ARouteServer

IXP automation made easy

Intro: route servers

- “A route server redistributes BGP routes received from its BGP clients to other clients according to a prespecified policy” (RFC7948)

- Benefits:
  - One BGP session, many peers;
  - Instantly get lot of routes.

- Drawbacks:
  - May be a SPoF (traffic blackhole);
  - May interfer with traffic engineering policies;
  - Outsourced routing / To trust or not to trust the route server.

Intro: route servers

- Lot of routes sent to clients: are they actually good?

```
inet.0: 630929 destinations, 967236 routes (630926 active, 1 holddown, 32 hidden)
  Prefix   Nexthop       MED     Lclpref    AS path
  37.18.122.0/23          80.249.212.148                  58291 174 42755 58272 I
  62.93.195.0/24          80.249.208.212                  13237 200093 1299 174 20756 I
  81.24.160.0/20          80.249.212.81                   15704 12321 3356 6739 20973 24799 I
```

source: “Moving to default Routeserver IRR filtering”

RIPE74 presentation by Erik Bais, A2B Internet

(Tier-1 ASNs in the middle of the AS_PATH)

Intro: route servers

- Lot of routes sent to clients: are they actually good?

```
4651 45328 23456
  80.249.210.15 from 80.249.209.0 (80.249.209.0)
    Origin IGP, metric 0, localpref 100, valid, external, best
    Community: 3856:53500
    Last update: Mon May  1 08:07:12 2017
```

(The reserved AS_TRANS ASN as origin)

Intro: route servers

- Lot of routes sent to clients: are they actually good?

```
12989 33724 64512
  80.249.209.70 from 80.249.208.255 (80.249.208.255)
    Origin IGP, metric 0, localpref 100, valid, external
    Community: 3856:53500
    Last update: Mon May  1 08:10:10 2017
```

(A private ASN as an Origin AS)

Intro: route servers

- Lot of routes sent to clients: are they actually good?

IXP megabit/sec cost & comparison (Job Snijders / Samer Abdel-Hafez)

File Modifica Visualizza Inserisci Formato Dati Strumenti Componenti aggiuntivi Guida

Solo visualizzazione

	B	C	D	E	F	G	H	I	J
1							Inspired by:	Dave Te[illegible]in's "[illegible]P Cost" talk	
2	**IXP Comparison April 2017 - all pricing normalised to Euro/Eurocent/Month**							(higher is m[illegible]sive)	
3	http://peering.exposed/								
4					85% util on 1x100GE	40% util on 1x100GE	85% util on 1x10GE	40% util on 1x10G[illegible]	
5	IXP	100GE price	10GE price	Ratio 100G:10G	Cent/Month/Mbps	Cent/Month/Mbps	Cent/Month/Mbps	Cent/Month/Mbps	Secure Route servers
6	MSK-IX Moscow	7339	1376	5.33	8.60	**18.30**	16.20	**34.40**	Yes
7	DO-IX	5167	1250	4.13	6.10	**12.90**	14.70	**31.30**	?
8	FranceIX Paris	5639	1128	5	6.60	**14.10**	13.30	**28.20**	No
9	LINX Juniper (LON1)	6285	1095	5.74	7.40	**15.70**	12.90	**27.40**	No
10	MIX	6042	1042	5.8	7.10	**15.10**	12.30	**26.10**	Yes
11	ECIX	5278	1028	5.13	6.20	**13.20**	12.10	**25.70**	Yes

http://peering.exposed

(Job Snijders / Samer Abdel-Hafez)

Intro: route servers

- Do I totally lose control of how my routes are propagated?

