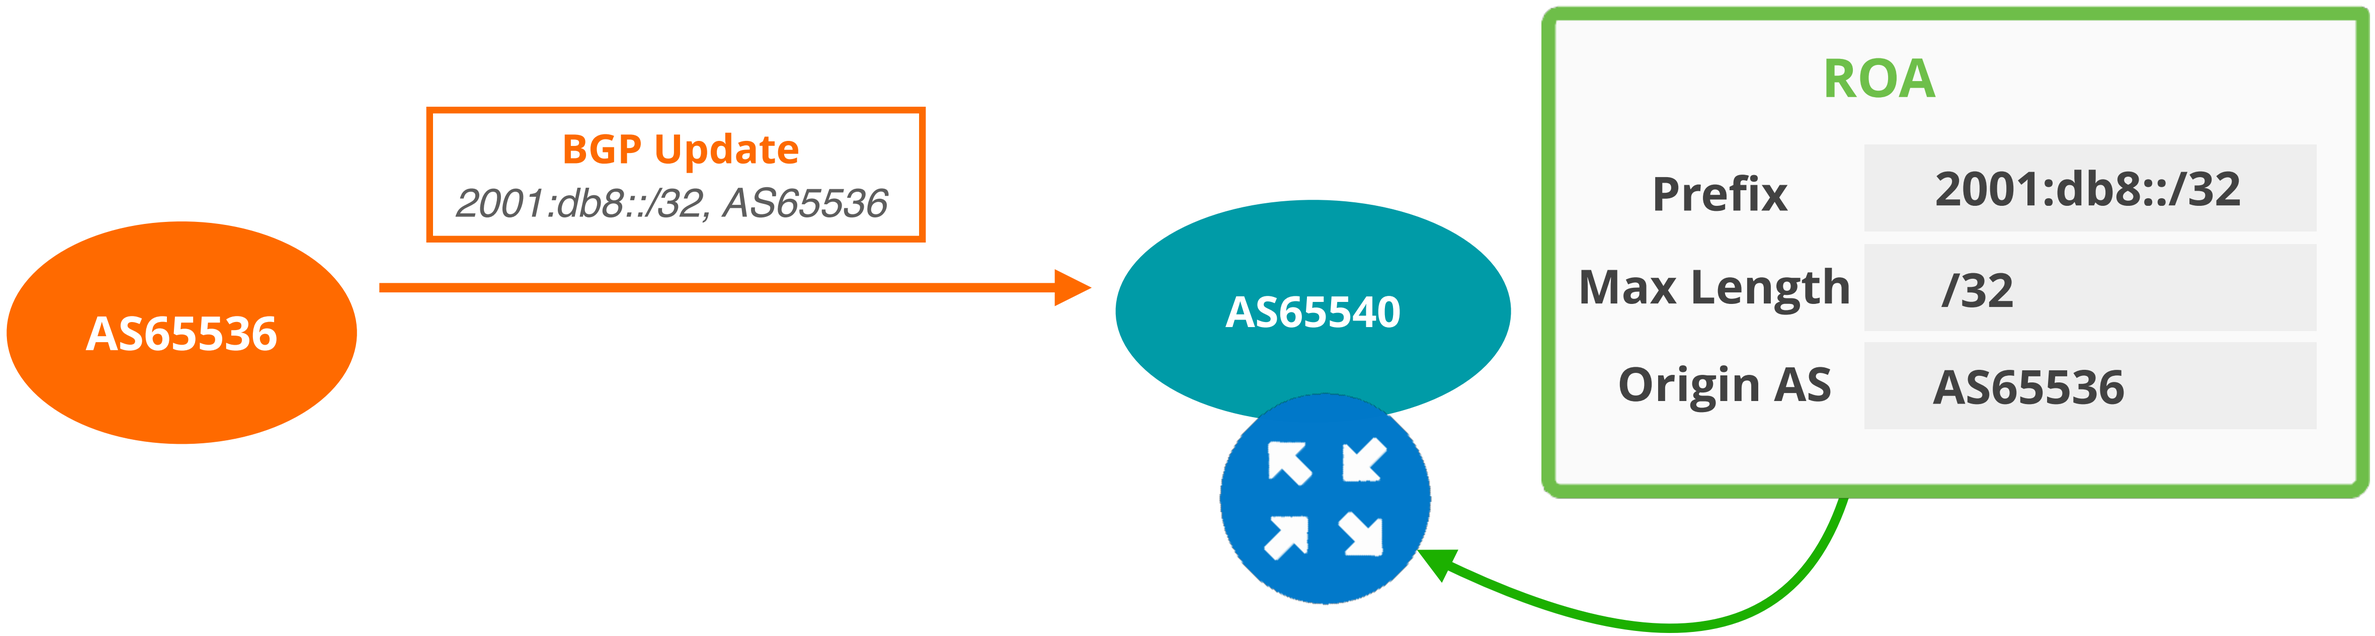

How does RPKI validate the origin of BGP routes?

