RIPE NCC
RIPE NETWORK COORDINATION CENTRE

The whois Protocol for Internet Routing Policy

or: how plaintext retrieved over TCP/43 ends up in router configurations

Vesna Manojlović & Ties de Kock

38C3, December 2024, Hamburg

https://events.ccc.de/congress/2024/hub/en/event/the-whois-protocol-for-internet-routing-policy-or-how-plaintext-retrieved-over-tcp-43-ends-up-in-router-configurations/

Dependencies...

- Whois is an **old** internet protocol.
- **Two** kinds of whois databases: domain names and internet numbers
- One of the databases for internet routing policy is operated by the RIPE NCC (“ripe database” / IRR)
- These databases feel kind of arcane…
- To use IRR routing policy, the RPSL information needs to end up in BGP router configurations

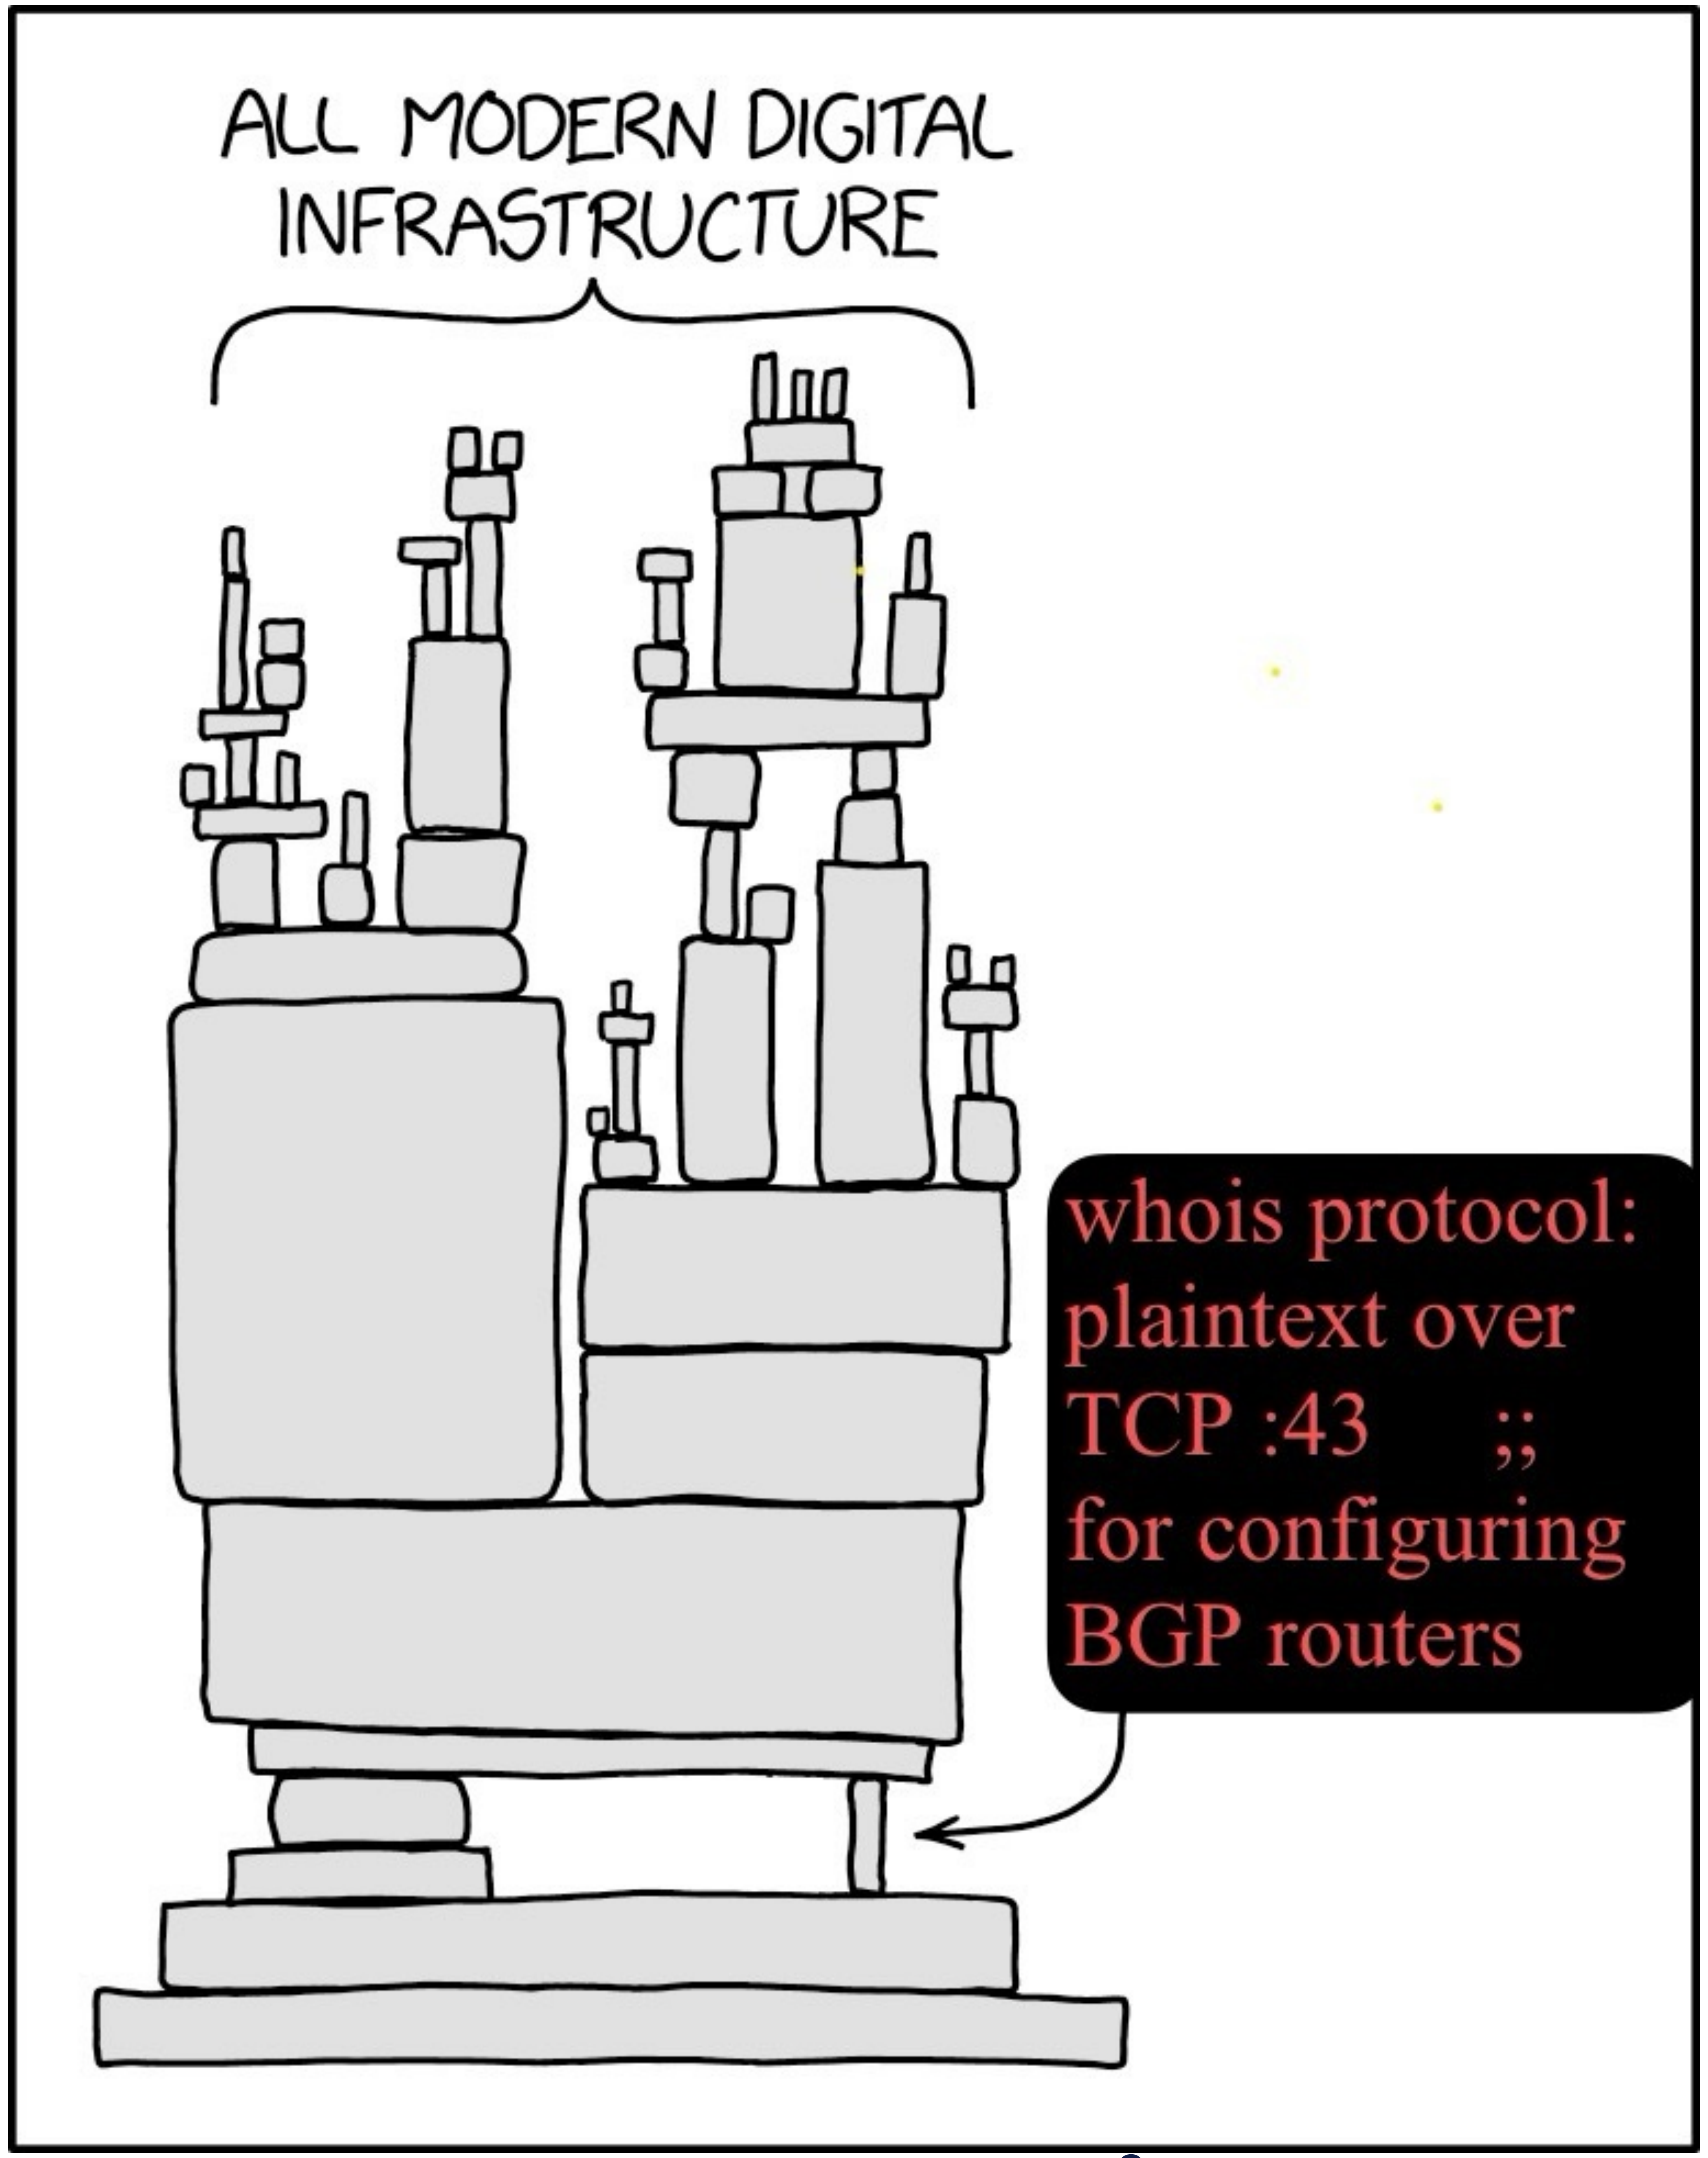

Previously, at CCC…

- CCC Camp 2007
  - "Using RIPE Routing Registry" workshop
  - *https://events.ccc.de/camp/2007/RoutingRegistry/index.html*
  - *https://becha.home.xs4all.nl/routing-registry-bgp-tutorial.pdf*
- #38C3
  - https://events.ccc.de/congress/2024/hub/en/event/bgp-enabled-hackerspaces-or-creatures/
  - https://events.ccc.de/congress/2024/hub/en/event/personal-autonomous-system-as-owner-operator-meetup/
  - https://events.ccc.de/congress/2024/hub/en/event/community-network-meetup/

whois Protocol History

- RFC3912
  - https://datatracker.ietf.org/doc/html/rfc3912
  - Original RFC812 (1982!)
- CLI clients included in every OS
- Servers / databases operated by registrars (domain names) & registries (RIR) & 3rd parties
- Port 43!

```
[becha@becha-pro ~ % whois -h whois.ripe.net 151.217.0.0
inetnum:        151.217.0.0 - 151.217.255.255
netname:        DE-CCC-20241127
country:        DE
geofeed:        https://geoloc.bad.network/as13020/geoloc.csv
remarks:        Geofeed https://geoloc.bad.network/as13020/geoloc.csv
remarks:        ==================================================
remarks:        === ===
remarks:        === If you have trouble with users from ===
remarks:        === this network, please contact ===
remarks:        === ===
remarks:        === ABUSE MAIL: abuse@ccc.de ===
remarks:        === ===
remarks:        === In case of urgency you can also ===
remarks:        === contact our abuse hotline: ===
remarks:        === ===
remarks:        === ABUSE HOTLINE: +49 40 401801-666 ===
remarks:        === ===
remarks:        ==================================================
abuse-c:        CCC-RIPE
org:            ORG-CCCE3-RIPE
admin-c:        CCC-RIPE
tech-c:         CCC-RIPE
status:         ASSIGNED PI
remarks:        Temporary assignment (start date: 2024/11/27, end date:
mnt-by:         CHAOS-MNT
mnt-by:         RIPE-NCC-END-MNT
created:        2024-11-27T08:43:49Z
last-modified:  2024-12-21T18:08:43Z
source:         RIPE
```

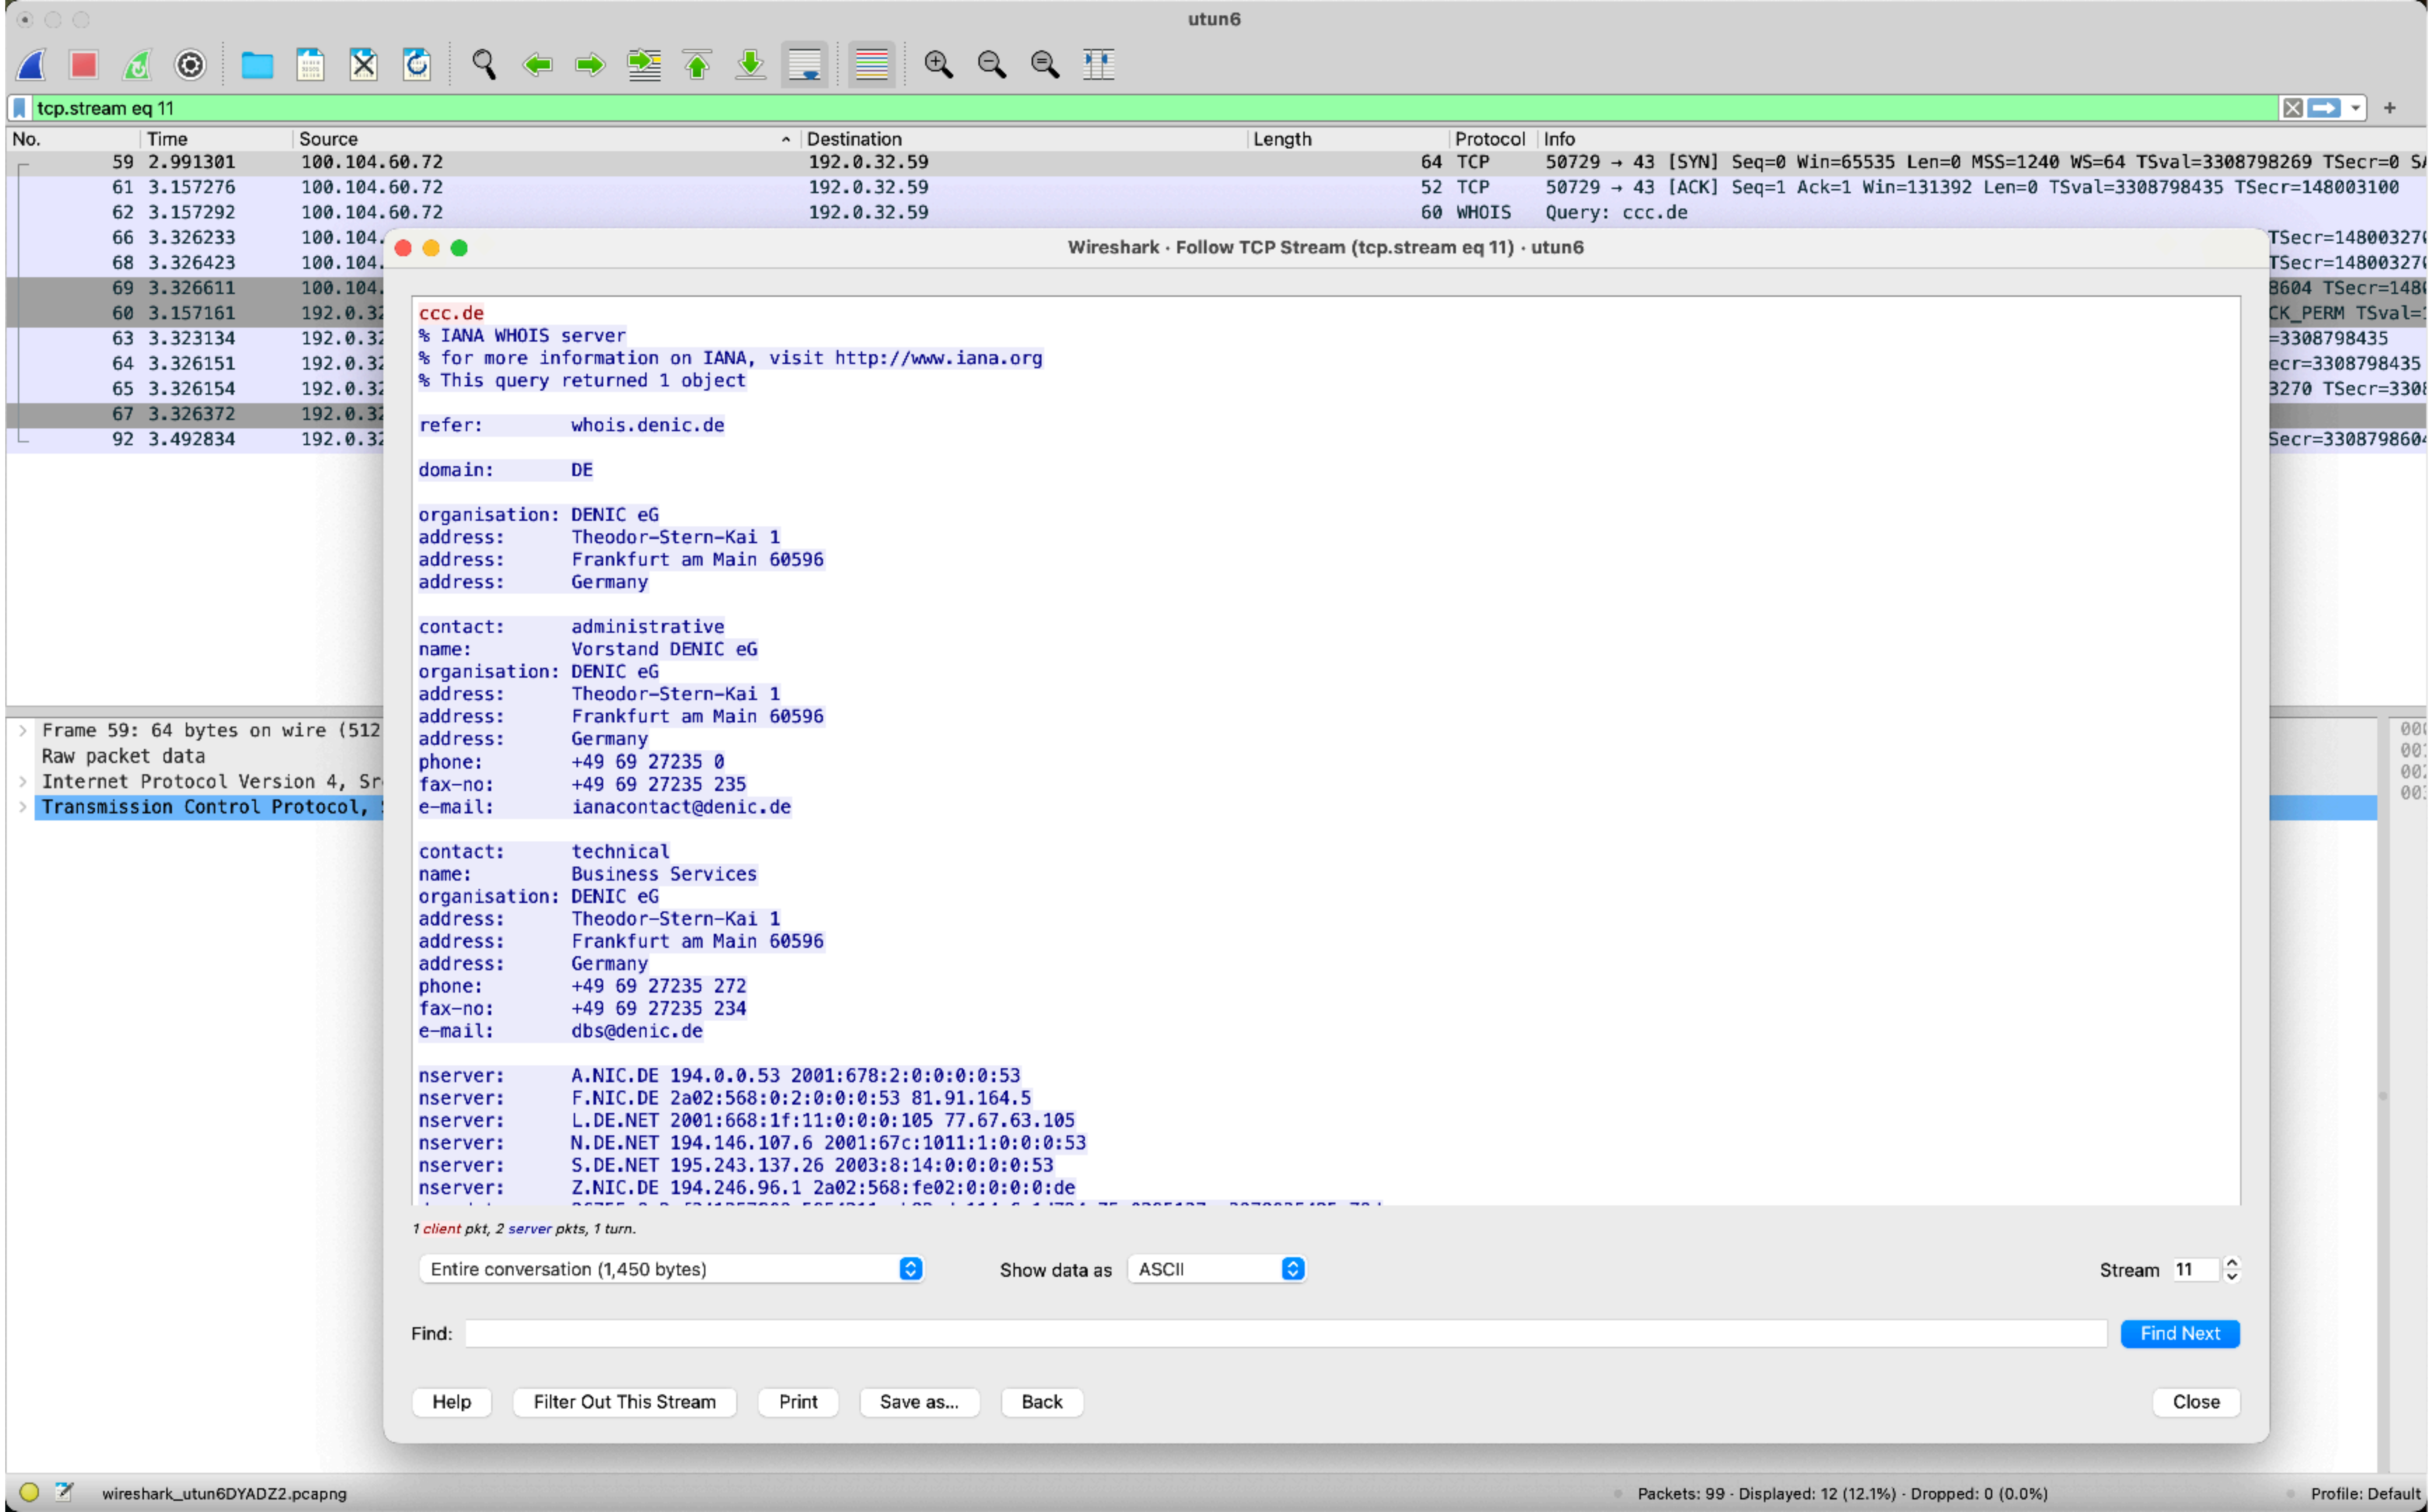
utun6
tcp.stream eq 11
No. Time Source Destination Length Protocol Info
59 2.991301 100.104.60.72 192.0.32.59 64 TCP 50729 → 43 [SYN] Seq=0 Win=65535 Len=0 MSS=1240 WS=64 TSval=3308798269 TSecr=0
61 3.157276 100.104.60.72 192.0.32.59 52 TCP 50729 → 43 [ACK] Seq=1 Ack=1 Win=131392 Len=0 TSval=3308798435 TSecr=148003100
62 3.157292 100.104.60.72 192.0.32.59 60 WHOIS Query: ccc.de
Frame 59: 64 bytes on wire
Raw packet data
Internet Protocol Version 4
Transmission Control Protocol
Wireshark · Follow TCP Stream (tcp.stream eq 11) · utun6
ccc.de
% IANA WHOIS server
% for more information on IANA, visit http://www.iana.org
% This query returned 1 object
refer: whois.denic.de
domain: DE
organisation: DENIC eG
address: Theodor-Stern-Kai 1
address: Frankfurt am Main 60596
address: Germany
contact: administrative
name: Vorstand DENIC eG
organisation: DENIC eG
address: Theodor-Stern-Kai 1
address: Frankfurt am Main 60596
address: Germany
phone: +49 69 27235 0
fax-no: +49 69 27235 235
e-mail: ianacontact@denic.de
contact: technical
name: Business Services
organisation: DENIC eG
address: Theodor-Stern-Kai 1
address: Frankfurt am Main 60596
address: Germany
phone: +49 69 27235 272
fax-no: +49 69 27235 234
e-mail: dbs@denic.de
nserver: A.NIC.DE 194.0.0.53 2001:678:2:0:0:0:0:53
nserver: F.NIC.DE 2a02:568:0:2:0:0:0:53 81.91.164.5
nserver: L.DE.NET 2001:668:1f:11:0:0:0:105 77.67.63.105
nserver: N.DE.NET 194.146.107.6 2001:67c:1011:1:0:0:0:53
nserver: S.DE.NET 195.243.137.26 2003:8:14:0:0:0:0:53
nserver: Z.NIC.DE 194.246.96.1 2a02:568:fe02:0:0:0:0:de
1 client pkt, 2 server pkts, 1 turn.
Entire conversation (1,450 bytes)
Show data as ASCII
Stream 11
Find:
Find Next
Help
Filter Out This Stream
Print
Save as...
Back
Close
wireshark_utun6DYADZ2.pcapng
Packets: 99 · Displayed: 12 (12.1%) · Dropped: 0 (0.0%)
Profile: Default

utun6

tcp.stream eq 6

No.	Time	Source	Destination	Length	Protocol	Info
40	1.381342	193.0.6.135	100.104.60.72	60	TCP	43 → 51013 [SYN, ACK] Seq=0 Ack=1 Win=28960 Len=0 MSS=1452 SACK_PERM TSval=
41	1.381568	100.104.60.72	193.0.6.135	52	TCP	51013 → 43 [ACK] Seq=1 Ack=1 Win=131392 Len=0 TSval=2822342794 TSecr=574348
42	1.381618	100.104.60.72	193.0.6.135	68	WHOIS	Query: 151.217.0.0/16
43	1.403465	193.0.6.				r=282234
44	1.403469	193.0.6.				98 TSecr=
45	1.403637	100.104.				TSecr=57
46	1.403767	193.0.6.				8599 TSe
47	1.403830	100.104.				TSecr=57
48	1.404503	193.0.6.				599 TSec
49	1.404541	100.104.				TSecr=57
50	1.405185	193.0.6.				600 TSec
51	1.405192	193.0.6.				48600 TS
52	1.405220	100.104.				TSecr=57
53	1.405233	100.104.				TSecr=5
54	1.405234	193.0.6.				0 TSecr=
55	1.405253	193.0.6.				348600 T
56	1.405258	100.104.				TSecr=5
57	1.405289	100.104.				TSecr=5
58	1.406449	193.0.6.				
59	1.406516	100.104.				TSecr=5
60	1.406650	100.104.				42819 TS
61	1.428675	193.0.6.				Secr=2822342819

Frame 58: 134 bytes on wire (10
Raw packet data
Internet Protocol Version 4, Sr
Transmission Control Protocol,
[8 Reassembled TCP Segments (40
WHOIS: Answer

Wireshark · Follow TCP Stream (tcp.stream eq 6) · utun6

```
151.217.0.0/16
% This is the RIPE Database query service.
% The objects are in RPSL format.
%
% The RIPE Database is subject to Terms and Conditions.
% See https://docs.db.ripe.net/terms-conditions.html

% Note: this output has been filtered.
%       To receive output for a database update, use the "-B" flag.

% Information related to '151.217.0.0 - 151.217.255.255'

% Abuse contact for '151.217.0.0 - 151.217.255.255' is 'abuse@ccc.de'

inetnum:        151.217.0.0 - 151.217.255.255
netname:        DE-CCC-20241127
country:        DE
geofeed:        https://geoloc.bad.network/as13020/geoloc.csv
remarks:        Geofeed https://geoloc.bad.network/as13020/geoloc.csv
remarks:        ==================================================
remarks:        === ===
remarks:        === If you have trouble with users from ===
remarks:        === this network, please contact ===
remarks:        === ===
remarks:        === ABUSE MAIL: abuse@ccc.de ===
remarks:        === ===
remarks:        === In case of urgency you can also ===
remarks:        === contact our abuse hotline: ===
remarks:        === ===
remarks:        === ABUSE HOTLINE: +49 40 401801-666 ===
remarks:        === ===
remarks:        ==================================================
abuse-c:        CCC-RIPE
org:            ORG-CCCE3-RIPE
admin-c:        CCC-RIPE
tech-c:         CCC-RIPE
status:         ASSIGNED PI
remarks:        Temporary assignment (start date: 2024/11/27, end date: 2024/12/31 and duration 34 days)
mnt-by:         CHAOS-MNT
mnt-by:         RIPE-NCC-END-MNT
```

Packet 42. 1 client pkt, 8 server pkts, 1 turn. Click to select.

Entire conversation (4,027 bytes) | Show data as ASCII | Stream 6

Find: | Find Next

Help | Filter Out This Stream | Print | Save as... | Back | Close

wireshark_utun6LVYCZ2.pcapng | Packets: 62 · Displayed: 23 (37.1%) · Dropped: 0 (0.0%) | Profile: Default

Mandatory Introduction to RIPE/NCC, RIRs & IRRs

Regional Internet Registries

Réseaux IP Européens (RIPE) & NCC

RIPE

Discussion forum open to everybody interested

The RIPE community

The RIPE Network Coordination Centre

RIPE NCC
RIPE NETWORK COORDINATION CENTRE

- ~200 employees
- Offices in Amsterdam and Dubai

Hierarchical Distribution of IP Numbers

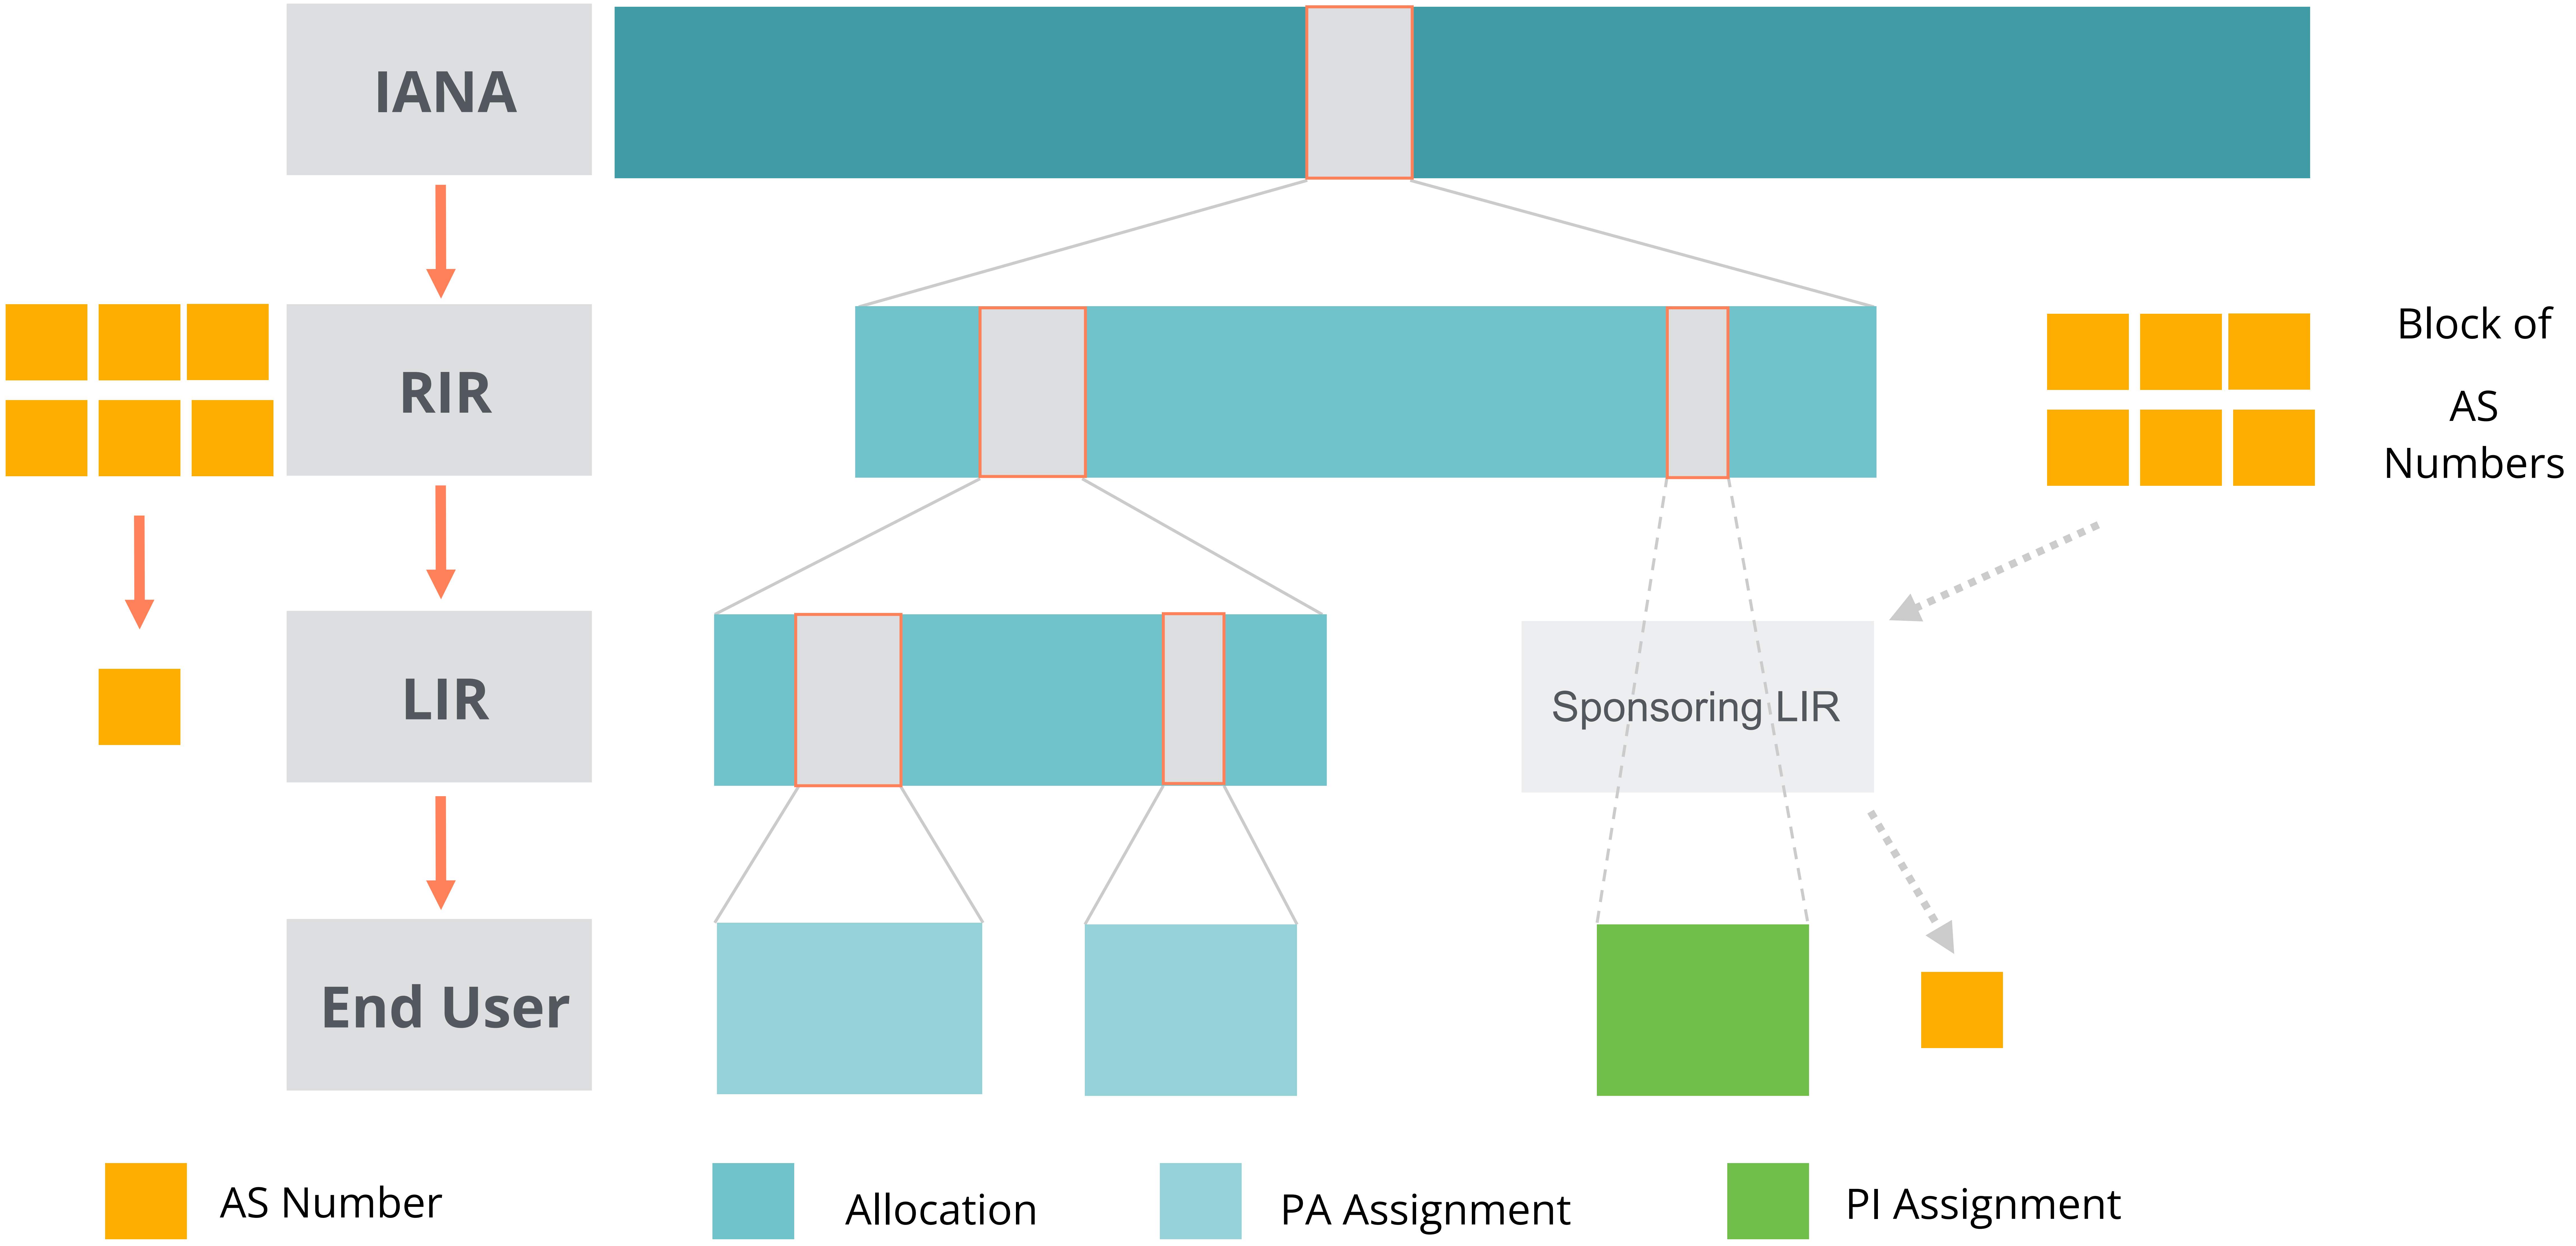

Purpose of the Internet Routing Registry (IRR)

- Registry of who holds IPs and ASNs
  - Part of whois database (“RIPE Database”)
- Keep **contact** information
  - For troubleshooting, notifying of outages, etc.
- Publishing **routing** policies