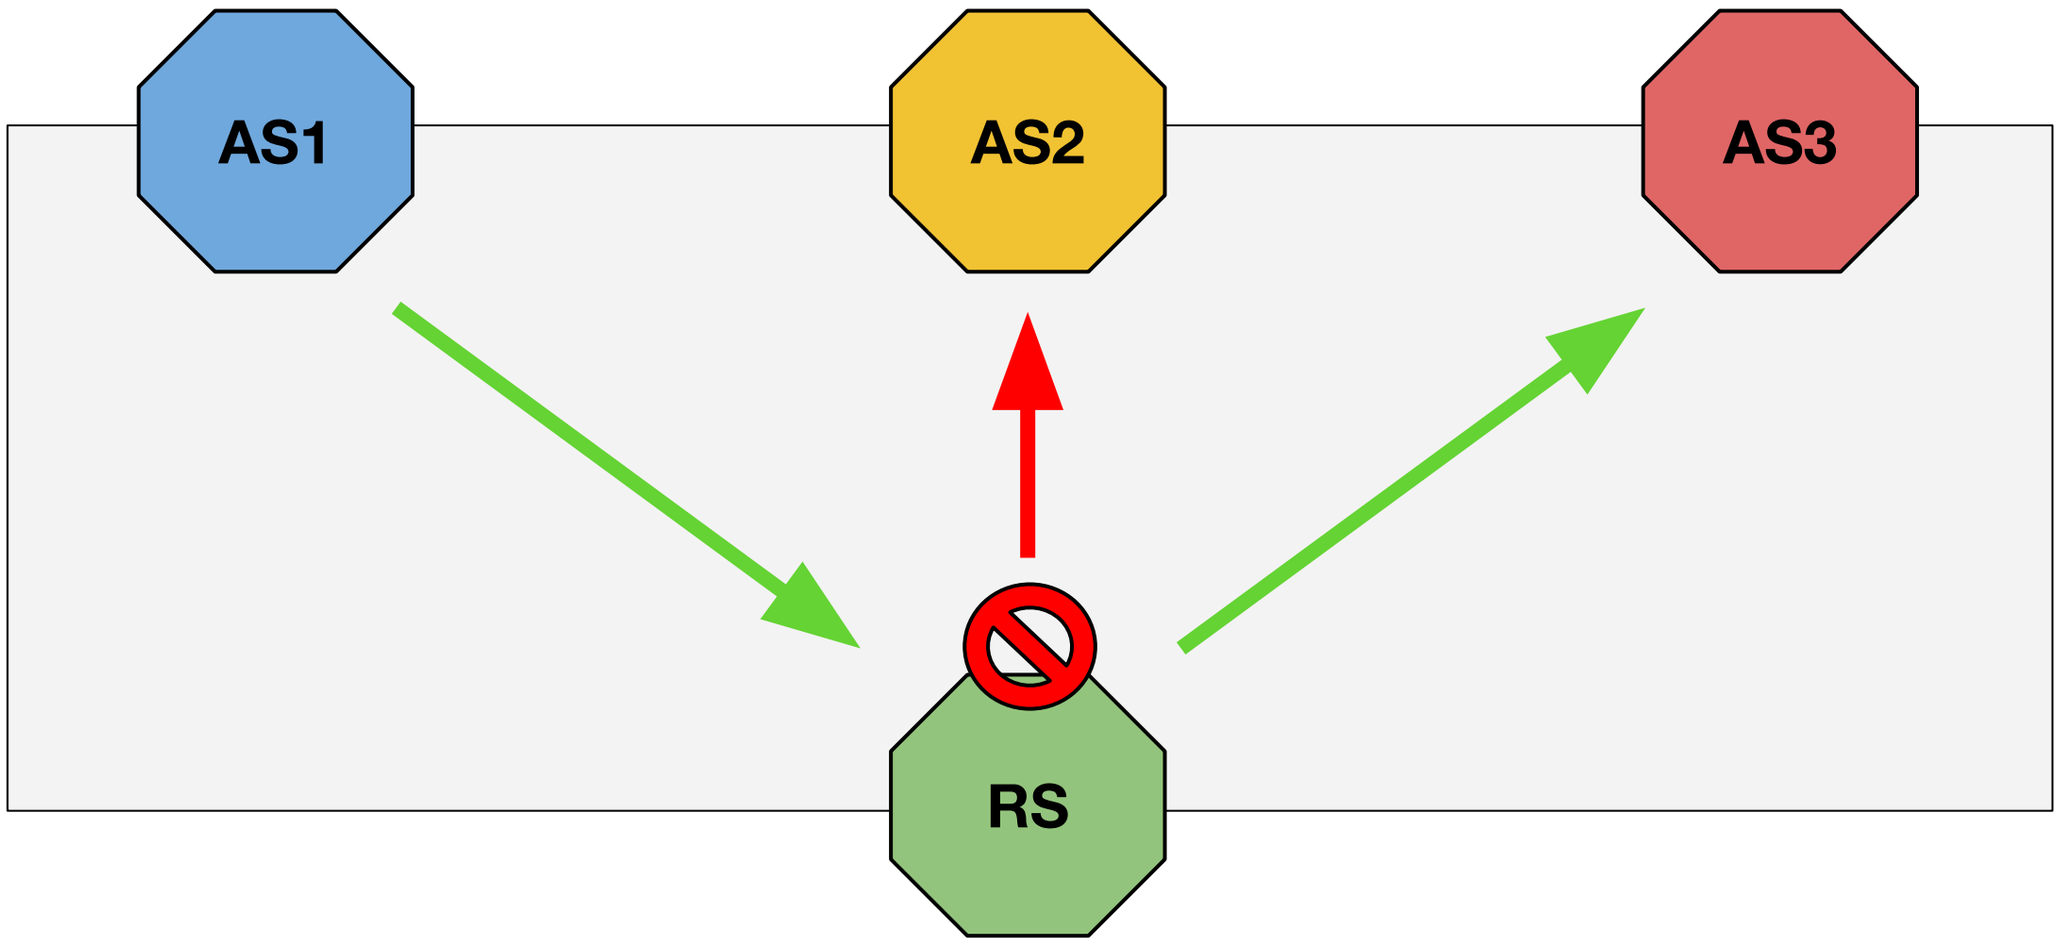

Overview: ARouteServer

- Goal: to help IXs to simplify operations needed to run useful and trustworthy route servers.
- ARouteServer automatically builds configs for route servers...
  - secure and feature-rich configurations!
- … and it also allows to validate them
  - Docker / KVM -based “live test” framework.
- Currently, BIRD and OpenBGPD are supported.

