MANRS (Mutually Agreed Norms for Routing Security)

Pacific Network Operators Group (PacNOG 29)

03, Dec 2021 | 1 pm Fiji Time (GMT+12)

Naveen Lakshman, MANRS Fellow
naveen.k.ipv6(at)gmail.com

Kunal Krishnil Raj, MANRS Fellow
kunalkrishnilraj(at)gmail.com

Agenda

- Routing Security
  - Route Hijacks | Route Leaks | IP Spoofing
- Any Solution/s?
- MANRS for
  - Network Operators | IXP | Cloud and CDNs | Network Equipment Vendors
- MANRS Actions
  - Filtering | Anti Spoofing (uRPF/ACL) | Coordination | Global Validation (IRR/RPKI)
- MANRS ROA Tools
- MANRS Observatory (Partner View)
- MANRS Lab (Cisco IOS)
  - Filtering
  - Anti-Spoofing (ACL/uRPF)
    - Access Control Lists (ACL)
    - Unicast Reverse Path Forwarding (uRPF)

Acknowledgements

This tutorial is made of notes, configurations and diagrams from contributions by Aftab Siddiqui, Dr. Philip Smith, Tashi Phuntsho, Md. Zobair Khan, Anurag Bhatia and Musa Stephen Honlue.

The Problem

A Routing Security Overview

Routing Incident: (Vodafone~Idea AS55410 Hijack)

Vodafone Idea (AS55410) started originating 31,000+ routes which don't belong to them.

- Main Upstream leakers AS9498 (Bharti Airtel) and AS1273 (Vodafone UK)
- Spread mostly via IX connections
- Some of which re-propagated to their Peers (AS6939 HE)

Prefixes belonged to Google, Microsoft, Akamai, Fastly Cloudflare, and many others were affected.

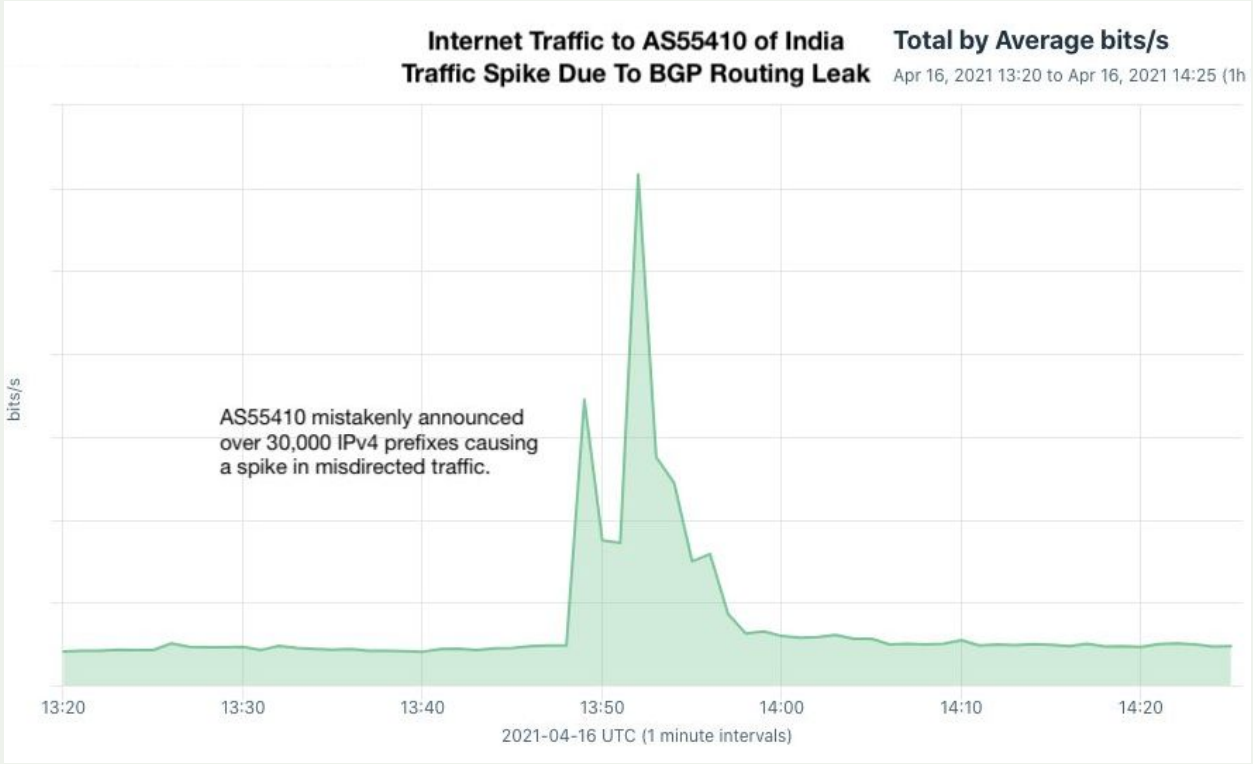

https://twitter.com/DougMadory/status/1383138595112955909

https://www.manrs.org/2021/04/a-major-bgp-hijack-by-as55410-vodafone-idea-ltd/
https://twitter.com/Qrator_Radar/status/1383383956511354882/photo/1

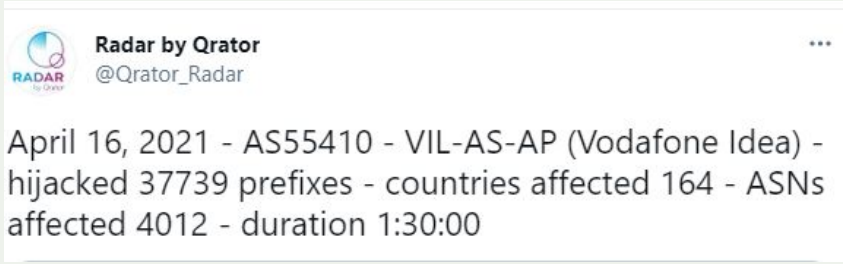

https://bgpstream.com/event/271636
https://bgpstream.com/event/271637

Series of Prefix Hijacks

Possible BGP hijack

Beginning at 2021-10-13 08:57:32 UTC, we detected a possible BGP hijack.
Prefix 45.128.160.0/22, is normally announced by AS47583 AS-HOSTINGER, CY.

But beginning at 2021-10-13 08:57:32, the same prefix (45.128.160.0/22) was also announced by ASN 212046.

This was detected by 4 BGPMon peers.

Expected
Start time: 2021-10-13 08:57:32 UTC
Expected prefix: 45.128.160.0/22
Expected ASN: 47583 (AS-HOSTINGER, CY)
Event Details
Detected advertisement: 45.128.160.0/22
Detected Origin ASN 212046 (MEZON-, LT)
Detected AS Path 51514 8455 13194 212046
Detected by number of BGPMon peers: 4

Detected by BGPStream

AS212046 - MEZON-LT - [LT] - Created Hijacks

RADAR by Qrator

Read FAQ

2021-10-13 08:58 UTC

You have received this letter because our system has detected **Created Hijacks** *possibly* global incidents for **AS212046**

Incident Type	**Created Hijacks**
Key ASN	AS212046 - MEZON-LT - [LT]
Overall Info	**Conflicts count all:** 1548 **ASNs affected:** 251 **Countries affected:** 16
Prefixes Info	**Prefixes created:** 1029 **Prefixes affected:** 1152
Propagation Info	**Max propagation:** 79%

October 13, 2021 — AS212046 — MEZON — hijacked 1029 prefixes, creating 1548 conflicts for 1152 prefixes and 251 ASNs in 16 countries. Maximum propagation: 79%. Duration: 1 hour.

Routing Incidents cause real world problems

Prefix/Route Hijacking

Route Leaks

IP address spoofing

Prefix/Route hijacking

Route hijacking, also known as “BGP hijacking,” is when a network operator or attacker (accidentally or deliberately) impersonates another network operator. This routes traffic to the wrong network operator, when another real route is available.

Example: The 2008 YouTube hijack; an attempt to block YouTube through route hijacking led to much of the traffic to YouTube being dropped around the world.

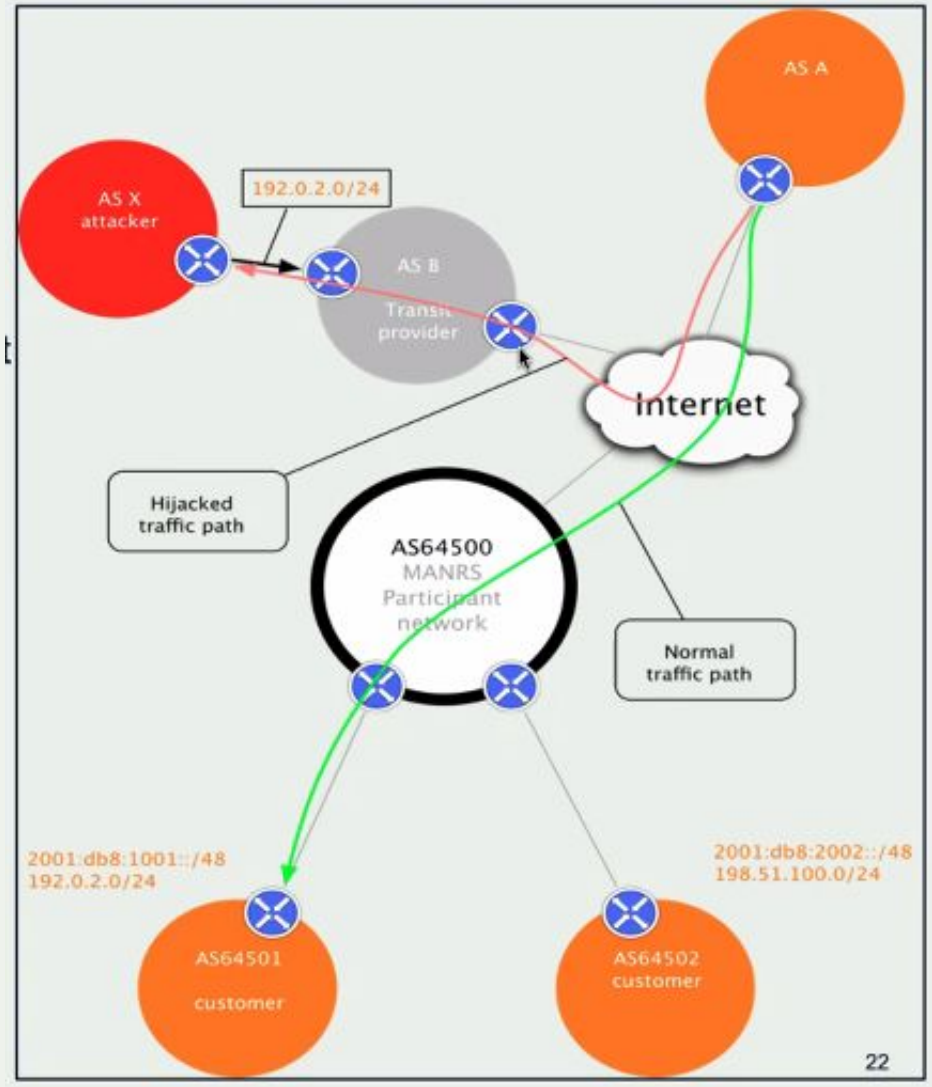

Route leak

When a network operator who is multi-homing (2 upstream) accidentally announces routes learned from one upstream to the other upstream. Customer AS becomes an intermediary (Hairpin turn leak), usually unintentional

Ex: June 2019. Allegheny leaked routes from another provider to Verizon, causing significant outage.

https://blog.cloudflare.com/how-verizon-and-a-bgp-optimizer- knocked-large-parts-of-the-internet-offline-today

Fix: Strong filtering policies (adjacent networks should strengthen their filtering policies to avoid accepting announcements that don't make sense).

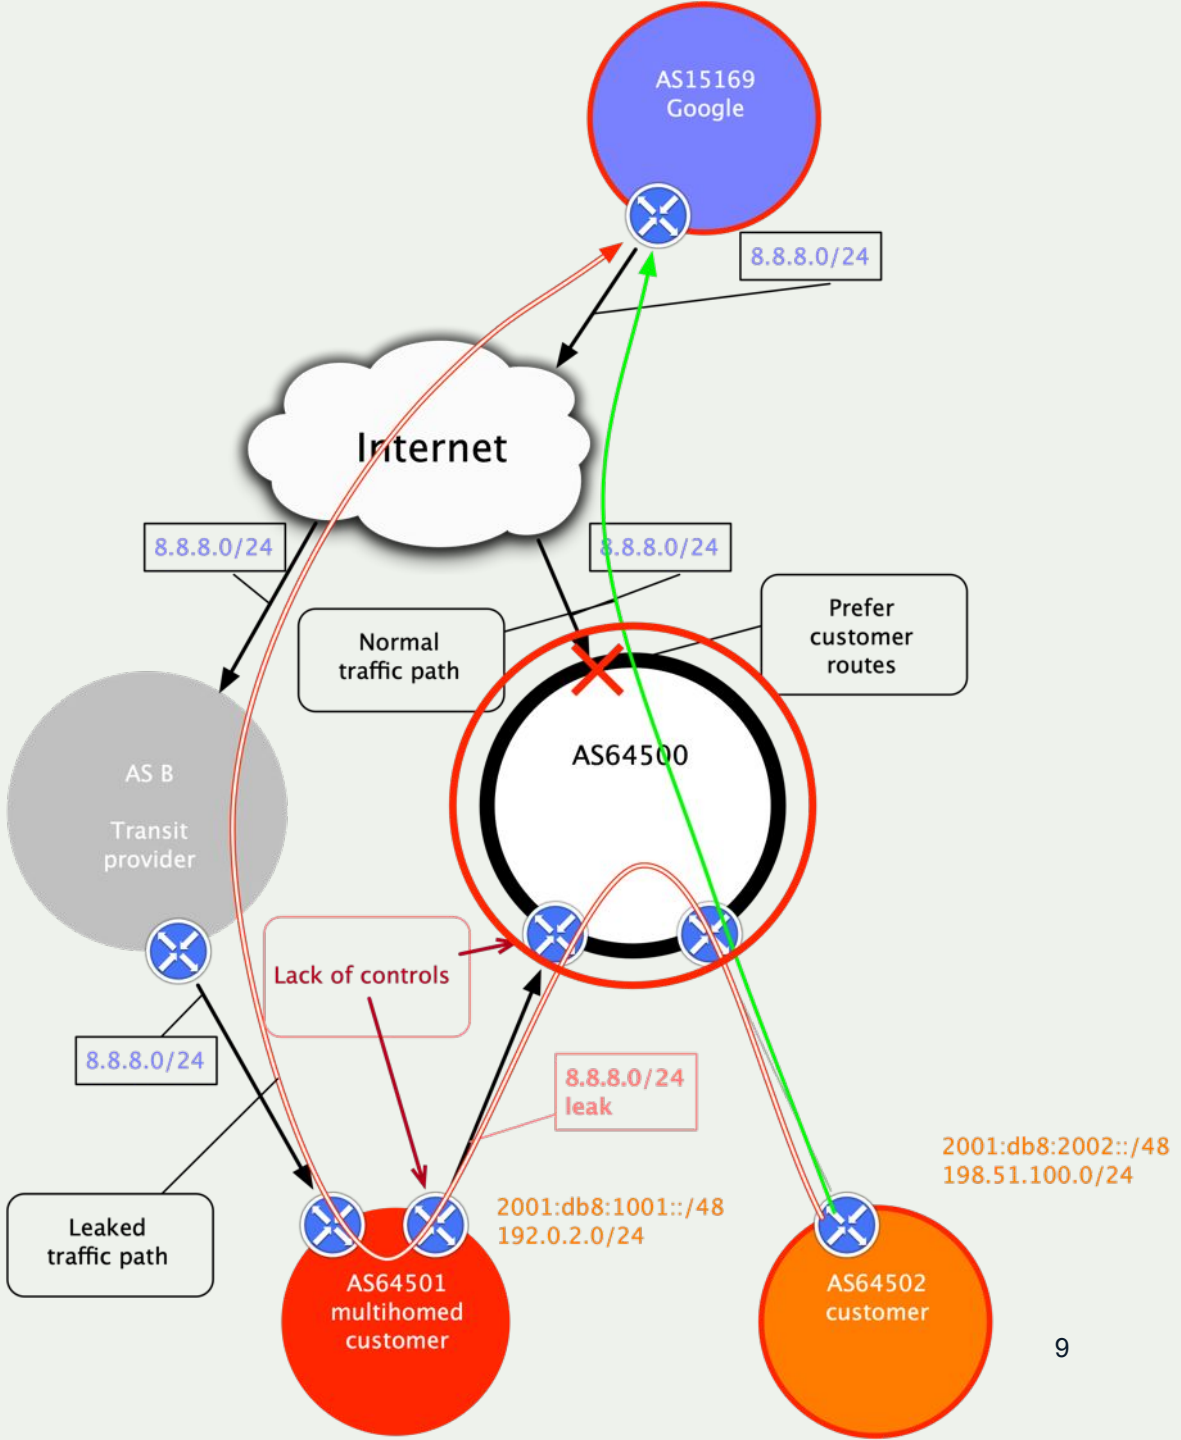

Routing Incidents (Pacific) May ~ November 2021

Event Details (Hijacks)	Prefixes affected	Reason
Expected Origin: (AS 7018) ATT-INTERNET4, US Detected Origin: (AS 18229) CTRLS-AS-IN CtrlS Datacenters Ltd., IN	172.0.0.0/12	More specific route 172.10.13.0/24
Expected Origin: (AS 21928) T-MOBILE-AS21928, US Detected Origin: (AS 9498) BBIL-AP BHARTI Airtel Ltd., IN	172.32.0.0/11	More specific route 172.32.0.0/23

Event Details (Leaks)	Prefixes affected
Origin AS: AS 132792 University of the Philippines, PH Leaked AS: AS 7473 SINGTEL-AS-AP, Singapore Leaked To: AS 6461 (Zayo, US)	202.92.152.0/24
Origin AS: AS 17639 Converge-AS, PH Leaked AS: AS7473 SINGTEL-AS-AP, Singapore Leaked To: AS 6461 (Zayo, US)	136.158.10.0/24

Source: bgpstream.com

Tools to Help

- Prefix and AS-Path filtering
- RPKI, IRR toolset, IRRPT, BGPQ3/4
- BGPSEC is standardized.

