Hands-on MANRS Tutorial

RIPE79

Rotterdam, October 2019

Massimiliano Stucchi

stucchi@isoc.org

There is a problem

- 12,600 total incidents (either outages or attacks, like route leaks and hijacks)
- About 4.4% of all Autonomous Systems on the Internet were affected
- 2,737 Autonomous Systems were a victim of at least one routing incident
- 1,294 networks were responsible for 4739 routing incidents

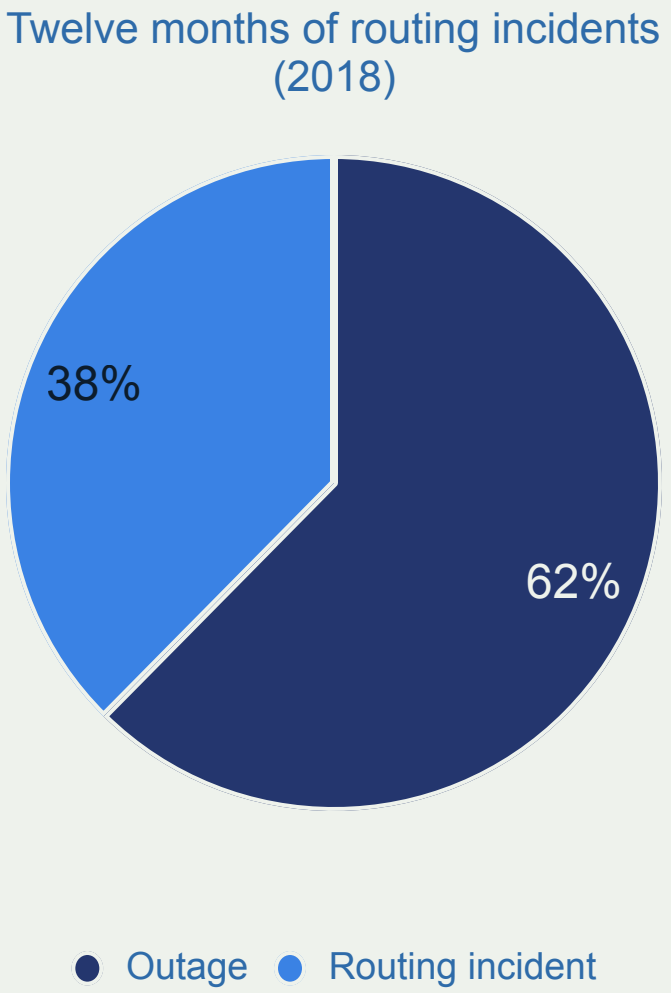

Source: https://www.bgpstream.com/

Routing Incidents Cause Real World Problems

Prefix/Route Hijacking

Route Leaks

IP address spoofing

We Are In This Together

Network operators have a collective responsibility to ensure a globally robust and secure routing infrastructure.

Your network's safety depends on a routing infrastructure that mitigates incidents from bad actors and accidental misconfigurations that wreak havoc on the Internet.

Security of your network depends on measures taken by other operators.

The more network operators work together, the fewer incidents there will be, and the less damage they can do.

Mutually Agreed Norms for Routing Security

MANRS provides baseline recommendations in the form of Actions

- Distilled from common behaviours – BCPs, optimised for low cost and low risk of deployment
- With high potential of becoming norms

MANRS builds a visible community of security minded operators

- Social acceptance and peer pressure

MANRS

MANRS for Network operators

Coordination

Facilitate global operational communication and coordination between network operators

Maintain globally accessible up-to-date contact information in common routing databases

Global Validation

Facilitate validation of routing information on a global scale

Anti-spoofing

Prevent traffic with spoofed source IP addresses

Enable source address validation for at least single-homed stub customer networks, their own end-users, and infrastructure

Filtering

Prevent propagation of incorrect routing information

Ensure the correctness of your own announcements and announcements from your customers to adjacent networks with prefix and AS-path granularity

MANRS for IXPs

Action 1
Prevent propagation of incorrect routing information

This mandatory action requires IXPs to implement filtering of route announcements at the Route Server based on routing information data (IRR and/or RPKI).

Action 2
Promote MANRS to the IXP membership

IXPs joining MANRS are expected to provide encouragement or assistance for their members to implement MANRS actions.

Action 3
Protect the peering platform

This action requires that the IXP has a published policy of traffic not allowed on the peering fabric and performs filtering of such traffic.

Action 4
Facilitate global operational communication and coordination

The IXP facilitates communication among members by providing necessary mailing lists and member directories.

Action 5
Provide monitoring and debugging tools to the members.

The IXP provides a looking glass for its members.

MANRS for CDN & Cloud: a draft action set

Action 1
Prevent propagation of incorrect routing information

Egress filtering

Ingress filtering – non-transit peers, explicit whitelists

Action 2
Prevent traffic with spoofed source IP addresses

Anti-spoofing controls to prevent packets with illegitimate source IP address

Action 3
Facilitate global operational communication and coordination

Contact information in PeeringDB

and relevant RIR databases

Action 4
Facilitate validation of routing information on a global scale

Publicly document ASNs and prefixes that are intended to be advertised to external parties.

Action 5
Encourage MANRS adoption

Actively encourage MANRS adoption among the peers

Action 6
Provide monitoring and debugging tools to peering partners

Provide monitoring tools to indicate incorrect announcements from peers that were filtered by the CDN&Cloud operator.

MANRS Implementation Guide

- Based on Best Current Operational Practices
- Published as RIPE-706

https://www.manrs.org/bcop/

Mutually Agreed Norms for Routing Security (MANRS) Implementation Guide

Version 1.0, BCOP series
Publication Date: 25 January 2017

MANRS – increasing adoption

209 ISPs

39 IXPs

Coordination and Global Validation

Data needs to be published in:

RIR Databases

PeeringDB

IRR Data and RPKI

RIR Database (RIPE NCC)

- You need to have:
  - A **Maintainer**
  - A **Person** object, not mandatory
  - A **Role** Object for your NOC or Team

- All of these need to be referenced in:
  - the **Organisation** object, and
  - the **Inetnums** or **Inet6nums** for your allocations

RIPE NCC Trainings

- Face to face courses available (Cyprus, Lyon, Doha coming up)
- Webinars
- Certification

https://www.ripe.net/training

Two Functions for the Role Object

Group of Persons

Abuse Contact

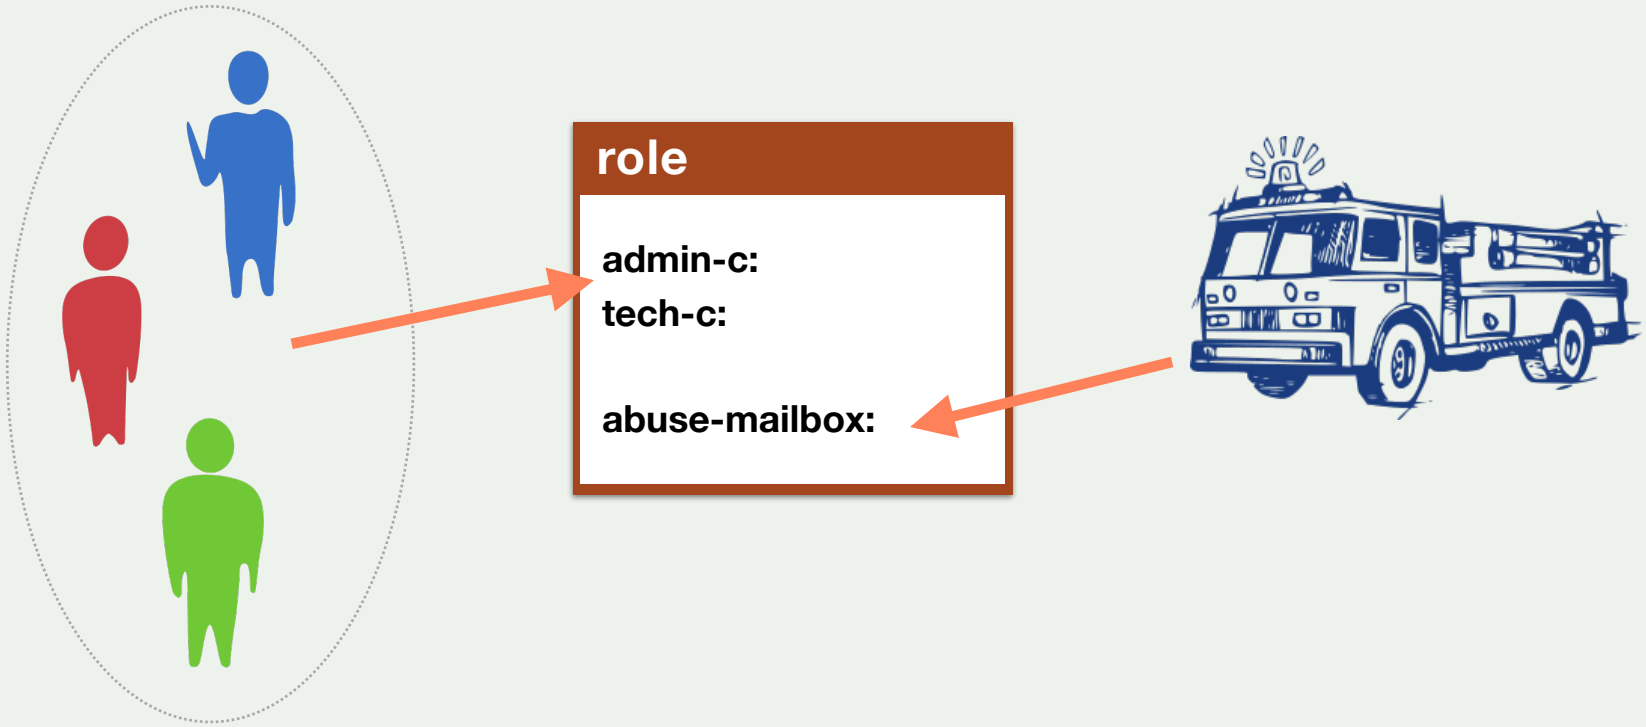

Role Object: Abuse Contact

- The **role** object contains the “abuse-mailbox:”
- Objects reference the **role** in “abuse-c:”
- RIPE Database shows the abuse contact in WHOIS query results

Role Object: Group of Persons

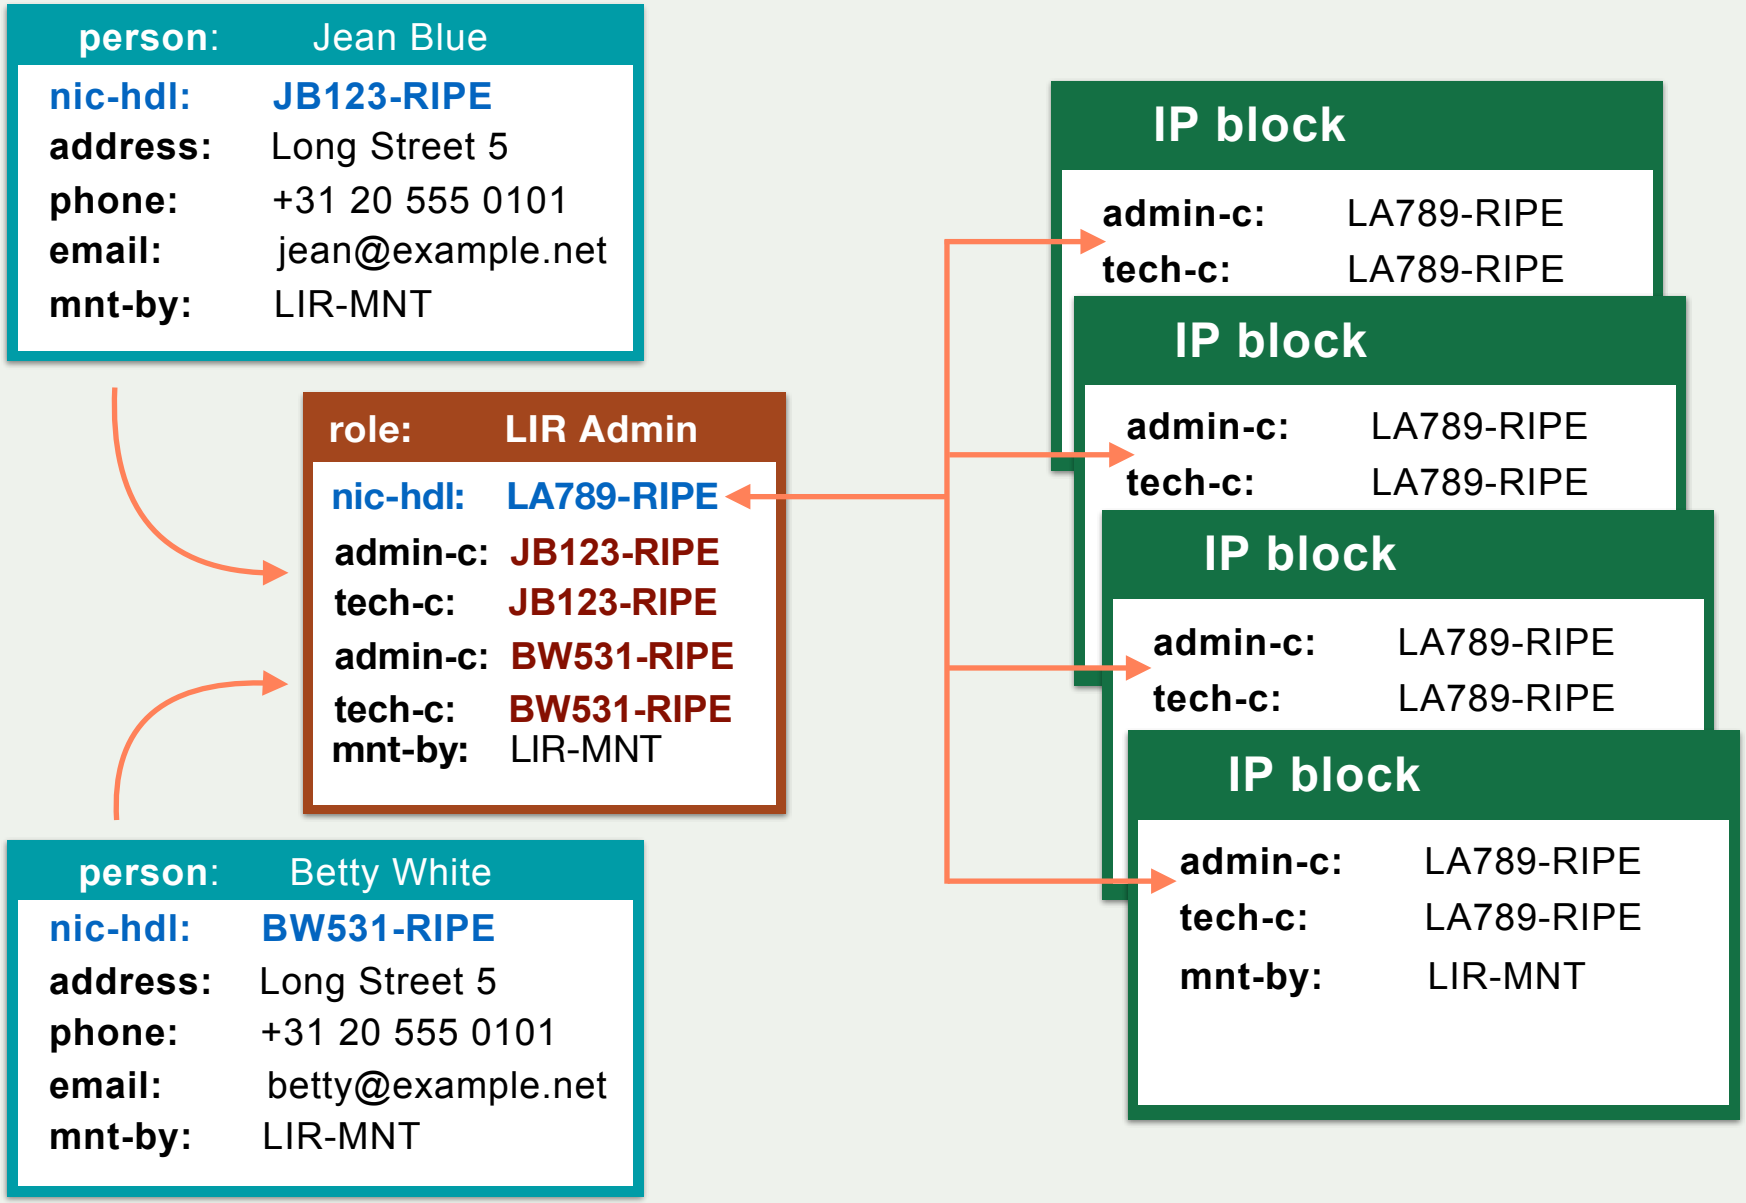

Inetnum and inet6num

IPv4 = inetnum

inetnum: 192.30.0.0 - 192.30.3.255

netname:	NL-NETWORK-20170101
country:	NL
org:	ORG-EE2-RIPE
admin-c:	DV789-RIPE
tech-c:	JS123-RIPE
status:	ALLOCATED PA
mnt-by:	RIPE-NCC-HM-MNT
mnt-by:	DEFAULT-LIR-MNT
source:	RIPE

