MANRS for Network Operators

Naveen Lakshman
Graphite Networks

Musa Stephen Honlue
AfriNIC

Agenda

- Routing Security
- Any Solution/s?
- MANRS Actions
  - Filtering
  - Anti Spoofing
  - Coordination
  - Global Validation (IRR/RPKI)
- MANRS Observatory
- MANRS Actions Hands-on Lab (Cisco IOS Platform)

Acknowledgements

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

The Problem

A Routing Security Overview

Routing Incidents are increasing (Vodafone Idea AS55410 Hijack)

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

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

https://www.manrs.org/2021/04/a-major-bgp-hijack-by-as55410-vodafone-idea-ltd/

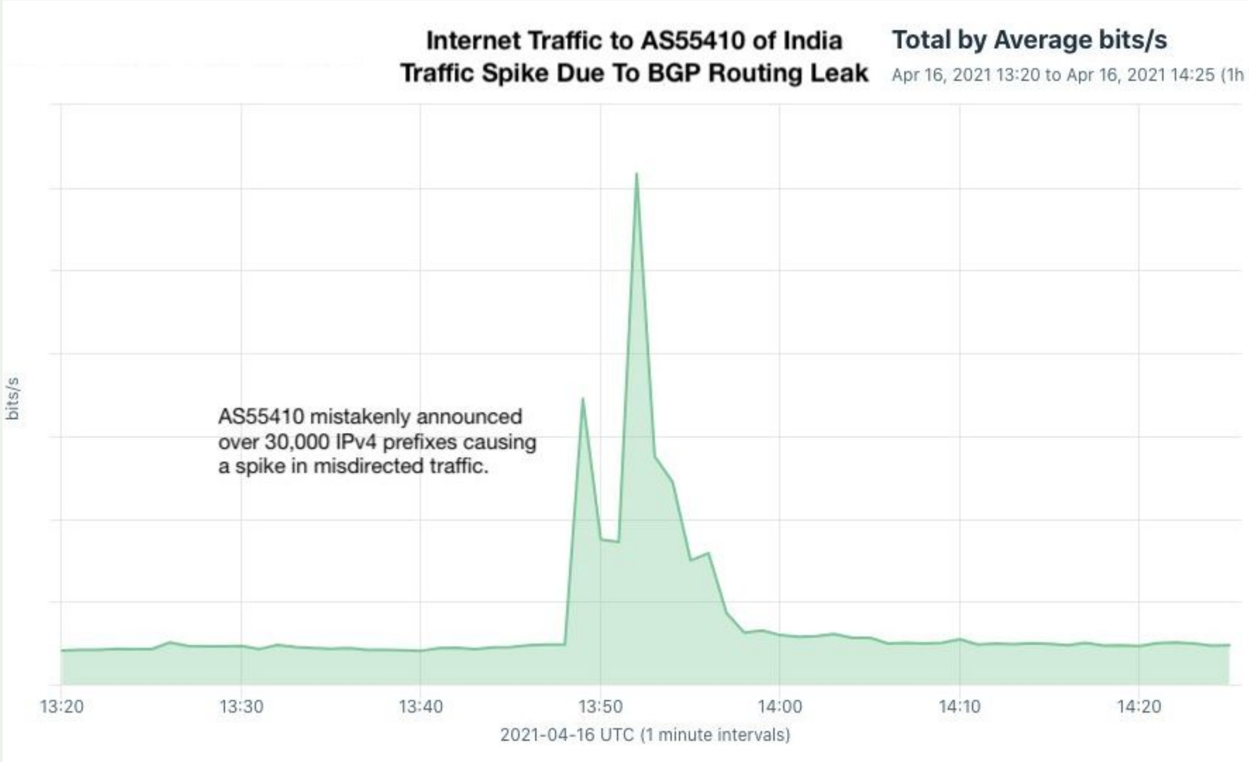

https://twitter.com/DougMadory/status/1383138595112955909

Routing Incidents Cause Real World Problems

- Insecure routing is one of the most common paths for malicious threats.
- Attacks can take anywhere from hours to months to recognize.
- Inadvertent errors can take entire countries offline, while attackers can steal an individual’s data or hold an organization’s network hostage.

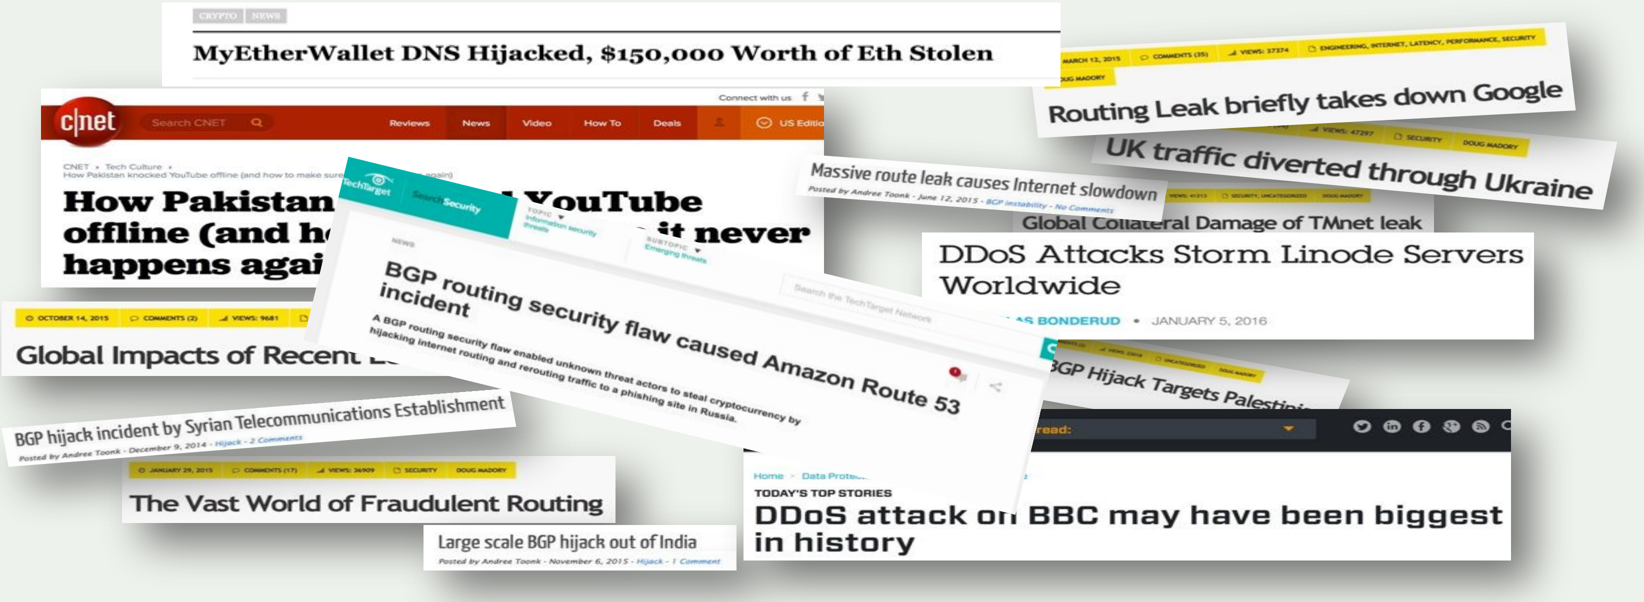

CRYPTO NEWS

MyEtherWallet DNS Hijacked, $150,000 Worth of Eth Stolen

MARCH 12, 2015 · COMMENTS (35) · VIEWS: 37374 · ENGINEERING, INTERNET, LATENCY, PERFORMANCE, SECURITY · DOUG MADORY

Routing Leak briefly takes down Google

VIEWS: 47297 · SECURITY · DOUG MADORY

UK traffic diverted through Ukraine

CNET › Tech Culture ›
How Pakistan knocked YouTube offline (and how to make sure ... again)

How Pakistan ... YouTube offline (and h... it never happens agai...

Massive route leak causes Internet slowdown
Posted by Andree Toonk - June 12, 2015 - BGP instability - No Comments

VIEWS: 41213 · SECURITY, UNCATEGORIZED · DOUG MADORY

Global Collateral Damage of TMnet leak

TechTarget · SearchSecurity · TOPIC Information security threats · SUBTOPIC Emerging threats

NEWS

BGP routing security flaw caused Amazon Route 53 incident

A BGP routing security flaw enabled unknown threat actors to steal cryptocurrency by hijacking internet routing and rerouting traffic to a phishing site in Russia.

DDoS Attacks Storm Linode Servers Worldwide

...AS BONDERUD • JANUARY 5, 2016

OCTOBER 14, 2015 · COMMENTS (2) · VIEWS: 9681

Global Impacts of Recent L...

...BGP Hijack Targets Palestin...

BGP hijack incident by Syrian Telecommunications Establishment
Posted by Andree Toonk - December 9, 2014 - Hijack - 2 Comments

JANUARY 29, 2015 · COMMENTS (17) · VIEWS: 36909 · SECURITY · DOUG MADORY

The Vast World of Fraudulent Routing

Large scale BGP hijack out of India
Posted by Andree Toonk - November 6, 2015 - Hijack - 1 Comment

Home › Data Prote...
TODAY'S TOP STORIES

DDoS attack on BBC may have been biggest in history

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 or pretends that a server or network is their client. 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.

Fix: Strong filtering policies (adjacent networks should strengthen their filtering policies to avoid accepting false announcements).