- Operated by:
  - IANA
  - RIRs

The RIPE (whois) Database

Public Internet resource and **routing registry** database

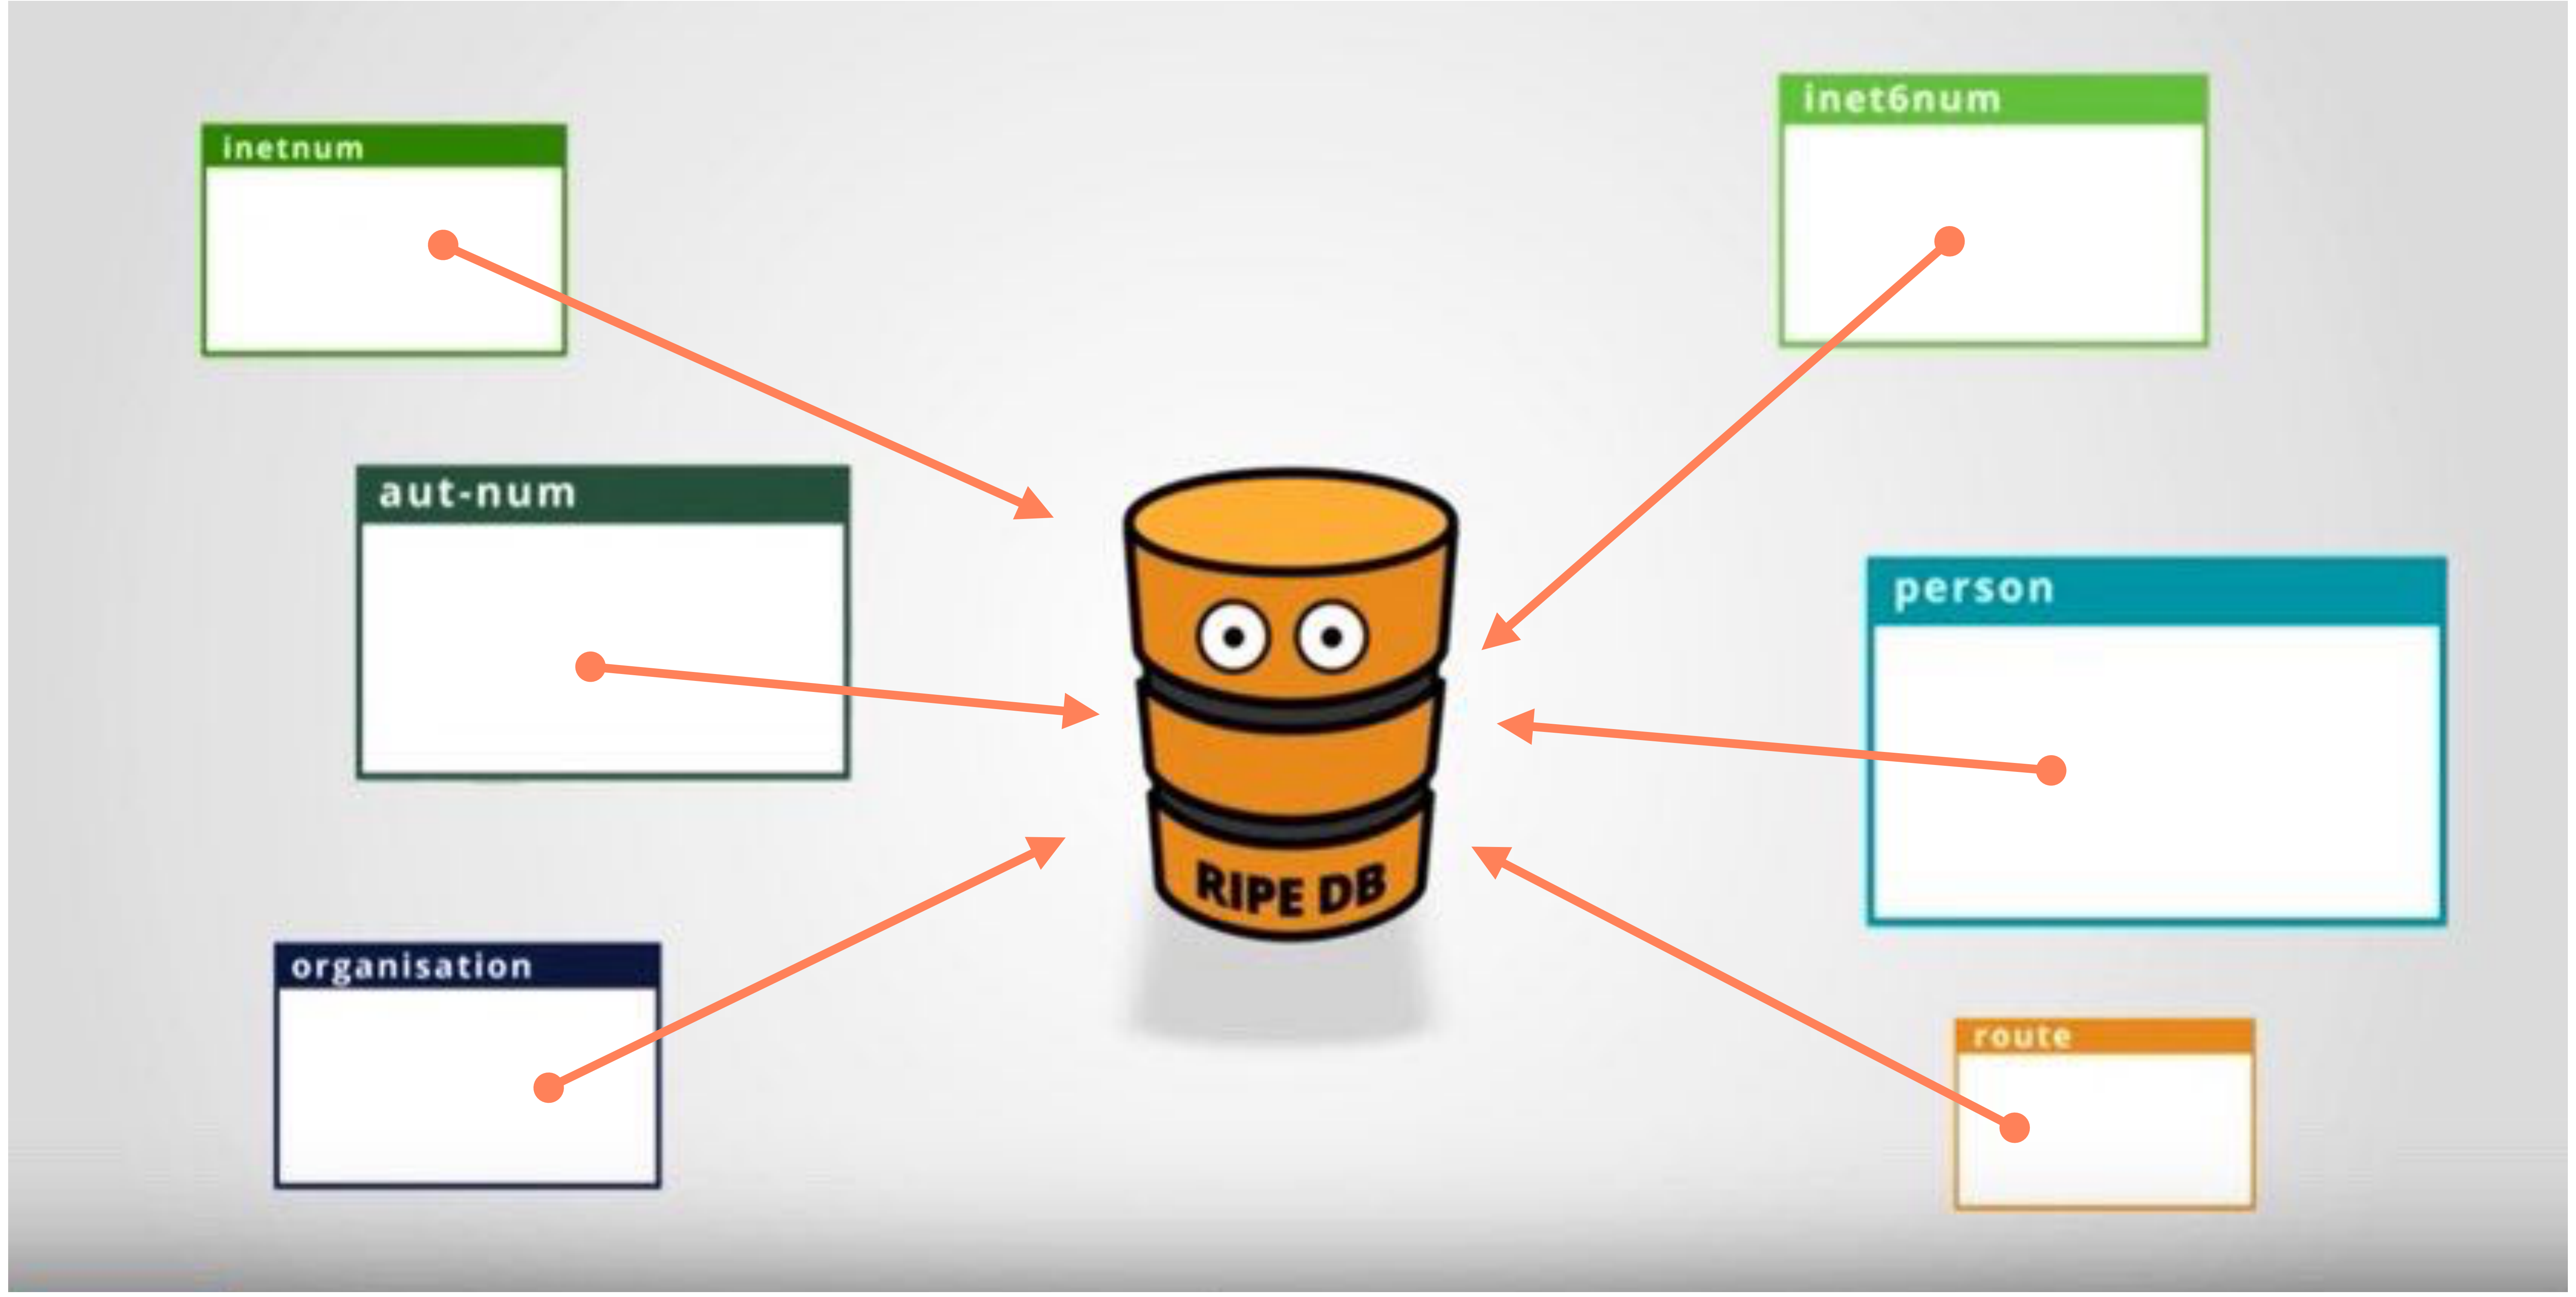

RIPE Database Objects

IPs and ASNs

- inetnum
- inet6num
- aut-num

Contact Information

- organisation
- person
- role

Routing

- route
- route6
- as-set

Reverse DNS

- domain

Object Protection

- mntner

Authentication in RIPE (whois) Database

Maintainers: Protecting DB Objects

person: Jean Blue
address: My Street 9876
address: Office 123
phone: +31 20 876 5432
e-mail: jean@example.net
nic-hdl: JB123-RIPE
mnt-by: LIR-MNT

mntner: LIR-MNT
admin-c: JB123-RIPE
notify: noc@example.org
upd-to: noc@example.org
auth: MD5-PW 1crypto-stuff
auth: SSO email@domain.com
auth: PGP-KEY-<key ID>
mnt-by: LIR-MNT

* MD5-PW will be deprecated in 2025

Maintainers: Authentication

- **SSO [RIPE]**
  - uses RIPE NCC Access account
  - for editing via a web interface (LIR Portal)

- **PGP / x509**
  - uses PGP key pair or x509 certificates
  - to authenticate: sign updates with private key

- **MD5-PW (will be deprecated in 2025)**
  - uses a MD5 hashed password
  - to authenticate: provide clear text password 🥹

https://docs.db.ripe.net/Authorisation/Using-the-Authorisation-Methods/

Authentication: History

- **Authentication differs per database**
- This is a historic design with many historic design issues, e.g.

- **MD5 hashes publicly available until 2011**
  - Passwords for leaked hashes reset in 2016

- **MAIL-FROM “authentication” (RADB)**
  - Yes: Authenticate by sending email from a specific address (including wildcards)
  - Deprecated in 2015

Where to Learn More?

- Webinars (for everyone)
  - ripe.net/training
  - http://youtube.com/ripencc
- Training Courses (only for LIRs)
  - ripe.net/training
- E-learning
  - academy.ripe.net
- Workshops & presentations: at your school?!
  - cd@ripe.net

Border Gateway Protocol (BGP)

BGP Routing Illustrated

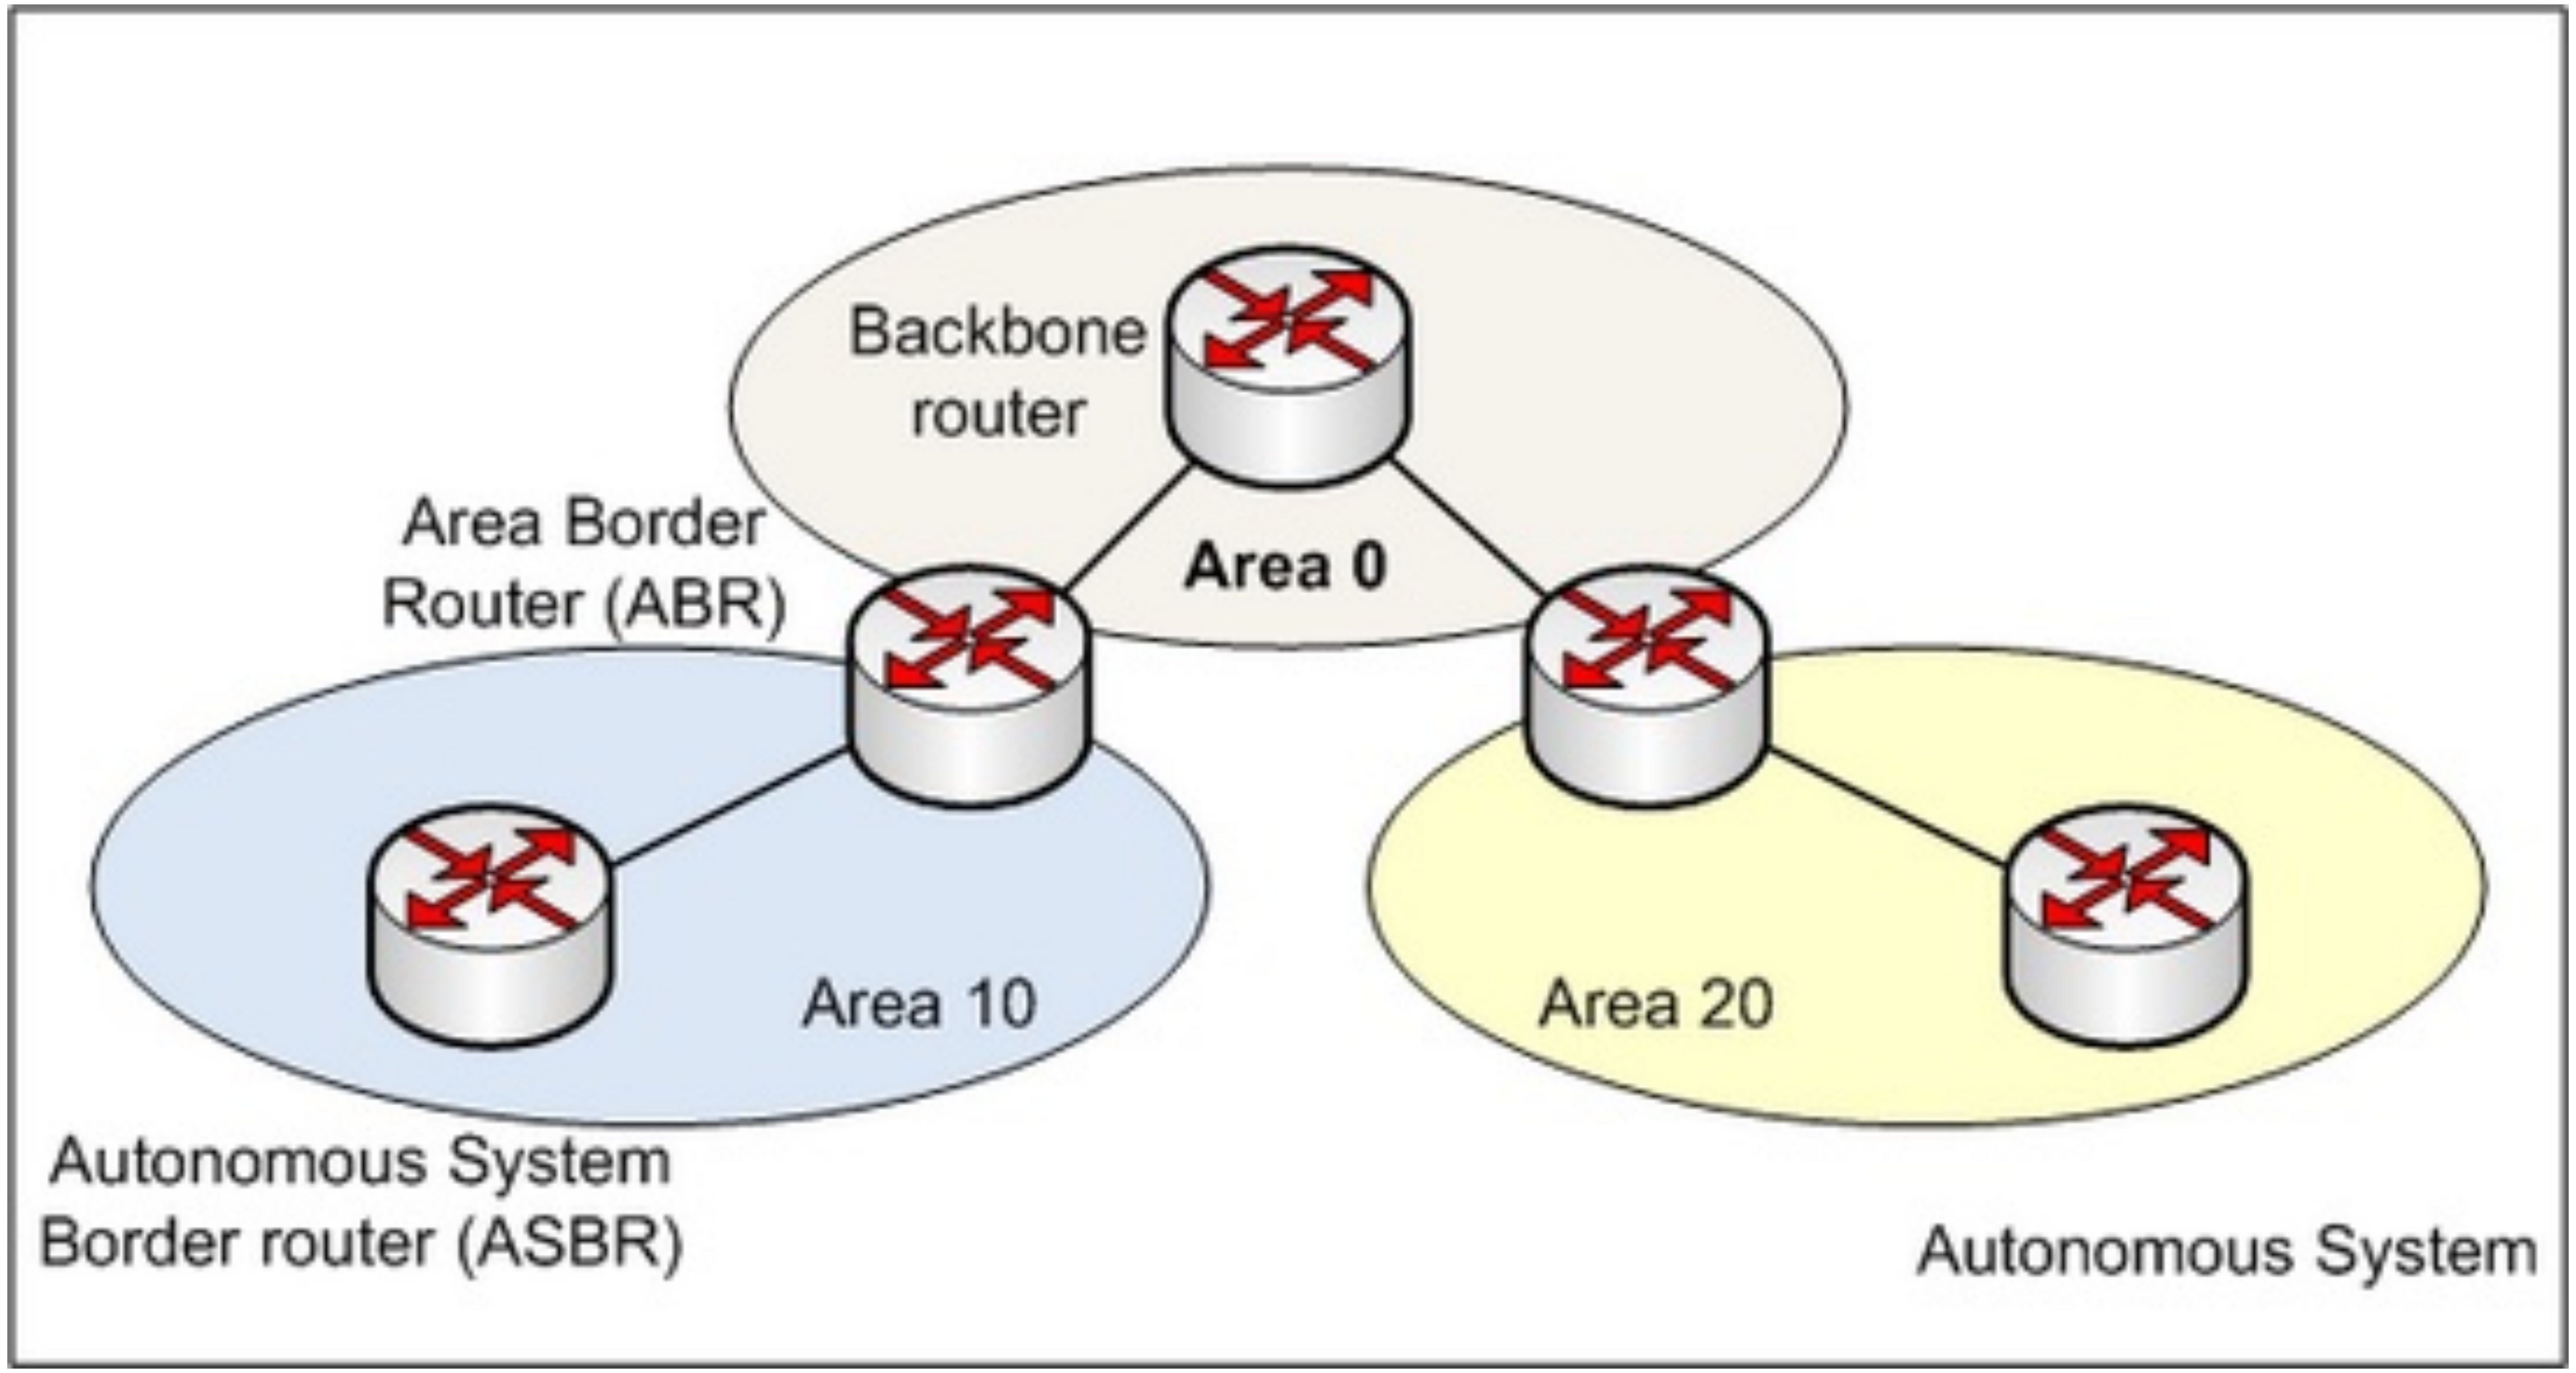

How Does it Work?

How Does it Work?

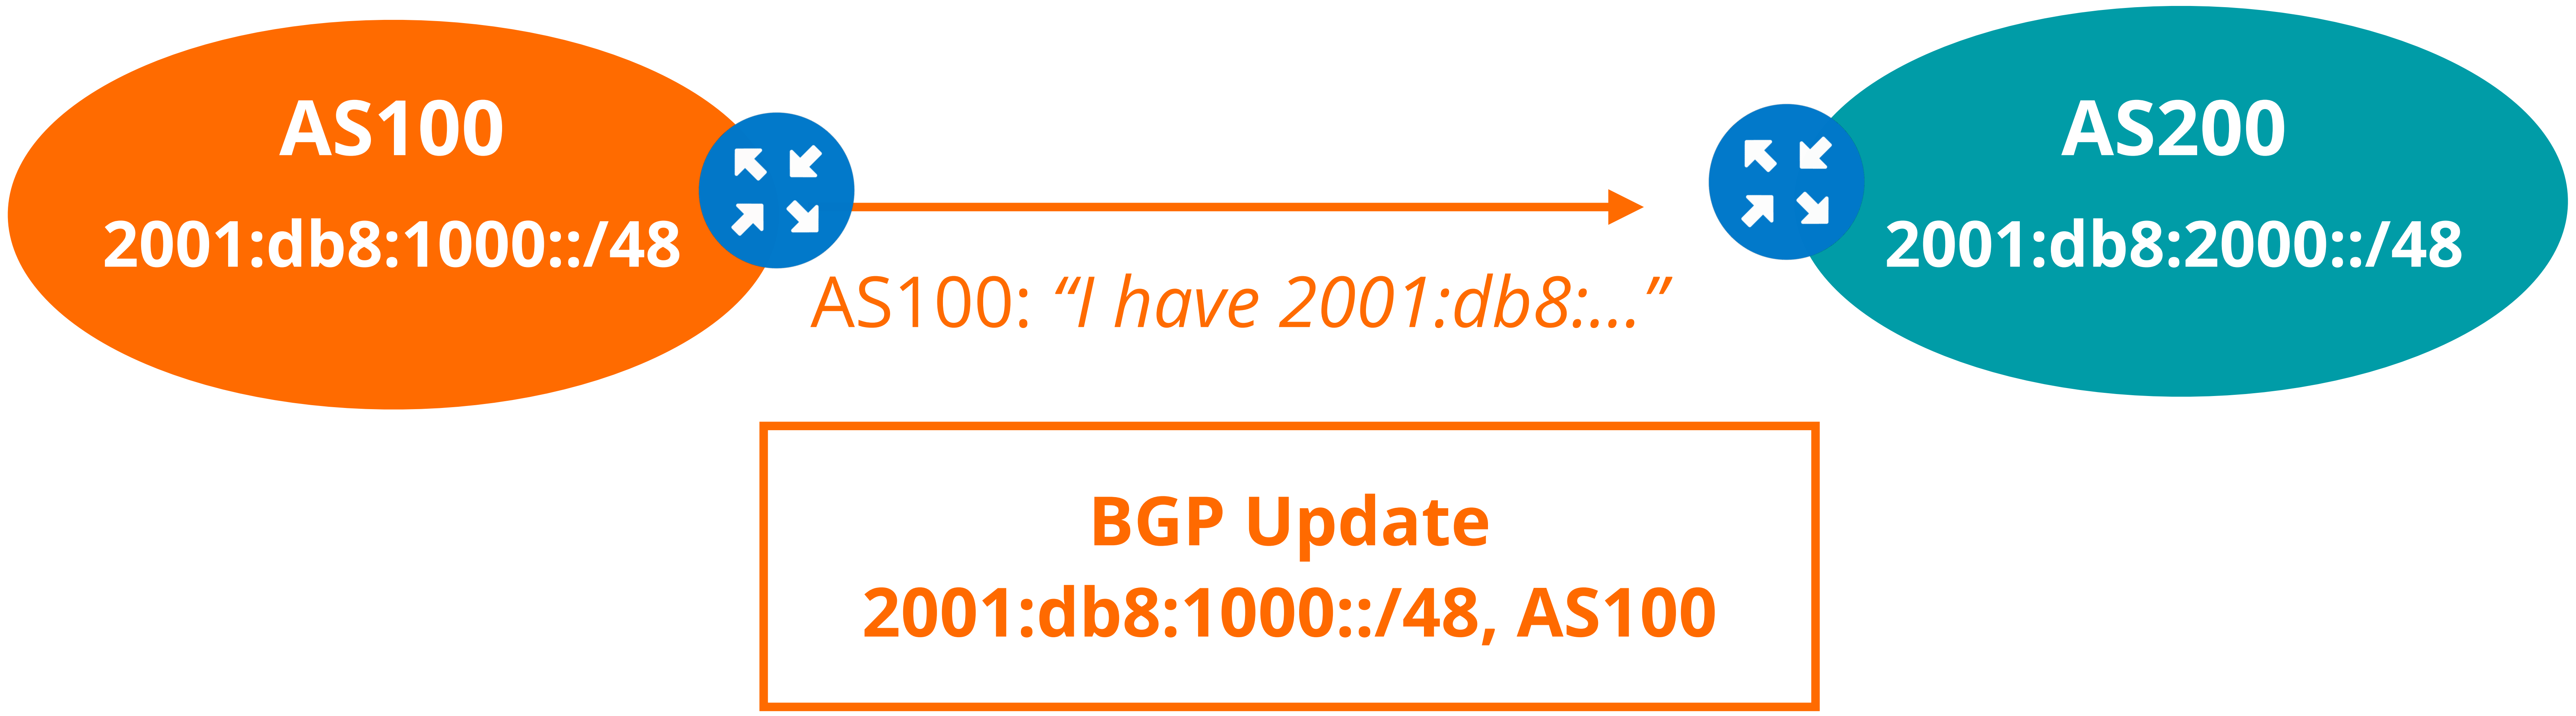

How Does it Work?

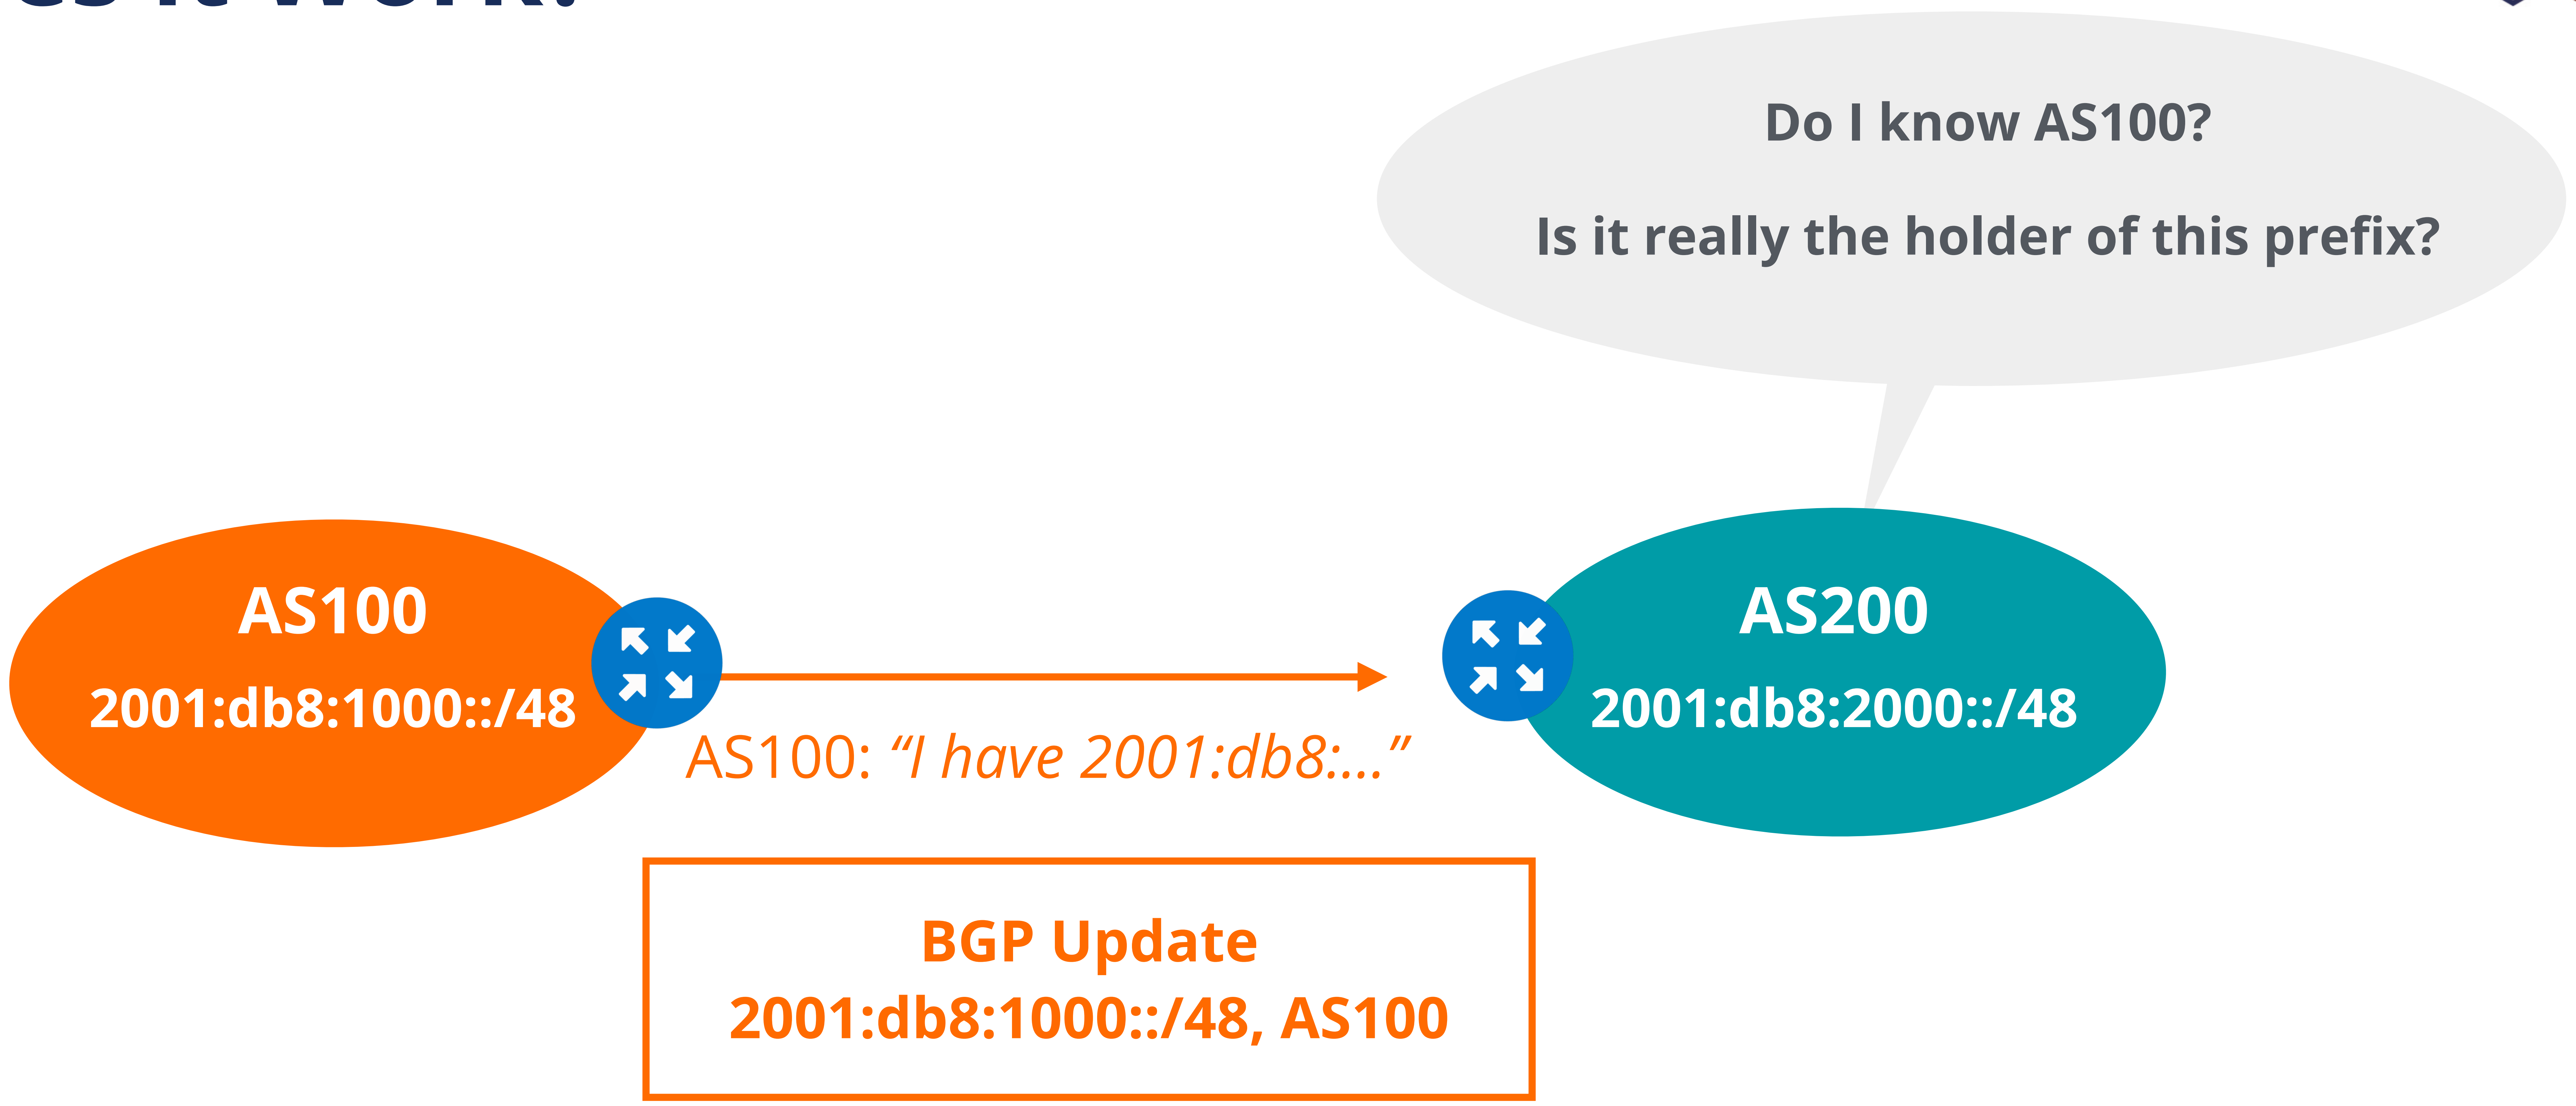

How Does it Work?

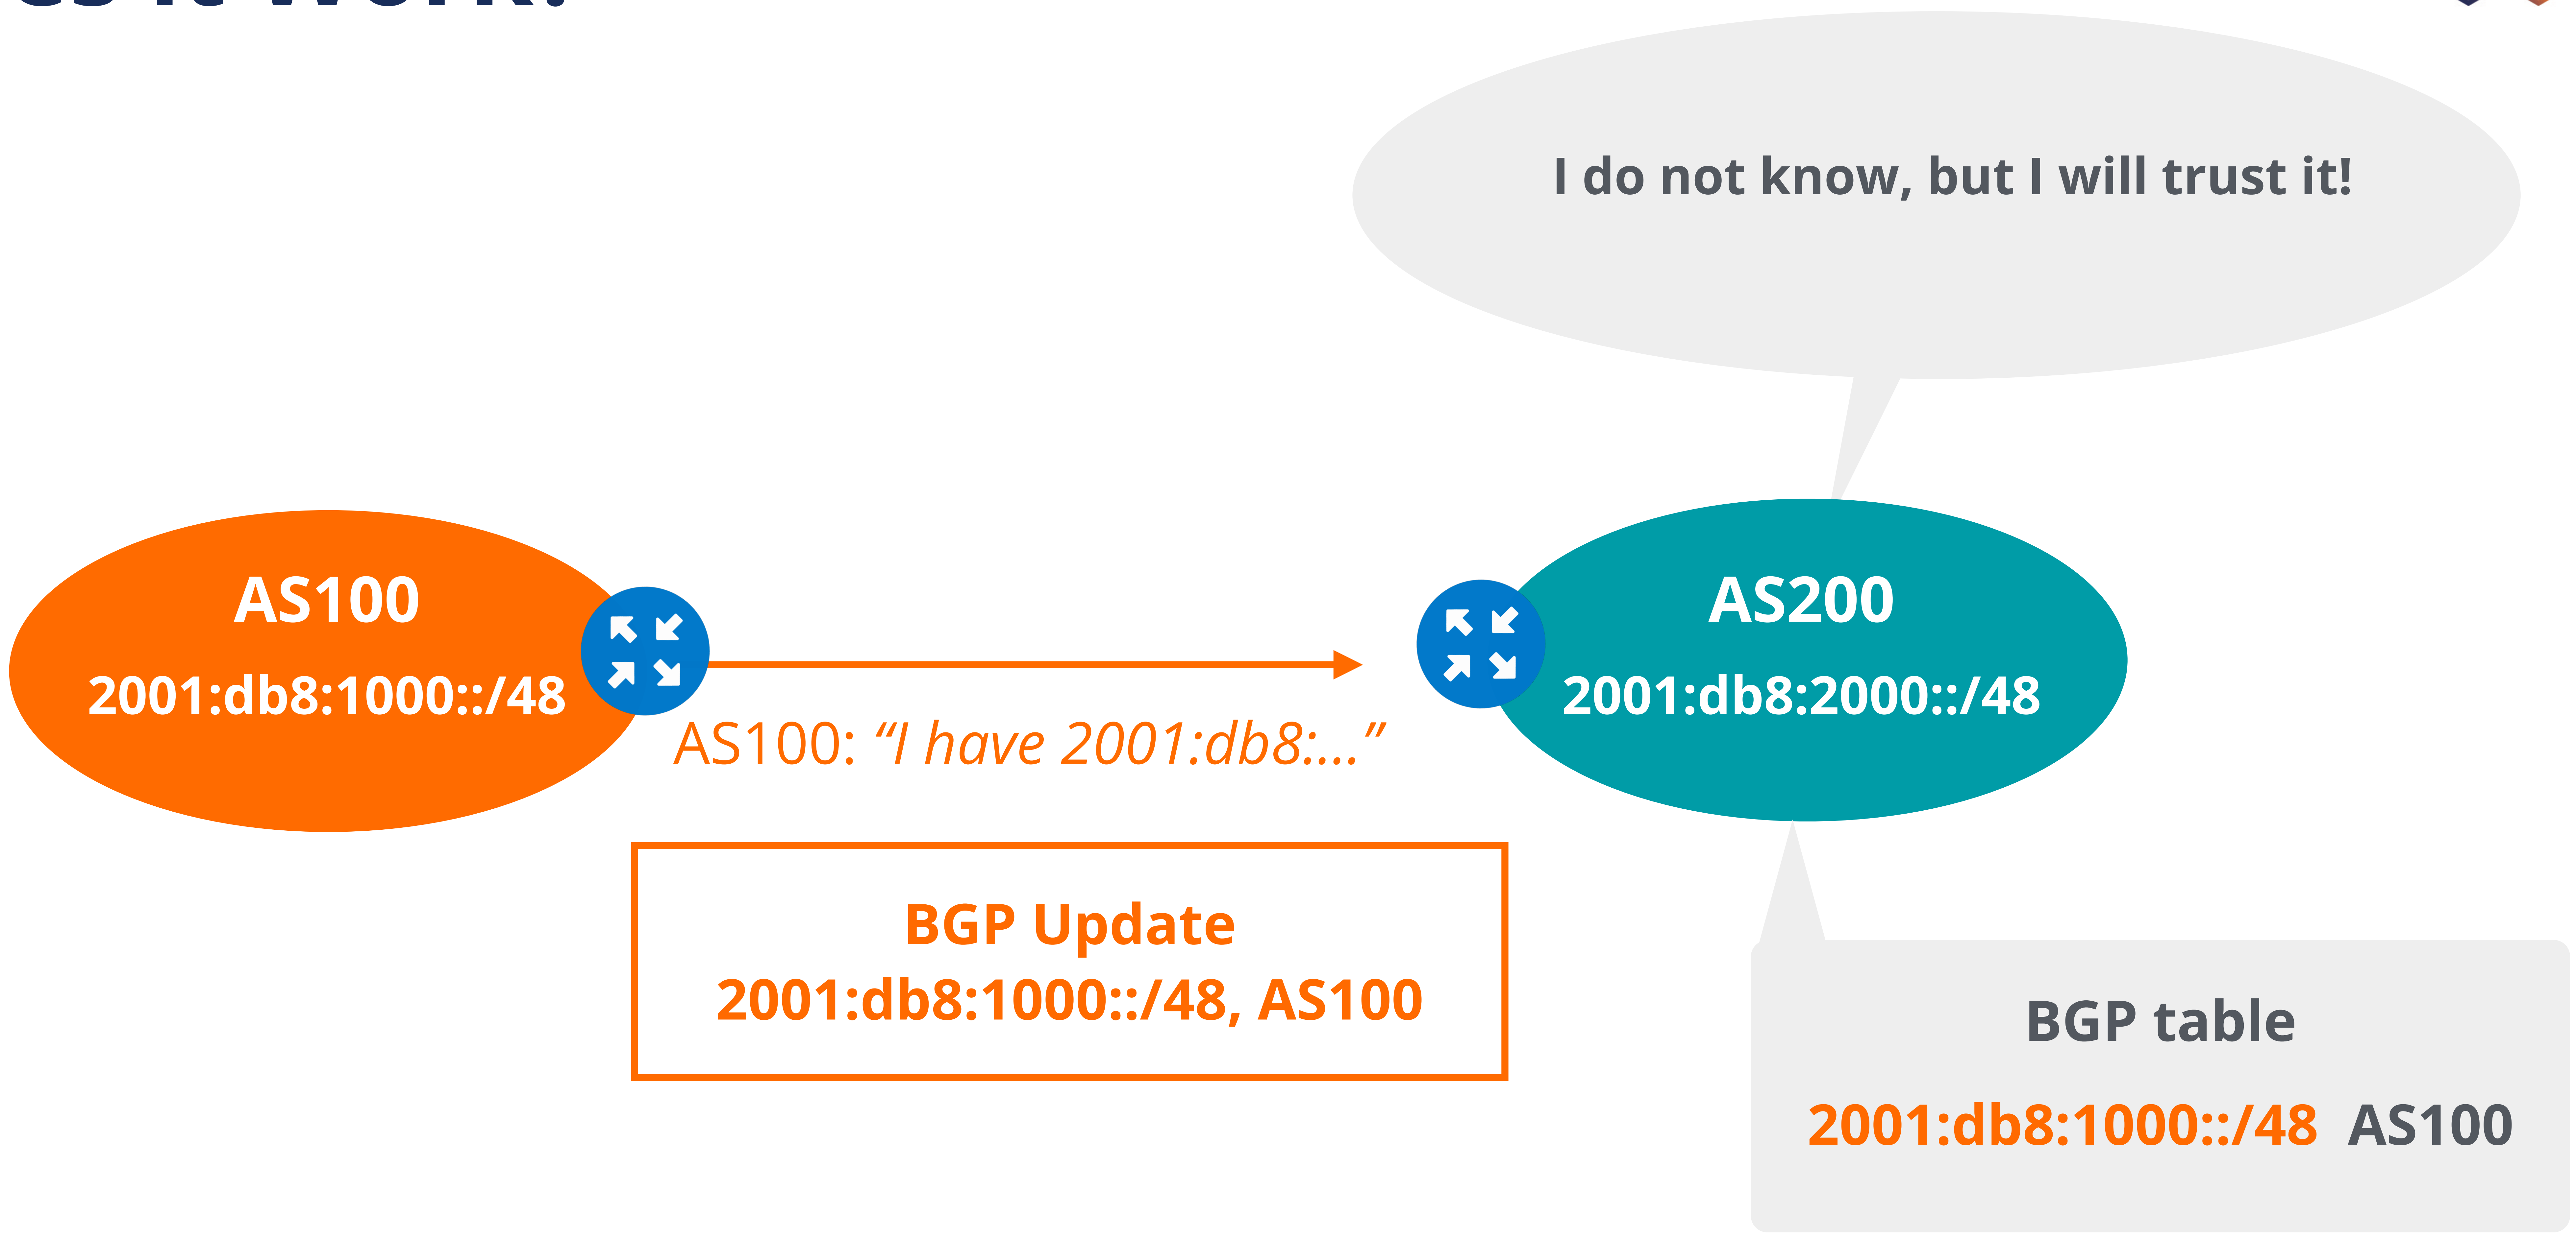

How Does it Work?