IPv6 = inet6num

inet6num: 2001:db8::/32

netname:	NL-NETWORK-20170101
country:	NL
org:	ORG-EE2-RIPE
admin-c:	DV789-RIPE
tech-c:	JS123-RIPE
status:	ALLOCATED-BY-RIR
mnt-by:	RIPE-NCC-HM-MNT
mnt-by:	DEFAULT-LIR-MNT
source:	RIPE

PeeringDB

- Register an account

- Associate it with your ASN and organisation

- Add all the information you can, especially:
  - Contacts
  - AS-Set
  - IXPs, Facilities where you peer/have a PoP

PeeringDB API

- Tools exist to leverage the PeeringDB API
- Helps deciding where to peer
- Helps understand which networks are available at an IXP or a facility
- Makes your life easier

Route(6) Objects

- **route(6)** objects register which IPv4/IPv6 prefix will be announced by which AS number
- Used for creating BGP filters

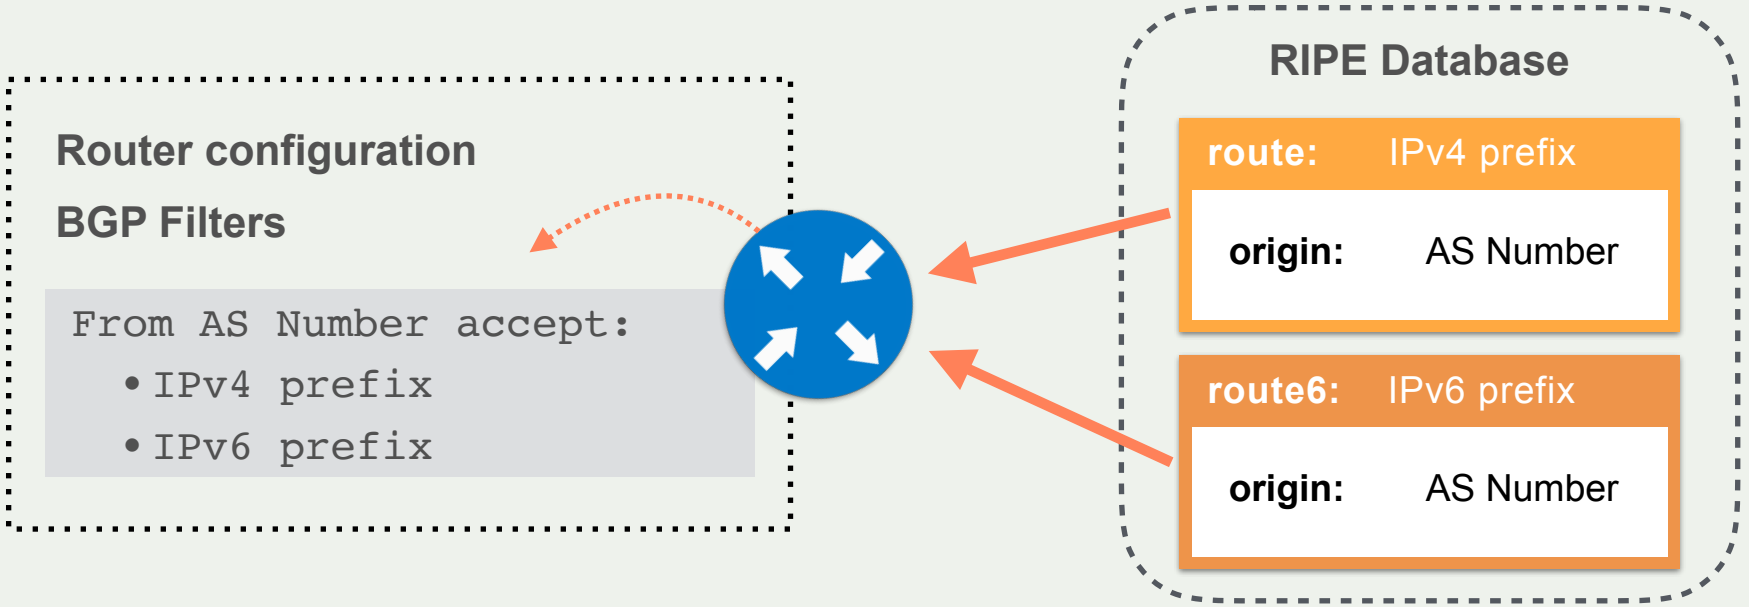

Registering IPv6 Routes

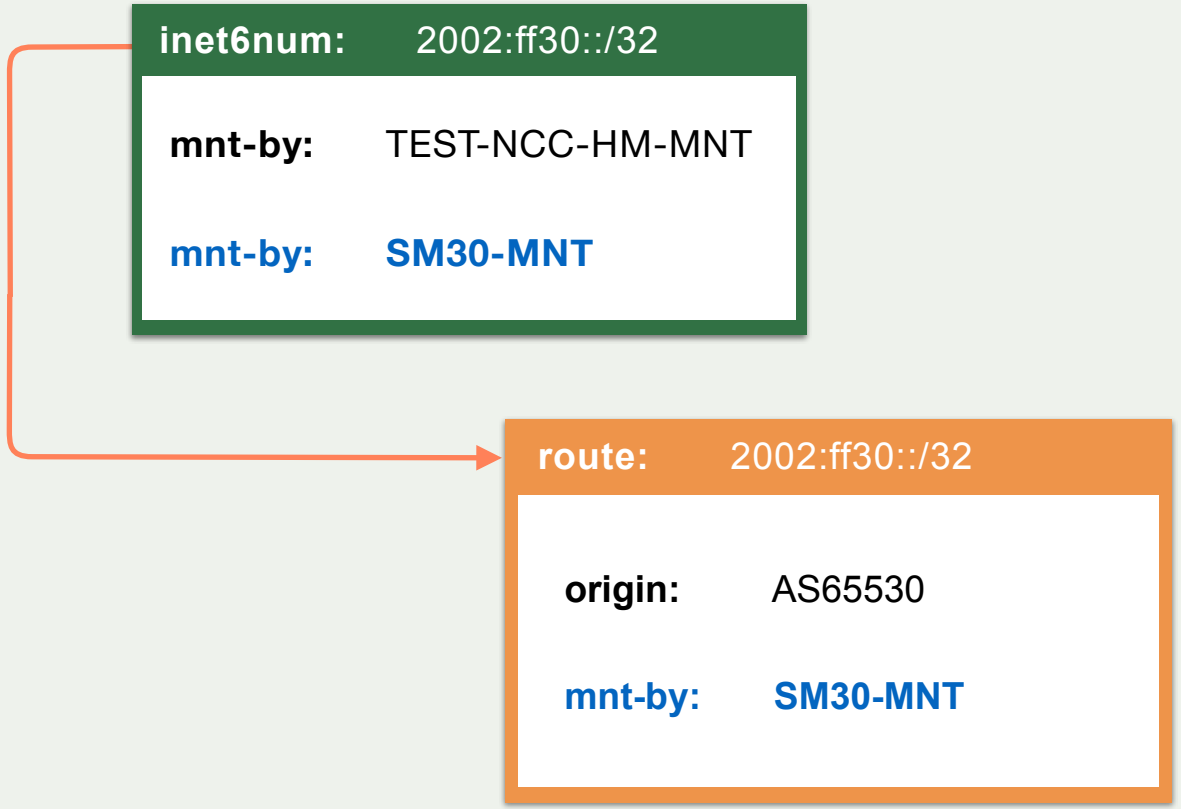

Registering IPv4 Routes

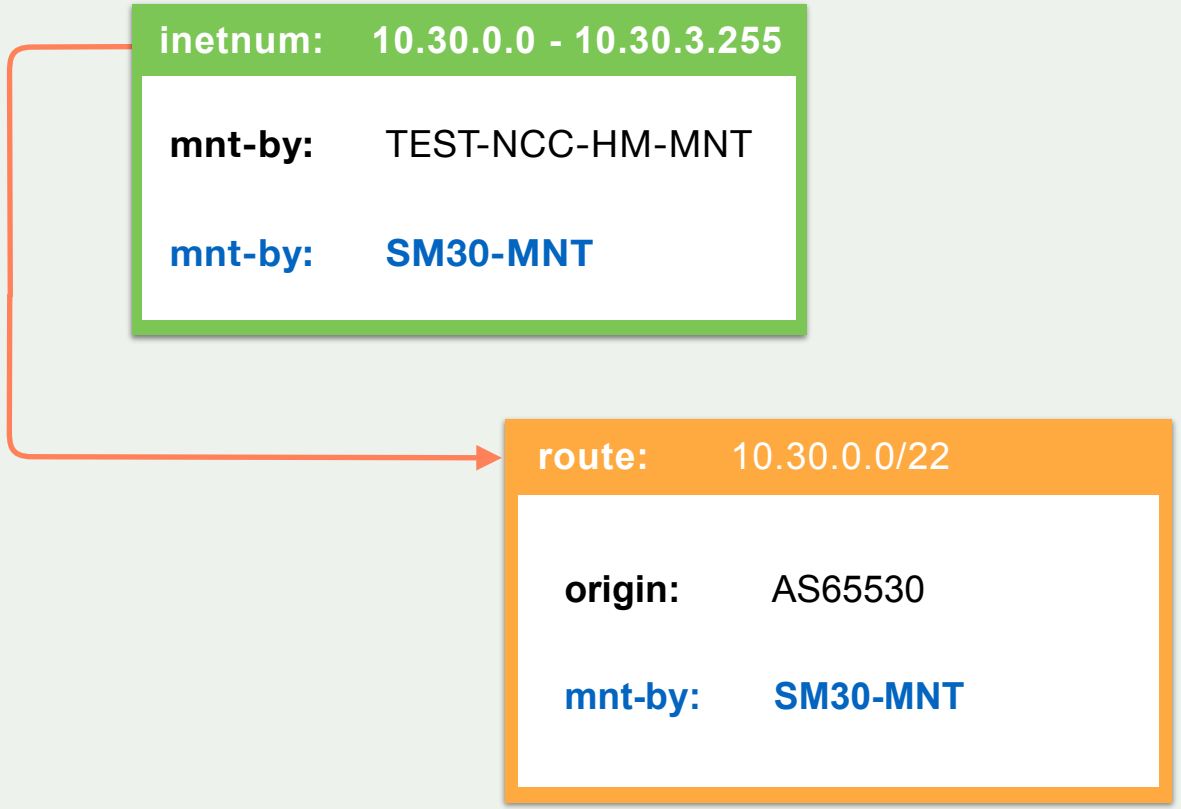

Route Objects and AS-Sets

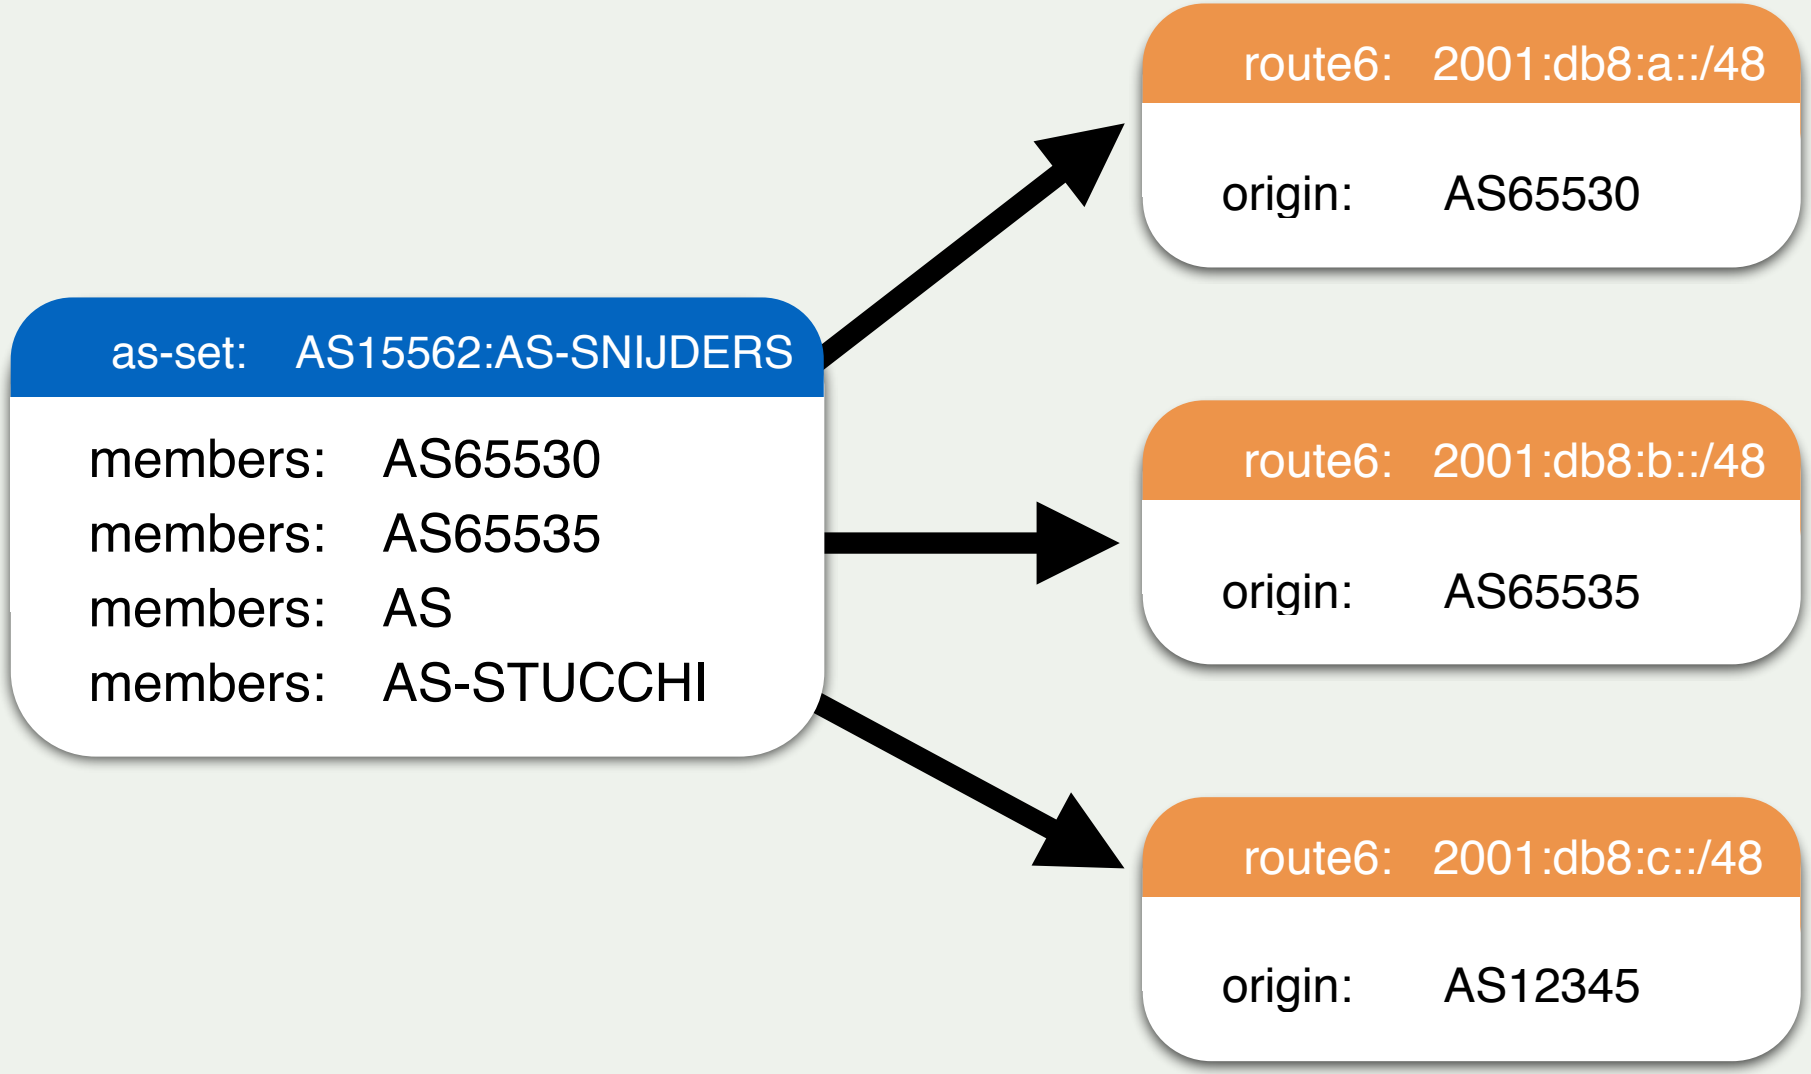

Limitation of AS-Sets

- Ask customers in service order form
- Look at PeeringDB
- Different trust levels based on the IRR

Limitations of AS-Sets

- Can exist in multiple IRRs

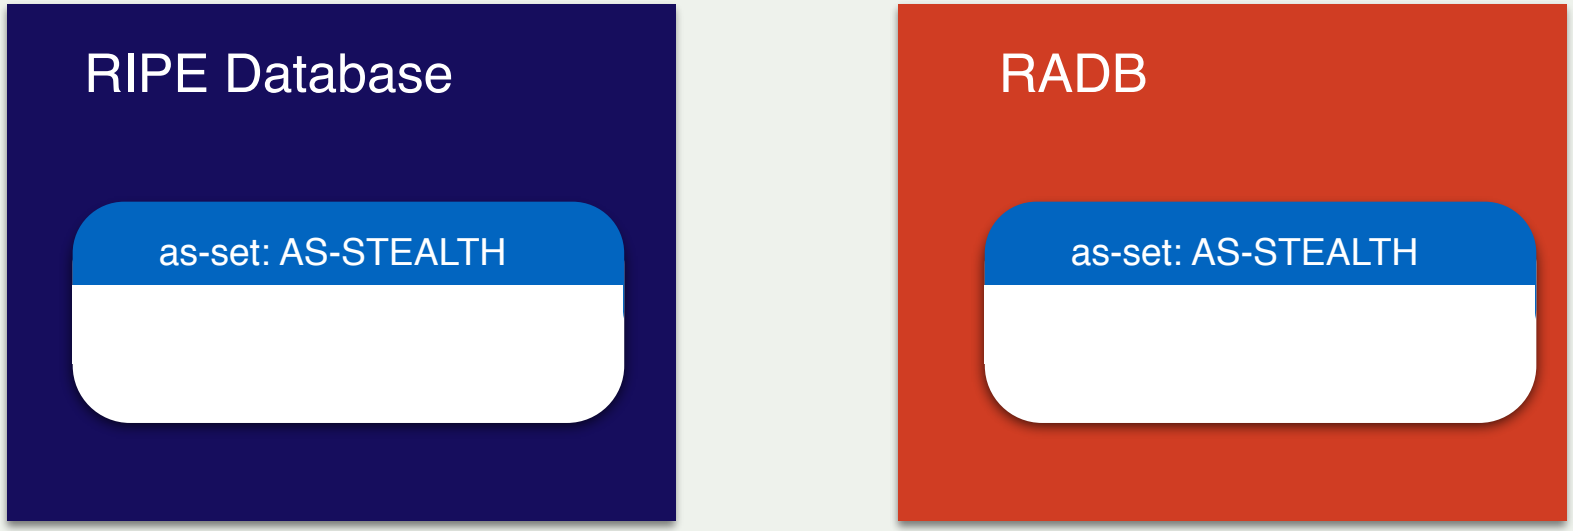

- AS-STEALTH exists in both the RIPE Database and RADB
- The two are not managed by the same organisation

Anti-Spoofing (RPF)

Reverse Path Forwarding

- Called uRPF (Unicast Reverse Path Forwarding)
- Checks if an entry exists in the routing table before accepting the packet and forwarding it

- Four modes
  - Loose
  - Strict
  - Feasible Path
  - VRF

uRPF Modes

Loose

Check that an entry exists in the routing table

Strict

Check that an entry exists in the routing table

and the route points to the receiving interface

Feasible

Check that an entry exists in the routing table

or any other route not installed/preferred

VRF

Check that an entry exists in the routing table

and the route points to the receiving interface

Cisco uRPF example

```
interface Gigabitethernet0/0
  ip verify unicast source reachable-via rx
```

Juniper uRPF example

```