How does it work

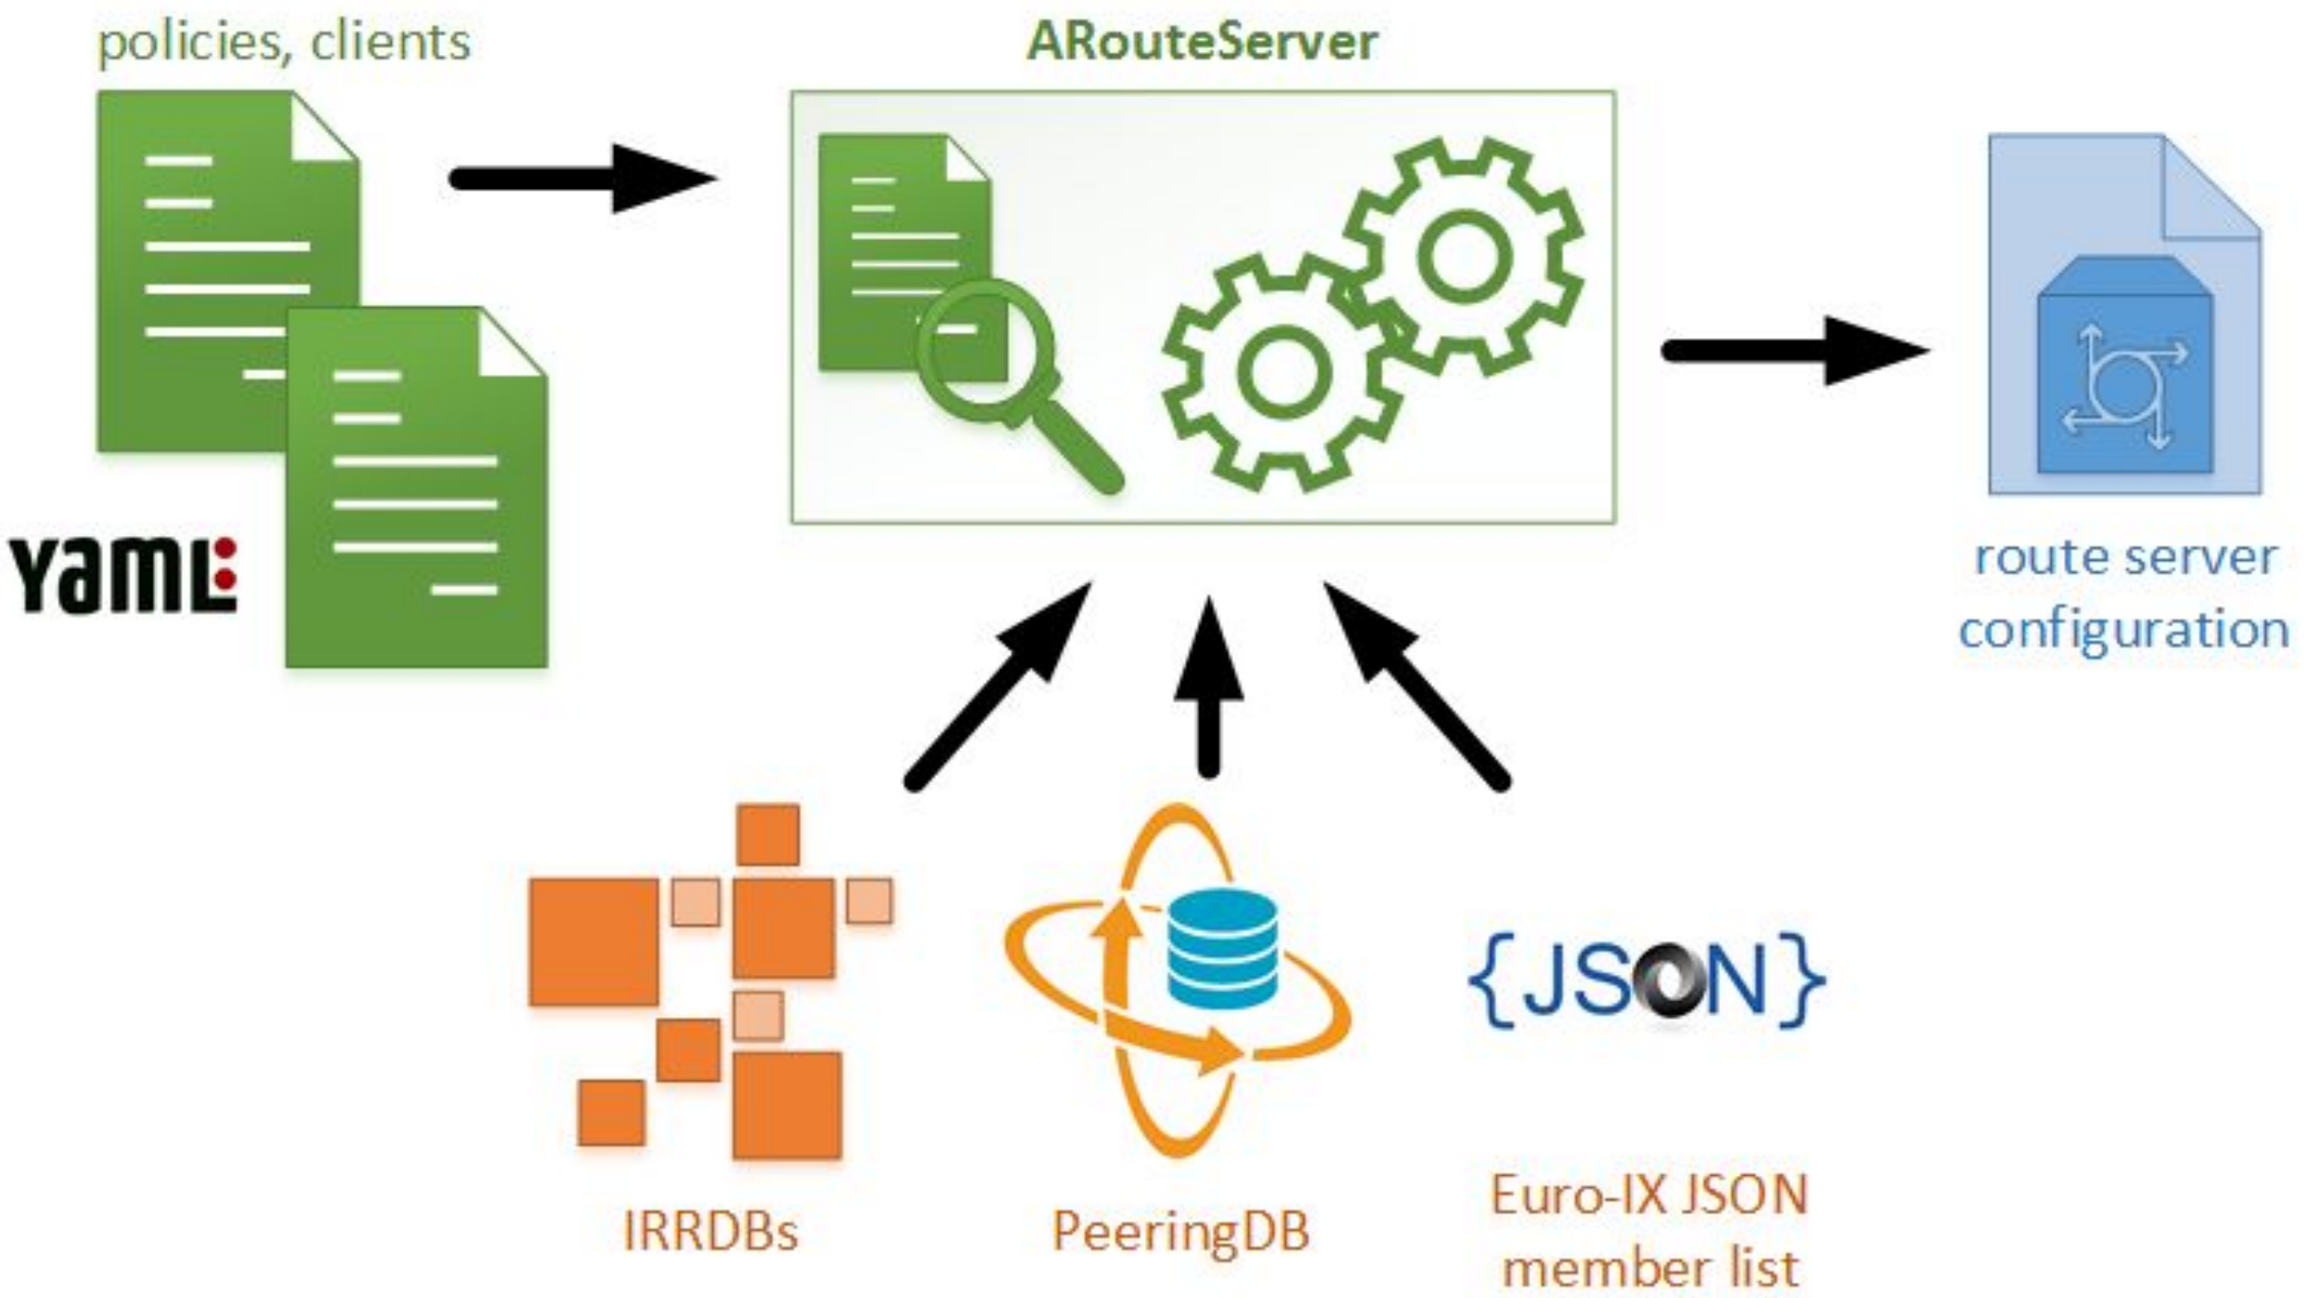

How does it work

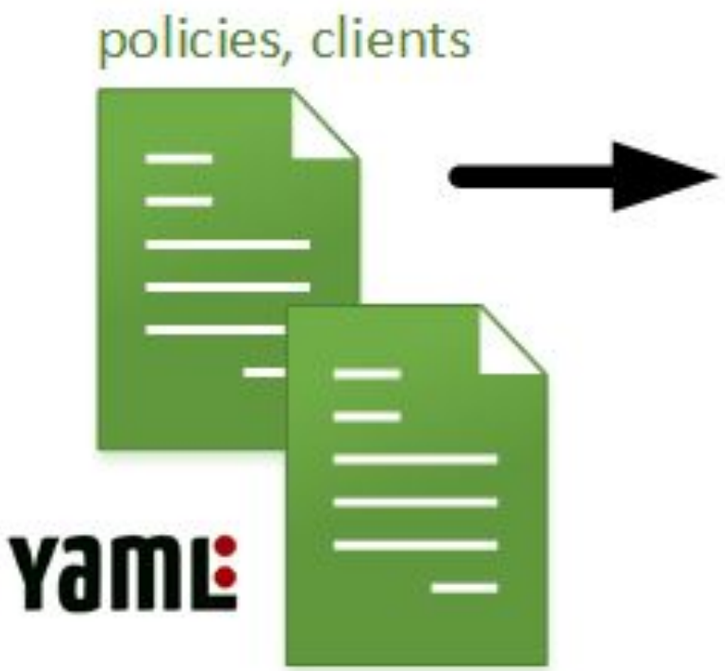

- Two YAML files are used to configure vendor-agnostic high-level policies and clients definitions.

```yaml
cfg:
  rs_as: 64496
  router_id: "192.0.2.2"
  filtering:
    next_hop:
      policy: "same-as"
  blackhole_filtering:
    policy_ipv4: "rewrite-next-hop"
    ...
```

```yaml
clients:
  - asn: 64511
    ip:
    - "192.0.2.11"
    - "2001:db8:1:1::11"
    irrdb:
      as_sets:
        - "RIPE::AS-FOO"
  ...
```

How does it work

- External sources are used to gather additional info:
  - IRRDBs for filters based on prefixes and origin ASNs;
  - PeeringDB for max-prefix limits and AS-SETs;
  - Euro-IX JSON member list files to build clients list automatically.

How does it work

- Finally, the route server configuration is automatically generated.
- Custom scripts can be used to validate it and to deploy it to the route server.

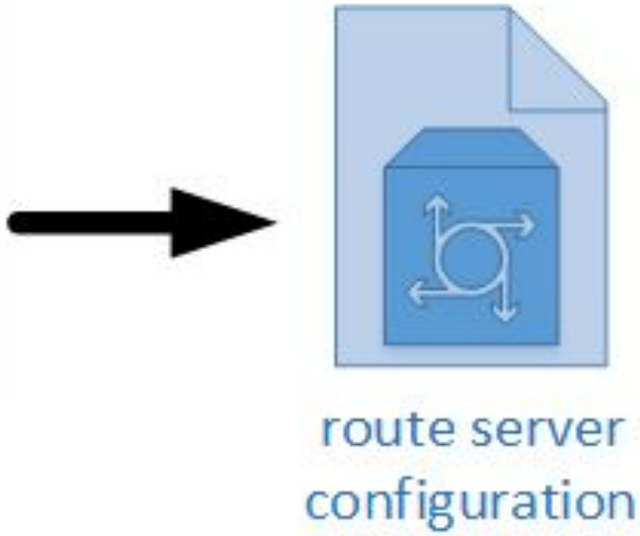


```
$ arouteserver bird --ip-ver 4 -o /etc/bird/bird4.new && \
    bird -p -c /etc/bird/bird4.new                    && \
    cp /etc/bird/bird4.new /etc/bird/bird4.conf       && \
    birdcl configure
```

Configuration

- Policies can be set on a global scope.
- Clients inherit global policies;
  - client-specific options override inherited ones.
- Client list can be automatically generated from Euro-IX JSON file;
  - this allows an easy integration with IXP-Manager.
- Custom, site-specific behaviours and configurations can be implemented using “hooks” and local files.