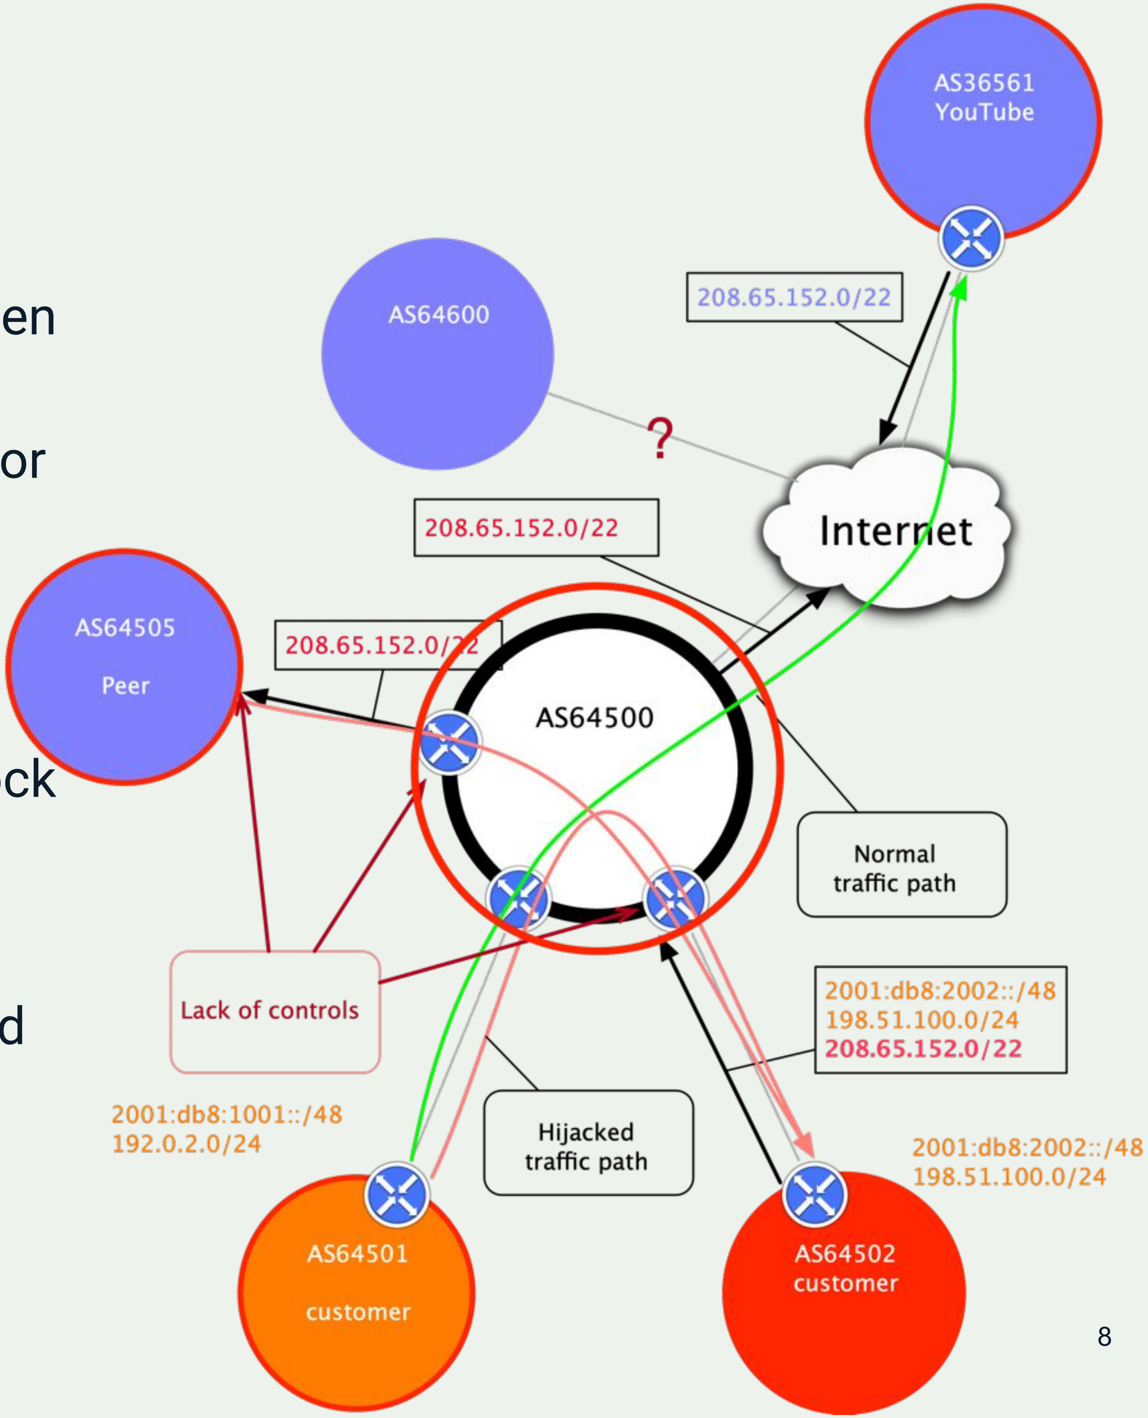

Prefix Hijacking

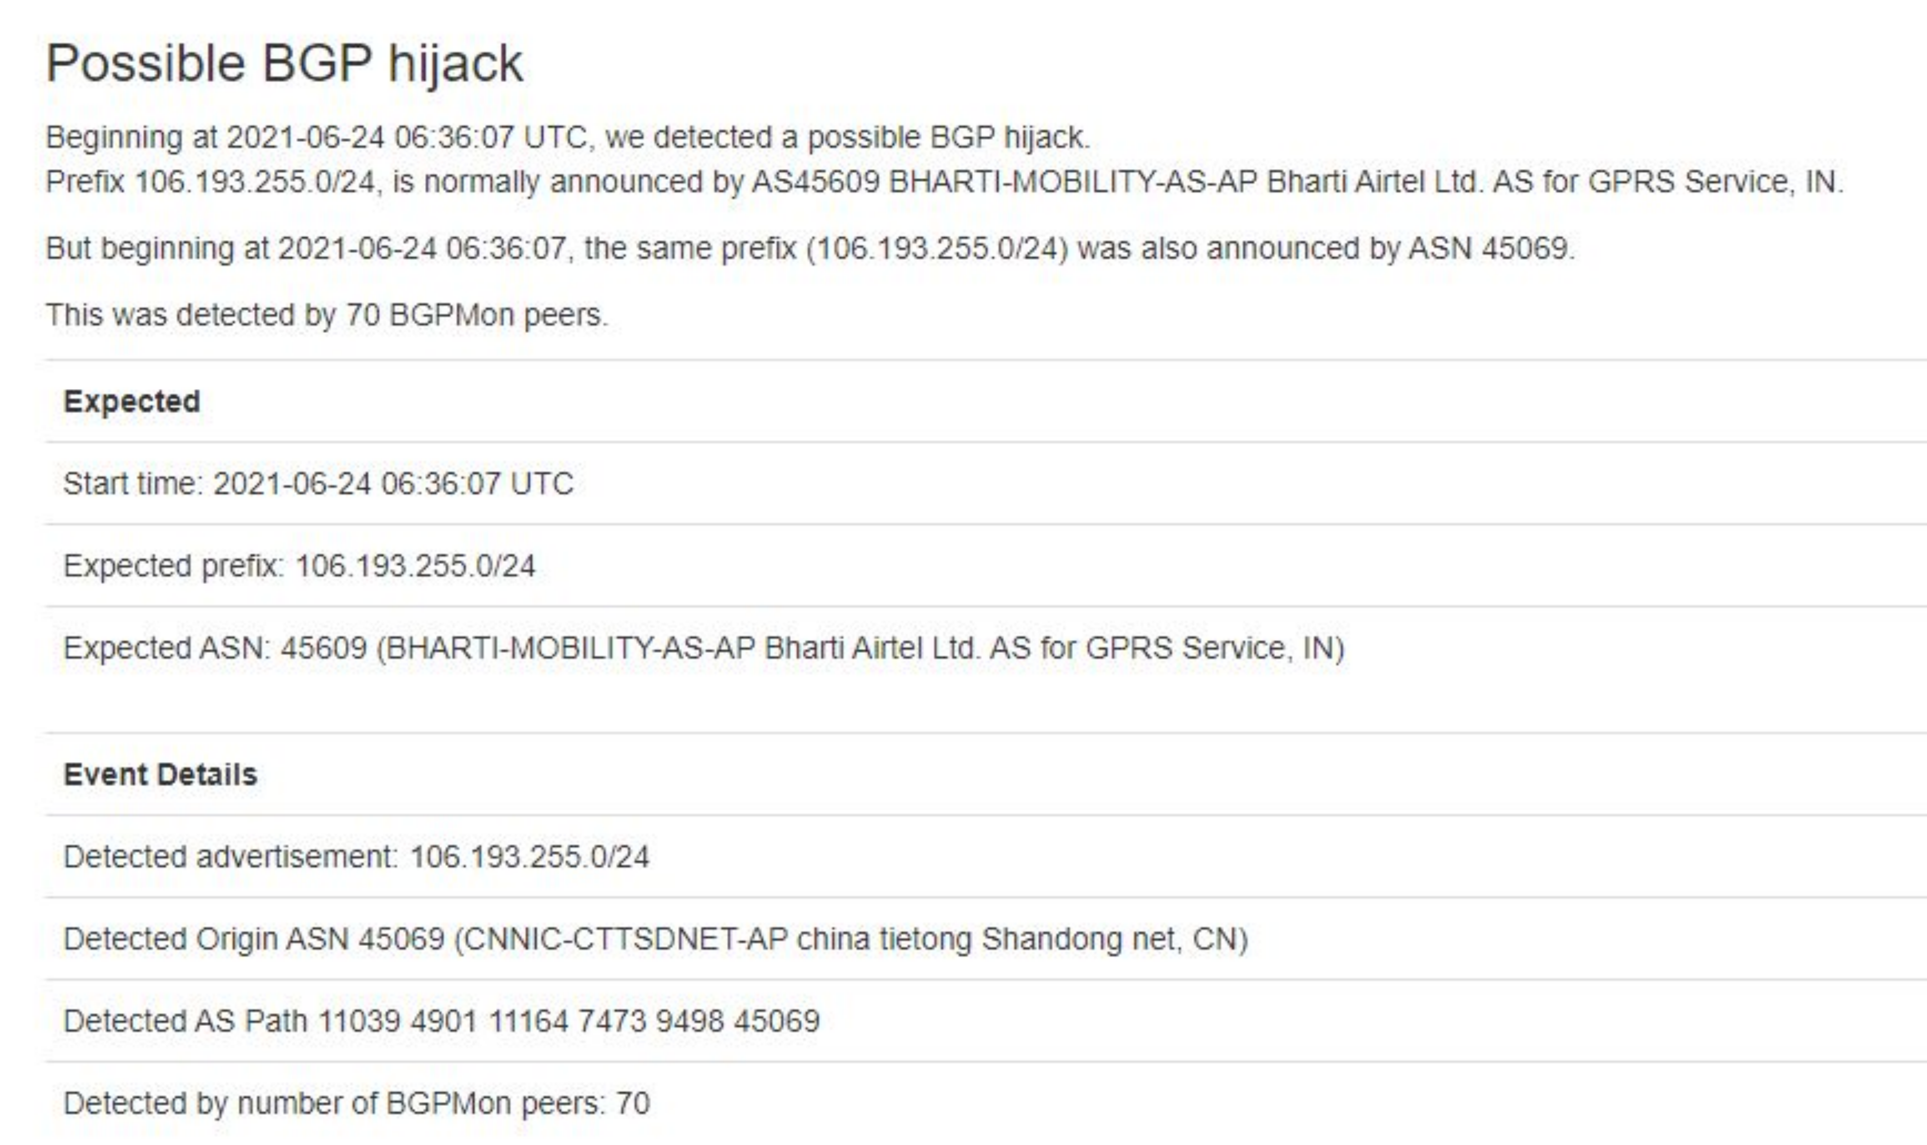

Possible BGP hijack

Beginning at 2021-06-24 06:36:07 UTC, we detected a possible BGP hijack.
Prefix 106.193.255.0/24, is normally announced by AS45609 BHARTI-MOBILITY-AS-AP Bharti Airtel Ltd. AS for GPRS Service, IN.

But beginning at 2021-06-24 06:36:07, the same prefix (106.193.255.0/24) was also announced by ASN 45069.

This was detected by 70 BGPMon peers.

Expected

Start time: 2021-06-24 06:36:07 UTC

Expected prefix: 106.193.255.0/24

Expected ASN: 45609 (BHARTI-MOBILITY-AS-AP Bharti Airtel Ltd. AS for GPRS Service, IN)

Event Details

Detected advertisement: 106.193.255.0/24

Detected Origin ASN 45069 (CNNIC-CTTSDNET-AP china tietong Shandong net, CN)

Detected AS Path 11039 4901 11164 7473 9498 45069

Detected by number of BGPMon peers: 70

Source: bgpstream.com

Route Leak

A route leak is where a network operator with multiple upstream providers accidentally announces to one of its upstream providers that is has a route to a destination through the other upstream provider. This makes the network an intermediary network between the two upstream providers, with one sending traffic through the network to get to the other.

Example: 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).

