The art of filtering

Leveraging IRRs and tools to better secure the Internet

Aftab Siddiqui - Massimiliano Stucchi
2022

APRICOT

The importance of filtering

- **Routing security is vital to the future and stability of the Internet.**
- The Internet's routing foundation has cracks, and they're growing. Not a single day goes by without dozens of incidents affecting the routing system. Route hijacking, route leaks, IP address spoofing, and other harmful activities can lead to DDoS attacks, traffic inspection, lost revenue, reputational damage, and more. These incidents are global in scale, with one operator's routing problems cascading to impact others.

Filtering - RFC7454

- The main aspect of securing BGP resides in controlling the prefixes that are received and advertised on the BGP peerings. Prefixes exchanged between BGP peers are controlled with inbound and outbound filters that can match on IP prefixes, AS paths or any other attributes of a BGP prefix (for example, BGP communities)

The process of filtering

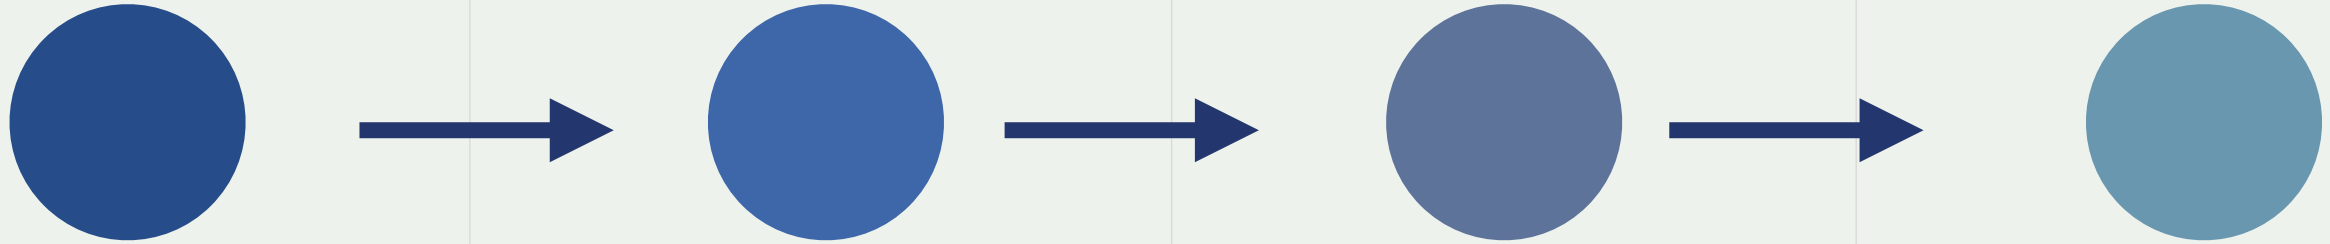

Collect all the ASNs in the customer/ neighbour cone

Collect all the route and route6 objects with origins from the customer/neighbor cone

Generate a prefix-list containing all the infos from route and route6 objects

Load the prefix list on the router peering with the customer or neighbour

Data Sources

- IRRs
- Bogons lists (IPv4 & IPv6)
- PeeringDB

The Internet Routing Registry

IRR Database objects

IPs and ASNs

inetnum

inet6num

aut-num

Routing

route

route6

As-set

Object Protection

mntner

Autonomous Number Objects

- Known as aut-num objects
- Register who holds an AS Number and the routing policy for that AS

aut-num:	AS12345
as-name:	YOUR-AS-NAME
org:	ORG-EE2-RIPE
import:	**from AS1010 accept ANY**
export:	**to AS1010 announce AS12345**
import:	**from AS987 accept ANY**
export:	**to AS987 announce AS12345**
admin-c:	DV789-RIPE
tech-c:	JS123-RIPE
status:	ASSIGNED
mnt-by:	RIPE-NCC-END-MNT
mnt-by:	DEFAULT-LIR-MNT
source:	RIPE

What are route(6) objects ?