But...

Not enough deployment

We need a standard approach to improving routing security.

The Solution:

Mutually Agreed Norms for Routing Security (MANRS)

MANRS improves the security and reliability of the global Internet routing system, based on collaboration among participants and shared responsibility for the Internet infrastructure.

MANRS sets a new norm for routing security.

MANRS Programmes

Network Operators

Internet eXchange Points (IXP)

Content Delivery Networks (CDNs) and Cloud Providers

Network Equipment Vendors

MANRS Actions for Network Operators

Action 1: 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

Action 2: 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

Action 3: Coordination

Facilitate global operational communication and coordination between network operators

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

Action 4: Global Validation

Facilitate validation of routing information on a global scale

Publish your data so others can validate IRR/RPKI

MANRS Implementation Guide https://www.manrs.org/isps/bcop/

Blue shading = Mandatory Action

MANRS Actions for Internet Exchange Points (IXP)

Action 1
Prevent propagation of incorrect routing information

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

Provide encouragement or assistance for IXP members to implement MANRS actions.

Action 3
Protect the peering platform

Have a published policy of traffic not allowed on the peering fabric and perform filtering of such traffic.

Action 4
Facilitate global operational communication and coordination

Facilitate communication among members by providing necessary mailing lists and member directories.

Action 5
Provide monitoring and debugging tools to the members.

Provide a looking glass for IXP members.

MANRS Implementation Guide https://www.manrs.org/ixps/

Blue shading = Mandatory Action

MANRS Actions for CDNs & Cloud Providers

Action 1
Prevent propagation of incorrect routing information

Ensure correctness of own announcements and of their peers (non-transit) by implementing explicit (whitelist) filtering with prefix granularity.

Action 2
Prevent traffic with illegitimate source IP addresses

Implement anti-spoofing controls to prevent packets with illegitimate source IP address from leaving the network (egress filters).

Action 3
Facilitate global operational communication and coordination

Maintain globally accessible, up-to-date 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 (IRR and/or RPKI)

Action 5
Encourage MANRS adoption

Actively encourage MANRS adoption among the peers.

Action 6
Provide monitoring and debugging tools to the peering partners

Provide a mechanism to inform peering partners if announcements did not meet the requirements of the peering policy.

Blue shading = Mandatory Action

MANRS Implementation Guide https://www.manrs.org/cdn-cloud-providers/

MANRS Actions for Network Equipment Vendors

Action 1

Provide solutions for the implementation of specific MANRS Actions

Prevent propagation of Incorrect Anti-Spoofing, Protecting peering platform (Layer2)

Action 2

Promote MANRS through training & technical content

Promote MANRS through training and technical content

MANRS Implementation Guide https://www.manrs.org/equipment-vendors/

Blue shading = Mandatory Action

MANRS for Network Operators

Action 1: Filtering

Action 2: Anti-Spoofing

Action 3: Coordination

Action 4: Global Validation

Action 1: Filtering

- Threats to routing: Any network (running BGP) can announce/accept any IP prefix.
- Your first line of defence.
- You control what routes you are announcing
  - You have no control over what other networks announce
- To avoid issues, you have to decide what routes to accept from other networks.
- Operator defines a clear routing policy and implements filter that ensures correctness of own announcements and customers prefixes to adjacent networks.
- Operator applies due diligence when checking the correctness of its customer announcements, specifically that the customer legitimately holds the ASN and the address space it announces.

https://www.manrs.org/isps/guide/filtering/

Reference Operator Networks

- AS64510 - Transit Provider
- AS64500 - MANRS Participant
- AS64511 - Peer
- AS64501 - Customer 1
- AS64502 - Customer 2

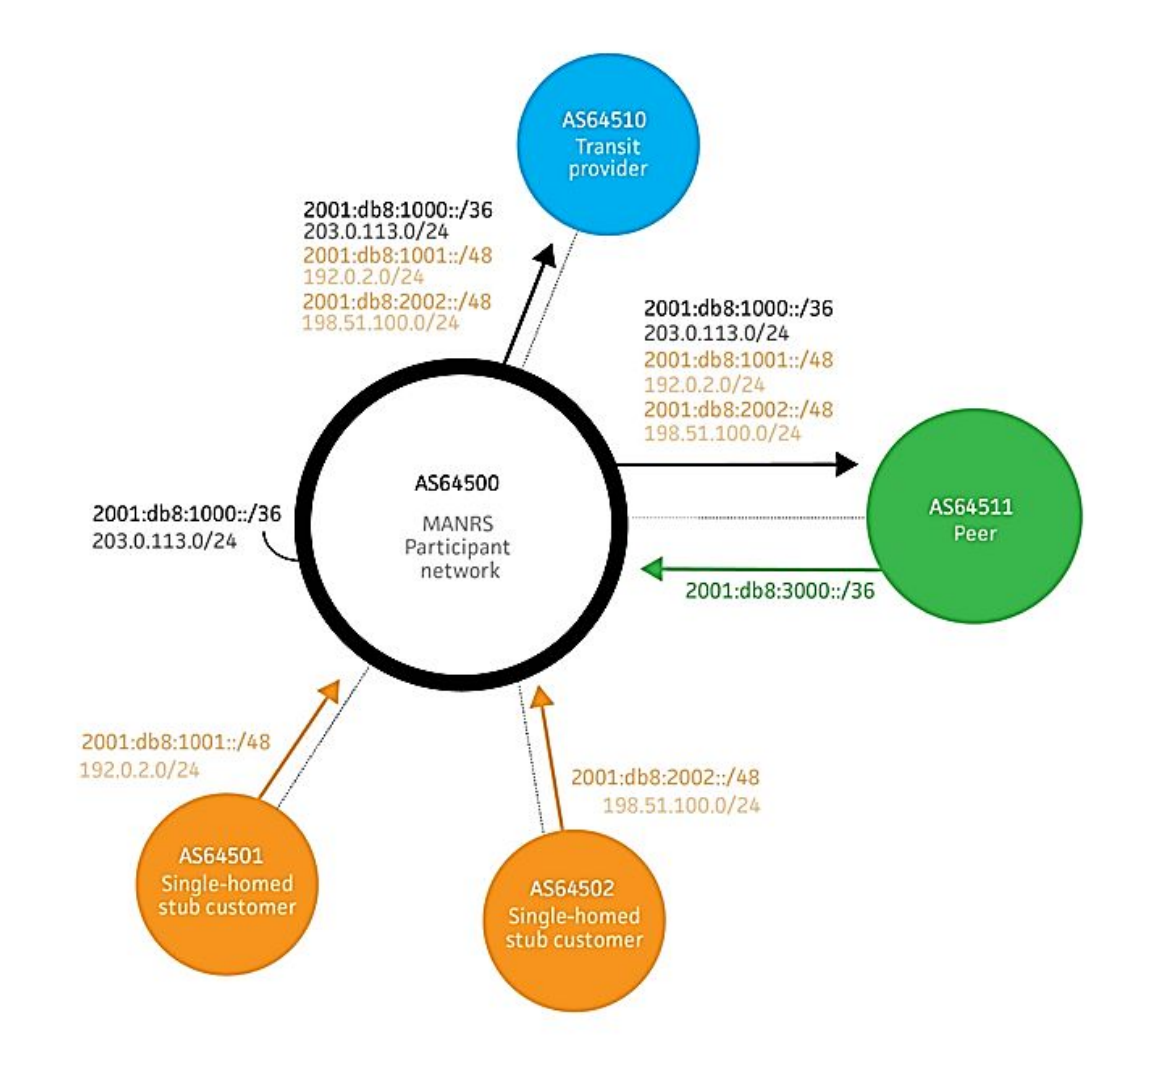

Implementing Filters

- In order to prevent propagation of incorrect routing information, MANRS participants are required to implement prefix filtering.
- By implementing prefix filtering, you can permit or deny announcements of certain prefixes from neighboring ASes.

Filtering

- Outbound filtering (prefixes you advertise to Transit, Peers and Customers)
- Inbound filtering (prefixes you receive from Transit, Peers and Customers)

Prefix Filtering (Outbound)

Outbound Filtering

- The configuration should ensure that only appropriate prefixes are sent.
- Prefixes belonging to both your network and its downstream/customer

Filter:

- Prefixes that are not globally routable
- Routes that are too specific (should neither be announced nor accepted by a BGP speaker /24 for IPv4, /48 for IPv6)
- The default route (not willing to receive it)

Prefix filtering (Outbound)

To Customer (Downstream)

- Default route
- Whole Internet routing table except default and bogons (special use addresses, rfc1918 unassigned blocks)

To Peers (other ISPs or Operators with whom you peer)

- Send your prefixes + your downstream customers or
- what you agreed to send

To Upstream/Transit Provider

- Your prefixes + your downstream customer prefixes

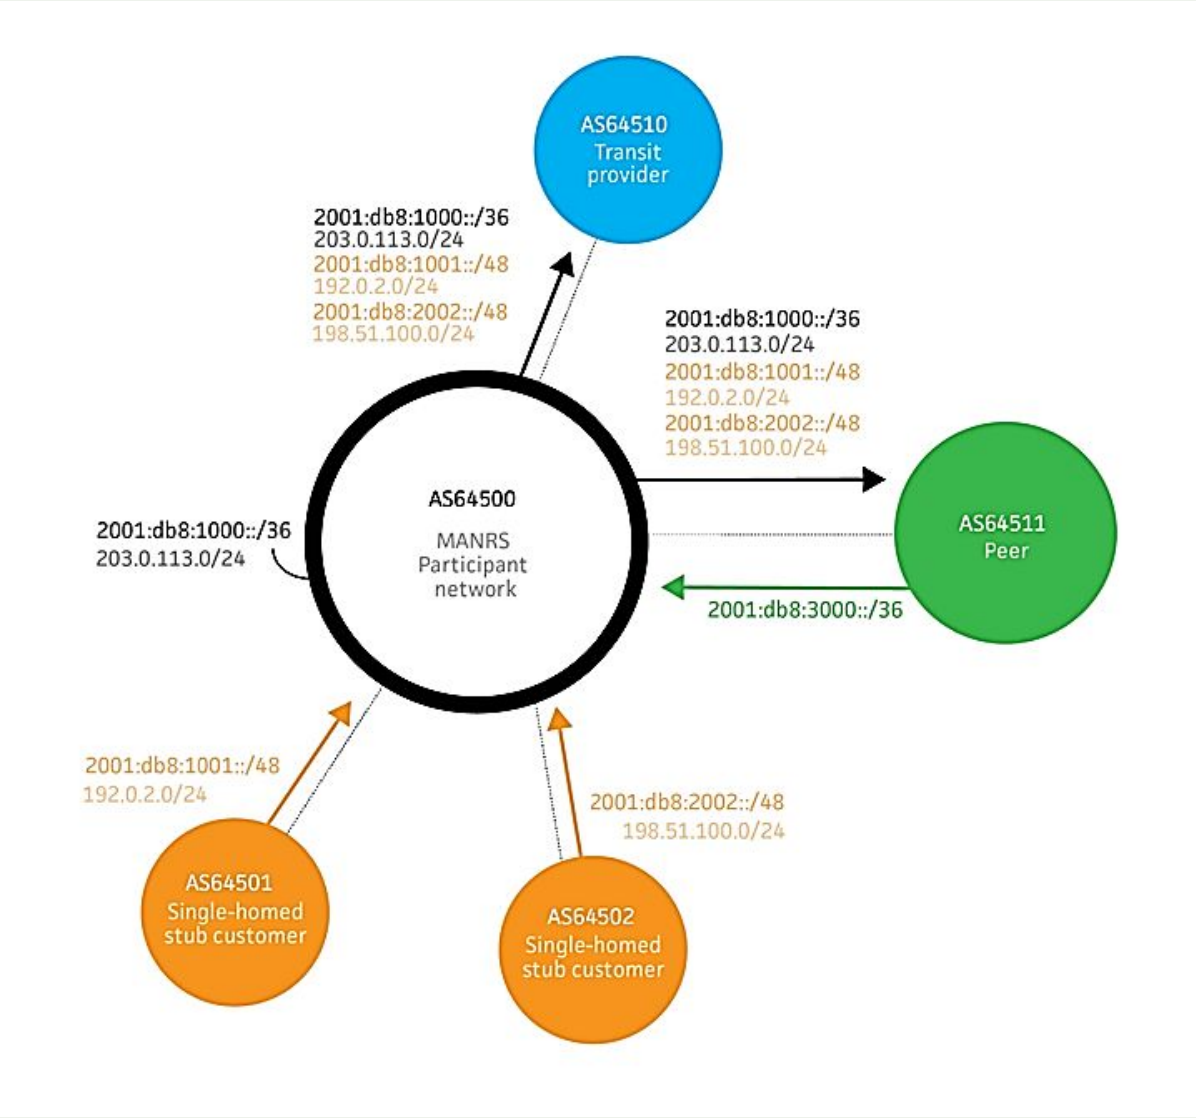

Prefix Filtering (Inbound)

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 (IRR)

Prefix filtering Inbound (Customers) ~(AS64501, AS64502)

- Customers speaking BGP, only accept prefixes registered to them, or prefixes registered to their customers
- If a provider has assigned address (PA) space to its customer, then the customer can announce back to his provider.
- If the Provider has NOT assigned address space to its customer, then:
  - Check in the five RIR databases to see if this address space really has been assigned to the customer

```
whois -h jwhois.apnic.net x.x.x.0/24
```

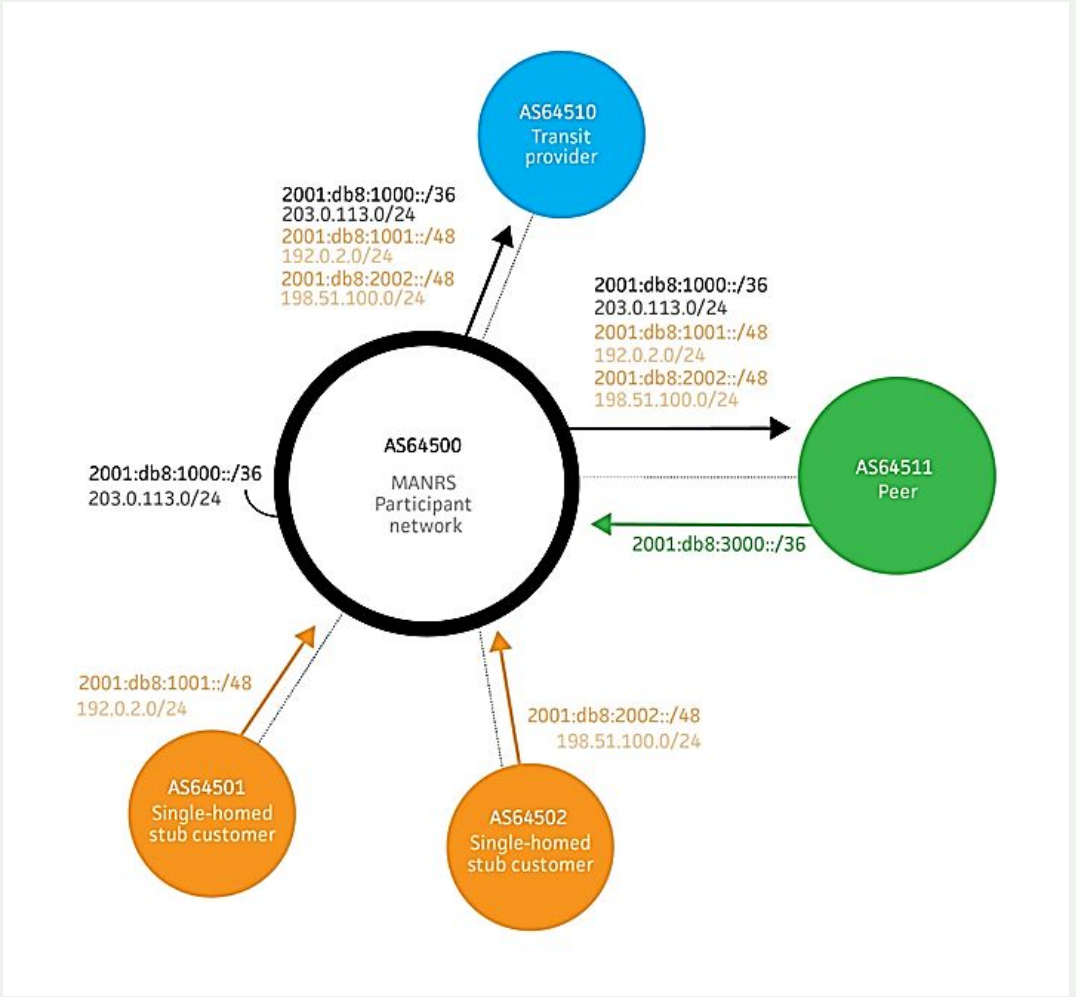

Prefix filtering Inbound (Customers)

Downstream/Customer has IPv4 192.0.2.0/24 and IPv6 2001:db8:1001::/48

Customer should only announce this to upstreams

```
router bgp 64500
address-family ipv4
  neighbor X.X.X.1 prefix-list AS64501-CUSTv4-IN in
address-family ipv6
  neighbor X:X:X:X::1 prefix-list AS64501-CUSTv6-IN in
!
ip prefix-list AS64501-CUSTv4-IN permit 192.0.2.0/24
ip prefix-list AS64501-CUSTv4-IN deny 0.0.0.0/0 le 32
!
ipv6 prefix-list AS64501-CUSTv6-IN permit 2001:db8:1001::/48
ipv6 prefix-list AS64501-CUSTv6-IN deny ::/0 le 128
```

Prefix filtering Inbound (Peers)

- Operators with whom you have agreed to exchange routes
- Only accept their prefixes and their downstream customers (or what was agreed)
- Verify they have the authority to route those prefixes (and their customers)
- Don't accept prefix length greater than /24 (IPv4) and /48 (IPv6)

Can use tools like bgpq3, bgpq4

https://github.com/snar/bgpq3

https://github.com/bgp/bgpq4

Configuration | Peer (Inbound)