Routes filtering

- NEXT_HOP enforcement: strict or same AS.
- Min/max prefix length & max AS_PATH length.
- AS_PATH sanitation:
  - leftmost ASN;
  - private/invalid ASNs;
  - “transit-free” ASNs.
- Bogons and IRR-based filters.
- RPKI BGP Origin Validation.

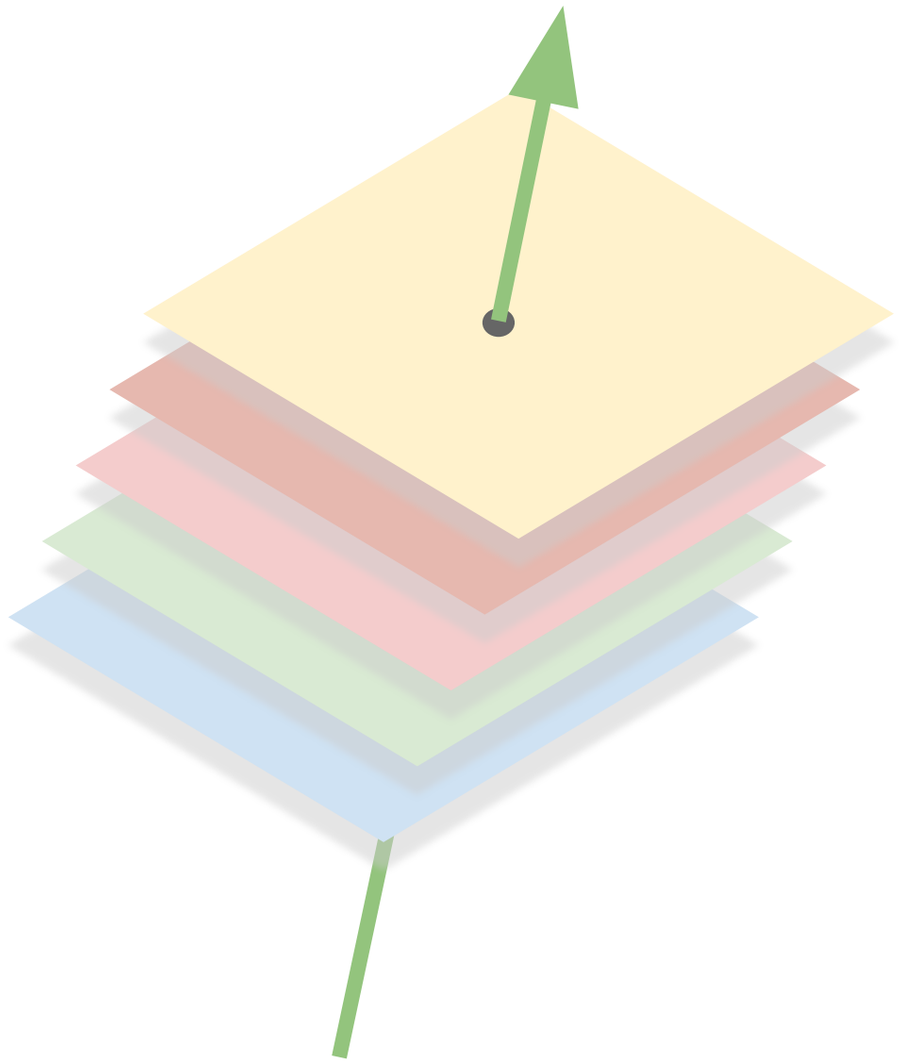

Filtering: NEXT_HOP enforcement

- “Strict” mode: accept only routes whose NEXT_HOP is the IP address of the announcing client.