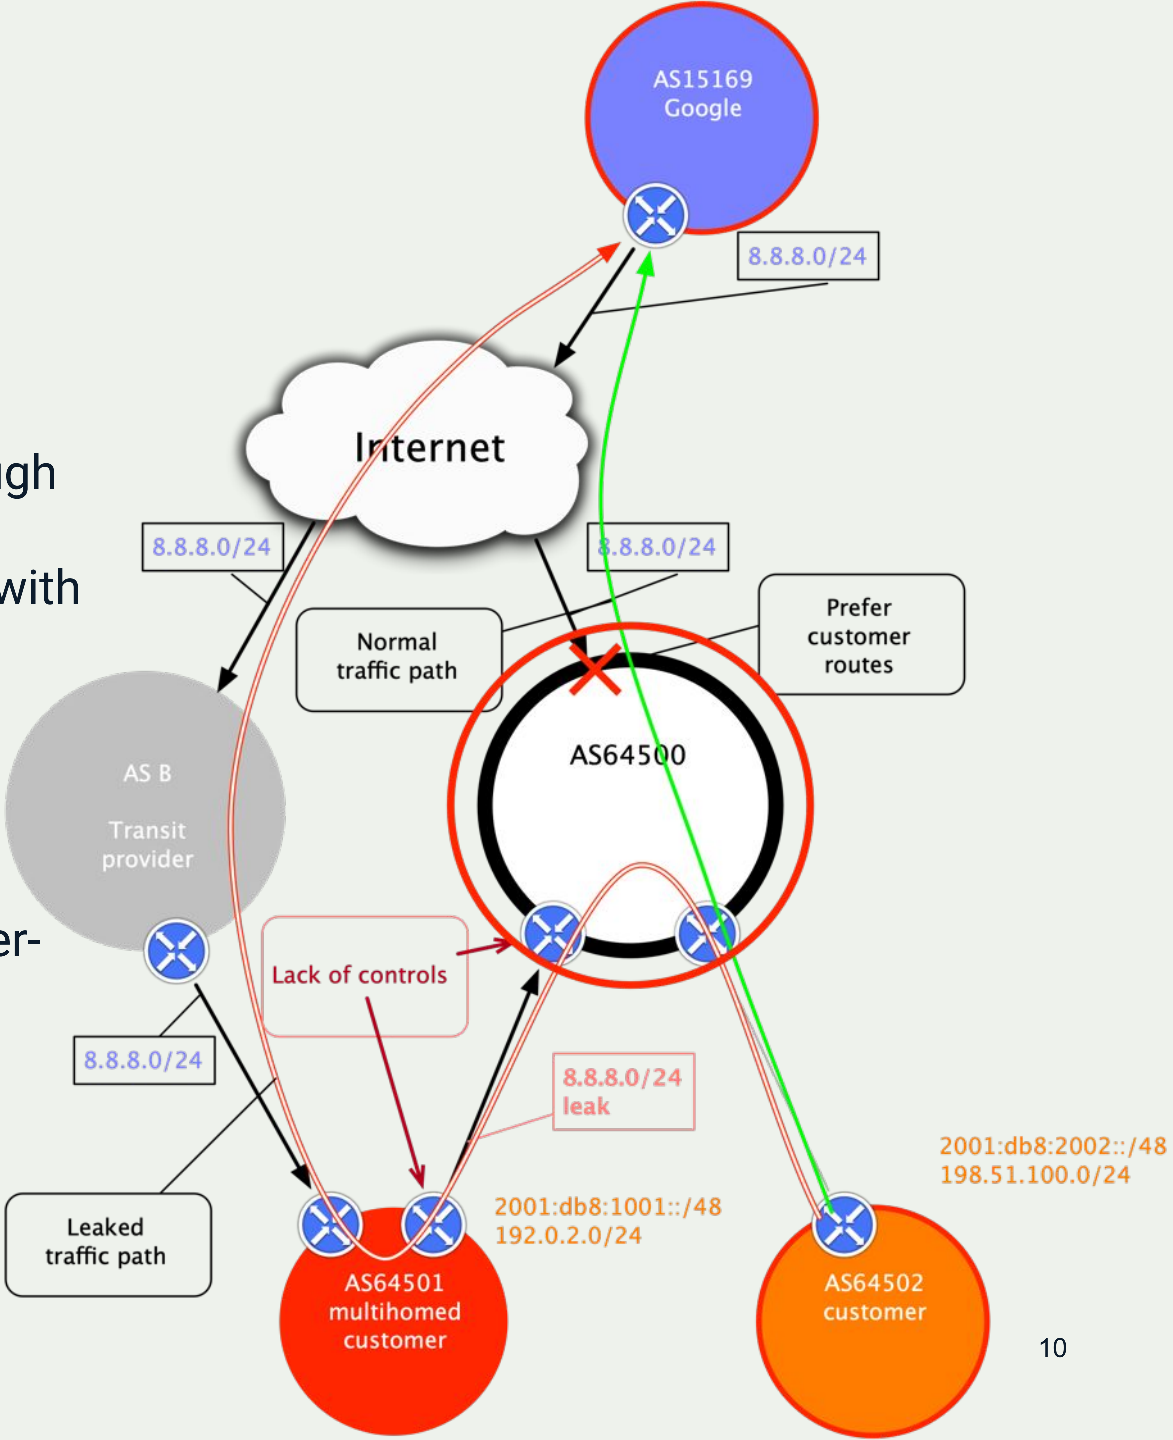

Route Leak

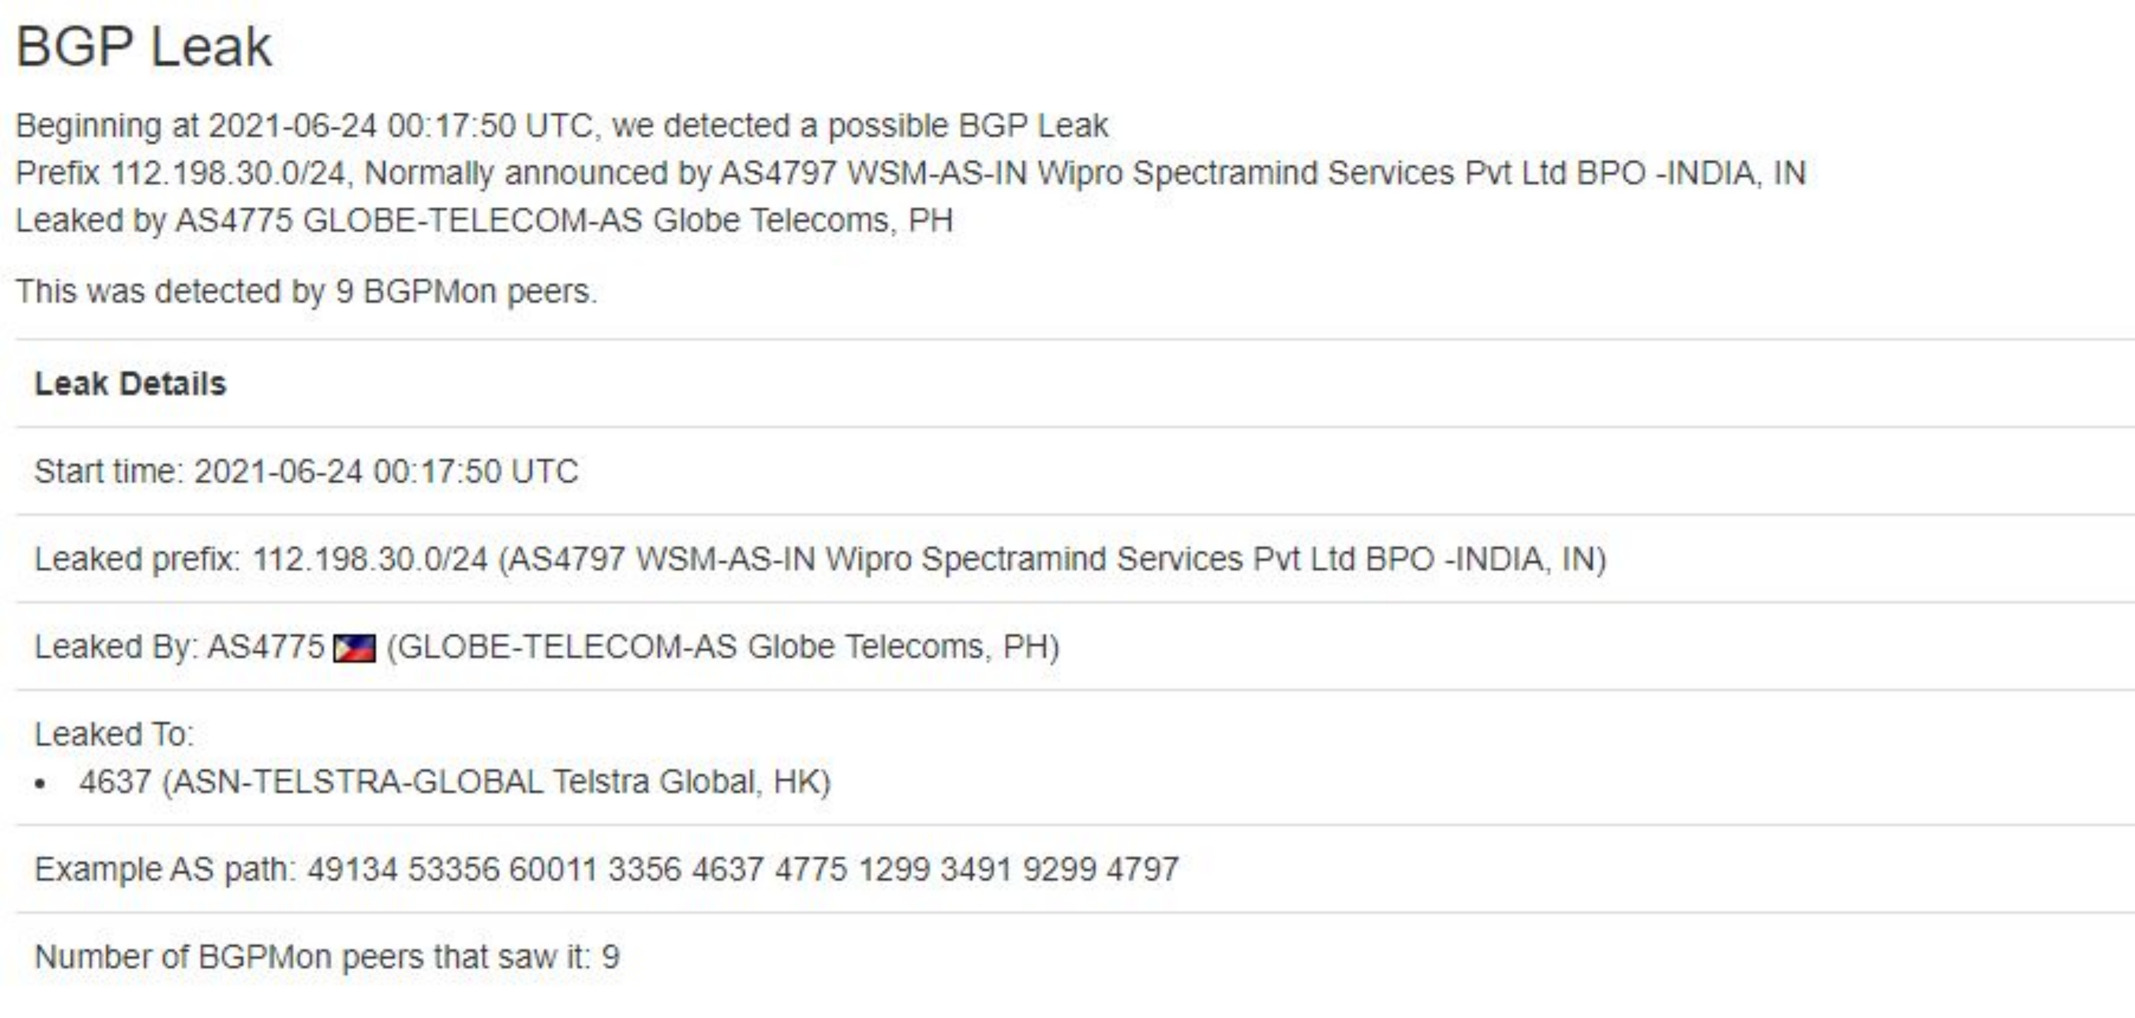

BGP Leak

Beginning at 2021-06-24 00:17:50 UTC, we detected a possible BGP Leak
Prefix 112.198.30.0/24, Normally announced by AS4797 WSM-AS-IN Wipro Spectramind Services Pvt Ltd BPO -INDIA, IN
Leaked by AS4775 GLOBE-TELECOM-AS Globe Telecoms, PH

This was detected by 9 BGPMon peers.

Leak Details

Start time: 2021-06-24 00:17:50 UTC

Leaked prefix: 112.198.30.0/24 (AS4797 WSM-AS-IN Wipro Spectramind Services Pvt Ltd BPO -INDIA, IN)

Leaked By: AS4775 🇵🇭 (GLOBE-TELECOM-AS Globe Telecoms, PH)

Leaked To:

- 4637 (ASN-TELSTRA-GLOBAL Telstra Global, HK)

Example AS path: 49134 53356 60011 3356 4637 4775 1299 3491 9299 4797

Number of BGPMon peers that saw it: 9

Source: bgpstream.com

Routing Incidents (South Asia) May ~ June 2021

Event Type	Event Details	Prefixes affected
BGP Hijack	Expected Origin: AS45609 BHARTI-MOBILITY-AS-AP Bharti Airtel Ltd Detected Origin: ASN 45069 CNNIC-CTTSDNET-AP China Tietong Shandong net, CN	106.193.255.0/24
BGP Leak	Origin AS: AS 4797 Wipro Spectramind Services Pvt Ltd, IN Leaker AS: AS4775 GLOBE-TELECOM-AS Globe Telecoms, PH Leaked to: AS 4637 (ASN-TELSTRA-GLOBAL Telstra Global, HK)	112.198.30.0/24
BGP Leak	Origin AS: AS132497 DNA-AS-AP DIGITAL NETWORK, IN Leaker AS: AS55644 VIL-AS-AP Vodafone Idea Ltd, IN Leaked to: AS3556 (Level3, US) AS3549 (LVLT-3549, US)	150.242.197.0/24
BGP Hijack	Expected Origin: AS328608 Africa-on-Cloud-AS, ZA Detected Origin: ASN 139879 GALAXY-AS-AP Galaxy Broadband, PK	156.241.0.0/16
BGP Hijack	Expected Origin: AS7018 ATT-INTERNET4, US Detected Origin: ASN18229 CTRLS-AS-IN CtrlS Datacenters Ltd., IN	172.0.0.0/12
BGP Hijack	Expected Origin: AS33567 TELECOM-LESOTHO, LS Detected Origin: ASN 55410 (VIL-AS-AP Vodafone Idea Ltd, IN)	41.203.176.0/20
BGP Leak	Origin AS: AS 132497 DIGITAL NETWORK ASSOCIATES, IN Leaker AS:AS 55644 Vodafone Idea Ltd, IN (AS 55644) Leaked to: AS3356 (LEVEL3, US)	103.245.69.0/24

Source: bgpstream.com

Tools to Help

- Prefix and AS-PATH filtering
- RPKI, IRR toolset, IRRPT, BGPQ3/Q4
- BGPSEC is standardized

But…

- Not enough deployment
- Lack of reliable data

We need a standard approach to improving routing security.