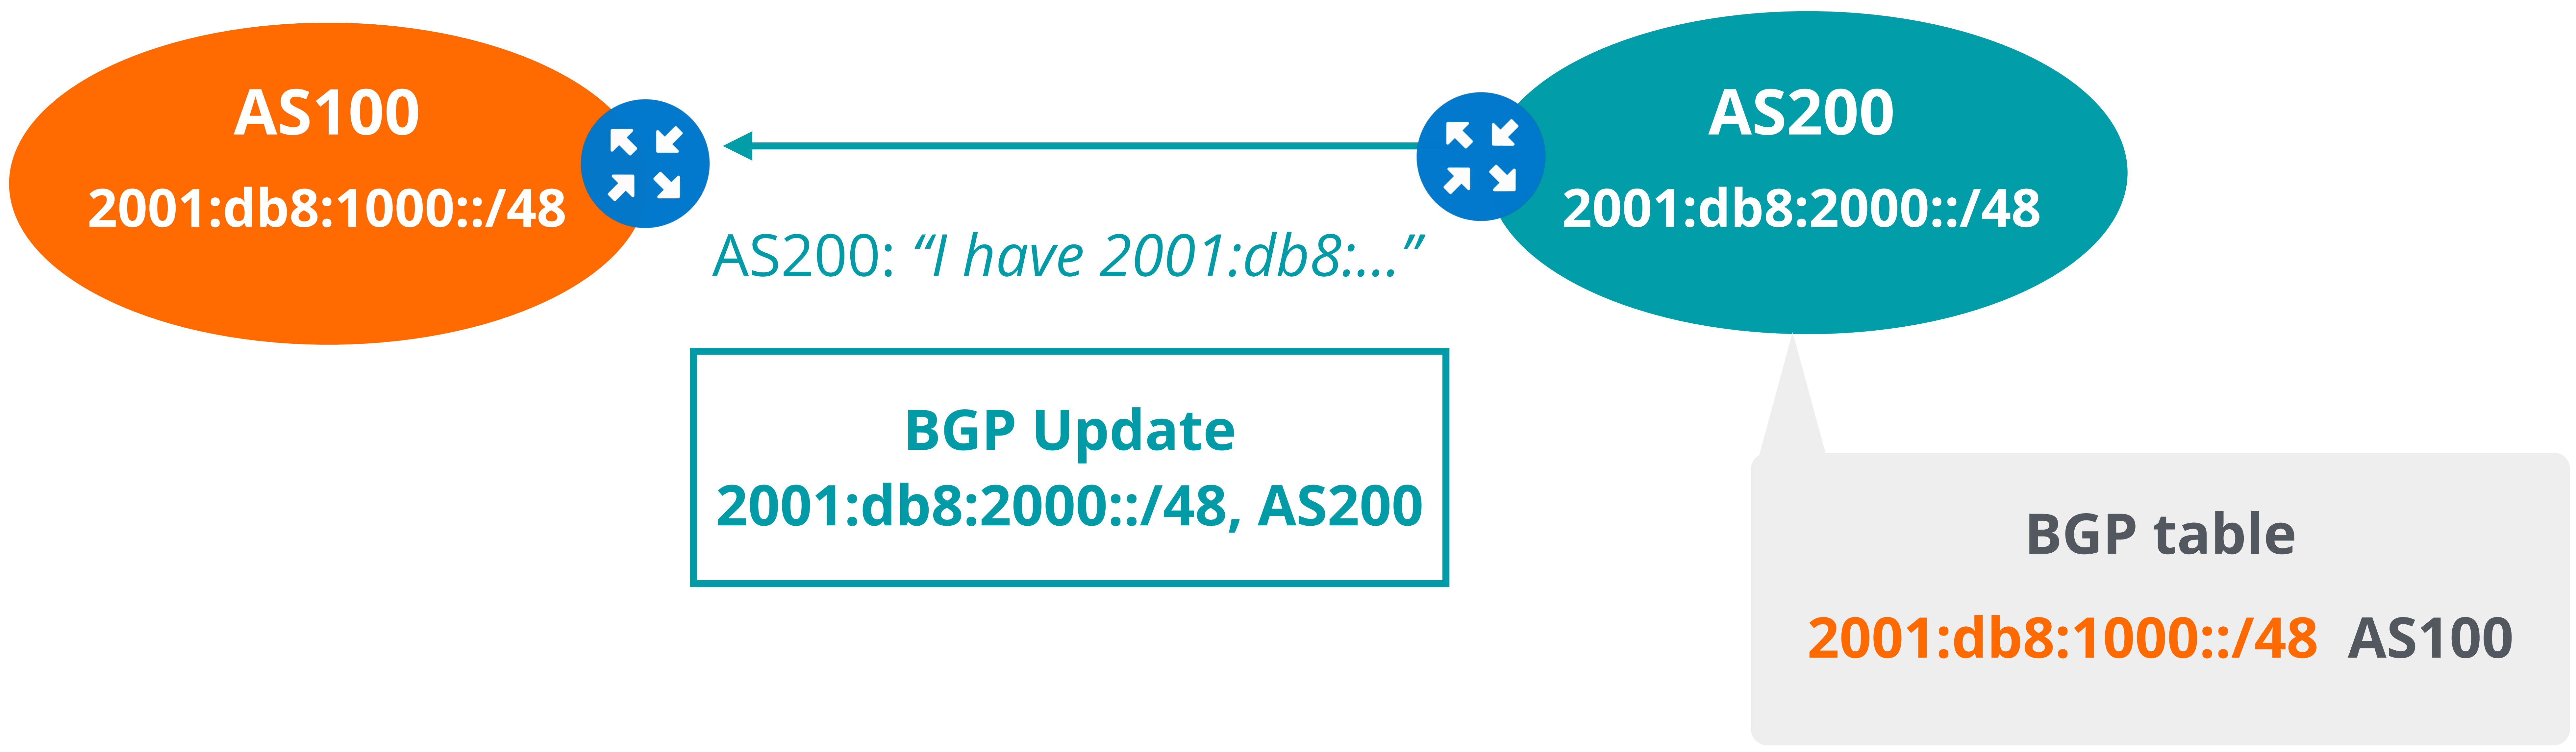

How Does it Work?

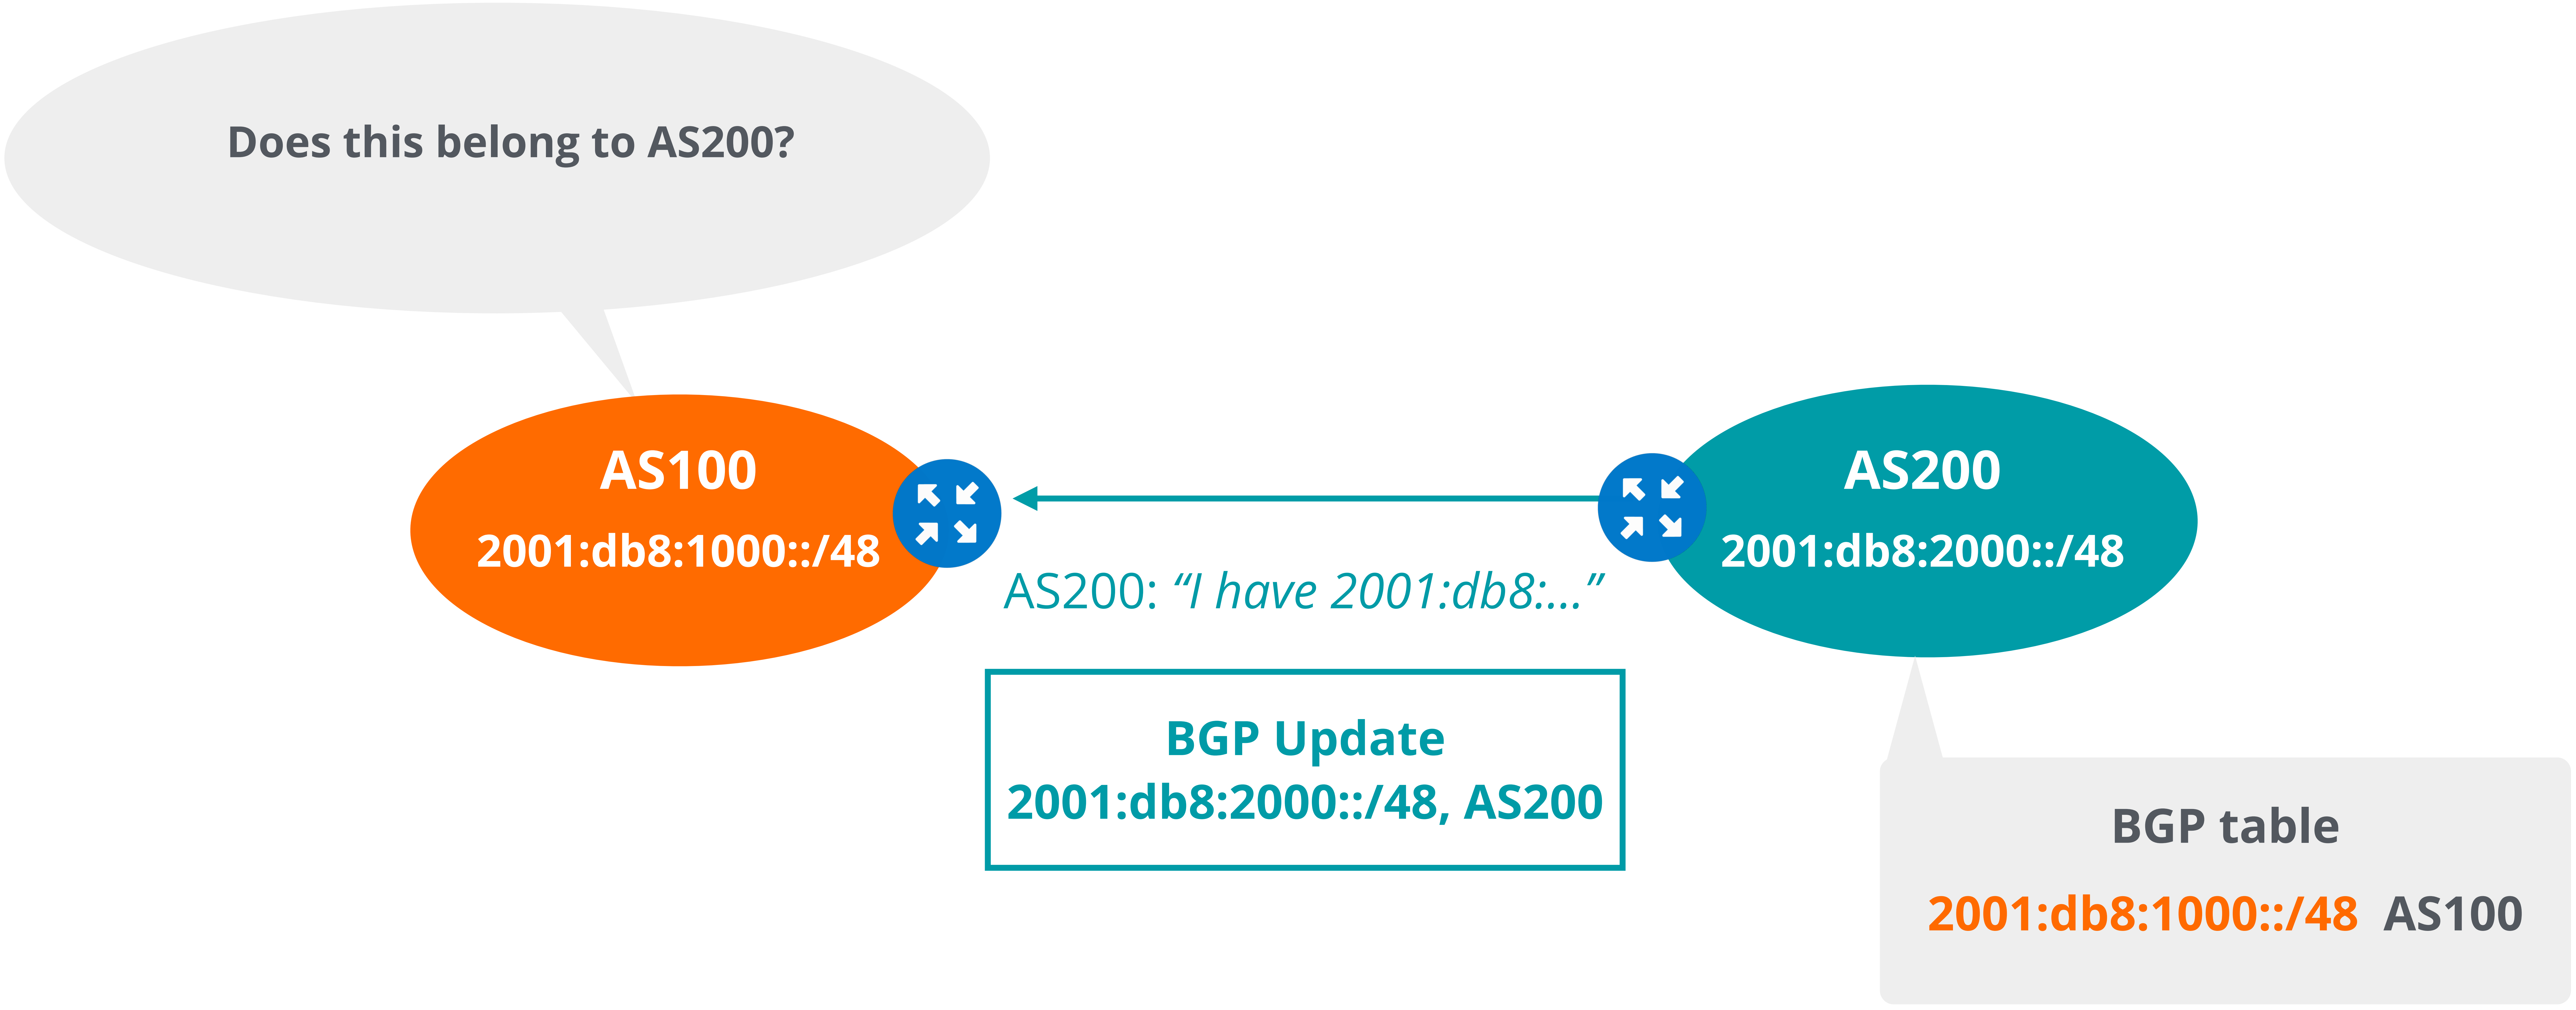

How Does it Work?

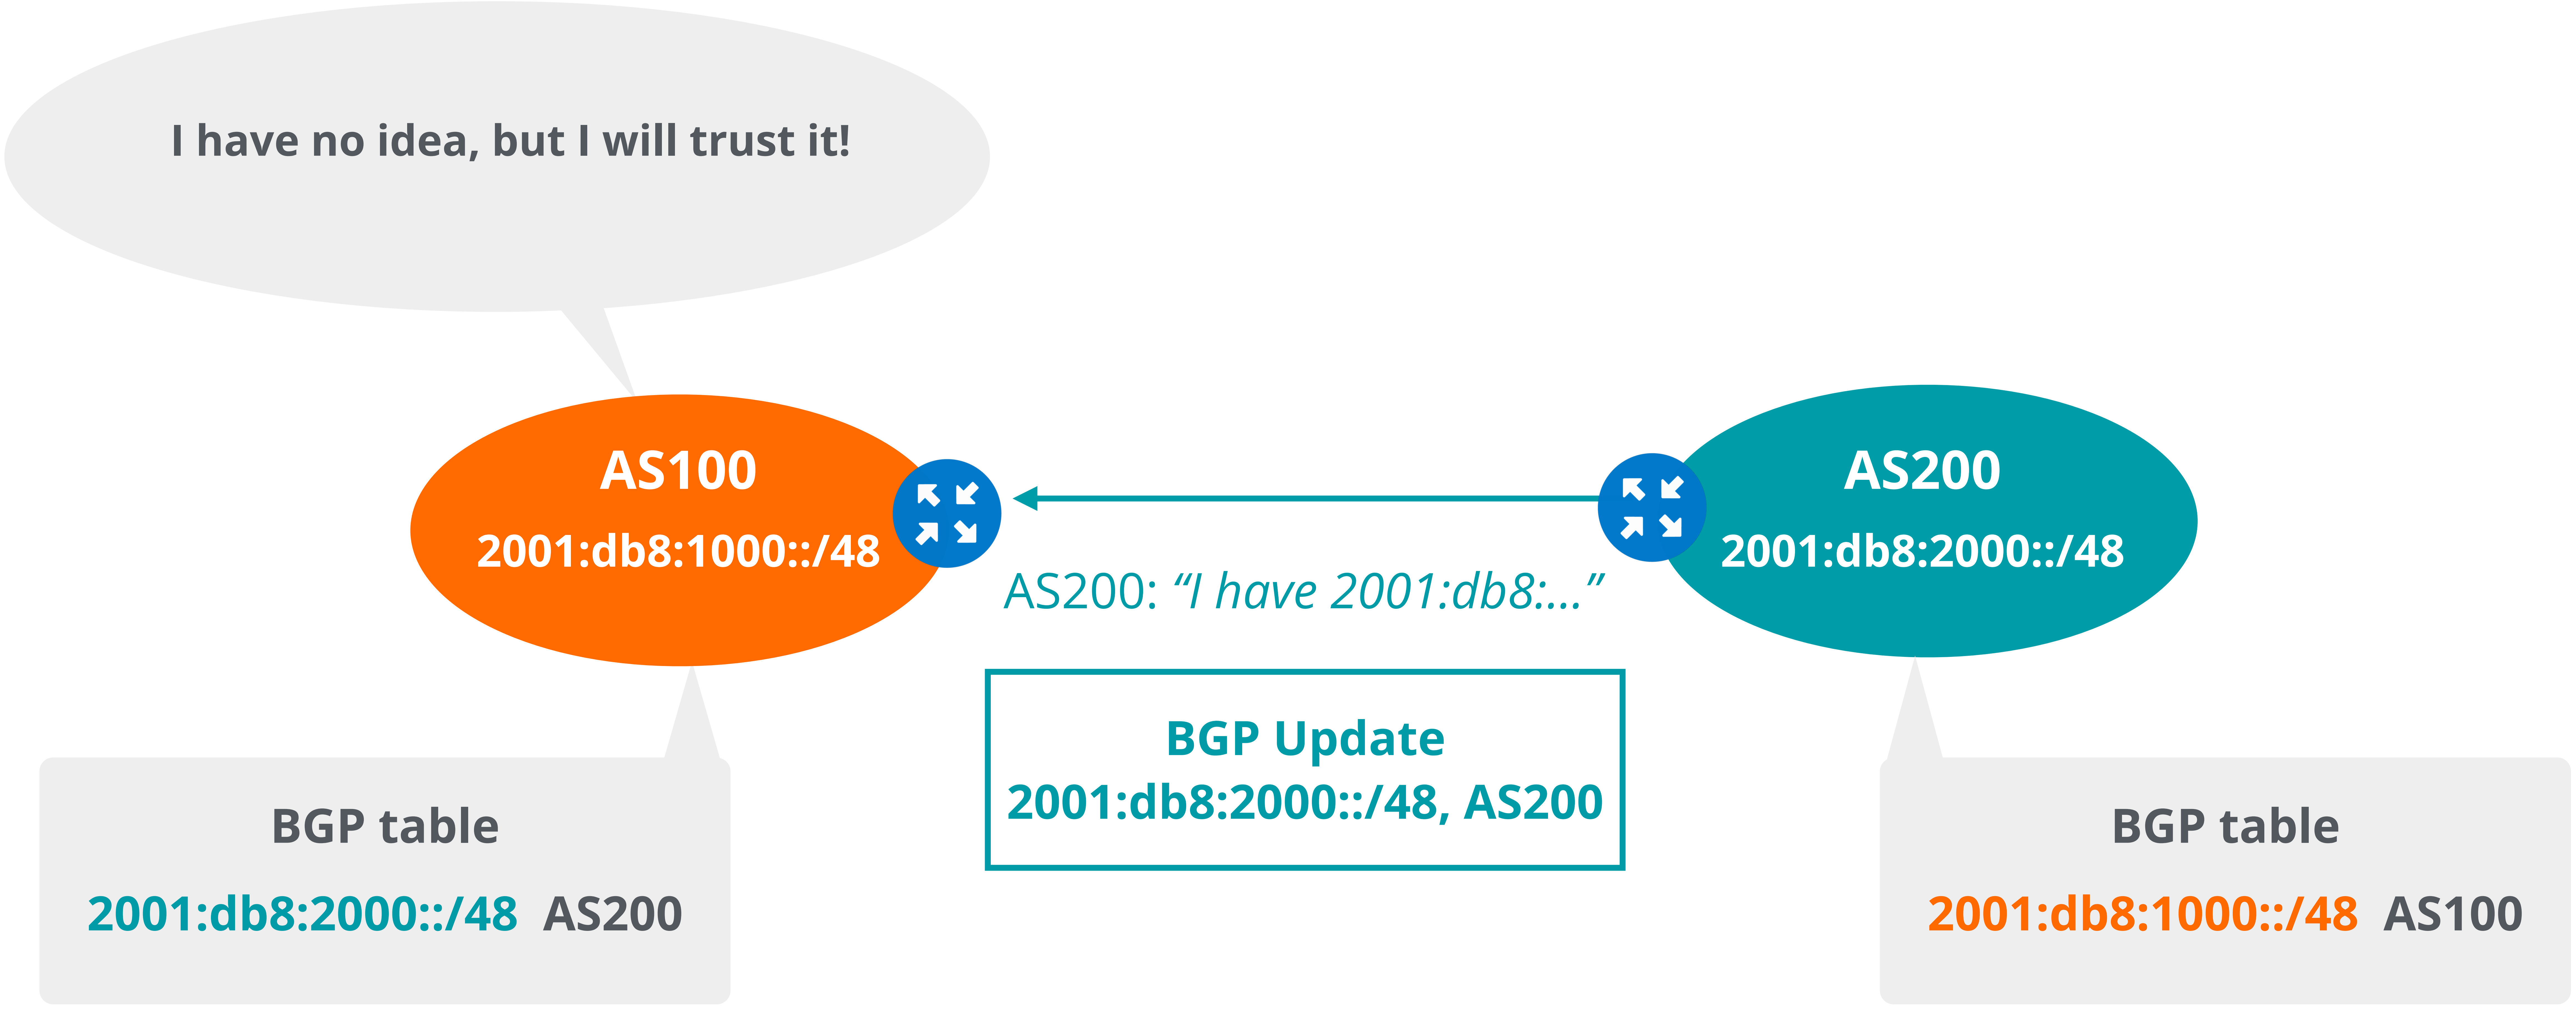

BGP assumes that everybody is telling the truth!

But what if someone lies?

Route Leaks

“The propagation of BGP announcements beyond their intended scope” [RFC7908]

- **Illegitimate propagation of legitimate prefixes (not bogus routes)**
- **Result from human errors or misconfigurations**
  - And/or improper or missing BGP route filters between BGP peers
- Leads to incorrect or suboptimal routing

Google Prefix leak - November 2018

- **What happened?**
- **MainOne leaked Google routes to CT and CT leaked them to other transits**
  - Google services (G Suite and Google Search) affected by the leak
- **Why?**
  - Due to misconfigured filters

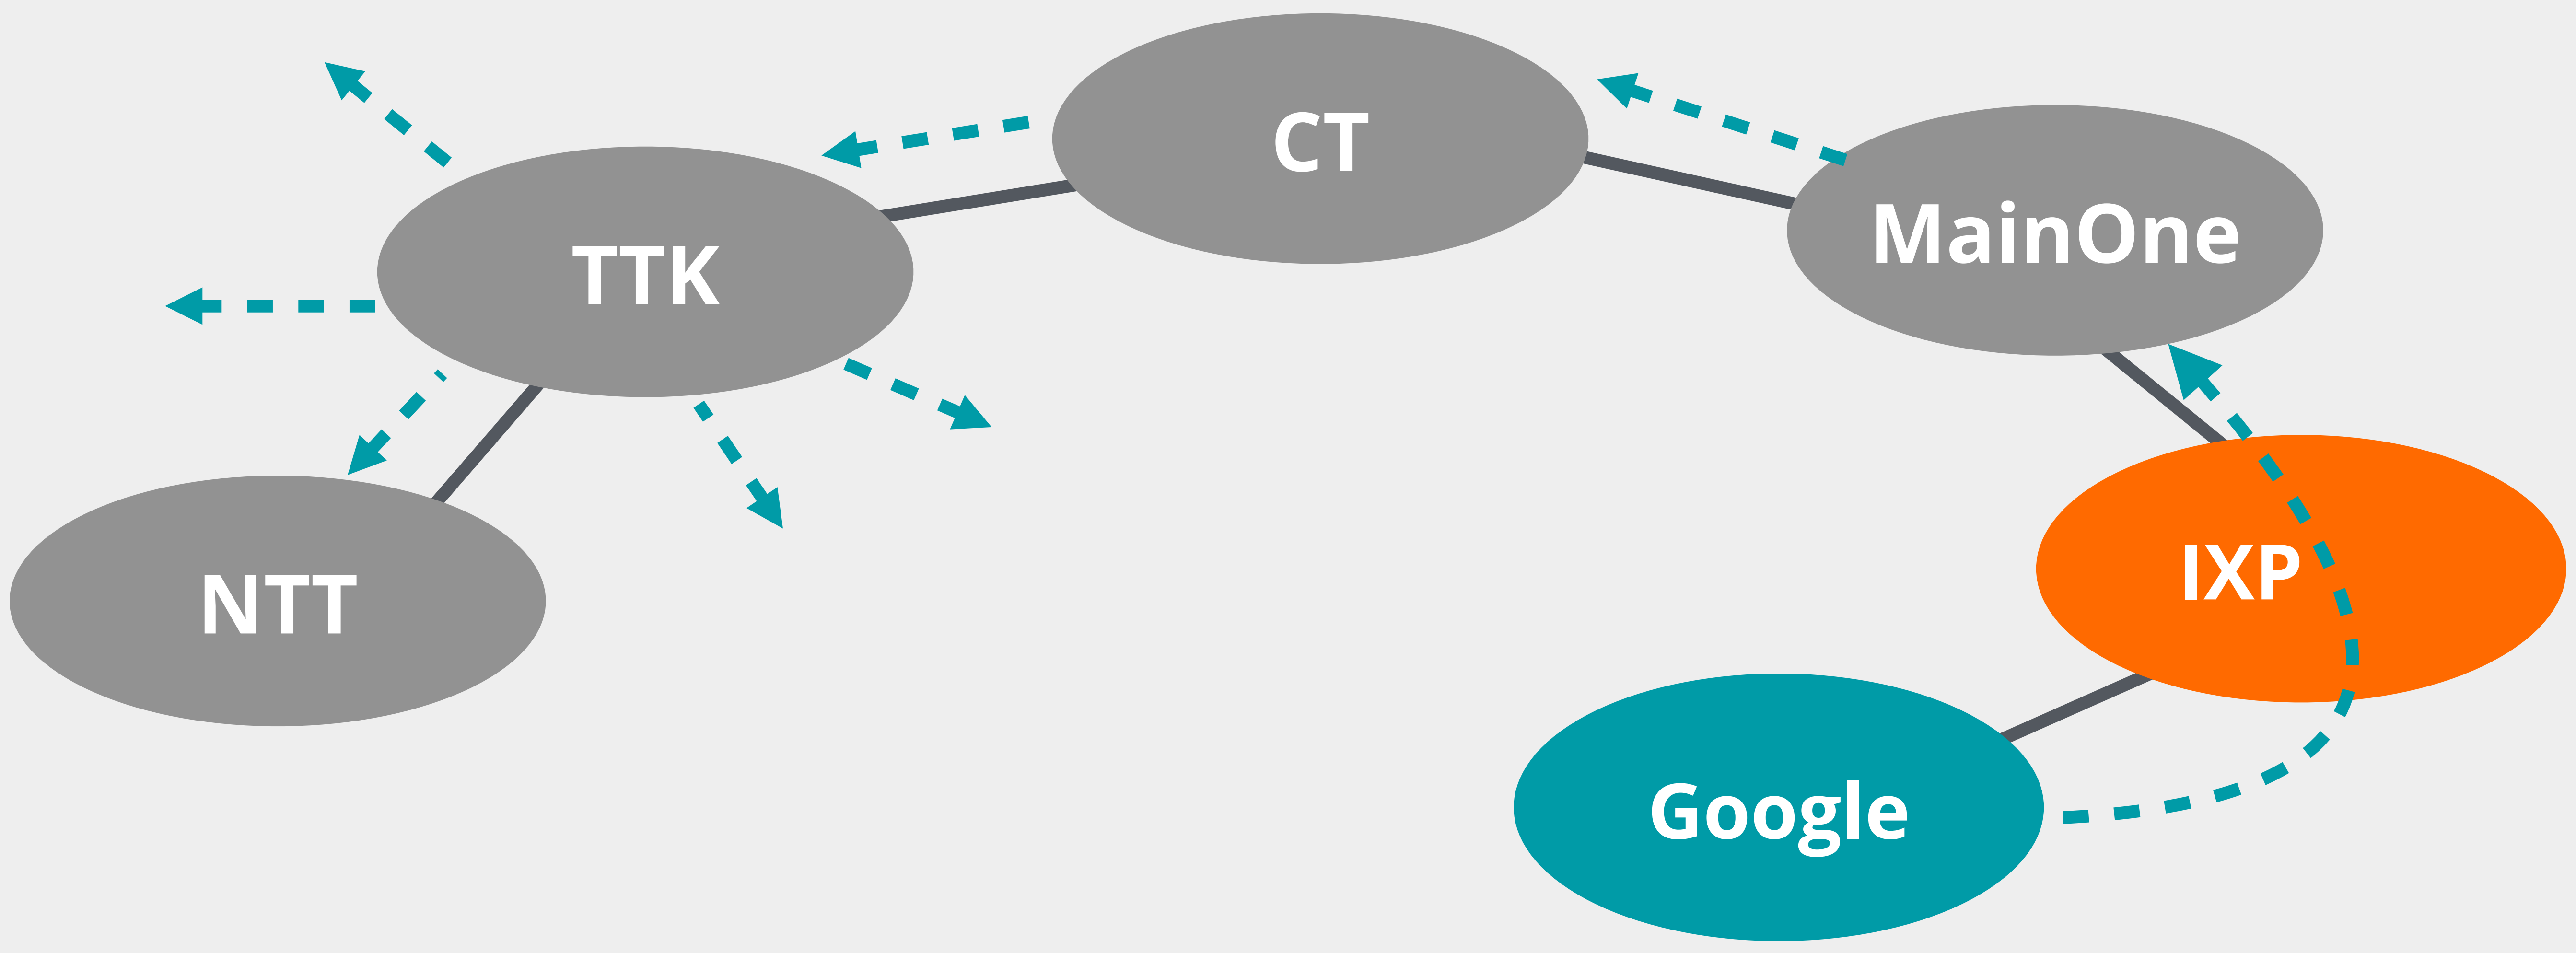

Google Prefix leak - November 2018

- **What's different with proper filters?**
  - Google's prefix wouldn't reach China Telecom
  - Proper outbound filters in MainOne, and/or
  - Proper inbound filters in CT

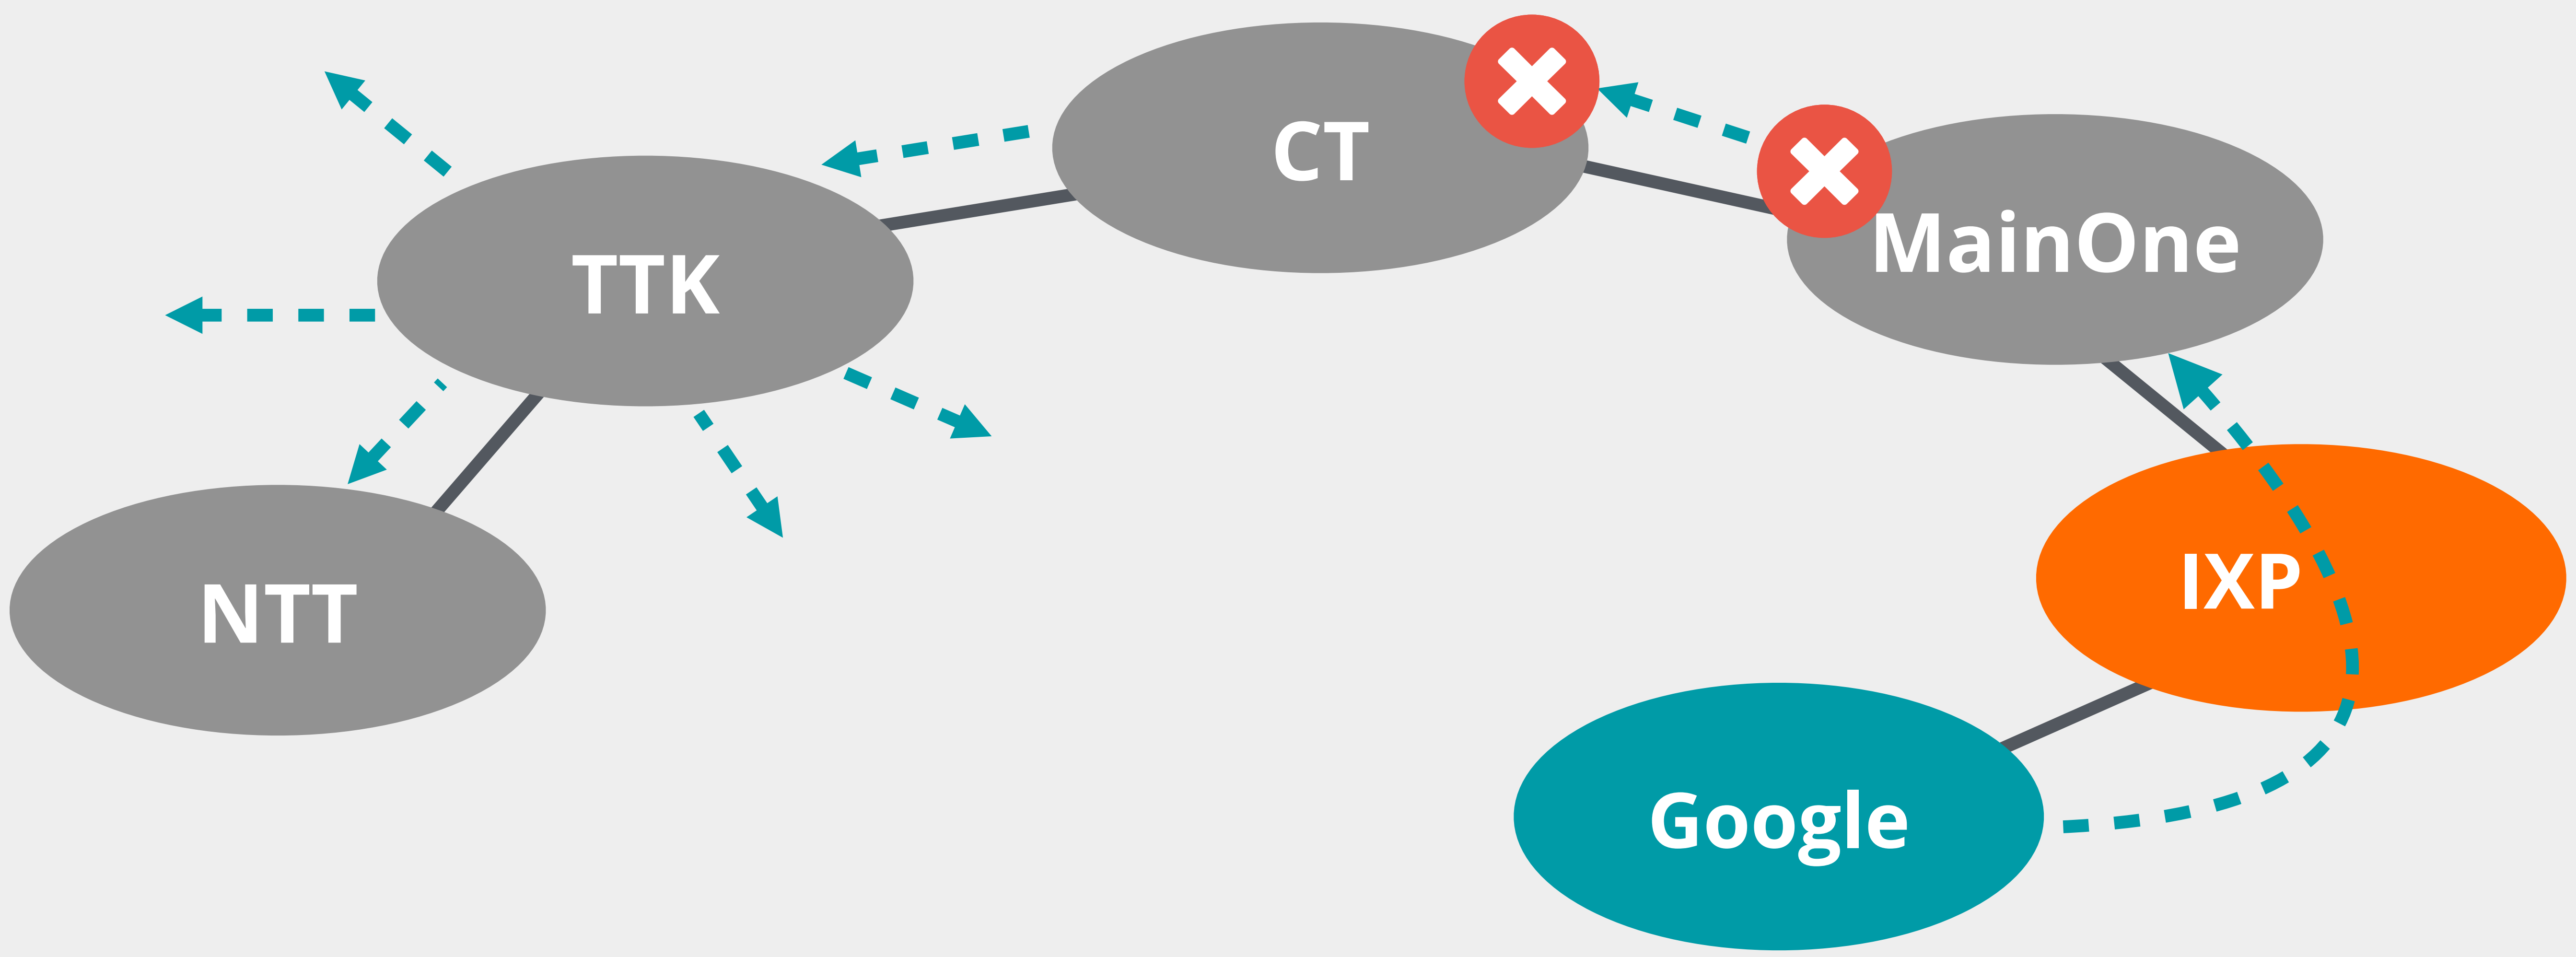

Implementing Route Filtering

How to Prevent Route Leaks?

Route filtering is the most powerful mechanism!