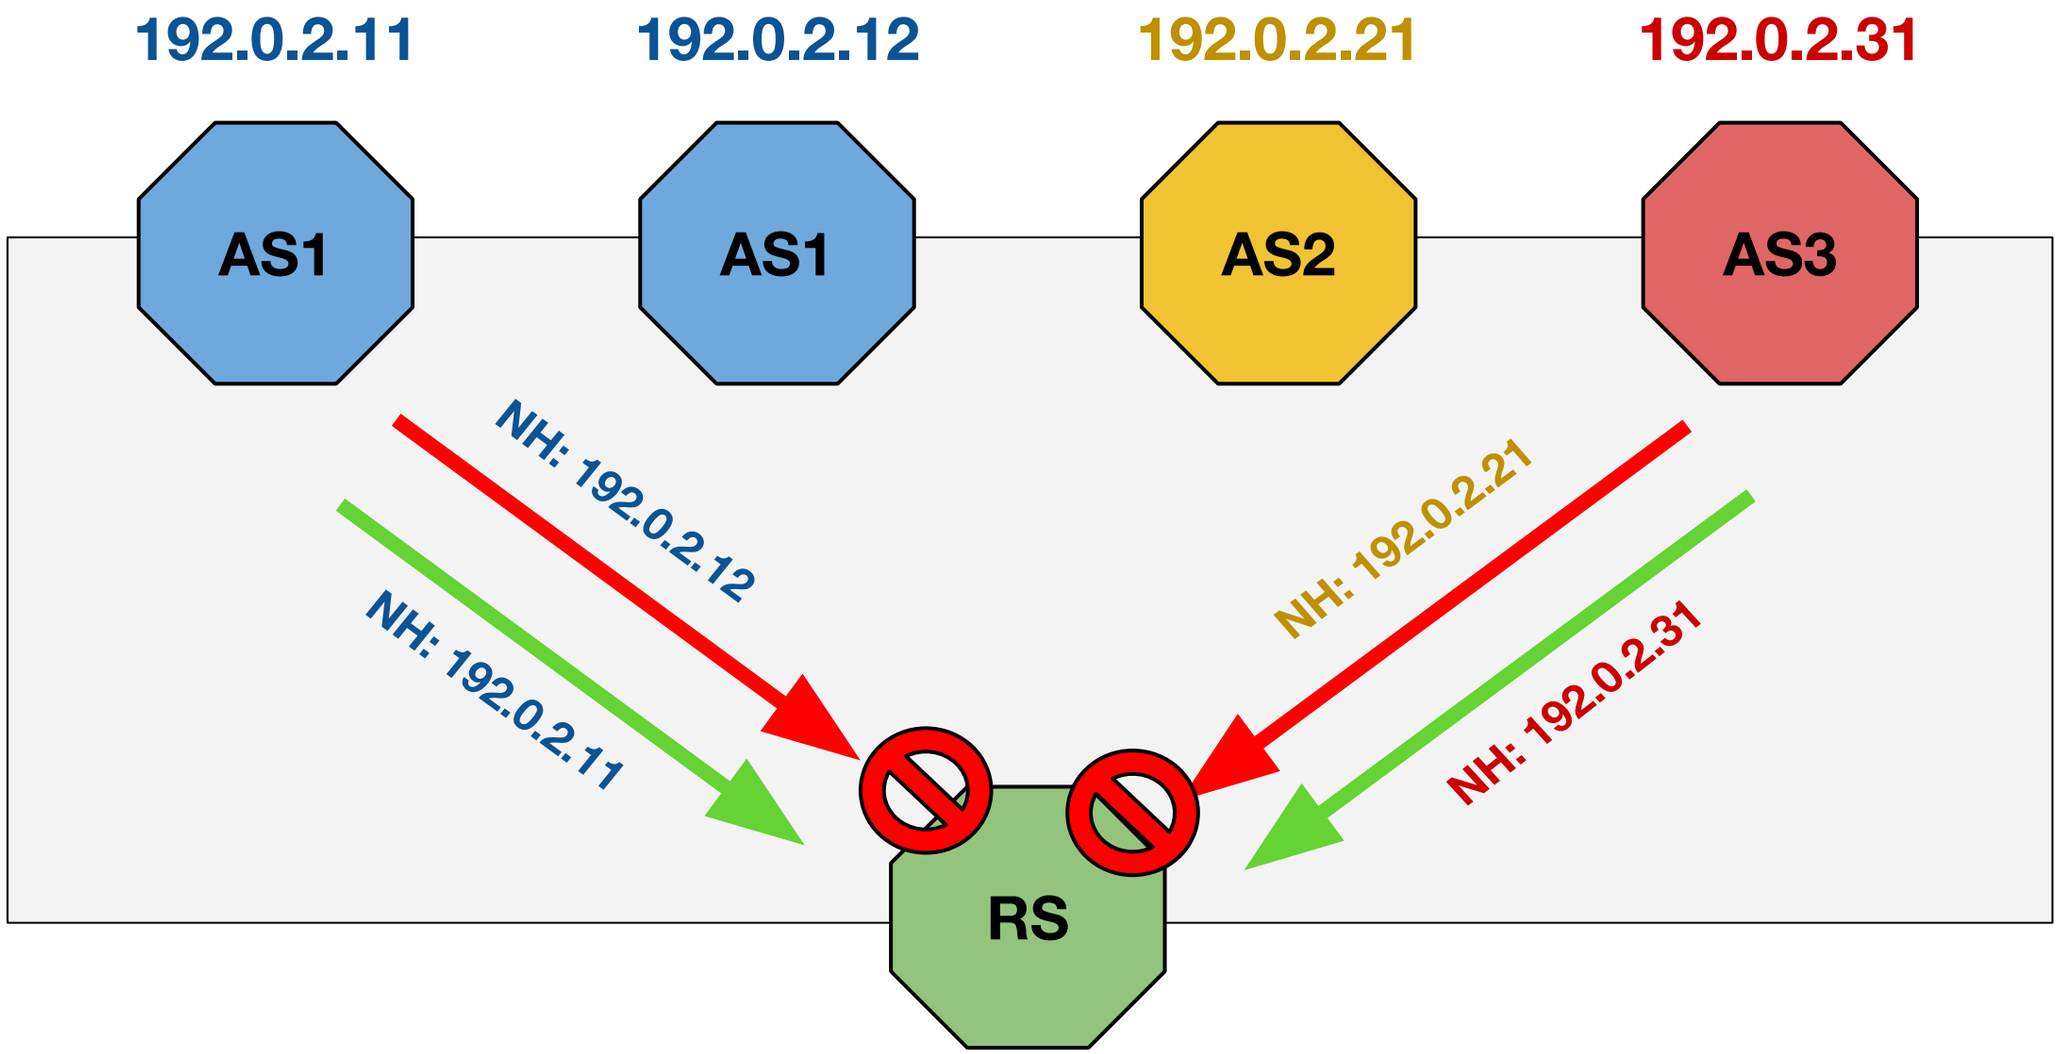

Filtering: NEXT_HOP enforcement

- “Same-AS” mode: allows the NEXT_HOP to be any address among those used by the announcing AS to connect to the IXP.