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

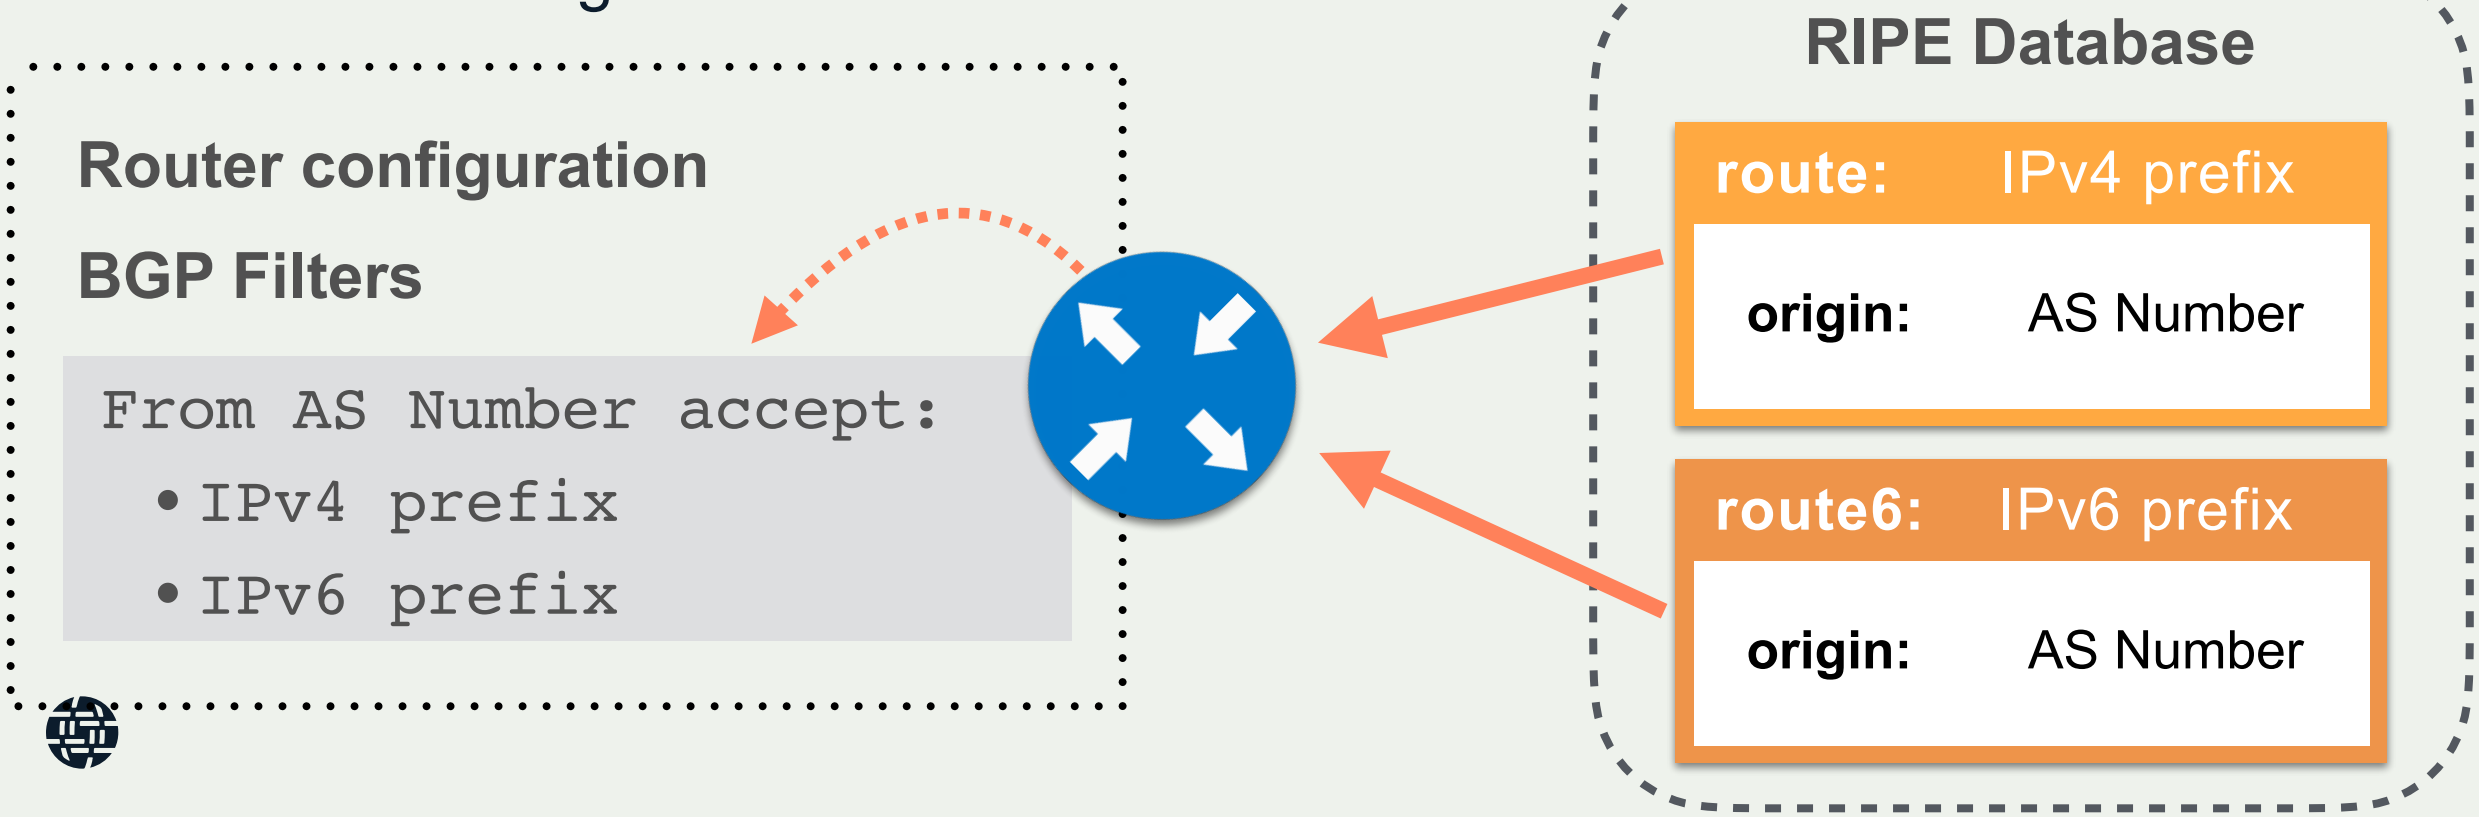

AS-Sets

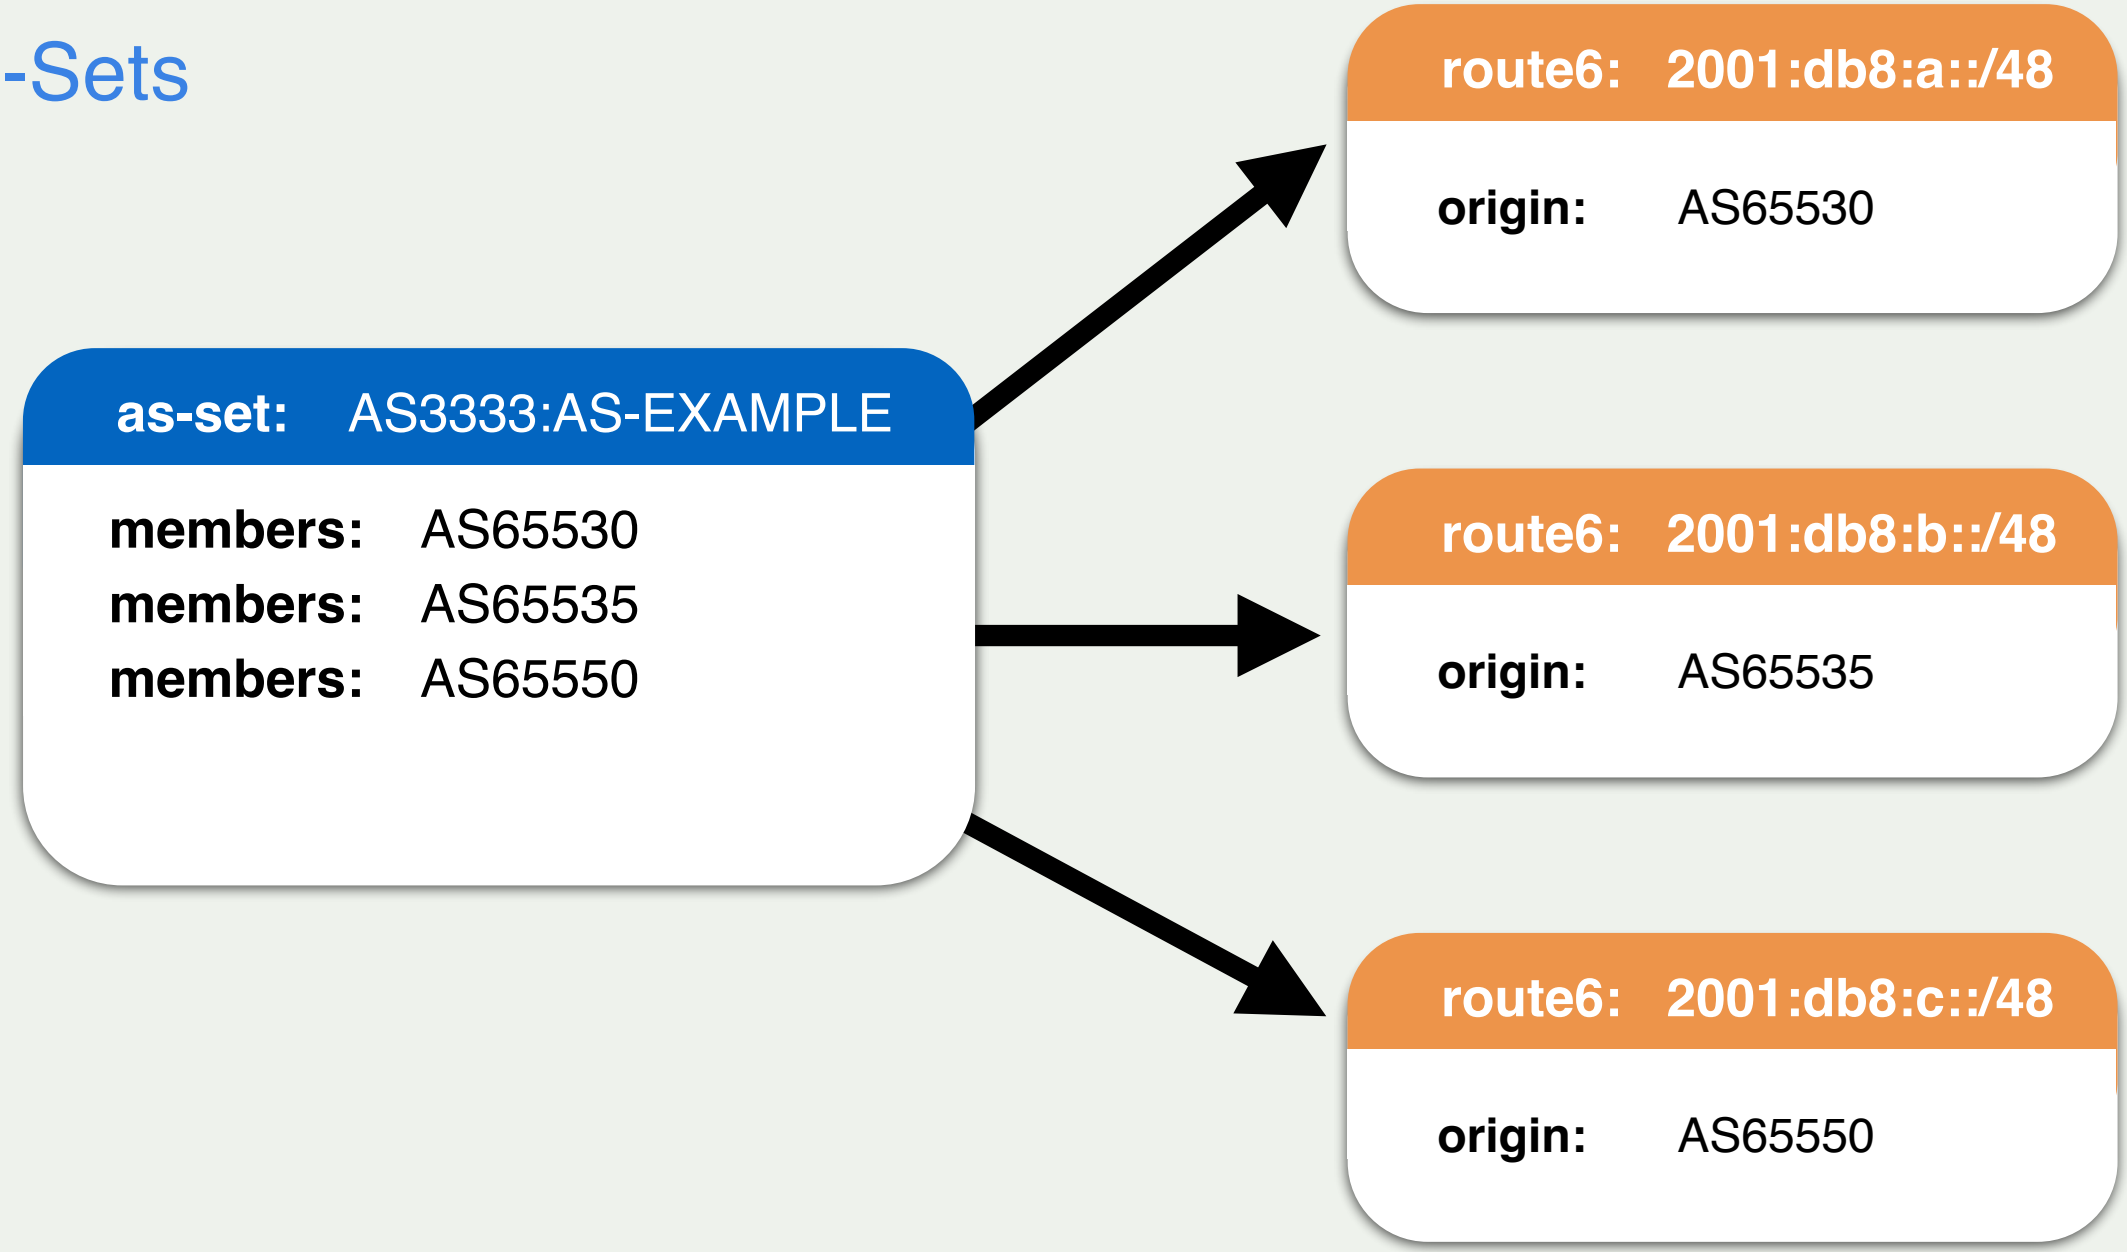
as-set: AS3333:AS-EXAMPLE
members: AS65530
members: AS65535
members: AS65550
route6: 2001:db8:a::/48
origin: AS65530
route6: 2001:db8:b::/48
origin: AS65535
route6: 2001:db8:c::/48
origin: AS65550