The Solution:

Mutually Agreed Norms for Routing Security (MANRS)

Provides crucial fixes to eliminate the most common routing threats

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

MANRS Network Operators Programme

Launched in 2014 by a handful of network operators with the following goals:

- Raise awareness of routing security problems and encourage the implementation of actions that can address them.
- Promote a culture of collective responsibility toward the security and resilience of the Internet’s global routing system.
- Demonstrate the ability of the Internet industry to address routing security problems.
- Provide a framework for network operators to better understand and address issues relating to the security and resilience of the Internet’s global routing system.

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

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

Blue shading = Mandatory Action

MANRS IXP Programme

Internet Exchange Points (IXPs) are a collaborative focal point to discuss and promote the importance of routing security.

Launched in 2018, the IXP Programme addresses the unique needs and concerns of IXPs with a separate set of MANRS actions.

IXPs can implement actions that demonstrate their commitment to routing security and bring significant improvement to the resilience and security of their peering relationships.

https://www.manrs.org/ixps/

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 CDN and Cloud Programme

Launched in 2020, the CDN and Cloud Provider Programme helps by requiring egress routing controls so networks can prevent incidents from happening.

Leveraging CDNs' and cloud providers' peering power can have significant positive spillover effect on the routing hygiene of networks they peer with.

Goals include:

- Create a secure network peering environment
- Encourage better routing hygiene from peering partners
- Demonstrate responsible behavior
- Improve operational efficiency for peering interconnections, minimizing incidents and providing more granular insight for troubleshooting

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/

Action 1: Filtering

(BCP 194 – RFC7454)

BGP Operations and Security

Filtering

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

- Operator defines a clear routing policy and implements the filter that ensures correctness of their own announcements and customers prefixes to adjacent networks.
- Operator applies due diligence when checking the correctness of its customer's announcements, specifically that the customer legitimately holds the ASN and the address space it announces.

RFC7454 (BCP 194)

- Generalized TTL Security Mechanism (GTSM)
- TCP Authentication Option (TCP-AO)
- Prefix filtering and automation of prefix filters
- Max-prefix filtering
- Autonomous System (AS) path filtering
- BGP community scrubbing

BCP 194 - Prefix Filtering

Inbound Filtering Loose Option

In this case, the following prefixes received from a BGP peer will be filtered:

- prefixes that are not globally routable
- prefixes not allocated by IANA (IPv6 only)
- routes that are too specific
- prefixes belonging to the local AS
- IXP LAN prefixes
- default route (depending on whether or not the route is requested)

BCP 194 - Prefix Filtering

Inbound Filtering Strict Option

In this case, filters are applied to make sure advertisements strictly conform to what is declared in routing registries. This varies across the registries and regions of the Internet.

Apply the following filters beforehand in case the routing registry that's used as the source of information by the script is not fully trusted:

- prefixes that are not globally routable
- routes that are too specific
- prefixes belonging to the local AS
- IXP LAN prefixes
- the default route (depending on whether or not the route is requested)

BCP 194 - Prefix Filtering

Outbound Filtering

The configuration should ensure that only appropriate prefixes are sent. These can be, for example, prefixes belonging to both the network in question and its downstream. This can be achieved by using BGP communities, AS paths, or both. It may be desirable to add the following filter:

- prefixes that are not globally routable
- routes that are too specific
- IXP LAN prefixes
- the default route (not willing to receive it)

BCP 194 - Max Prefix Filtering

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, it is RECOMMENDED to 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.

Data Sources

IRRs

https://wq.apnic.net/static/search.html

PeeringDB - For AS-Sets

https://www.peeringdb.com/

Bogons lists (IPv6 & IPv4)