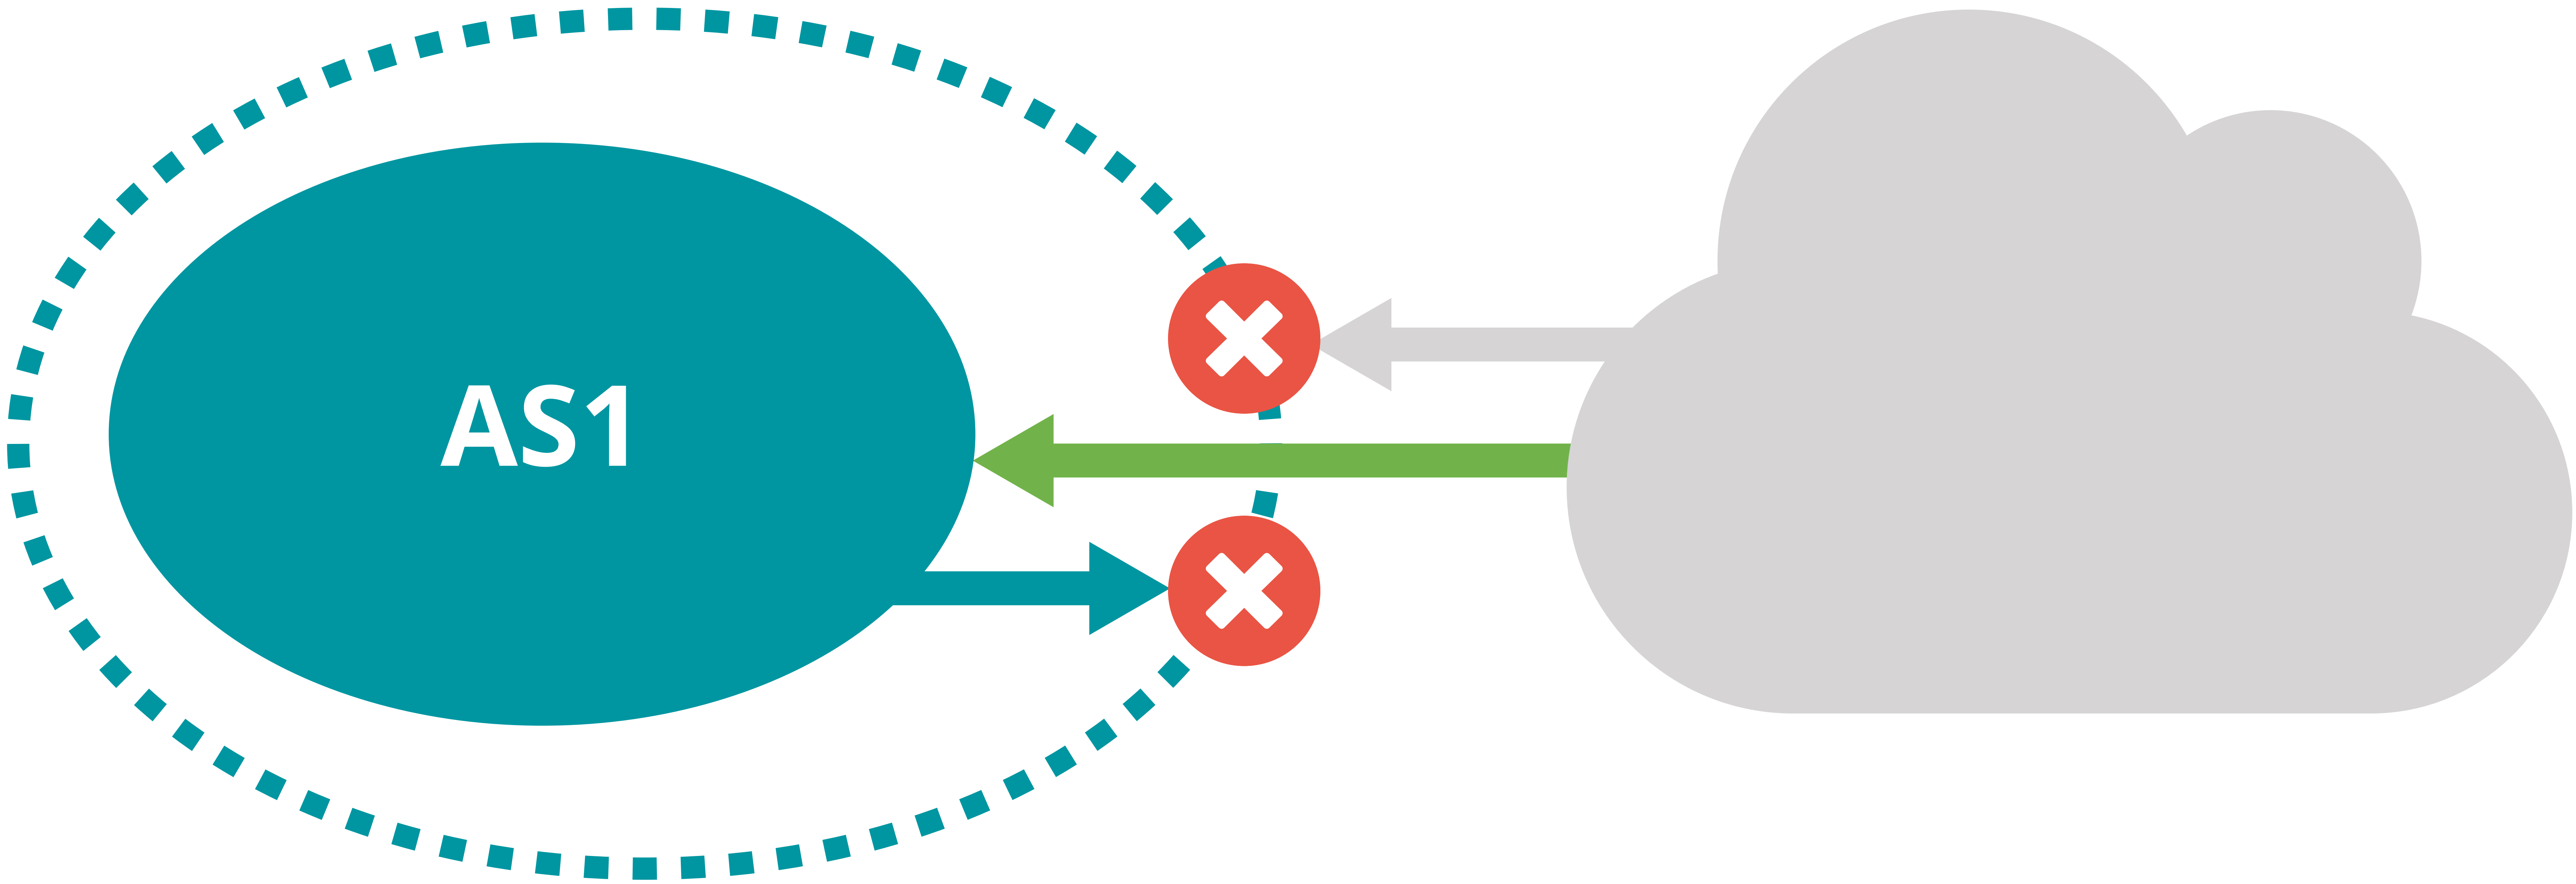

What is BGP route Filtering?

- **The most basic protection mechanism against malicious or accidental BGP incidents:**
  - Prevents **route leaks**
  - Mitigates the impact of **BGP hijacks**
- **Technique used to control prefixes on the BGP peering**
  - Which prefixes will you **advertise** to your peers?
  - Which prefixes will you **accept** into your network?

Essential for routing security!

Other Reasons for Filtering

- **Business relationships**
  - Customer-provider, peer-peer
- **Technical reasons**
  - Reduce memory utilisation, scalability
- **Traffic engineering**
  - Manipulate traffic flows and influence best path selection

BGP Filters (BGP Policies)

- **Used to filter prefixes exchanged between BGP peers**
- **Describe BGP peers and routing relationships with them**
- **Filters can match on**
  - IP prefixes
  - AS paths
  - Or any other BGP attributes (e.g. MED, BGP communities, etc)

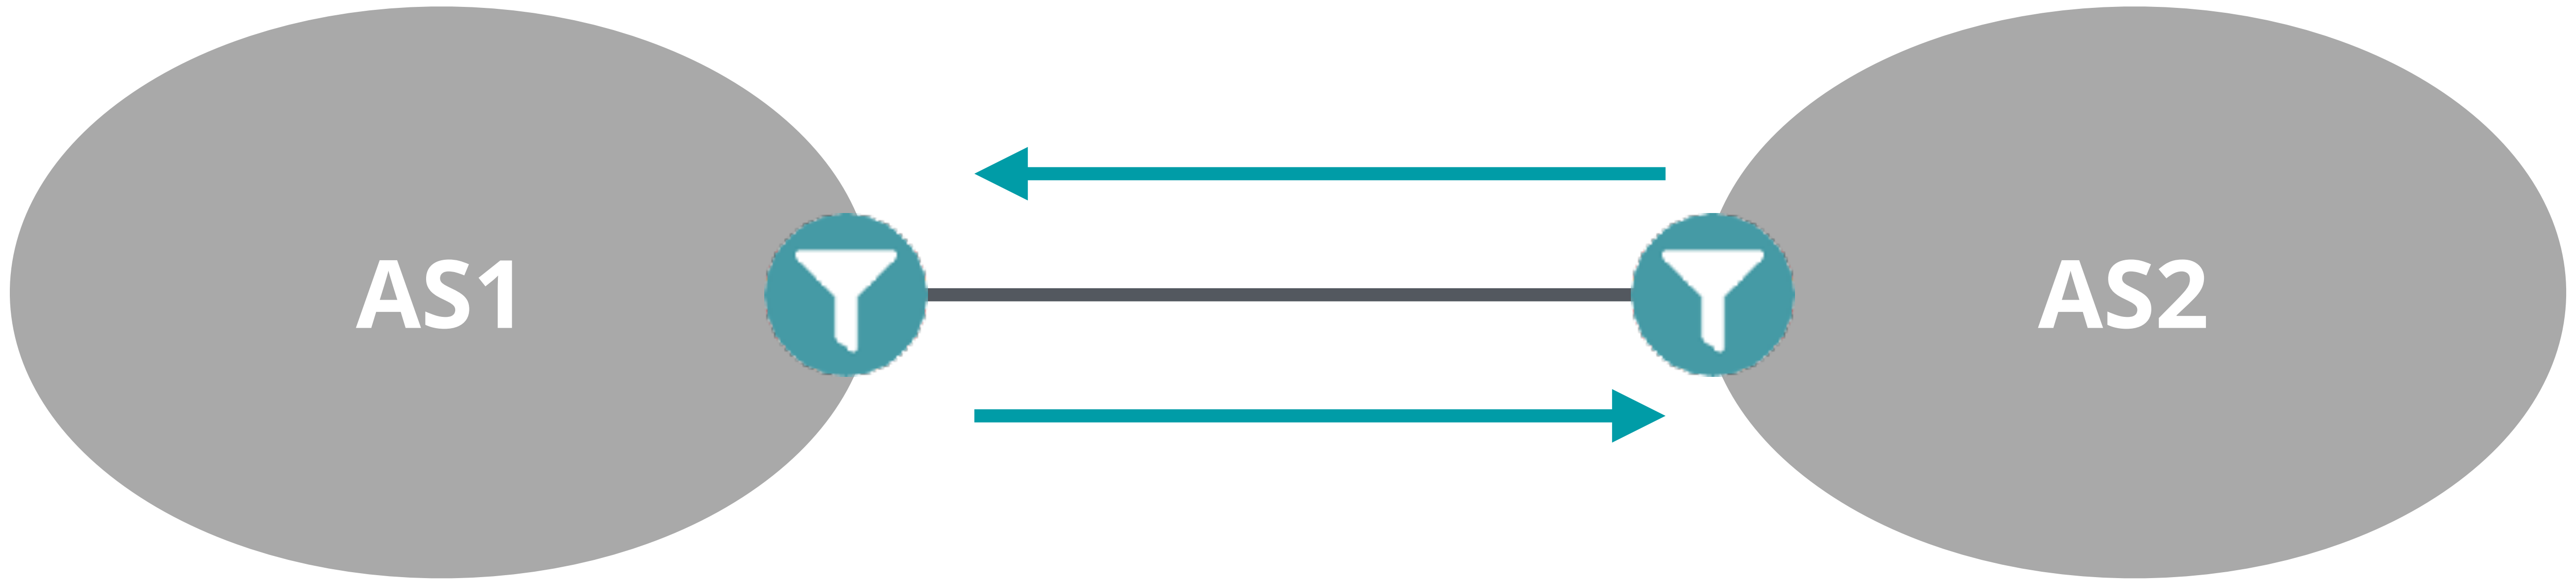

BGP Filters (BGP Policies)

- **Inbound policy:**
  - For **incoming** (received) routes
  - Detects configuration mistakes and attacks
- Should be applied on each eBGP peer
  - Both on ingress and egress
- **Outbound policy:**
  - For **outgoing** (advertised) routes
  - Limits propagation of routing information

Filtering Principles

- **Filter as close to the edge as possible**
- Filter **as precisely** as possible
- Two filtering approaches:
  - **Explicit Permit** (permit then deny any)
  - **Explicit Deny** (deny then permit any)

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 tools:**
  - Filtergen (Level3)
  - bgpq4
  - IRRToolSet
  - IRR Power Tools

Easy to use, but not highly scalable

Which Routes Should be Filtered Out?

- Special-purpose prefixes (IPv4/IPv6) (Martians)
- Unallocated prefixes
- Routes that are too specific
- Prefixes belonging to the local AS
- IXP LAN prefixes
- The default route (0.0.0.0/0, ::/0)

RFC 7454 - "BGP Operations and Security"

- lists the prefixes to be filtered out -

Registering in the IRR System

IRR Support Routing Security

- **The Internet Routing Registry (IRR) composed of many databases:**
  - RIPE NCC, APNIC, RADB, JPIRR, Level3, NTTCom, etc.
  - Operators / tools often take the *union* of entries over the databases
- **Their information can be used to:**
  - Improve stability and consistency of routing
  - Provide global view of routing policies
  - Automation of creating BGP filters
  - Network Troubleshooting

Source: http://www.irr.net

Why Register Routing Information?

- **Document your routing policy**
  - Associate network prefixes with an **origin AS**
- **Helps to filter unauthorised announcements**
  - **Mitigates route hijacks and denial-of-service**
- **Many transit providers and IXPs require it**
  - **They build their filters based on the Routing Registry**

The RIPE Routing Registry

- A subset of the RIPE Database and part of the global IRR
- Used for registering routing policy information
- Includes several objects

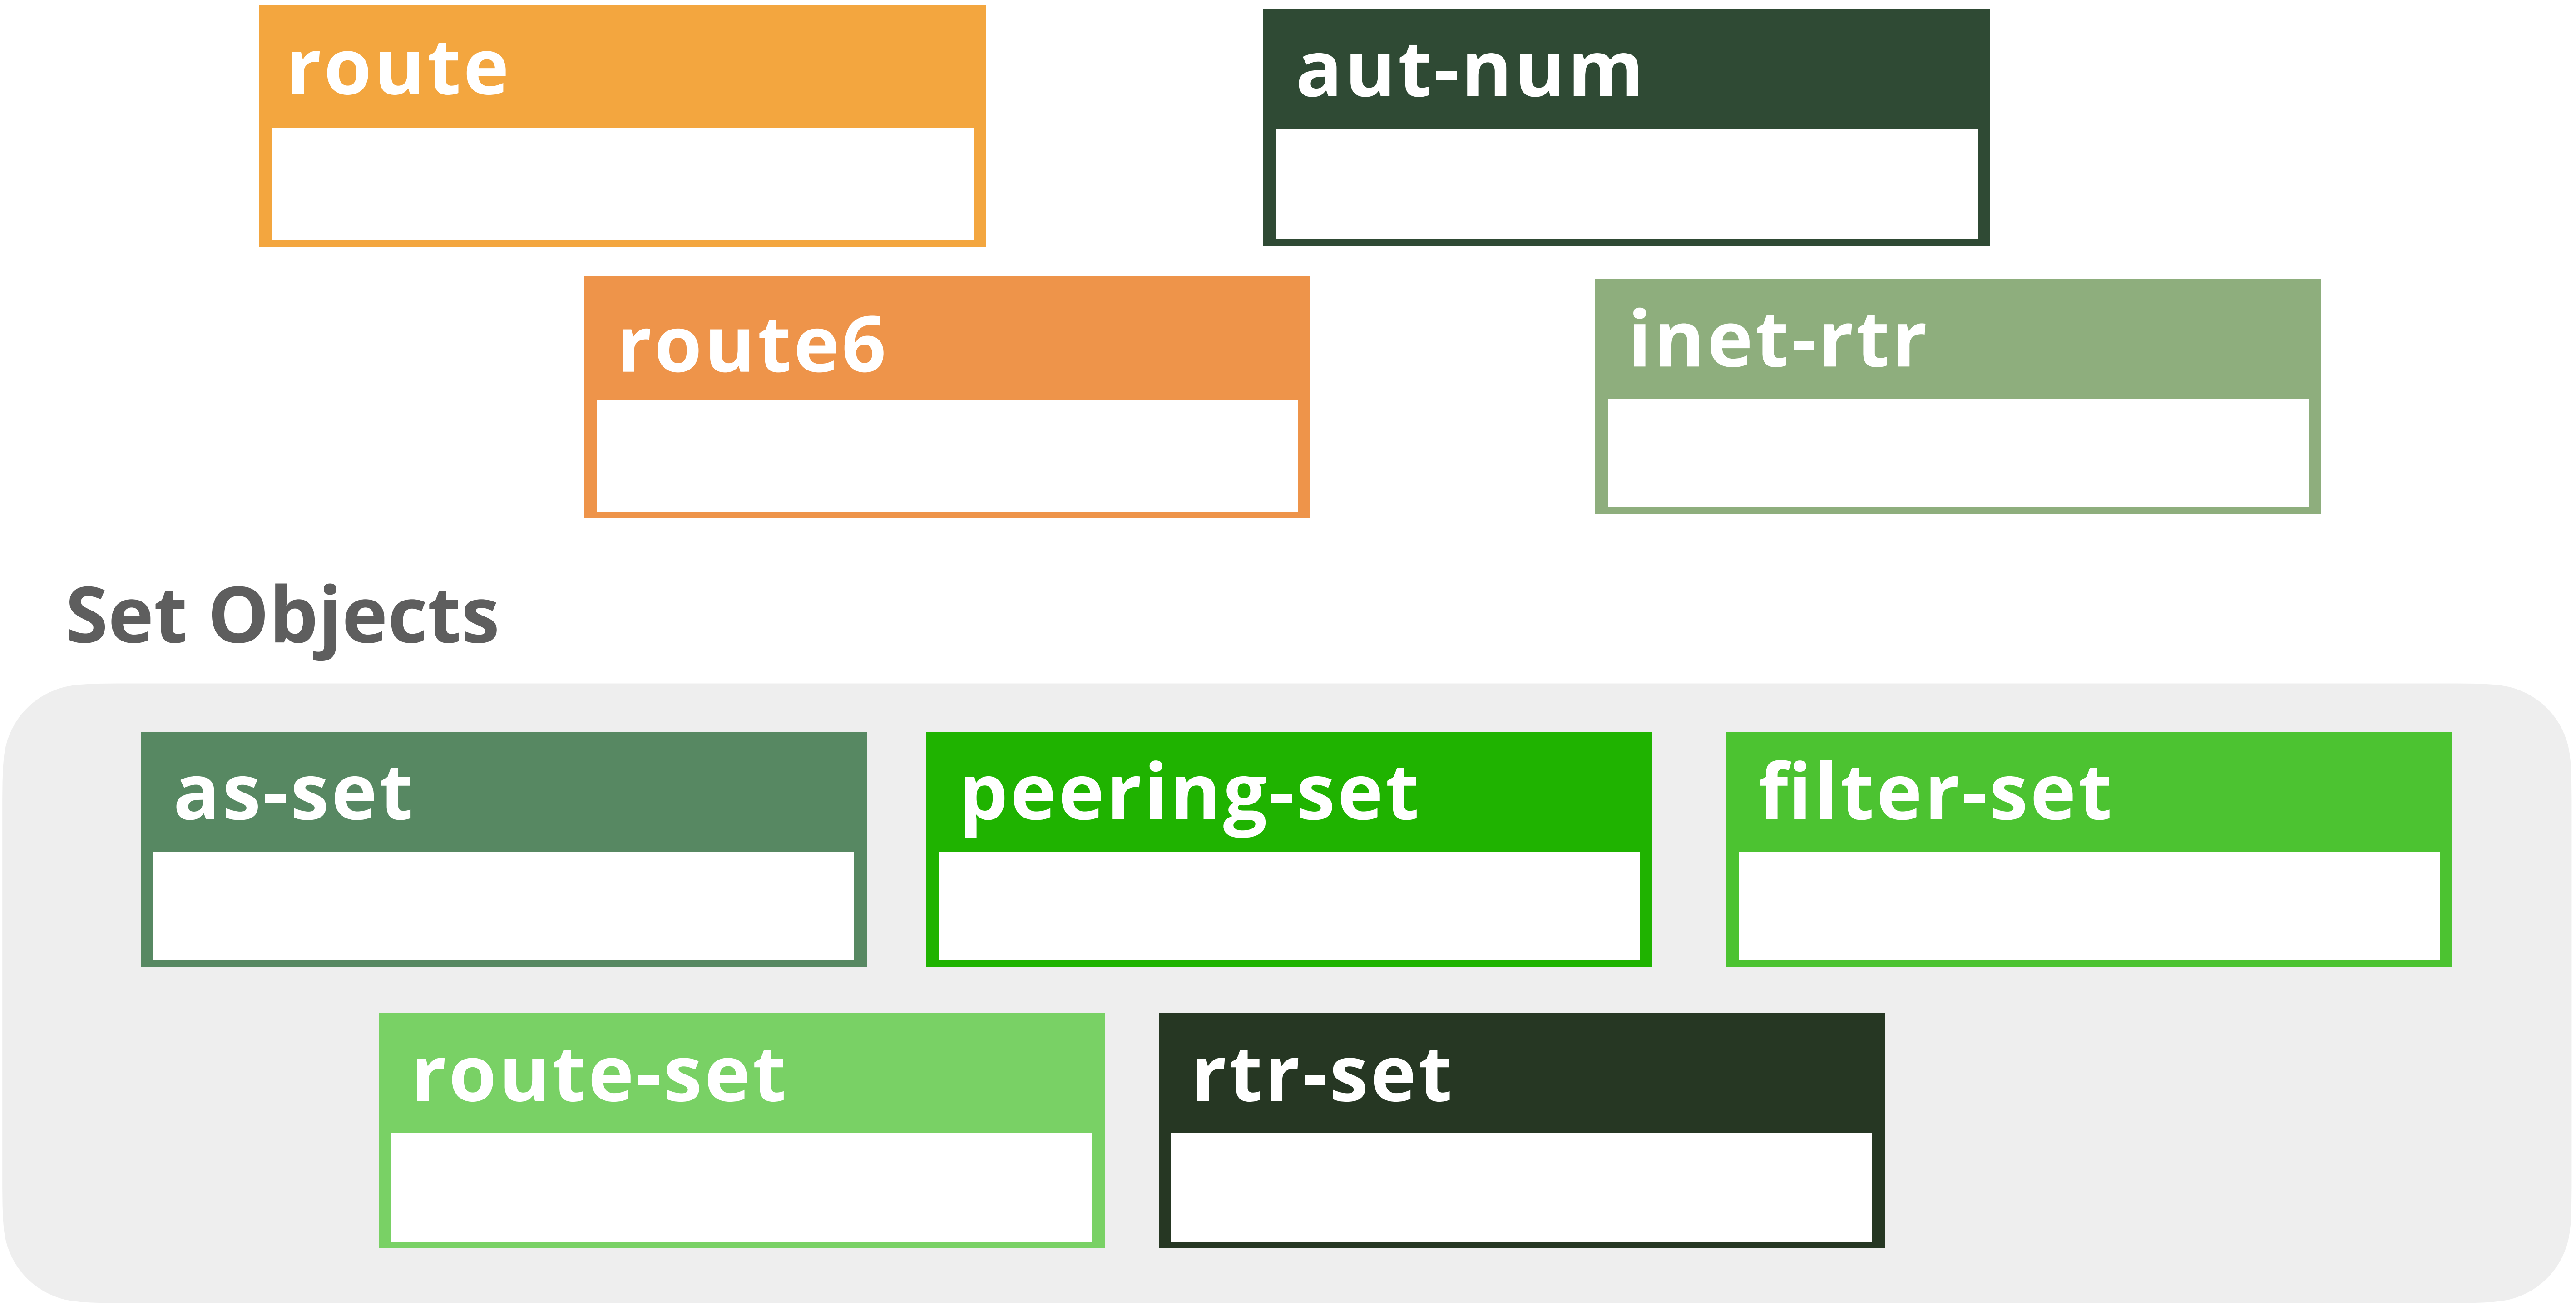

Source: https://apps.db.ripe.net/docs/RPSL-Object-Types/

Route & Route6 Objects

- Contains routing information for IPv4/IPv6 address space
- Specifies from which AS a certain prefix may be originated
- Used for creating BGP filters

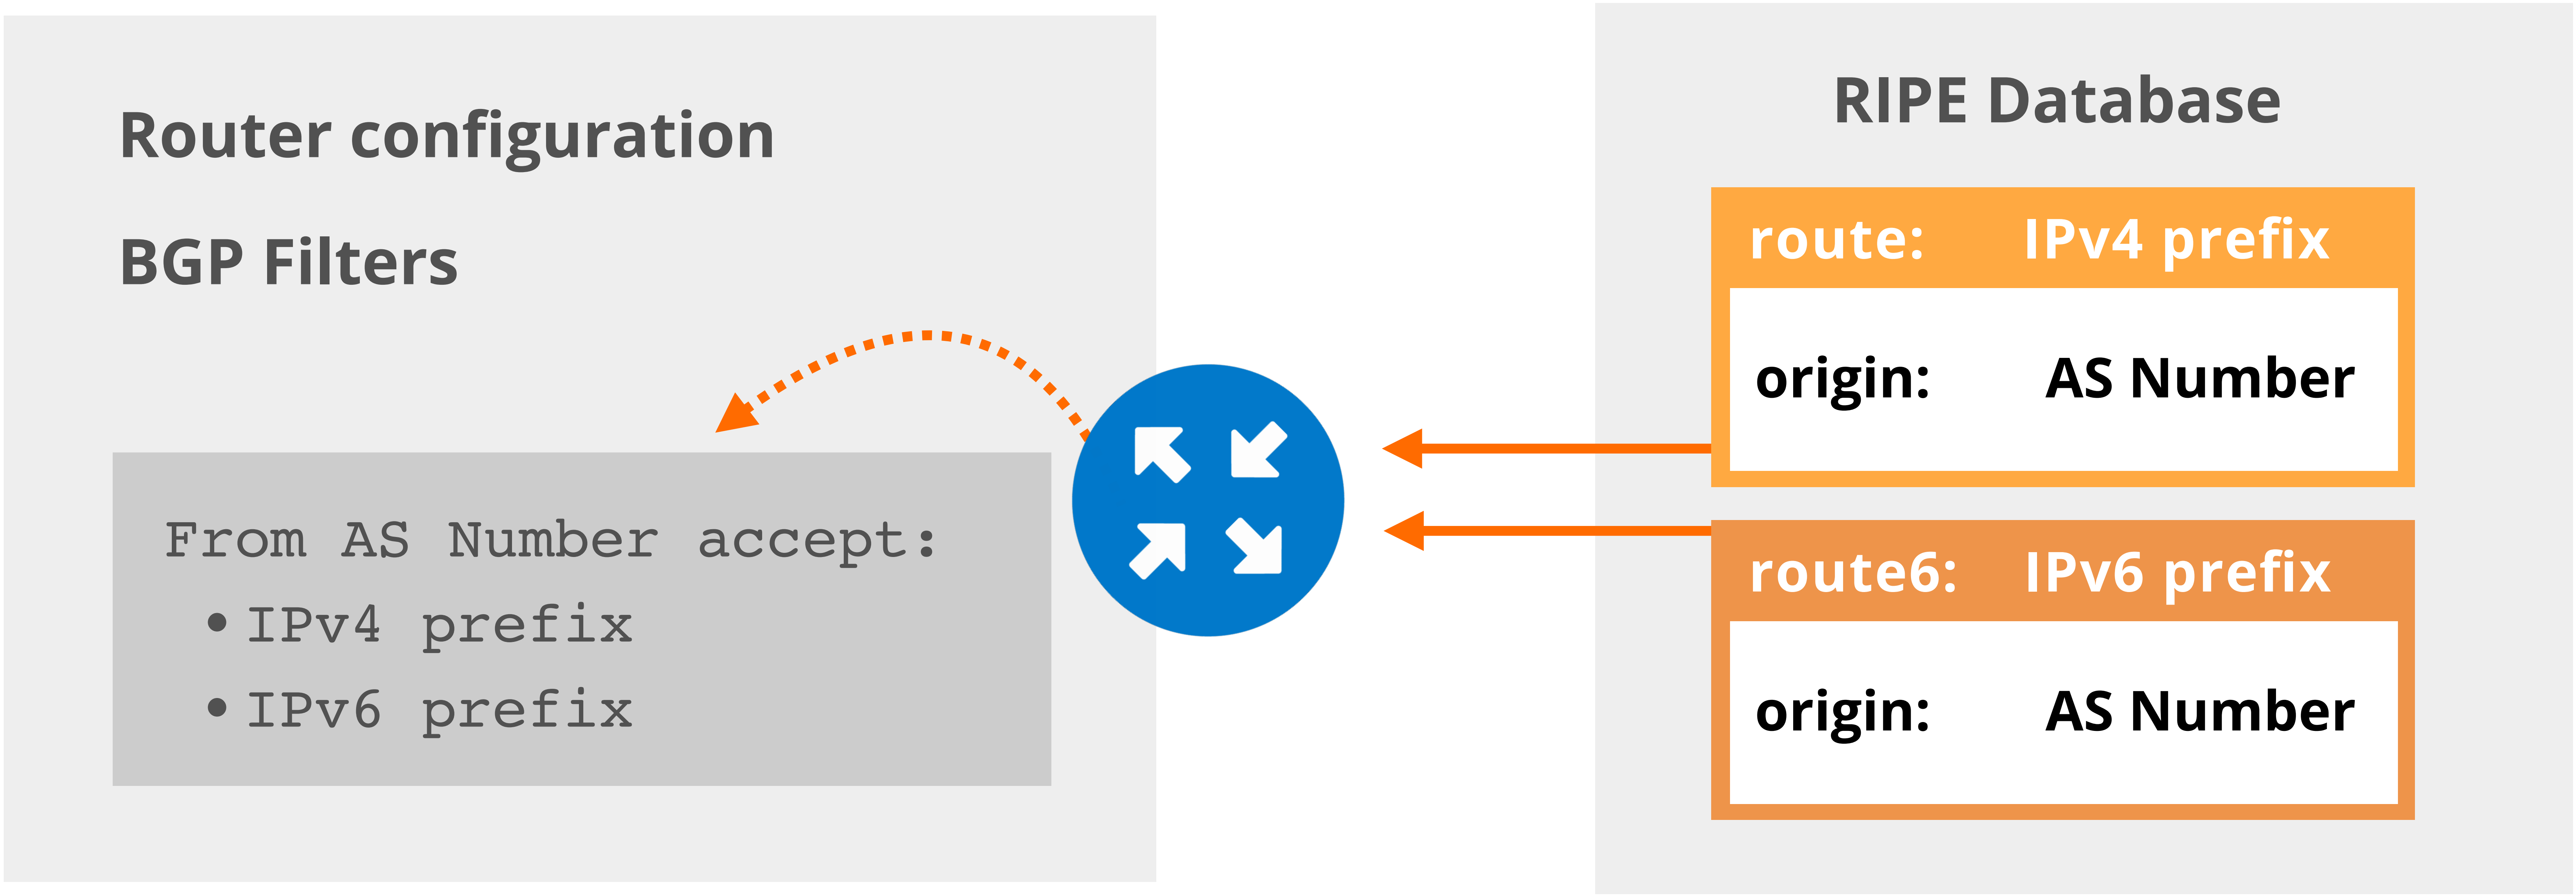

Authorisation Rules for Route(6)

- **You need permission from:**
  1. **inetnum or inet6num**
  2. **route or route6**

1

Allocation

mnt-by: RIPE-NCC-HM-MNT
mnt-by: **DEFAULT-LIR-MNT**
mnt-routes: **ANOTHER-MNT**

2

route(6)

origin: **AS65550**
mnt-by: **ANOTHER-MNT**

* **mnt-routes** delegates the creation of route(6) objects

Registering IP Routes

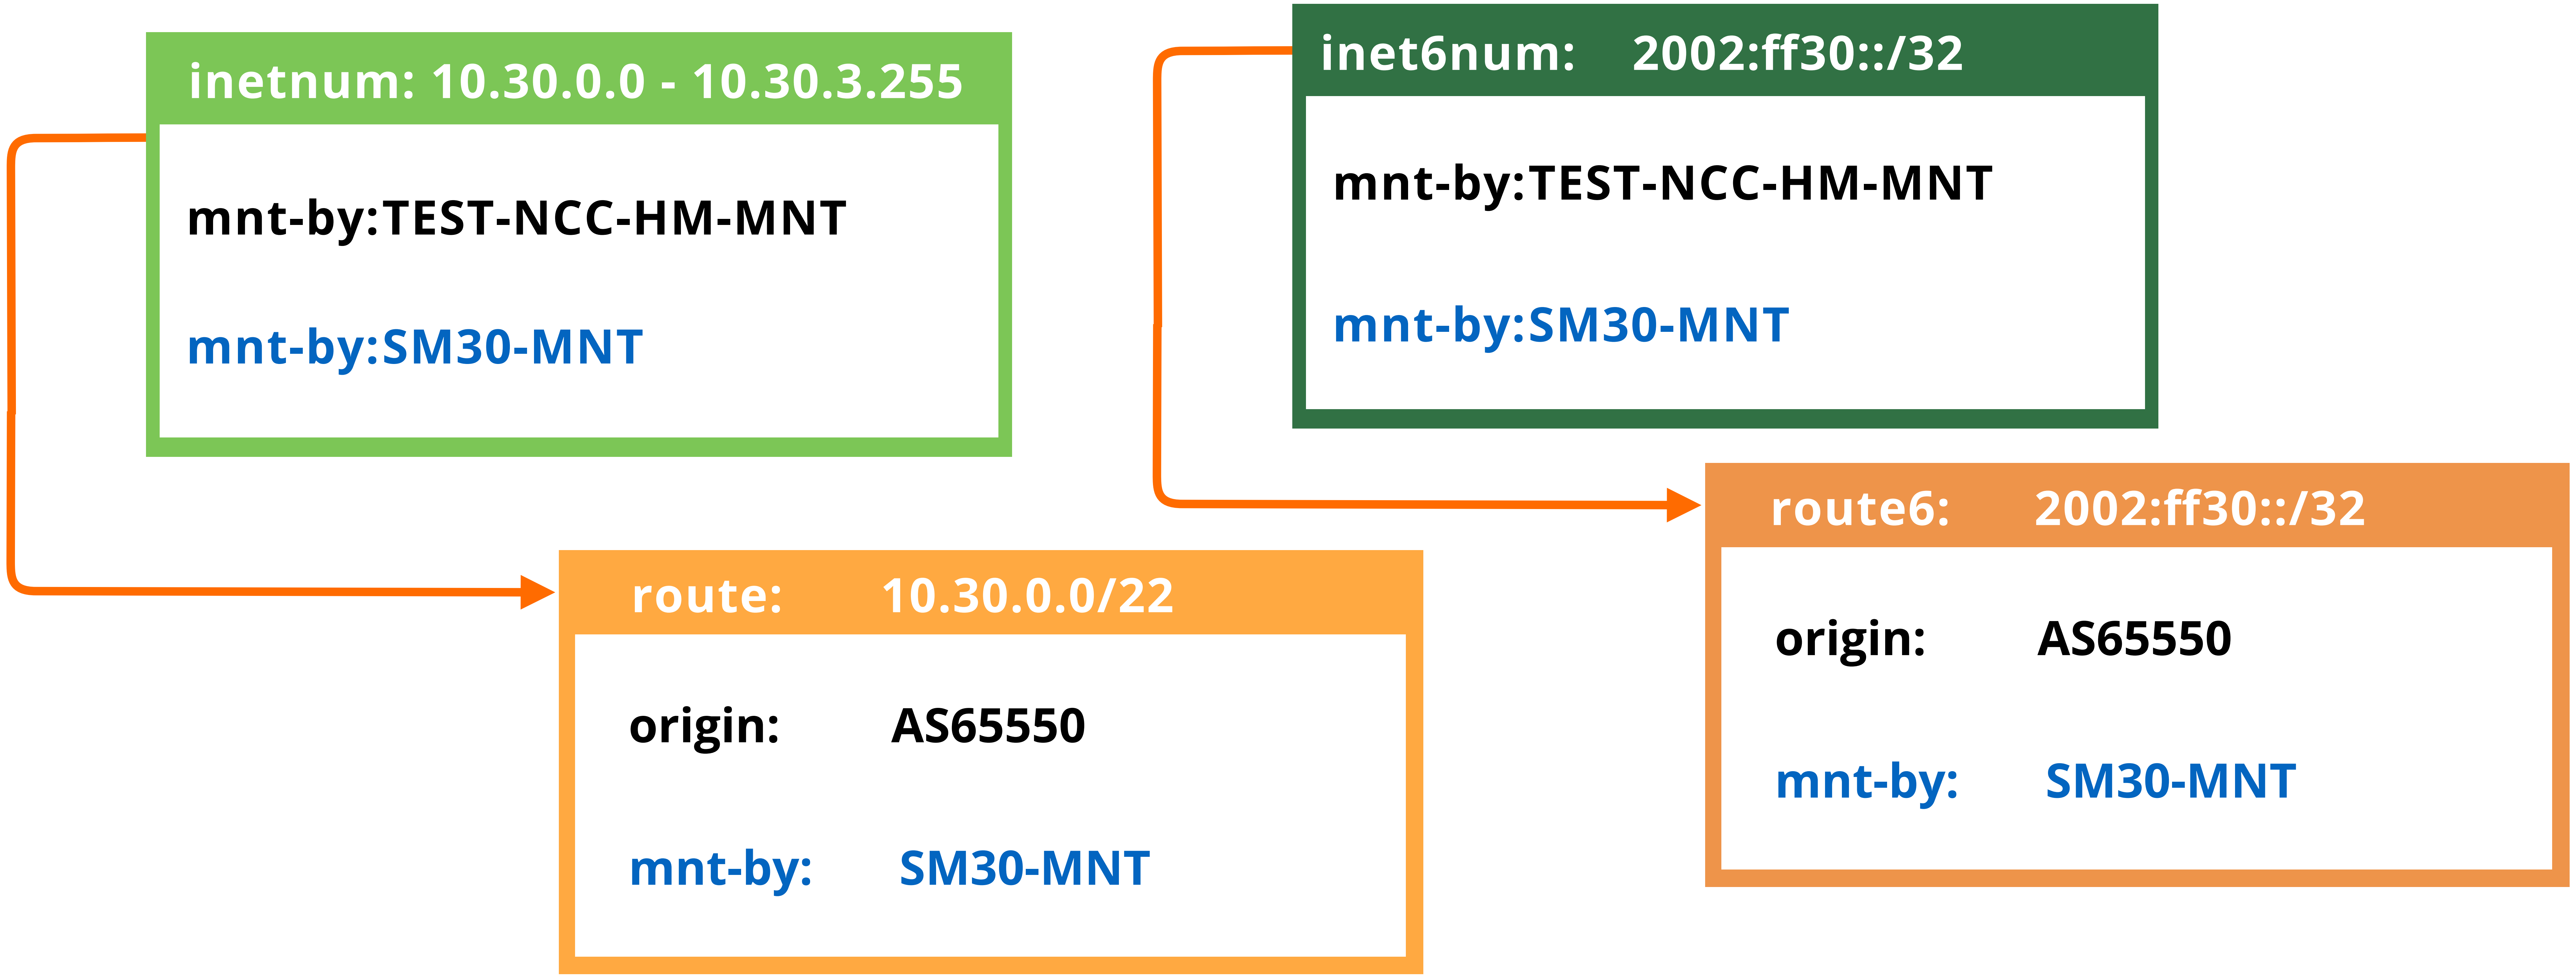

https://docs.db.ripe.net/Appendices/Appendix-D--Route-Object-Creation-Flowchart/#route-object-creation-flowchart

aut-num

aut-num:	AS64500
as-name:	YOUR-AS-NAME
org:	ORG-EE2-RIPE
import:	**from AS65550 accept ANY**
export:	**to AS65550 announce AS64500**
import:	**from AS64496 accept ANY**
export:	**to AS64496 announce AS64500**
admin-c:	DV789-RIPE
tech-c:	JS123-RIPE
status:	ASSIGNED
mnt-by:	RIPE-NCC-END-MNT
mnt-by:	**DEFAULT-LIR-MNT**
source:	**RIPE**

Registers **who** holds that AS Number

Defines the routing policy for an AS

- **Import** - specifies which routes you accept
- **Export** - specifies which routes you announce

BGP Routing Policy

- **Who are your BGP peers? Which ASes**
- **What is your BGP relationship with them?**
  - Customer, Provider, Peer
- **What are your routing decisions?**
  - Which prefixes to accept?
  - Which prefixes to announce?
  - Which prefixes will be preferred in case of multiple routes?