If a peer has 2001:db8:3000::/36 and 198.51.0.0/22 prefixes

```
router bgp 64500
address-family ipv4
   neighbor X.X.X.1 prefix-list PEER-v4-IN in
address-family ipv6
   neighbor X:X:X:X::1 prefix-list PEER-v6-IN in
   exit

prefix list PEER-v4-IN permit 198.51.0.0/22 le 24
prefix list PEER-v4-IN deny 0.0.0.0/0 le 32
!
ipv6 prefix-list PEER-v6-IN permit 2001:db8:3000::/36 le 48
ipv6 prefix-list PEER-v6-IN deny ::/0 le 128
```

Prefix filters Inbound (Transit | Upstream)

If we want to receive just a Default Route

```
router bgp 64500
address-family ipv4
   neighbor X.X.X.1 prefix-list TRANSIT-DEFv4-IN in
address-family ipv6
   neighbor X:X:X::1 prefix-list TRANSIT-DEFv6-IN in
   exit
!
ip prefix-list TRANSIT-DEFv4-IN permit 0.0.0.0/0
!
ipv6 prefix-list TRANSIT-DEFv6-IN permit ::/0
```

From Upstream | Transit Provider

Full Internet prefixes/routes

- Don't accept your prefixes
- Don't accept bogon prefixes.
  - rfc1918, rfc6890 and unassigned address blocks
- Don't accept prefixes with length greater than /24 (IPv4) and /48 (IPv6)
- Don’t accept default

From Upstream | Transit (Configuration)

```
router bgp 64500
address-family ipv4
   neighbor X.X.X.1 prefix-list TRANSIT-FULL-v4 in
address-family ipv6
   neighbor X:X:X::1 prefix-list TRANSIT-FULL-v6 in
!
ip prefix-list TRANSIT-FULL-v4 deny 203.0.113.0/24 le 32
ip prefix-list TRANSIT-FULL-v4 deny 192.0.2.0/24 le 32
ip prefix-list TRANSIT-FULL-v4 deny 198.51.100.0/24 le 32
!
ip prefix-list TRANSIT-FULL-v4 permit 0.0.0.0/0 le 24

ipv6 prefix-list TRANSIT-FULL-v6 deny 2001:db8:1000::/36 le 128
ipv6 prefix-list TRANSIT-FULL-v6 deny  2001:db8:1001::/48 le 128
ipv6 prefix-list TRANSIT-FULL-v6 deny 2001:db8:2002::/48 le 128
!
ipv6 prefix-list TRANSIT-FULL-v6 permit ::/0 le 48
```

Prefix Filters - Inbound from Transit Provider (IPv4)

```
router bgp 64500
address-family ipv4
    neighbor X.X.X.1 prefix-list TRANSITv4-IN in
!
ip prefix-list TRANSITv4-IN deny 0.0.0.0/0                ! Default
ip prefix-list TRANSITv4-IN deny 0.0.0.0/8 le 32          ! Network Zero
ip prefix-list TRANSITv4-IN deny 10.0.0.0/8 le 32         ! RFC1918
ip prefix-list TRANSITv4-IN deny 100.64.0.0/10 le 32      ! RFC6598 shared address
ip prefix-list TRANSITv4-IN deny <your prefix>/X le 32    ! Your address space
ip prefix-list TRANSITv4-IN deny 127.0.0.0/8 le 32        ! Loopback
ip prefix-list TRANSITv4-IN deny 169.254.0.0/16 le 32     ! APIPA
ip prefix-list TRANSITv4-IN deny 172.16.0.0/12 le 32      ! RFC1918
ip prefix-list TRANSITv4-IN deny 192.0.0.0/24 le 32       ! IETF Protocol
ip prefix-list TRANSITv4-IN deny 192.0.2.0/24 le 32       ! TEST1
ip prefix-list TRANSITv4-IN deny 192.168.0.0/16 le 32     ! RFC1918
ip prefix-list TRANSITv4-IN deny 198.18.0.0/15 le 32      ! Benchmarking
ip prefix-list TRANSITv4-IN deny 198.51.100.0/24 le 32    ! TEST2
ip prefix-list TRANSITv4-IN deny 203.0.113.0/24 le 32     ! TEST3
ip prefix-list TRANSITv4-IN deny 224.0.0.0/4 le 32        ! Multicast
ip prefix-list TRANSITv4-IN deny 240.0.0.0/4 le 32        ! Future Use
ip prefix-list TRANSITv4-IN deny 0.0.0.0/0 ge 25          ! Prefixes longer than /24
ip prefix-list TRANSITv4-IN permit 0.0.0.0/0 le 32 (/24)
```

Prefix Filters - Inbound from Transit Provider (IPv6)

```
router bgp 64500
address-family ipv6
    neighbor X:X:X:X::1 prefix-list TRANSITv6-IN in
!
ipv6 prefix-list TRANSITv6-IN deny 2001::/32 le 128          ! Teredo subnets (rfc4380)
ipv6 prefix-list TRANSITv6-IN deny 2001:db8::/32 le 128      ! Documentation (rfc3849)
ipv6 prefix-list TRANSITv6-IN deny 2002::/16 le 128          ! 6to4 subnets (rfc3056)
ipv6 prefix-list TRANSITv6-IN deny <your::/32> le 128        ! Your prefix
ipv6 prefix-list TRANSITv6-IN deny 3ffe::/16 le 128          ! Old 6bone
ipv6 prefix-list TRANSITv6-IN deny fc00::/7 le 128           ! ULA (rfc4193, rfc8190)
ipv6 prefix-list TRANSITv6-IN deny fe00::/9 le 128           ! Reserved IETF
ipv6 prefix-list TRANSITv6-IN deny fe80::/10 le 128          ! Link-local (rfc4291)
ipv6 prefix-list TRANSITv6-IN deny fec0::/10 le 128          ! Reserved IETF
ipv6 prefix-list TRANSITv6-IN deny ff00::/8 le 128           ! Multicast
ipv6 prefix-list TRANSITv6-IN permit 2000::/3 le 48          ! Global Unicast
ipv6 prefix-list TRANSITv6-IN deny ::/0 le 128
```

AS Path filtering

We can limit the AS PATH in announced prefixes using BGP AS path filter. The regular expression ^$ in ACL statement matches empty AS_PATH thus it allows only locally announced prefixes being sent to ISP.

Customer1 (AS64501)

```
ip as-path access-list 10 permit ^$
!
router bgp 64501
neighbor x.x.x.x filter-list 10 out
```

Provider (MANRS Participant)

```
ip as-path access-list 10 permit ^64501$
!
router bgp 64500
neighbor x.x.x.x filter-list 10 in
```

Max Prefix Filtering (BCP 194)

It is recommended to configure a limit on the number of routes to be accepted from a peer. The following rules are generally RECOMMENDED:

- From peers, have a limit lower than the number of routes in the Internet. This will shut down the BGP peering if the peer suddenly advertises the full table.
- From upstreams that provide full routing, it is RECOMMENDED to have a limit higher than the number of routes in the Internet. A limit is still useful in order to protect the network (and in particular, the routers' memory) if too many routes are sent by the upstream.

Maximum Prefix Limit

```
router bgp ASN
address family ip [v4|v6]
neighbor <peer addr|group > maximum prefix <max value> [threshold][restart N]
[warning only]
```

- Drop the peering if more than 3000 prefixes received

```
address-family ipv4
neighbor X.X.X.1 maximum prefix 3000
!
address-family ipv6
neighbor X:X:X:X::1 maximum-prefix 3000
```

- Log a warning when it receives more than 3000 prefixes

```
neighbor X.X.X.1 maximum prefix 3000 warning only
```

- Restart the peering session automatically after 30 minutes

```
neighbor X.X.X.1 maximum prefix 3000 restart 30
```

Prefix lists - Tools

Tools are there to help you

- bgpq3/bgpq4
- Level3 Filtergen

```
┌──(naveen㉿ LAPTOP-6VNOIOAD)-[~]
└─$ bgpq4 -S APNIC -l AS10075-v4-in AS10075
no ip prefix-list AS10075-v4-in
ip prefix-list AS10075-v4-in permit 103.7.248.0/22
ip prefix-list AS10075-v4-in permit 103.7.250.0/24
ip prefix-list AS10075-v4-in permit 103.7.251.0/24
ip prefix-list AS10075-v4-in permit 103.131.156.0/22
ip prefix-list AS10075-v4-in permit 103.131.156.0/24
ip prefix-list AS10075-v4-in permit 103.131.157.0/24
ip prefix-list AS10075-v4-in permit 103.131.158.0/24
ip prefix-list AS10075-v4-in permit 103.131.159.0/24
ip prefix-list AS10075-v4-in permit 103.229.82.0/23
ip prefix-list AS10075-v4-in permit 163.47.156.0/22
ip prefix-list AS10075-v4-in permit 163.47.156.0/23
ip prefix-list AS10075-v4-in permit 163.47.156.0/24
ip prefix-list AS10075-v4-in permit 163.47.157.0/24
ip prefix-list AS10075-v4-in permit 163.47.158.0/24
ip prefix-list AS10075-v4-in permit 163.47.159.0/24
```

bgpq3: Cisco, Juniper, Bird

by default bgpq3 generates configuration based on RADB data

bgpq4: Nokia/SR, Arista, Mikrotik, Huawei

by default bgpq4 generates configuration based on NTT's IRR

RIR maintained databases (AFRINIC, ARIN, APNIC, LACNIC and RIPE) shall be trusted more than the others because they have latest update about which address space allocated to ASes. Encouraged to use **'-S**' flag to limit database sources to only ones they trust

https://github.com/snar/bgpq3 https://github.com/bgp/bgpq4

Tools (BGPQ4) ~ (Juniper, Huawei, Nokia)

```
┌──(naveen㉿ LAPTOP-6VNOIOAD)-[~]
└─$ bgpq4 -l AS10074-v4-in AS10075 -J -S APNIC
policy-options {
replace:
 prefix-list AS10074-v4-in {
    103.7.248.0/22;
    103.7.250.0/24;
    103.7.251.0/24;
    103.131.156.0/22;
    103.131.156.0/24;
    103.131.157.0/24;
    103.131.158.0/24;
    103.131.159.0/24;
    103.229.82.0/23;
    163.47.156.0/22;
    163.47.156.0/23;
    163.47.156.0/24;
    163.47.157.0/24;
    163.47.158.0/24;
    163.47.159.0/24;
 }
}
```

```
┌──(naveen㉿ LAPTOP-6VNOIOAD)-[~]
└─$ bgpq4 -l AS10074-v4-in AS10075 -U -S APNIC
undo ip ip-prefix AS10074-v4-in
ip ip-prefix AS10074-v4-in permit 103.7.248.0 22
ip ip-prefix AS10074-v4-in permit 103.7.250.0 24
ip ip-prefix AS10074-v4-in permit 103.7.251.0 24
ip ip-prefix AS10074-v4-in permit 103.131.156.0 22
ip ip-prefix AS10074-v4-in permit 103.131.156.0 24
ip ip-prefix AS10074-v4-in permit 103.131.157.0 24
ip ip-prefix AS10074-v4-in permit 103.131.158.0 24
ip ip-prefix AS10074-v4-in permit 103.131.159.0 24
ip ip-prefix AS10074-v4-in permit 103.229.82.0 23
ip ip-prefix AS10074-v4-in permit 163.47.156.0 22
ip ip-prefix AS10074-v4-in permit 163.47.156.0 23
ip ip-prefix AS10074-v4-in permit 163.47.156.0 24
ip ip-prefix AS10074-v4-in permit 163.47.157.0 24
ip ip-prefix AS10074-v4-in permit 163.47.158.0 24
ip ip-prefix AS10074-v4-in permit 163.47.159.0 24
```

```
┌──(naveen㉿ LAPTOP-6VNOIOAD)-[~]
└─$ bgpq4 -l AS10074-v4-in AS10075 -N -S APNIC
configure router policy-options
begin
no prefix-list "AS10074-v4-in"
prefix-list "AS10074-v4-in"
    prefix 103.7.248.0/22 exact
    prefix 103.7.250.0/24 exact
    prefix 103.7.251.0/24 exact
    prefix 103.131.156.0/22 exact
    prefix 103.131.156.0/24 exact
    prefix 103.131.157.0/24 exact
    prefix 103.131.158.0/24 exact
    prefix 103.131.159.0/24 exact
    prefix 103.229.82.0/23 exact
    prefix 163.47.156.0/22 exact
    prefix 163.47.156.0/23 exact
    prefix 163.47.156.0/24 exact
    prefix 163.47.157.0/24 exact
    prefix 163.47.158.0/24 exact
    prefix 163.47.159.0/24 exact
exit
commit
```

BGPQ4 (Mikrotik, Arista, Bird)

```
┌──(naveen㉿ LAPTOP-6VNOIOAD)-[~]
└─$ bgpq4 -l AS10074-v4-in AS10075 -K -S APNIC
/routing filter add action=accept chain="AS10074-v4-in-V4" prefix=103.7.248.0/22
/routing filter add action=accept chain="AS10074-v4-in-V4" prefix=103.7.250.0/24
/routing filter add action=accept chain="AS10074-v4-in-V4" prefix=103.7.251.0/24
/routing filter add action=accept chain="AS10074-v4-in-V4" prefix=103.131.156.0/22
/routing filter add action=accept chain="AS10074-v4-in-V4" prefix=103.131.156.0/24
/routing filter add action=accept chain="AS10074-v4-in-V4" prefix=103.131.157.0/24
/routing filter add action=accept chain="AS10074-v4-in-V4" prefix=103.131.158.0/24
/routing filter add action=accept chain="AS10074-v4-in-V4" prefix=103.131.159.0/24
/routing filter add action=accept chain="AS10074-v4-in-V4" prefix=103.229.82.0/23
/routing filter add action=accept chain="AS10074-v4-in-V4" prefix=163.47.156.0/22
/routing filter add action=accept chain="AS10074-v4-in-V4" prefix=163.47.156.0/23
/routing filter add action=accept chain="AS10074-v4-in-V4" prefix=163.47.156.0/24
/routing filter add action=accept chain="AS10074-v4-in-V4" prefix=163.47.157.0/24
/routing filter add action=accept chain="AS10074-v4-in-V4" prefix=163.47.158.0/24
/routing filter add action=accept chain="AS10074-v4-in-V4" prefix=163.47.159.0/24
```

```
┌──(naveen㉿ LAPTOP-6VNOIOAD)-[~]
└─$ bgpq4 -l AS10074-v4-in AS10075 -e -S APNIC
no ip prefix-list AS10074-v4-in
ip prefix-list AS10074-v4-in
    seq 1 permit 103.7.248.0/22
    seq 2 permit 103.7.250.0/24
    seq 3 permit 103.7.251.0/24
    seq 4 permit 103.131.156.0/22
    seq 5 permit 103.131.156.0/24
    seq 6 permit 103.131.157.0/24
    seq 7 permit 103.131.158.0/24
    seq 8 permit 103.131.159.0/24
    seq 9 permit 103.229.82.0/23
    seq 10 permit 163.47.156.0/22
    seq 11 permit 163.47.156.0/23
    seq 12 permit 163.47.156.0/24
    seq 13 permit 163.47.157.0/24
    seq 14 permit 163.47.158.0/24
    seq 15 permit 163.47.159.0/24
```

BGPQ3 (IPv6 Filters) ~ (Bird, Cisco, Juniper)

Bird

```
naveen@manrs:~$
naveen@manrs:~$ bgpq3 -bl AS10075-v6-in AS10075 -6
AS10075-v6-in = [
    2403:cd40::/32
];
naveen@manrs:~$
```

Cisco

```
naveen@manrs:~$
naveen@manrs:~$ bgpq3 -l AS10075-v6-in AS10075 -6
no ipv6 prefix-list AS10075-v6-in
ipv6 prefix-list AS10075-v6-in permit 2403:cd40::/32
naveen@manrs:~$
```

Juniper

```
naveen@manrs:~$
naveen@manrs:~$ bgpq3 -Jl AS10075-v6-in AS10075 -6
policy-options {
replace:
 prefix-list AS10075-v6-in {
    2403:cd40::/32;
 }
}
naveen@manrs:~$
```

BGPQ4

(AS Path access-list using AS-SET)

```
┌──(naveen LAPTOP-6VNOIOAD)-[~]
└─$ bgpq4 -f 10075 -l AS10075-in AS-FGL
no ip as-path access-list AS10075-in
ip as-path access-list AS10075-in permit ^10075(_10075)*$
ip as-path access-list AS10075-in permit ^10075(_[0-9]+)*_(7565|7690|9230|9288)$
ip as-path access-list AS10075-in permit ^10075(_[0-9]+)*_(9441|9451|9651|9723)$
ip as-path access-list AS10075-in permit ^10075(_[0-9]+)*_(9825|9832|13335|17469)$
ip as-path access-list AS10075-in permit ^10075(_[0-9]+)*_(17471|17641|17806|17819)$
ip as-path access-list AS10075-in permit ^10075(_[0-9]+)*_(18022|18109|18230|18715)$
ip as-path access-list AS10075-in permit ^10075(_[0-9]+)*_(23688|23893|23923|23956)$
ip as-path access-list AS10075-in permit ^10075(_[0-9]+)*_(23991|24050|24122|24342)$
ip as-path access-list AS10075-in permit ^10075(_[0-9]+)*_(24389|24432|24481|24556)$
ip as-path access-list AS10075-in permit ^10075(_[0-9]+)*_(37972|38011|38017|38023)$
ip as-path access-list AS10075-in permit ^10075(_[0-9]+)*_(38026|38030|38031|38036)$
ip as-path access-list AS10075-in permit ^10075(_[0-9]+)*_(38054|38067|38069|38071)$
ip as-path access-list AS10075-in permit ^10075(_[0-9]+)*_(38137|38138|38192|38200)$
ip as-path access-list AS10075-in permit ^10075(_[0-9]+)*_(38203|38210|38212|38256)$
ip as-path access-list AS10075-in permit ^10075(_[0-9]+)*_(38267|38313|38315|38493)$
ip as-path access-list AS10075-in permit ^10075(_[0-9]+)*_(38555|38556|38558|38562)$
ip as-path access-list AS10075-in permit ^10075(_[0-9]+)*_(38588|38592|38614|38712)$
ip as-path access-list AS10075-in permit ^10075(_[0-9]+)*_(38721|38744|45176|45245)$
ip as-path access-list AS10075-in permit ^10075(_[0-9]+)*_(45273|45276|45326|45766)$
ip as-path access-list AS10075-in permit ^10075(_[0-9]+)*_(45904|45905|45925|45951)$
ip as-path access-list AS10075-in permit ^10075(_[0-9]+)*_(55344|55406|55473|55492)$
ip as-path access-list AS10075-in permit ^10075(_[0-9]+)*_(55531|55550|55708|55733)$
ip as-path access-list AS10075-in permit ^10075(_[0-9]+)*_(55828|56054|56115|56121)$
ip as-path access-list AS10075-in permit ^10075(_[0-9]+)*_(56138|56264|58445|58508)$
ip as-path access-list AS10075-in permit ^10075(_[0-9]+)*_(58527|58581|58587|58599)$
ip as-path access-list AS10075-in permit ^10075(_[0-9]+)*_(58615|58616|58623|58629)$
ip as-path access-list AS10075-in permit ^10075(_[0-9]+)*_(58655|58656|58657|58662)$
ip as-path access-list AS10075-in permit ^10075(_[0-9]+)*_(58665|58668|58673|58682)$
ip as-path access-list AS10075-in permit ^10075(_[0-9]+)*_(58684|58688|58689|58691)$
ip as-path access-list AS10075-in permit ^10075(_[0-9]+)*_(58695|58701|58704|58705)$
ip as-path access-list AS10075-in permit ^10075(_[0-9]+)*_(58715|58736|58737|58749)$
```