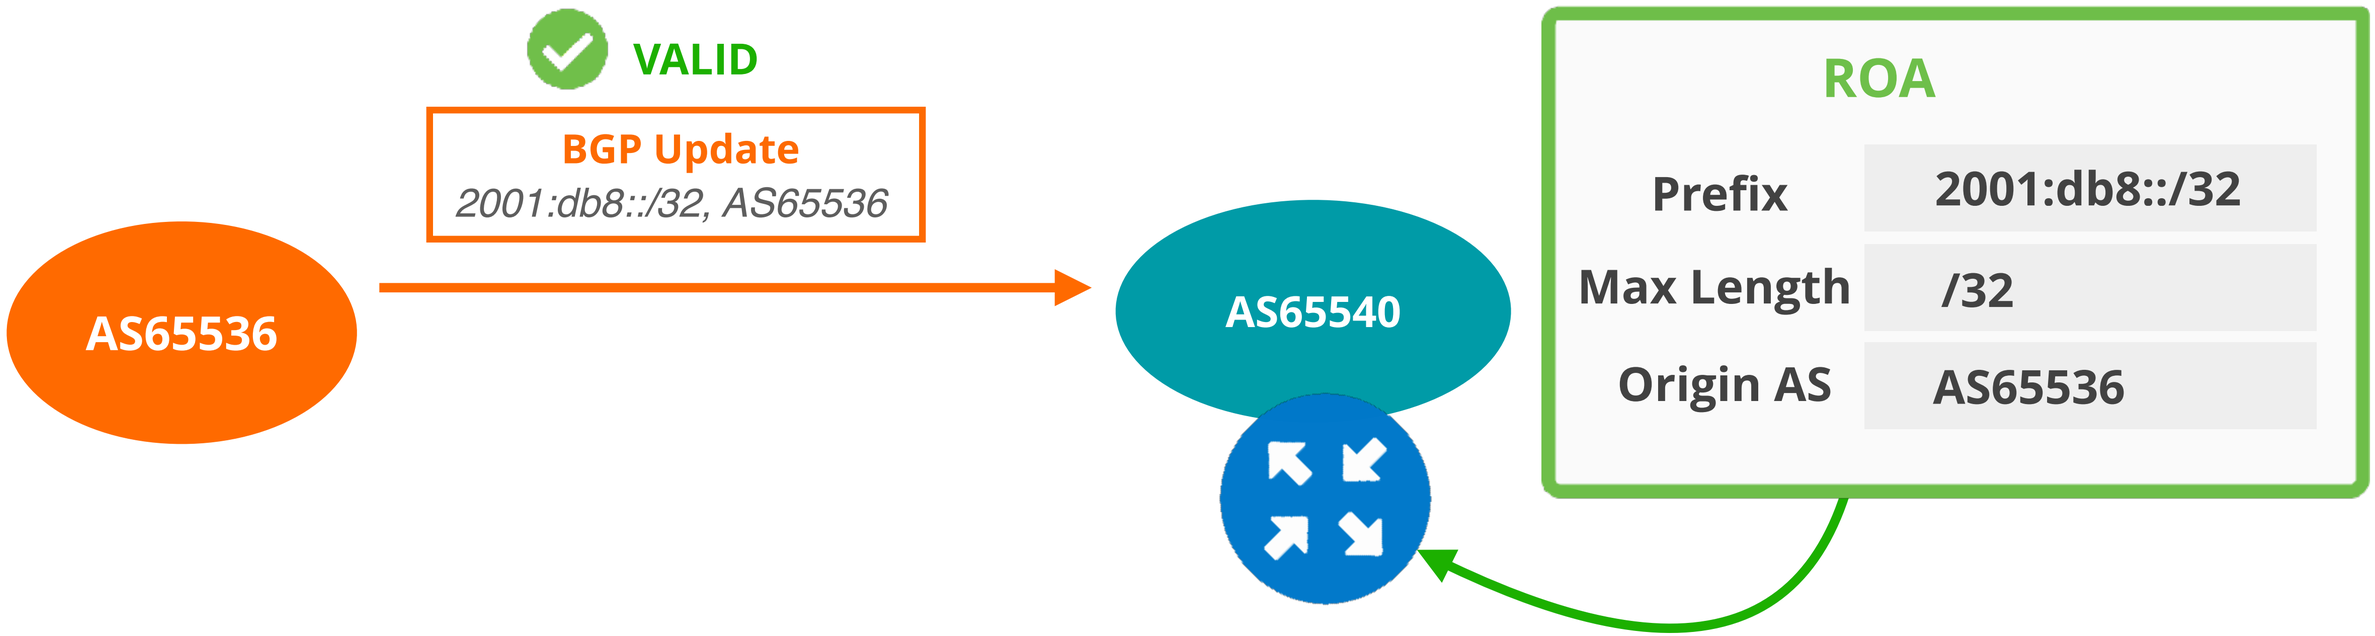

How does RPKI validate the origin of BGP routes?