Customer Cones

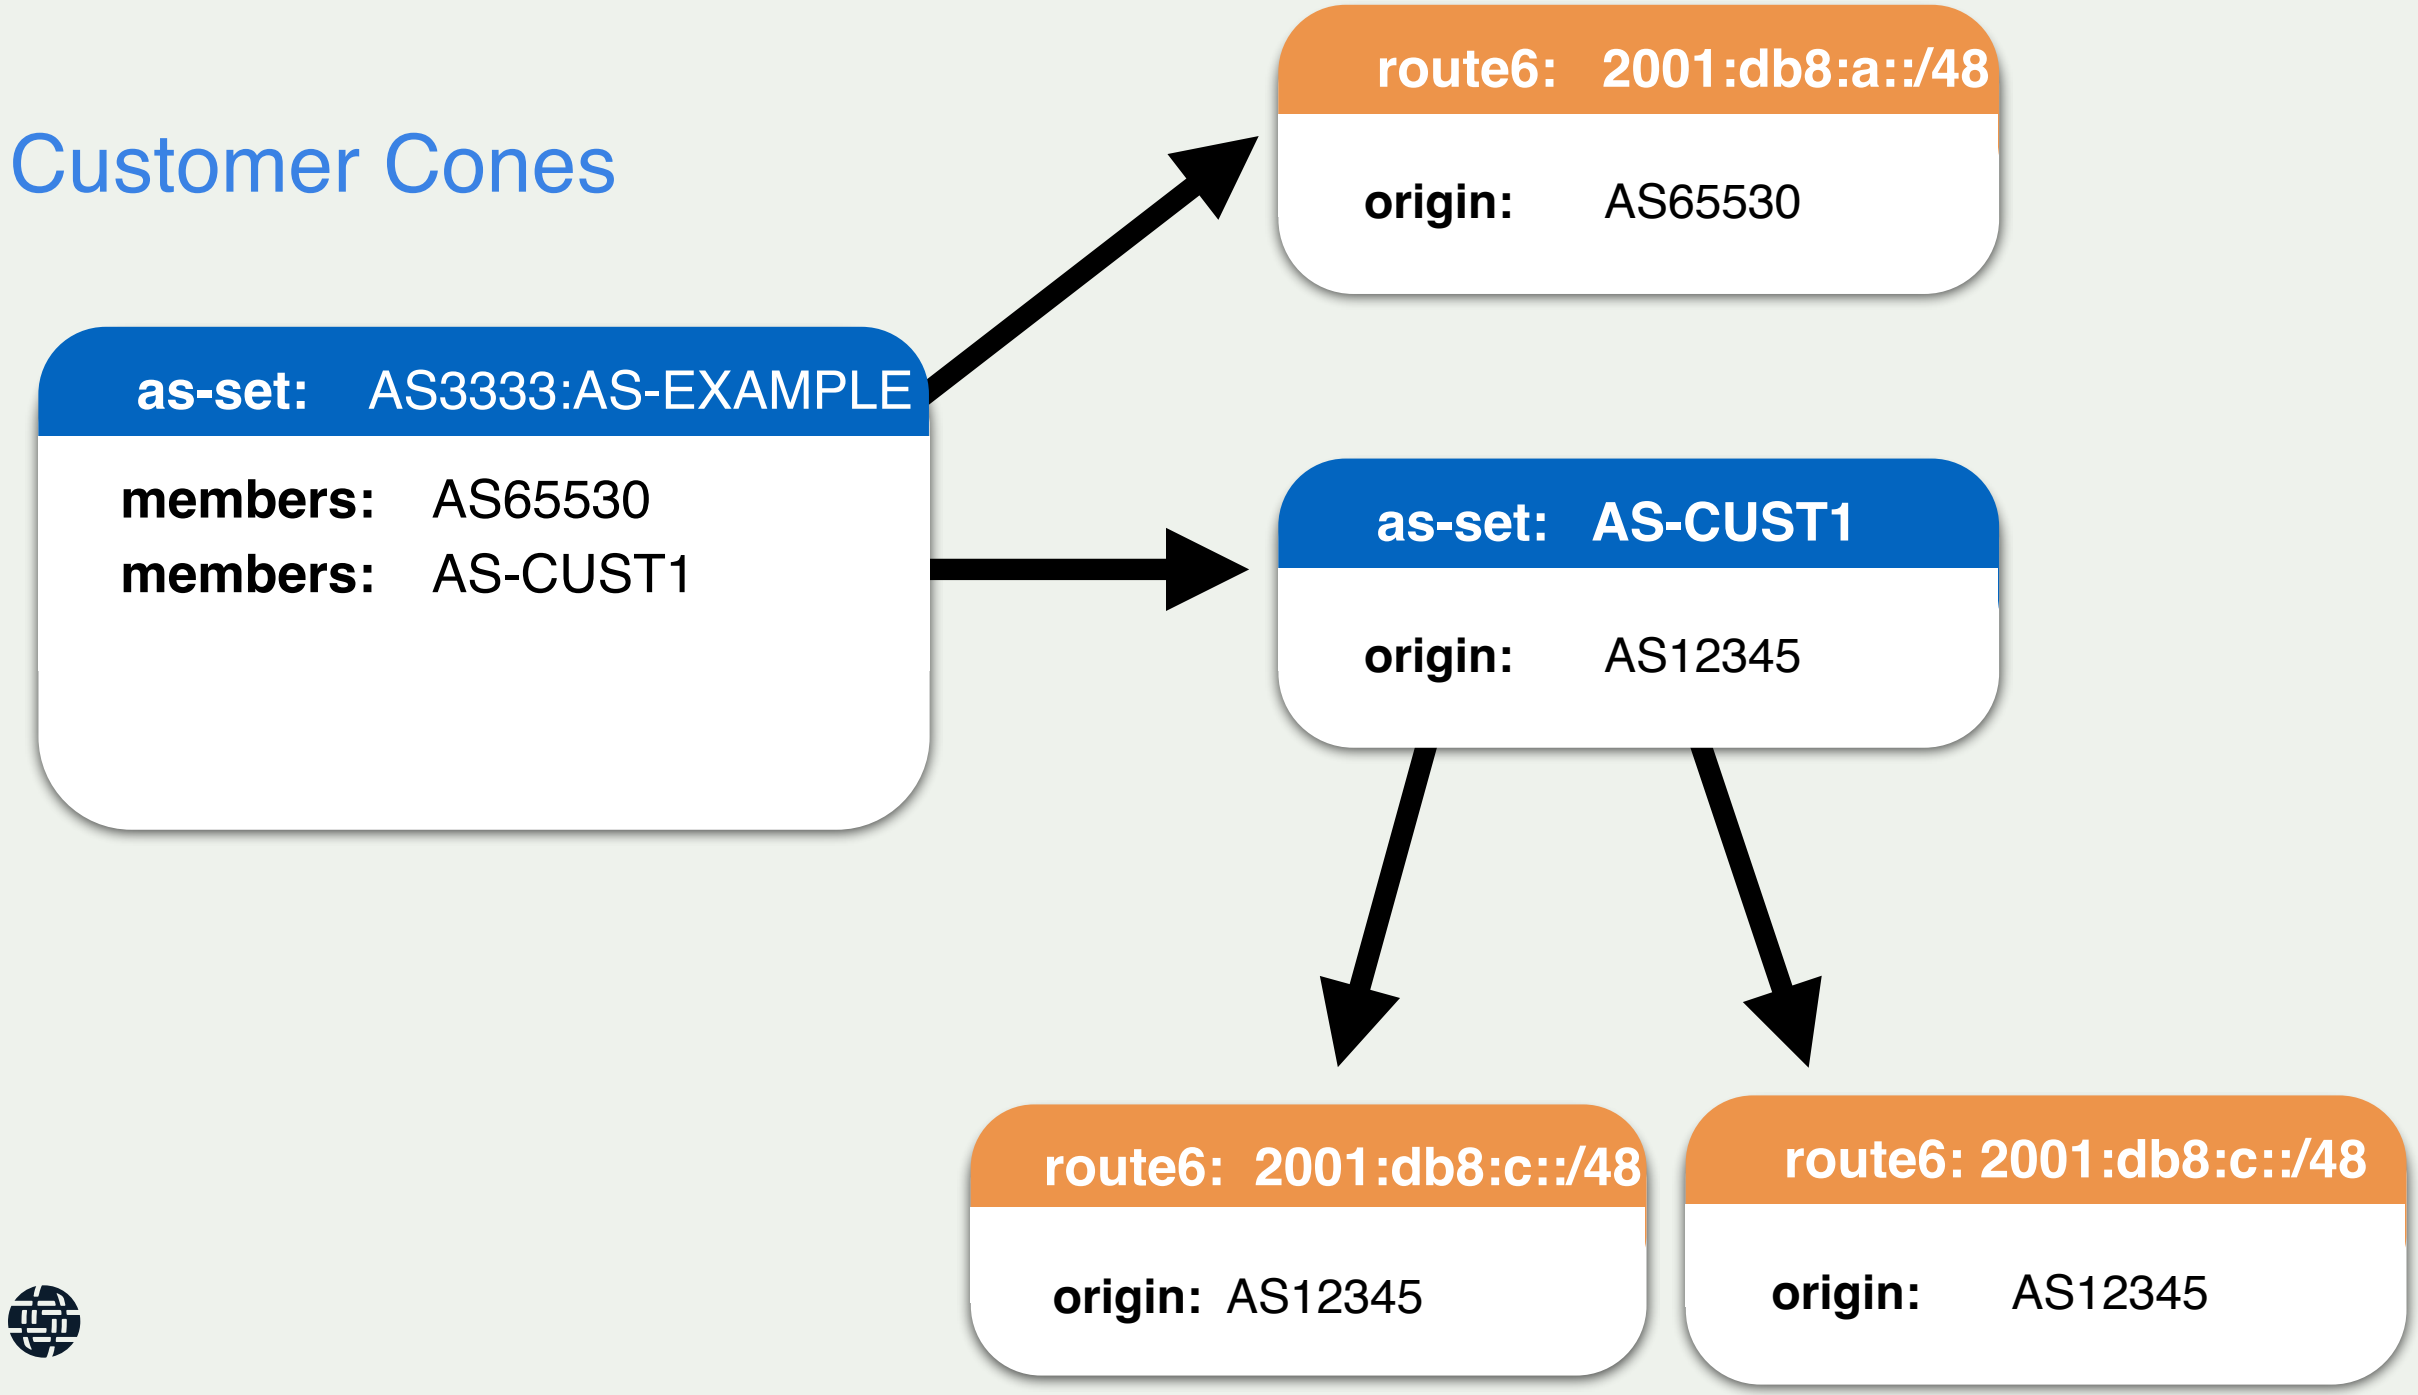

Prefix Lists

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
  - Level3 Filtergen —> whois -h filtergen.level3.net (resource limitations!!!!)
  - Bgpq3/4
  - peval —> not working
  - https://www.dan.me.uk/filtergen —> Web version