https://www.team-cymru.org/Services/Bogons/fullbogons-ipv4.txt

https://www.team-cymru.org/Services/Bogons/fullbogons-ipv6.txt

ASN Bogons

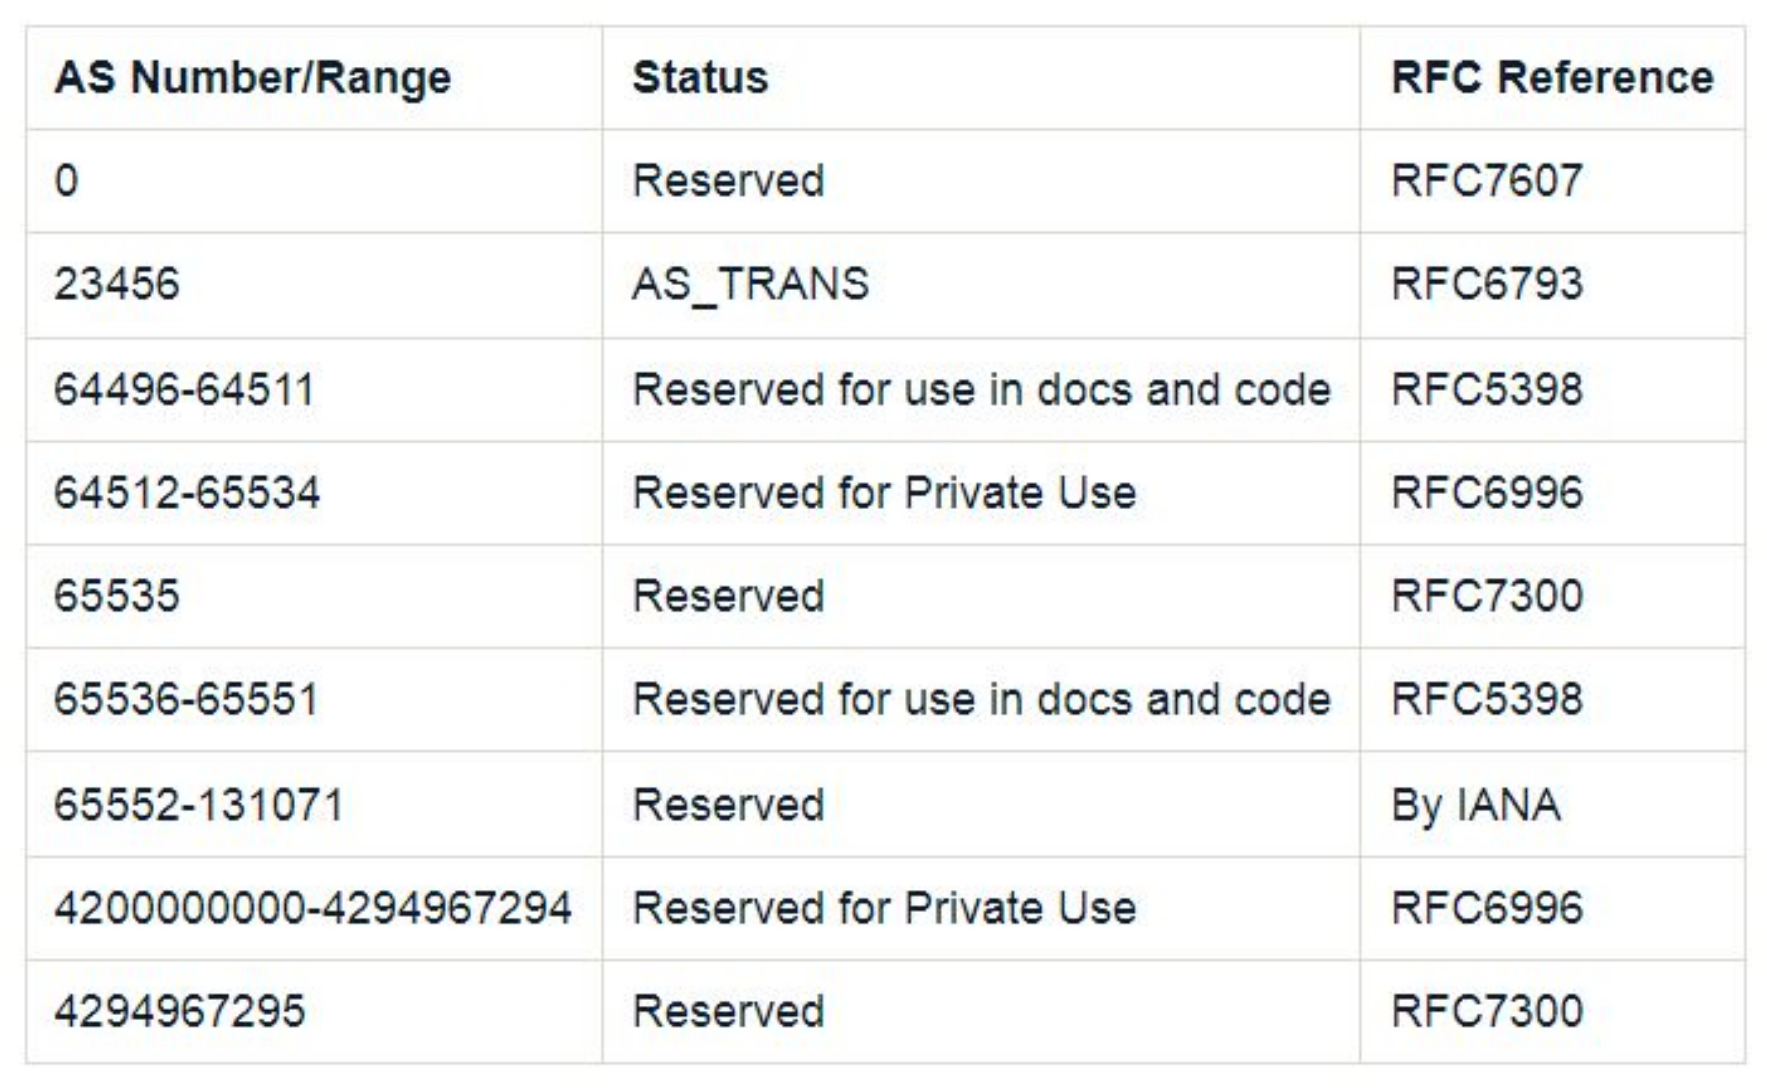

AS Number/Range	Status	RFC Reference
0	Reserved	RFC7607
23456	AS_TRANS	RFC6793
64496-64511	Reserved for use in docs and code	RFC5398
64512-65534	Reserved for Private Use	RFC6996
65535	Reserved	RFC7300
65536-65551	Reserved for use in docs and code	RFC5398
65552-131071	Reserved	By IANA
4200000000-4294967294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