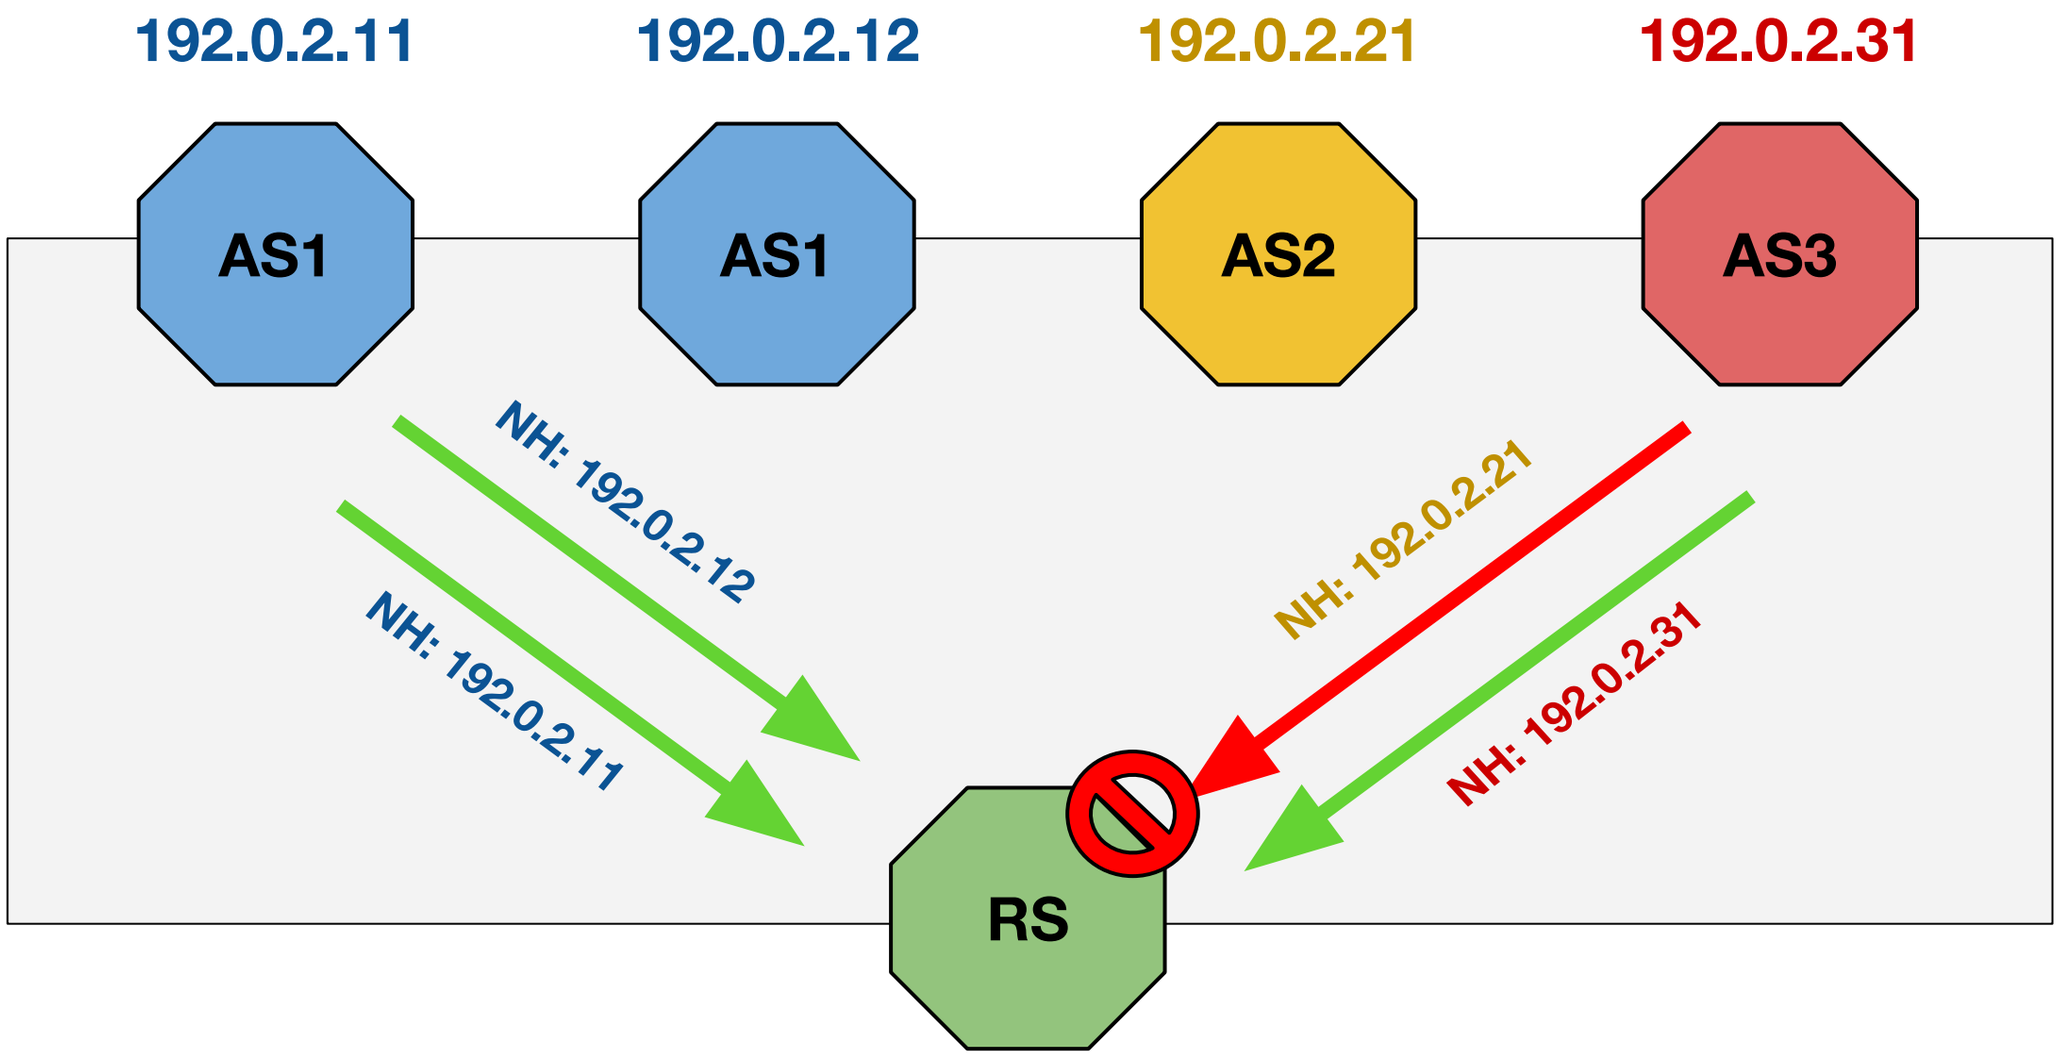

Filtering: AS_PATH sanitation

- Leftmost ASN in the AS_PATH must match the ASN of the client.

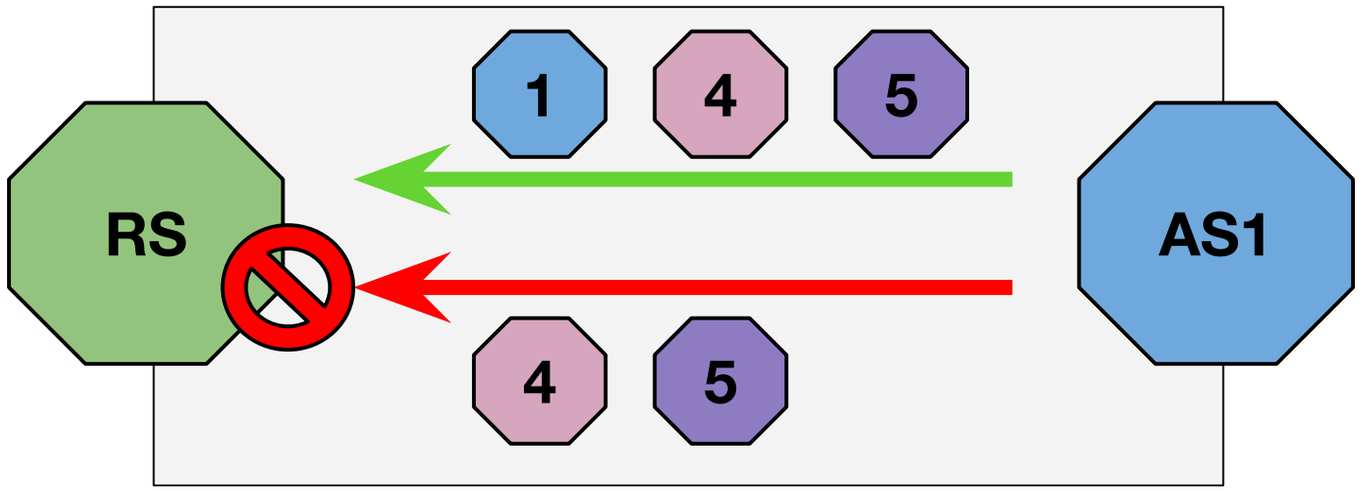

- Private/invalid ASNs must not appear in the AS_PATH.

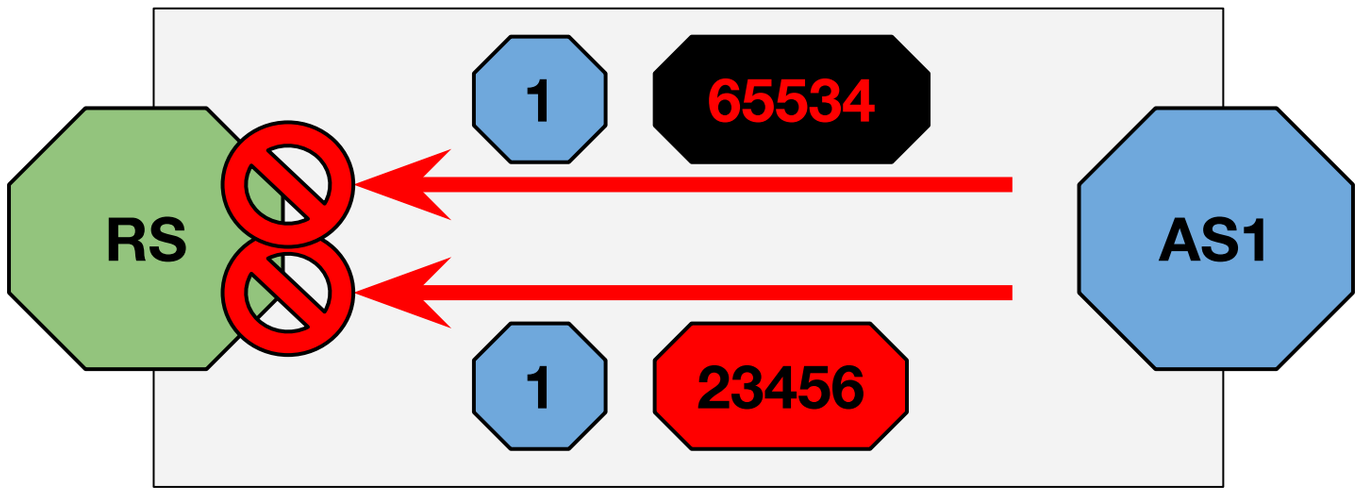

Filtering: AS_PATH sanitation

- “Transit-free” ASNs must not appear in the AS_PATH - unless they are in the first position.

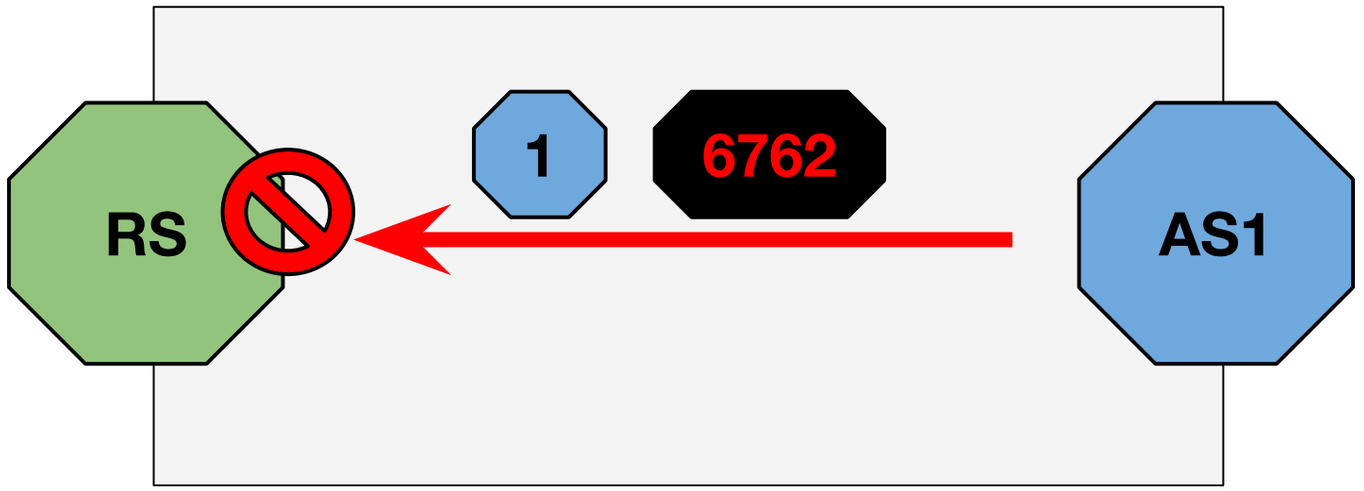

Filtering: prefix validation

- IRR databases are used to get the list of allowed prefixes and origin ASNs on the basis of client's AS-Macros (bgpq3).

```
$ whois AS-ONE
as-set:         AS-ONE
members:        AS4
```

```
$ whois -i origin AS4
route:    192.0.2.0/24
origin:   AS4
```