Prefix Filtering

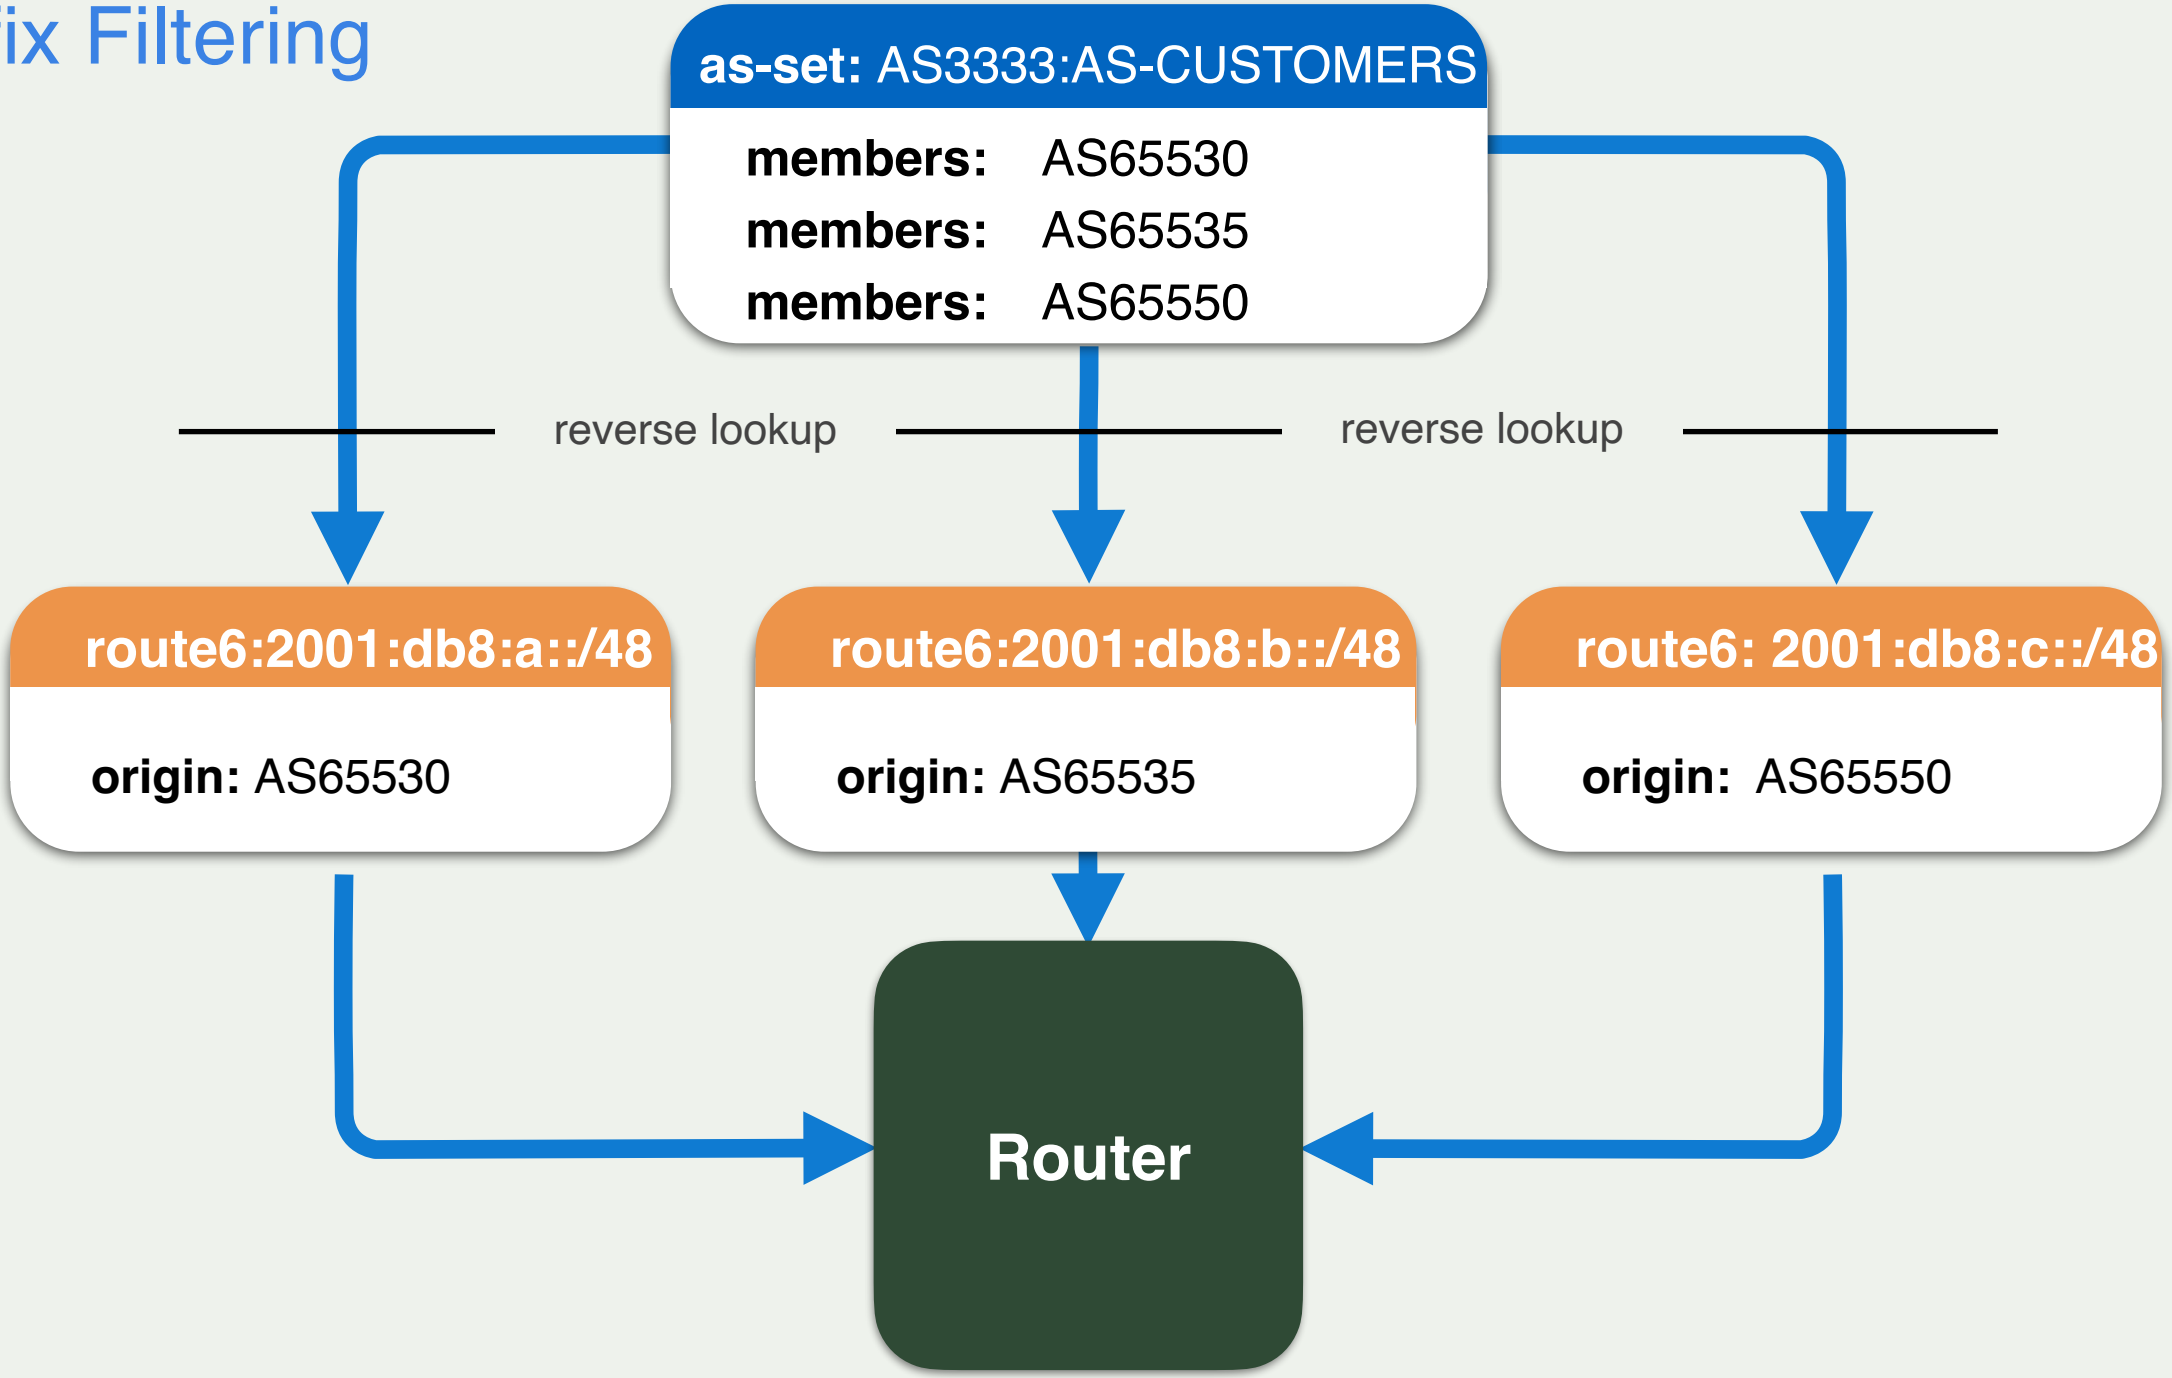

Generating a prefix filter

$ **bgpq4 -h rr.ntt.net -l AS15562-in AS15562:AS-SNIJDERS**
no ip prefix-list AS15562-in
ip prefix-list AS15562-in permit 67.221.245.0/24
ip prefix-list AS15562-in permit 165.254.255.0/24
ip prefix-list AS15562-in permit 165.254.255.0/26
ip prefix-list AS15562-in permit 165.254.255.26/32
ip prefix-list AS15562-in permit 165.254.255.64/26
ip prefix-list AS15562-in permit 165.254.255.132/32
ip prefix-list AS15562-in permit 165.254.255.133/32
ip prefix-list AS15562-in permit 165.254.255.144/28
ip prefix-list AS15562-in permit 165.254.255.149/32
ip prefix-list AS15562-in permit 165.254.255.160/28
ip prefix-list AS15562-in permit 192.147.168.0/24
ip prefix-list AS15562-in permit 204.2.30.0/23
ip prefix-list AS15562-in permit 204.42.254.192/26
ip prefix-list AS15562-in permit 209.24.0.0/16

Ingress filtering 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

Bogons

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

- Team Cymru provides lists of bogons
  - http://www.team-cymru.com/bogon-reference.html

Bogon/Martian ASN Filters

- 0, # RFC 7607
- 23456, # RFC 4893 AS_TRANS
- 64496..64511, # RFC 5398 and documentation/example ASNs
- 64512..65534, # RFC 6996 Private ASNs
- 65535, # RFC 7300 Last 16 bit ASN
- 65536..65551, # RFC 5398 and documentation/example ASNs
- 65552..131071, # RFC IANA reserved ASNs

```
bgp as-path access-list bogon-asns deny 23456
bgp as-path access-list bogon-asns deny 64496-131071
bgp as-path access-list bogon-asns deny 4200000000-4294967295
```

https://bgpfilterguide.nlnog.net/guides/bogon_asns/

Bogon/Martian Prefixes (IPv4)

- 0.0.0.0/8, # RFC 1122 'this' network
- 10.0.0.0/8, # RFC 1918 private space
- 100.64.0.0/10, # RFC 6598 Carrier grade nat space
- 127.0.0.0/8, # RFC 1122 localhost
- 169.254.0.0/16, # RFC 3927 link local
- 172.16.0.0/12, # RFC 1918 private space
- 192.0.2.0/24, # RFC 5737 TEST-NET-1
- 192.88.99.0/24, # RFC 7526 6to4 anycast relay
- 192.168.0.0/16, # RFC 1918 private space
- 198.18.0.0/15, # RFC 2544 benchmarking
- 198.51.100.0/24, # RFC 5737 TEST-NET-2
- 203.0.113.0/24, # RFC 5737 TEST-NET-3
- 224.0.0.0/4, # multicast
- 240.0.0.0/4 # reserved

Bogon/Martian Prefix Filters (FRR)

- ip prefix-list BOGONS_v4 deny 0.0.0.0/8 le 32
- ip prefix-list BOGONS_v4 deny 10.0.0.0/8 le 32
- ip prefix-list BOGONS_v4 deny 100.64.0.0/10 le 32
- ip prefix-list BOGONS_v4 deny 127.0.0.0/8 le 32
- ip prefix-list BOGONS_v4 deny 169.254.0.0/16 le 32
- ip prefix-list BOGONS_v4 deny 172.16.0.0/12 le 32
- ip prefix-list BOGONS_v4 deny 192.0.2.0/24 le 32
- ip prefix-list BOGONS_v4 deny 192.88.99.0/24 le 32
- ip prefix-list BOGONS_v4 deny 192.168.0.0/16 le 32
- ip prefix-list BOGONS_v4 deny 198.18.0.0/15 le 32
- ip prefix-list BOGONS_v4 deny 198.51.100.0/24 le 32
- ip prefix-list BOGONS_v4 deny 203.0.113.0/24 le 32
- ip prefix-list BOGONS_v4 deny 224.0.0.0/4 le 32
- ip prefix-list BOGONS_v4 deny 240.0.0.0/4 le 32