https://www.manrs.org/2021/01/routing-security-terms-bogons-vogons-and-martians/

Generating a Prefix Filter

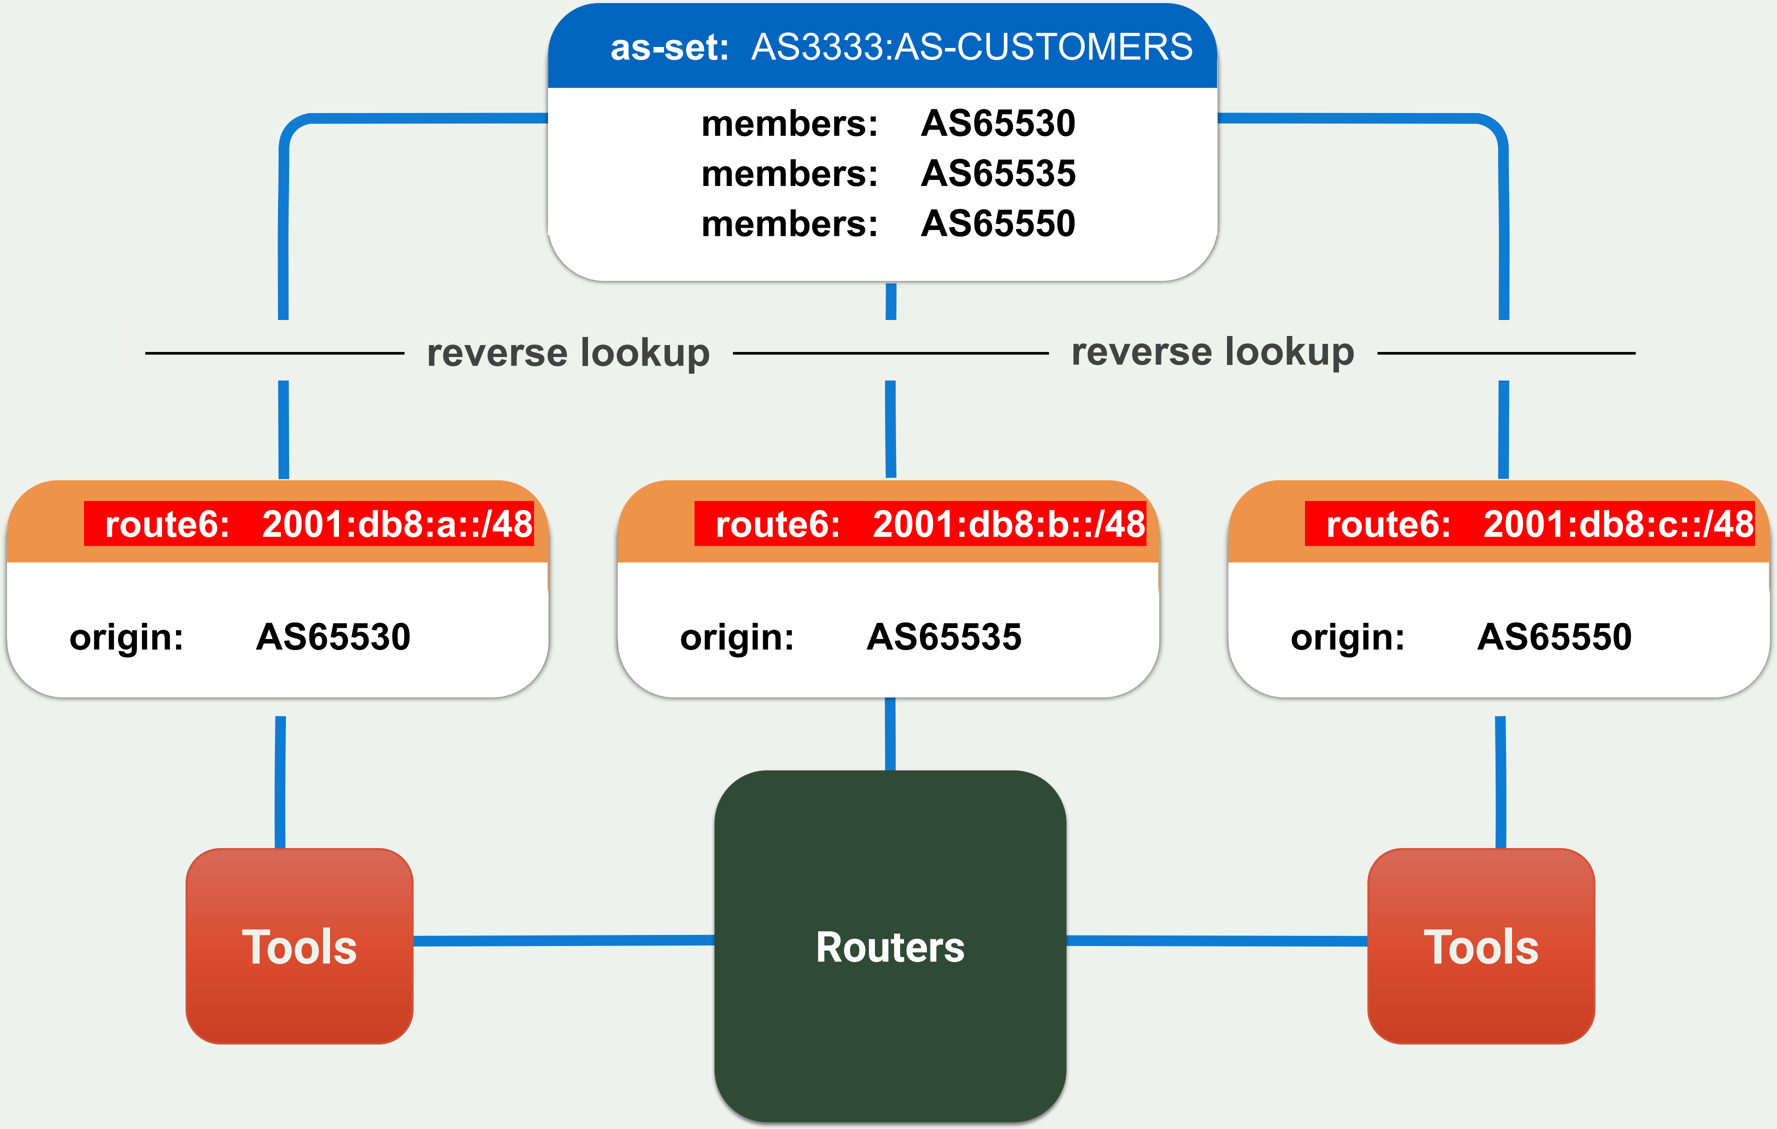

Generating a Prefix list

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

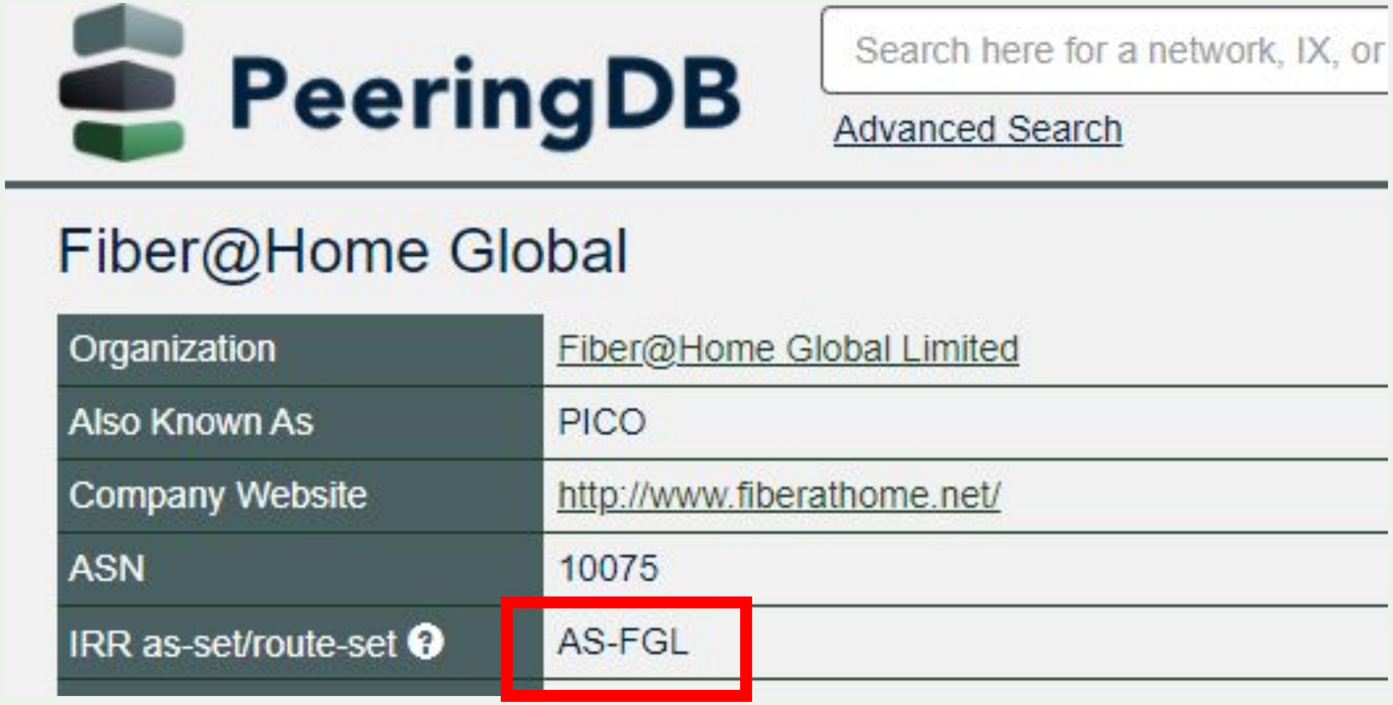

Fiber@Home Global

Organization	Fiber@Home Global Limited
Also Known As	PICO
Company Website	http://www.fiberathome.net/
ASN	10075
IRR as-set/route-set ?	AS-FGL

AS-Set of AS10075 from peeringdb

Generating a prefix list

With the list of ASNs, run an inverse query for each one

- Get the route objects where they are listed as Origin

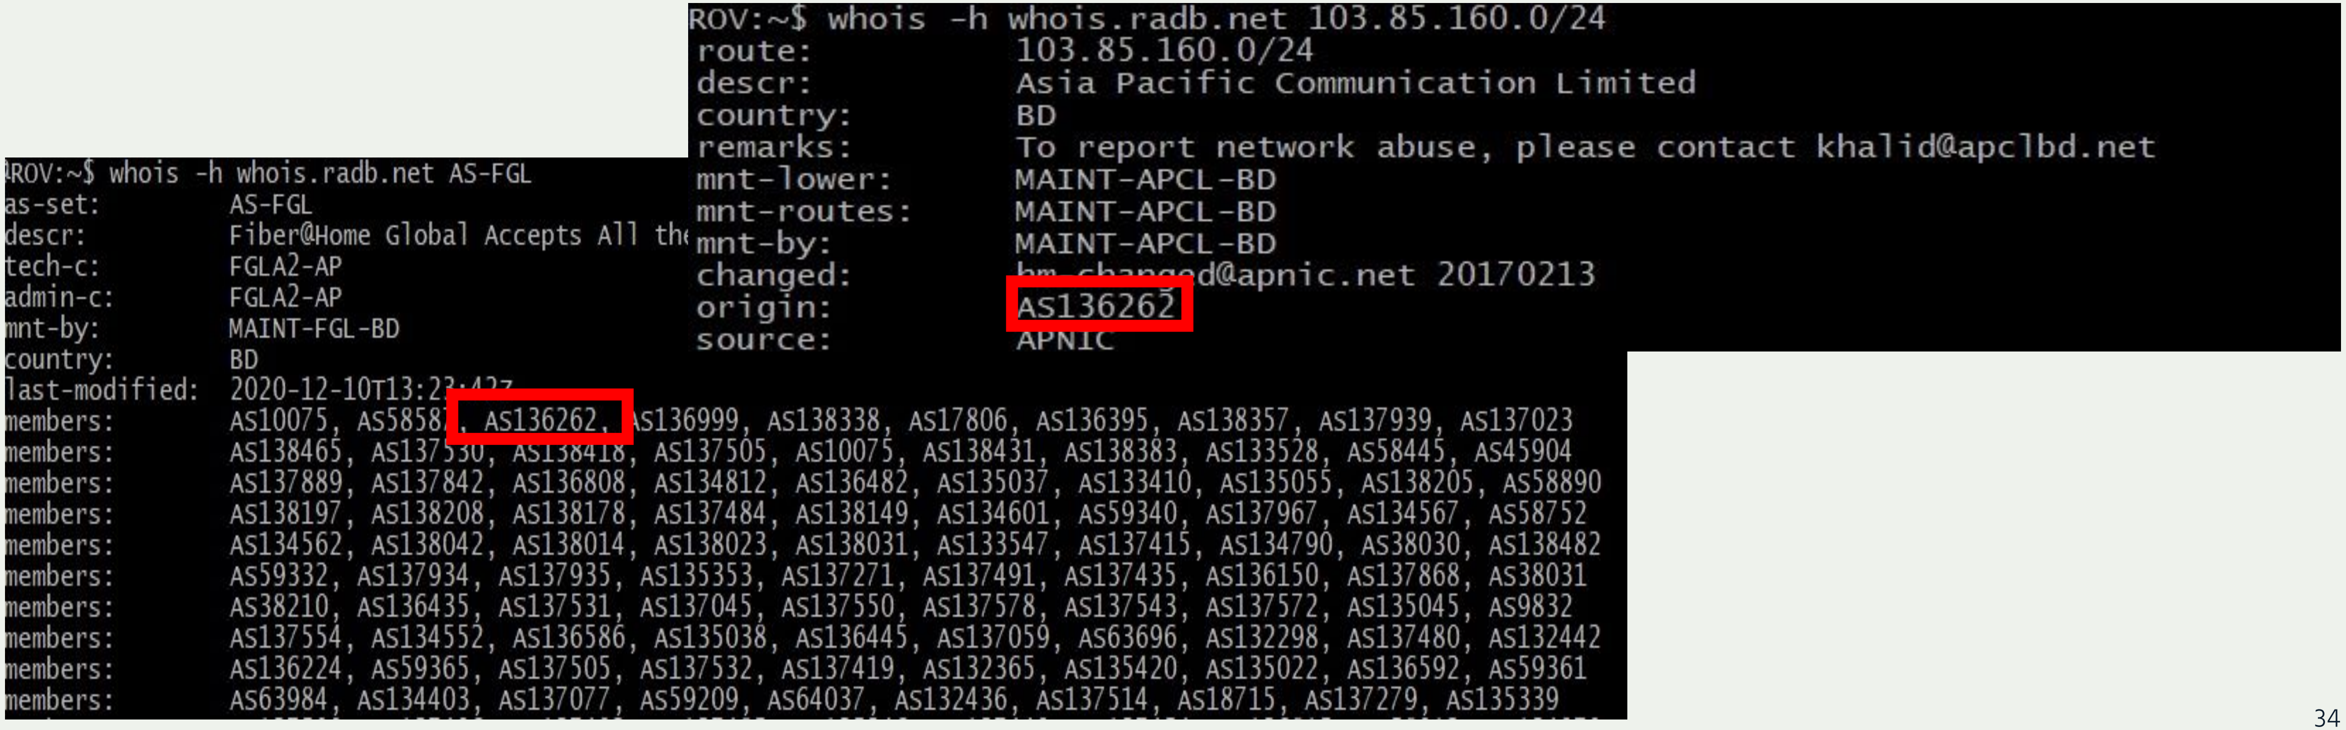

```
ROV:~$ whois -h whois.radb.net 103.85.160.0/24
route:          103.85.160.0/24
descr:          Asia Pacific Communication Limited
country:        BD
remarks:        To report network abuse, please contact khalid@apclbd.net
mnt-lower:      MAINT-APCL-BD
mnt-routes:     MAINT-APCL-BD
mnt-by:         MAINT-APCL-BD
changed:        hm-changed@apnic.net 20170213
origin:         AS136262
source:         APNIC
```

```
ROV:~$ whois -h whois.radb.net AS-FGL
as-set:         AS-FGL
descr:          Fiber@Home Global Accepts All the
tech-c:         FGLA2-AP
admin-c:        FGLA2-AP