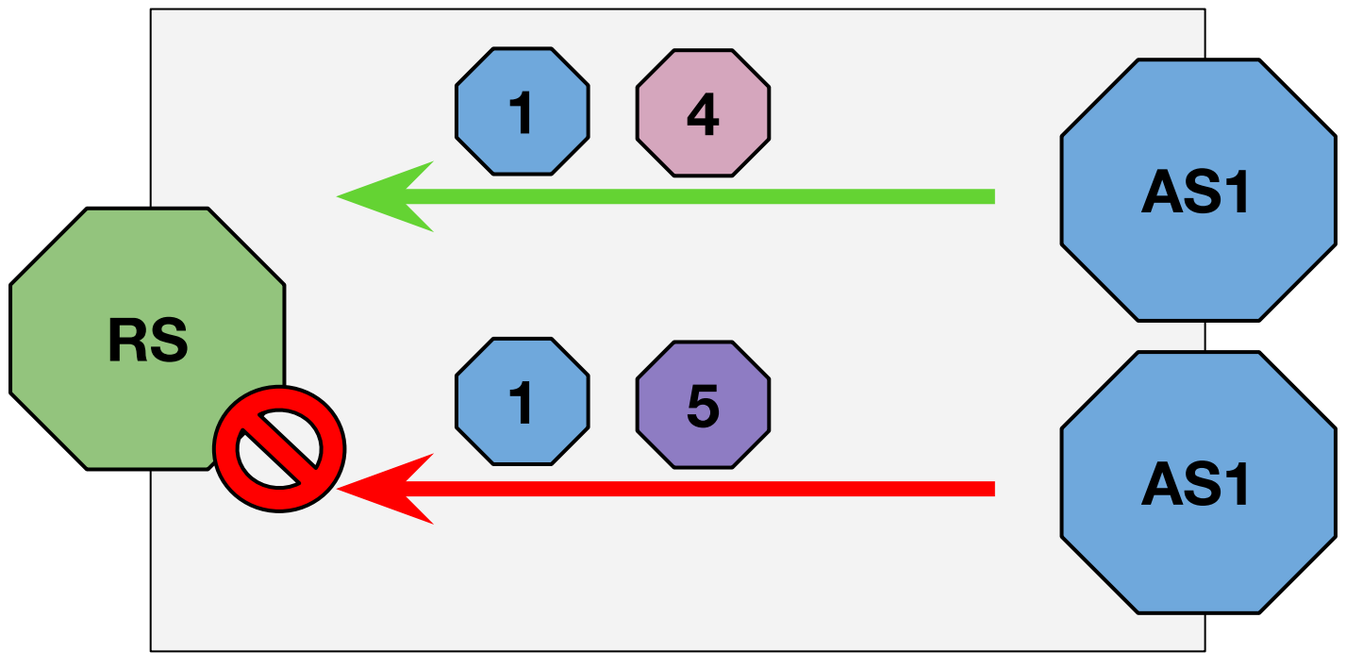

Filtering: prefix validation

- Authoritative data from the ARIN Whois database is also used to get a list of prefixes for which a specific ASN has been authorized to announce routes for.
- Improvement on filters accuracy: 23% (YYCIX data).

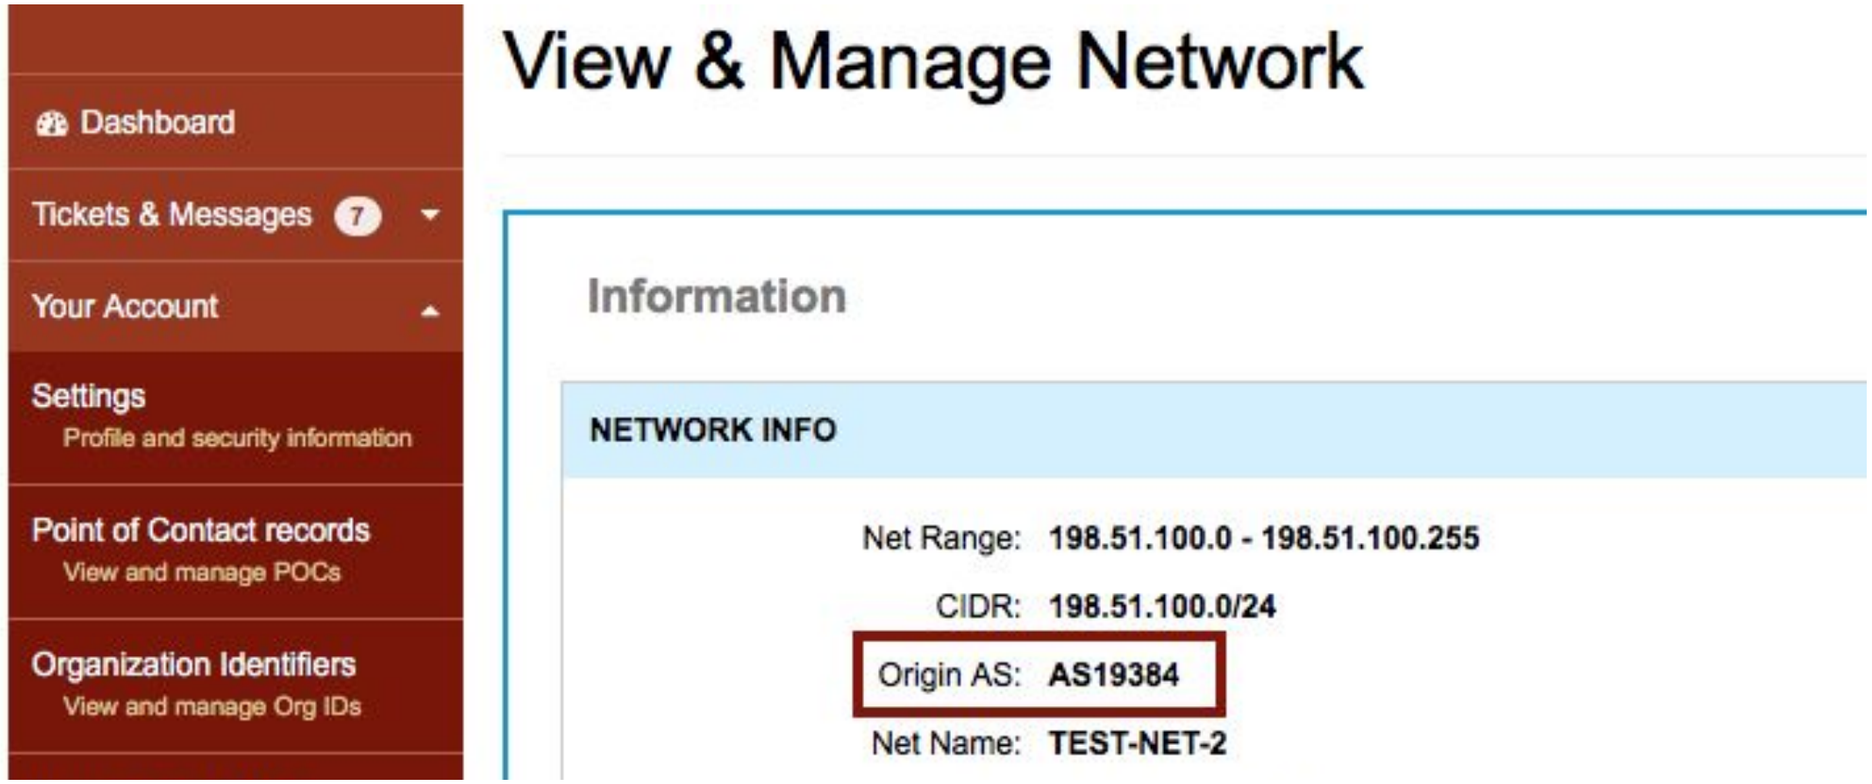

source: https://teamarin.net

Filtering: prefix validation

- RPKI ROAs (Route Origin Authorisations) are used for Prefix Origin Validation.