[edit interface ge-0/0/0 unit 0 family inet]
rpf-check;
```

Using ACLs for source validation

- ACLs can also be used
  - Towards a provider’s servers
  - Towards Infrastructure networks
  - When uRPF cannot be used because of platform limitations

Cisco ACL example

```
ip access-list extended fromCUSTOMER
 permit ip 192.168.0.0 0.0.255.255 any
 permit ip 10.0.0.0 0.0.0.3 any
 deny   ip  any any
!
interface Gigabitethernet0/0
 ip access-group fromCUSTOMER in
!
```

Juniper ACL example

```
 firewall family inet {
filter fromCUSTOMER {
 term CUSTOMER {
  from source-address {
   192.168.0.0/16;
   10.0.0.0/30;
  }
  then accept;
 }
 term Default {
  then discard;
 }
}
}
[edit interface ge-0/0/0 unit 0 family inet]
filter {
 input fromCUSTOMER;
 }
```

Filtering

What is filtering

- Techniques used to decide which routes to allow inside your routing table or network
  - and also what you announce to your neighbours

Why is filtering important ?

- Your first line of defence

- You control what you are announcing
  - You have **no control** over what **other networks** announce

- To avoid issues, **you have to decide** what to accept from other networks

Data sources

- IRRs
- Bogons lists (IPv6 & IPv4)
- PeeringDB
  - For AS-Sets

Generating a Prefix Filter

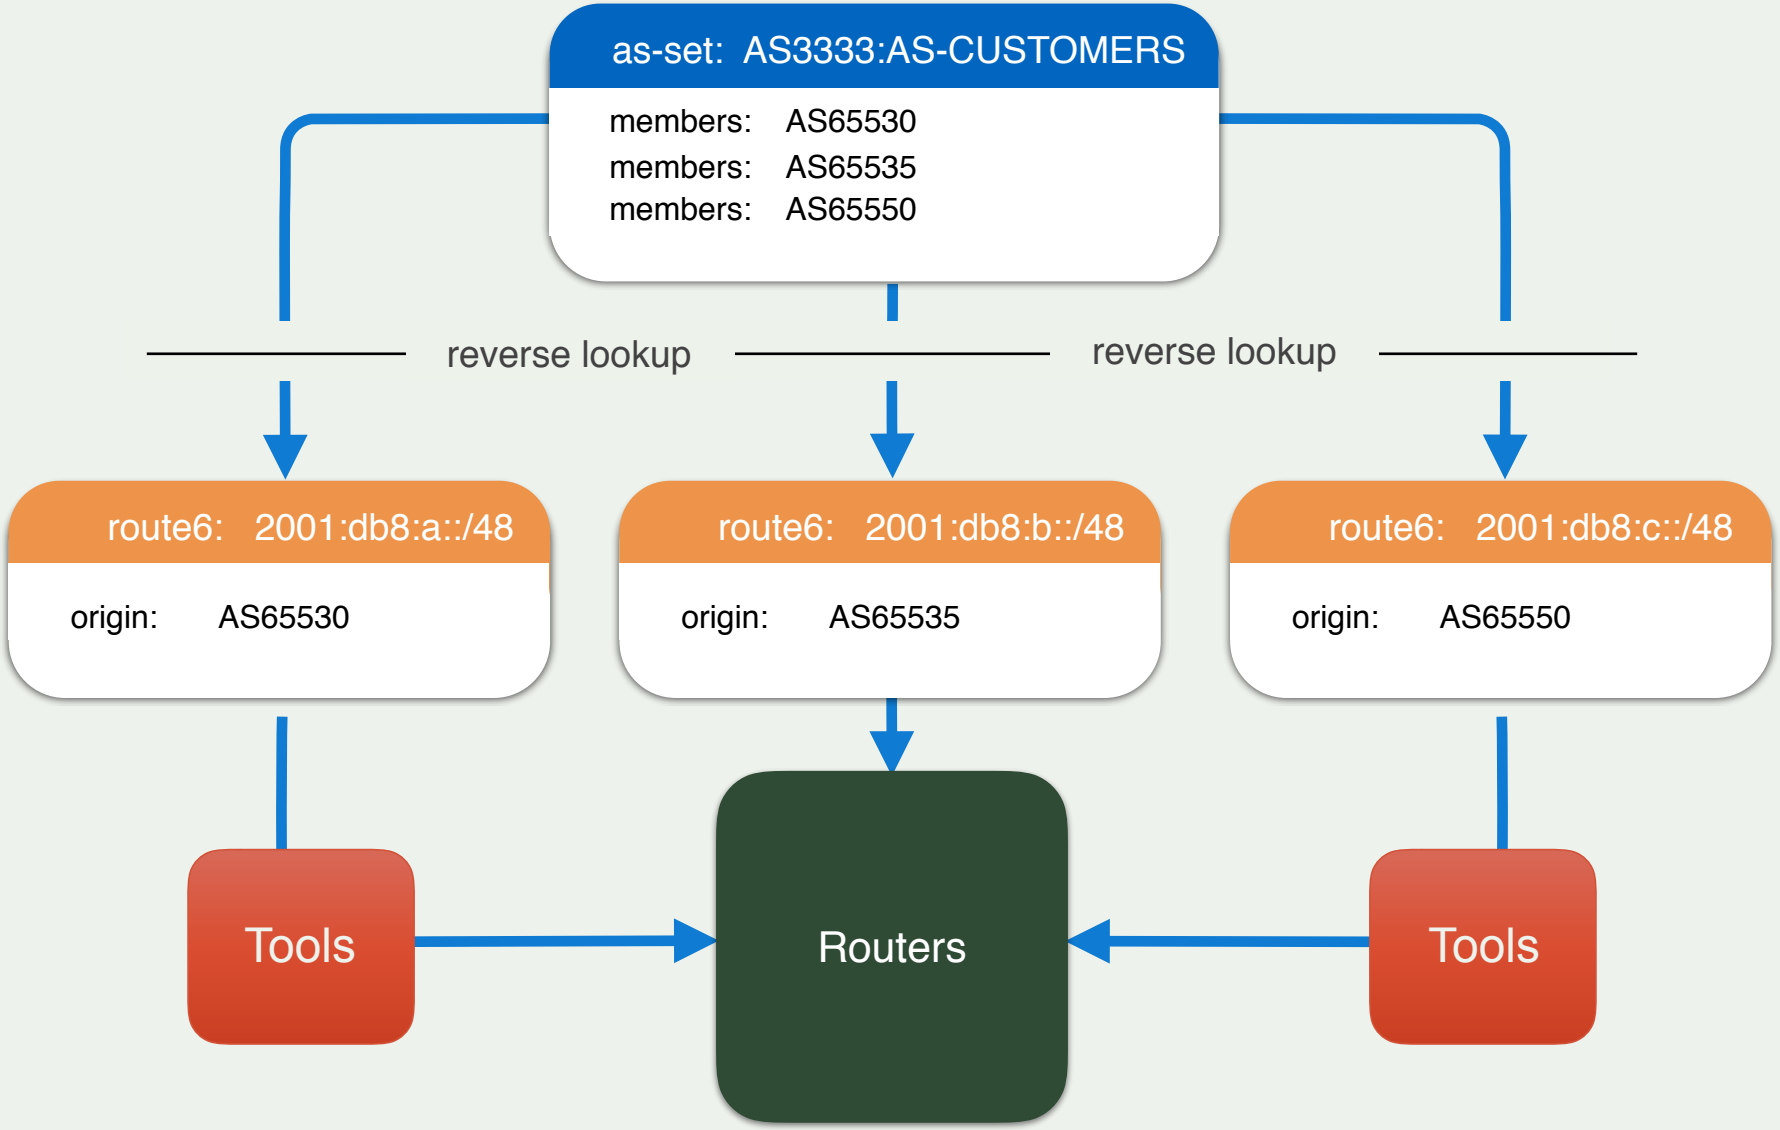

Generating a prefix list

- Check the AS-Set
  - Walk the AS-Set and prepare a list of all the ASNs contained
  - If another AS-Set is contained, recursively walk it

- With the list of ASNs, run an inverse query for each one
  - Get the route objects where they are listed as Origin:

Ingress filters

- Best Practices:
  - Don’t accept BOGON ASNs
  - Don’t accept BOGON prefixes
  - Don’t accept your own prefix
  - Don’t accept default (unless you requested it)
  - Don’t accept prefixes that are too specific
  - Don’t accept if AS Path is too long
  - Create filters based on Internet Routing Registries

Bogons

- Routes you shouldn't see in the routing table
  - Private addresses
  - Unallocated space
  - Reserved space (Documentation, Multicast, etc.)