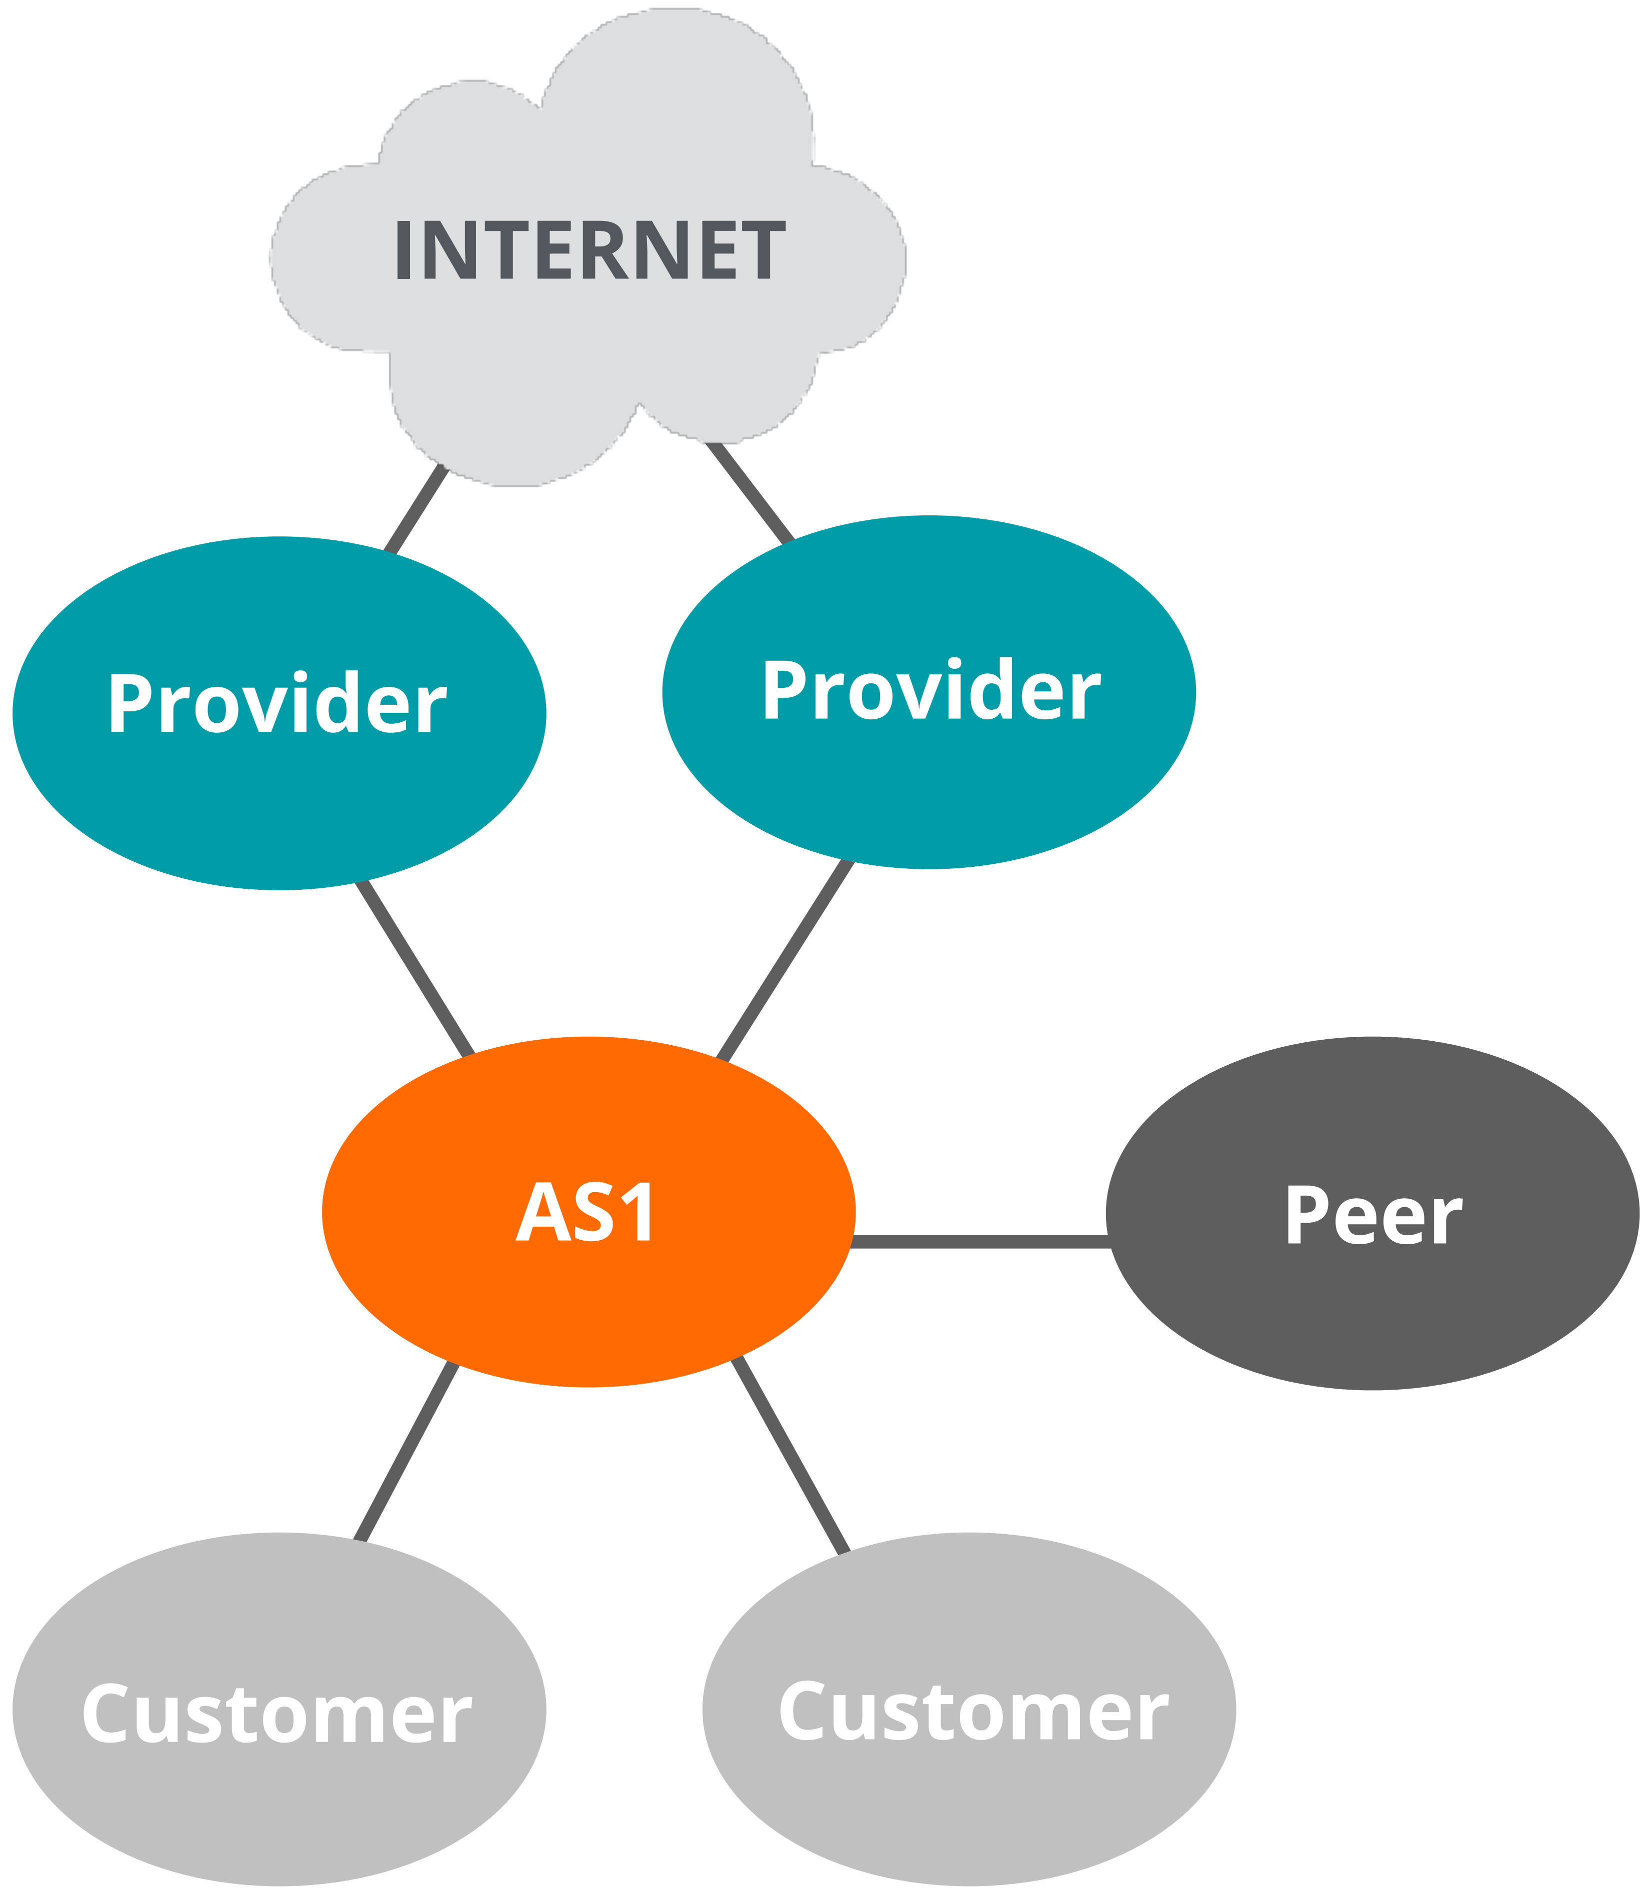

IRRs Use RPSL Language

- **RPSL - Routing Policy Specification Language**
- **Allows network operators to specify their routing policies**
  - Generic way to describe BGP configuration in the IRR
  - Not vendor-specific
- Originated from a RIPE Document (RIPE-181)
- Can be translated into router configuration

RFC 2622 - Routing Policy Specification Language

RFC 2650 - Using RPSL in Practice

Defining Routing Policy in RPSL

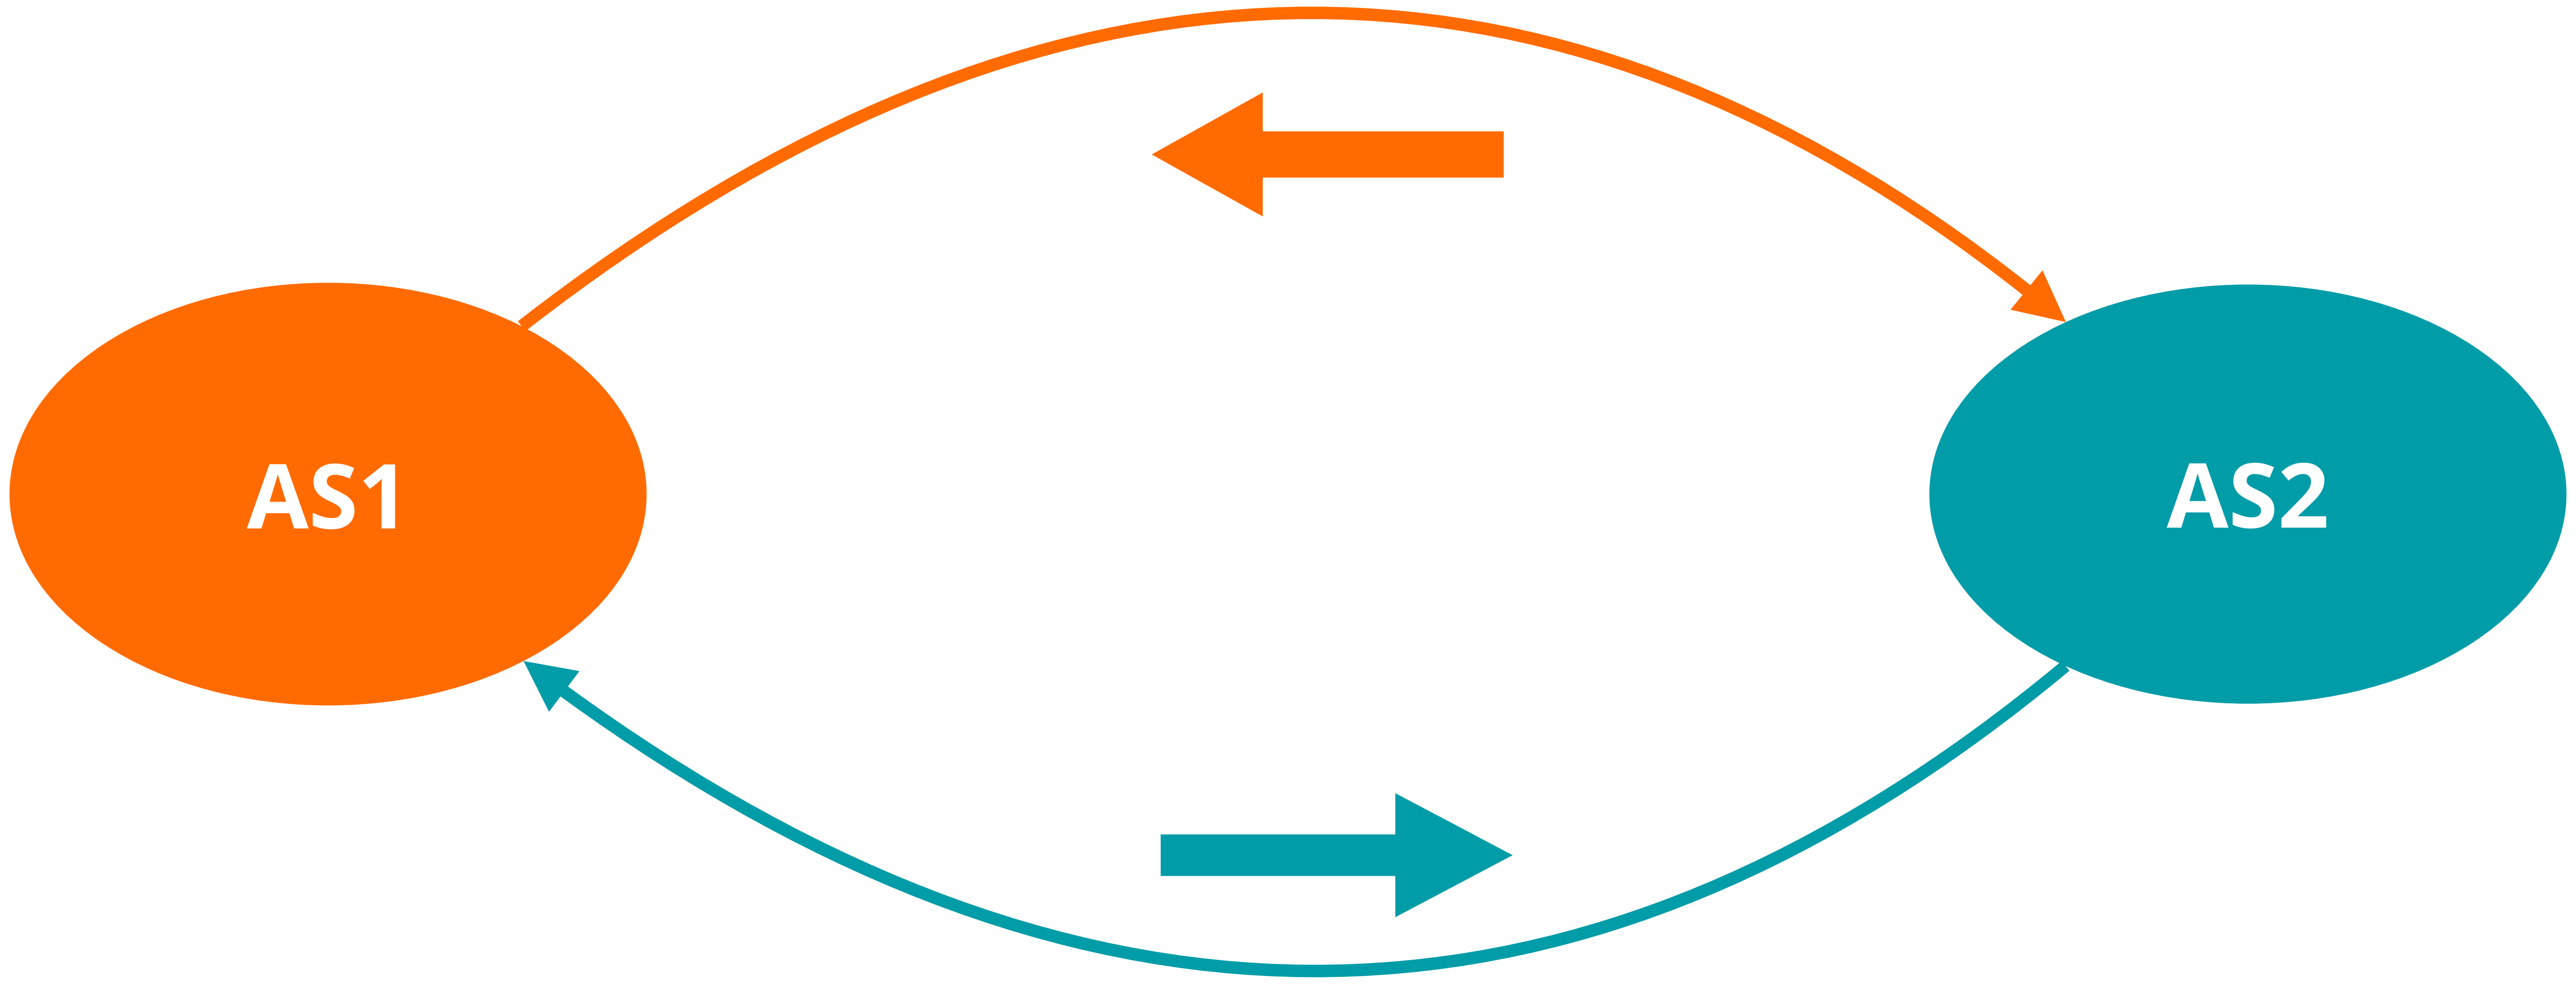


```
aut-num:  AS1
import: from AS2 accept AS2
export: to AS2 announce AS1
```

Routing Policy Example

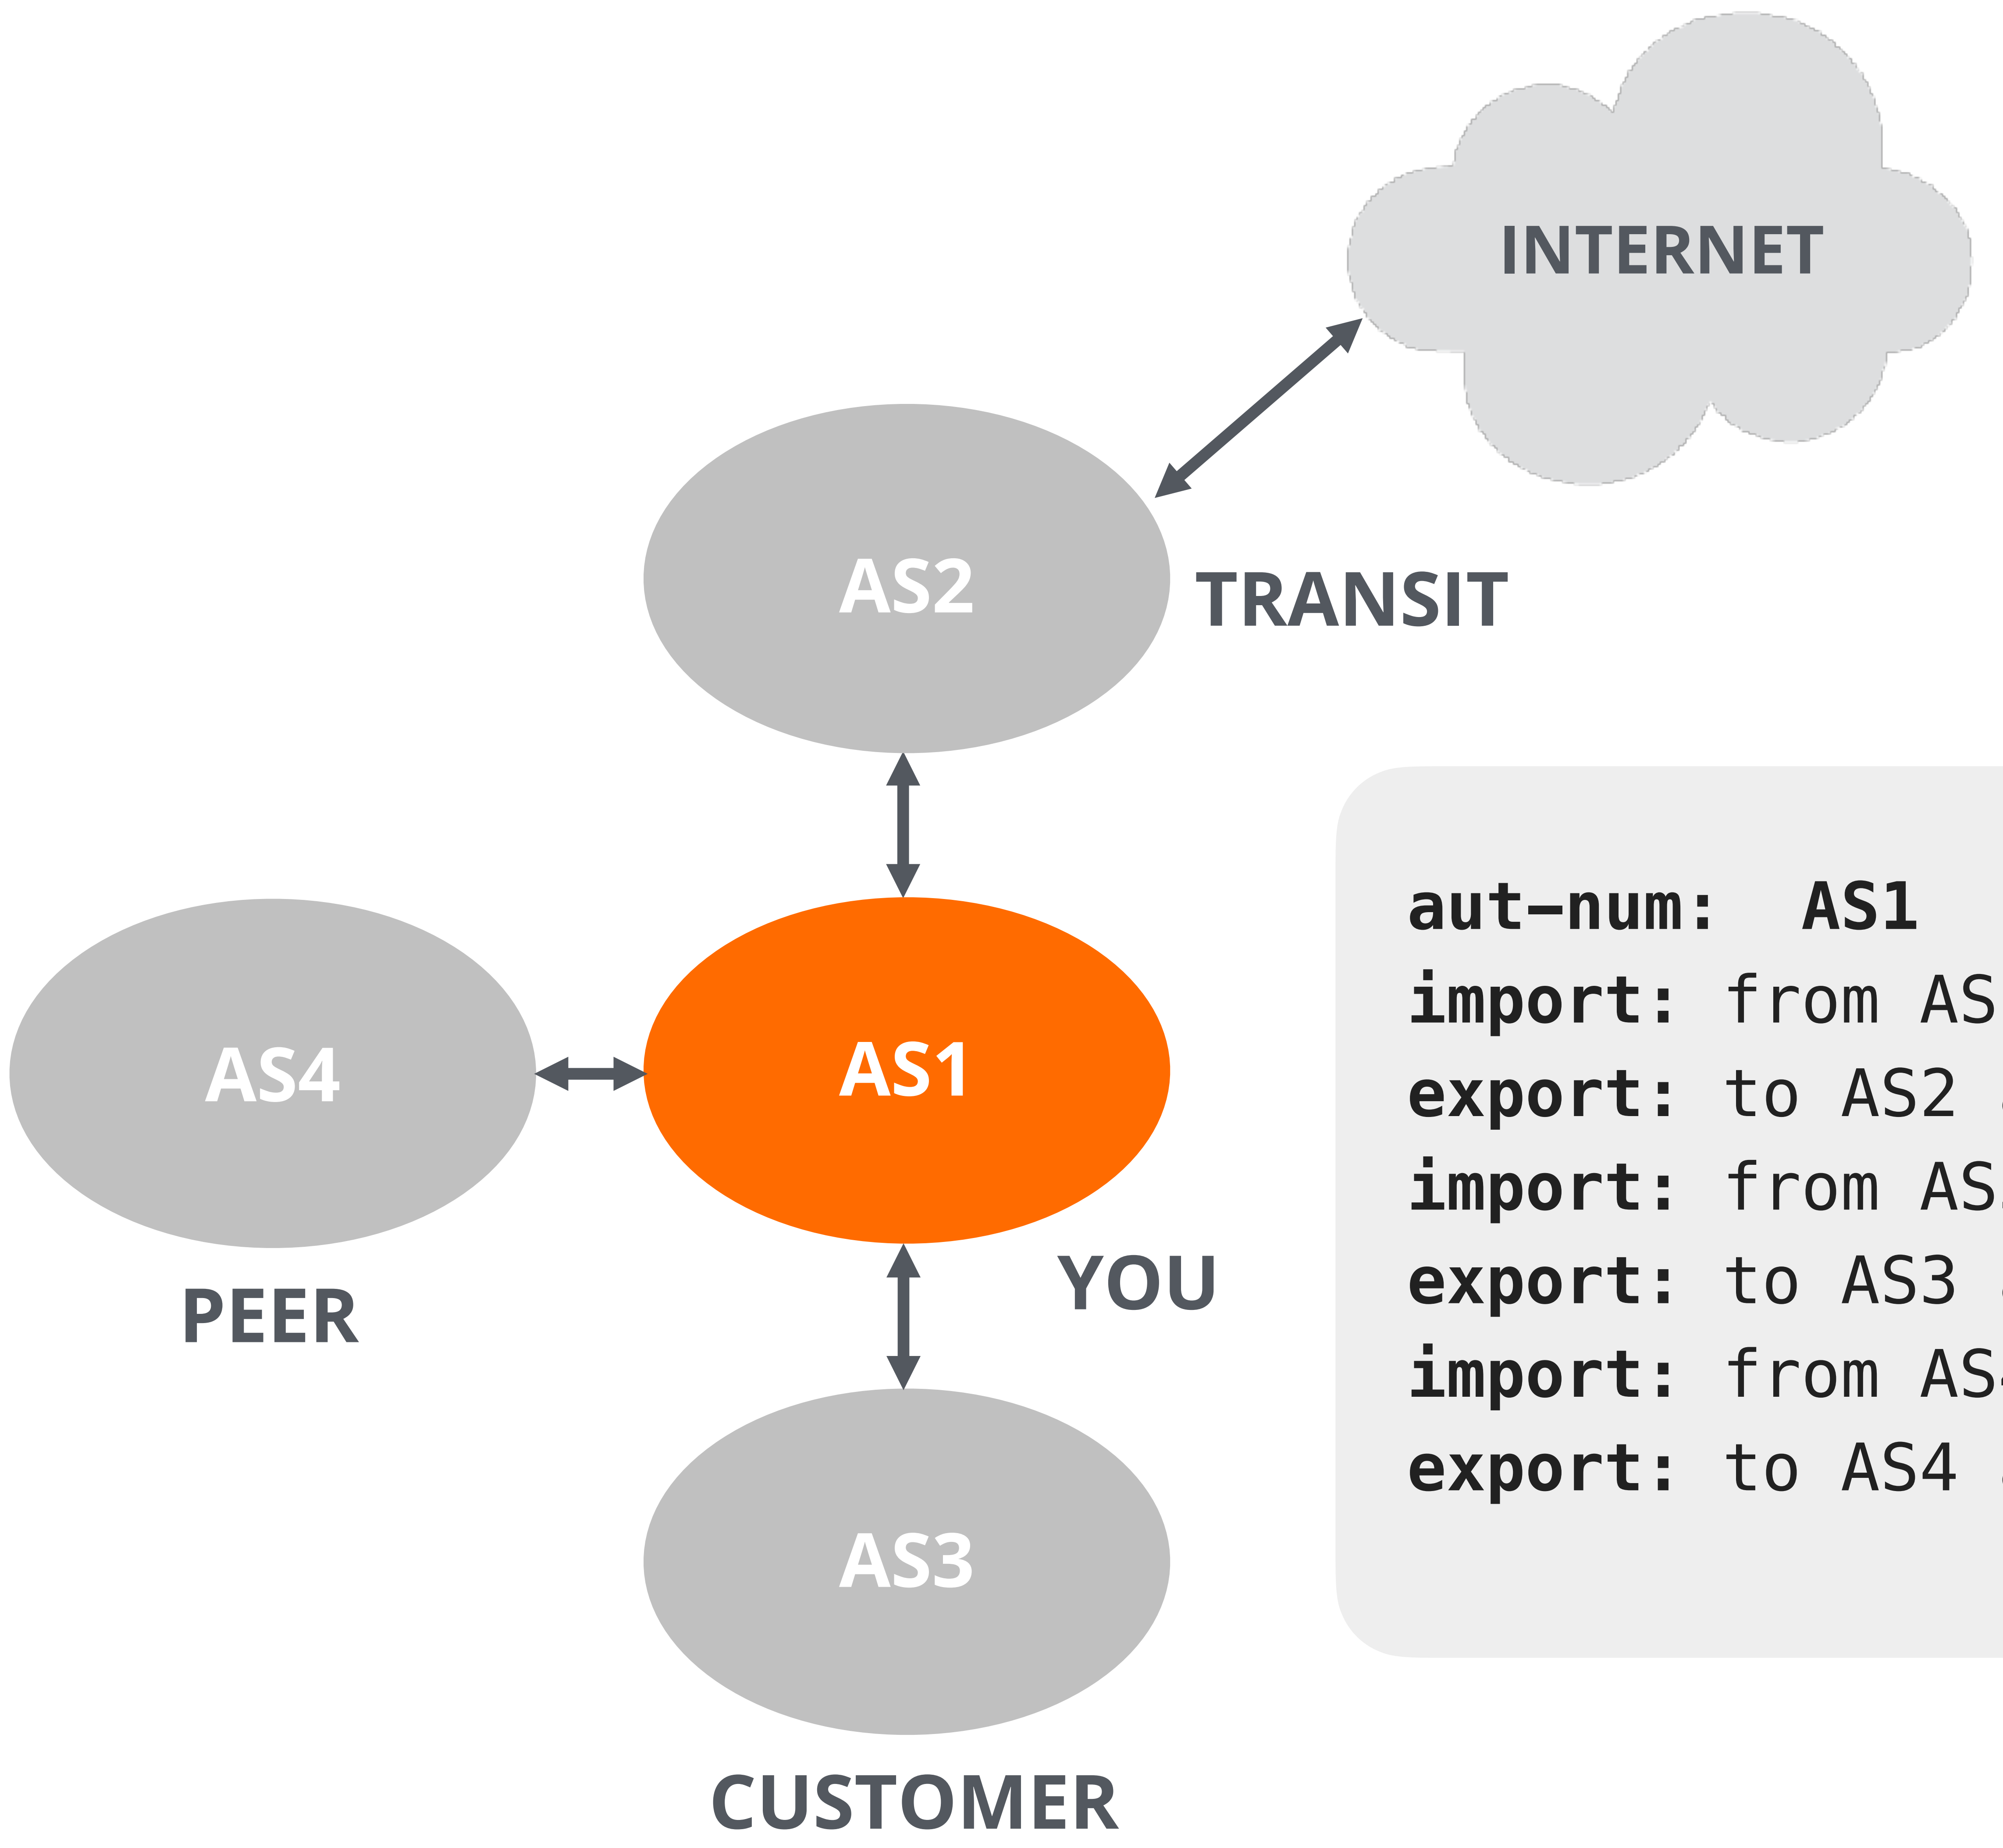


```
aut-num:  AS1
import: from AS2 accept ANY
export: to AS2 announce AS1 AS3
import: from AS3 accept AS3
export: to AS3 announce ANY
import: from AS4 accept AS4
export: to AS4 announce AS1 AS3
```

RPSL Structure in Practice

BGPq4: a CLI Tool for Creating Prefix Filters

BGPq4: Prefix Filter for a AS

```
$ bgpq4 -J AS3333
policy-options {
replace:
 prefix-list NN {
    193.0.0.0/21;
    193.0.10.0/23;
    193.0.12.0/23;
    193.0.18.0/23;
    193.0.20.0/23;
    193.0.22.0/23;
    193.230.194.0/24;
 }
}
```

BGPq4: Prefix Filter for a AS-SET: it Expands

```
$ bgpq4 -J AS-RIPENCC
policy-options {
replace:
 prefix-list NN {
    23.128.24.0/24;
    [...32 lines...]
    193.0.0.0/21;
    193.0.10.0/23;
    193.0.12.0/23;
    193.0.18.0/23;
    193.0.20.0/23;
    193.0.22.0/23;
    193.0.24.0/21;
    193.230.194.0/24;
 }
}
```

BGPq4: Prefix Filter for a Network: This Recurses

```
$ bgpq4 -J AS3320:AS-DTAG
policy-options {
replace:
 prefix-list NN {
    0.242.236.0/23;
    [...1.872.303 lines...]
    223.255.254.0/24;
    230.22.60.0/24;
    233.27.98.0/24;
    233.31.187.0/24;
    233.160.91.0/24;
    233.184.222.0/24;
    233.191.108.0/24;
    233.199.75.0/24;
    233.227.187.0/24;
    233.236.58.0/24;
 }
```

Reality Check

- **The IRR system has limitations**
  - Conflicting data, no central authority, no holdership checks, not updated
- **It is still widely used**
- **Improving IRR accuracy**
  - Keep your IRR information up to date
  - Route filtering using IRRdv4 (validates against RPKI)
  - IRR databases should remove inconsistent records regularly

IRRd https://github.com/irrdnet/irrd/

The Knowledge is in the Community

RIPE Meetings

- Five-day event where ISPs, network operators and other interested parties gather to:
  - **Discuss** policies and procedures to allocate IP addresses and ASNs
  - **Learn about** current technical and policy issues
  - **Share** experiences, latest developments and best common practices
  - **Network** with peers
- Usually held twice a year

- Student tickets available

RIPE Fellowship

- Aims to increase the diversity within the RIPE community:
  - A good geographical spread
  - Diversity of stakeholder groups and interests
  - Gender balance
- Funding to attend RIPE and regional meetings
  - ripe.net/fellowship

RIPE Academic Cooperation Initiative

- Connects academia with the RIPE community
- Funding to attend RIPE and all regional meetings (SEE, CAPIF, MENOG)
- Join the mailing list:
  - https://www.ripe.net/mailman/listinfo/raci-list
- Past RACI attendees:
  - https://www.ripe.net/participate/ripe/raci/alumni
- **ripe.net/raci**

Network Operators Groups (NOGs)

- Informal groups of local Network Operators
- Forum for exchange between operators about issues/problems/ current events in the networking world
- Communication via mailing lists, IMs, meetings
- Documentation of best practices (e.g https://bgpfilterguide.nlnog.net/)
- labs.ripe.net/nogs

RIPE NCC Hackathons

- Hackathons info, calendar & list:
  - https://www.ripe.net/meetings/hackathons/
  - labs.ripe.net/hackathons
- Upcoming: DNS hackathon, March 2025, Stockholm
  - https://labs.ripe.net/author/becha/join-the-dns-hackathon-2025/
- Just finished: Green Tech hackathon
  - https://labs.ripe.net/author/becha/approaching-the-green-tech-h
  - http://github.com/RIPE-Atlas-Community/Green-Tech

Community Communication

- Upcoming events:
  - SEE-13 meeting, 7-8 April, Sofia : ripe.net/see-13
  - RIPE90, 12-16 May, Lisbon : RIPE90.ripe.net
  - RIPE91, October, Bucharest
- https://forum.ripe.net
- @ripencc@mastodon.social

RPKI
Basics

What is RPKI?

- A security framework for the Internet
- Verifies the association between resource holders and their resources
  - Attaches digital certificate to IP addresses and AS numbers
- Used to validate the origin of BGP announcements (BGP OV)
  - Is the originating ASN authorised to originate a particular prefix?
  - Helps to mitigate **BGP Origin Hijacks** and **Route leaks**

Motivation: Sub-Prefix Hijacks are/were Common

- 2008: Pakistan Telecom hijacked YouTube
- 2013: Ukranian Telco (Vega) hijacked BT prefixes
- 2014: STE hijacked many prefixes including US DoD, YouTube, Level 3, Akamai…
- 2016: AWS Route Leak
- 2017: Google, Apple, FB, Microsoft, Riot Games prefixes were hijacked
- 2018: Amazon Route 53 Hijack myetherwallet.com cryptocurrency hijack
- 2018: Google Prefix Leak
- 2019: Taiwan Network Information Center Public DNS in Taiwan was hijacked
- 2020: Telstra hijacked 500 prefixes
- 2020: >8k BGP prefixes hijacked, affected companies such as Amazon, Akamai, Alibaba
- 2021: AS55410 (Vodafone Idea) hijacked almost 40k prefixes, belonging to 4k+ ASNs in 164 countries.

RPKI Chain of Trust

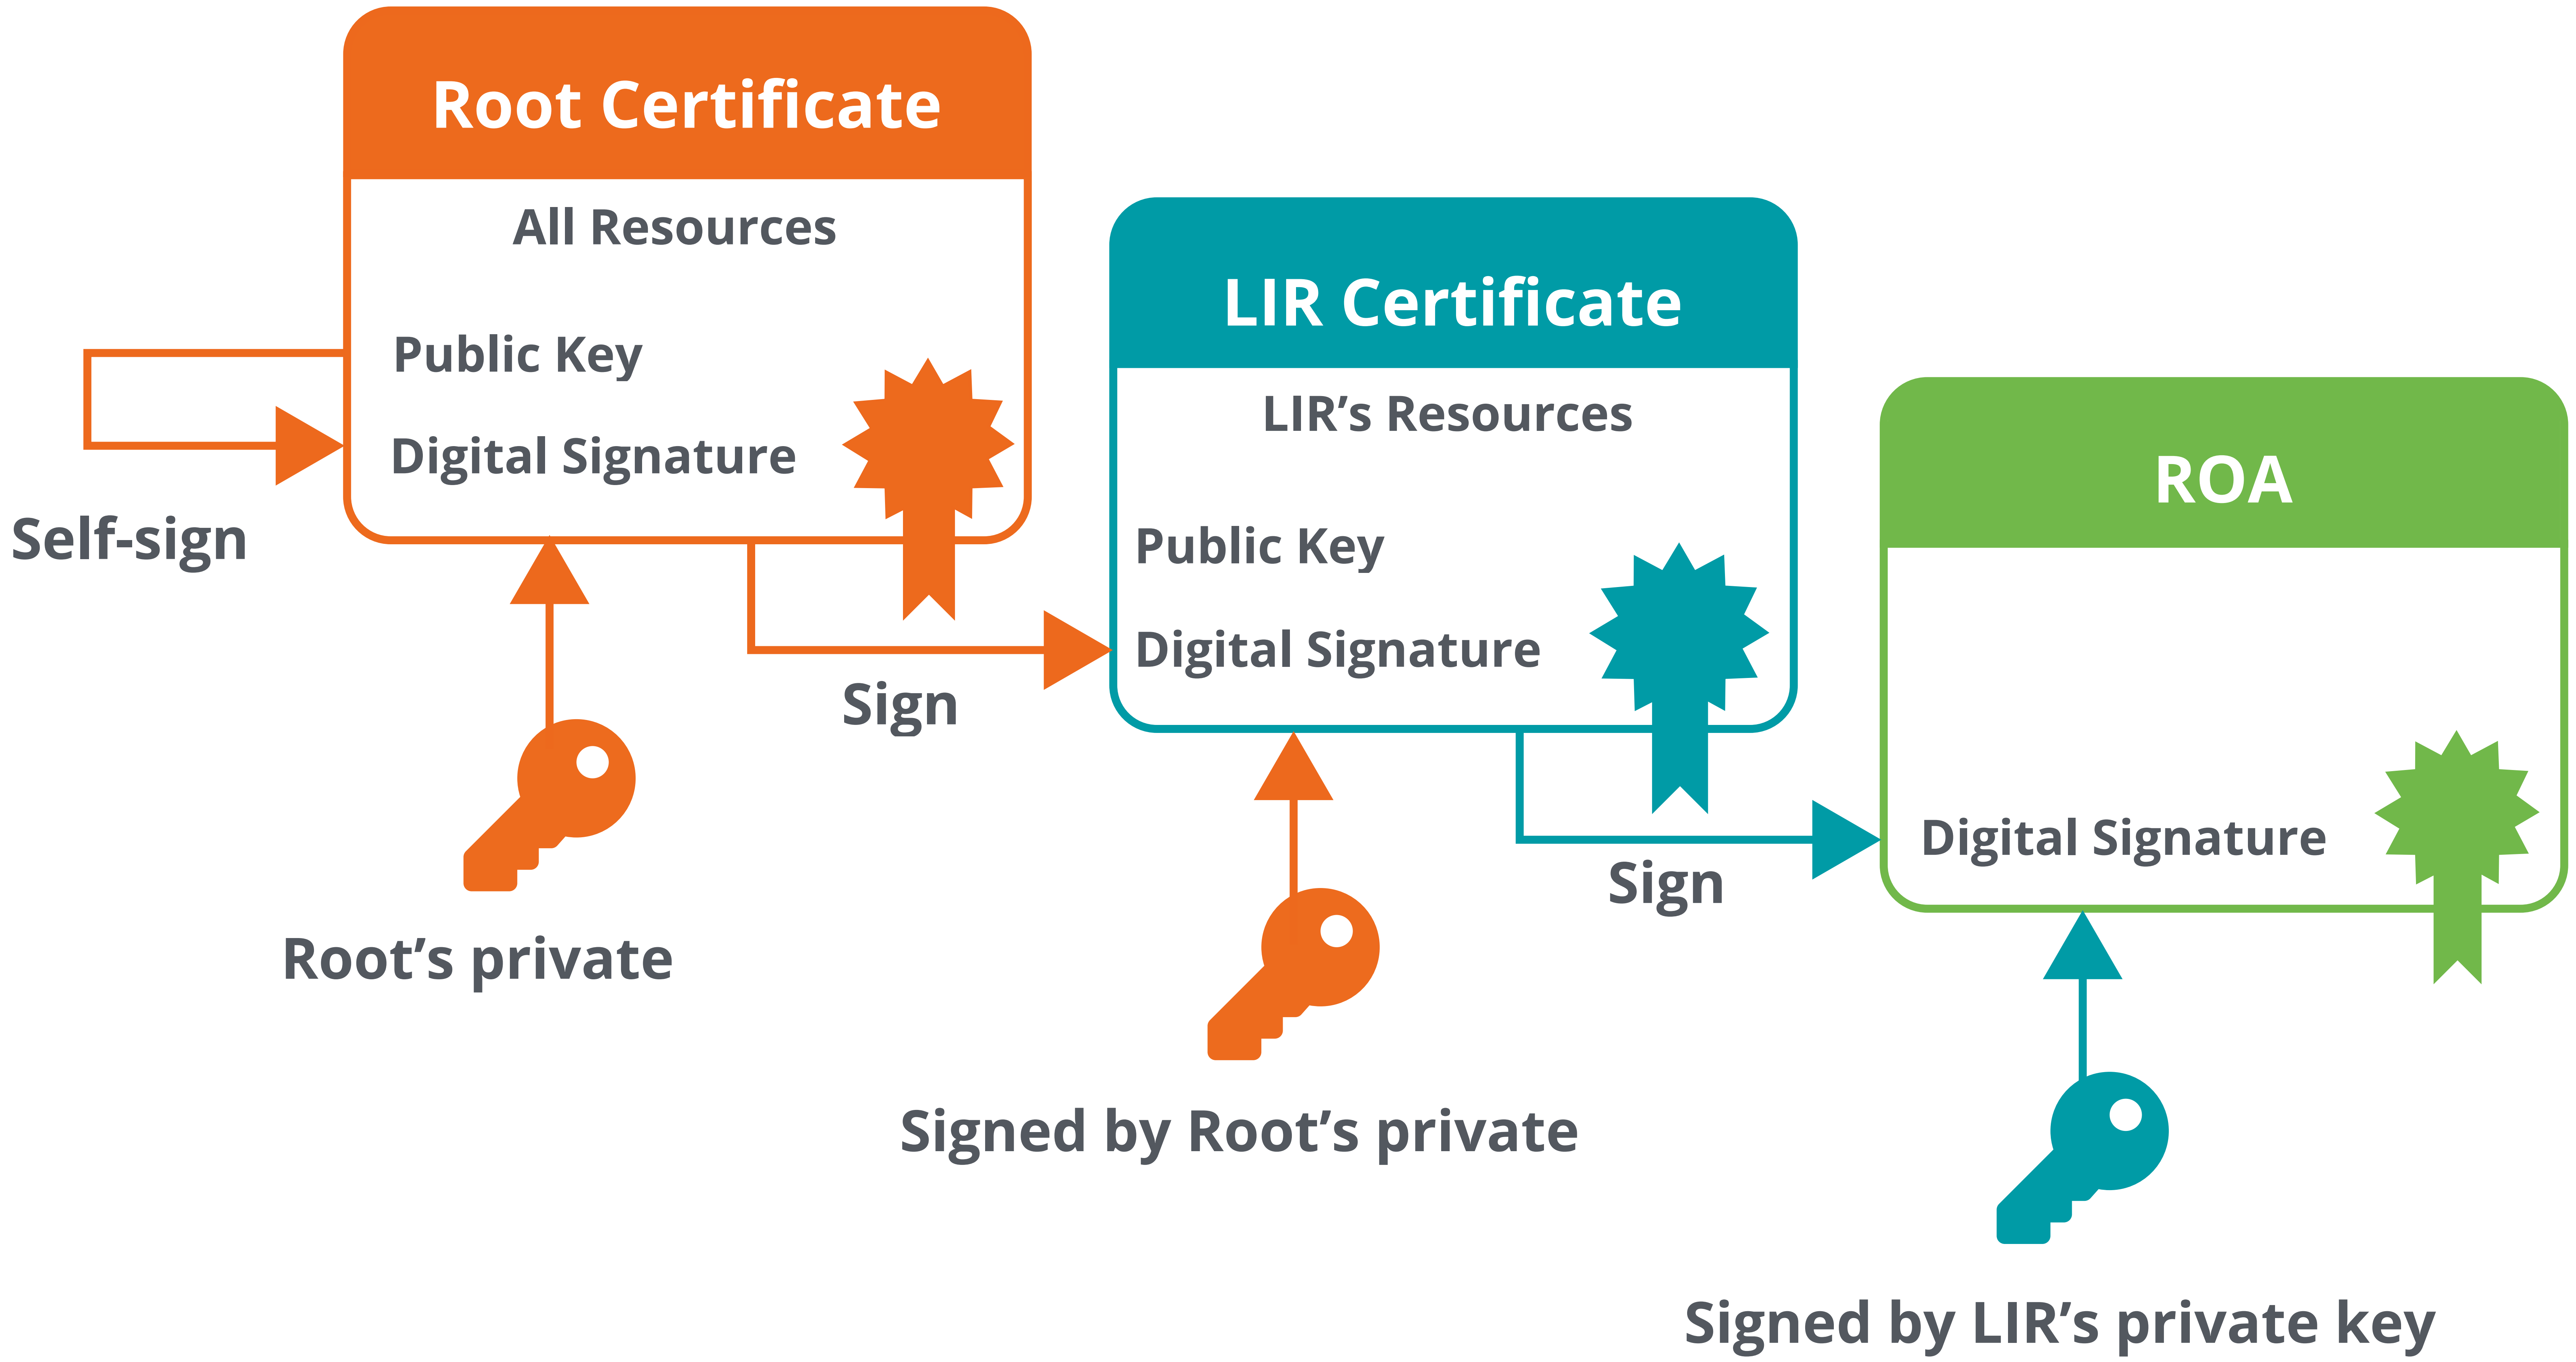

How Does RPKI Work?