https://bgpfilterguide.nlnog.net/guides/bogon_prefixes/

Live Demo

Live Demo

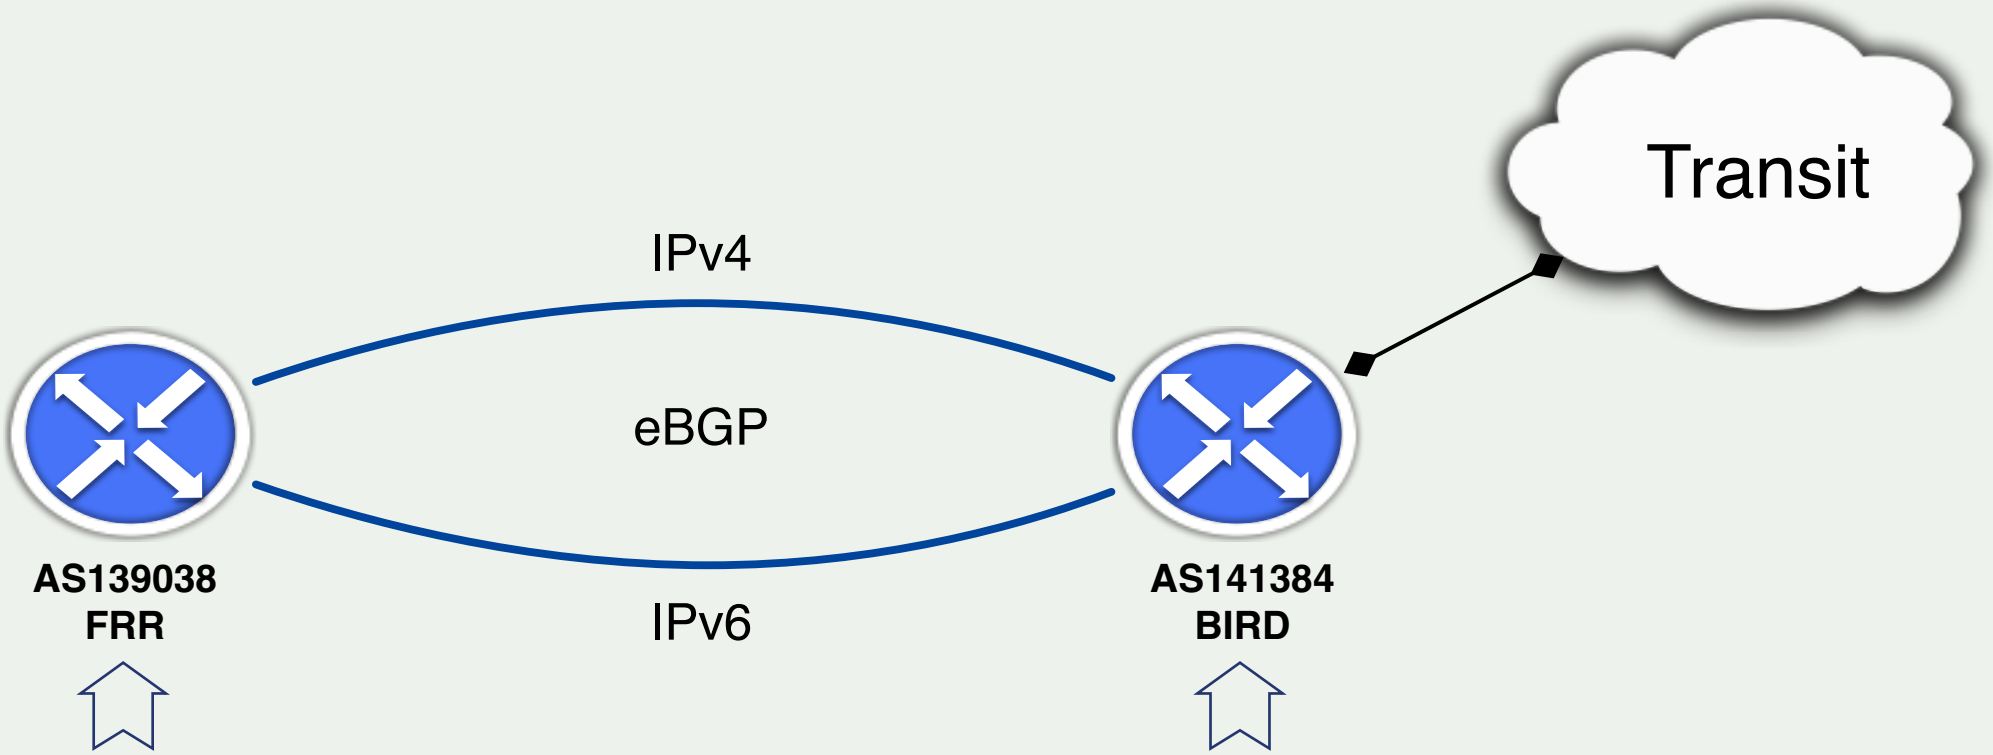

- Automatically implement ingress/ egress prefix filters using bgpq3/4
- Add basic bogon filters on ingress/ egress
- Add ASN filtering
- Add prefix limit

- Automatically implement ingress/ egress prefix filters using bgpq3/4
- Add basic bogon filters on ingress/ egress
- Add ASN filtering
- Add prefix limit

Antispoofing and RPF

Anti-Spoofing

- What is spoofing
  - What are the consequences of spoofed traffic
  - Tracking/Tracing a spoofed packet

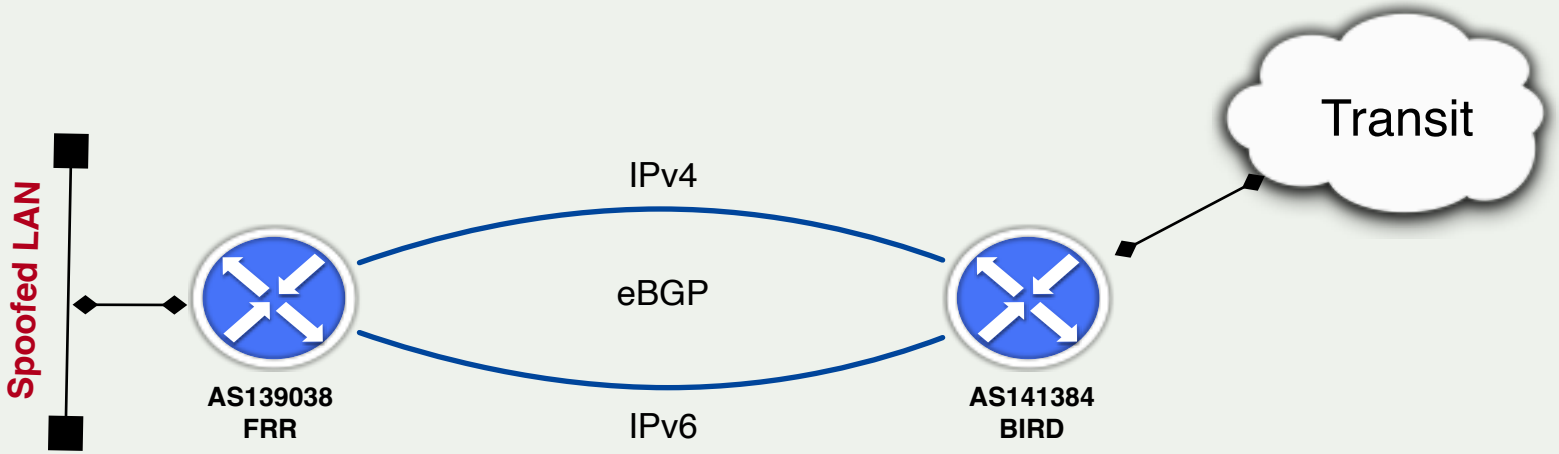

Reverse Path Forwarding

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

- Two main modes
  - Loose
  - Strict
- Additional modes (feasible-path and vrf) not supported everywhere

Strict and loose RPF

- Strict
  - Checks if the entry is in the routing table
  - and the route points to the receiving interface

- Loose
  - Simply checks that an entry exists for the route in the routing table

PeeringDB

PeeringDB

- What is Peeringdb?
  - PeeringDB is a publicly available network database that is the go-to location for interconnection data. The database facilitates global network connections at Internet Exchange Points (IXPs), data centers, and other interconnection facilities, and it serves as a starting point for interconnection decisions.

PeeringDB

- Is it for me? Why should I put my network details?
  - Almost one-third of Autonomous System Numbers (ASNs) register their interconnection data in the PeeringDB database. That means, by using PeeringDB and adding your own interconnection data, you'll be able to confidently find information about networks looking to interconnect, where and how to connect with them, and they'll be able to find the same information about your network. Since the database is user-maintained and validated by our volunteers, you can trust that the information is accurate and up-to-date.
  - This data will help you to accelerate the process of finding and connecting with other networks while supporting a faster and more decisive deployment of your own network expansion and development plans.
  - Many networks are building automation that relies on PeeringDB. If you don't have an up to date PeeringDB record this might stop their automation configuring sessions.

PeeringDB

- How can I use Peeringdb for network operation
  - Create several database records, known as objects, to establish your presence in PeeringDB.
  - Start with entering minimum required data and add and update the information over time. To maximize the value of these entry in PeeringDB it is recommended to add more than the minimum required information. This information is required:
    - Company name
    - AS number
    - Contact information (mandatory for networks with a connection to an Internet Exchange)
  - Database records to create
    - User - Start the process by creating a user account. It is recommended to use email that already exist in IRR databases for ASN
    - Org - This is core record of Organisation
    - Net - Basic network information is automatically retrieved from IRR/NIR database on the ASN.

PeeringDB - Examples

- Demo
- Setup/Update Network in PeeringDB
- Fetch relevant information using GUI/Web Interface
  - List of Networks peering in IX Australian Sydney
- Fetch relevant information using CLI (cURL)
  - List of Networks peering in IX Australian Sydney
  - curl -s -X GET "https://www.peeringdb.com/api/net?ix=716&__in=Sydney" | jq '.data[]|.name,.asn' | paste - -

RPKI

RPKI

A security framework for verifying the association between resource holders and their Internet resources

Attaches digital certificates to network resources upon request that lists all resources held by the member

- AS Numbers
- IP Addresses

Operators associate those two resources

- Route Origin Authorisations (ROAs)

Two elements of RPKI

What is in a ROA ?

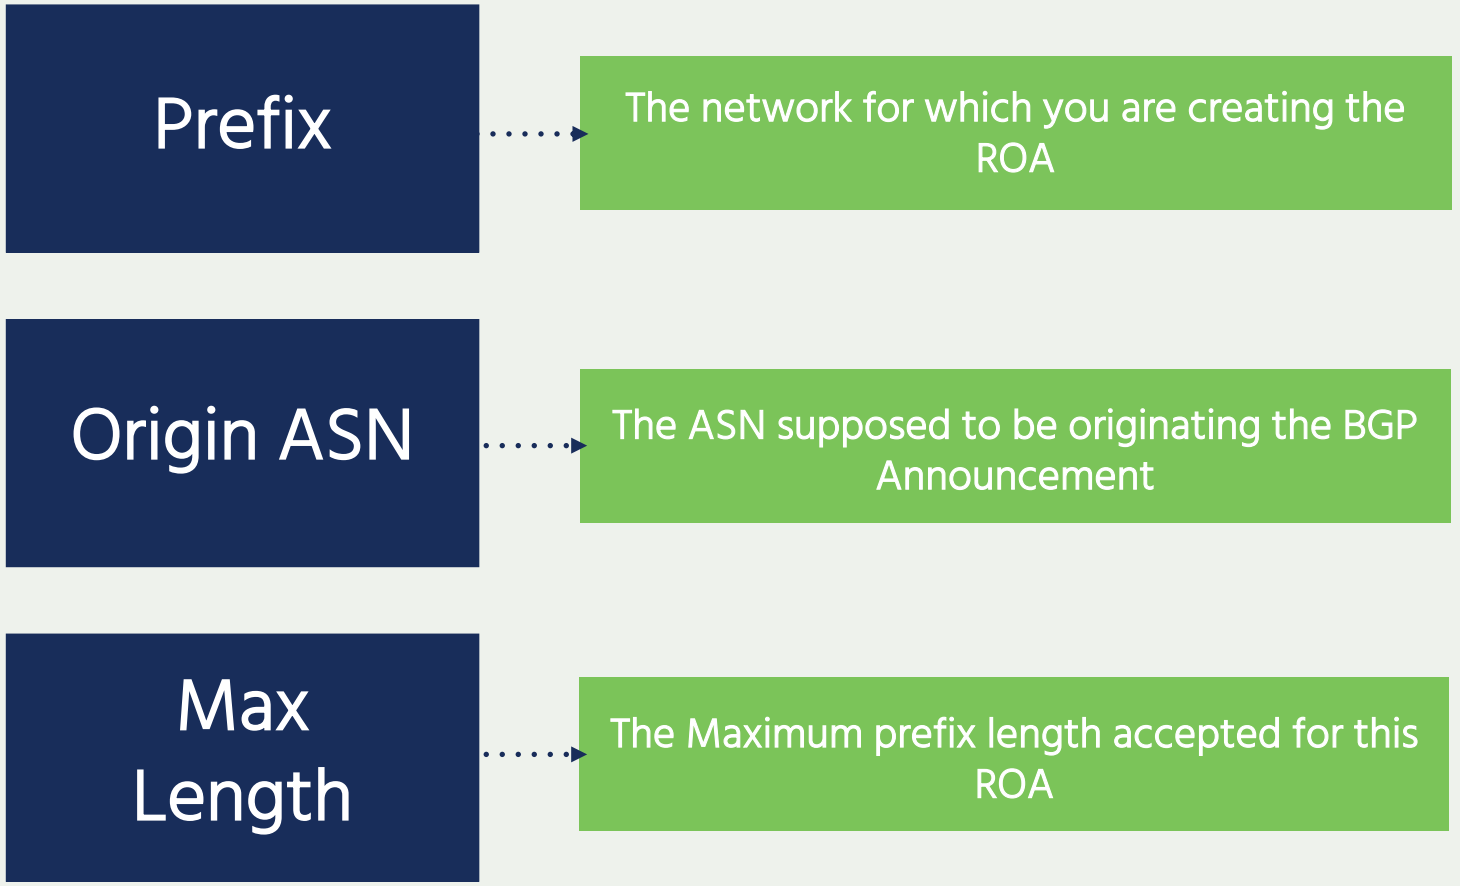

What is in a ROA

```
# python3 print_roa.py 6B4669CE7B7D11E99F99E776C4F9AE02.roa
ROA Version:    0
SigningTime:    2021-12-01T07:14:52Z
asID:           0
 addressFamily: 1
     IPAddress: 103.138.210.0/24
 addressFamily: 2
     IPAddress: 2001:df0:5580::/48
     IPAddress: 2001:df0:5580:400::/54
     IPAddress: 2001:df0:5580:c00::/54
```