- Team Cymru provides lists for both IPv6 and IPv4, updated daily
  - http://www.team-cymru.com/bogon-reference.html

ASN Bogons

- 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

Prefix-lists

- Lists of routes you want to accept or announce
- You can create them manually or automatically
  - With data from IRRs
- Or using a tool
  - bgpq3
  - peval
  - Level3 Filtergen

Bgpq3

- De-facto standard for generating filters

- Written in C, Uses RADB as primary data source

- Generates filters for Cisco, Juniper, Bird, OpenBGPd
  - Patches exist for Mikrotik

https://github.com/snar/bgpq3

Router provisioning

- Prefix lists should be updated every day
  - Or upon customer/peer request
- Automated procedures should be in place
  - Using NETCONF or better Ansible/Salt
- Pay attention to consider all the cases in your procedures
  - Databases could be down, generating empty lists…

RPKI

RPKI Chain of Trust

RIPE NCC Root Certificate

Self-signed

RPKI Chain of Trust

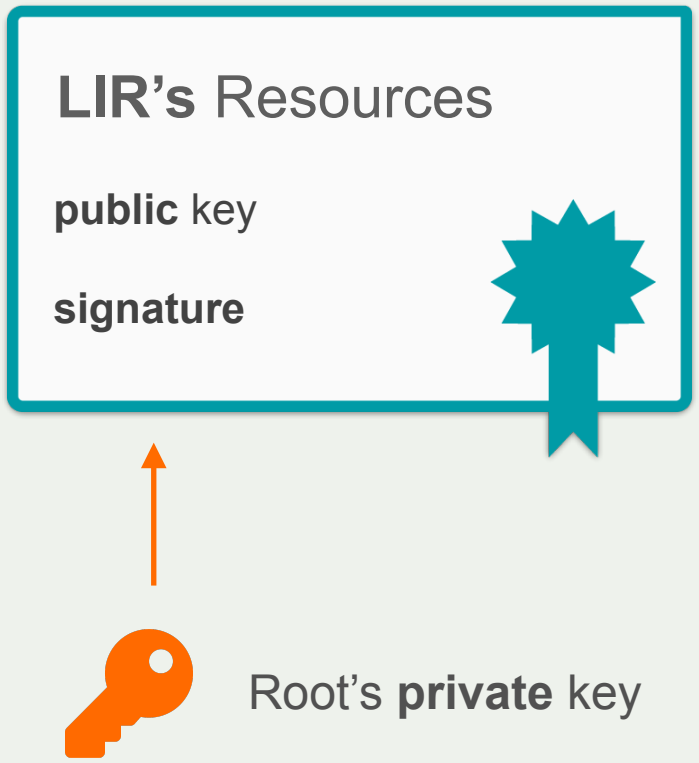

LIR Certificate

Signed by the Root private key

RPKI Chain of Trust

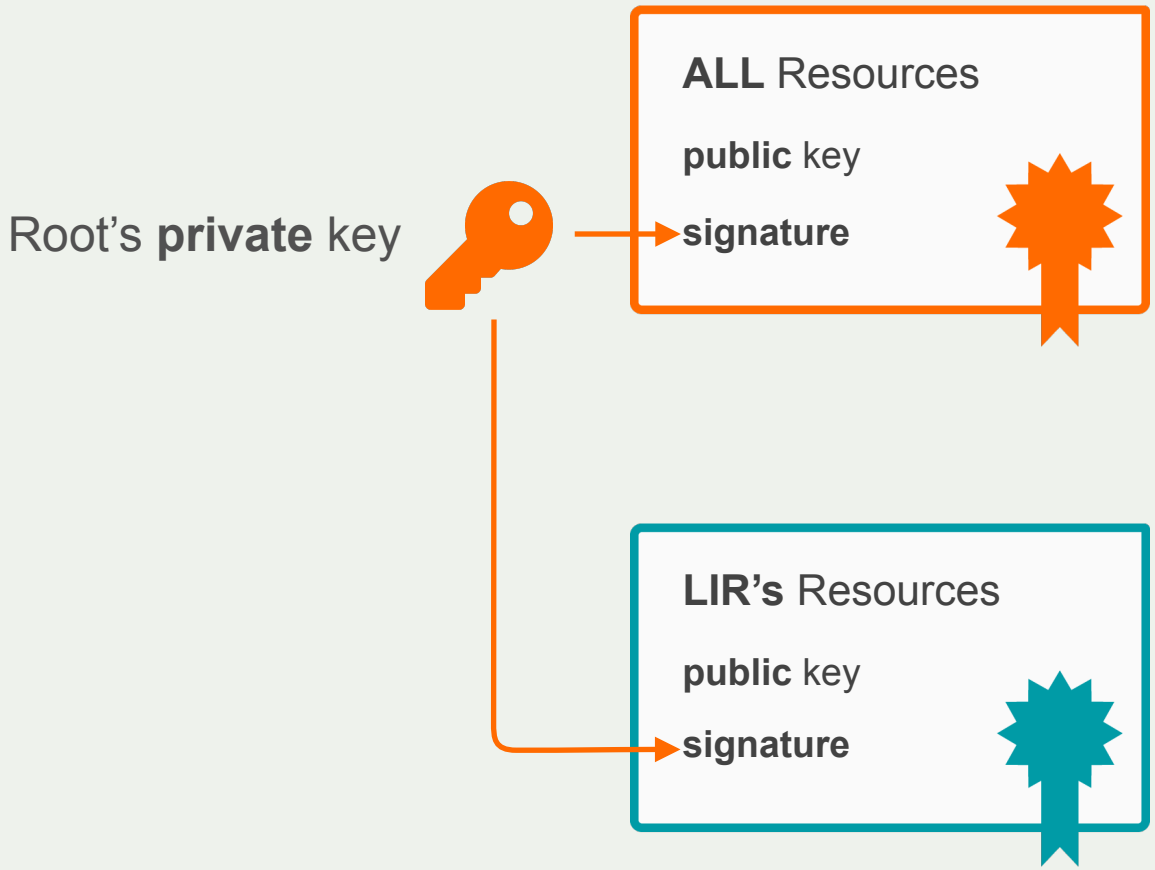

Two elements of RPKI

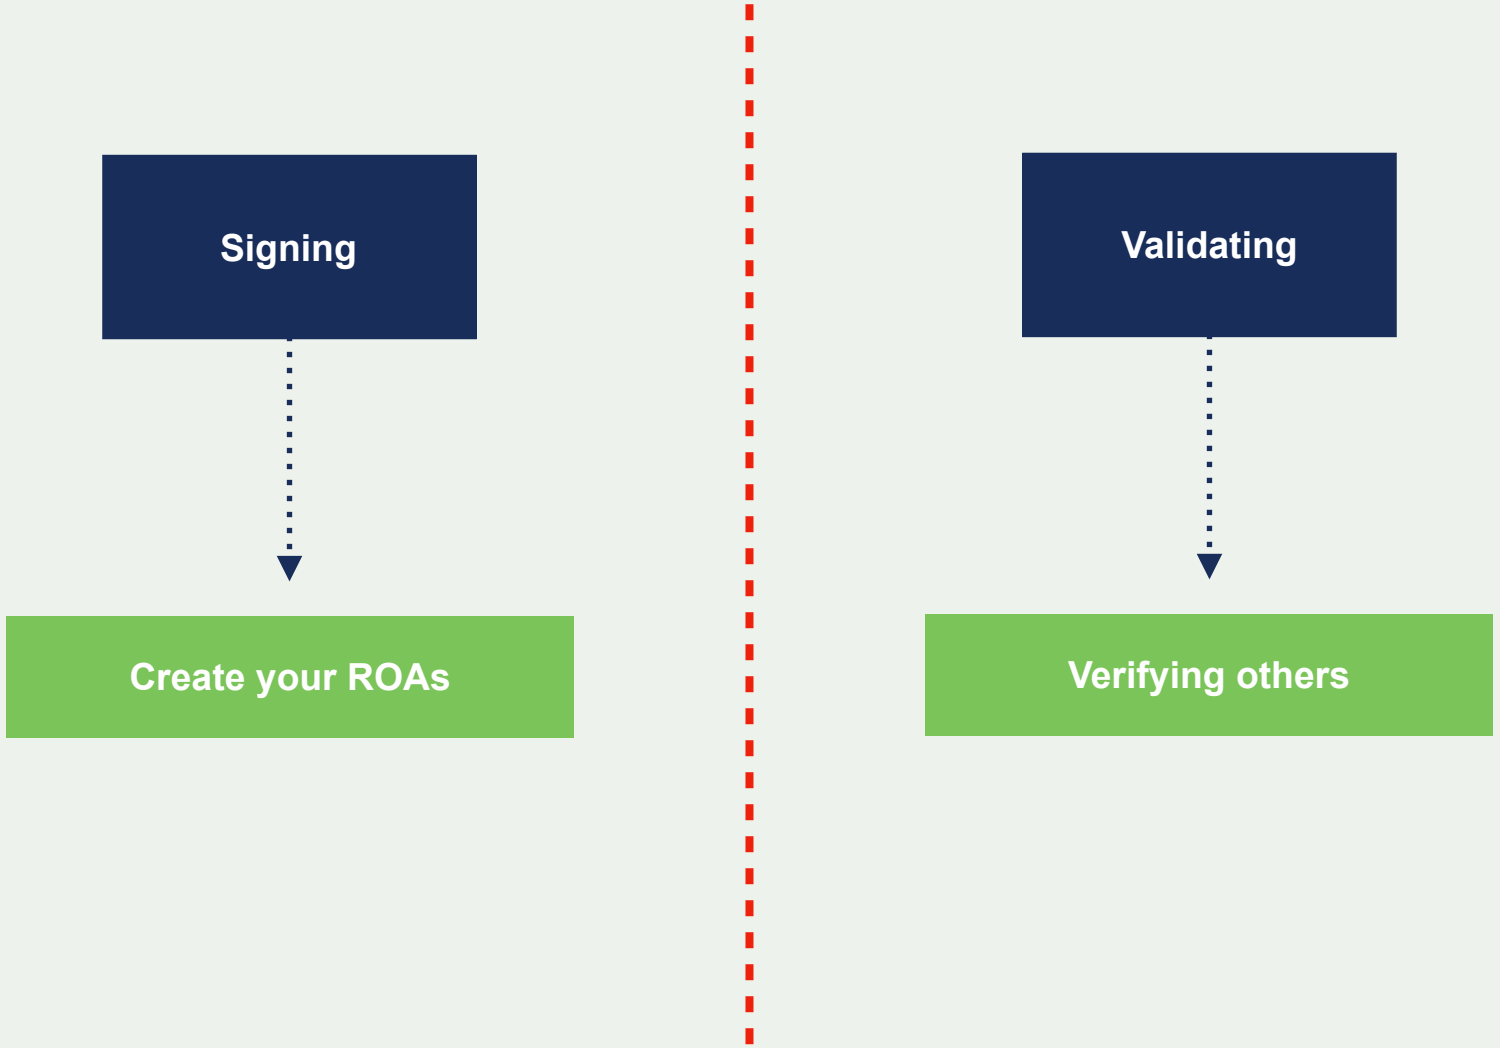

ROA (Route Origin Authorisation)

- LIRs can create a ROA for each one of their resources (IP address ranges)
- Multiple ROAs can be created for an IP range
- ROAs can overlap

What is in a ROA ?