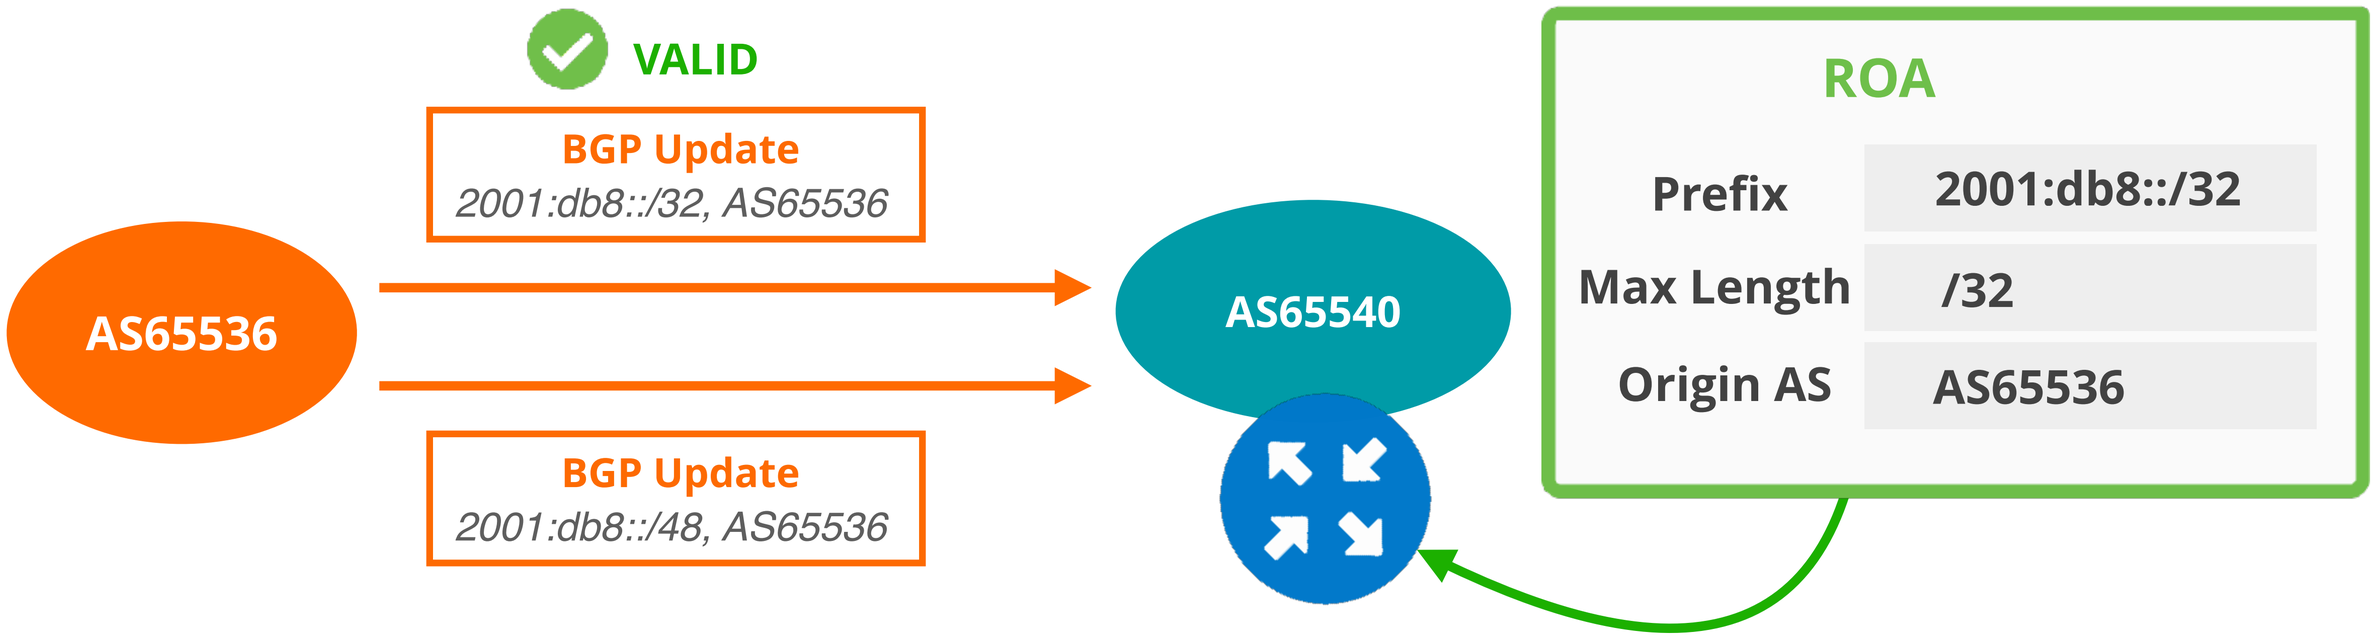

How does RPKI validate the origin of BGP routes?