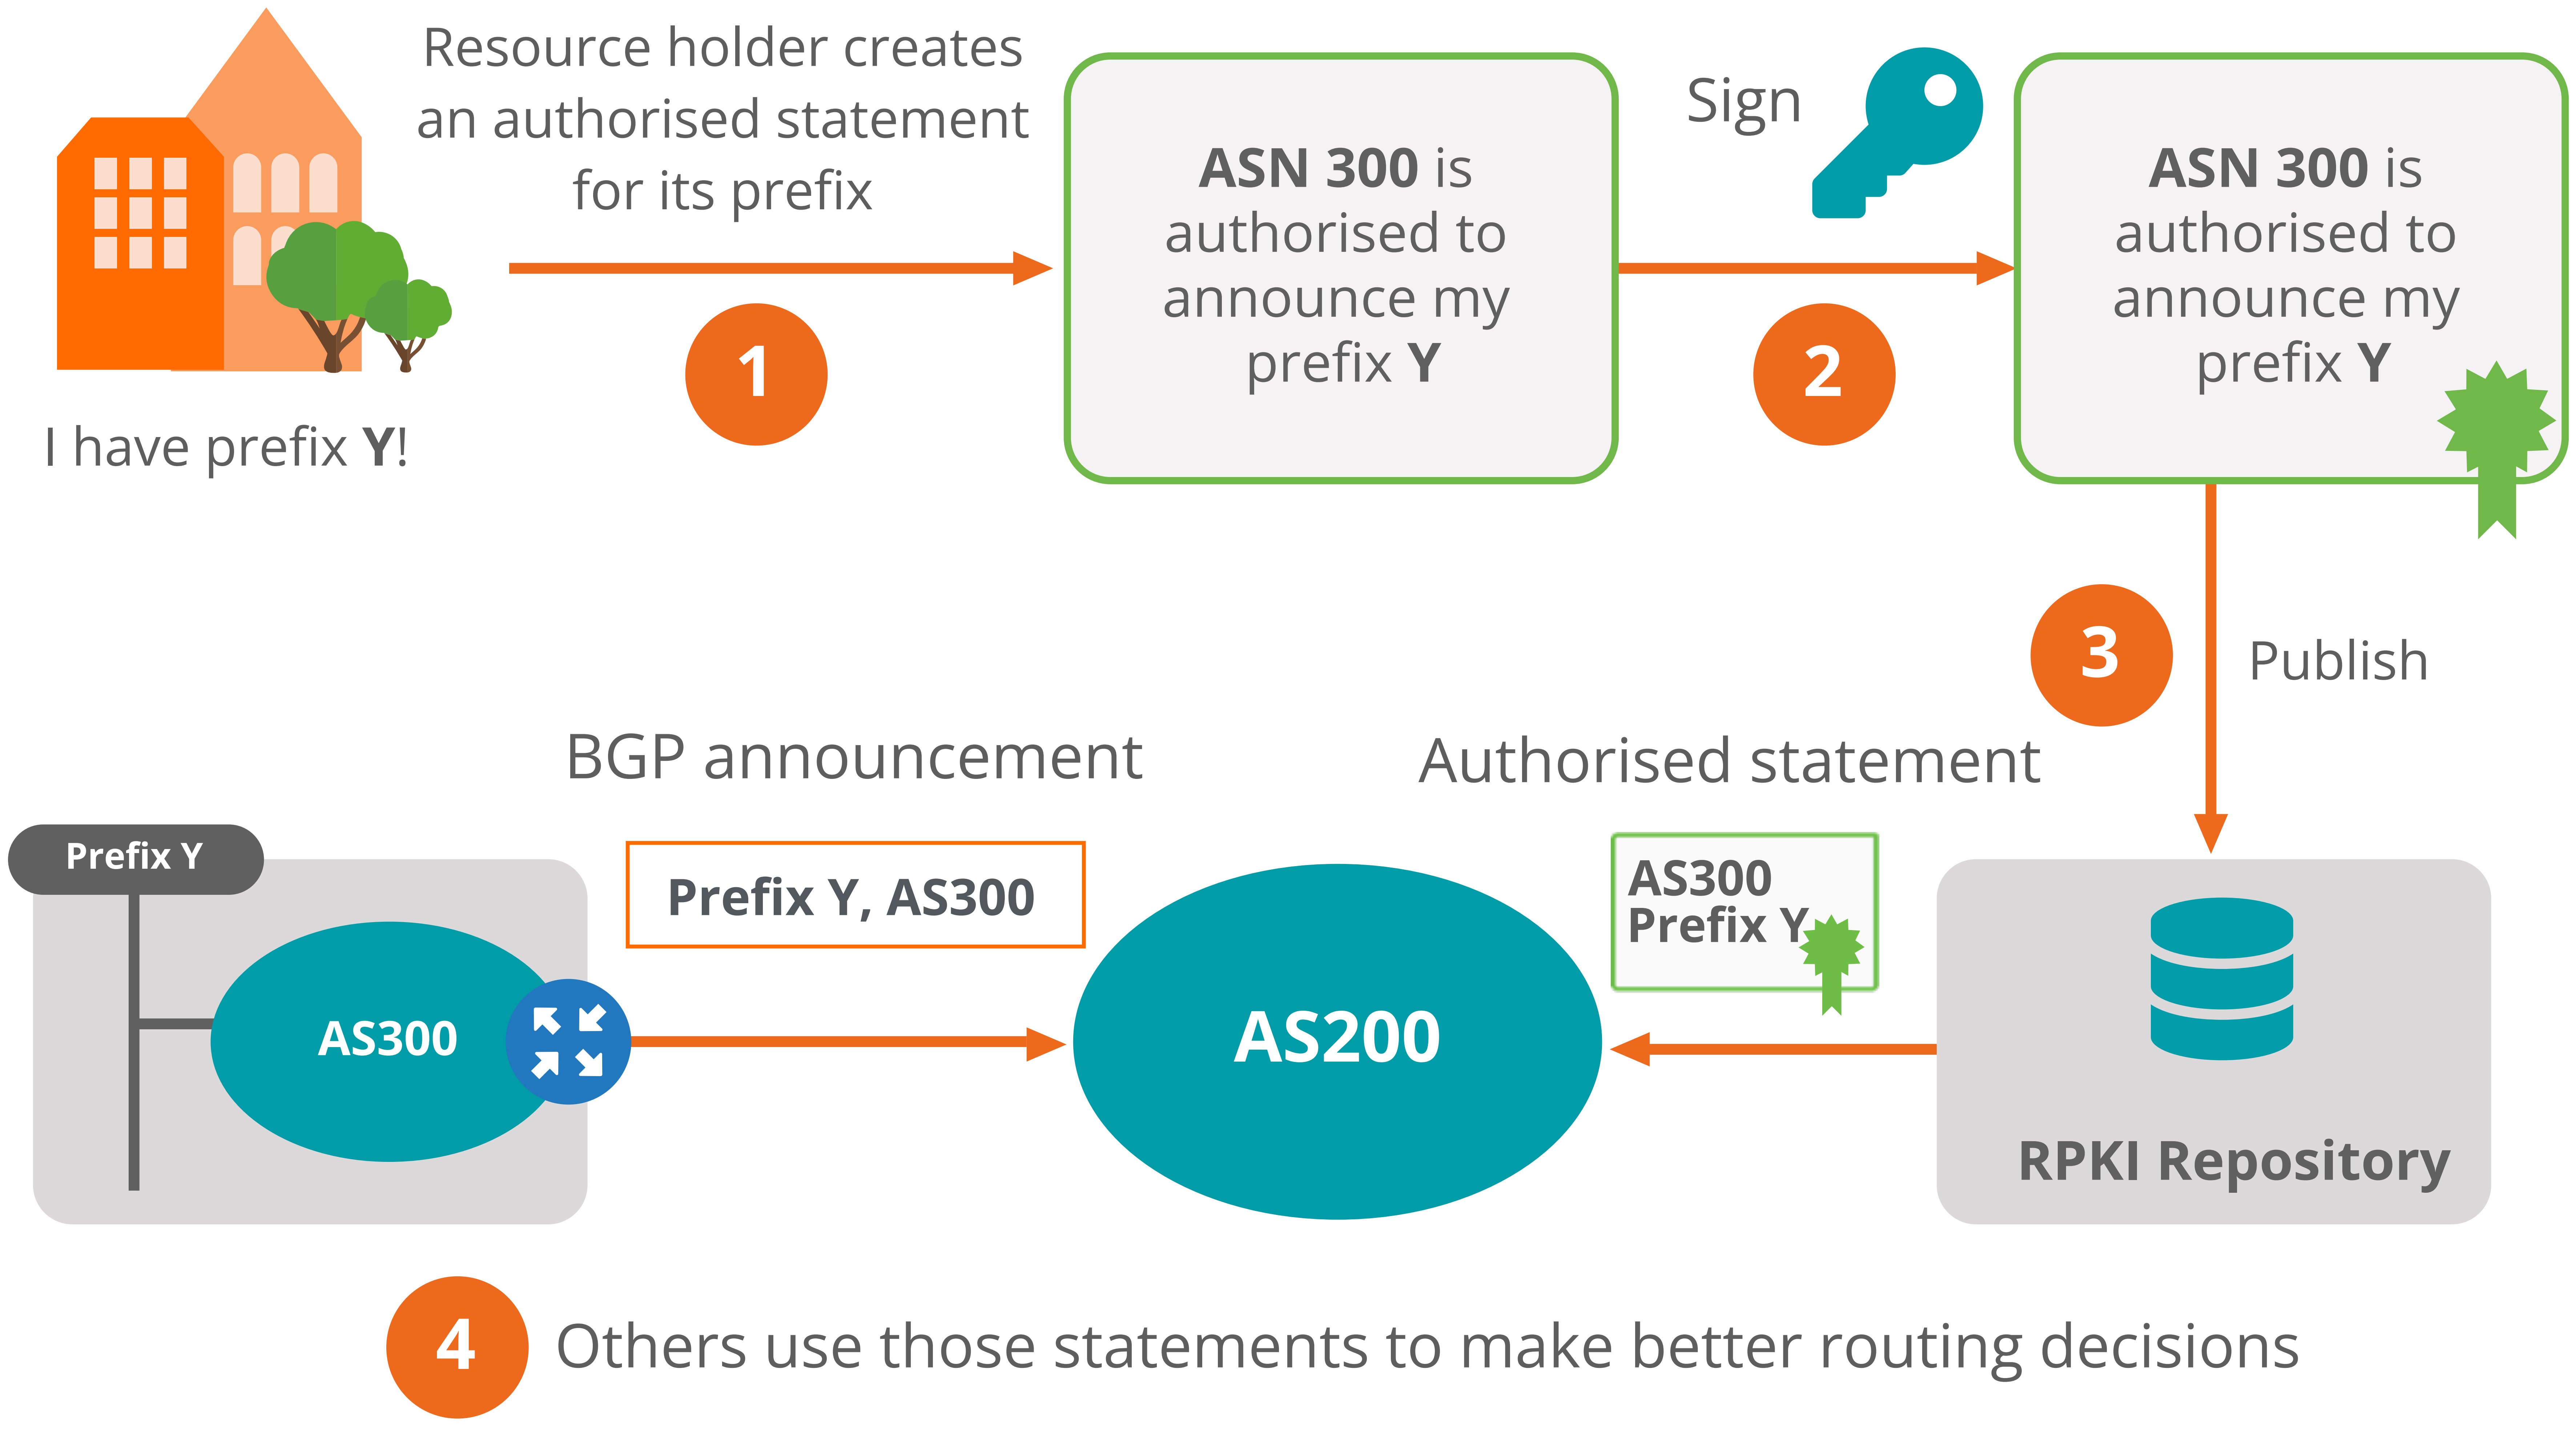

RPSL: Imperfect. RPKI: Incomplete.

- **The best practice in configuring BGP is to secure it by generating router configuration from RPLS policy retrieved over unauthenticated channels.**
- **Multiple IRR databases can contain objects for the same resource**
- **Many networks do not configure this kind of policy**
  - A provider then _adds_ this in another database. Problem solved.
- **RPKI is under development, and is not yet a replacement for the IRR system**
- **RPKI can improve the data quality in the IRR**

BECHA@ripe.net
tdekock@ripe.net

Bonus Slides

RPKI

BGP Origin Hijacks

- **An AS originates a prefix that is not authorised to originate**
- **Hijacker impersonates the legitimate holder**
  - May hijack an **allocated** or **unallocated** address space
- **It may announce the exact same prefix or more specifics**
  - Prefix Hijack
  - **Sub-prefix Hijack** (De-aggregation hijack or subnet attack)

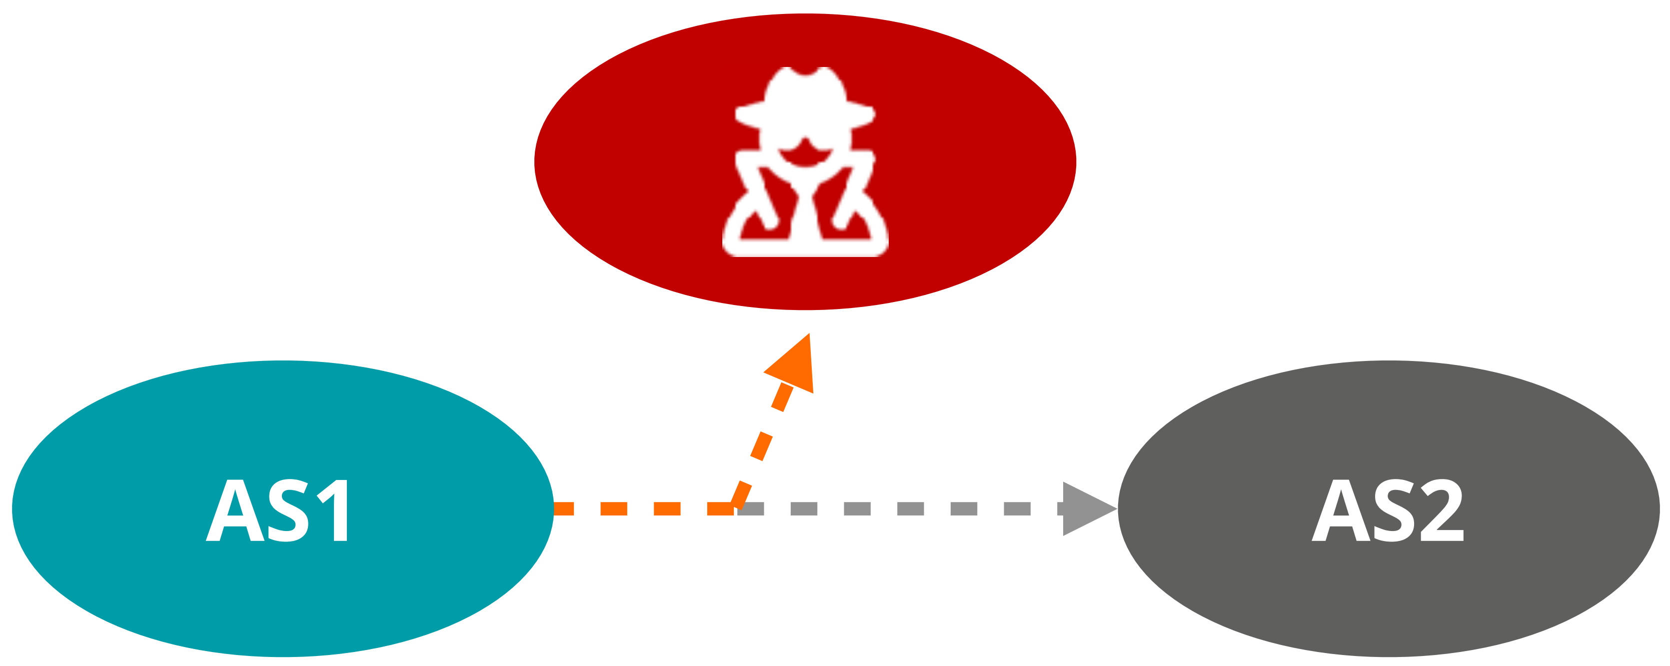

Prefix Hijack

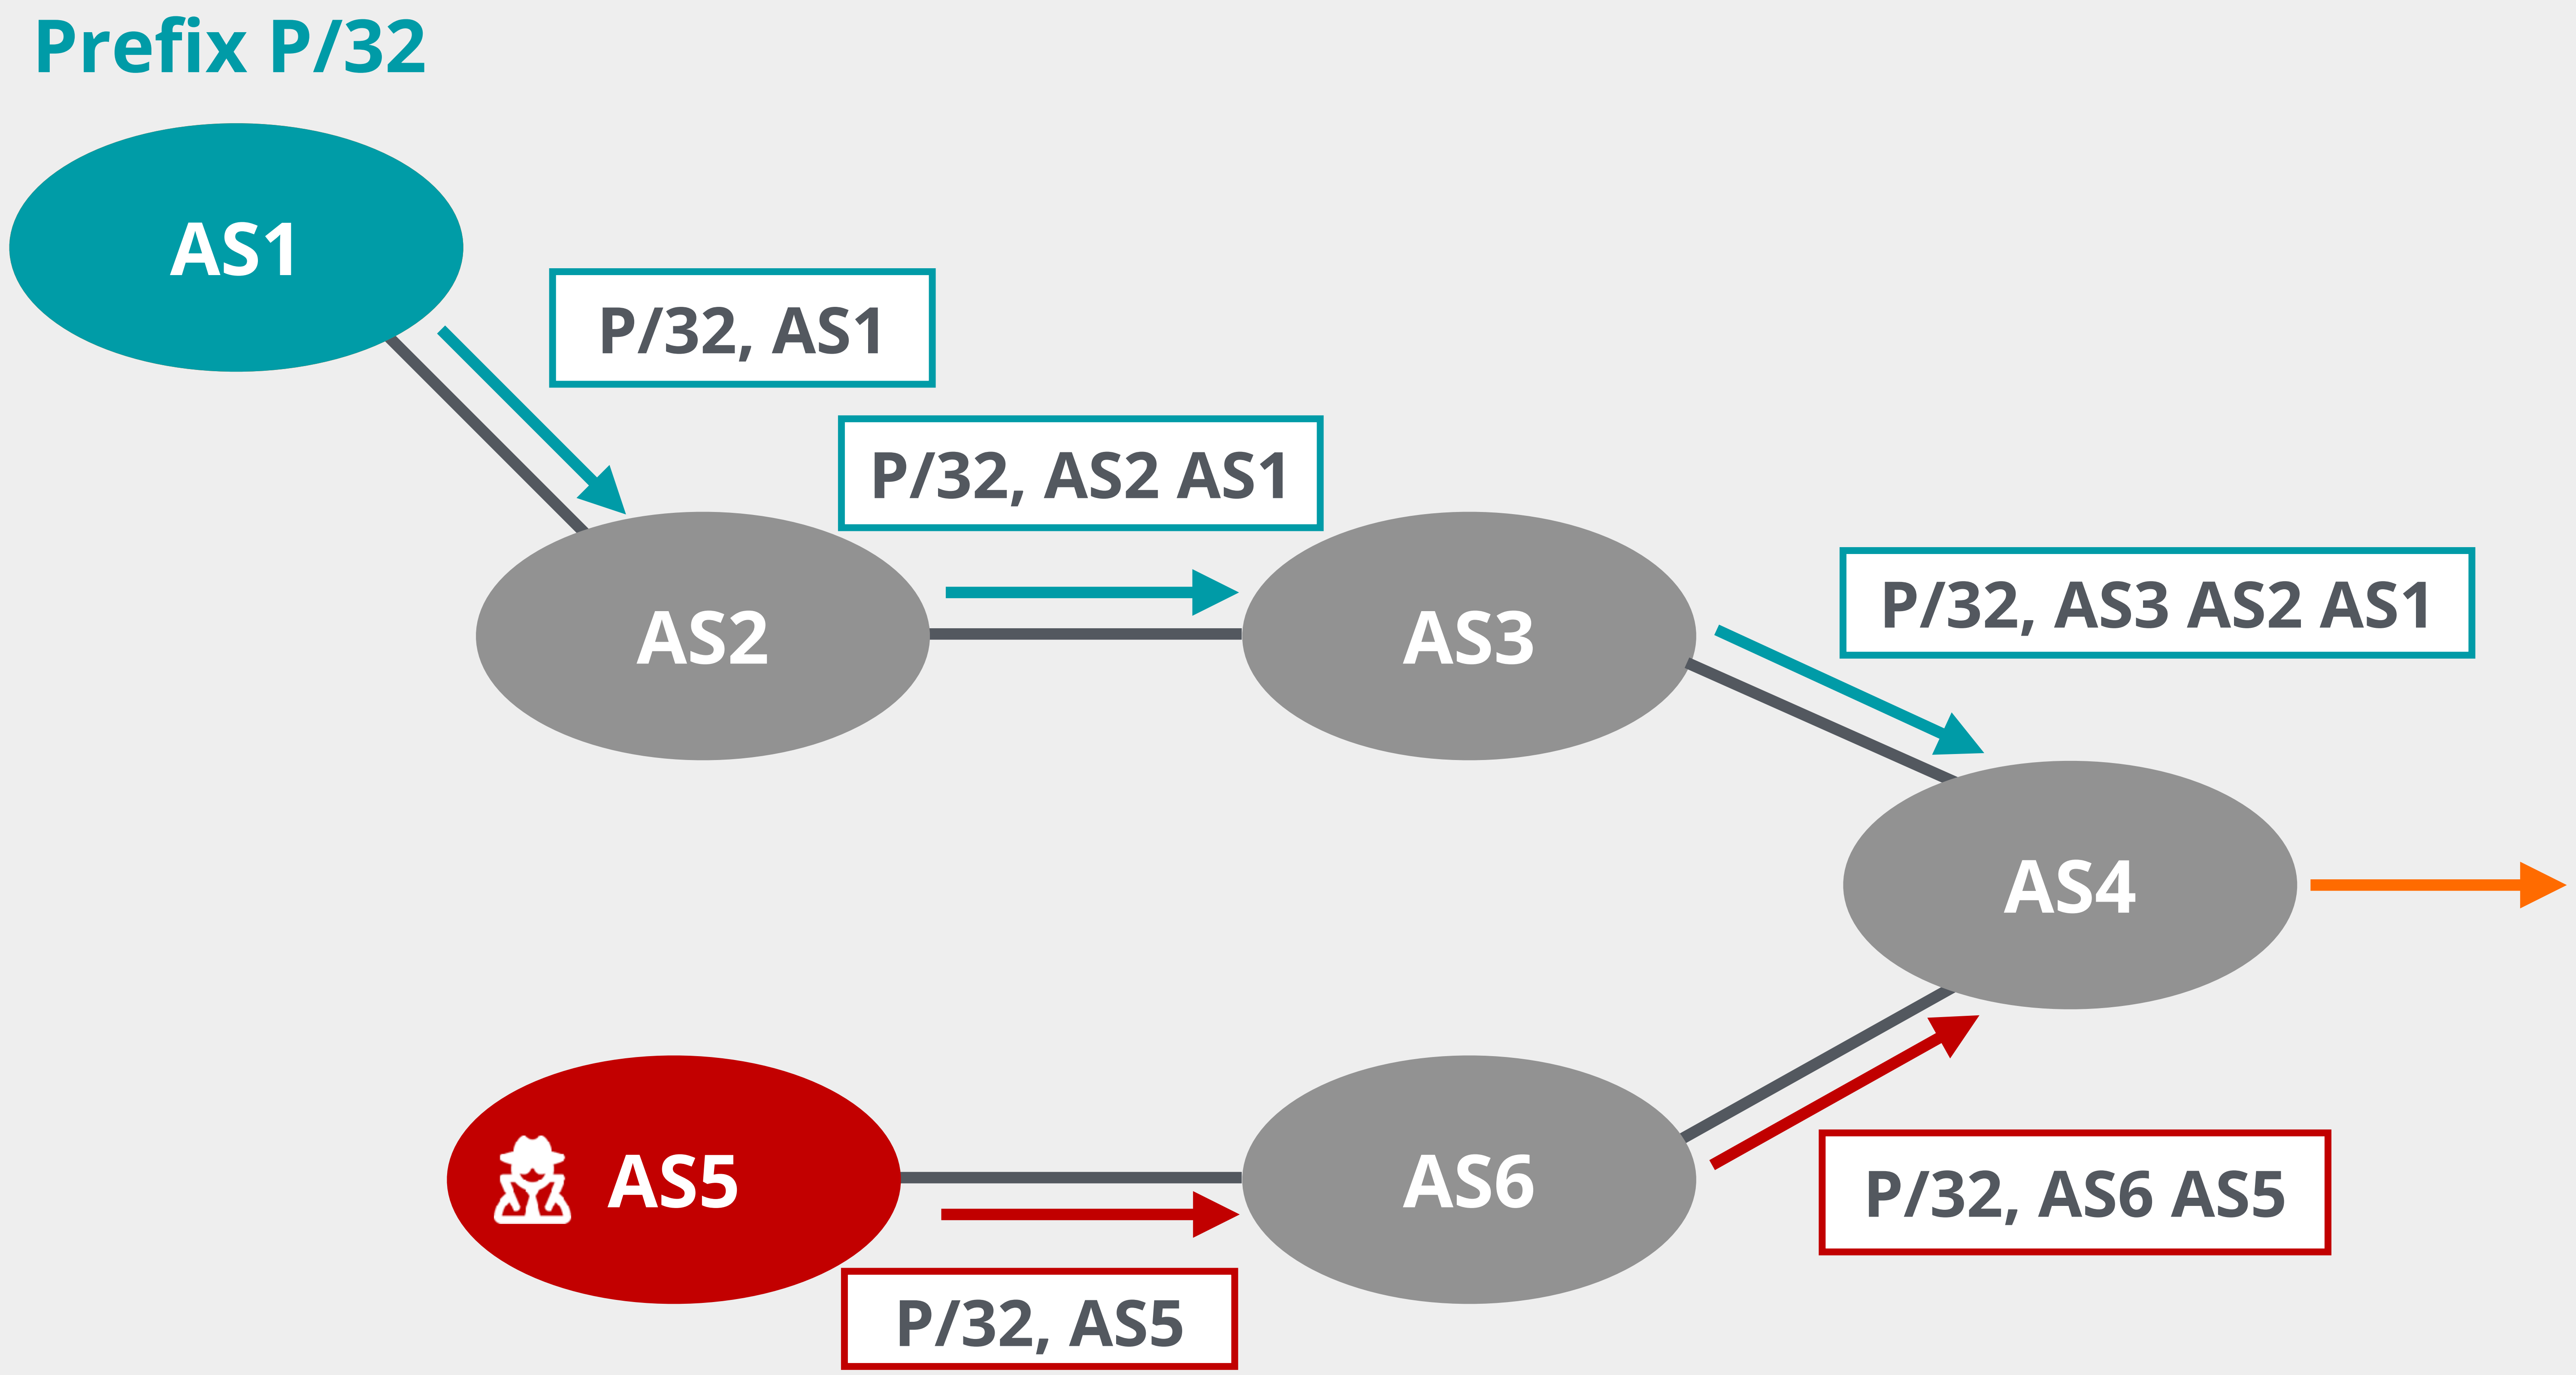

This is a **local hijack!**

Only some networks are affected based on BGP path selection process

Sub-prefix Hijack (Subnet Attack)

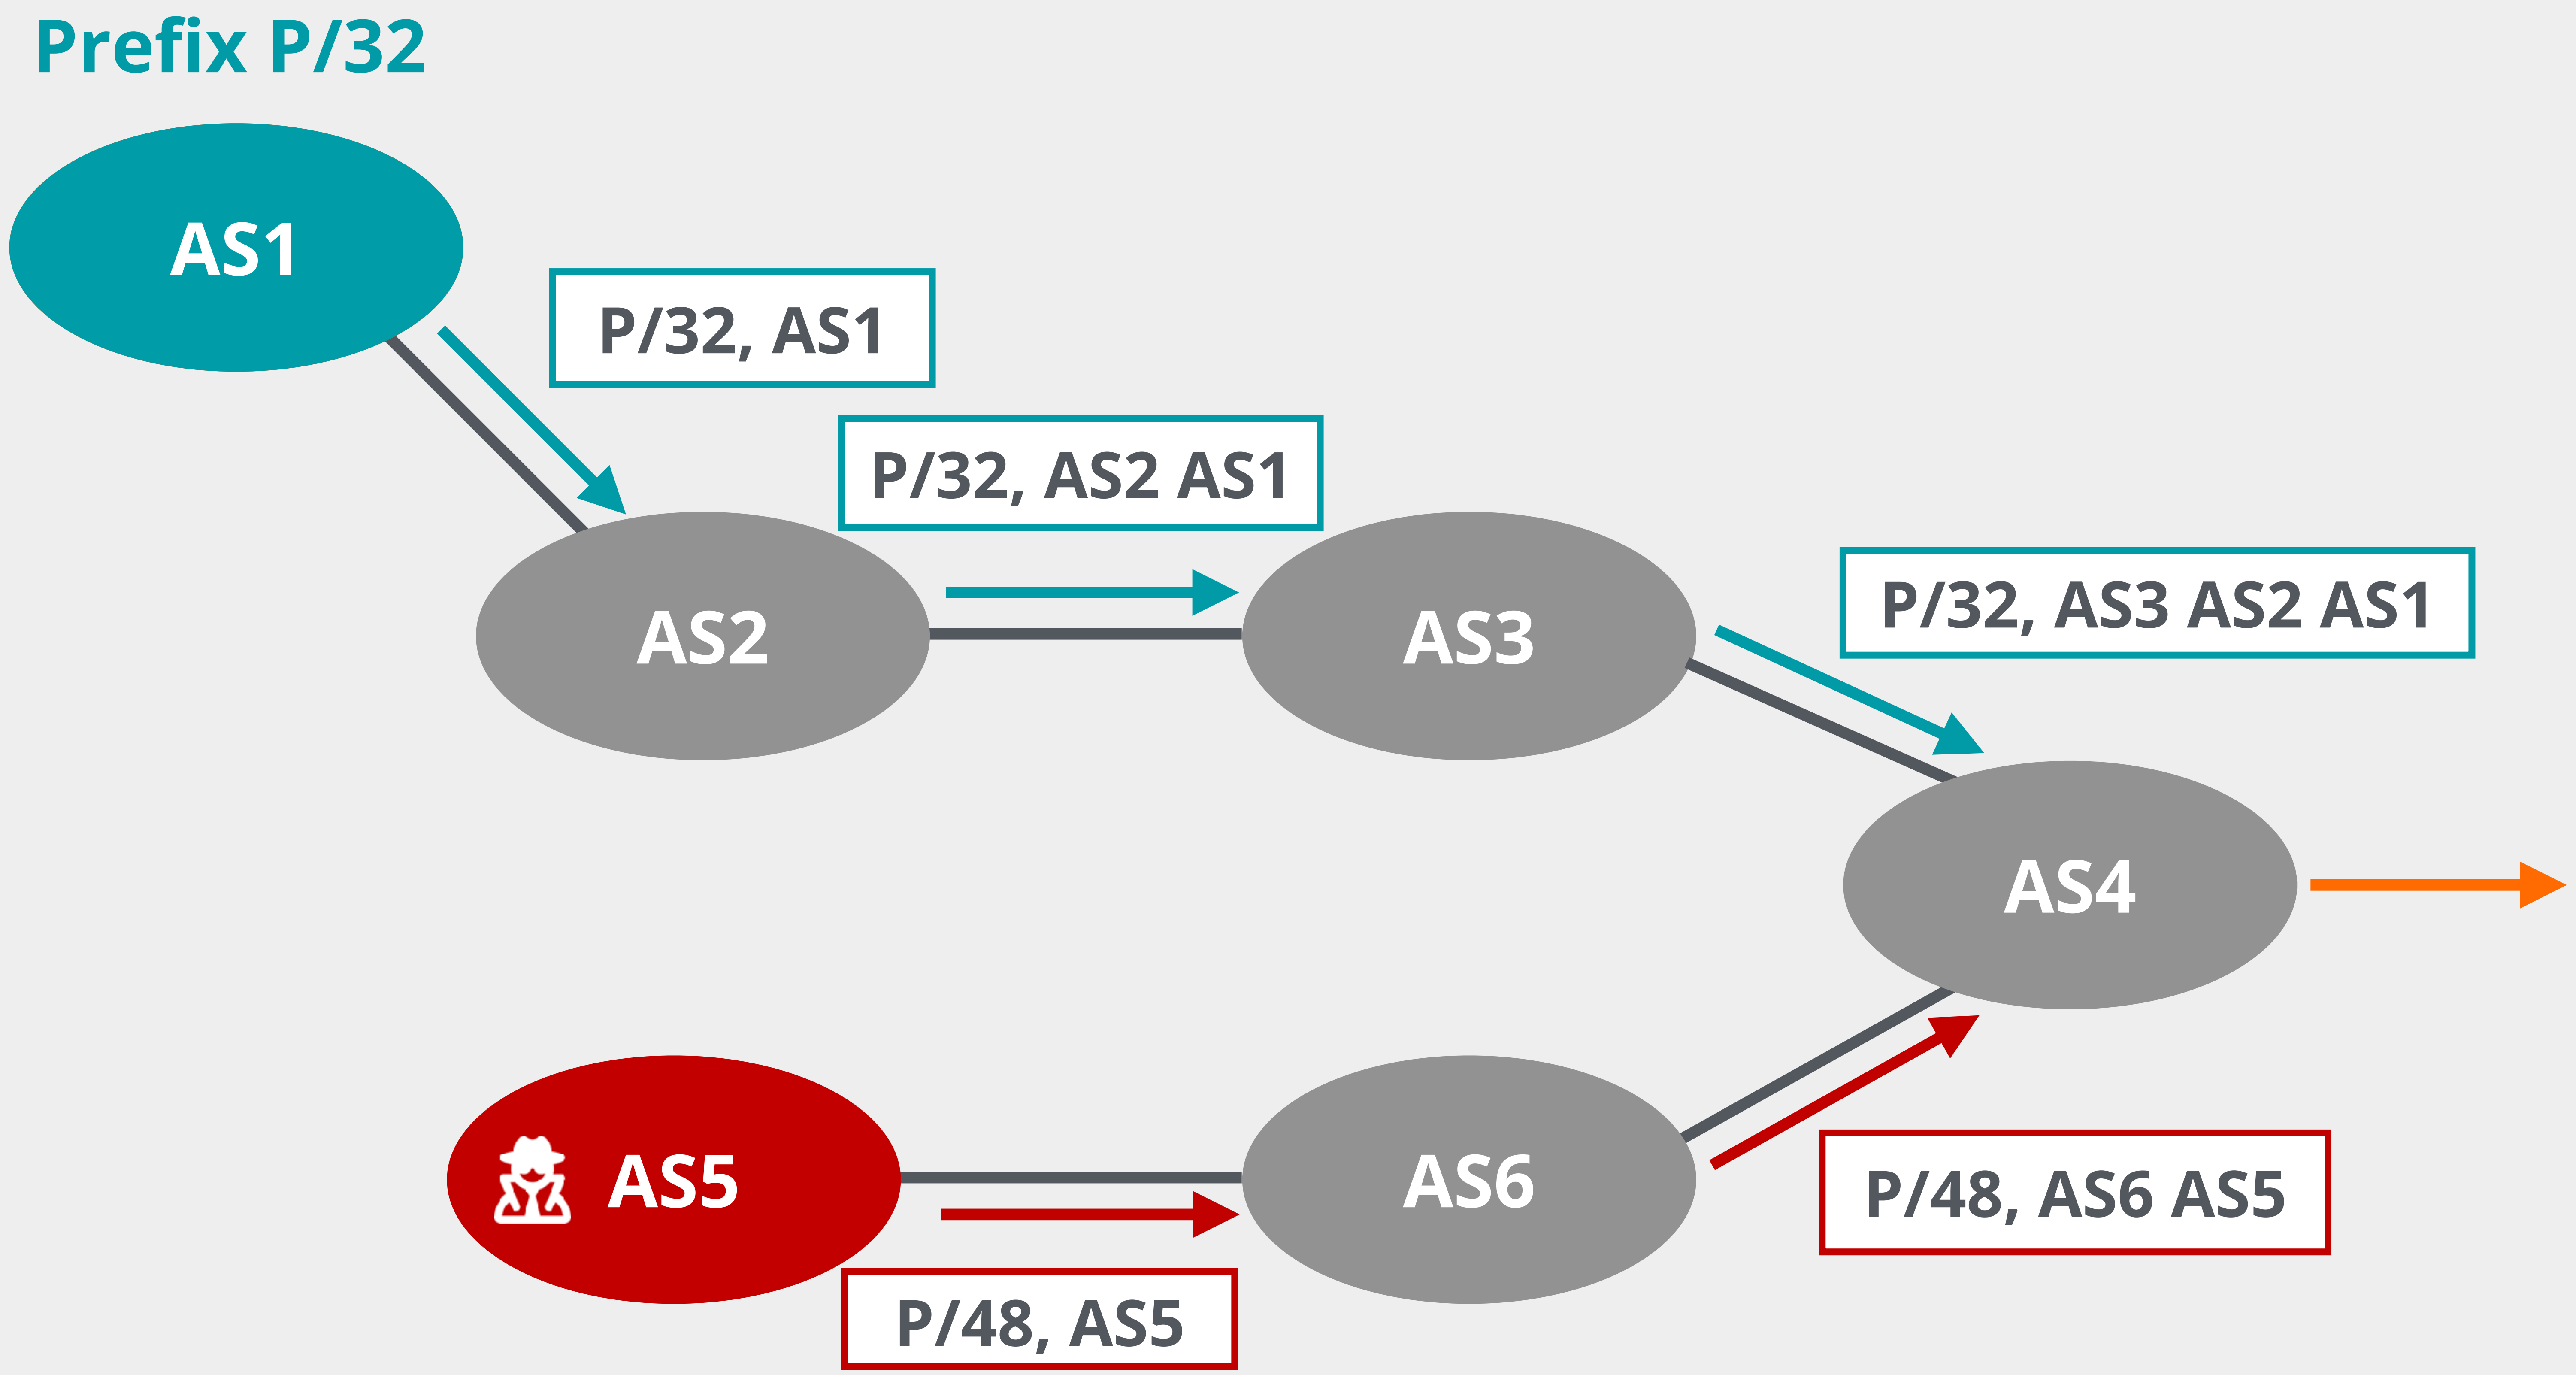

This is a **global hijack!**

All traffic for more specific prefix will be forwarded to the hijacker's network

April 2018: Amazon MyEtherWallet

- **BGP hijack of Amazon DNS**
- **What happened?**
- **Why?**
  - Attack to steal cryptocurrency

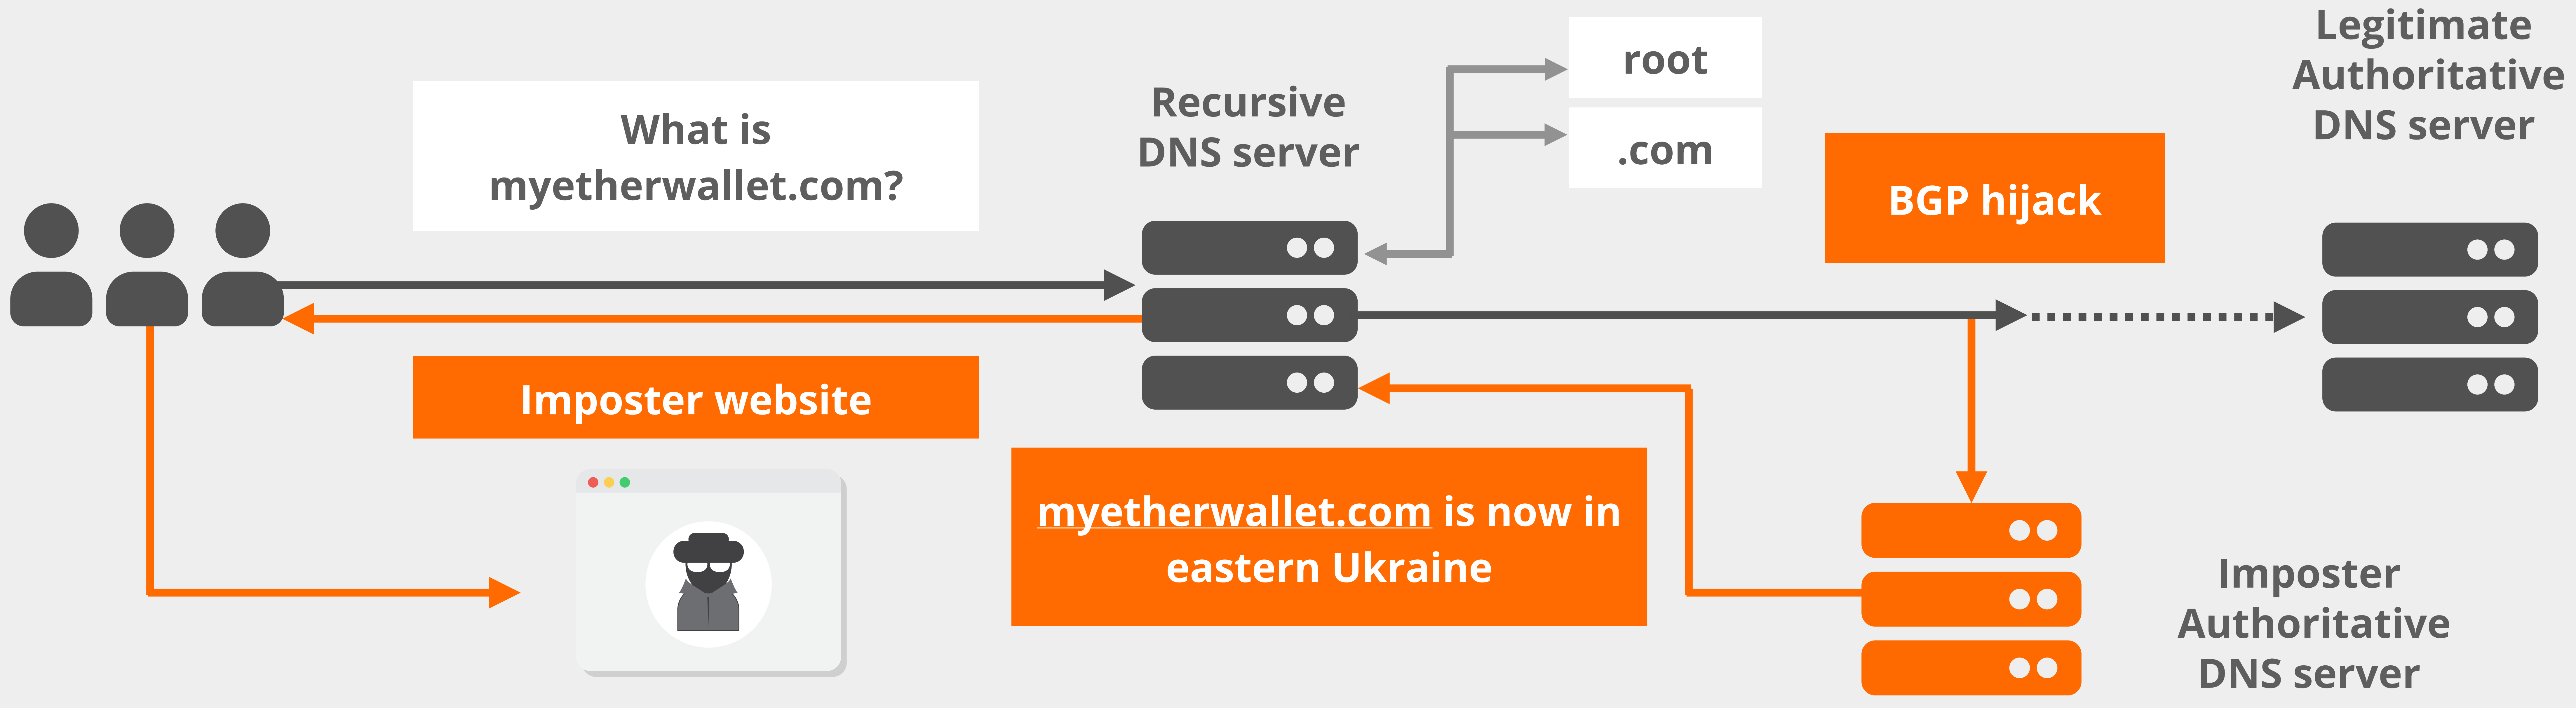

What is RPKI?