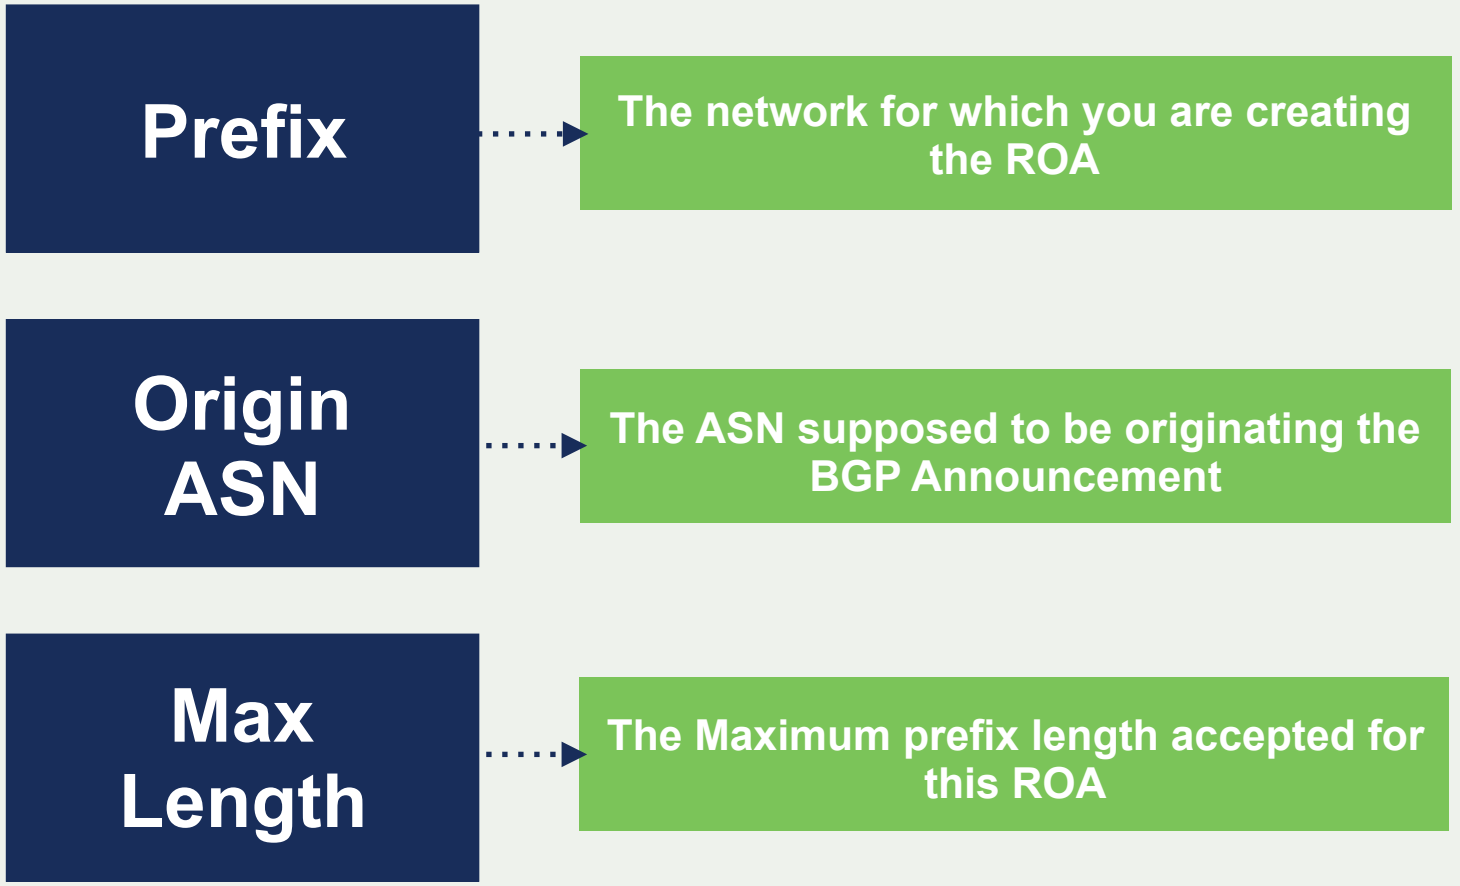

Route Origin Authorisation

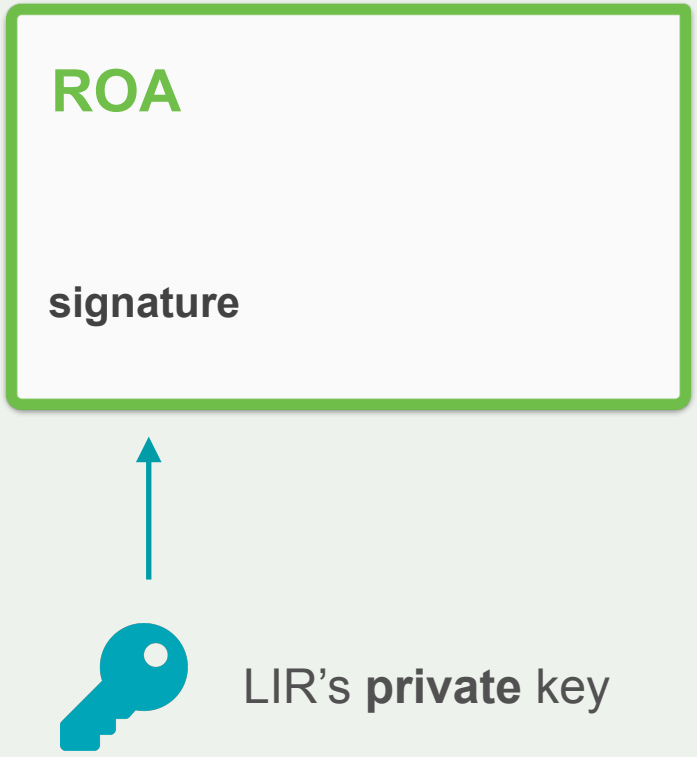

Prefix

is authorised to be announced by

AS Number

RPKI Chain of Trust

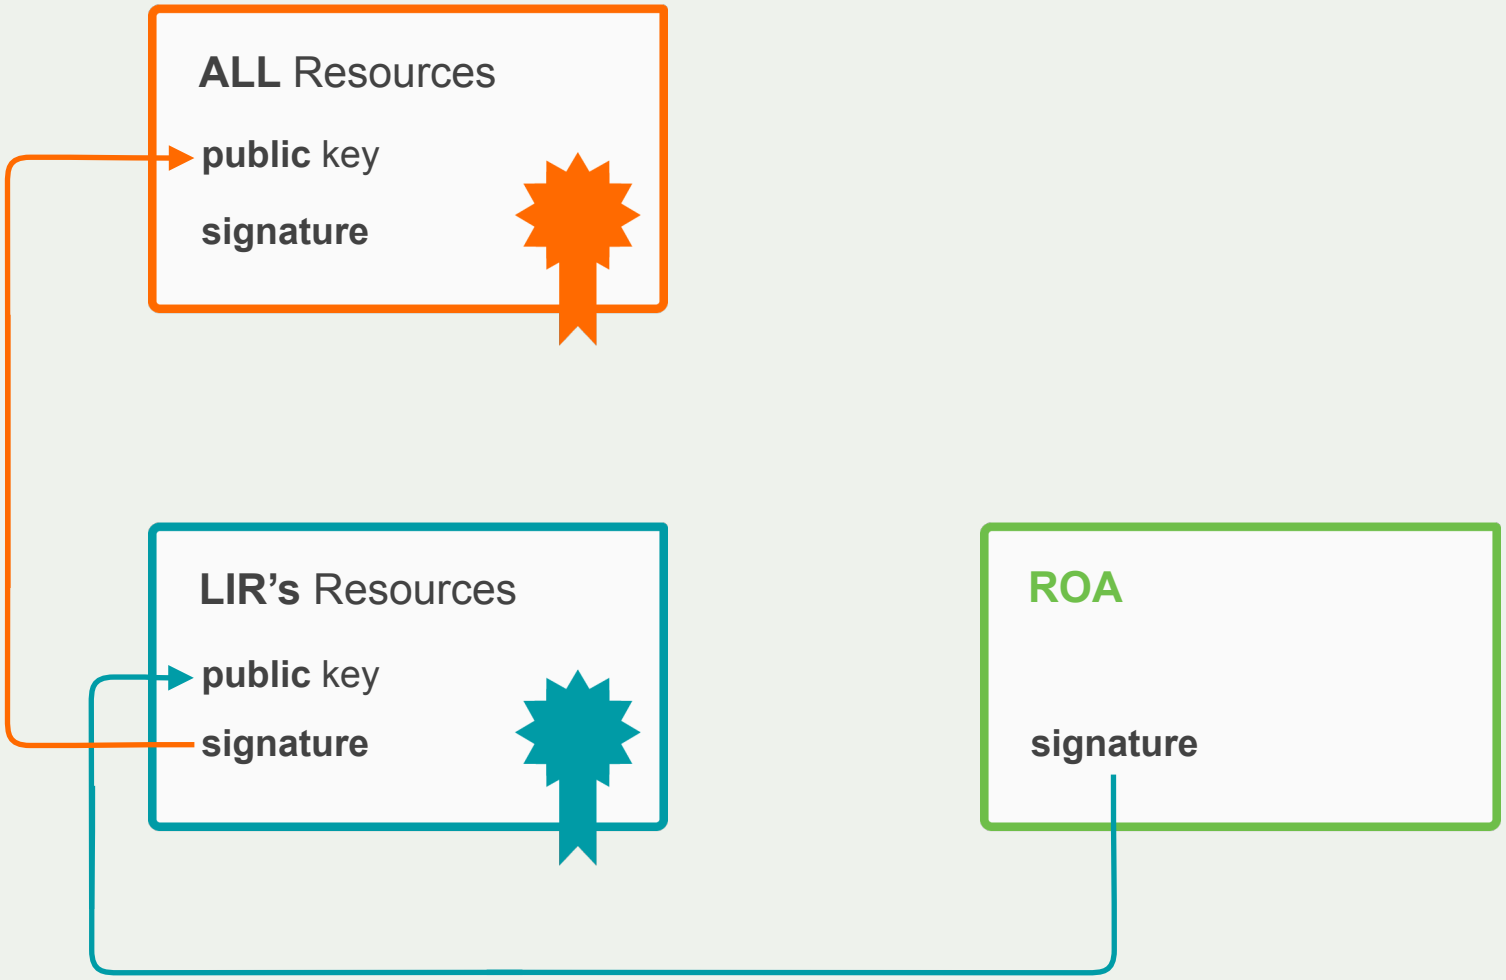

Hosted RPKI

- Automate signing and key roll overs
  - One click setup of resource certificate
  - User has a valid and published certificate for as long as they are the holder of the resources
  - All the complexity is handled by the hosted system

- Lets you focus on creating and publishing ROAs
  - Match your intended BGP configuration

Non-hosted RPKI

- Run your own Certificate Authority

- With your own software

- At the moment, **not advised,** because of lack of software and options
  - Krill is _almost_ there