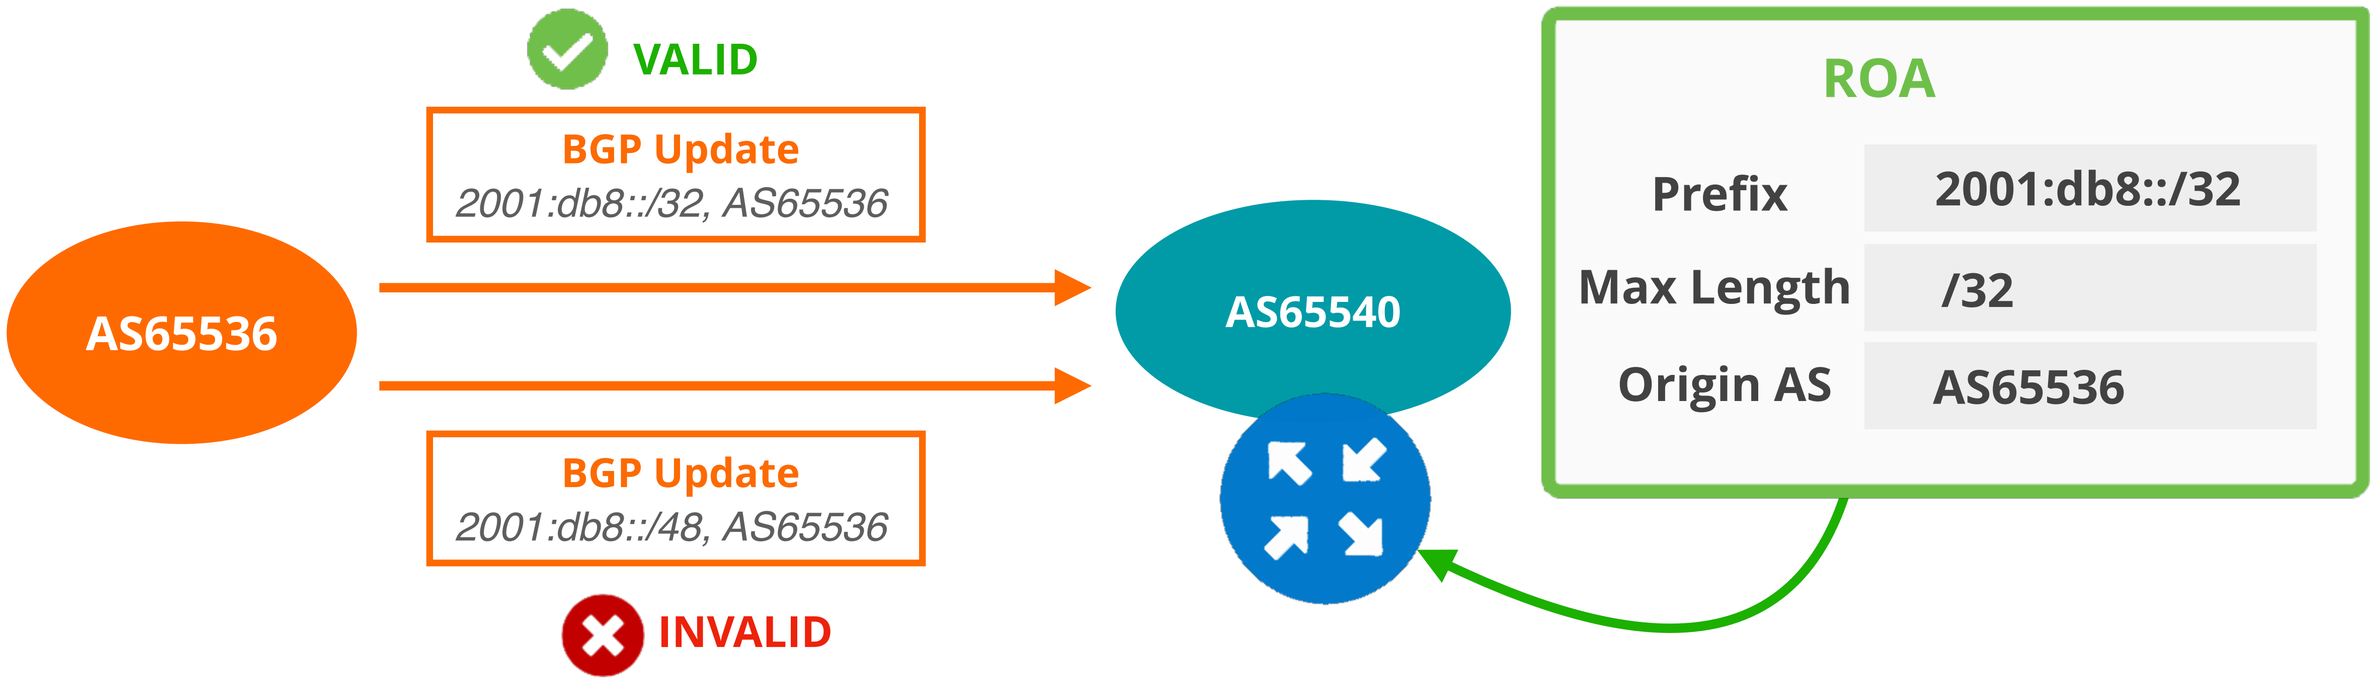

How does RPKI validate the origin of BGP routes?