- A security framework for the Internet
- Verifies the association between resource holders and their resources
  - Attaches digital certificate to IP addresses and AS numbers
- Used to validate the origin of BGP announcements (BGP OV)
  - Is the originating ASN authorised to originate a particular prefix?
  - Helps to mitigate **BGP Origin Hijacks** and **Route leaks**

How Does RPKI Work?

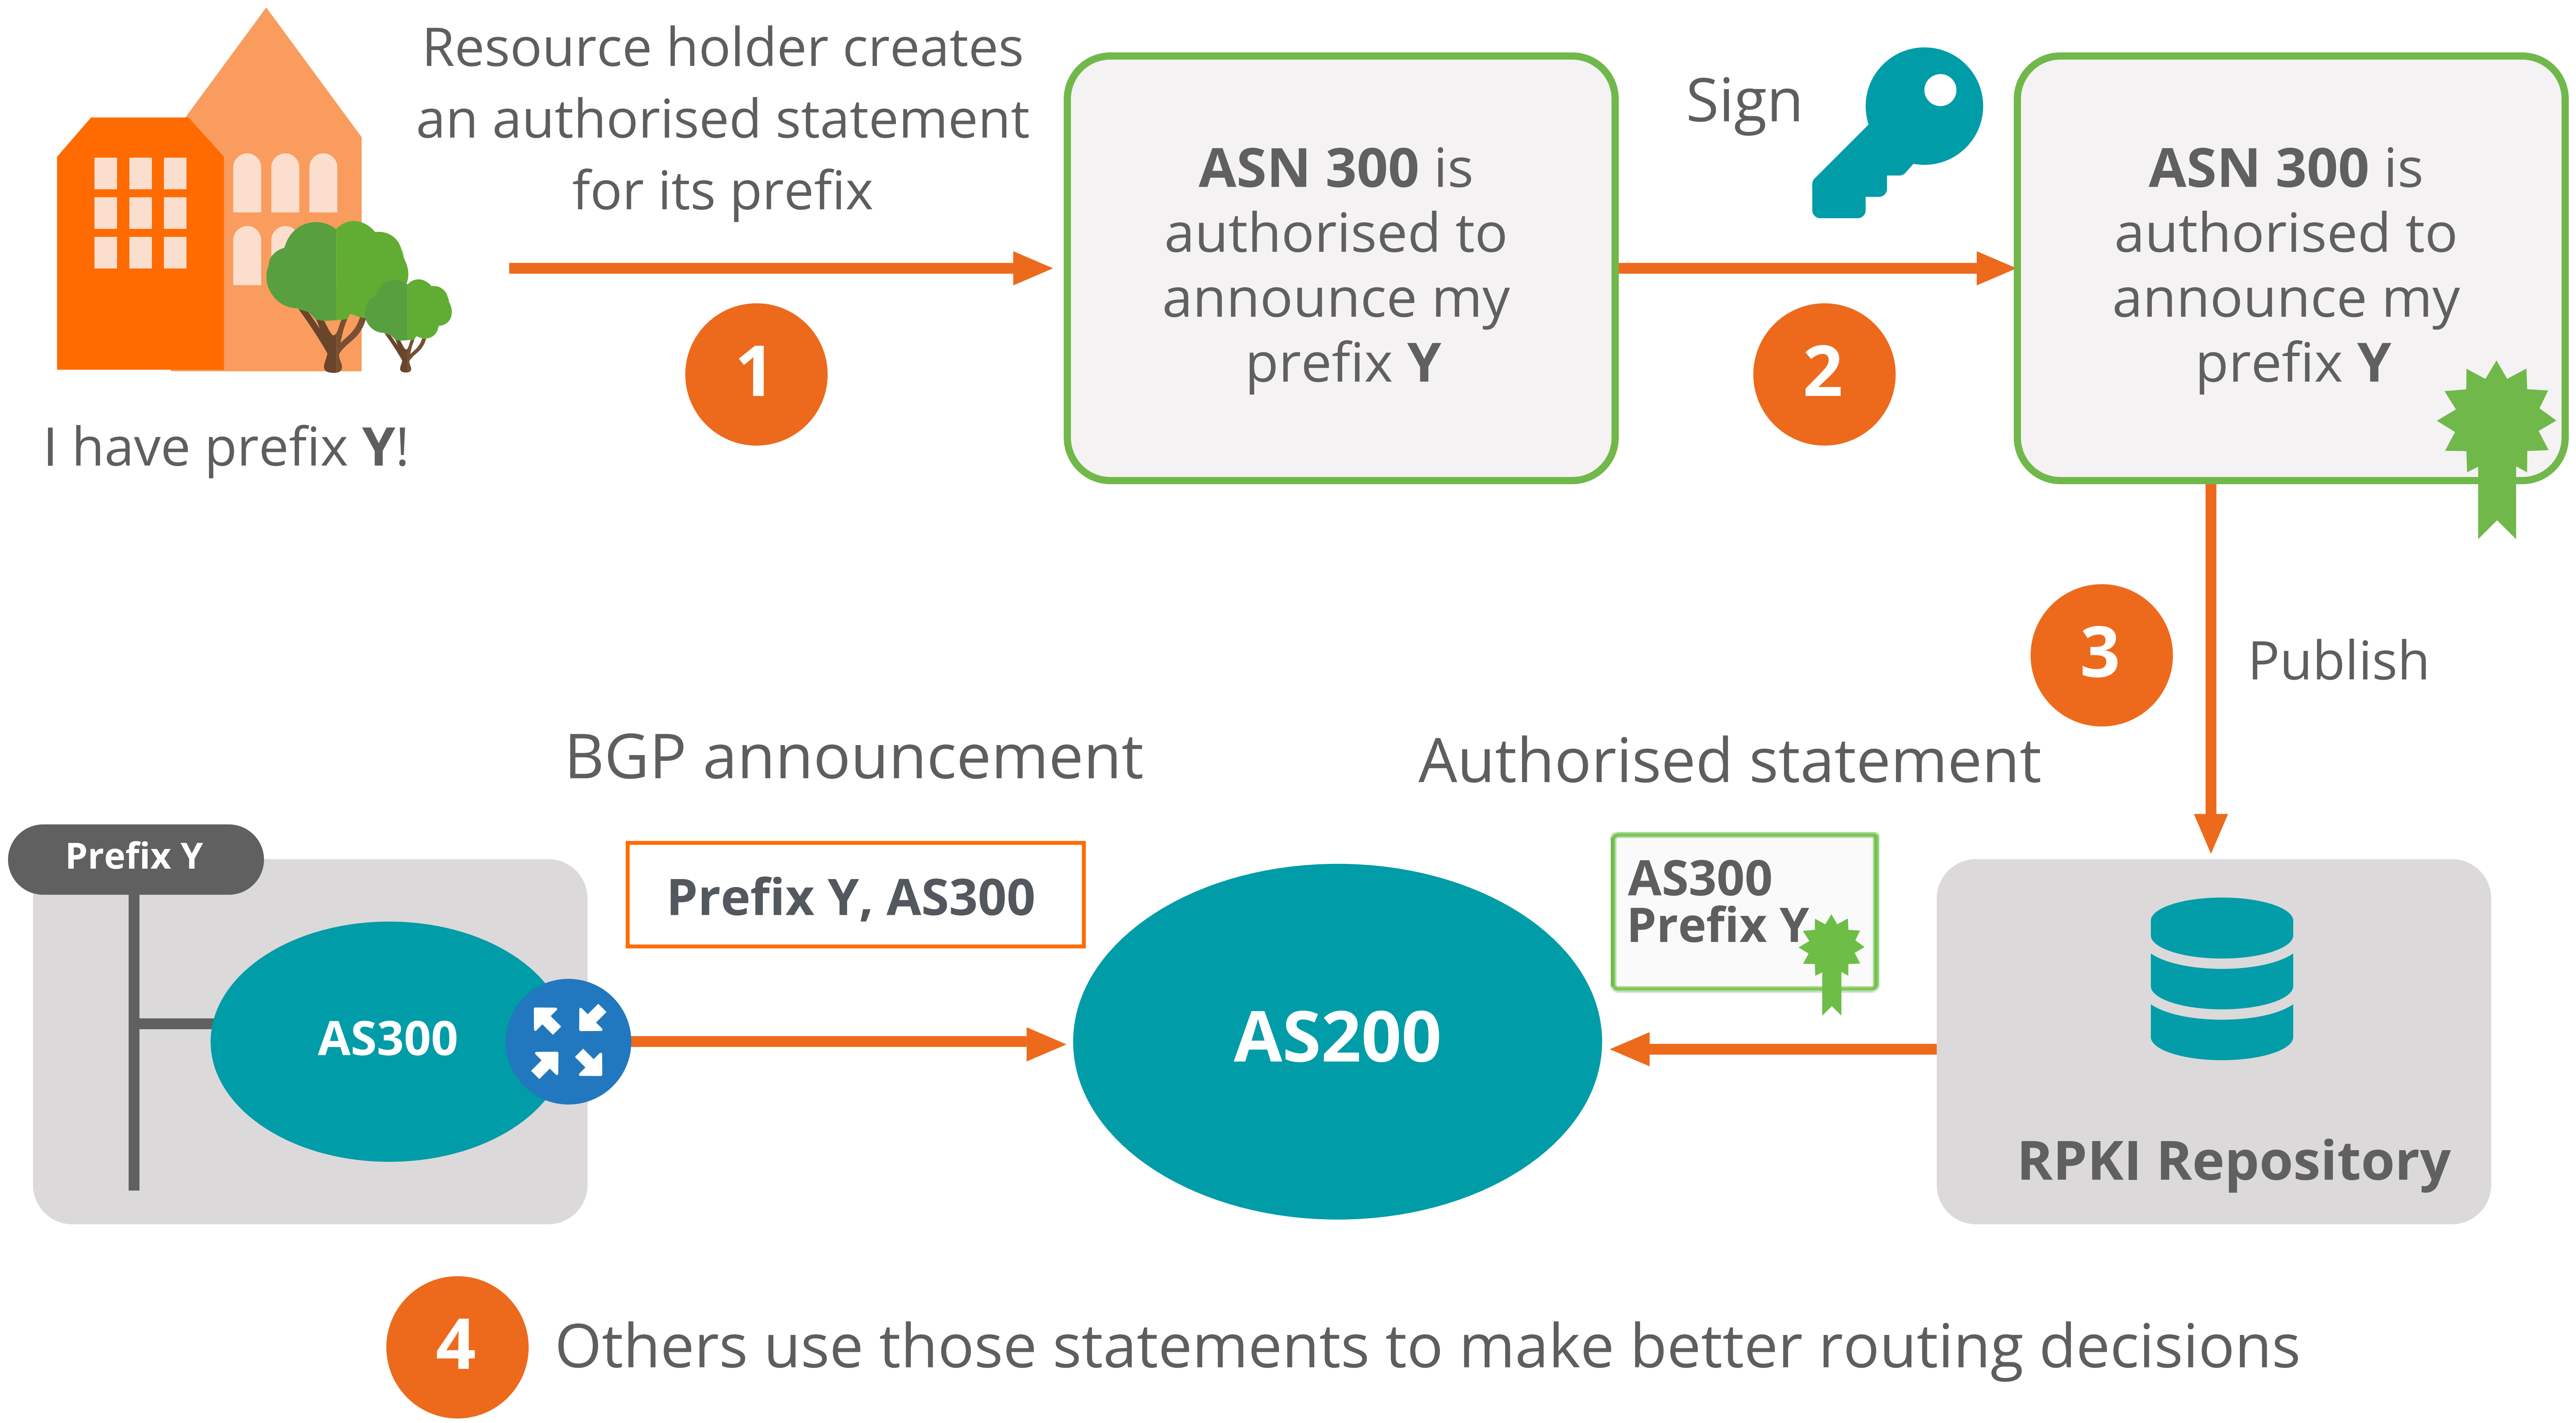

Trust in RPKI

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

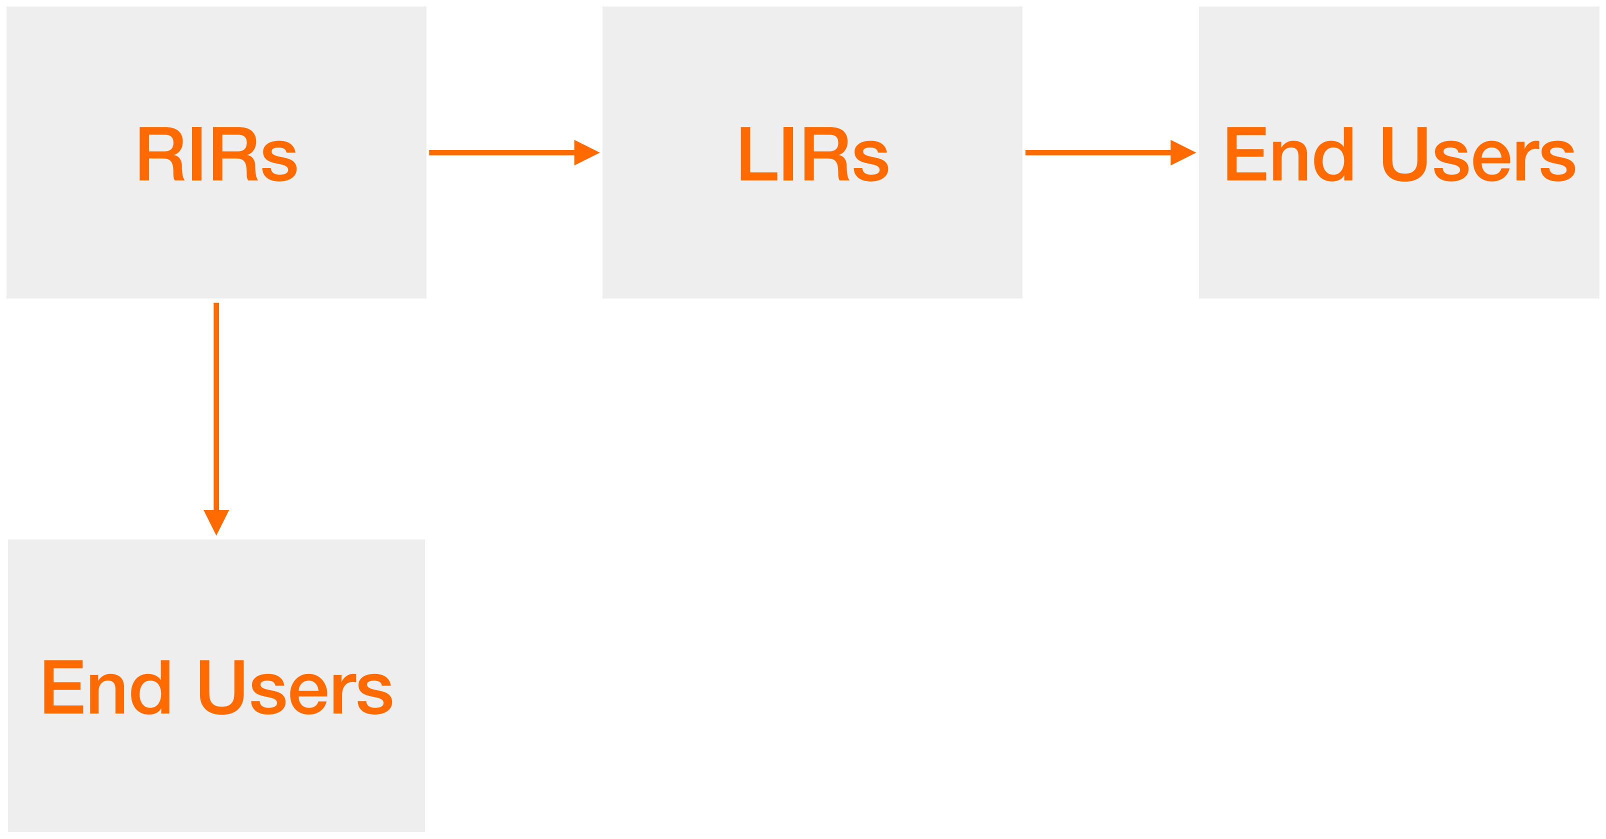

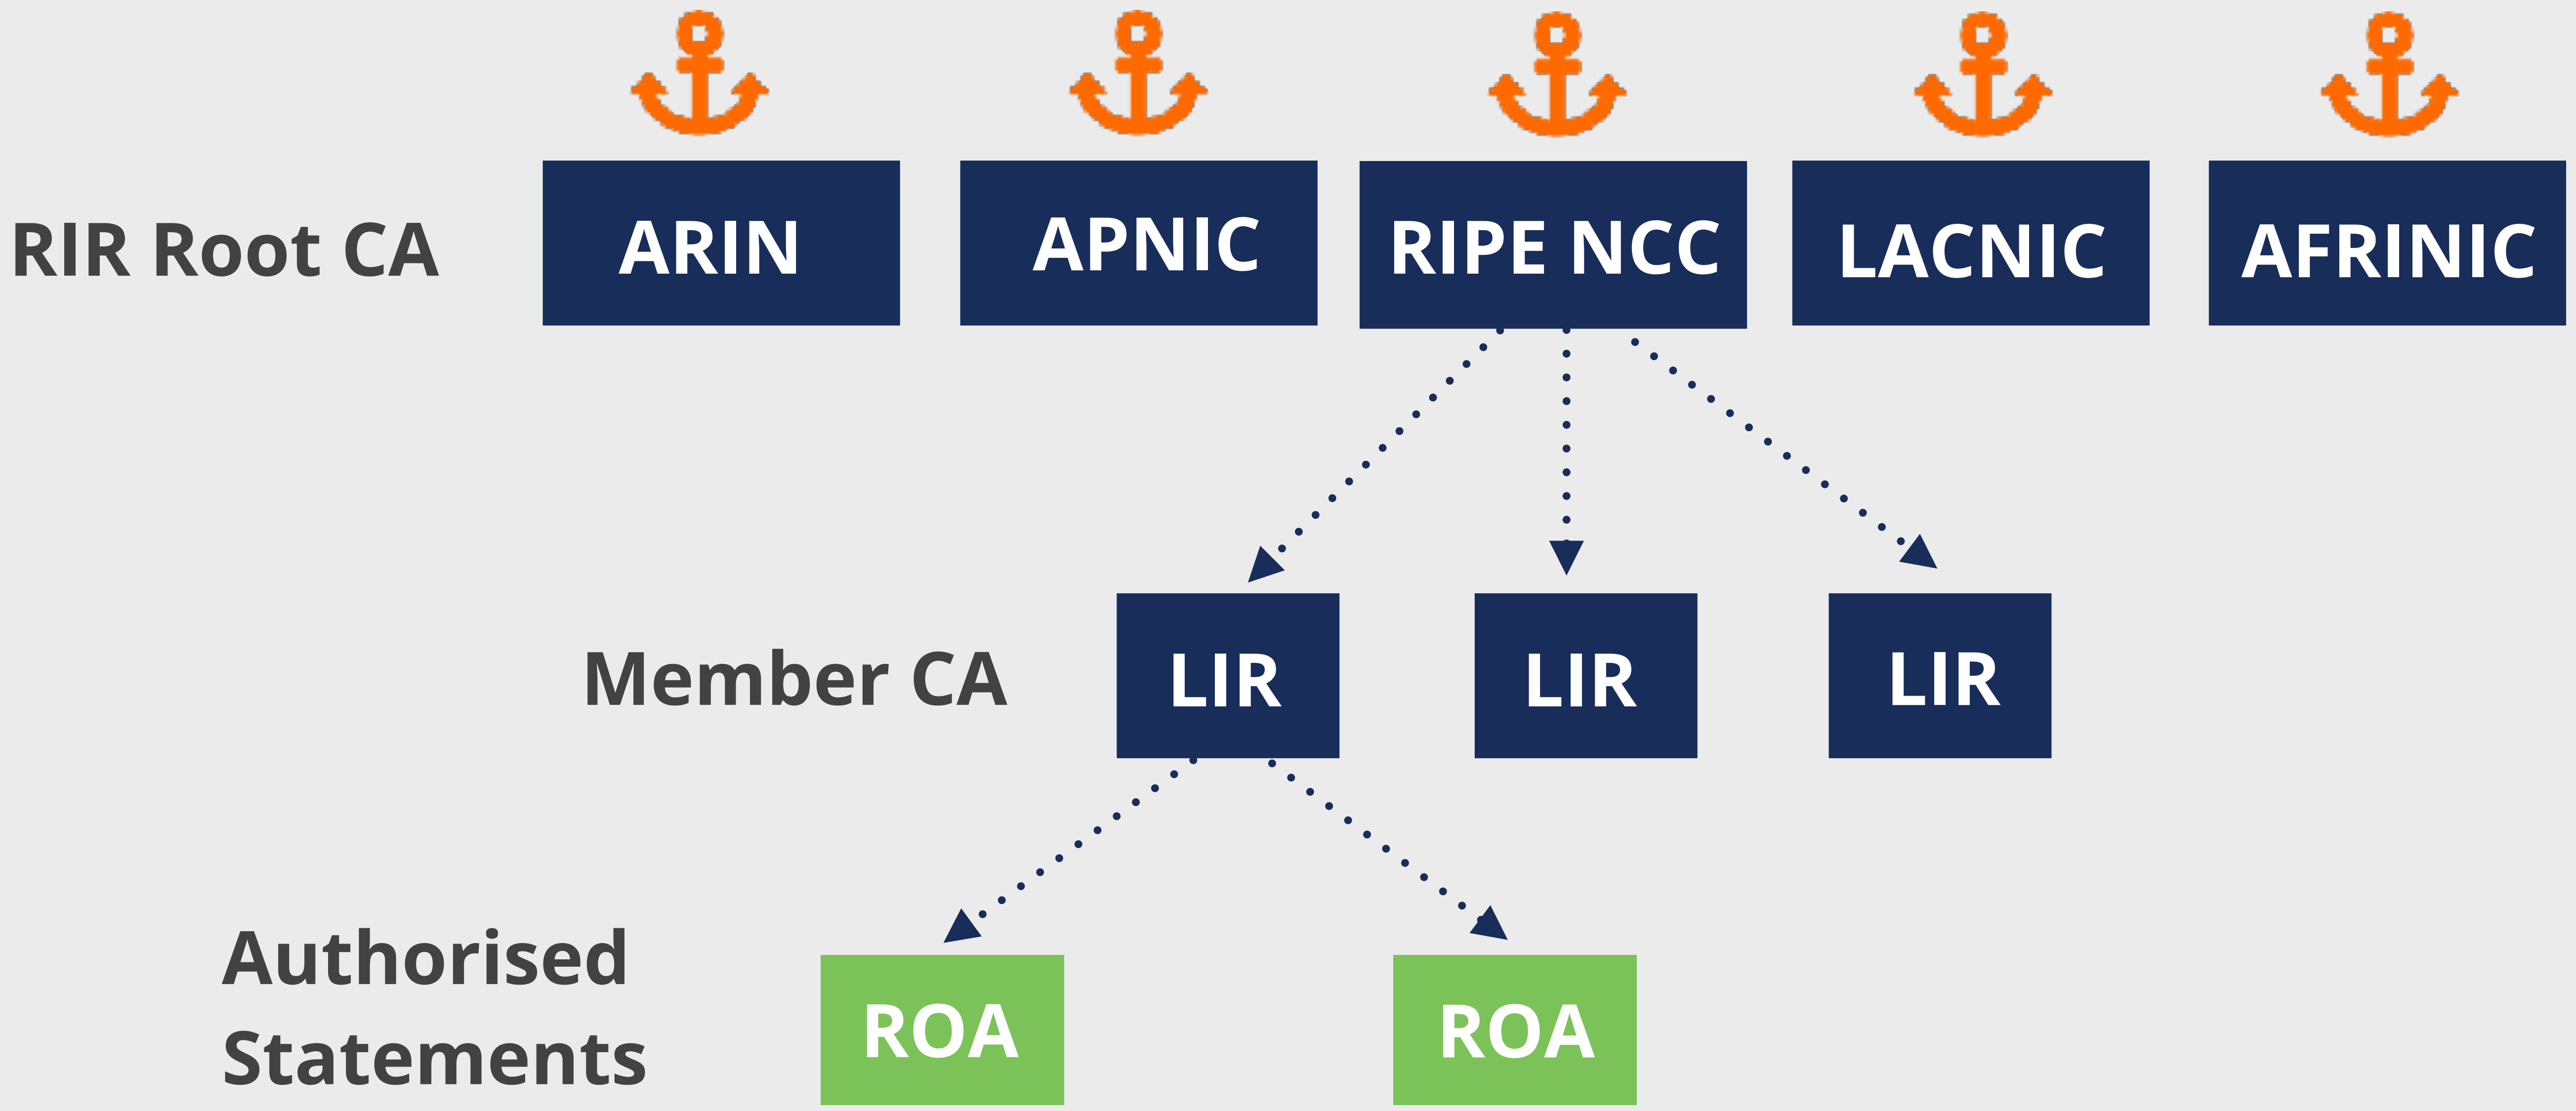

RPKI Chain of Trust

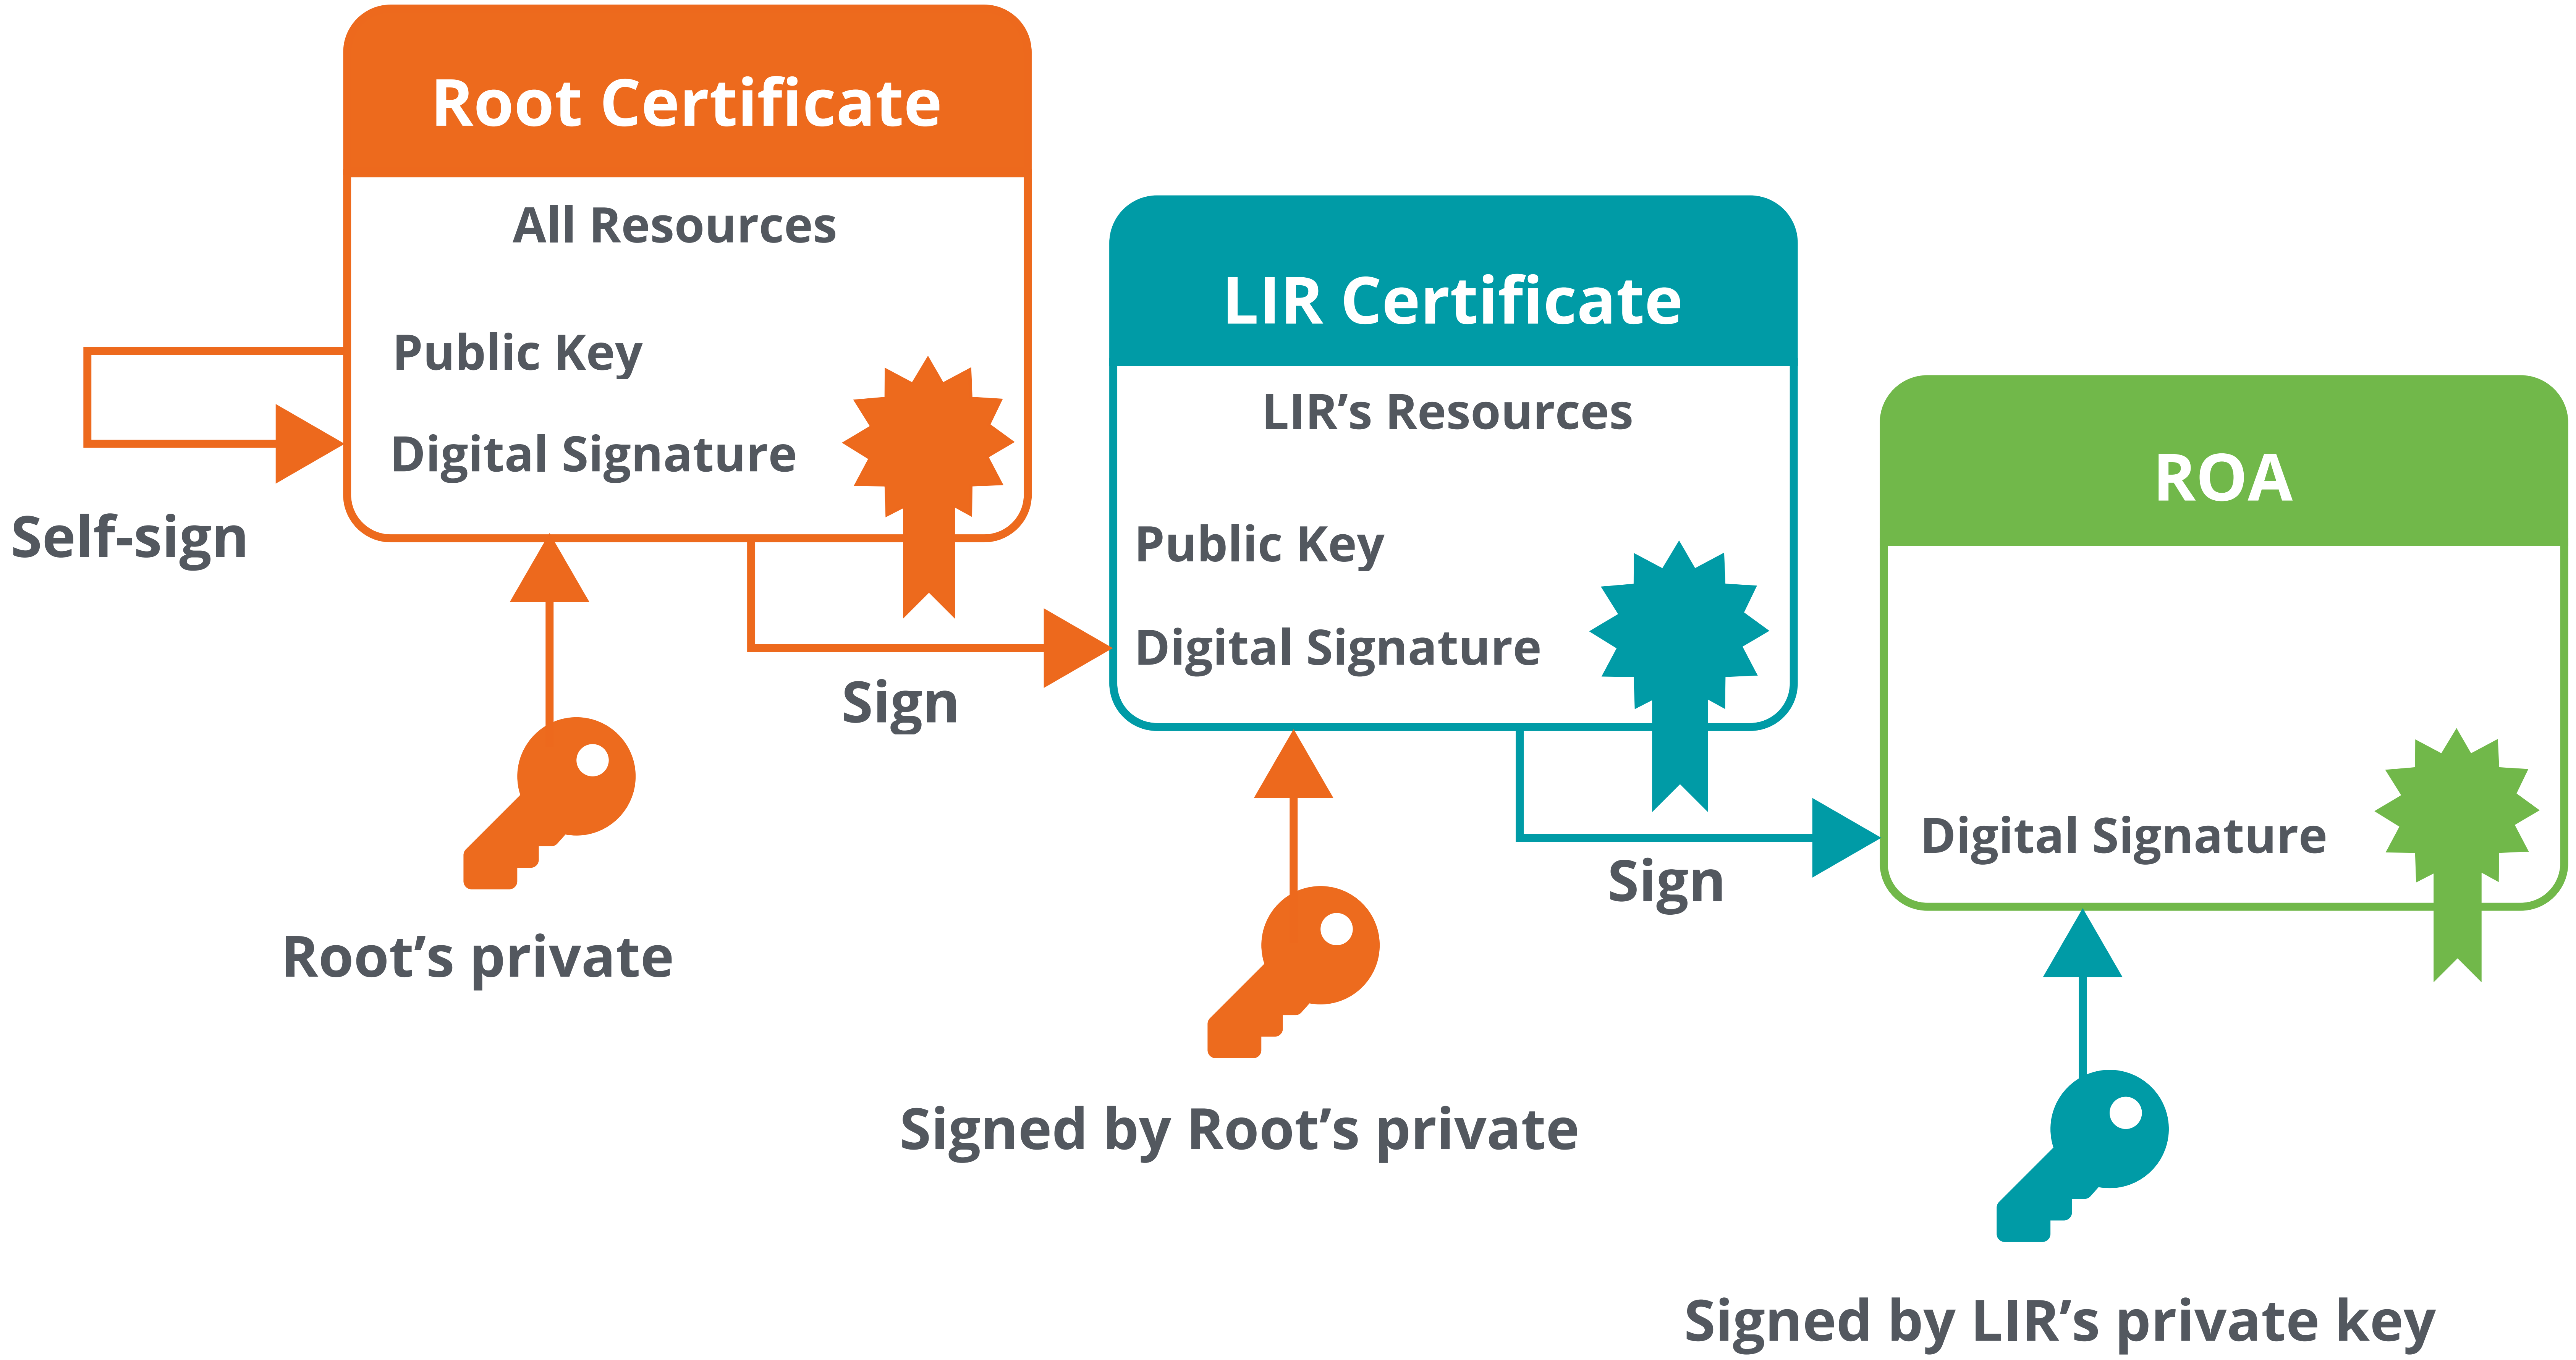

Elements of RPKI

- The RPKI system consists of two parts:

SIGNING		VALIDATION
Create ROAs for your prefixes in the RPKI system	+	Verify the information provided by others

Hosted RPKI

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

RIPE NCC Hosted System

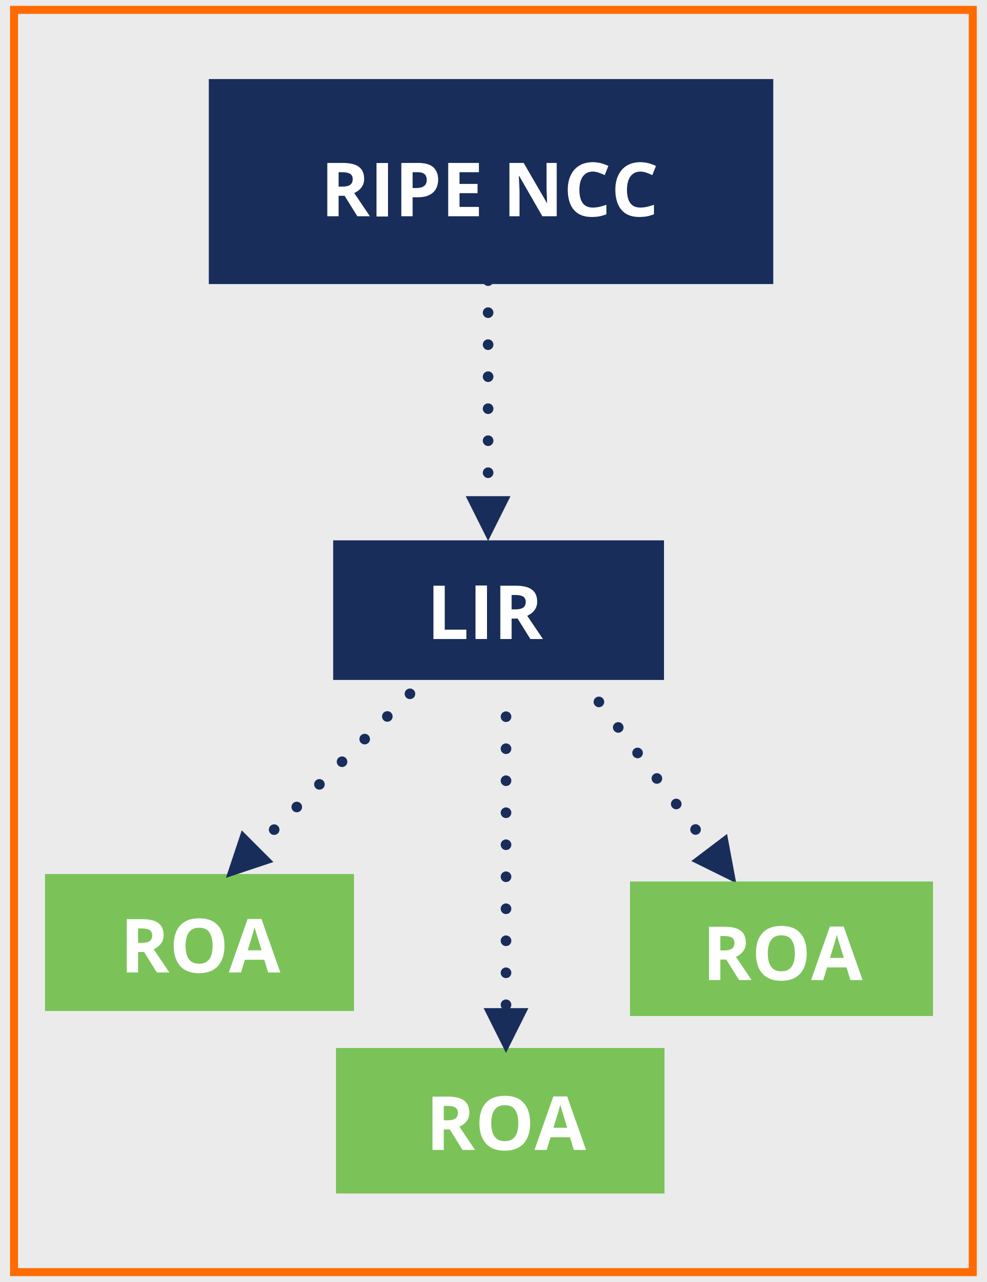

Delegated RPKI

- **Each LIR manages its part of the RPKI system:**
  - Runs its own CA as a child of the RIR
  - Manages keys/key rollovers
  - Creates, signs and publishes ROAs

- **Certificate Authority (CA) Software**
  - **Krill** (NLnet Labs)
  - **rpkid** (Dragon Research Labs)