First login to the dashboard

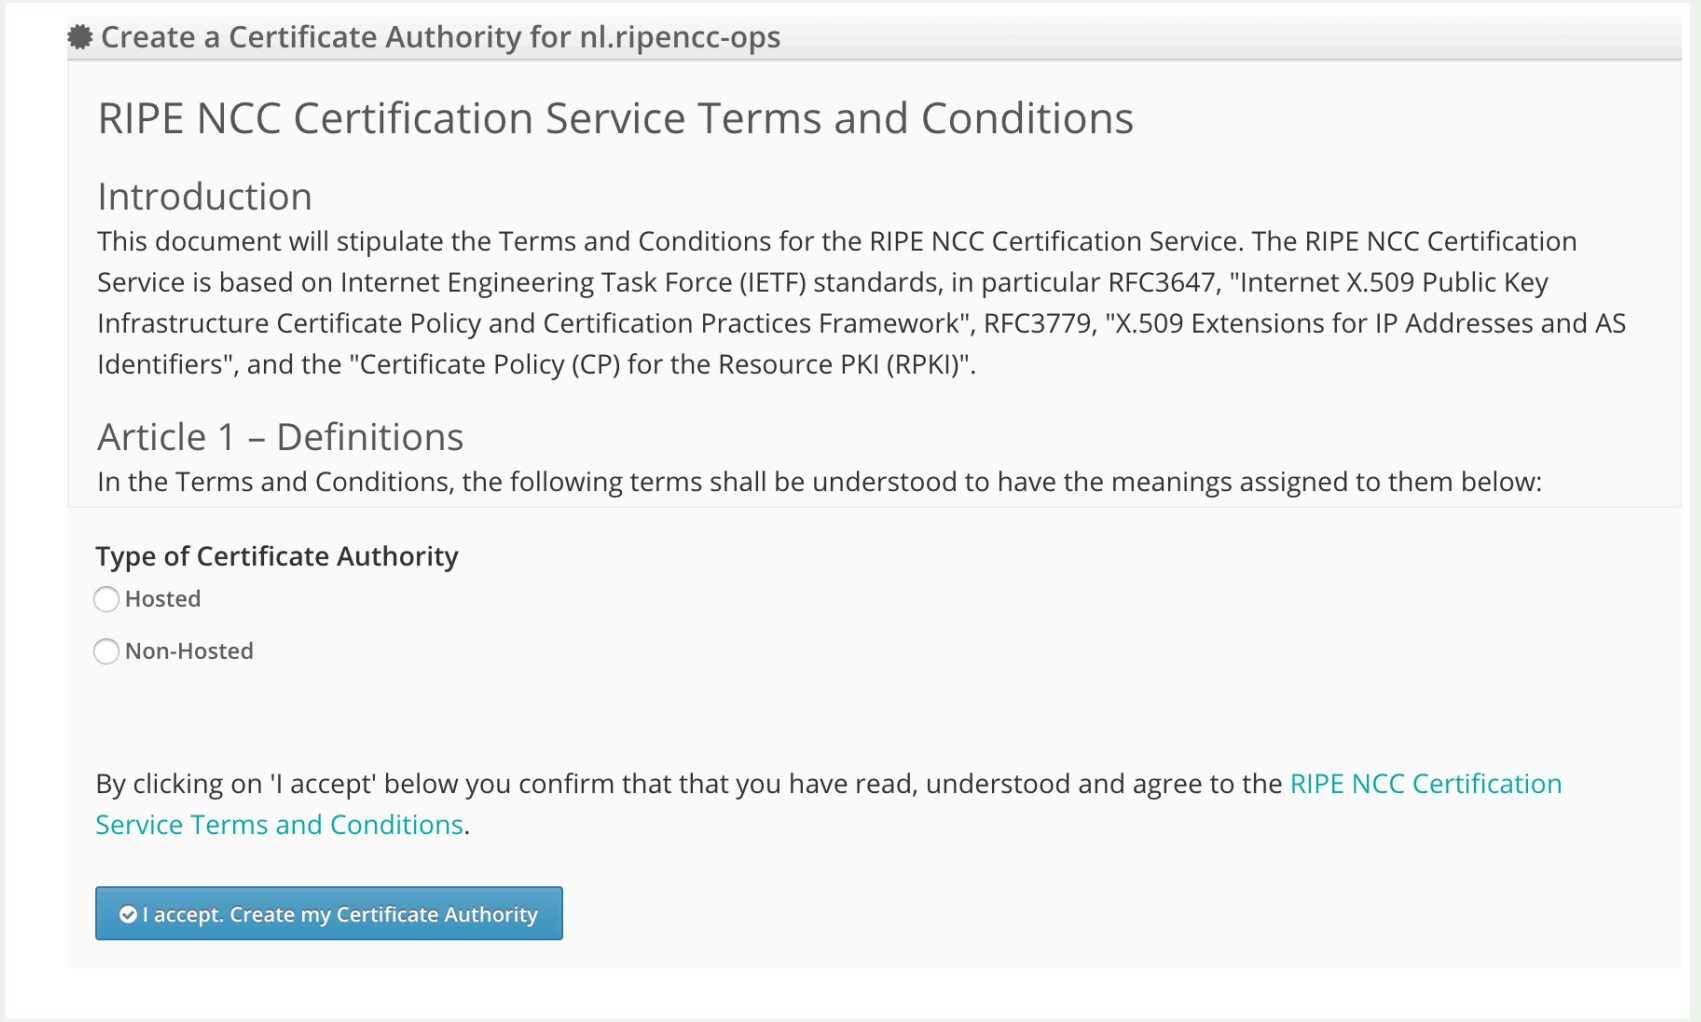

Creating ROAs

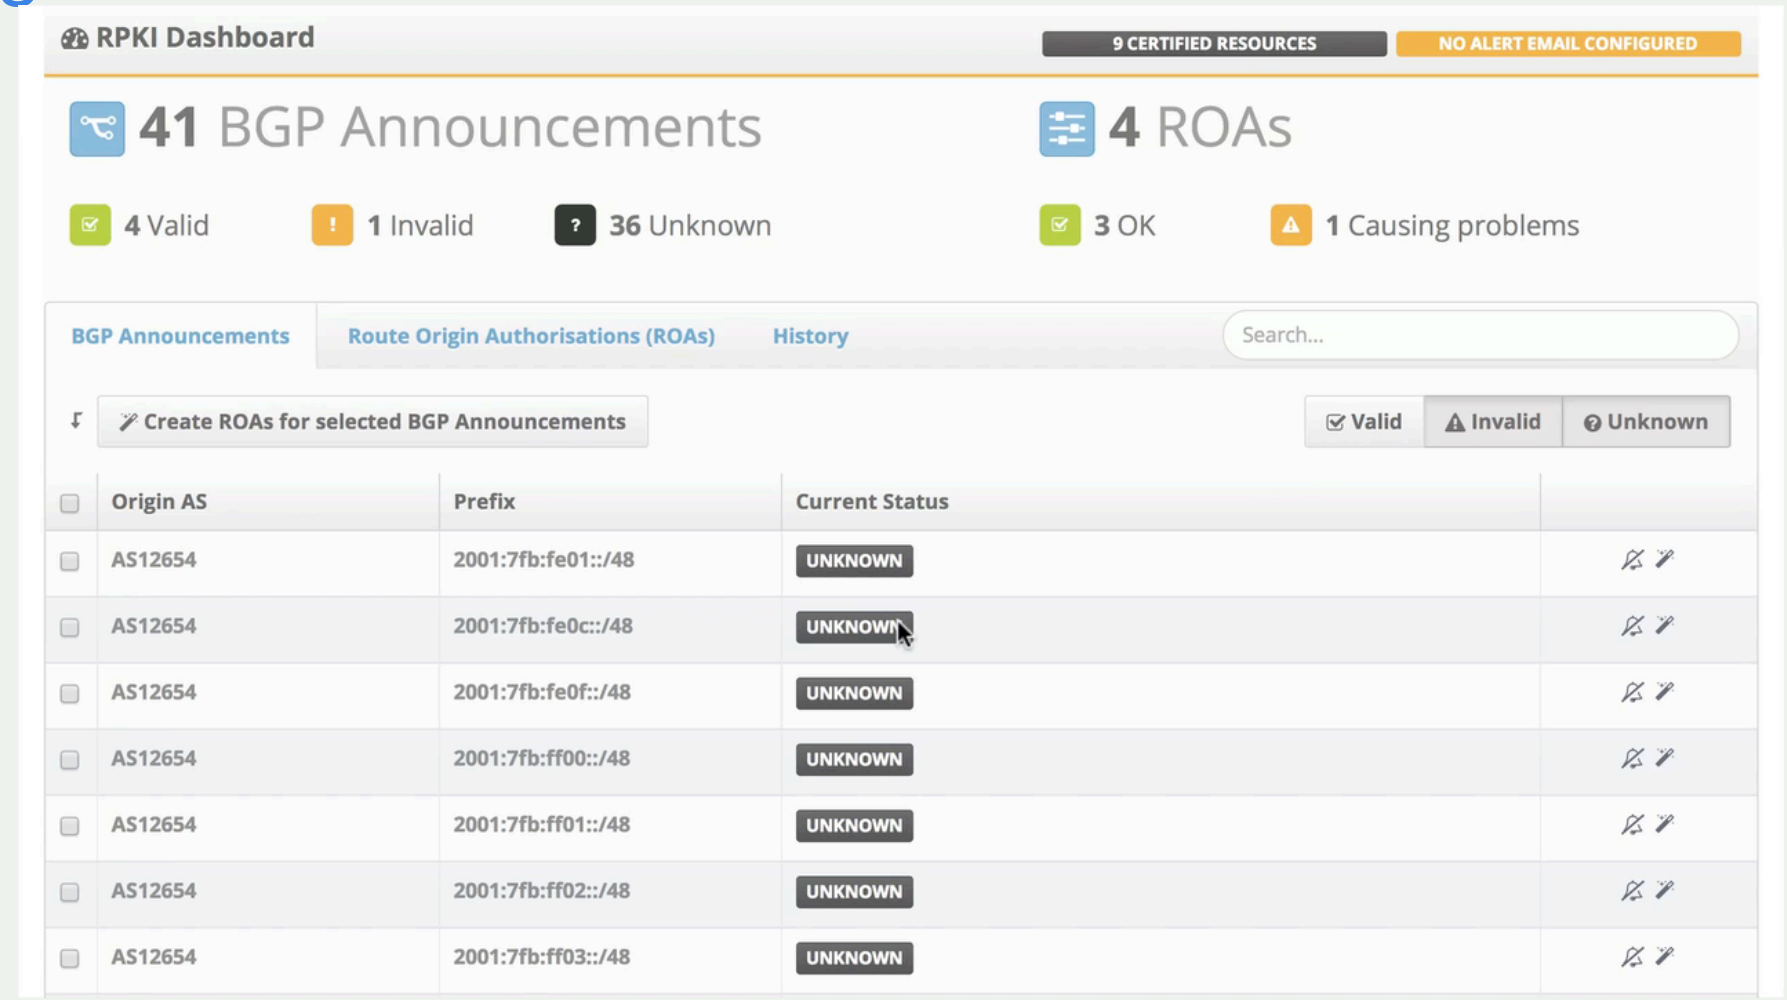

Reviewing changes

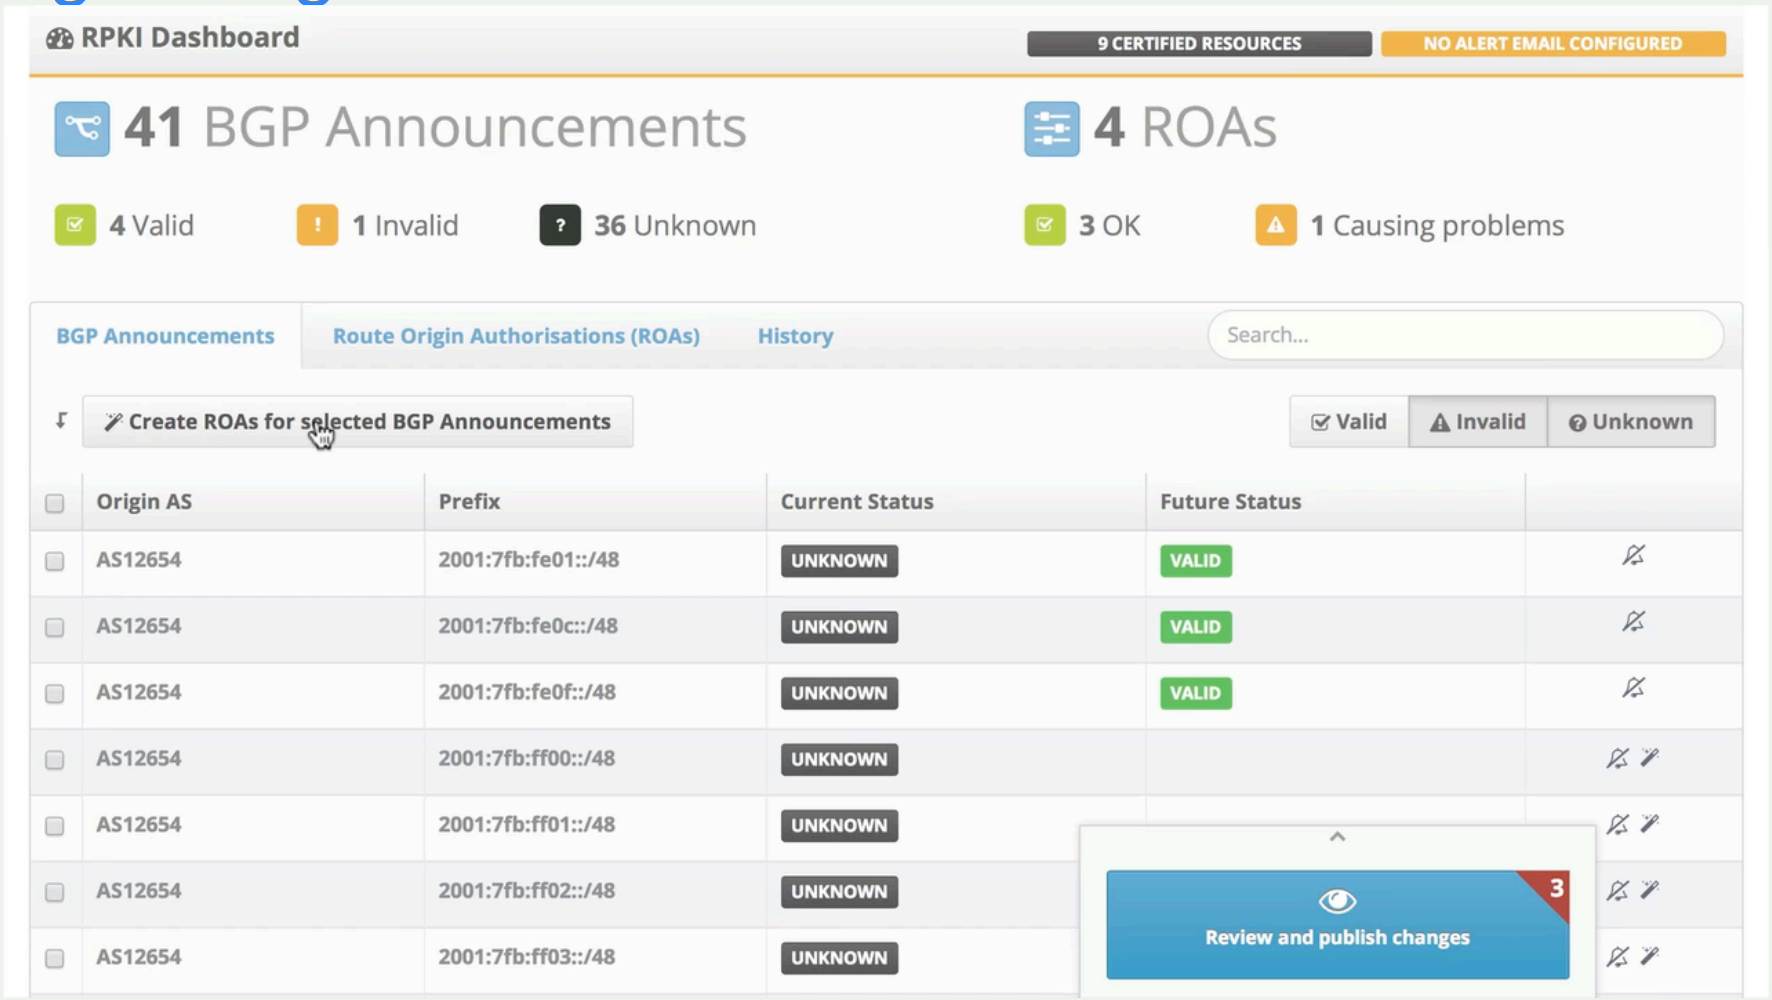

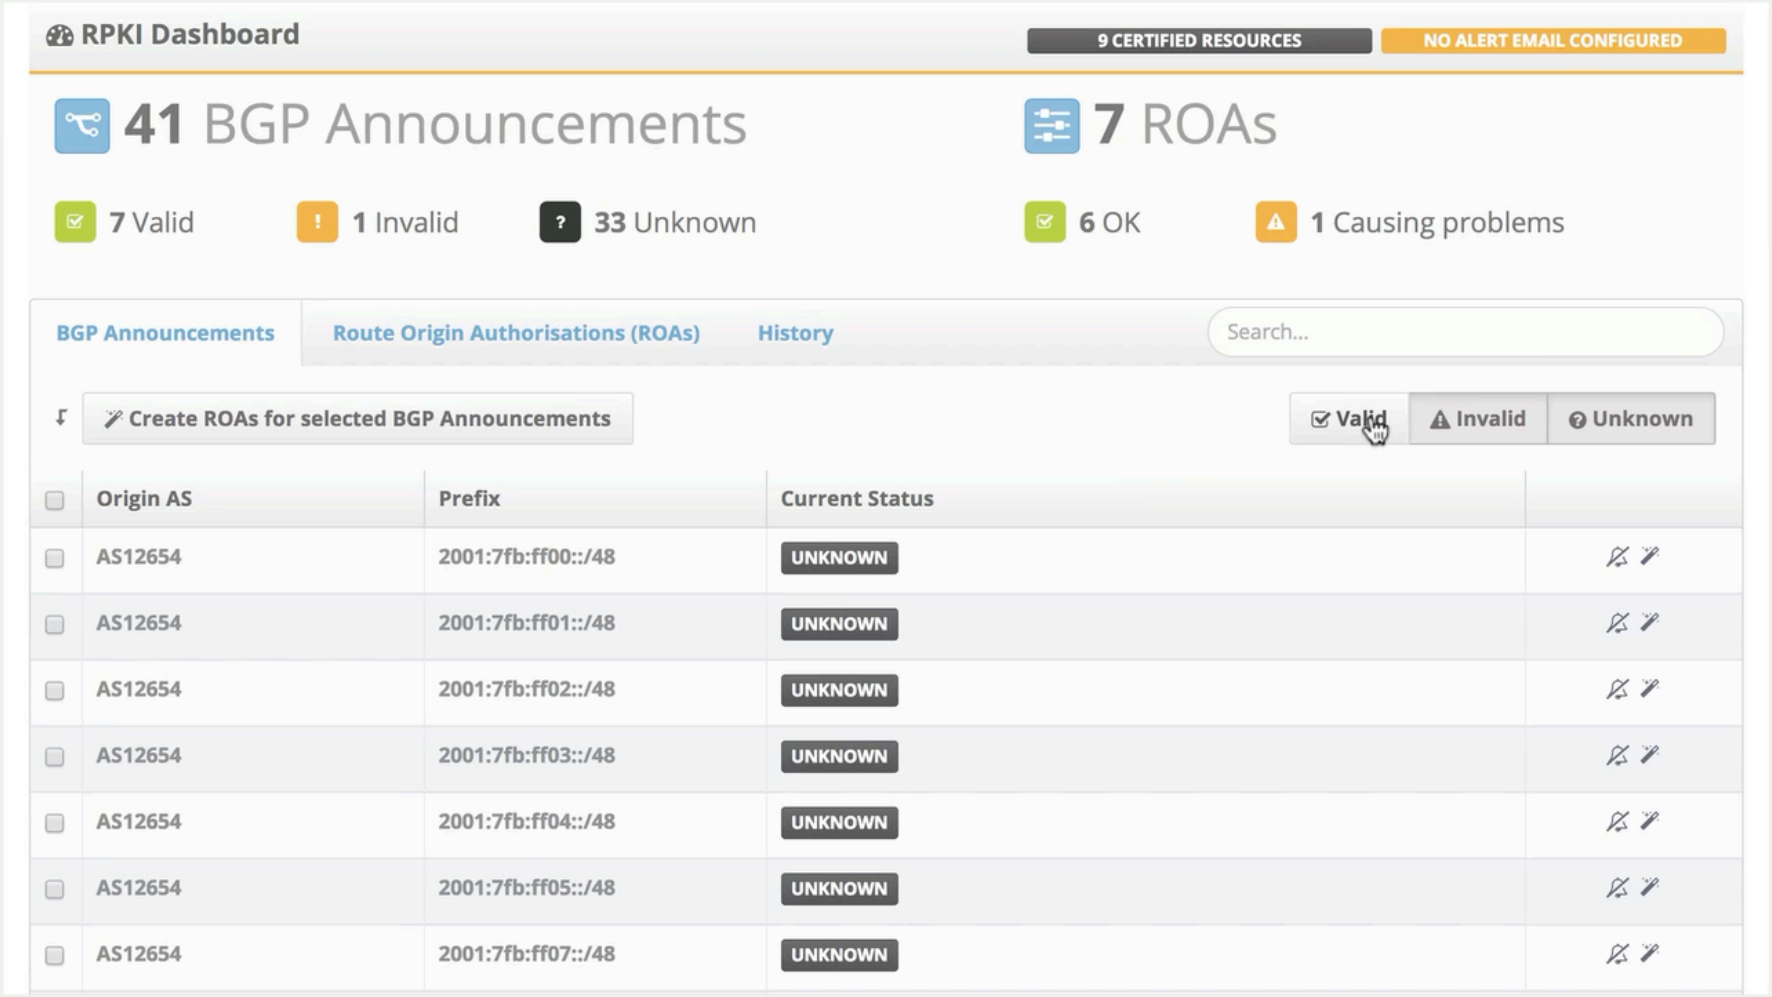
RPKI Dashboard
9 CERTIFIED RESOURCES
NO ALERT EMAIL CONFIGURED
41 BGP Announcements
7 Valid
1 Invalid
33 Unknown
7 ROAs
6 OK
1 Causing problems
BGP Announcements
Route Origin Authorisations (ROAs)
History
Search...
Create ROAs for selected BGP Announcements
Valid
Invalid
Unknown
Origin AS
Prefix
Current Status
AS12654 2001:7fb:ff00::/48 UNKNOWN
AS12654 2001:7fb:ff01::/48 UNKNOWN
AS12654 2001:7fb:ff02::/48 UNKNOWN
AS12654 2001:7fb:ff03::/48 UNKNOWN
AS12654 2001:7fb:ff04::/48 UNKNOWN
AS12654 2001:7fb:ff05::/48 UNKNOWN
AS12654 2001:7fb:ff07::/48 UNKNOWN