mnt-by:         MAINT-FGL-BD
country:        BD
last-modified:  2020-12-10T13:23:42Z
members:        AS10075, AS58587, AS136262, AS136999, AS138338, AS17806, AS136395, AS138357, AS137939, AS137023
members:        AS138465, AS137530, AS138418, AS137505, AS10075, AS138431, AS138383, AS133528, AS58445, AS45904
members:        AS137889, AS137842, AS136808, AS134812, AS136482, AS135037, AS133410, AS135055, AS138205, AS58890
members:        AS138197, AS138208, AS138178, AS137484, AS138149, AS134601, AS59340, AS137967, AS134567, AS58752
members:        AS134562, AS138042, AS138014, AS138023, AS138031, AS133547, AS137415, AS134790, AS38030, AS138482
members:        AS59332, AS137934, AS137935, AS135353, AS137271, AS137491, AS137435, AS136150, AS137868, AS38031
members:        AS38210, AS136435, AS137531, AS137045, AS137550, AS137578, AS137543, AS137572, AS135045, AS9832
members:        AS137554, AS134552, AS136586, AS135038, AS136445, AS137059, AS63696, AS132298, AS137480, AS132442
members:        AS136224, AS59365, AS137505, AS137532, AS137419, AS132365, AS135420, AS135022, AS136592, AS59361
members:        AS63984, AS134403, AS137077, AS59209, AS64037, AS132436, AS137514, AS18715, AS137279, AS135339
```