What is in a Manifest

```
# python3 print_rpki_manifest.py QtPsuchXCrCQ62Ae2zN5wNPYptA.mft
Manifest Version: 0
SigningTime:      2022-01-31T19:11:40Z
Number:           5686
thisUpdate:       2022-01-31T19:11:39Z
nextUpdate:       2022-02-02T19:11:39Z
fileHashAlg:      id-sha256
fileList[  0]:    B3:A5:F0:0D:05:3F:AF:55:3A:CD:CC:63:9B:55:C4:2E:FE:7B:C2:83:DC:47:86:8E:45:40:72:55:50:7A:BD:69 QtPsuchXCrCQ62Ae2zN5wNPYptA.crl
fileList[  1]:    99:3A:CB:3F:6F:5A:0D:47:E2:31:4B:26:B4:8E:B4:49:64:6C:6C:94:19:9F:3A:5E:8B:4E:F9:FE:C5:82:B5:94 C8D6C4CAD1EF11EAA9944238C4F9AE02.roa
fileList[  2]:    61:2A:54:B6:4E:88:EA:7E:92:98:38:47:66:3B:A3:09:0E:C0:9B:C1:C4:05:4C:03:5F:D5:93:DC:0A:A9:EA:11 AEEF6AFEF2C711EB9D849B85C4F9AE02.roa
fileList[  3]:    AA:A5:B3:4F:6E:12:F5:BF:6F:F0:E5:DD:27:5D:C2:07:21:48:E5:F4:06:4C:11:83:AE:D3:58:EB:FD:5C:9C:2D 201867E6378B11EC8BB75C6CC4F9AE02.roa
fileList[  4]:    FF:BF:FC:CA:48:8E:37:F6:6F:75:A5:DD:2D:FE:DD:CA:72:BA:6B:F1:B1:7A:CC:CC:96:AB:30:8A:B6:D3:51:1D 6B4669CE7B7D11E99F99E776C4F9AE02.roa
fileList[  5]:    79:52:78:30:4B:8F:C0:DB:DC:9C:B1:C8:DE:CF:64:B0:9C:37:6B:CA:B3:DA:D3:98:39:9A:FB:C5:AF:68:14:E6 132F2CD0F5BA11EB90F4A13DC4F9AE02.roa
fileList[  6]:    5D:72:BD:74:DC:F1:A4:B7:FF:33:88:D2:E9:A9:ED:F8:D3:B4:BC:48:78:D9:D9:BF:36:2D:FE:EB:D1:E2:AA:E9 B53FD334575411ECAA094876C4F9AE02.roa
```

What is in a CRL

openssl crl -inform DER -text -noout -in QtPsuchXCrCQ62Ae2zN5wNPYptA.crl

```
Certificate Revocation List (CRL):
        Version 2 (0x1)
        Signature Algorithm: sha256WithRSAEncryption
        Issuer: CN = A913B805, serialNumber = 42D3ECB9C8570AB090EB601EDB3379C0D3D8A6D0
        Last Update: Jan 31 19:11:38 2022 GMT
        Next Update: Feb  2 19:11:38 2022 GMT
        CRL extensions:
            X509v3 Authority Key Identifier:
                keyid:42:D3:EC:B9:C8:57:0A:B0:90:EB:60:1E:DB:33:79:C0:D3:D8:A6:D0

            X509v3 CRL Number:
                11370
Revoked Certificates:
    Serial Number: 2004
        Revocation Date: Oct 28 02:45:07 2021 GMT
    Serial Number: 2005
        Revocation Date: Oct 28 01:44:52 2021 GMT
    Serial Number: 207A
        Revocation Date: Oct 28 06:25:12 2021 GMT
    Serial Number: 207C
        Revocation Date: Oct 28 02:25:15 2021 GMT
    Serial Number: 207E
        Revocation Date: Oct 28 02:05:13 2021 GMT
    Serial Number: 2080
        Revocation Date: Oct 28 02:45:08 2021 GMT
    Serial Number: 2083
        Revocation Date: Oct 28 06:25:12 2021 GMT
    Serial Number: 214F
        Revocation Date: Jan 28 19:07:34 2022 GMT
    Serial Number: 2150
        Revocation Date: Jan 29 07:01:01 2022 GMT
    Serial Number: 2151
        Revocation Date: Jan 29 19:11:34 2022 GMT
    Serial Number: 2152
        Revocation Date: Jan 30 07:04:26 2022 GMT
    Serial Number: 2153
        Revocation Date: Jan 30 19:12:06 2022 GMT
    Serial Number: 2154
        Revocation Date: Jan 31 07:09:46 2022 GMT
    Serial Number: 2155
        Revocation Date: Jan 31 19:11:38 2022 GMT
    Signature Algorithm: sha256WithRSAEncryption
         31:49:64:fc:76:f6:1a:67:21:1e:95:0f:85:72:b5:cf:db:cb:
         ff:75:47:cd:51:fe:56:63:fd:62:95:80:45:bf:d0:69:63:0f:
         90:d4:66:1a:d1:31:89:f1:b2:91:75:b9:2e:84:13:3a:3f:21:
         2b:40:c3:65:01:54:d2:2a:9a:f1:23:66:5b:b4:ce:48:d5:96:
         81:54:28:36:a0:e4:af:2d:51:c5:e0:e0:84:82:03:b4:c2:74:
         d9:66:f6:f7:a4:23:e1:8a:f0:10:06:7a:ab:ae:76:31:c4:79:
         3b:ba:69:41:c3:11:d4:5e:60:00:41:da:36:98:cd:14:0c:c2:
         0e:23:fc:3a:0a:9c:5a:33:b3:8a:81:4c:ca:16:15:22:13:b3:
         c9:32:66:16:42:0c:a2:66:99:0a:5c:84:aa:f3:38:5b:a2:d3:
         be:30:92:7f:bd:7a:67:cb:b2:3e:85:da:02:52:2b:29:ff:5f:
         ce:d2:b2:f8:74:34:16:5d:89:59:39:93:0b:82:d2:29:d9:e7:
         a0:3b:7c:4e:f8:51:0d:6d:93:5b:5e:fa:36:a5:72:ed:58:f3:
         5d:f8:7a:bf:fa:23:a0:08:1f:91:6b:bc:cb:6c:15:d9:b6:64:
         c2:9d:c7:72:2b:25:09:95:36:d1:36:40:a4:40:13:e7:3b:85:
         21:68:a2:1b
```

Global Validation

Providing information through the RPKI system

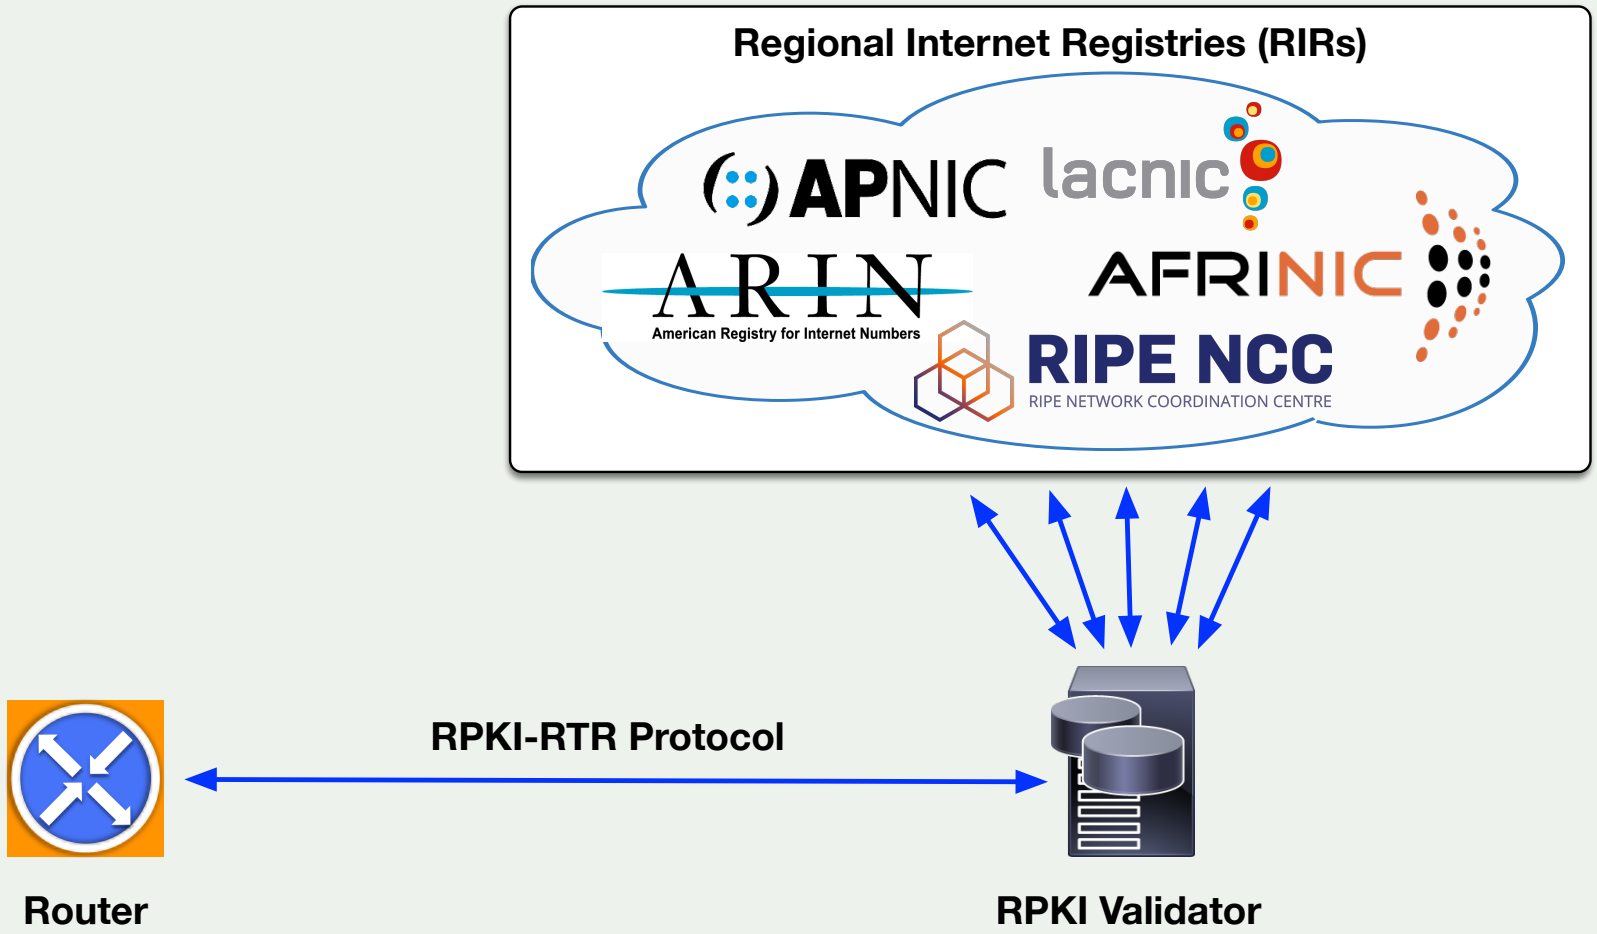


ROA
193.0.24.0/21
AS2121
Max Length: _
193.0.24.0/21
193.0.24.0/22
193.0.28.0/22
ROA
193.0.24.0/23
AS2121
Max Length: /24
ROA
193.0.30.0/23
AS2121
Max Length: _
/23
/23
/23
/23
/24
/24
/24
/24
/24
/24
/24
/24