BGPQ4 - Extended Access List ~ (Cisco, Juniper)

```
┌──(naveen LAPTOP-6VNOIOAD)-[~]
└─$ bgpq4 -E AS10075 -S APNIC -l AS10075-v4-in
no ip access-list extended AS10075-v4-in
ip access-list extended AS10075-v4-in
permit ip host 103.7.248.0 host 255.255.252.0
permit ip host 103.7.250.0 host 255.255.255.0
permit ip host 103.7.251.0 host 255.255.255.0
permit ip host 103.131.156.0 host 255.255.252.0
permit ip host 103.131.156.0 host 255.255.255.0
permit ip host 103.131.157.0 host 255.255.255.0
permit ip host 103.131.158.0 host 255.255.255.0
permit ip host 103.131.159.0 host 255.255.255.0
permit ip host 103.229.82.0 host 255.255.254.0
permit ip host 163.47.156.0 host 255.255.252.0
permit ip host 163.47.156.0 host 255.255.254.0
permit ip host 163.47.156.0 host 255.255.255.0
permit ip host 163.47.157.0 host 255.255.255.0
permit ip host 163.47.158.0 host 255.255.255.0
permit ip host 163.47.159.0 host 255.255.255.0
```

```
naveen@LAPTOP-6VNOIOAD:~$
naveen@LAPTOP-6VNOIOAD:~$ bgpq3 -E AS10075 -S APNIC -l AS10075-v4-in -J
policy-options {
 policy-statement AS10075-v4-in {
replace:
  from {
    route-filter 103.7.248.0/22 exact;
    route-filter 103.7.250.0/24 exact;
    route-filter 103.7.251.0/24 exact;
    route-filter 103.131.156.0/22 exact;
    route-filter 103.131.156.0/24 exact;
    route-filter 103.131.157.0/24 exact;
    route-filter 103.131.158.0/24 exact;
    route-filter 103.131.159.0/24 exact;
    route-filter 103.229.82.0/23 exact;
    route-filter 163.47.156.0/22 exact;
    route-filter 163.47.156.0/23 exact;
    route-filter 163.47.158.0/24 exact;
    route-filter 163.47.159.0/24 exact;
  }
 }
}
naveen@LAPTOP-6VNOIOAD:~$
```

ASN Bogons

AS Number/Range	Status	RFC Reference
0	Reserved	RFC7607
23456	Transition_AS	RFC6793
64496 - 64511	Reserved for use in docs and code	RFC5398
64512 - 65534	Reserved for Private Use	RFC6996
65535	Reserved	RFC7300
65536 - 65551	Reserved for use in docs and code	RFC5398
65552 - 131071	Reserved	By IANA
4200000000 - 4294967294	Reserved for Private Use	RFC6996
4294967295	Reserved	RFC7300

Filtering AS Bogons using AS-Path Access Lists

```
ip as-path access-list 99 permit _0_
   ip as-path access-list 99 permit _23456_
   ip as-path access-list 99 permit _(6449[6-9])_|_(6450[0-9])_|_(6451[0-1])_|_(6553[6-9])_|_(6554[0-9])_|_(6555[0-1])_
   ip as-path access-list 99 permit _6(4(5(1[2-9]|[2-9][0-9])|[6-9][0-9][0-9])|5([0-4][0-9][0-9]|5([0-2][0-9]|3[0-5])))_
   ip as-path access-list 99 permit _6555[2-9]_|_655[6-9][0-9]_|_65[6-9][0-9][0-9]_|_6[6-9][0-9][0-9][0-9]_
   ip as-path access-list 99 permit _[7-9][0-9][0-9][0-9][0-9]_|_1[0-2][0-9][0-9][0-9][0-9]_|_130[0-9][0-9][0-9]_
   ip as-path access-list 99 permit _1310[0-6][0-9]_|_13107[0-1]_
   ip as-path access-list 99 permit _42[0-8][0-9][0-9][0-9][0-9][0-9][0-9][0-9]_
   ip as-path access-list 99 permit _(429[0-3][0-9][0-9][0-9][0-9][0-9][0-9])_|_(4294[0-8][0-9][0-9][0-9][0-9][0-9])_
   ip as-path access-list 99 permit _(42949[0-5][0-9][0-9][0-9][0-9])_|_(429496[0-6][0-9][0-9][0-9])_
   ip as-path access-list 99 permit _(4294967[0-1][0-9][0-9])_|_(42949672[0-8][0-9])_|_(429496729[0-4])_
   !
   route-map PEER-IN deny 1
     match as-path 99
```

Source: http://as2914.net/bogon_asns/configuration_examples.txt

Action 2: Anti-Spoofing

(BCP 38 – RFC2827 and more) Network Ingress Filtering

Source Address Validation (SAV)

Source Address Validation (SAV) is the best current practice (BCP 38/RFC 2827) aimed at filtering packets based on source IP addresses at the network edges.

Check the source IP address of IP packets

- filter invalid source address
- filter close to the packets origin as possible
- filter precisely as possible

If no networks allow IP spoofing, we can eliminate these kinds of Spoofing/DoS attacks

Source Address Validation (SAV)

ACL (Access Control Lists)

- Create an access-list that lists all customer IP blocks and use ingress filtering to filter packets that are sourced from spoofed IP address

uRPF (Unicast Reverse Path Forwarding)

- uRPF is a technique where the router can discard packets with invalid or fake or incorrect source addresses by a simple check against the Forwarding table (FIB).
- Designed to help mitigate attacks based on source address spoofing.
- Interface configured for uRPF drops packets if they are from spoofed IP source address
- Check incoming packets using ‘routing table’ against FIB/CEF table
- Uses less CPU and RAM, compared to implementing access-lists

Unicast Reverse Path Forwarding (uRPF)

Loose Mode

Check that an entry exists in the routing table

Strict Mode

Check that an entry exists in the routing table

and the route points to the receiving interface

BCP38: Network ingress filtering https://tools.ietf.org/html/bcp38

BCP84: Ingress filtering for multi-homed networks https://tools.ietf.org/html/bcp84

uRPF: Strict Mode

Router compares source address of incoming packet with FIB entry

- If FIB entry interface matches incoming interface, the packet is forwarded
- If FIB entry interface does not match incoming interface, the packet is dropped

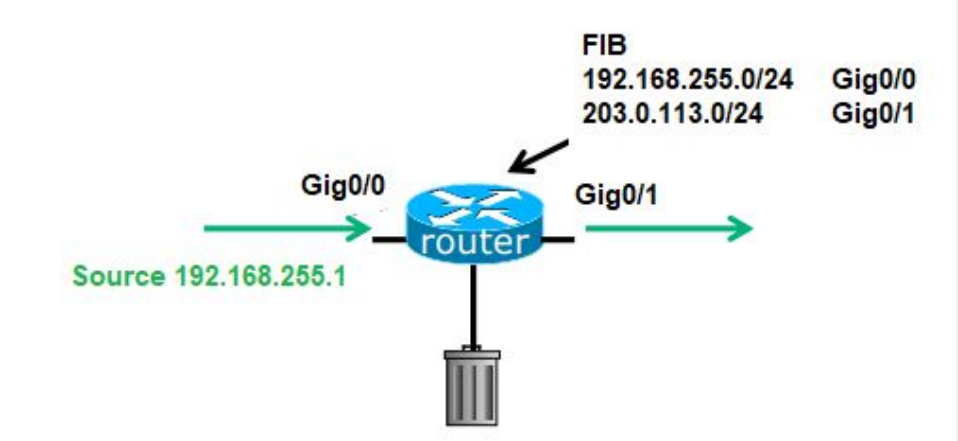

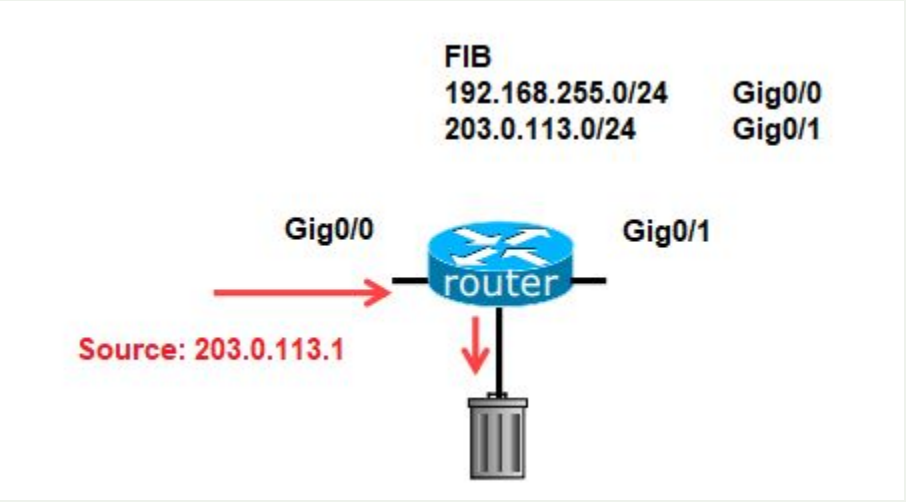

Image credit: NSRC uRPF

uRPF: Strict Mode

Configuration in operator's router

Cisco

```
int gig 0/0
ip verify unicast source reachable-via rx
ipv6 verify unicast source reachable-via rx
```

Juniper

```
family inet {
      rpf-check;
  }
family inet6 {
      rpf-check;
  }
```

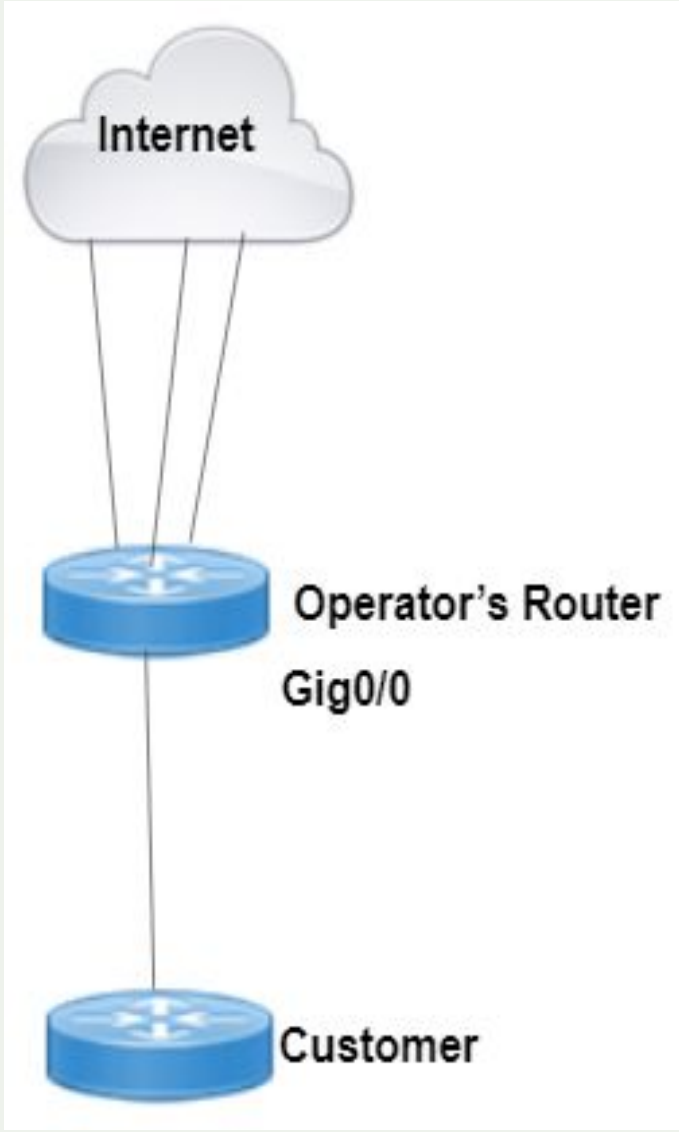

- Implement uRPF on all single-homed customer facing interfaces

Verification (uRPF)

```
as64500#show ip interface GigabitEthernet 0/0 | begin uRPF
  Input features: uRPF, MCI Check
  IPv4 WCCP Redirect outbound is disabled
  IPv4 WCCP Redirect inbound is disabled
  IPv4 WCCP Redirect exclude is disabled
  IP verify source reachable-via RX
    5166 verification drops
    0 suppressed verification drops
    12 verification drop-rate
```

```
as64500#show ipv6 interface GigabitEthernet 0/0 | begin Unicast
  Input features: Verify Unicast Reverse-Path
 IPv6 verify source reachable-via rx
    0 verification drop(s) (process), 19371 (CEF)
    0 suppressed verification drop(s) (process), 0 (CEF)
  !
```

uRPF: Loose Mode

- Source address must be in the FIB
- Typically used to drop non-routed address space
- Can be used when asymmetric traffic flows are present (for multihoming)

Cisco

```
interface gig 0/0
ip verify unicast source reachable-via any
ipv6 verify unicast source reachable-via any
```

ACL - SAV example (Cisco)

ACLs can also be used on devices where automatic filtering features are not available you can use ACLs to manually implement equivalent filtering.

- Towards a provider's servers
- Towards infrastructure networks
- Deployed on the PE/CE boundary

Configuration in operator's router

```
ip access-list extended fromCUSTOMER1
 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 fromCUSTOMER1 in
!
```

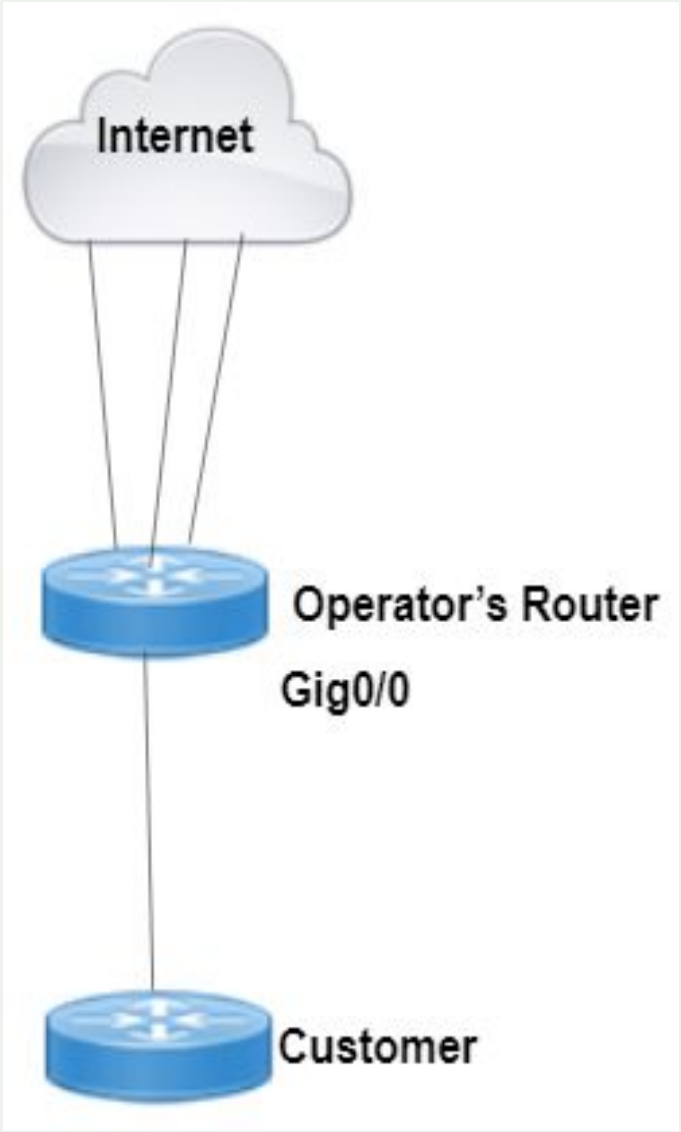

ACL example (Juniper)

Configuration in operator's router

```
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;
}
```

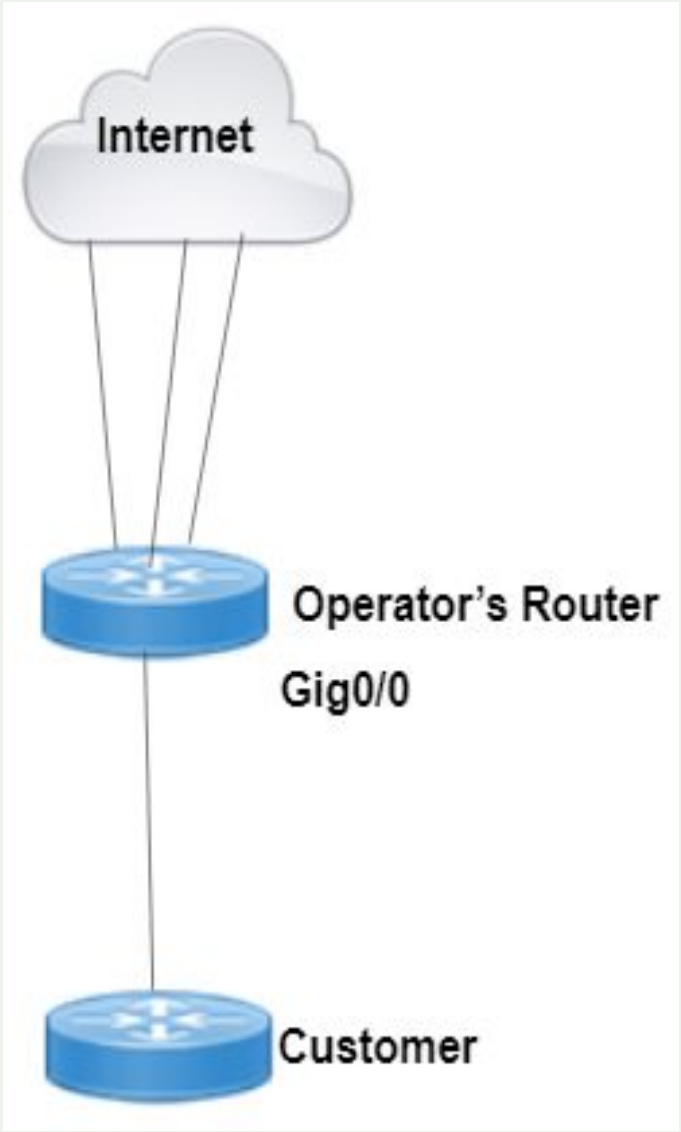

Action 3: Coordination

Facilitating global operational communication and coordination between network operators

Coordination

MANRS participants should perform the following action in order to facilitate global communication between network operators.