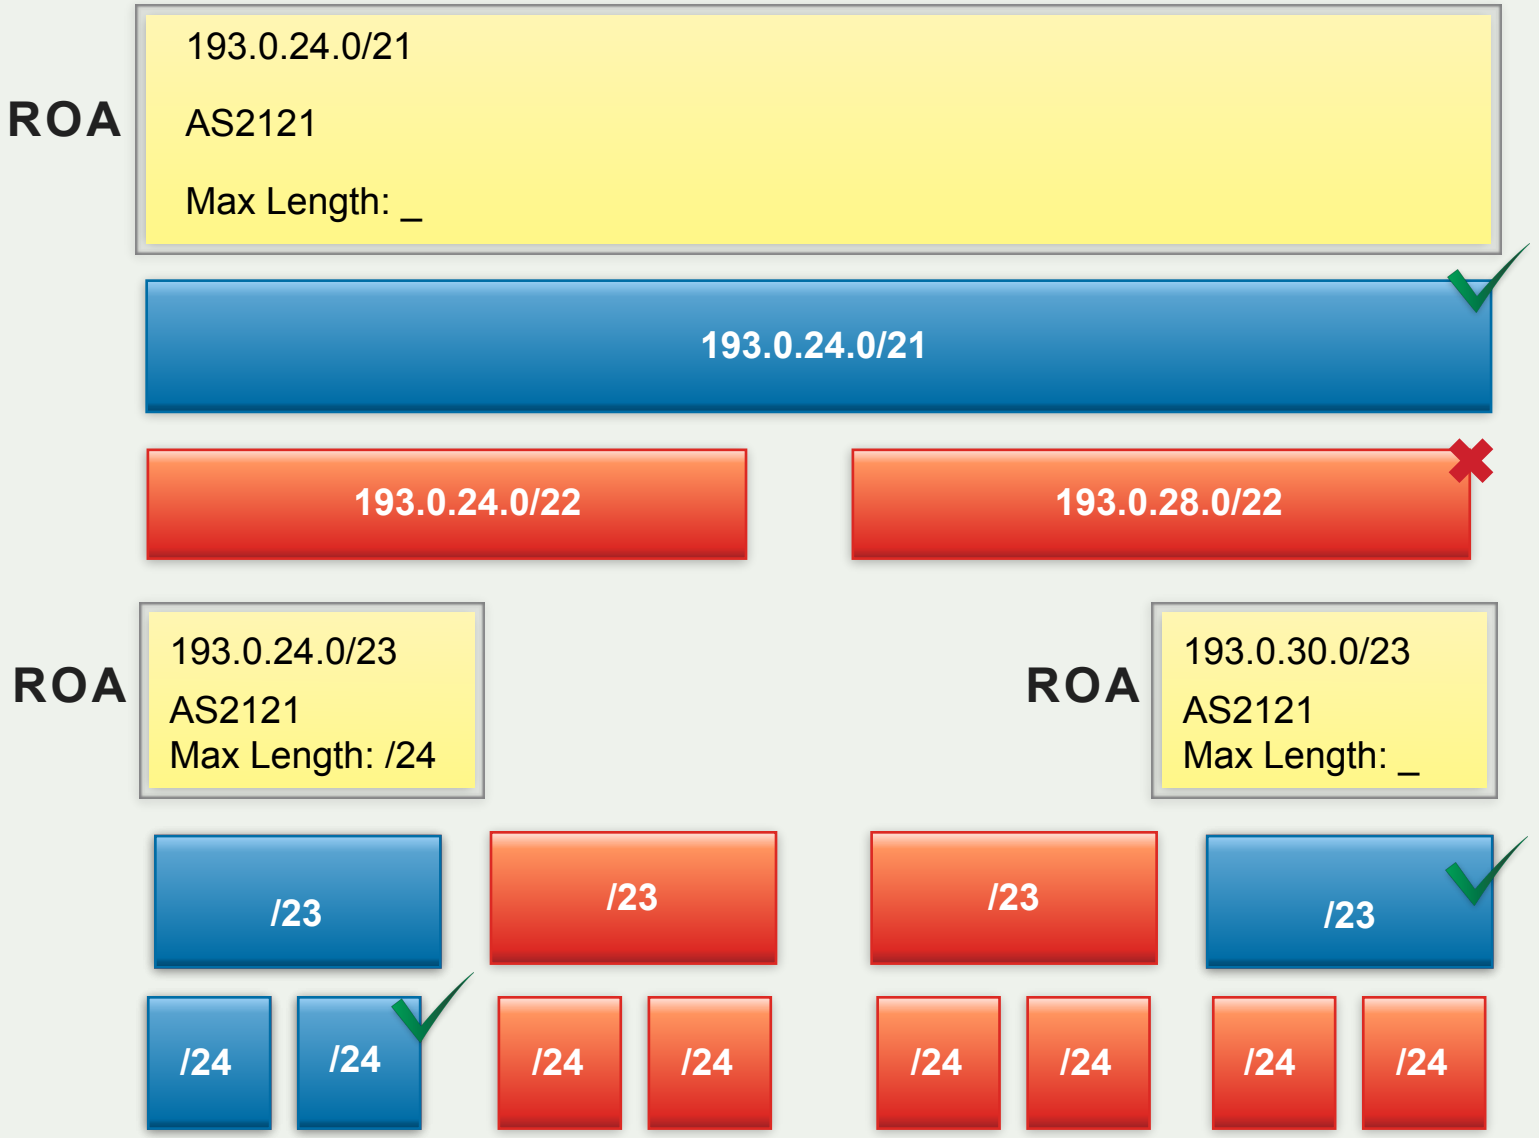
ROA
193.0.24.0/21
AS2121
Max Length: _
193.0.24.0/21
193.0.24.0/22
193.0.28.0/22
ROA
193.0.24.0/23
AS2121
Max Length: /24
ROA
193.0.30.0/23
AS2121
Max Length: _
/23
/23
/23
/23
/24
/24
/24
/24
/24
/24
/24
/24

Relying Party

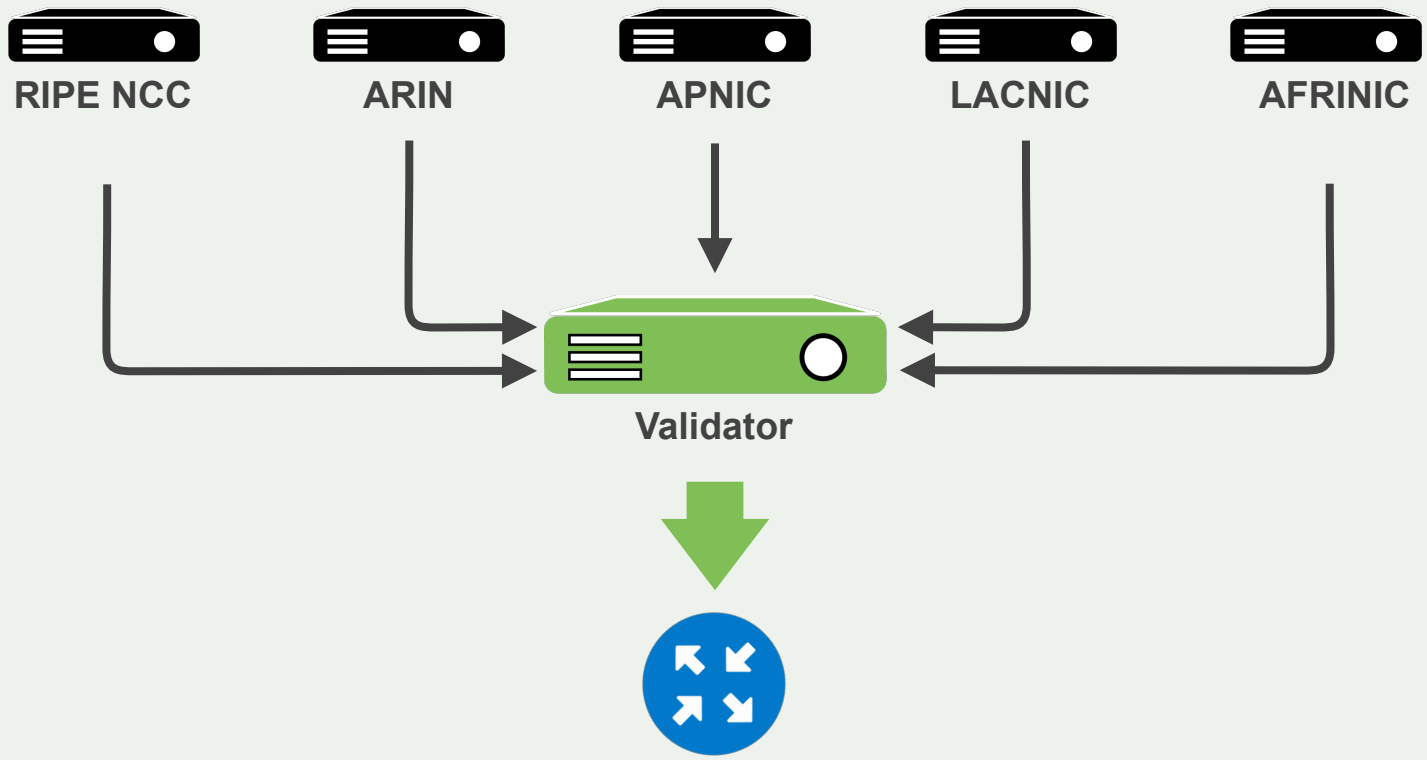

RPKI-RTR

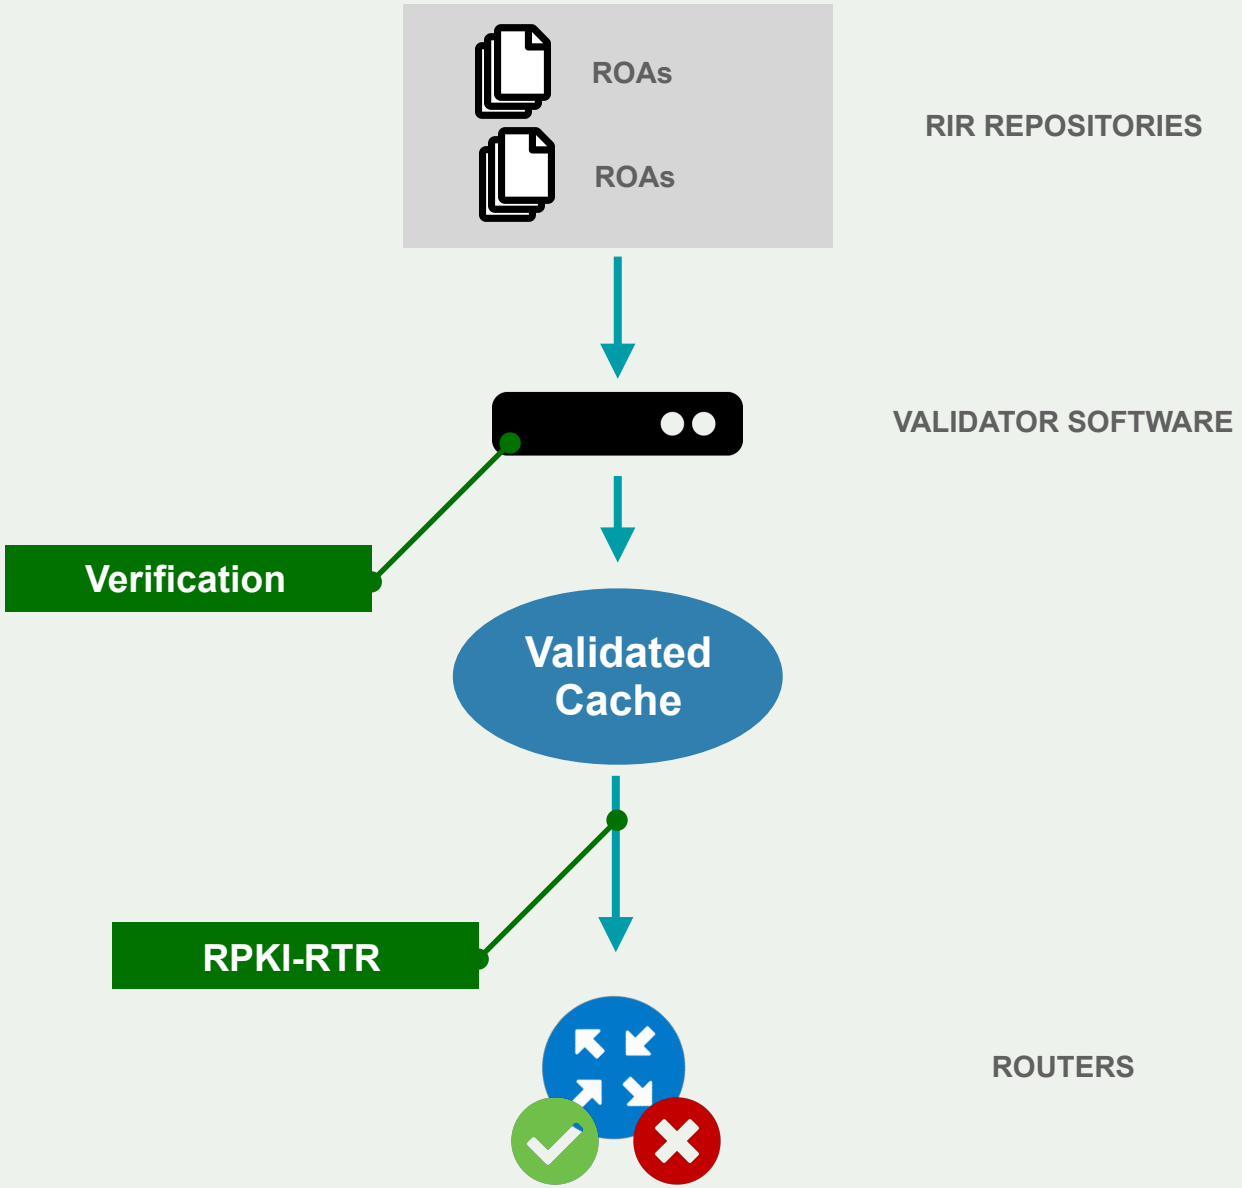
ROAs
ROAs
RIR REPOSITORIES
VALIDATOR SOFTWARE
Verification
Validated Cache
RPKI-RTR
ROUTERS