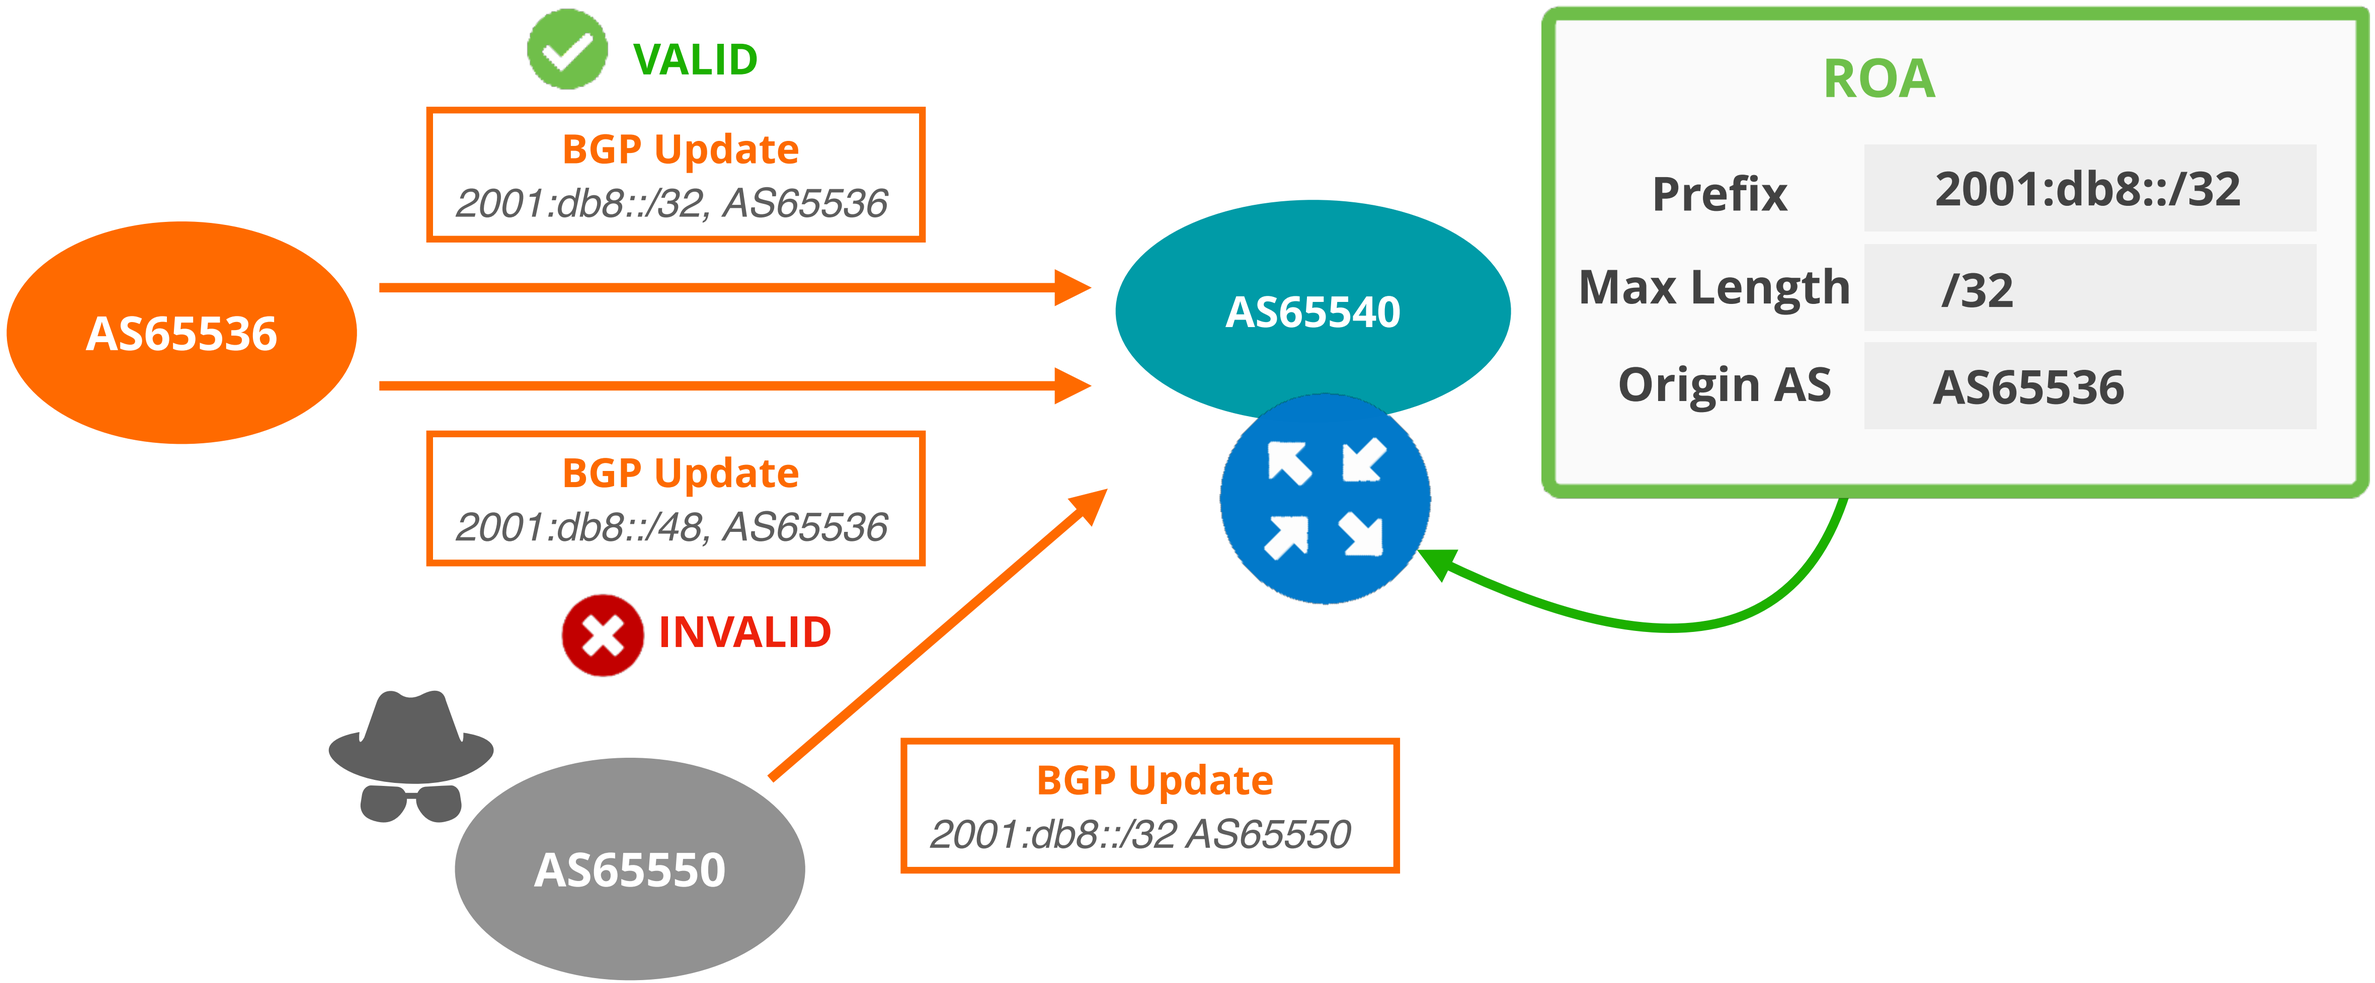

How does RPKI validate the origin of BGP routes?