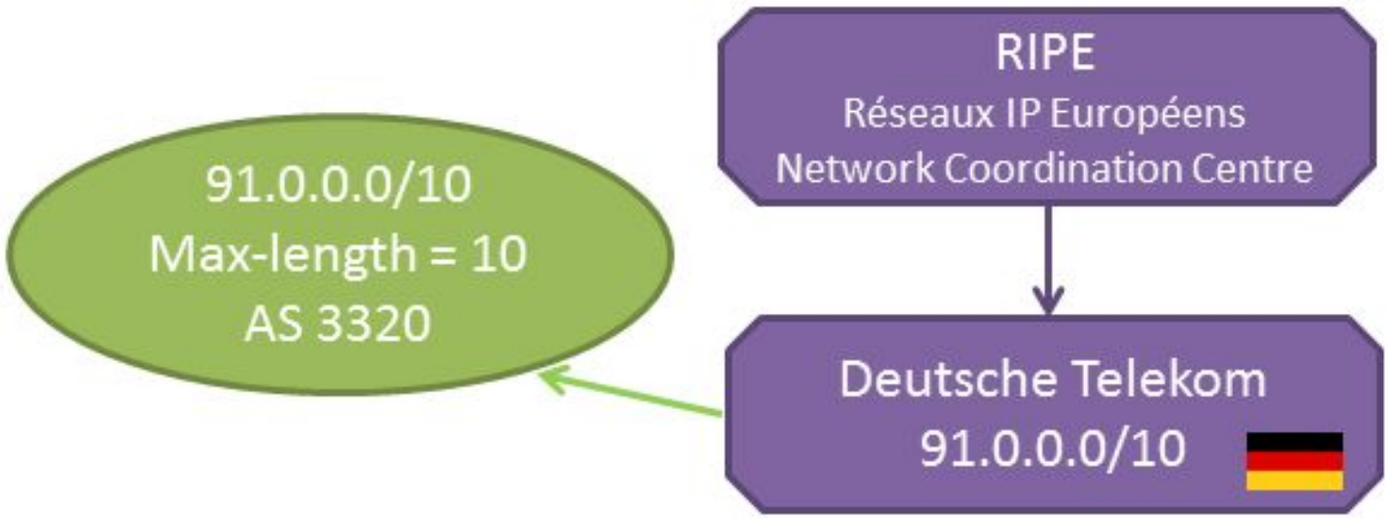

source: http://roalert.org

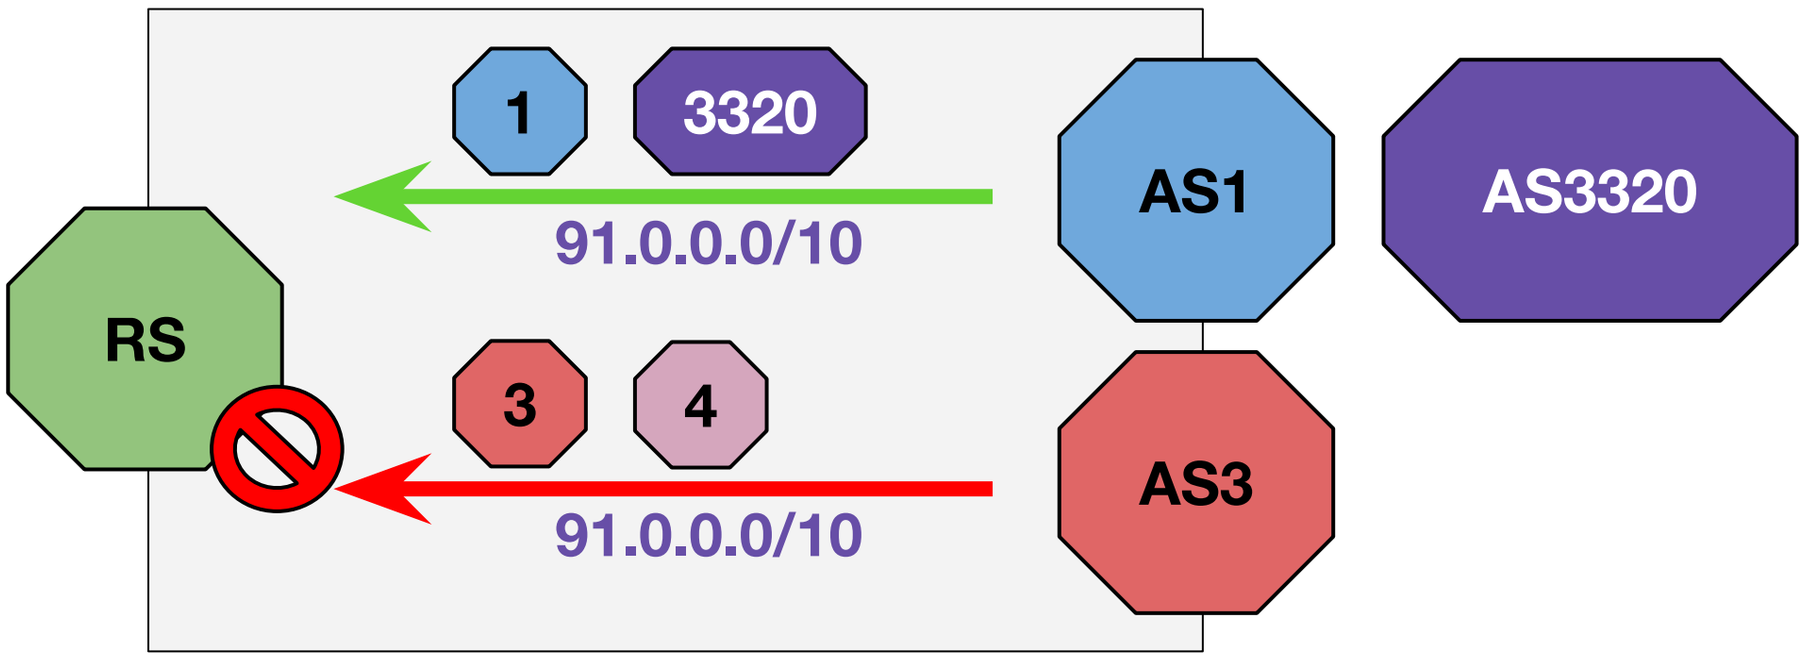

Services for the clients

- Blackhole filtering support:
  - optional NEXT_HOP rewriting to allow Layer-2 level filtering.
- Route propagation control (via BGP communities):
  - announce / do not announce to any / specific peer;
  - prepend to any / specific peer;
  - add NO_EXPORT / NO_ADVERTISE to any / specific peer.
- GRACEFUL_SHUTDOWN support (RFC8326).
- Full list of features available on GitHub.

Project status

- Actively developed. Both BIRD and OpenBGPD users.
- Feedback from real life is strongly needed and encouraged.

Source code and examples available on GitHub:

https://github.com/pierky/arouteserver

Full documentation:

https://arouteserver.readthedocs.io/

Pier Carlo Chiodi - https://pierky.com/

@pierky

Questions?