Relying Party

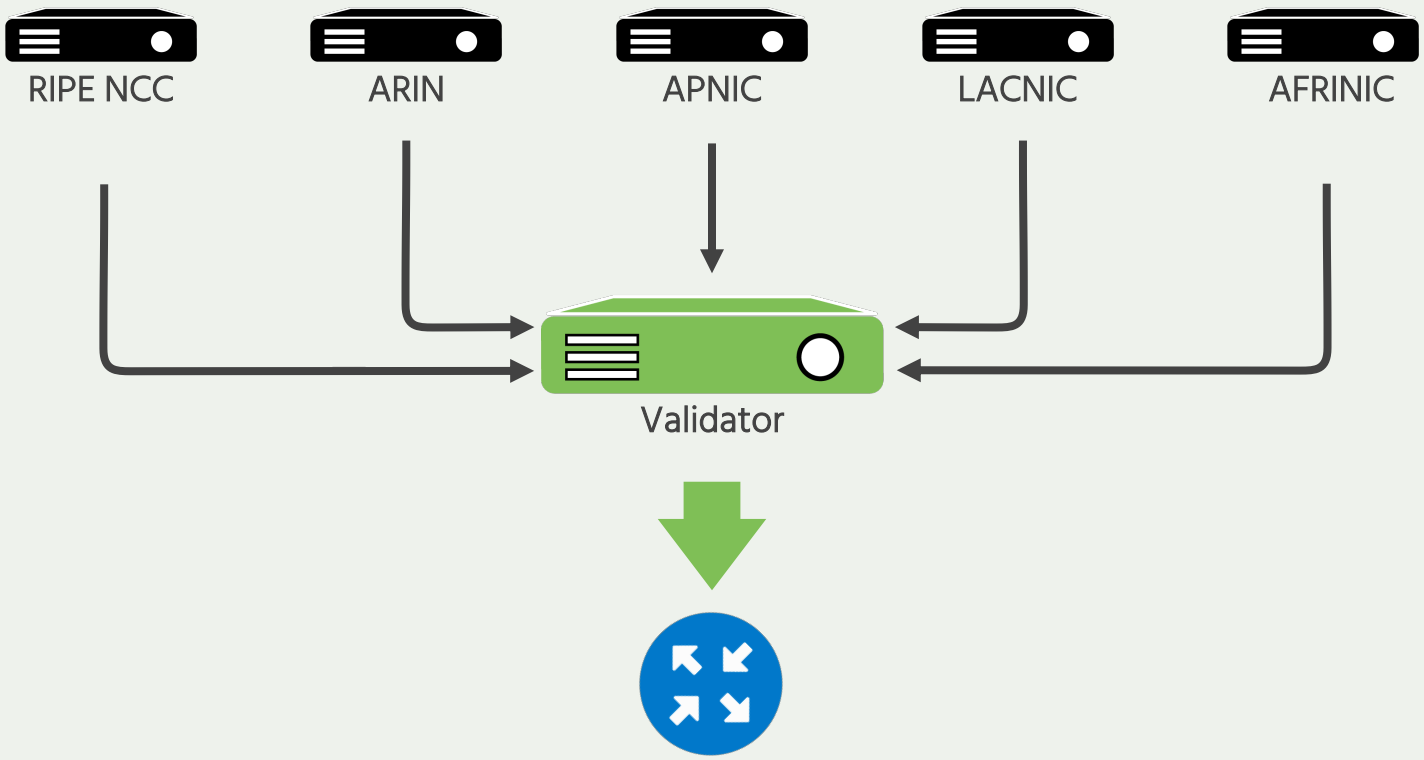

RPKI-RTR

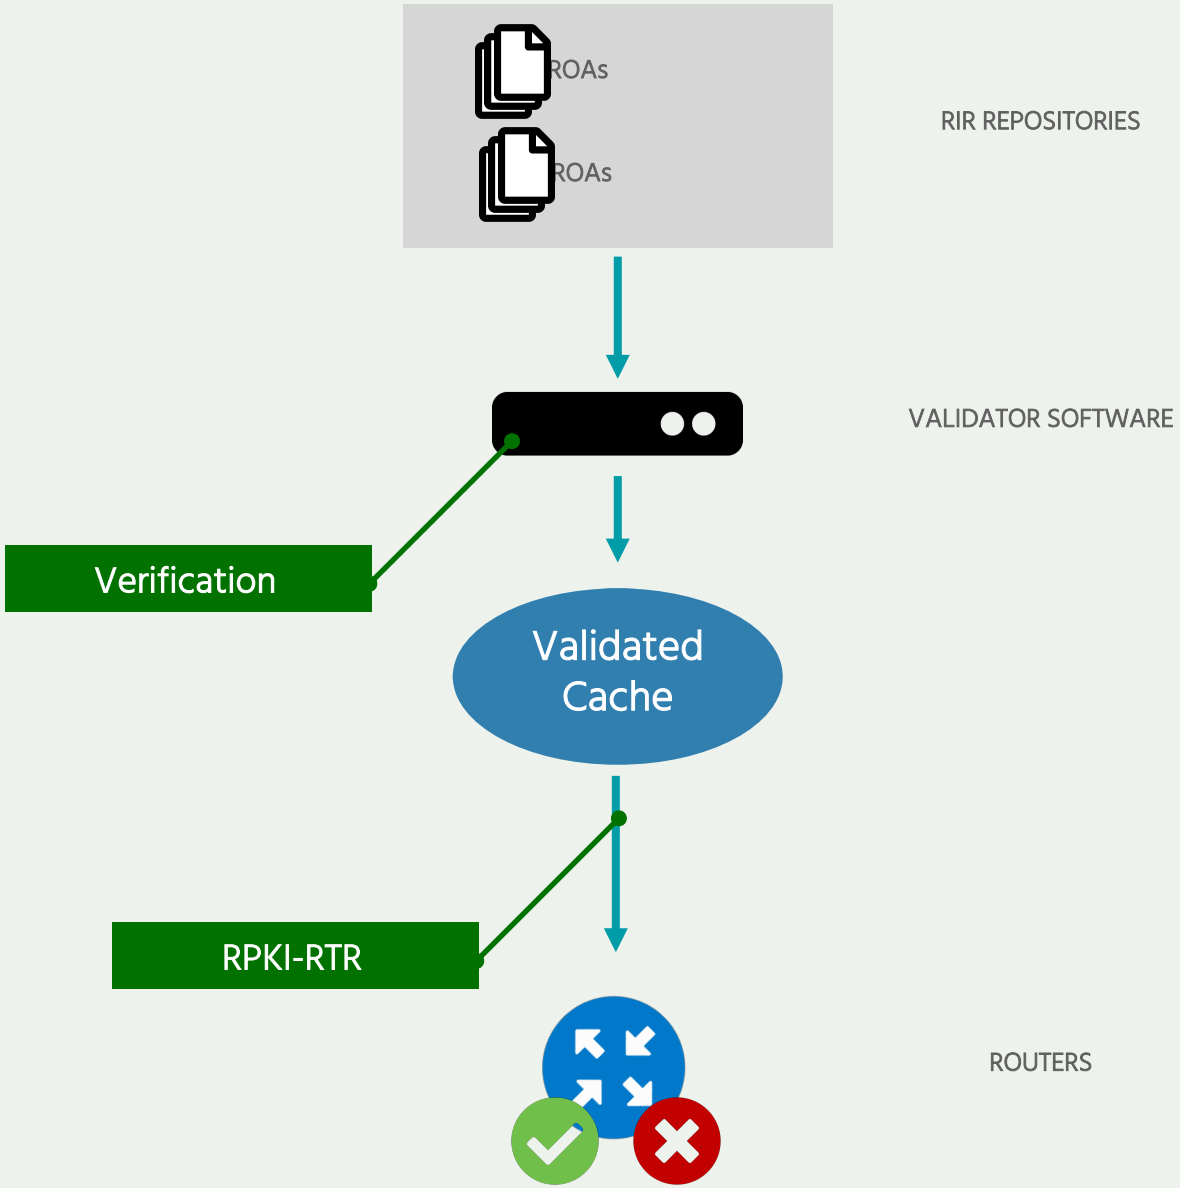
ROAs
ROAs
RIR REPOSITORIES
VALIDATOR SOFTWARE
Verification
Validated Cache
RPKI-RTR
ROUTERS