- Maintain globally accessible and up-to-date contact information.
- Successful prevention and mitigation of routing incidents strongly depends on effective operational communication and coordination between network operators globally.
- In order to achieve this, it is essential that you maintain your up-to-date contact information in databases that are publicly accessible.
- Network operators are advised to maintain their contact information in several databases (ex. an IRR)

Maintaining Contact Information

MANRS participants should publish and maintain their contact information to their region's RIR Whois Database, IRR Database, PeeringDB and their company website.

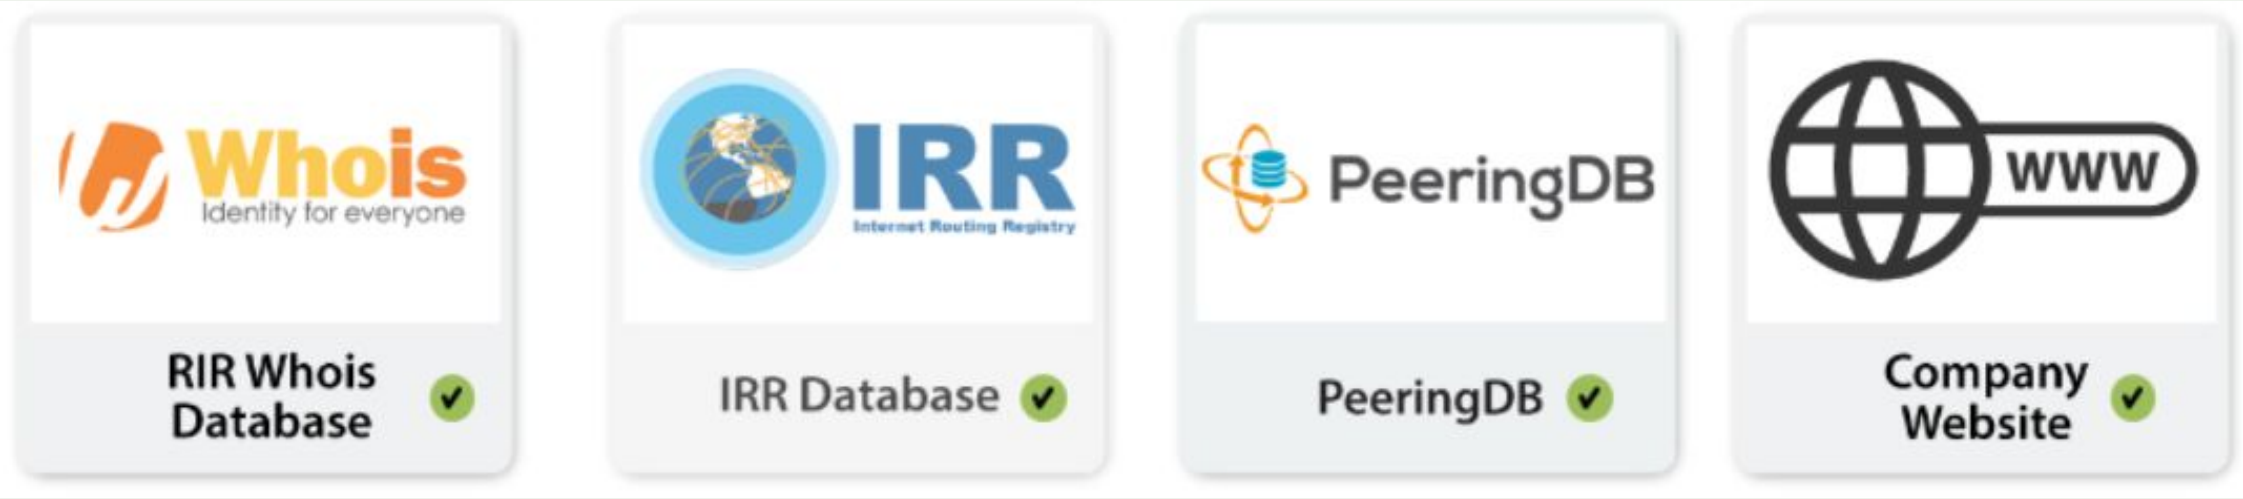

Coordination

PeeringDB

Search here for a network, IX, or facility.

Advanced Search

Fiber@Home Global

Contact Information

Role	Name	Phone E-Mail
NOC	NOC	+8801841158587 iig@fiberathome.net
Policy	SUMON AHMED SABIR	+8801711527065 sumon@fiberathome.net
Technical	CHINMAY BISWAS	+8801716463150 chinmay.biswas@fiberathome.net
Technical	ANIRBAN DATTA	+8801847102419 anirban@fiberathome.net

Organization	Fiber@Home Global Limited
Also Known As	PICO
Company Website	http://www.fiberathome.net/
ASN	10075
IRR as-set/route-set	AS-FGL
Route Server URL	
Looking Glass URL	
Network Type	NSP
IPv4 Prefixes	4500
IPv6 Prefixes	2000
Traffic Levels	200-300Gbps
Traffic Ratios	Mostly Inbound
Geographic Scope	Asia Pacific
Protocols Supported	Unicast IPv4 Multicast IPv6 Never via route servers
Last Updated	2020-12-01T18:20:15Z
Notes	

Coordination

Maintaining Contact Information in RIRs: AFRINIC, APNIC, RIPE, LACNIC, ARIN

whois -h whois.apnic.net AS10075

naveen@LAPTOP-6VNOIOAD: ~

```
aut-num:        AS10075
as-name:        FGL-AS-BD
descr:          Fiber@Home Global Limited
country:        BD
org:            ORG-FGL3-AP
admin-c:        FGLA2-AP
tech-c:         FGLA2-AP
abuse-c:        AF576-AP
mnt-lower:      MAINT-FGL-BD
mnt-routes:     MAINT-FGL-BD
mnt-by:         APNIC-HM
mnt-irt:        IRT-FGL-BD
last-modified:  2020-10-06T14:13:34Z
source:         APNIC

irt:            IRT-FGL-BD
address:        House # 8/B, Road1, Gulshan-1, Dhaka Dhaka 1212
e-mail:         iig@fiberathome.net
abuse-mailbox:  iig@fiberathome.net
admin-c:        FGLA2-AP
tech-c:         FGLA2-AP
auth:           # Filtered
remarks:        iig@fiberathome.net was validated on 2021-04-13
mnt-by:         MAINT-FGL-BD
last-modified:  2021-04-13T13:16:00Z
source:         APNIC
```

naveen@LAPTOP-6VNOIOAD: ~

```
organisation:   ORG-FGL3-AP
org-name:       Fiber@Home Global Limited
country:        BD
address:        House # 8/B, Road1, Gulshan-1
phone:          +8801817022207
fax-no:         +88028815010
e-mail:         iig@fiberathome.net
mnt-ref:        APNIC-HM
mnt-by:         APNIC-HM
last-modified:  2018-09-17T12:57:28Z
source:         APNIC

role:           ABUSE FGLBD
address:        House # 8/B, Road1, Gulshan-1, Dhaka Dhaka 1212
country:        ZZ
phone:          +000000000
e-mail:         iig@fiberathome.net
admin-c:        FGLA2-AP
tech-c:         FGLA2-AP
nic-hdl:        AF576-AP
remarks:        Generated from irt object IRT-FGL-BD
abuse-mailbox:  iig@fiberathome.net
mnt-by:         APNIC-ABUSE
last-modified:  2020-10-06T14:13:34Z
source:         APNIC

role:           FiberHome Global Limited administrator
address:        House#8/B, Road#1, Gulshan-1, Dhaka Dhaka 1212
country:        BD
phone:          +8801817022207
fax-no:         +8801817022207
e-mail:         iig@fiberathome.net
admin-c:        FGLA2-AP
tech-c:         FGLA2-AP
nic-hdl:        FGLA2-AP
mnt-by:         MAINT-FGL-BD
```

Action 4: Global Validation (IRR/RPKI)

Facilitating validation of routing information on a global scale

Global Validation

There are 2 ways to provide the validation information (IRR and/or RPKI)

Providing information through the IRR system

Internet Routing Registries (IRRs) contain information—submitted and maintained by ISPs or other entities—about Autonomous System Numbers (ASNs) and routing prefixes. IRRs can be used by ISPs to develop routing plans.

The global IRR is comprised of a network of distributed databases maintained by Regional Internet Registries (RIRs) such as APNIC, service providers (such as NTT), and third parties (such as RADB).

Global Validation

Routing information should be made available on a global scale to facilitate validation, which includes routing policy, ASNs and prefixes that are intended to be advertised to third parties. Since the extent of the internet is global, information should be made public and published in a well known place using a common format.

Object	Source	Description
aut-num	IRR	Policy documentation
route/route6	IRR	NLRI/origin
as-set	IRR	Customer cone
ROA	RPKI	NLRI/origin

Global Validation

```
$ whois -h whois.apnic.net 1.1.1.0/24

route:          1.1.1.0/24
origin:         AS13335
descr:          APNIC Research and Development, 6 Cordelia St
mnt-by:         MAINT-AU-APNIC-GM85-AP
last-modified:  2018-03-16T16:58:06Z
source:         APNIC
```

Global Validation

```
$ whois -h whois.radb.net 1.1.1.0/24

route:          1.1.1.0/24
origin:         AS13335
descr:          APNIC Research and Development, 6 Cordelia St
mnt-by:         MAINT-AU-APNIC-GM85-AP
last-modified:  2018-03-16T16:58:06Z
source:         APNIC

route:          1.1.1.0/24
descr:          Cloudflare, Inc.
descr:          101 Townsend Street, San Francisco, California 94107, US
origin:         AS13335
mnt-by:         MNT-CLOUD14
notify:         rir@cloudflare.com
```

Global Validation

Internet Routing Registry (IRR)

- Network operators can document which AS is originating their IPv4/IPv6 prefixes
- Used by operators to filter prefixes received from their customers and peers
- Third party databases need to be used (RADB, Operators/NTT)
  - RADB comes with a recurring yearly subscription costs
  - For RADB, a commercial relationship with merit is required. (lacks accuracy of data)
  - For RADB any paid member can update/delete information for their resources (lots of junk data)
  - For NTTCOM, a customer relationship with them is required.

Resource Public Key Infrastructure (RPKI)

Global Validation

Providing information through the RPKI system

- Store information about prefixes originated by your network in the form of Route Origin Authorization (ROA) objects.
- Only prefixes that belong to our ASN is covered.
- Only the origin ASN is verified, not the full path.
- All Regional Internet Registries (RIR) offers a hosted Resource Certification service.

Resource Public Key Infrastructure (RPKI)

- A security framework for verifying the association between resource holders and their Internet resources
- RPKI is a way to define data in an out-of-band system such that the information that are exchanged by BGP can be validated to be correct.
- RPKI is used to make Internet routing more secure.

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

- AS Numbers
- IP Addresses

ROA (Route Origin Authorization)

Legitimate holder of a block of IP addresses can use their resource certificate to make an authoritative, signed statement about which BGP AS is authorized to originate their prefix in BGP

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

Prefix	103.229.0.0/23
Max-Length	/24
Origin ASN	AS10075

What can RPKI do?

Authoritatively proof:

- Who is the legitimate owner of an address, and
- Identify which ASNs have the permission from the holder to originate the address

Prevents route hijacking

- A prefix originated by an AS without authorization
- Reason: Malicious intent

Prevents mis-orgination

- A prefix that is mistakenly originated by an AS which does not own it, also route leaks
- Reason: configuration mistake / fat finger

RPKI Validation States

Valid

- the prefix (prefix length) and AS pair found in the database.

Invalid

- prefix is found, but origin AS is wrong
- the prefix length is longer than the maximum length

Not Found

- No valid ROA found
- Neither valid nor invalid (perhaps not created)

RPKI Components

Issuing Party – Internet Registries (*IRs)

- Certificate Authority (CA) that issues resource certificates to end-holders
- Publishes the objects (ROAs) signed by the resource certificate holders