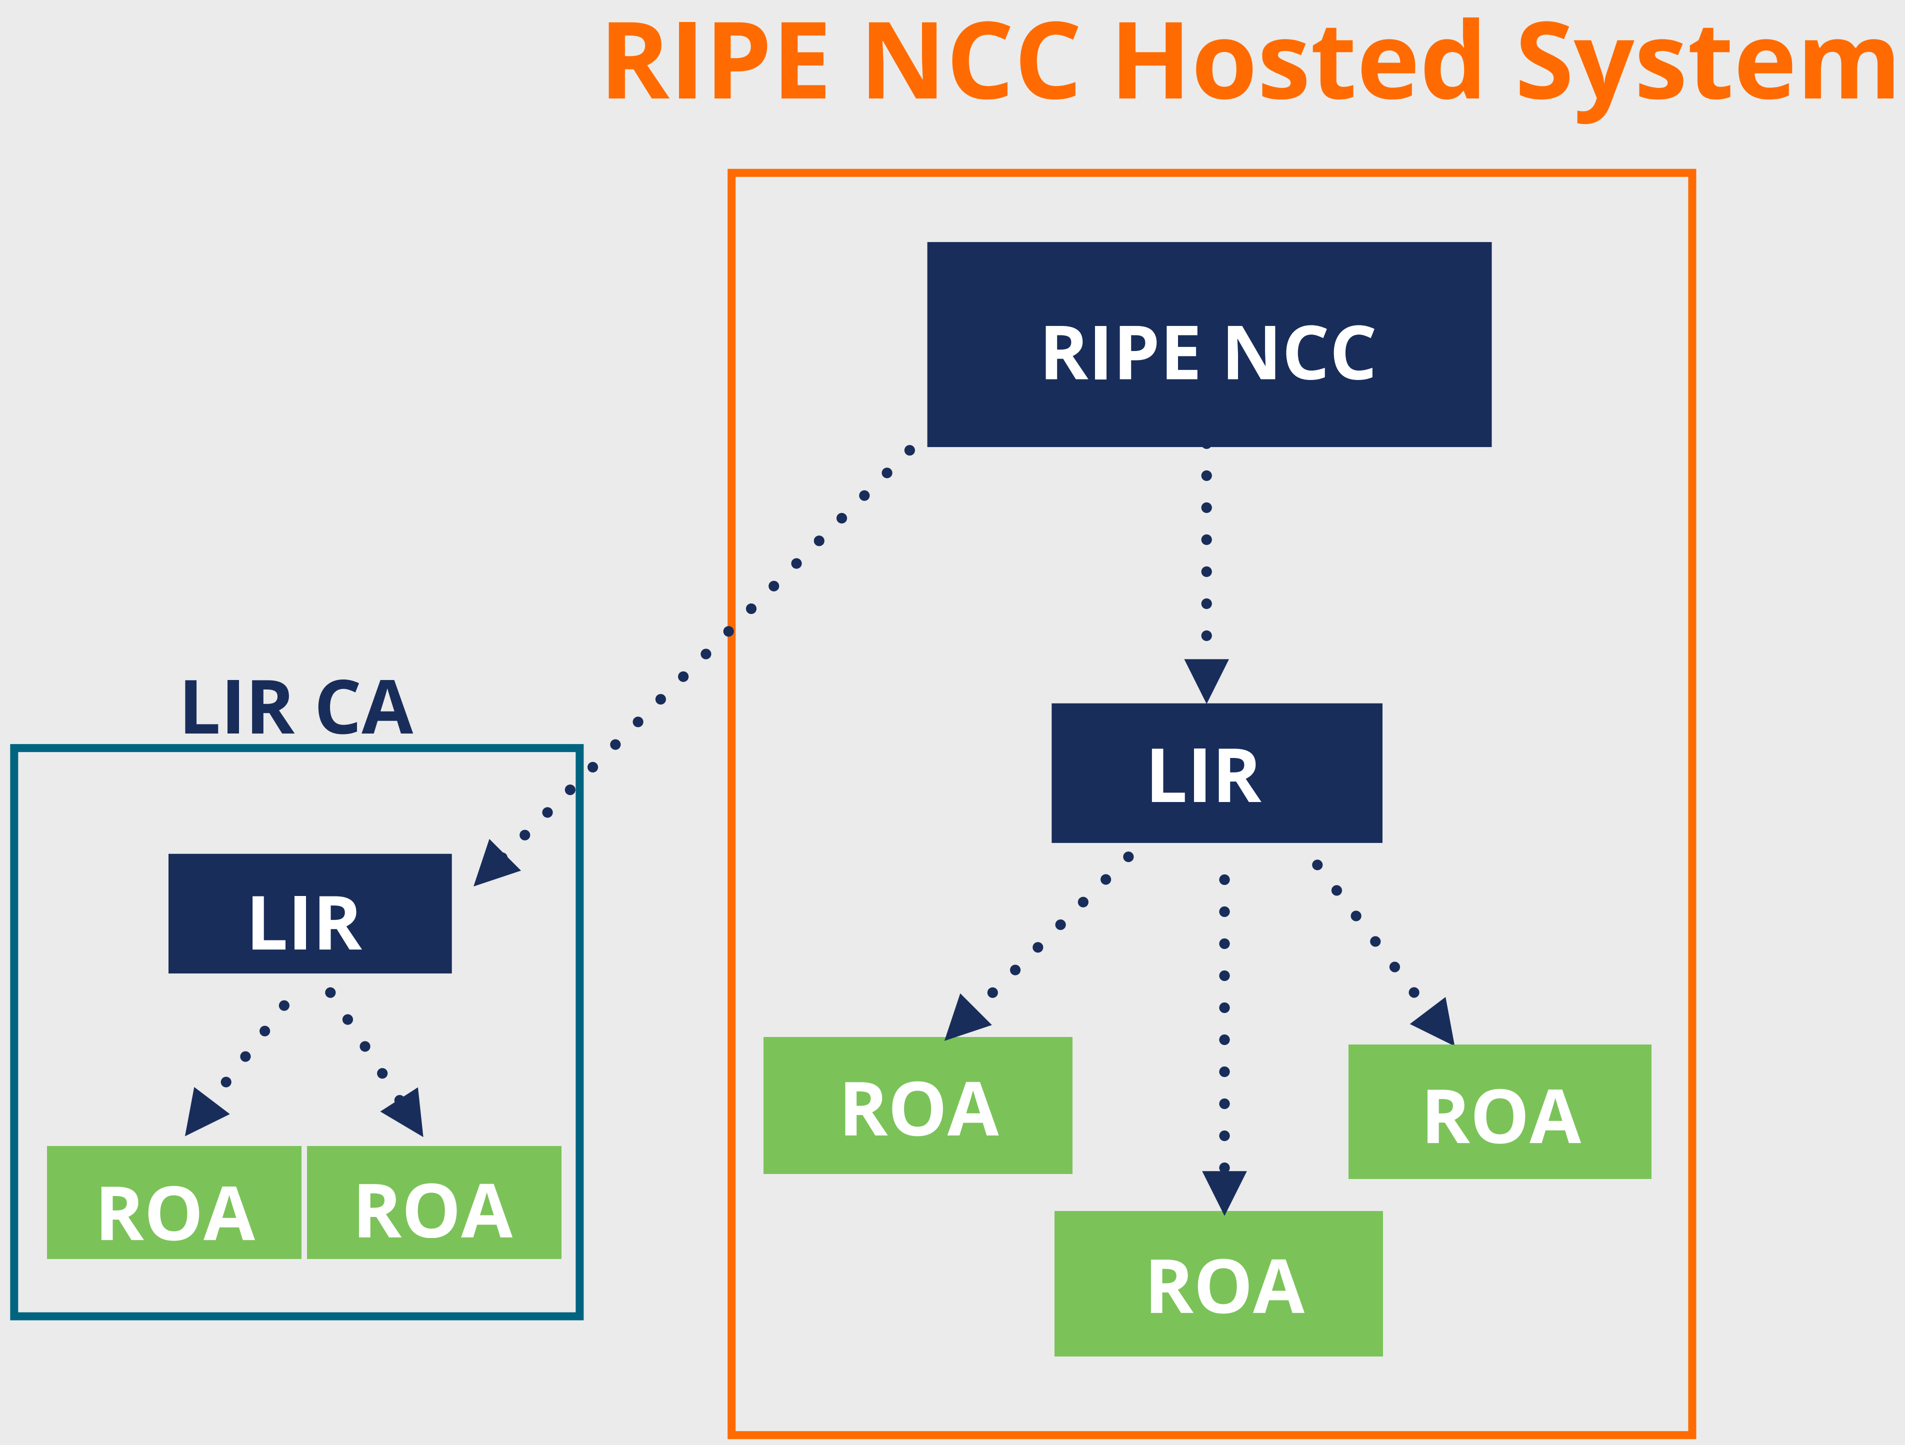

Elements of RPKI

- The RPKI system consists of two parts:

SIGNING	+	VALIDATION
Create ROAs for your prefixes in the RPKI system		Verify the information provided by others

RPKI Validation

- **Verifying the information provided by others**
- **First, validate the RPKI data**
  - Install a validator software locally in your network
  - Verify holdership through a public key and certificate infrastructure
- **Second, validate the origin of BGP announcements**
  - **Known as BGP Origin Validation (BGP OV) or Route Origin Validation (ROV)**
  - This is done in a BGP router in your network

RPKI Validator

- Also known as **Relying Party (RP)** software
- Connects to RPKI repositories via rsync or RRDP protocol
- Uses information in TALs to connect to the repositories

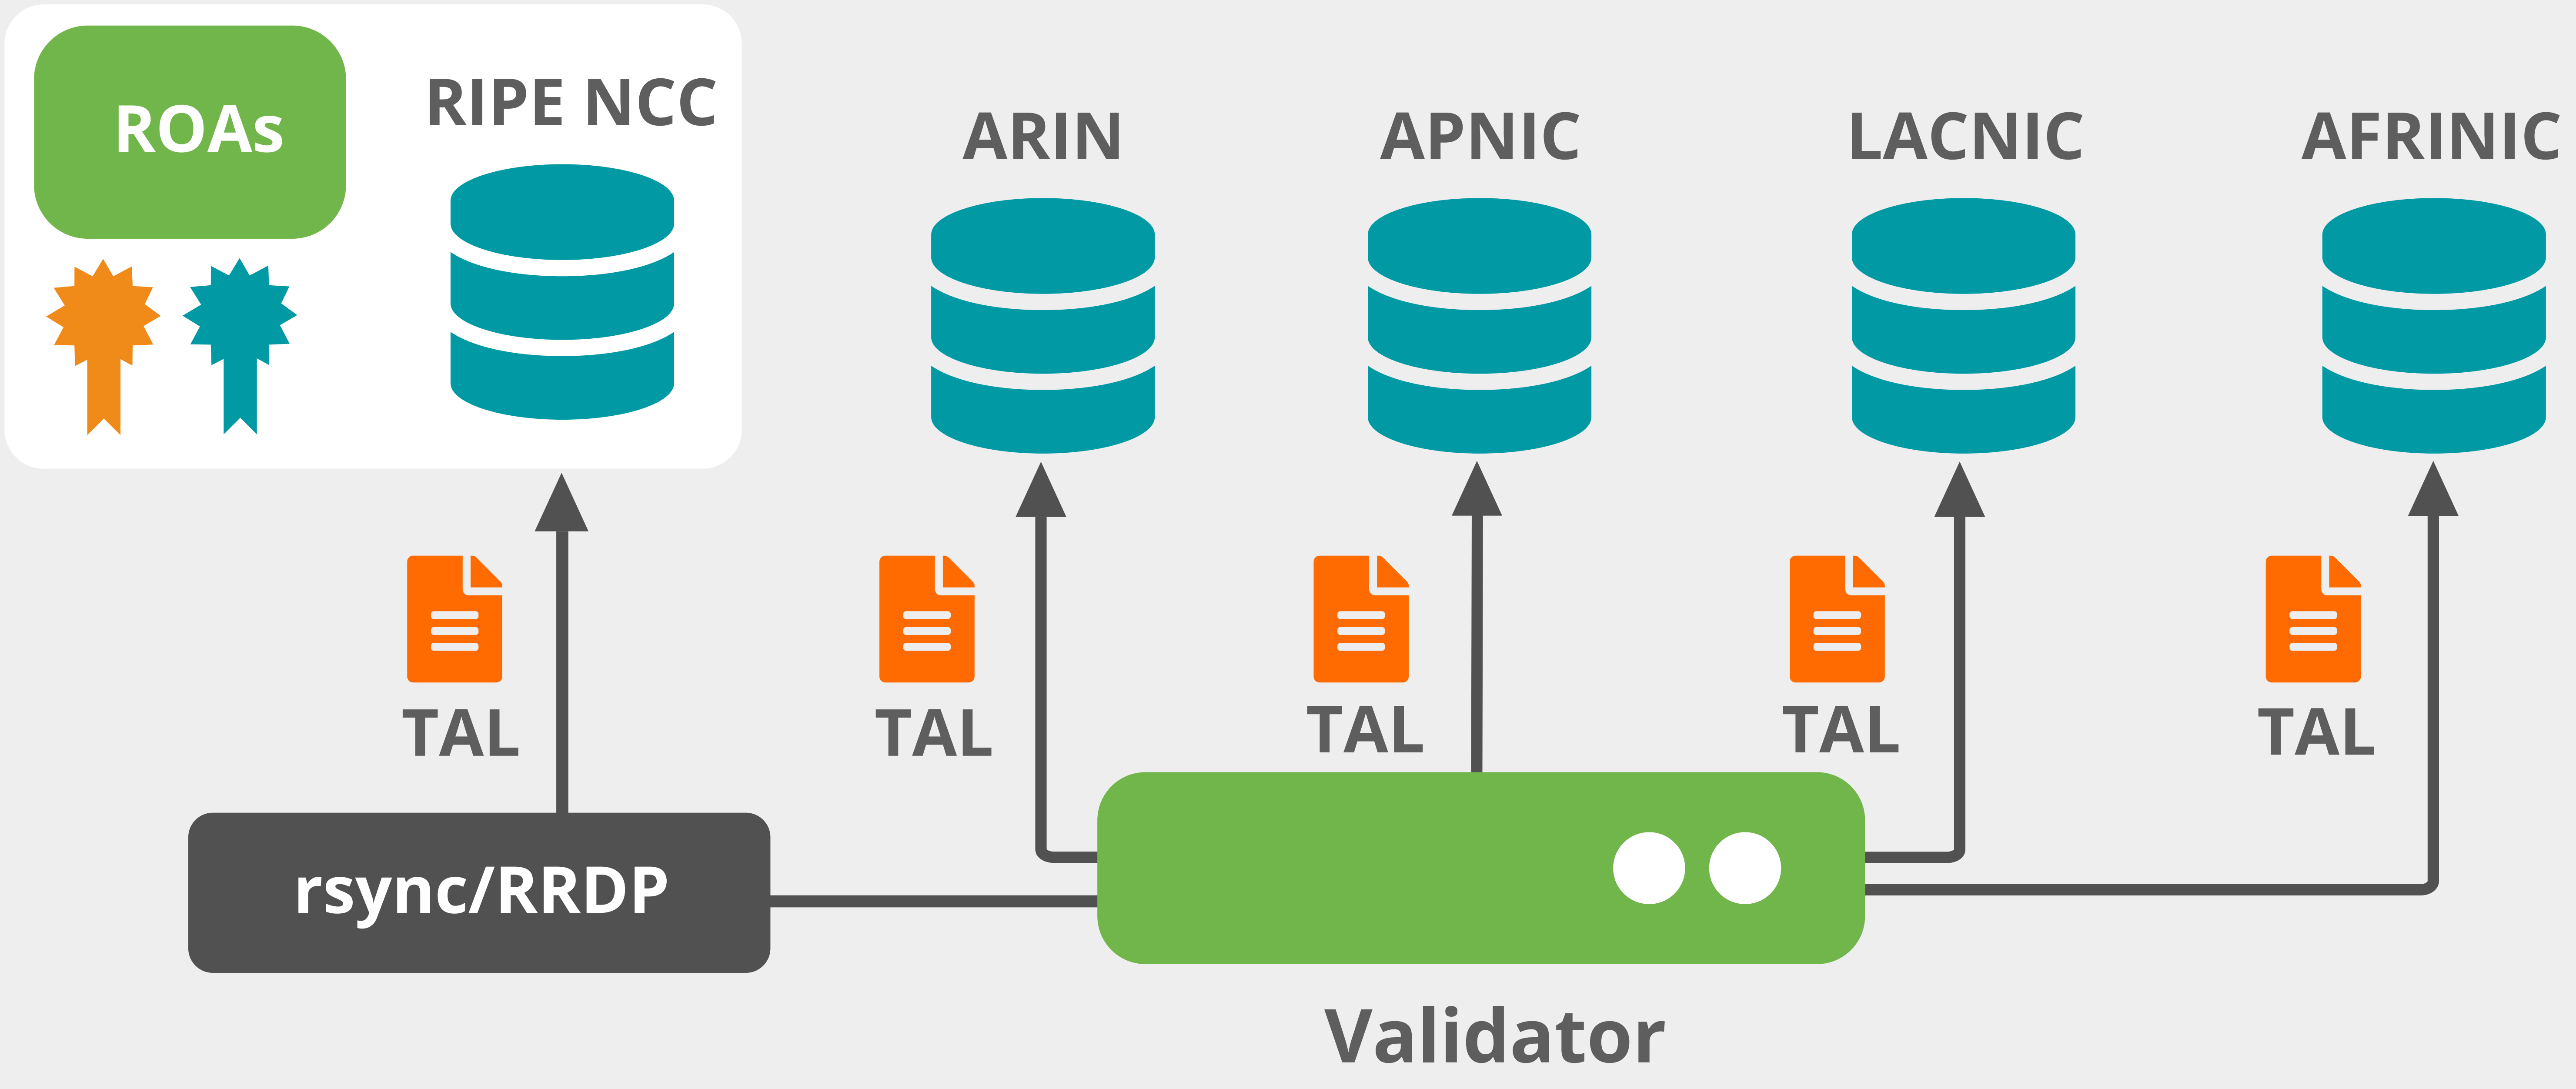

RPKI Validator

- **Downloads 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

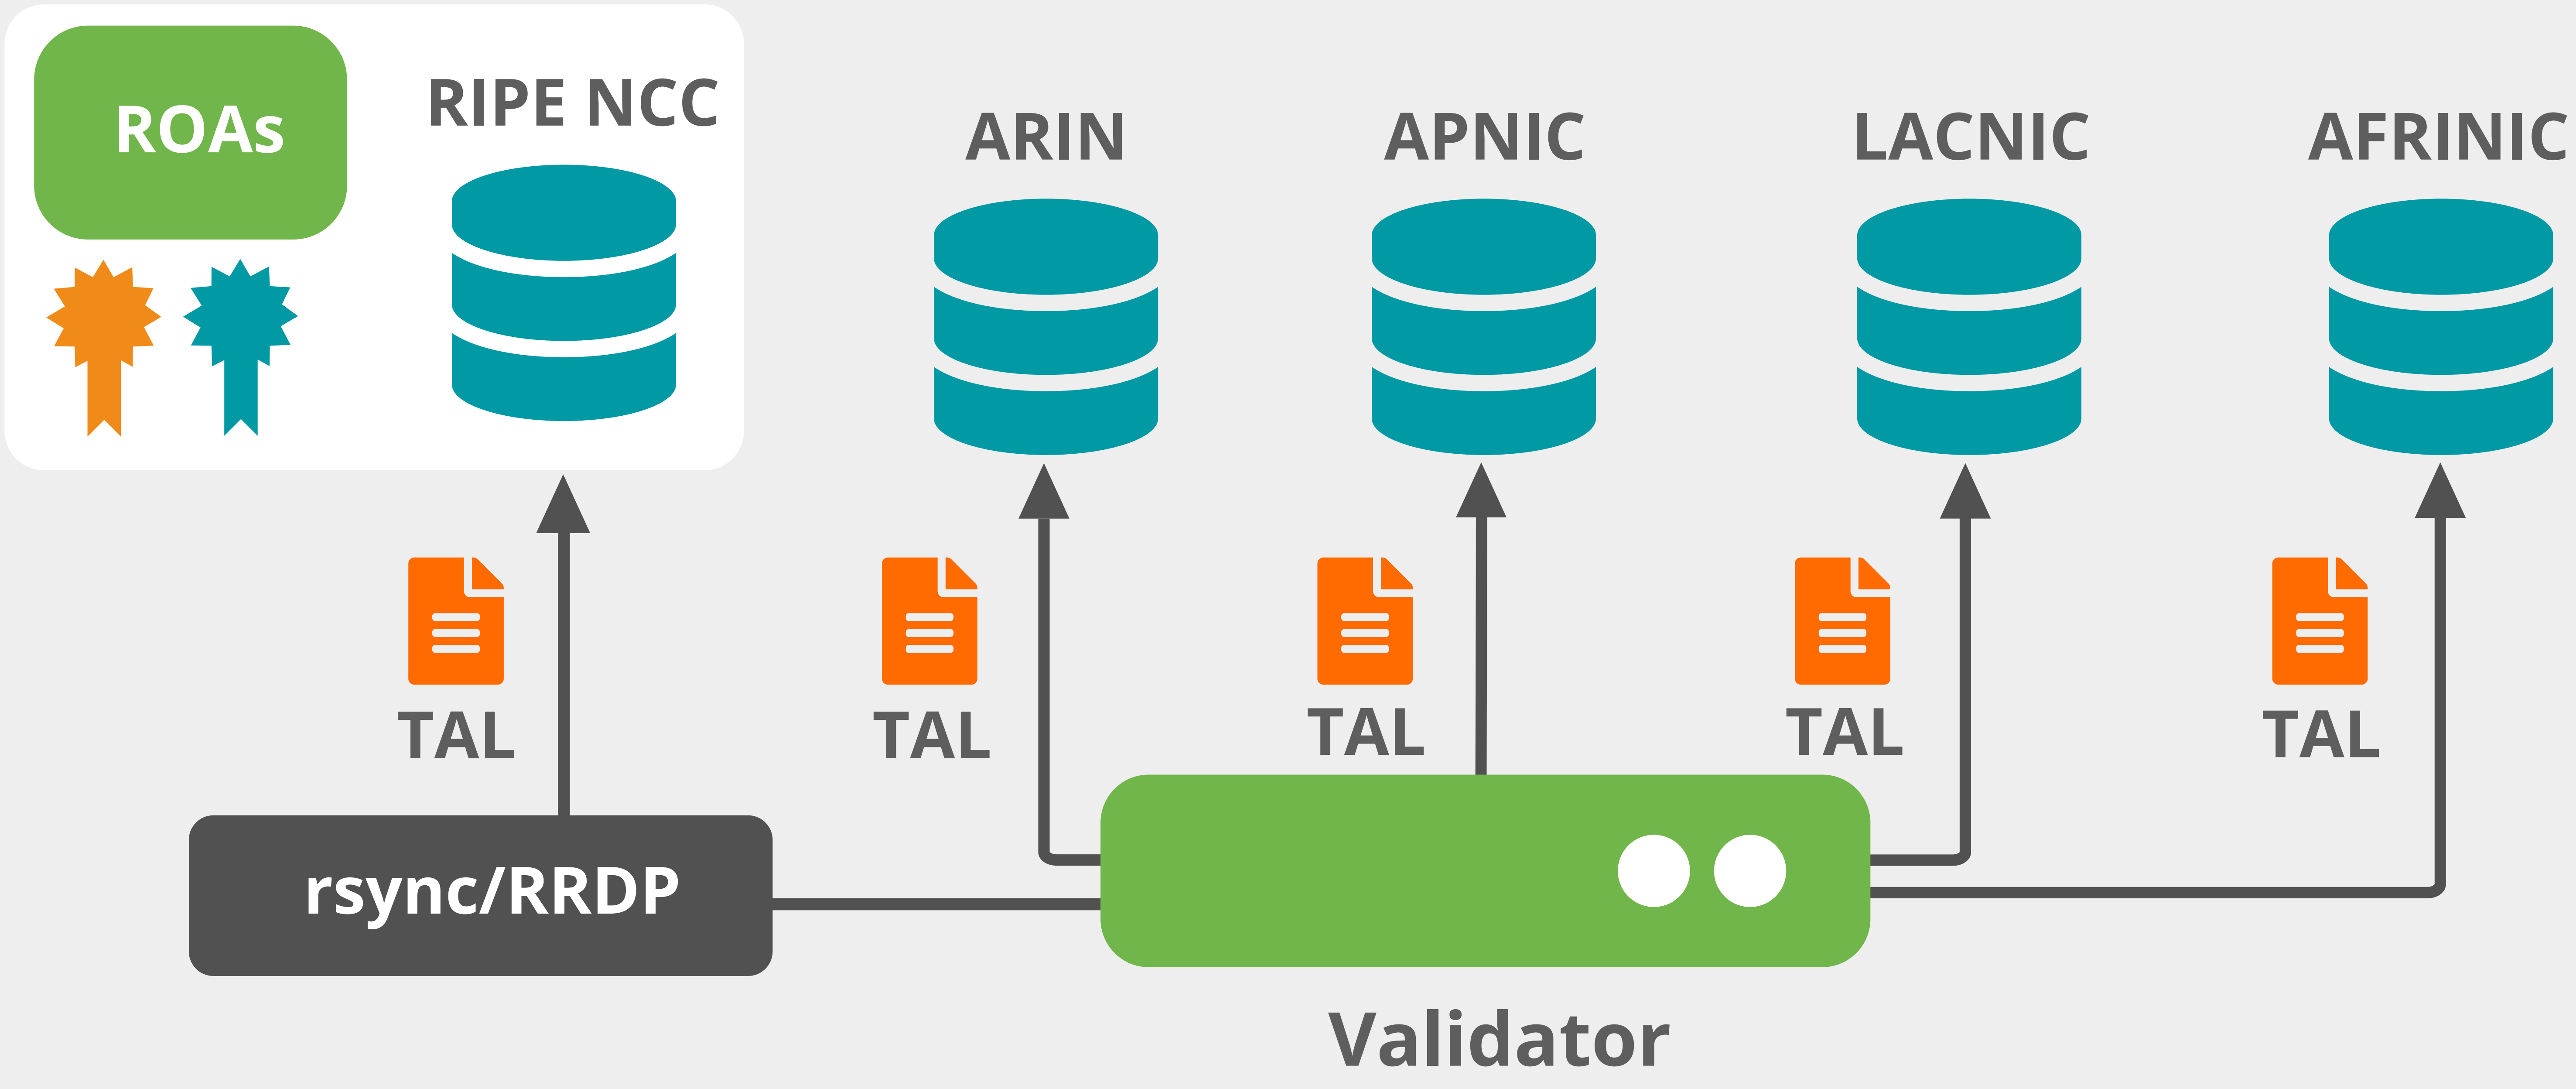

ROA Validation Process

IF chain is complete ==> ROA is **VALID!**

ELSE validation is unsuccessful ==> ROA is **INVALID!**

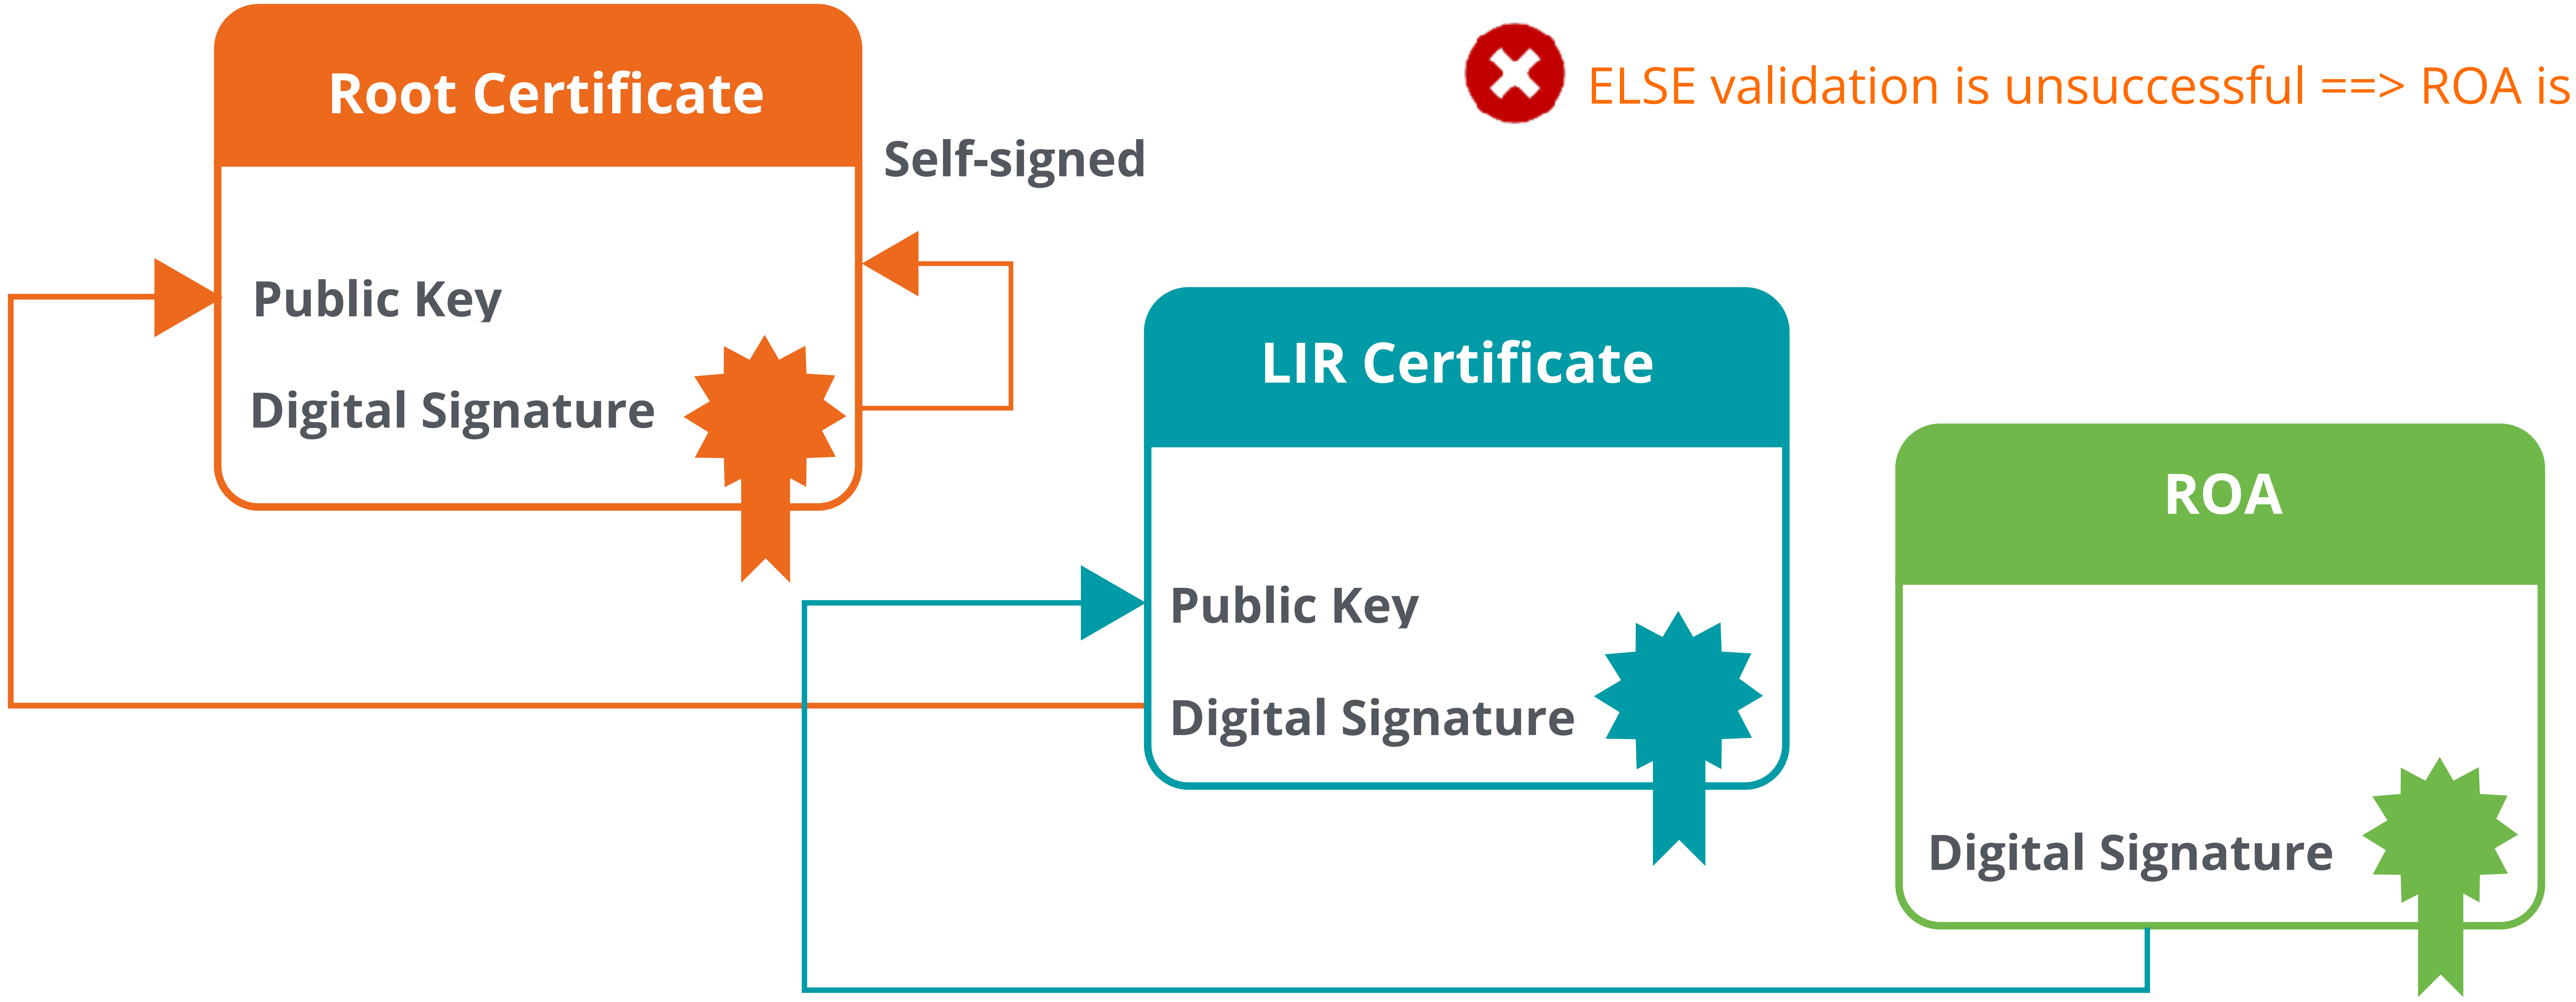

Valid ROAs are sent to the router

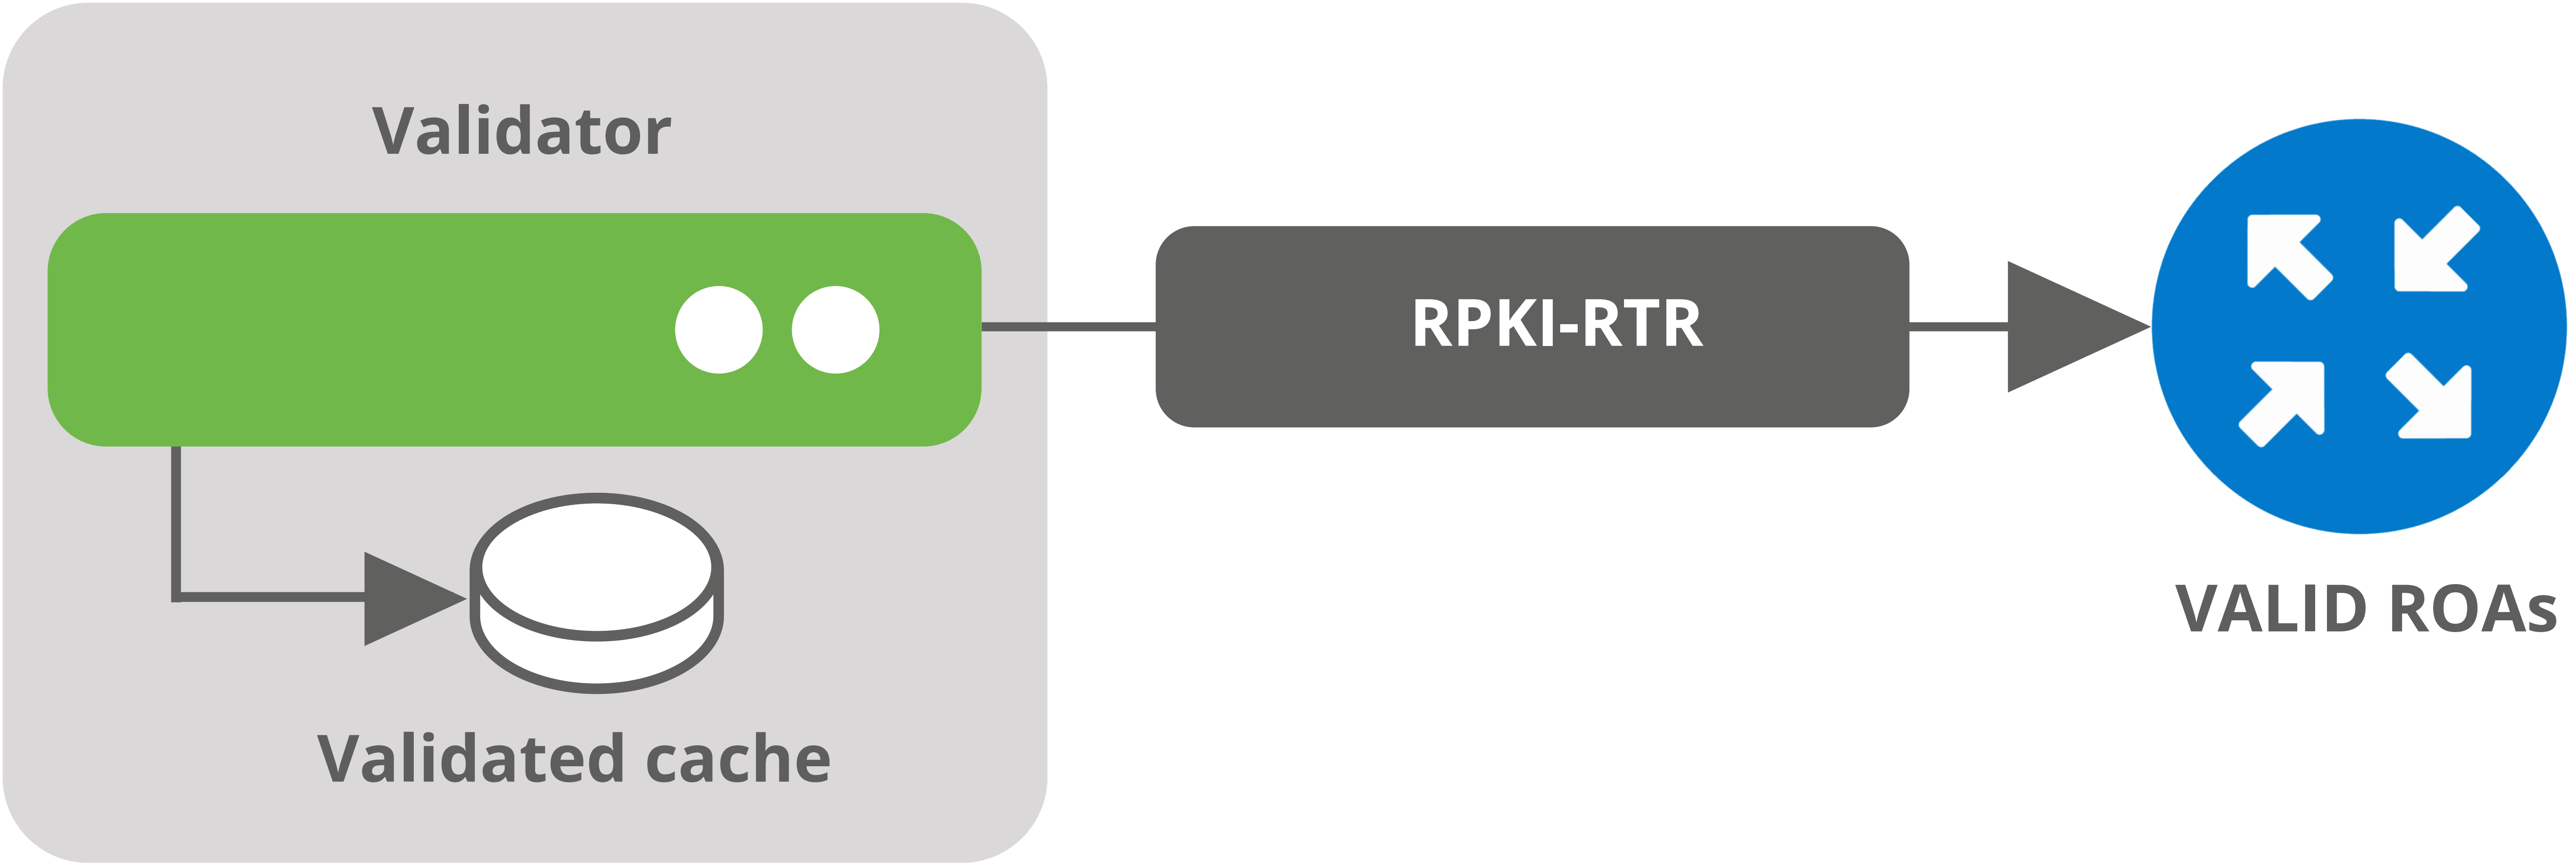

Router uses this information to make better routing decisions!

BGP Origin Validation (BGP OV)

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

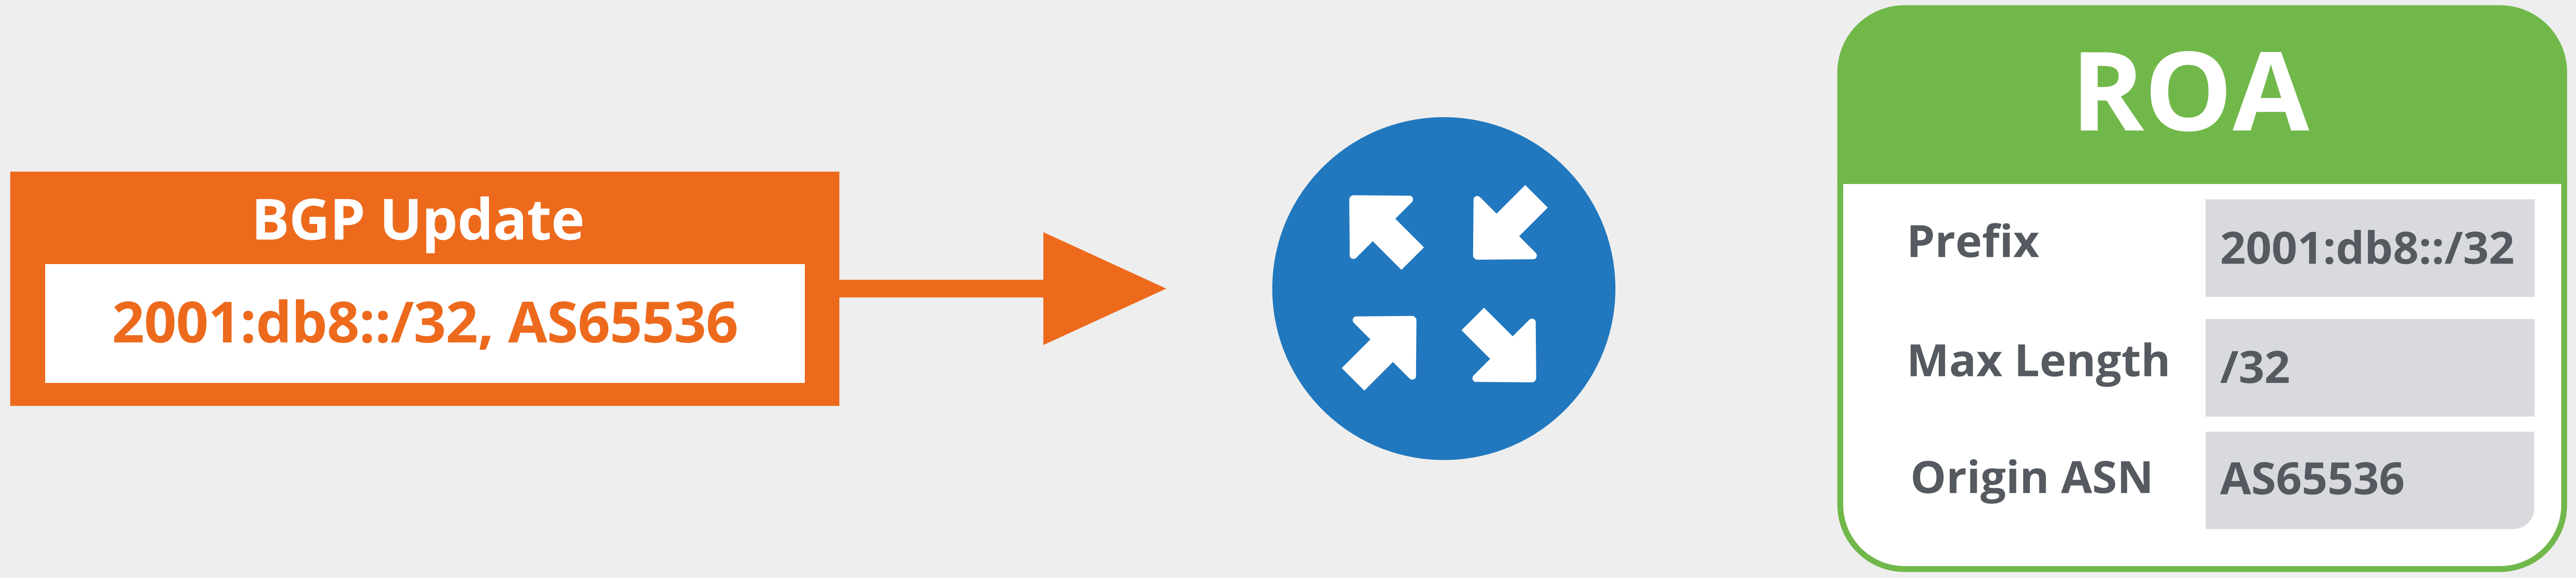

RFC 6811 - BGP Prefix Origin Validation https://datatracker.ietf.org/doc/html/rfc6811

Current Limitations of RPKI

- **RPKI now implements IRR route objects**
- **IRR contains more data**
  - **as-sets: this is used to generate filters**
- **Coming up:**
  - path security (ASPA)
  - Mapping Origin Authorizations
  - BGPsec