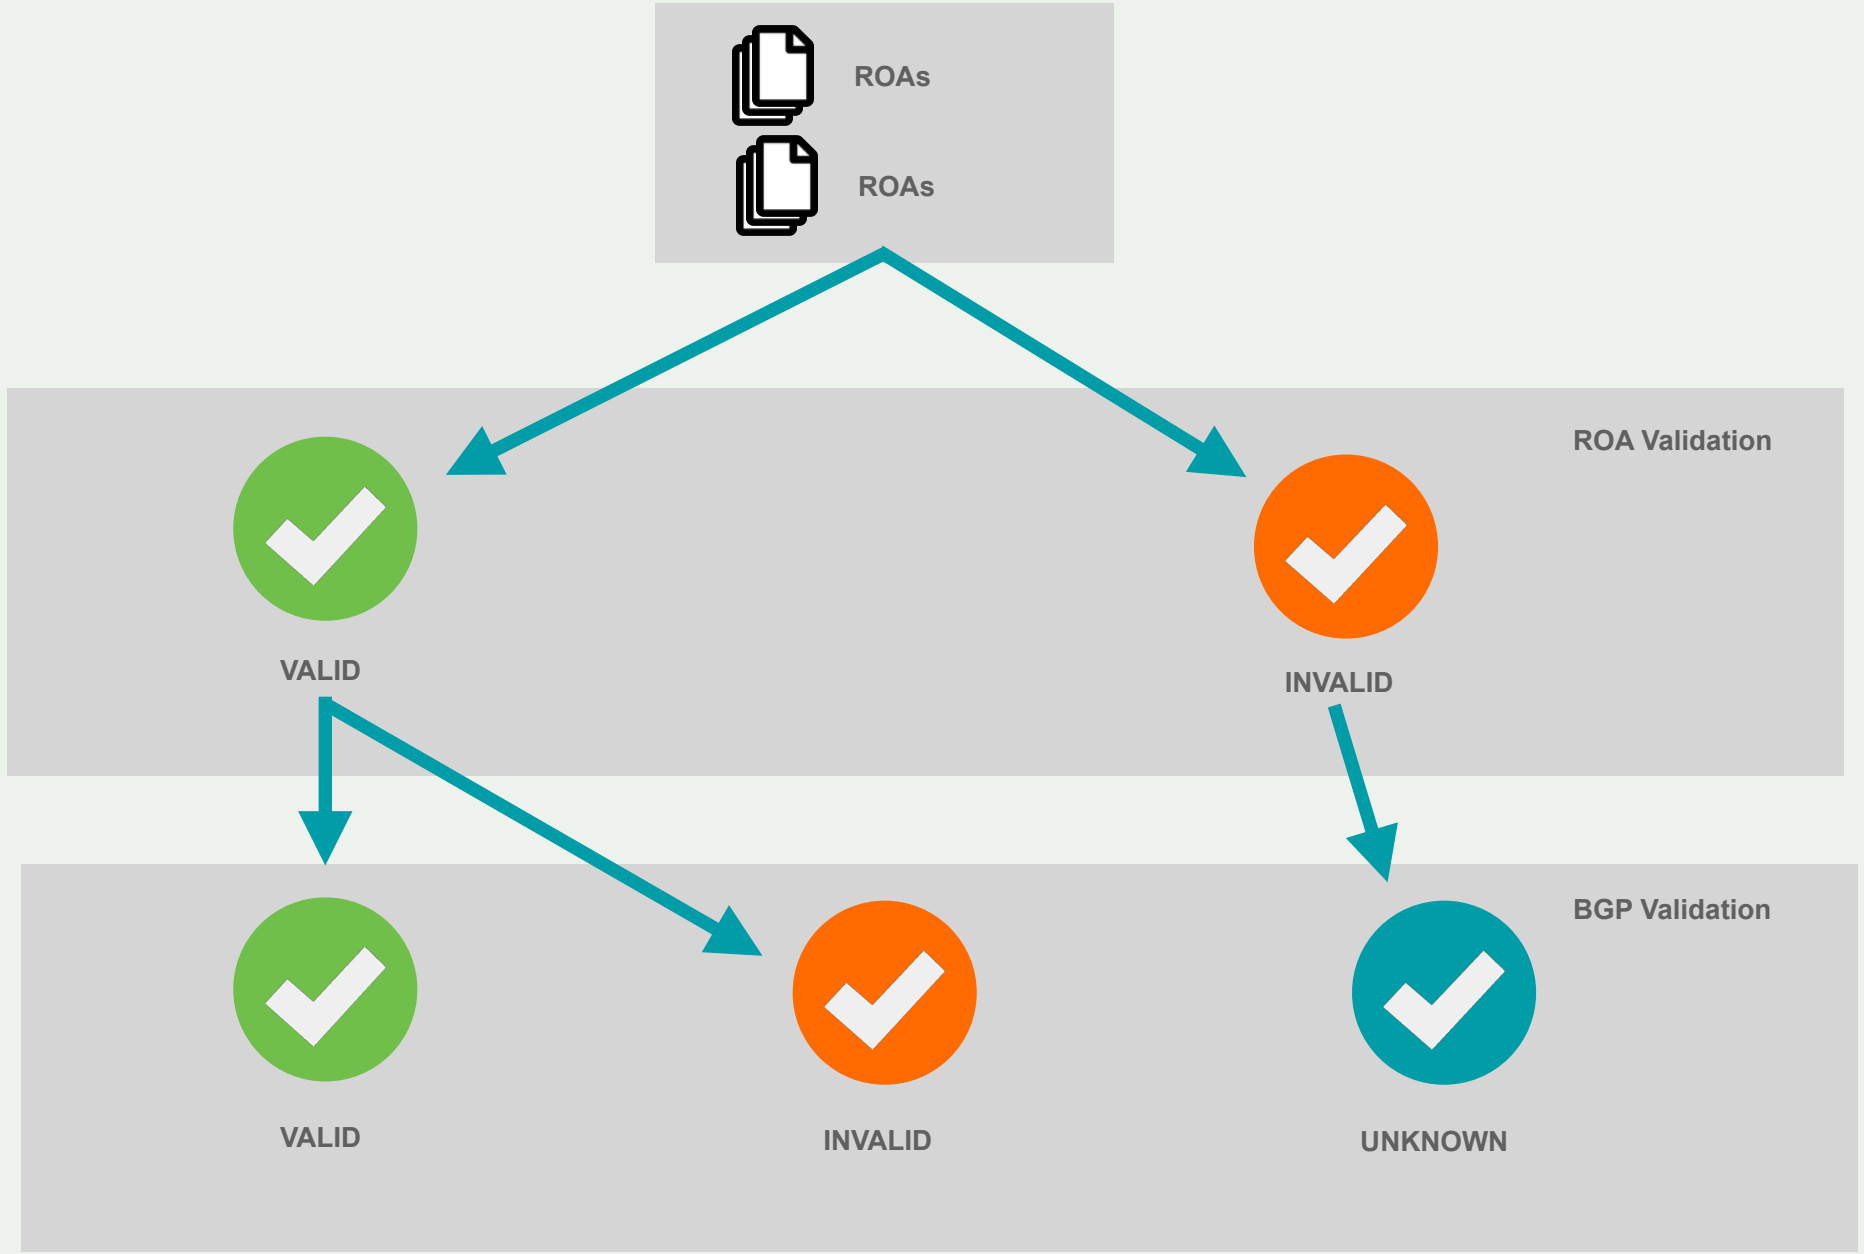
ROAs
ROAs
ROA Validation
VALID
INVALID
BGP Validation
VALID
INVALID
UNKNOWN

RPKI - The next steps

- ROA vs ROV
- What are the consequences of not having a valid ROA? Examples
- Why maintaining a valid ROA is important?
- Demo (APNIC Portal): Create and Delete ROA
- ROV - drop or not to drop invalids

Demo and Exercise

RPKI - The next steps

- Demo FRR/Junos implementing ROV with Routinator and Fort as RP

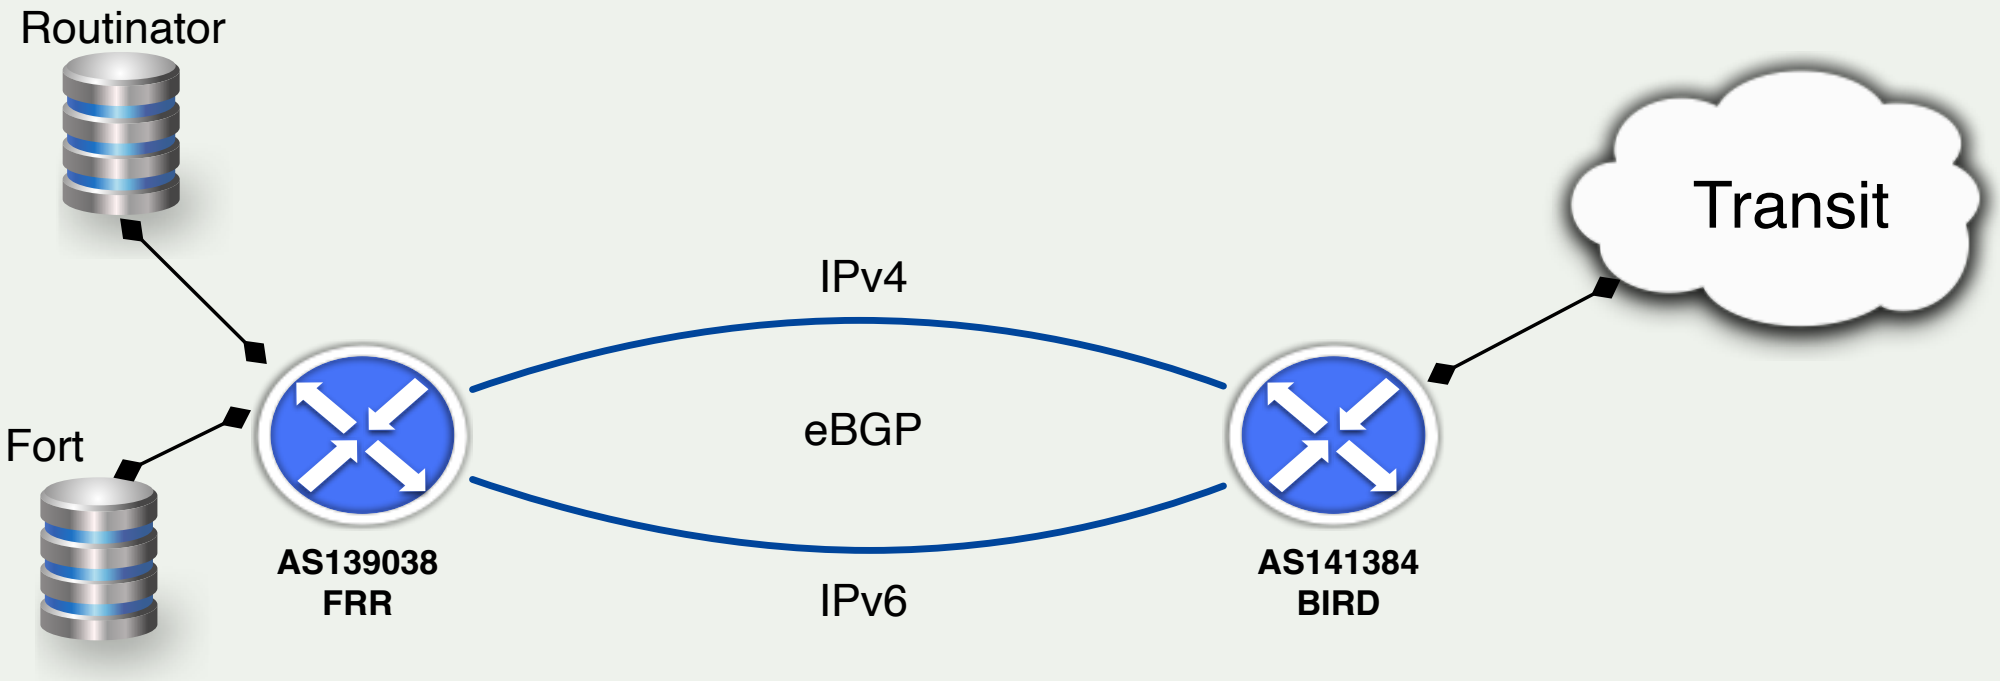

Wrapping up

Best Current Practice 38

- Defines some steps to take in order to have a “cleaner” routing table
- Restricting forged traffic (TCP and UDP)
- Implies the use of:
  - Prefix filters
  - Bogon filters
  - uRPF
- http://tools.ietf.org/html/bcp38

Future additional tools for enhancing filtering

- BGPSec
- ASPA (Draft)
- AS-Cones (Draft)

Additional resources and tools

- https://bgpfilterguide.nlnog.net/
- IRRToolkit (written in C++)
  - https://github.com/irrtoolset/irrtoolset
- Rpsltool (perl)
  - http://www.linux.it/~md/software
- Coloclue's Kees
  - https://github.com/coloclue/kees/
- IRR Explorer (web)
  - http://irrexplorer.nlnog.net
- IRR Power Tools (PHP)
  - http://sourceforge.net/projects/irrpt/

Questions ?