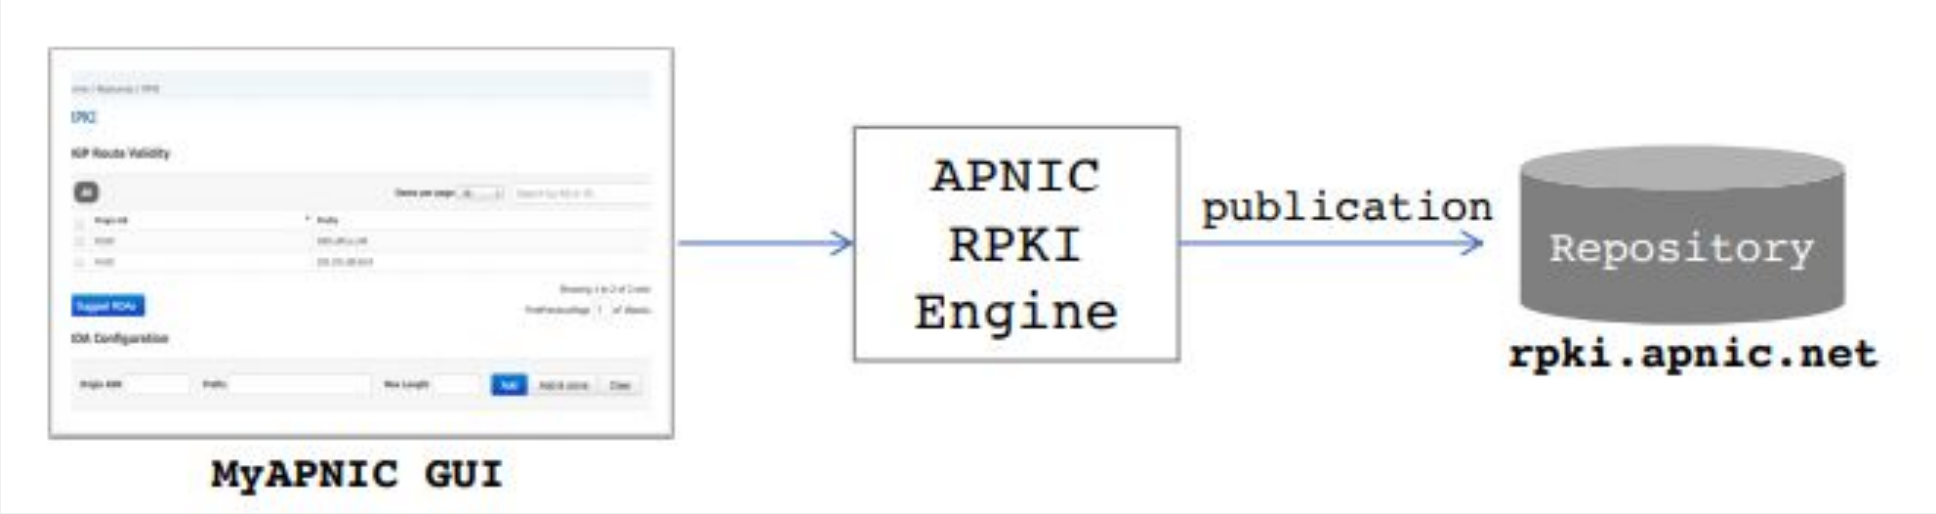

Source: APNIC

RPKI-RTR

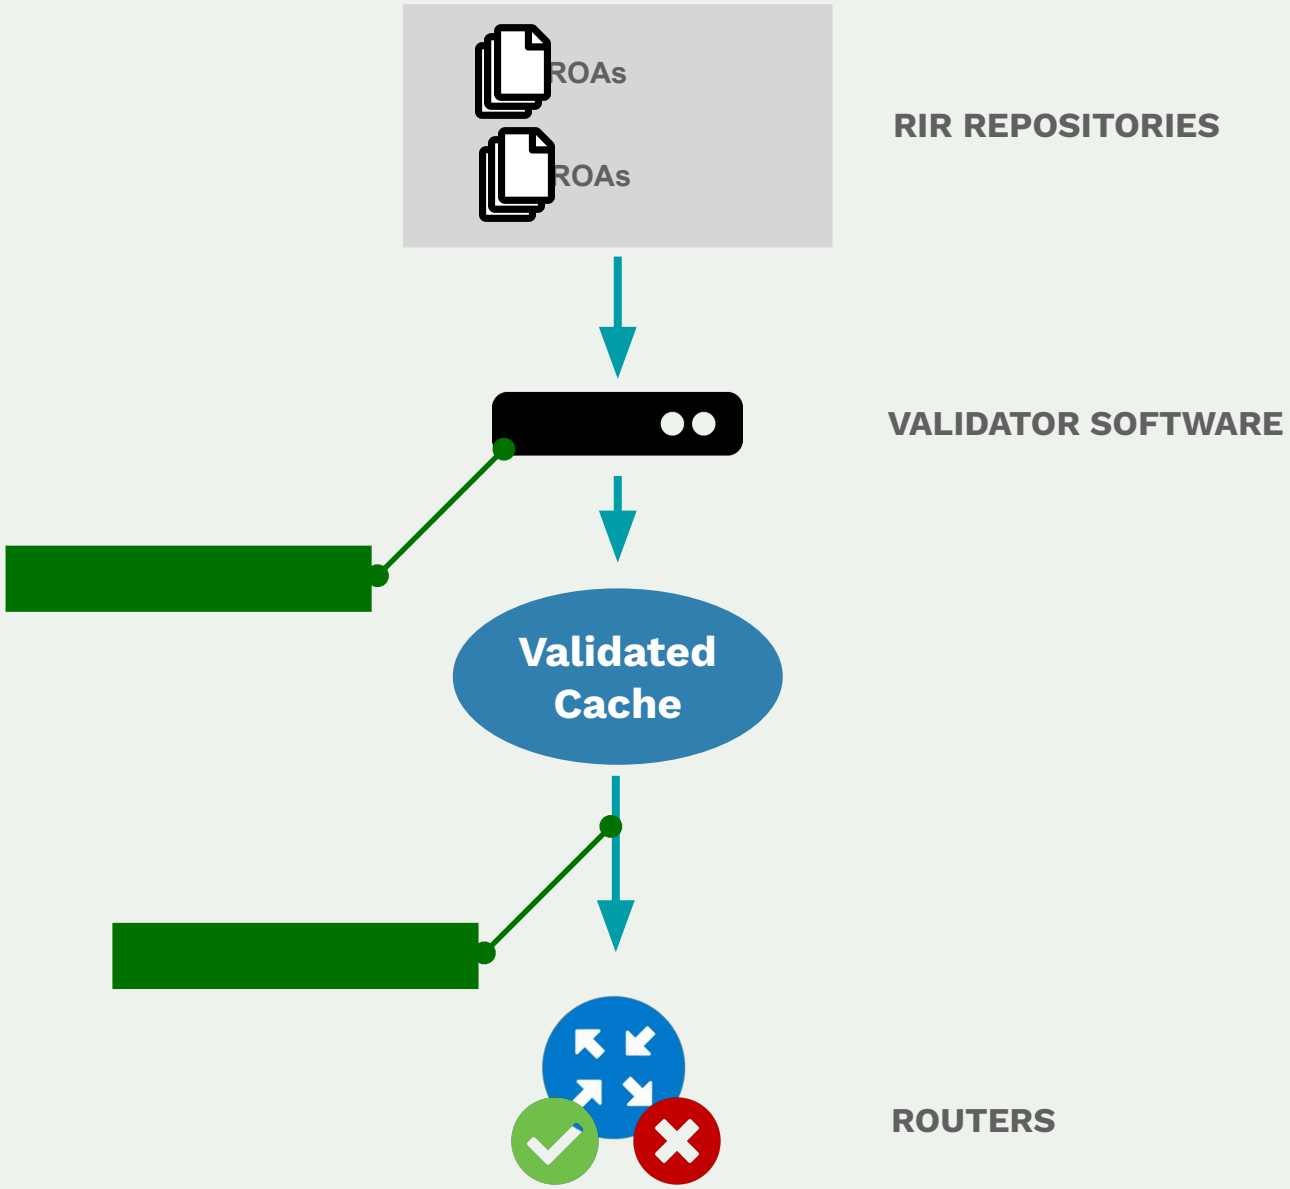
ROAs
ROAs
RIR REPOSITORIES
VALIDATOR SOFTWARE
Validated Cache
ROUTERS

What is in a ROA ?

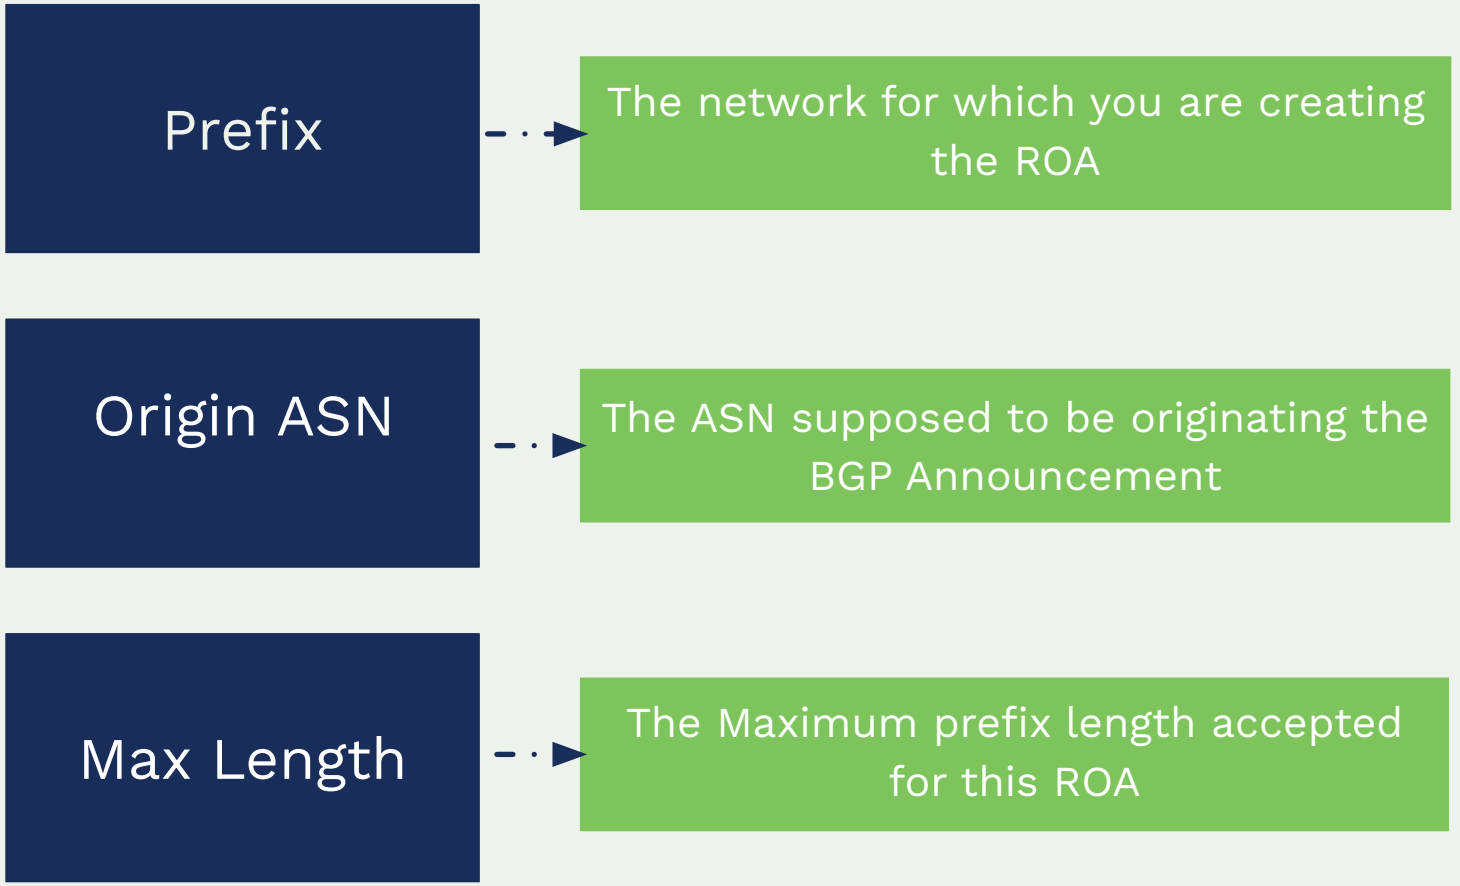

Prefix Validation Status

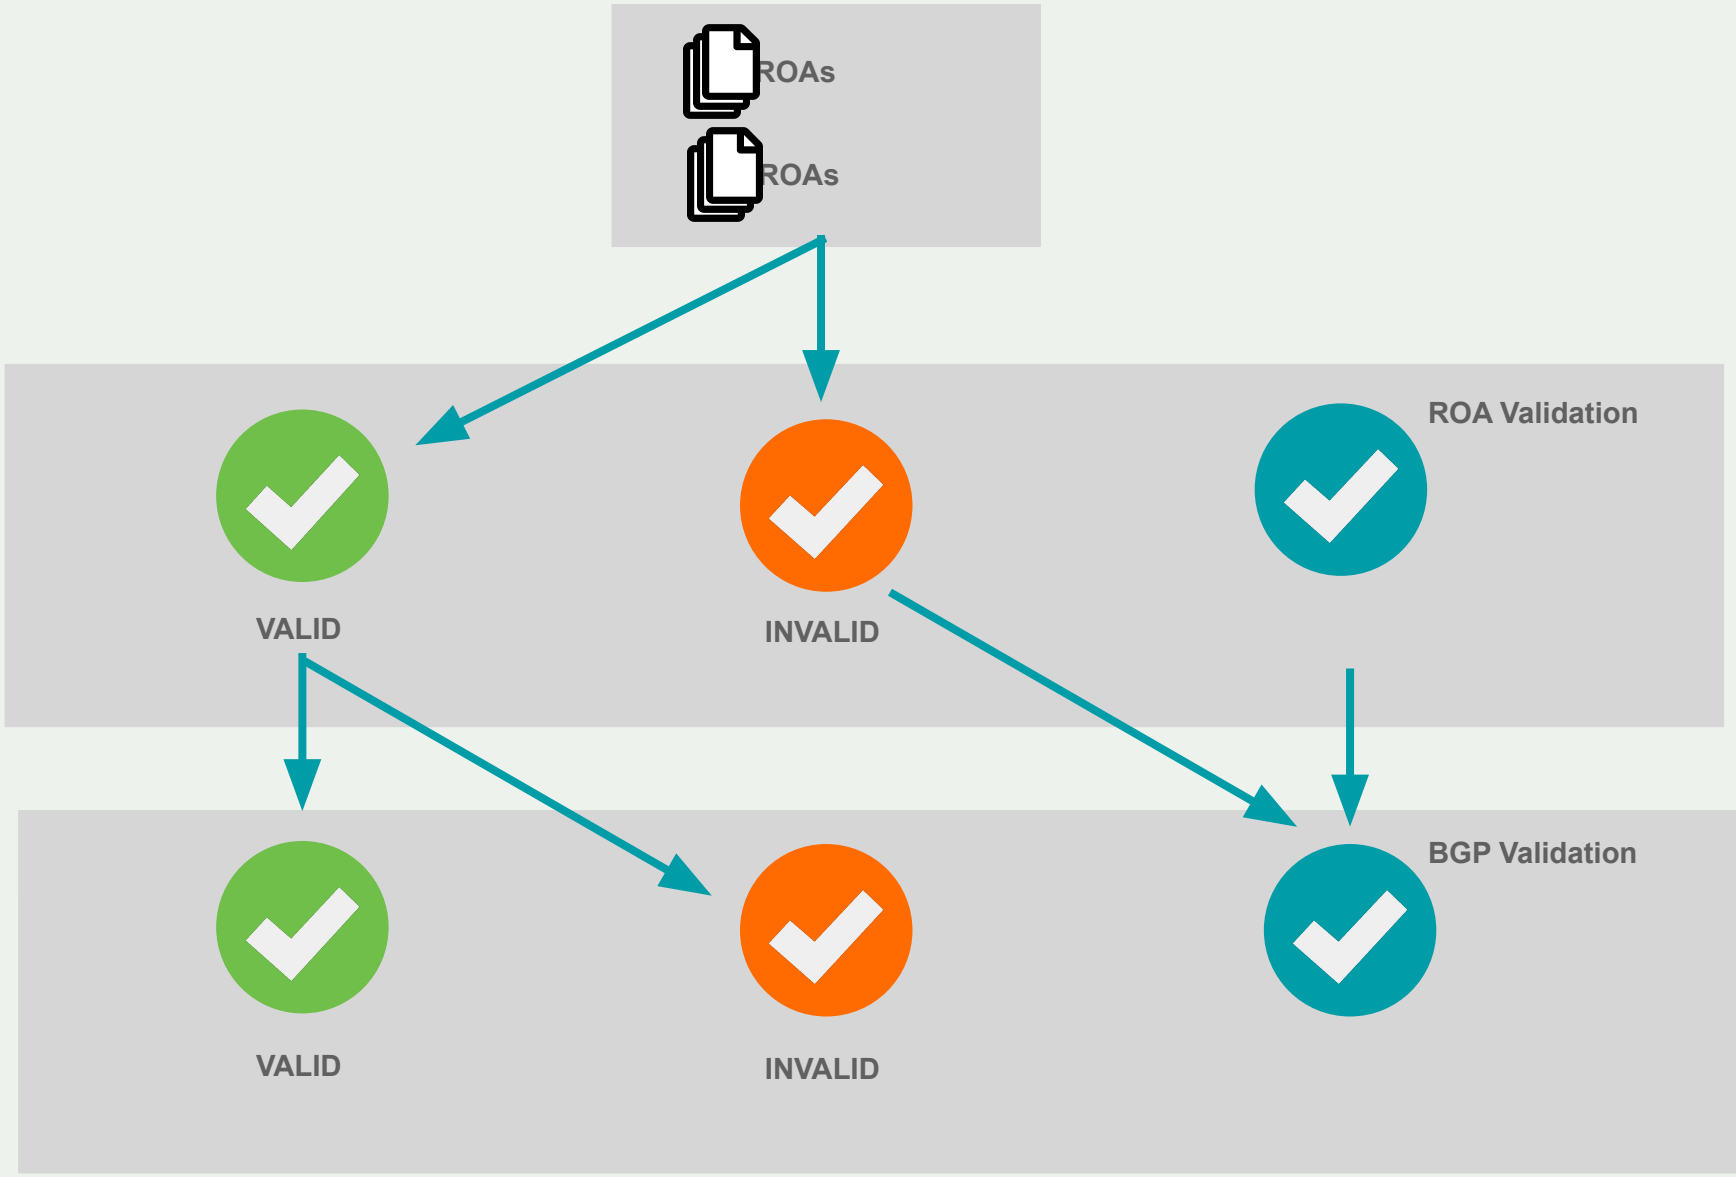

Validator Software

- NLNetLabs Routinator - https://github.com/NLnetLabs/routinator/
- FORT Validator - https://github.com/NICMx/FORT-validator/
- Cloudflare OctoRPKI - https://github.com/cloudflare/cfrpki
- RPKI-Client - https://rpki-client.org/
- Prover - https://github.com/lolepezy/rpki-prover
- Rpstir2 - https://github.com/bgpsecurity/rpstir2

How to Install Relying Party Software

MANRS - How to Videos

OctoRPKI
https://youtu.be/3Lx5wL7oG0c

Routinator
https://youtu.be/0dpmejgkTcs

Fort Validator
https://youtu.be/lCfUbJhnq3Q

Creating ROAs in MyAPNIC

Resources

Internet Resources

Summary

View all of your resource holdings.

IPv4

View your IPv4 resource holdings.

IPv6

View your IPv6 resource holdings.

AS Numbers

View your ASN resource holdings.

Reverse DNS Delegations

Add Reverse Delegations

Add new reverse delegations.

Reverse Delegation Summary

View and manage reverse delegations

Whois Updates

Whois Updates

Add, update, and delete individual Whois objects.

Bulk Whois Updates

Add, update, and delete multiple Whois objects.

Contact Details Update

Update contact details of the internet resources associated with yc

Maintainers

View your registered maintainers, and register new maintainers.

IRTs

View your registered IRT objects, and register new IRT objects.

Resource certification

RPKI

Set up your RPKI engine, and manage your Route Origin Authorizat objects.

Route management

Routes

Add, update, delete and view routes. Create Route Origin Authoris routes.

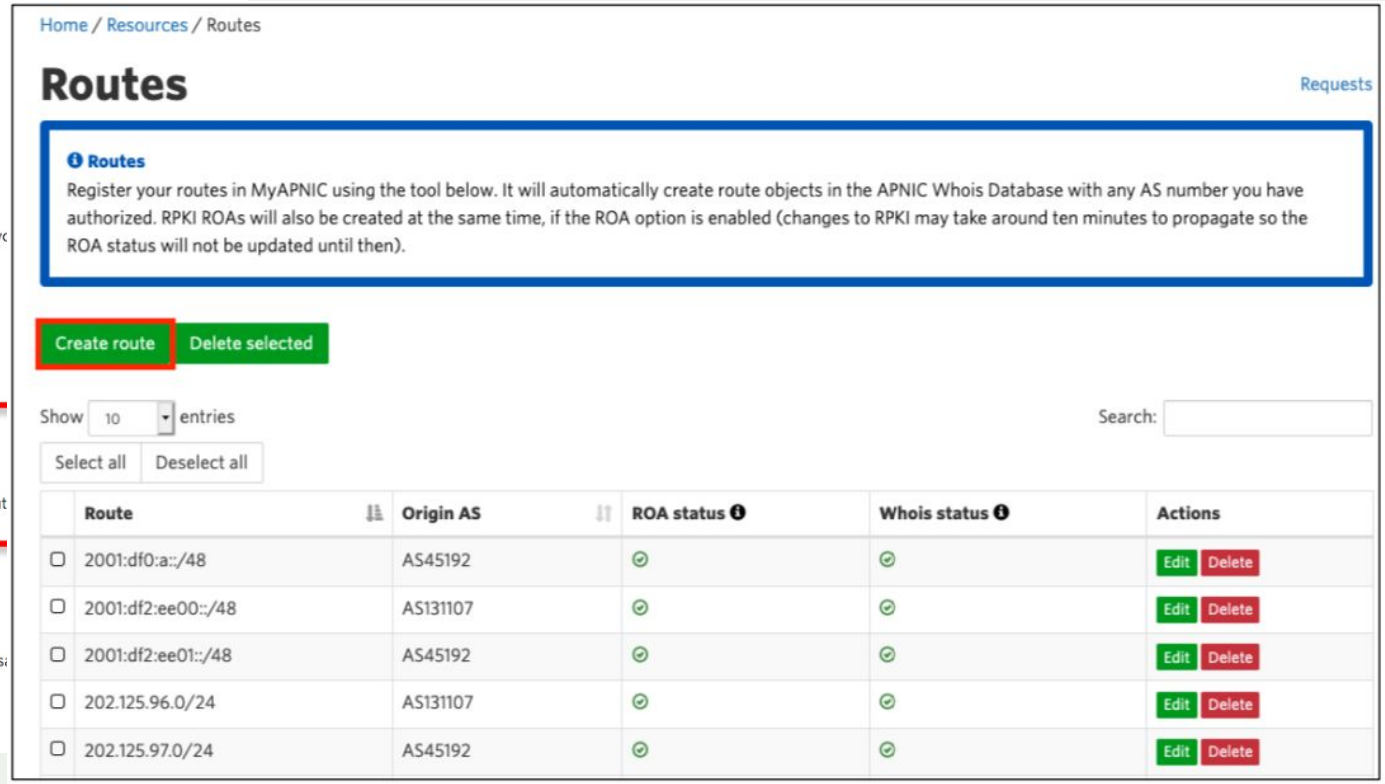

MyAPNIC Resources page.