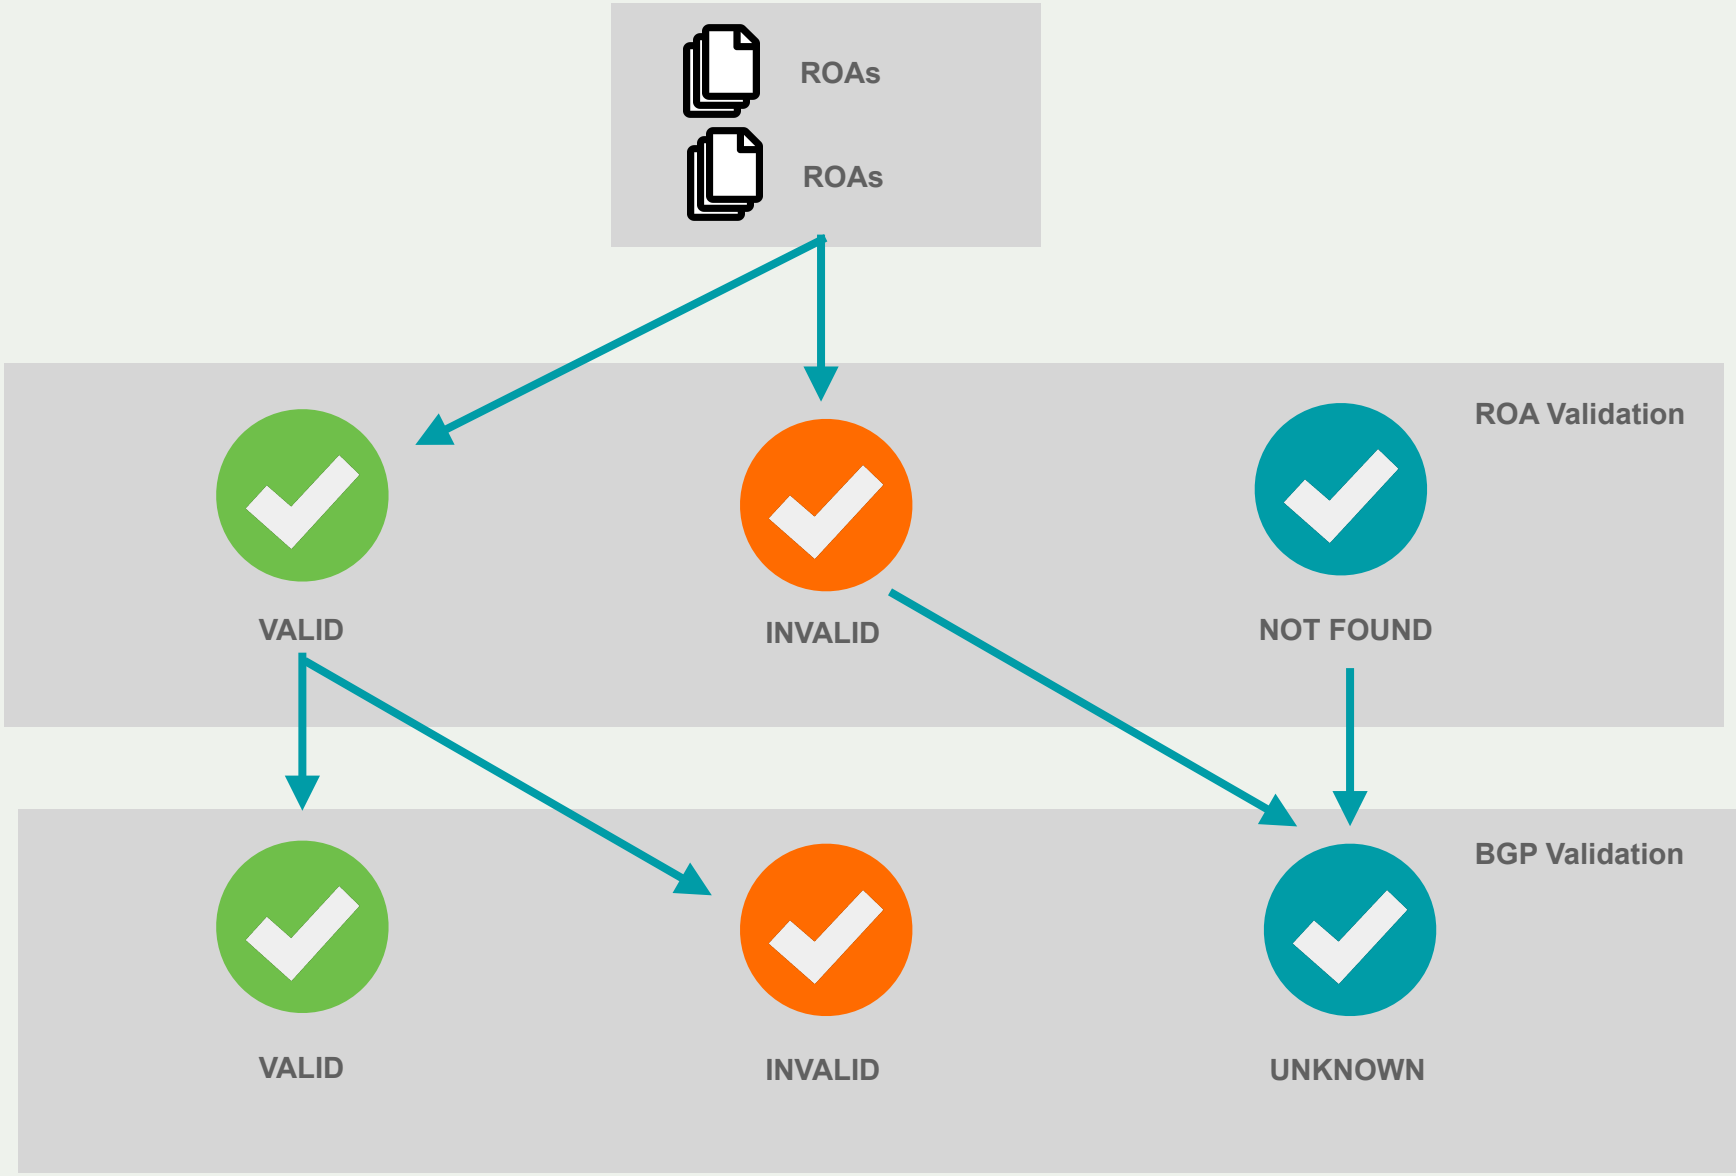
ROAs
ROAs
ROA Validation
VALID
INVALID
NOT FOUND
BGP Validation
VALID
INVALID
UNKNOWN

Validator Software

- RIPE NCC Validator
- NLNetLabs Routinator
- Cloudflare OctoRPKI
- NIC MX Fort

Cisco Origin Validation configuration

```
(config)# conf t
(config)# router bgp $ASN
(config-router)# bgp rpki server tcp 100.64.1.1 port 8323 refresh 300
(config-router)# bgp rpki server tcp 100.64.1.1 port 3323 refresh 300
```

Cisco Origin Validation configuration

```
(config-router)# route-map rpki-accept permit 10
(route-map)# match rpki valid
(route-map)# set local-preference 100
(route-map)# route-map rpki-accept permit 20
(route-map)# match rpki not-found
(route-map)# set local-preference 80
```

Cisco Origin Validation configuration

```
(config)# router bgp $ASN
(config)# address-family ipv4
(config)# neighbor 192.168.1.254 route-map rpki-accept in
(config)# address-family ipv6
(config)# neighbor 2002:eeee:ffff::a route-map rpki-accept in
```

Juniper Origin Validation configuration

```
routing-options {
    autonomous-system 64511;
    validation {
        group rpki-validator {
            session 100.64.1.1 {
                refresh-time 120;
                hold-time 180;
                port 8282;
                local-address 100.64.1.2;
}}}}
```

Juniper Origin Validation configuration

```
policy-statement send-direct {
  from protocol direct;
  then accept;}
policy-statement validation {
  term valid {
    from {
      protocol bgp;
      validation-database valid; }
    then {
      local-preference 110;
      validation-state valid;
      community add origin-validation-state-valid;
      accept;
      }}
```

Juniper Origin Validation configuration

```
term invalid {
   from {
      protocol bgp;
      validation-database invalid;}
   then {
      local-preference 90;
      validation-state invalid;
      community add origin-validation-state-invalid;
      accept;
}}
```

Juniper Origin Validation configuration

```
term unknown {
   from protocol bgp;
   then {
   validation-state unknown;
   community add origin-validation-state-unknown;
   accept;
}}
```

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)

BGPSec

- RPKI does not protect against path redirection attacks

- We need a way to verify the AS-Path of a given BGP Announcement
  - And understand if anyone tampered with the data on the way to our routers

- With BGPSec, the AS-Path attribute is cryptographically signed
  - Using the operator’s certificate from RPKI

- In order to validate an AS-Path, routers verify the chain of trust of all the signatures of the AS-Path

Origin Validation Check

- Go with your browser to

http://www.ripe.net/s/rpki-test

- And check if your network applies Origin Validation

Wrapping up

MANRS Training Tutorials

6 training tutorials based on information in the Implementation Guide.

A test at the end of each tutorial.

About to begin training moderators for online classes (43 applications received!)

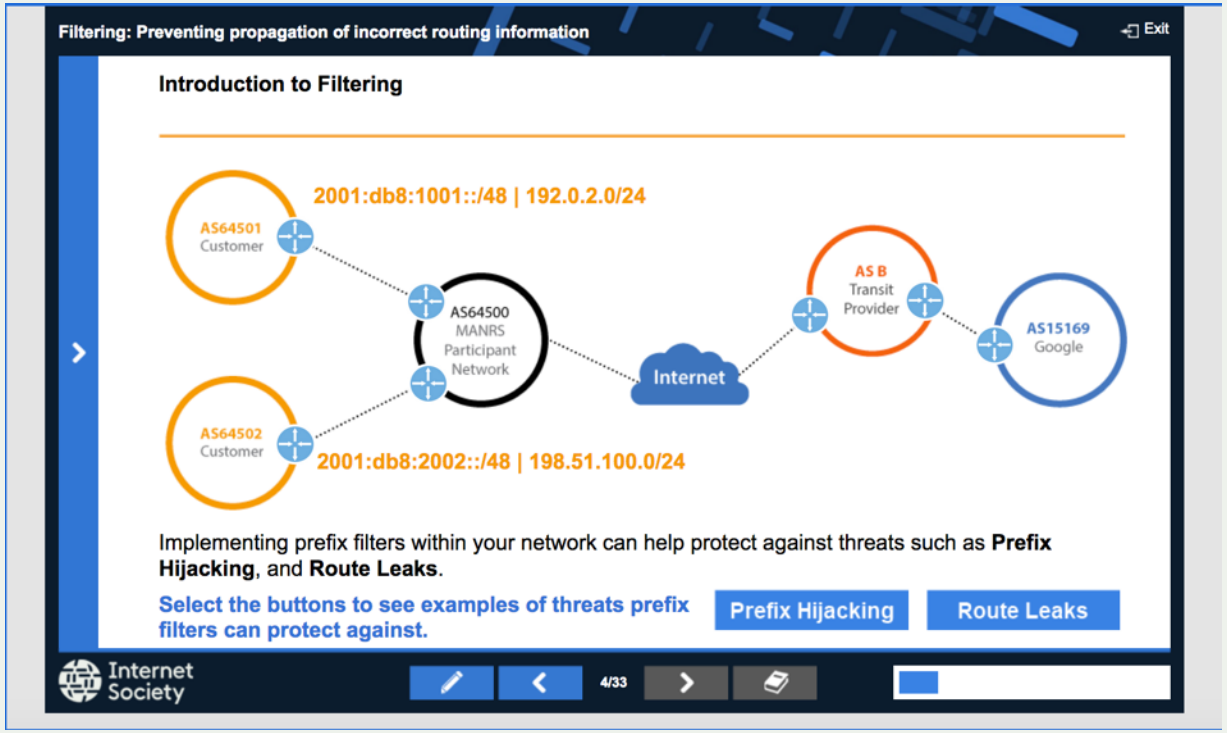

https://www.manrs.org/tutorials

MANRS Hands-on Lab

The prototype lab is ready, finalising the production version.

- Cisco
- Juniper
- Mikrotik

Can be used as a standalone lab or as a final exam

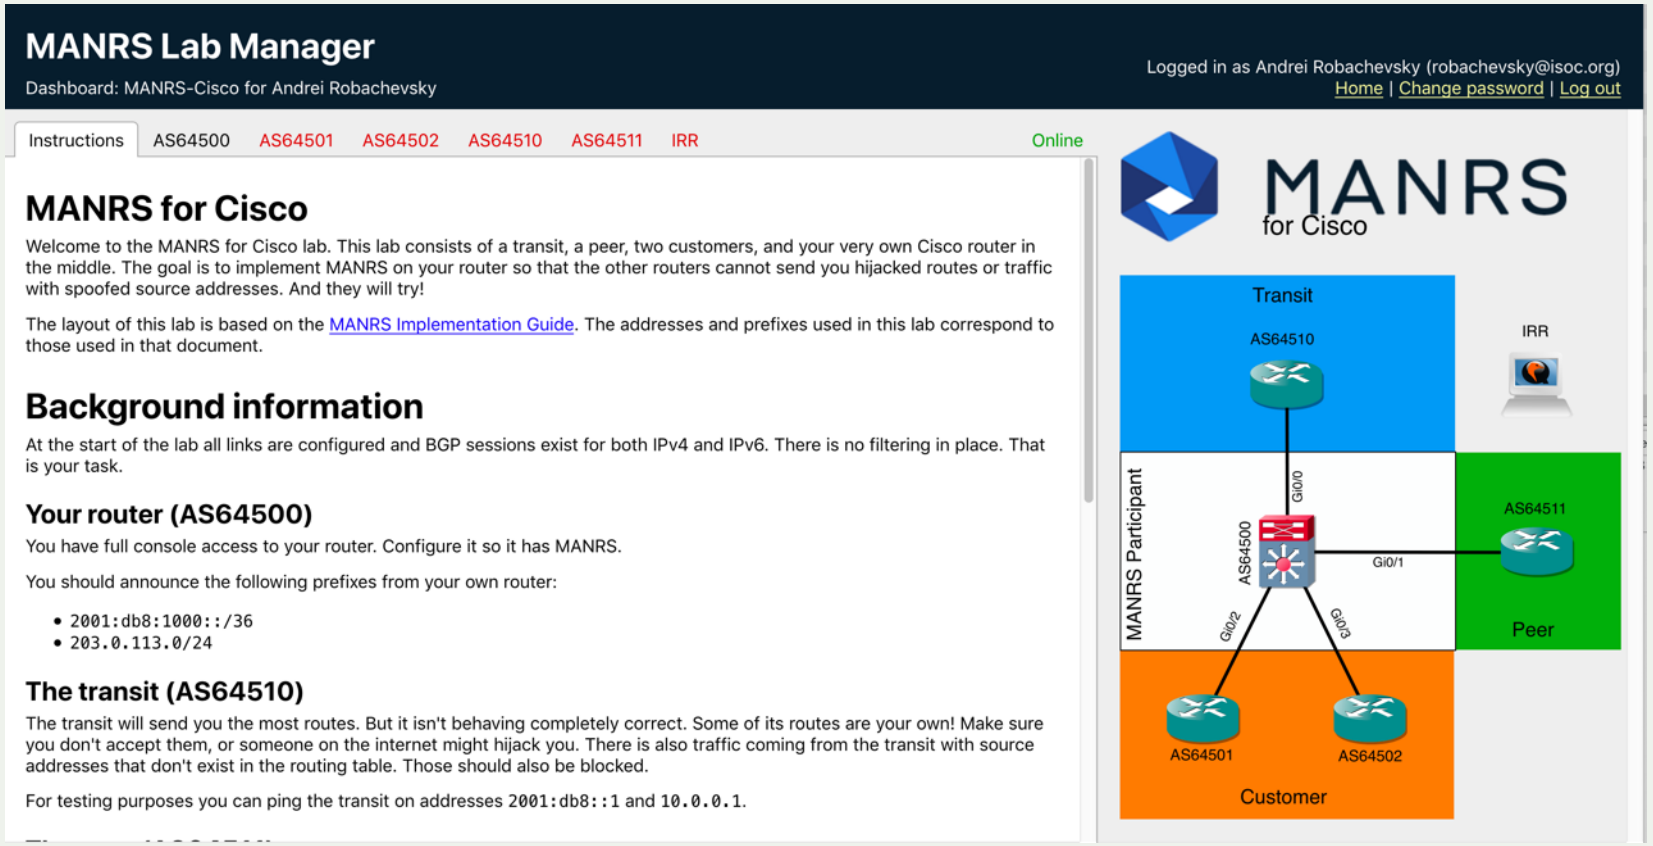

Join MANRS

Visit https://www.manrs.org

- Fill out the sign up form with as much detail as possible.
- We may ask questions and request tests

Get Involved in the Community

- Participants support the initiative and implement the actions in their own networks and encouraging MANRS adoption
- Participants are engaged in substantive activities – developing MANRS requirements and guidance, assisting with capacity and awareness building activities

Questions ?

stucchi@isoc.org
@stucchimax