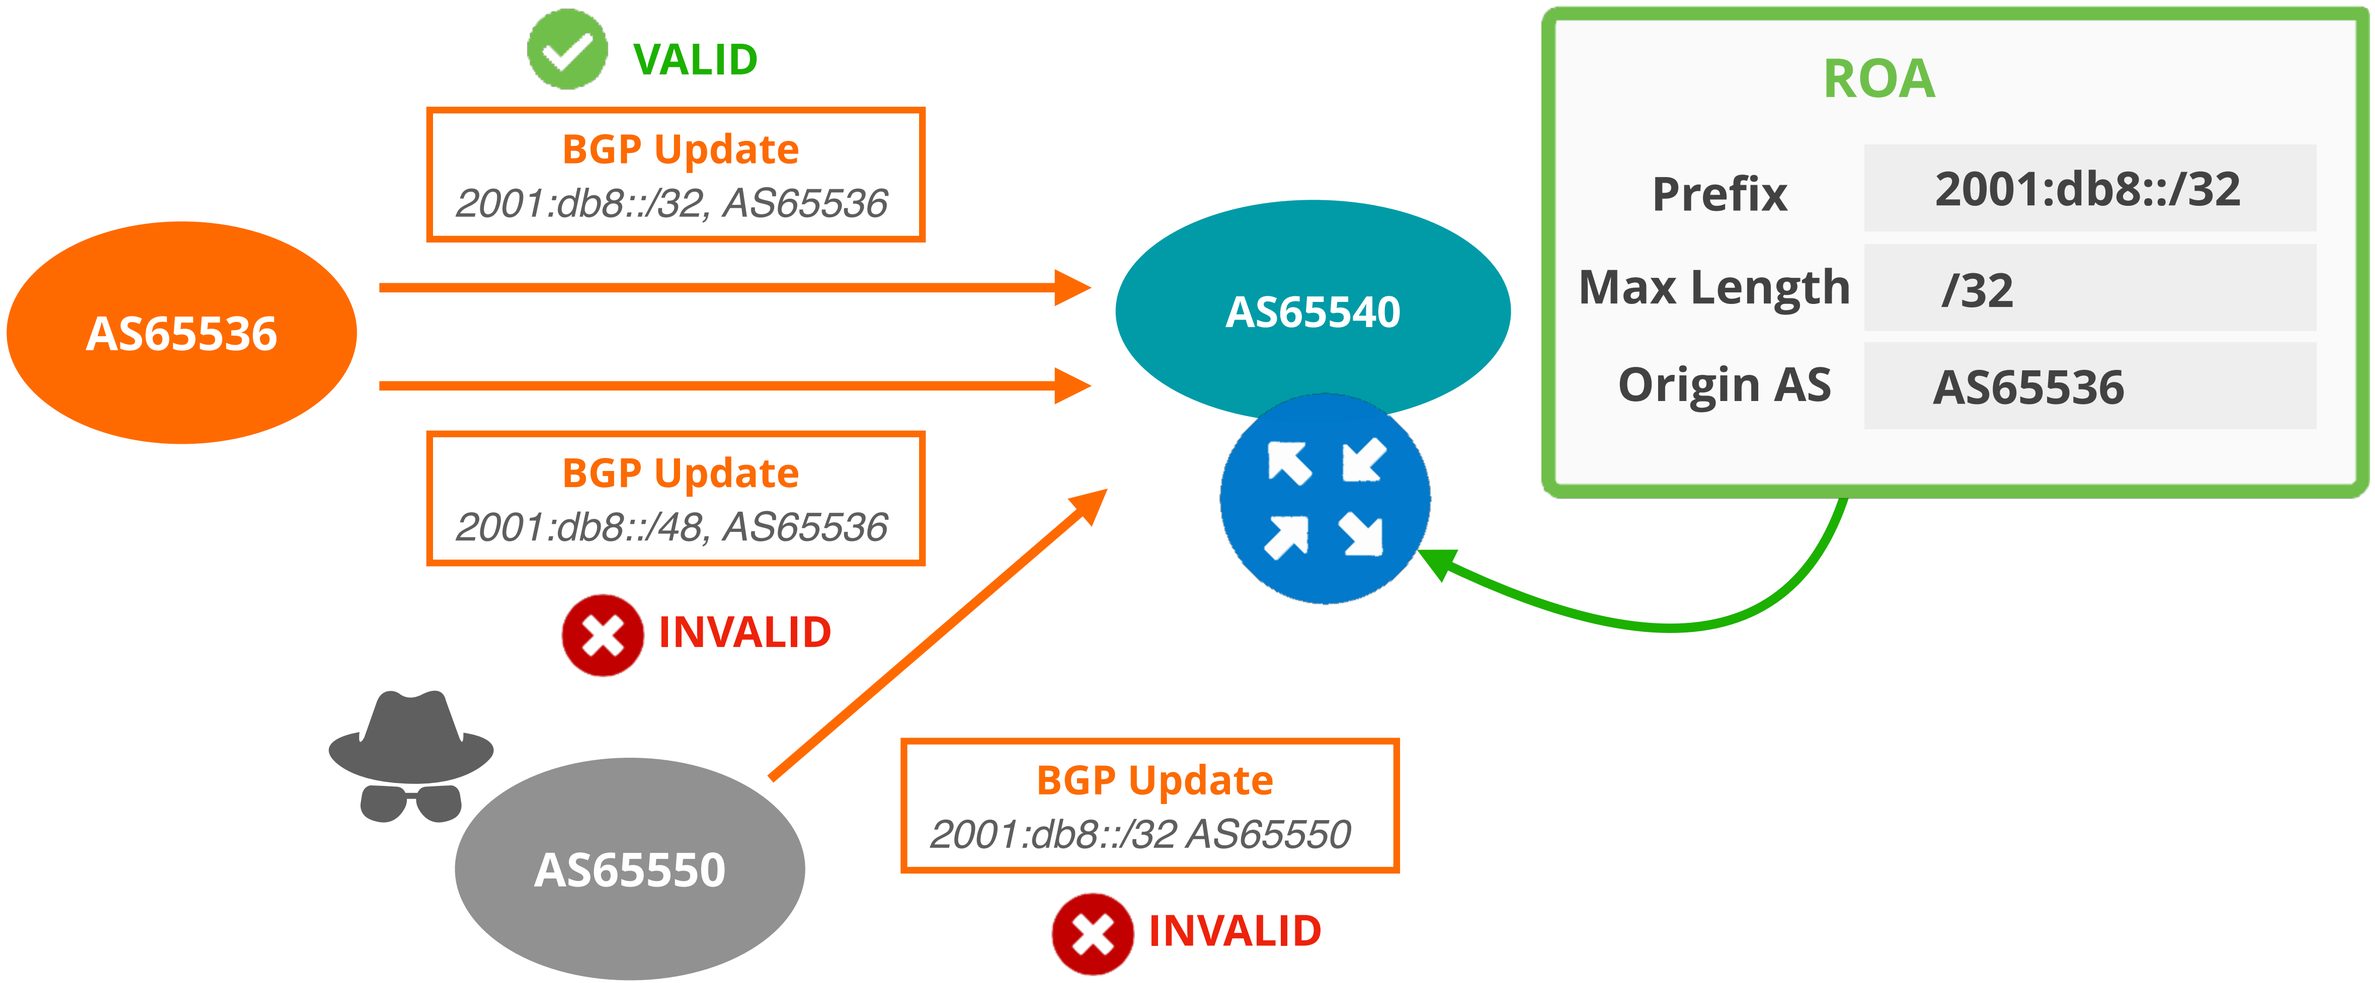

What to do with invalids?

For BGP origin validation to achieve its goal...

- Invalids **should be dropped**!
- Big networks are **already doing it**

Is BGP OV with RPKI enough for BGP security?

- It is only the **first step**
  - can not help if the AS Path is modified (forged origin attacks)
- It is a stepping stone to “**Path Validation**”
- The ultimate goal is to validate the full BGP path by using **RPKI certificates**
  - BGPsec (RFC 8205)
  - ASPA (draft)

A global RPKI ecosystem enhances routing security!

- RPKI is a **powerful mechanism**
  - Prevents BGP hijacks, mis-originations and route leaks
  - Currently used for validating the origin AS
  - Stepping stone to BGP path validation
- RPKI is opt-in
  - It will only work if every network agrees to abide by it
- Currently ~53% of the prefixes on the Internet have ROAs
  - BGP hijacking may cause significant damage unless the majority implements it

Let’s deploy RPKI today!

Give support for secure Internet routing
and
help to mitigate routing incidents globally!

RPKI Test Dashboard

https://localcert.ripe.net/

- 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

Ondrej.Caletka@ripe.net
https://ondřej.caletka.nl

RIPE NCC
Academy

Learn something new today!

academy.ripe.net

RIPE NCC Certified Professionals

https://getcertified.ripe.net/

Ënn

Соңы

An Críoch

پایان

Ende

Y Diwedd

Vége

Endir

Finvezh

վերջ

Koniec

Son

დასასრული

הסוף

Кінець

Кінець

Finis

Lõpp

Amaia

Loppu

Tmiem

Liðugt

Kpaj

Kraj

Sfârşit

النهاية

Конец

Slutt

Konec

Fund

Fine

Fin

Einde

Fí

Край

Beigas

Τέλος

Fim

Slut

Pabaiga