Creating ROA (IPv6 Prefix)

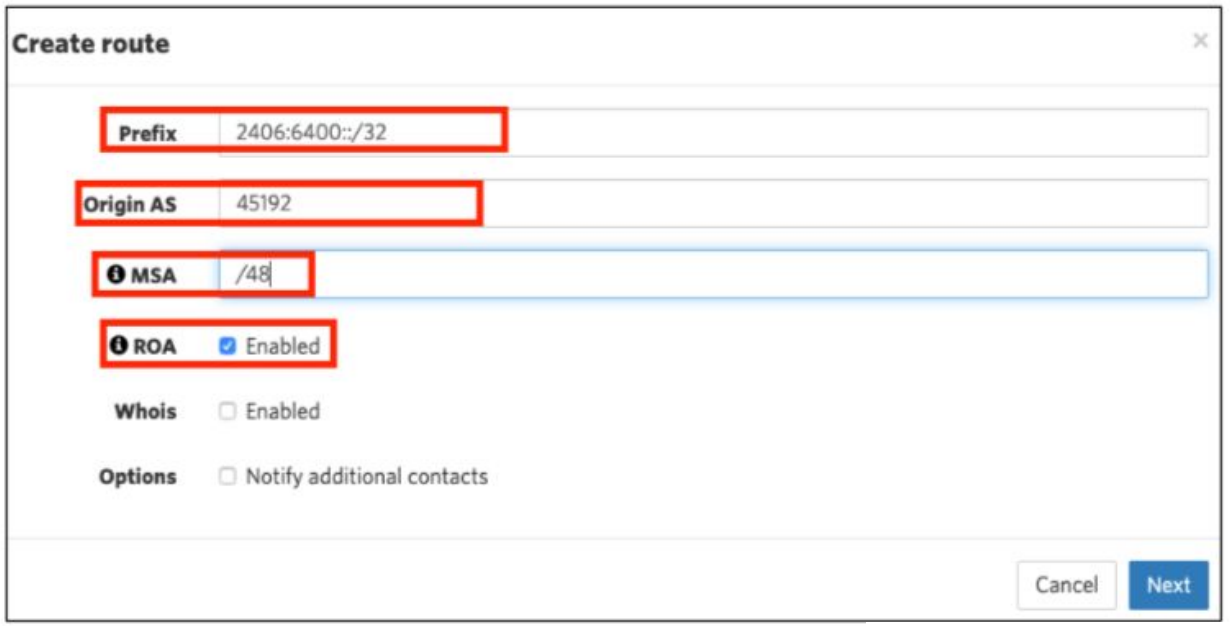

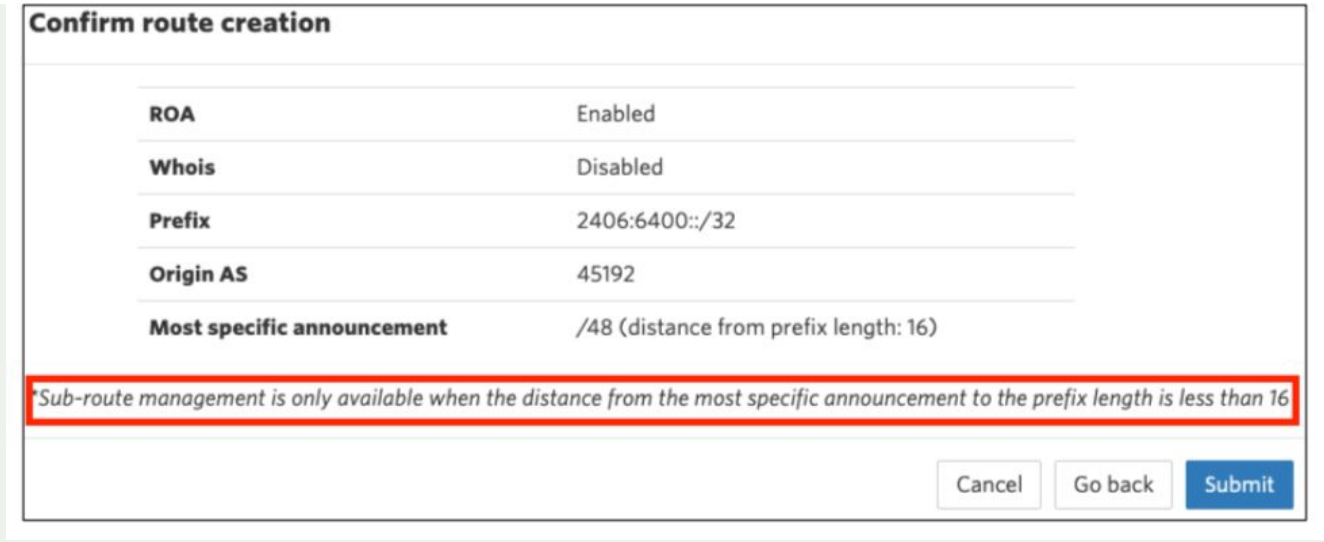

Creating ROA (IPv4 Prefix)

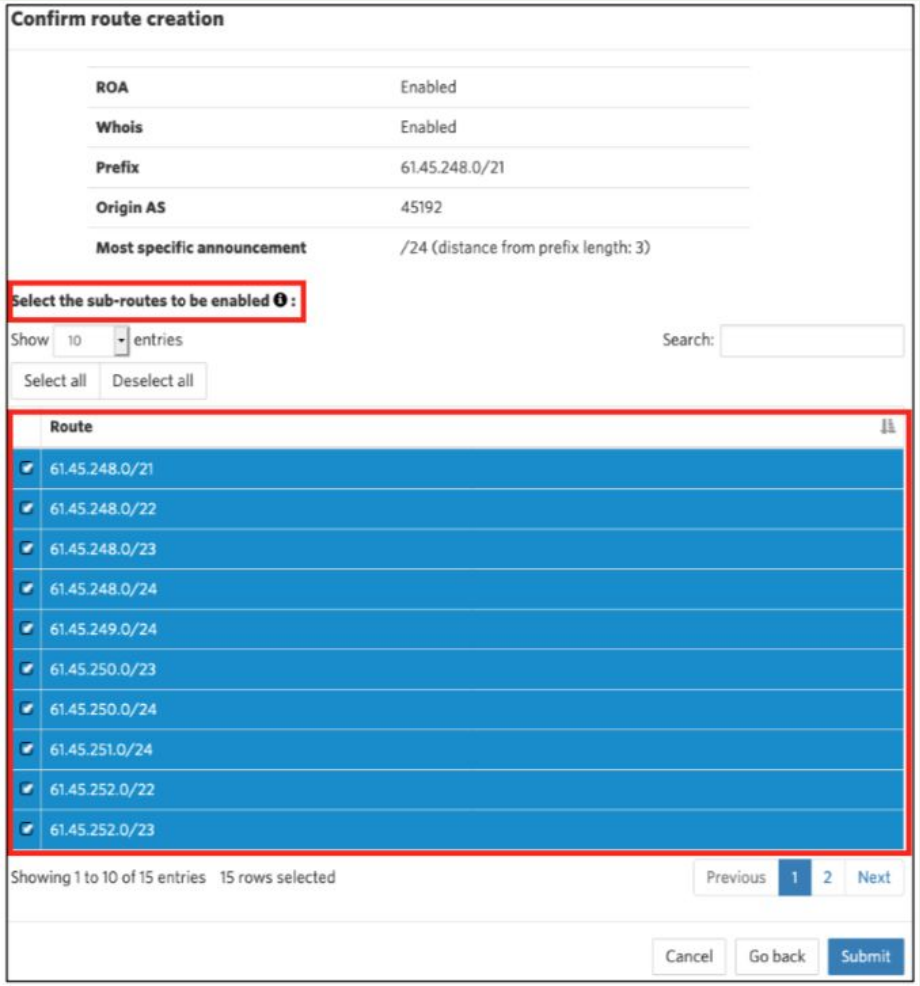

Route Origin Validation (ROV)

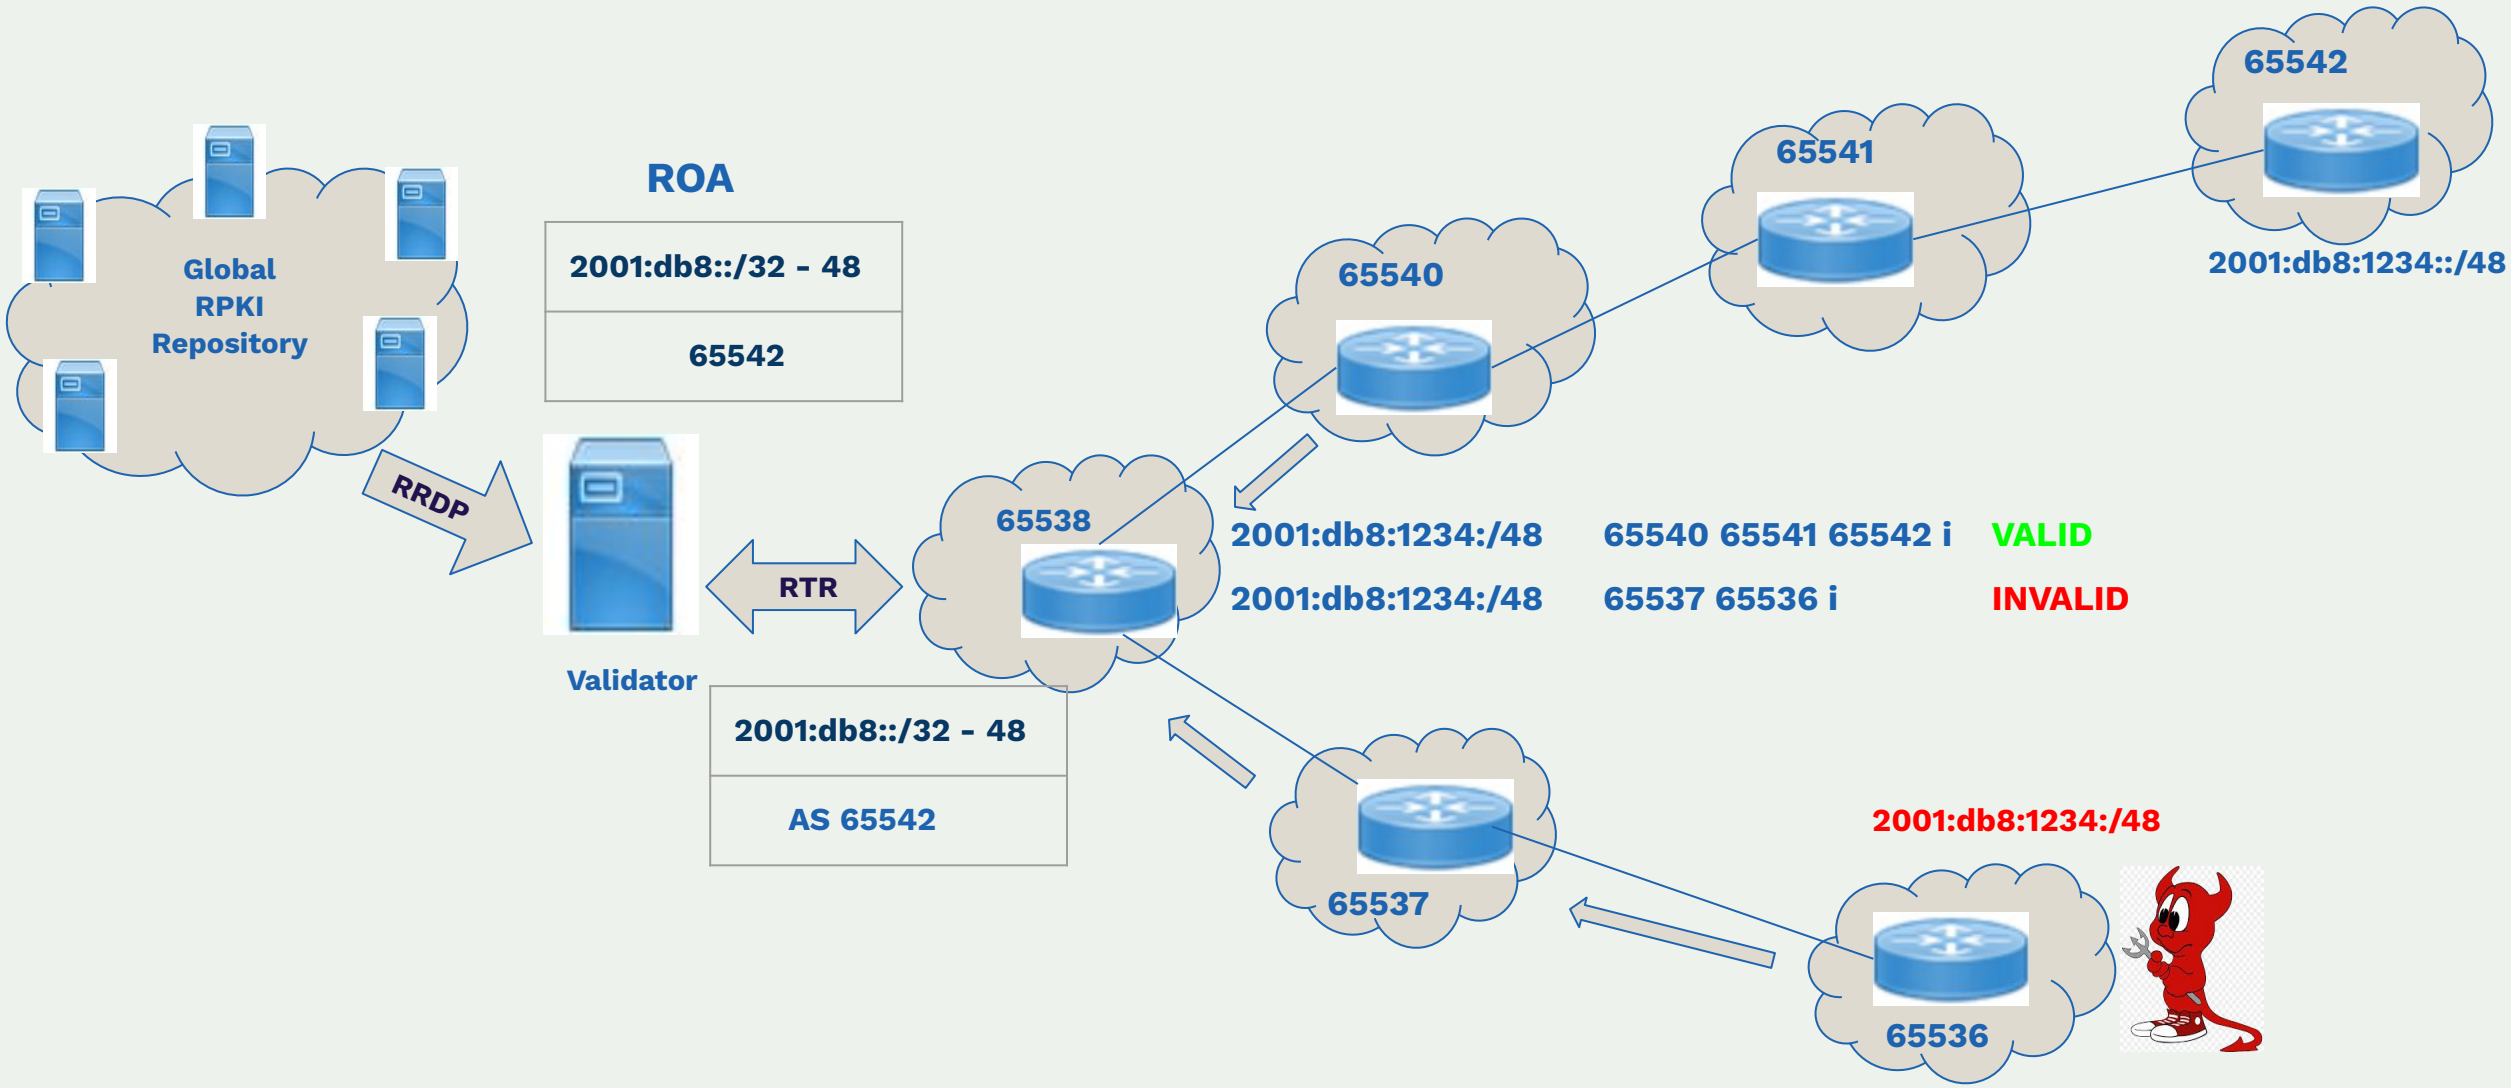

RFC 3849 : 2001:db8::/32

Image Credit: APNIC

How to Videos ~ Creating ROAs

MyAPNIC:

https://youtu.be/NLG2siznuu4

AFRINIC:

https://youtu.be/jBWCdfM0jcM

ARIN:

https://youtu.be/dueSmJwWzQ4

LACNIC:

https://youtu.be/VLcvfEJ8T4Y

RIPE-NCC:

https://youtu.be/KgUXsTKW2b4

https://www.manrs.org/resources/how-to-videos/

Route Origin Validation (ROV) – Implementations

- Cisco IOS – available from release 15.2
- Cisco IOS/XR – available from release 4.3.2
- Juniper – available from release 12.2
- Nokia – available from release R12.0R4
- Huawei – available from release V800R009C10
- FRR – available from release 4.0
- BIRD – available from release 1.6
- OpenBGPD – available from OpenBSD release 6.4
- GoBGP – available since 2018
- VyOS – available from release 1.2.0-RC11
- Mikrotik ROS – available from release v7

source: nsrc

Implement Origin Validation (ROV) ~ Cisco (IOS-XE, IOS-XR), Juniper, Arista

MANRS - How to Videos

Cisco Router ROV configuration
https://youtu.be/82O_CvW6T8c

Juniper JunOS ROV configuration
https://youtu.be/NtQ4sqLmw18

Arista ROV configuration
https://youtu.be/rXLtZcSY6gc

https://www.manrs.org/resources/how-to-videos/

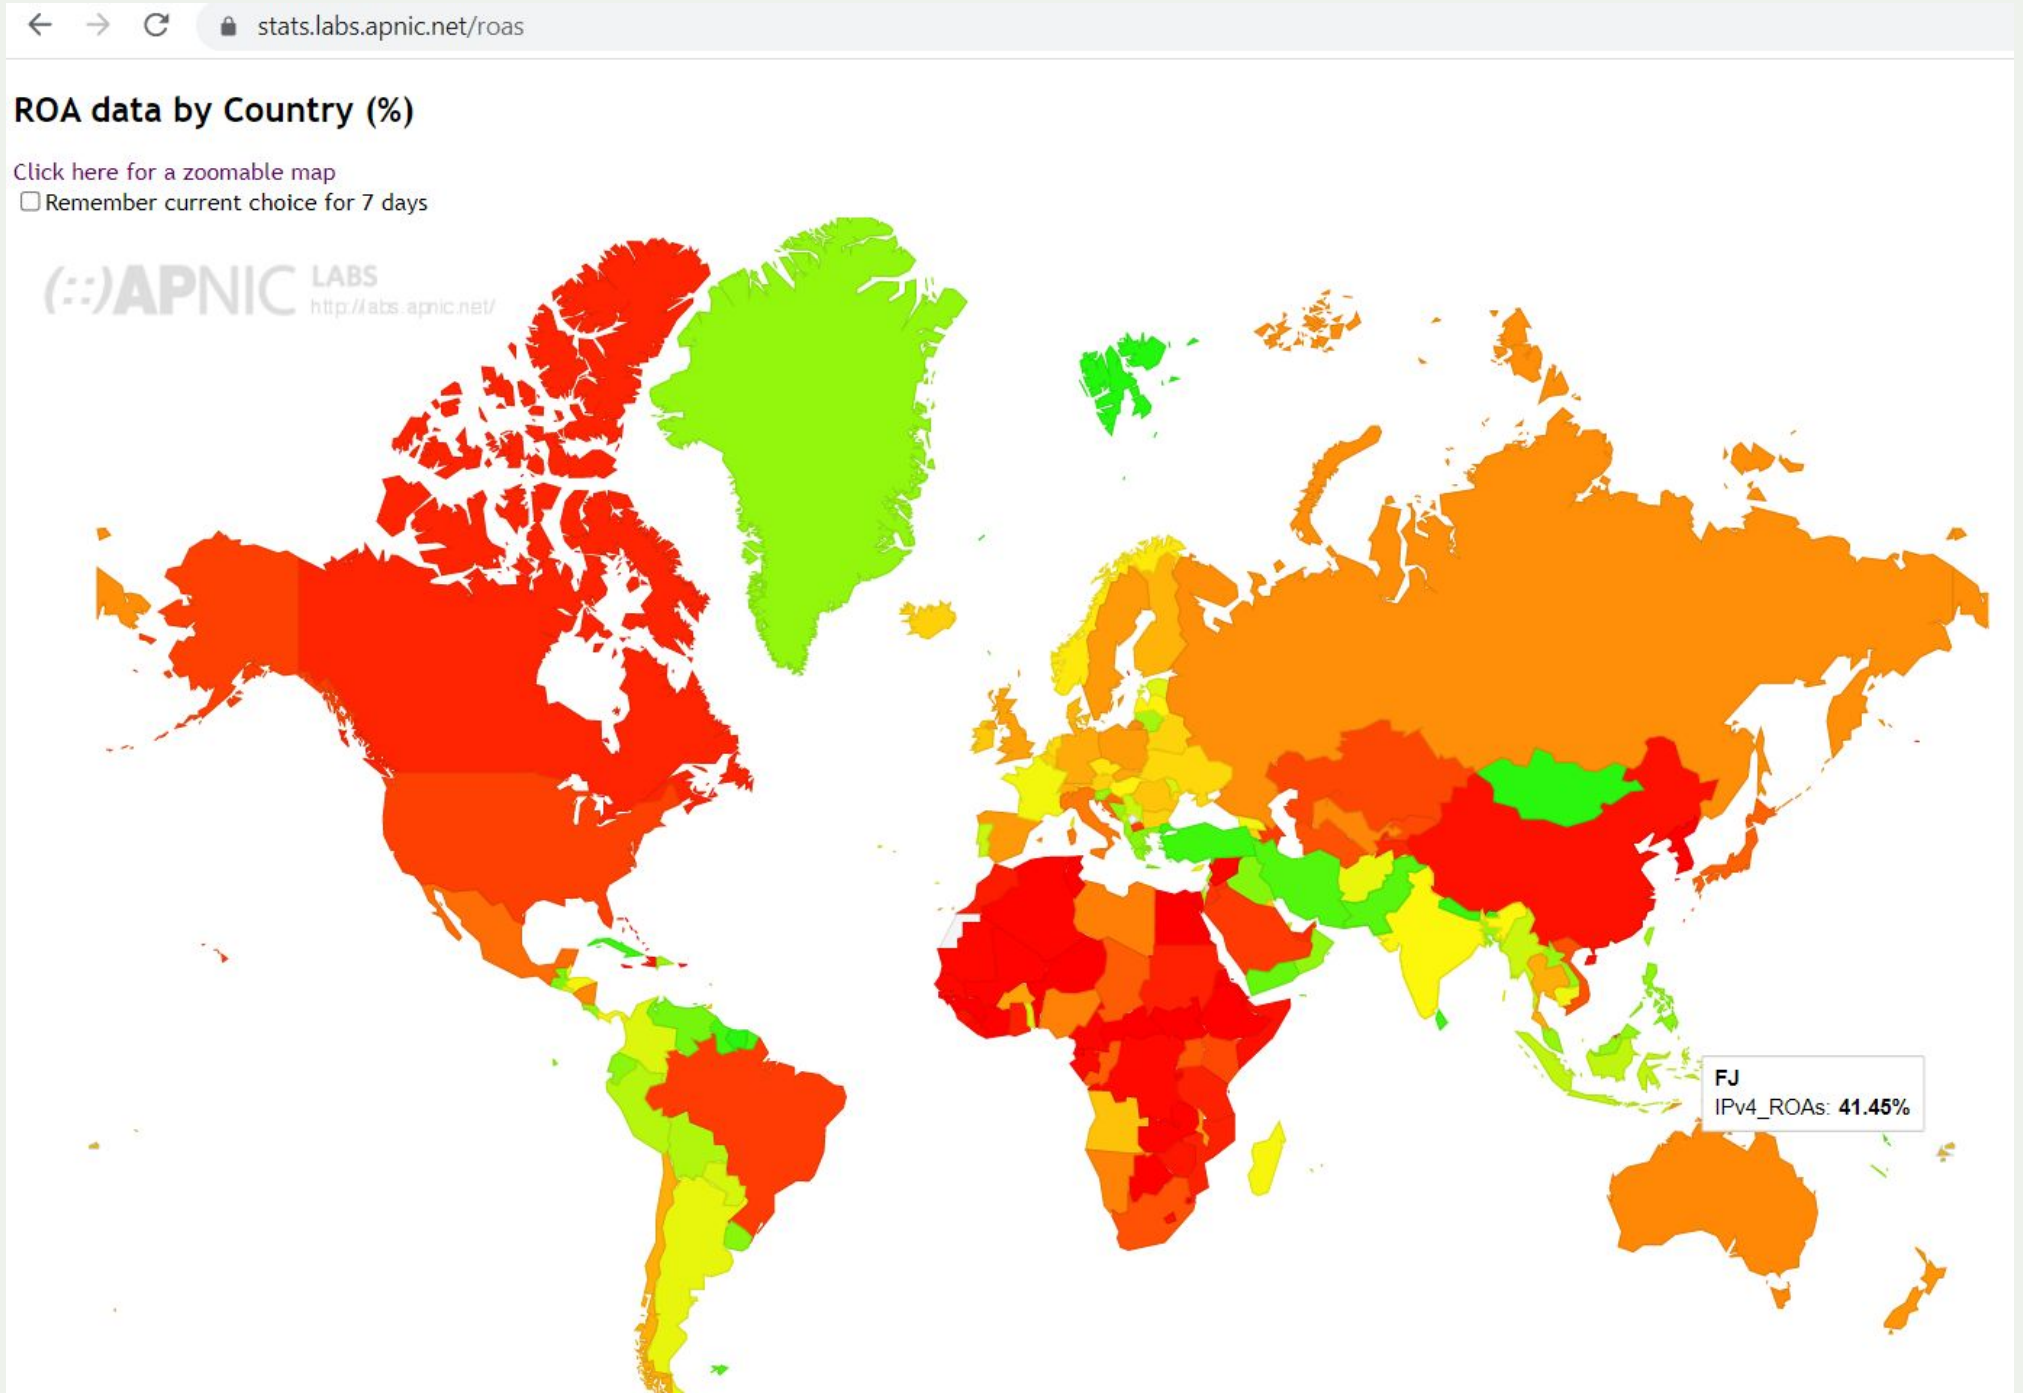
stats.labs.apnic.net/roas
ROA data by Country (%)
Click here for a zoomable map
Remember current choice for 7 days
(::)APNIC LABS
http://labs.apnic.net/
FJ
IPv4_ROAs: 41.45%

Where to find ROA? | https://roa-stats.manrs.org/

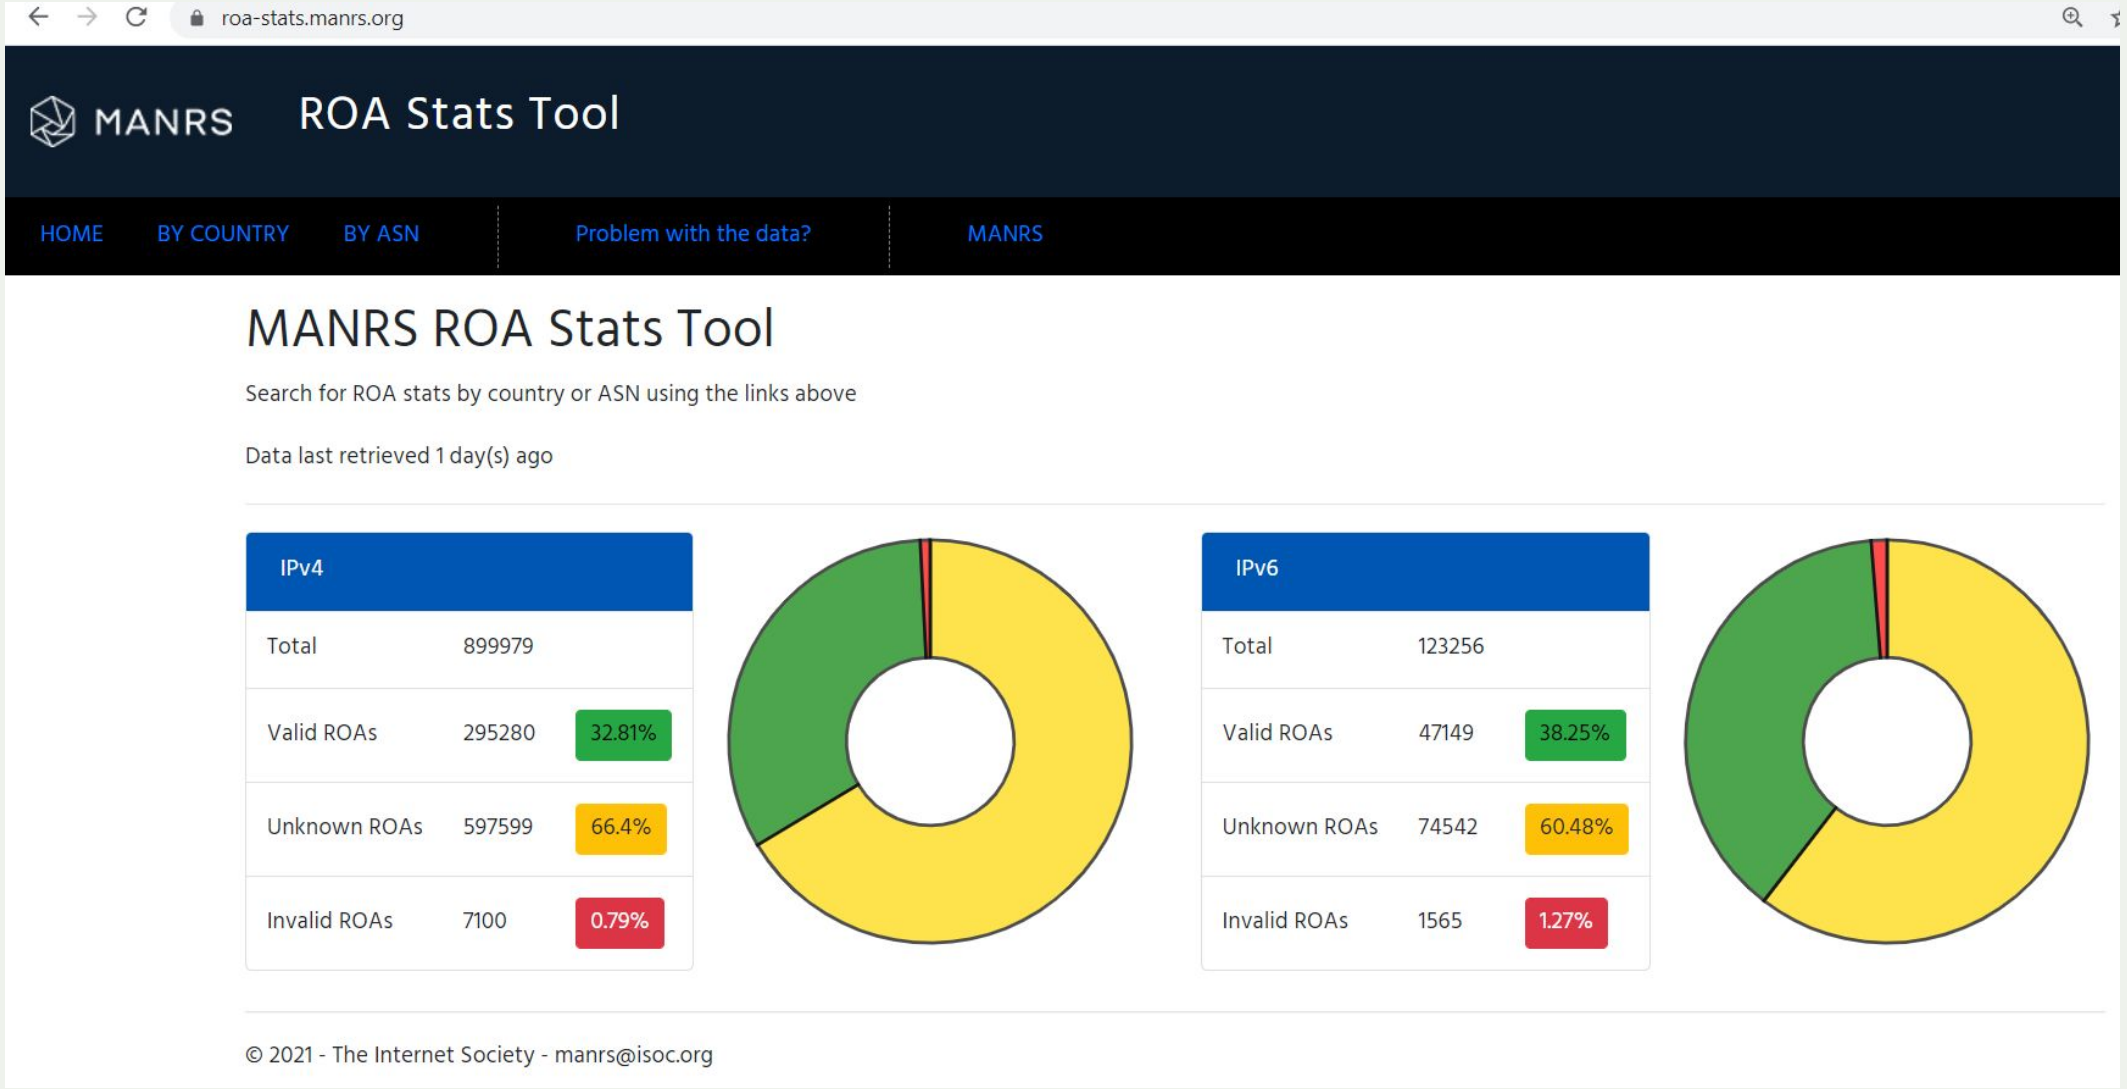

Search Country, ASN

Country report for Fiji

Data last retrieved 1 day(s) ago

IPv4

Prefix	ASN	ASN Name	Status
45.117.240.0/22	45355	DIGICELPACIFIC-1-AP Digicel Fiji Limited	Valid
27.123.128.0/18	38442	VODAFONEFIJI-AS-FJ Vodafone Fiji Limited	Valid
45.112.224.0/22	4638	IS-FJ-AS Telecom Fiji Limited	Valid
45.117.240.0/24	45355	DIGICELPACIFIC-1-AP Digicel Fiji Limited, FJ	Valid
103.58.20.0/22	45355	DIGICELPACIFIC-1-AP Digicel Fiji Limited	Valid

#Protect the Core

LEARN MORE: https://www.manrs.org

https://www.manrs.org/join/