Prefix-lists - IRR

Lists of routes you want to accept or announce

You can create them manually or automatically

- With data from IRRs

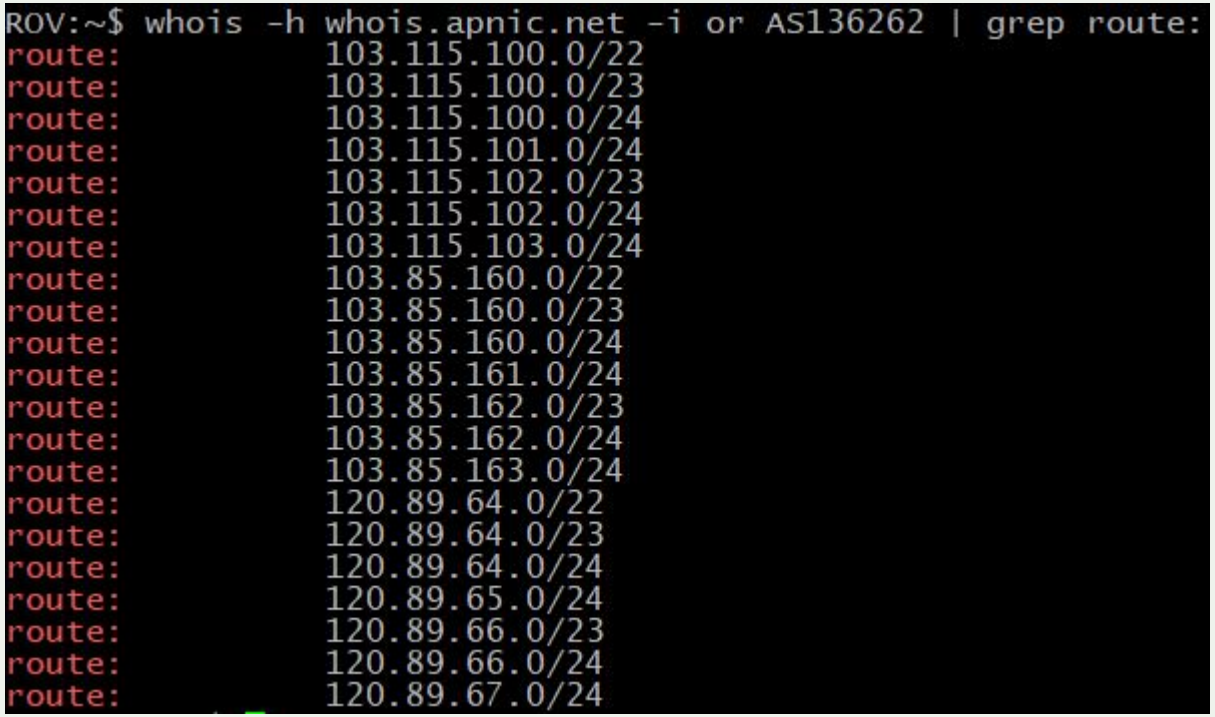

```
ROV:~$ whois -h whois.apnic.net -i or AS136262 | grep route:
route:          103.115.100.0/22
route:          103.115.100.0/23
route:          103.115.100.0/24
route:          103.115.101.0/24
route:          103.115.102.0/23
route:          103.115.102.0/24
route:          103.115.103.0/24
route:          103.85.160.0/22
route:          103.85.160.0/23
route:          103.85.160.0/24
route:          103.85.161.0/24
route:          103.85.162.0/23
route:          103.85.162.0/24
route:          103.85.163.0/24
route:          120.89.64.0/22
route:          120.89.64.0/23
route:          120.89.64.0/24
route:          120.89.65.0/24
route:          120.89.66.0/23
route:          120.89.66.0/24
route:          120.89.67.0/24
```

Prefix-lists - Tools

Tools are there to help you

- bgpq3/bgpq4
- Level3 Filtergen

```
$ bgpq3 -4 -l AS64500-v4 AS64500:AS-ALL
 no ip prefix-list AS64500-v4
 ip prefix-list AS64500-v4 permit 203.0.113.0/24
 ip prefix-list AS64500-v4 permit 192.0.2.0/24
 ip prefix-list AS64500-v4 permit 198.51.100.0/24
```

https://github.com/snar/bgpq3

AS-Filter - Tools

Tools are there to help you

- bgpq3/bgpq4
- Level3 Filtergen

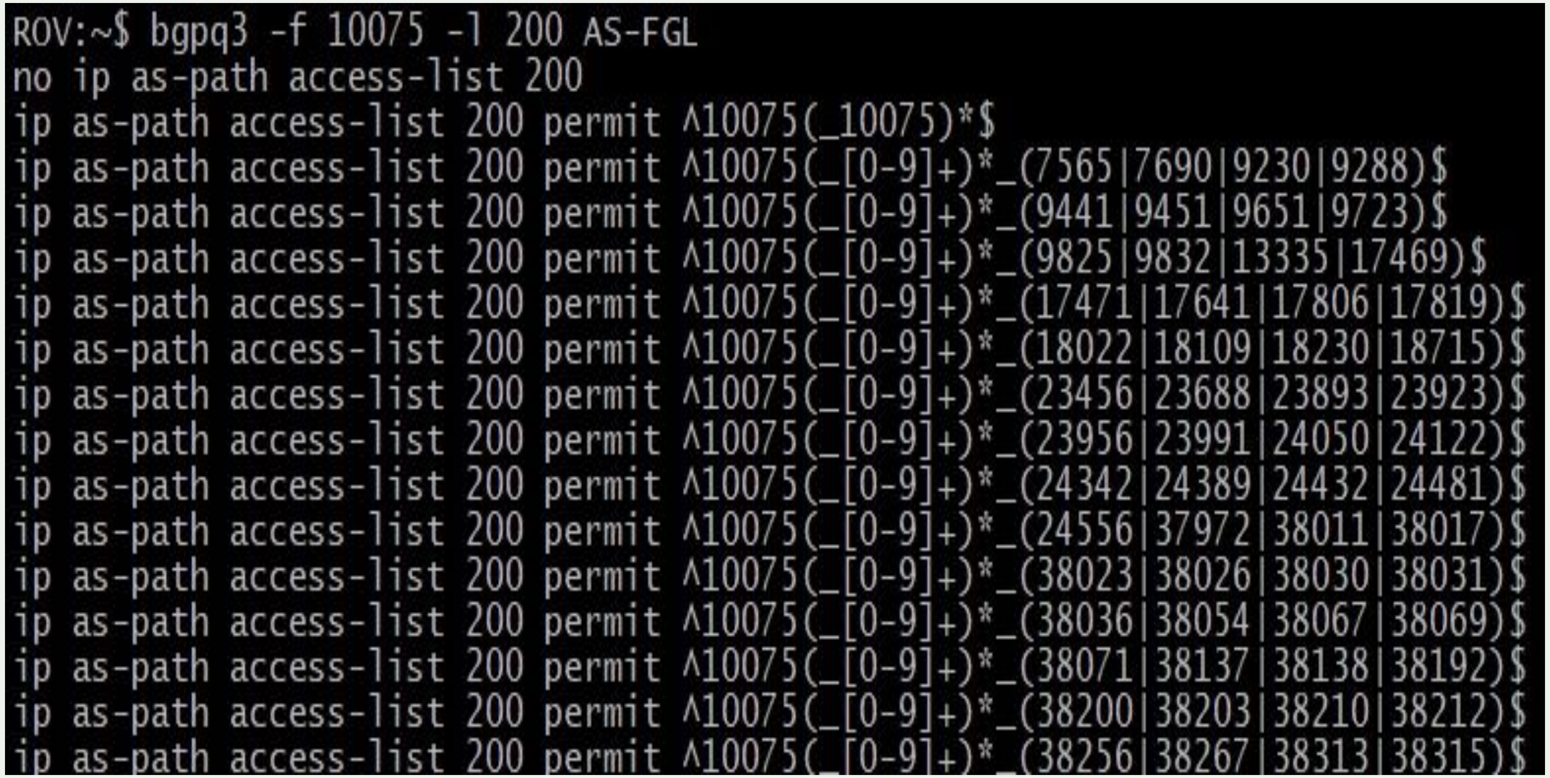

```
ROV:~$ bgpq3 -f 10075 -l 200 AS-FGL
no ip as-path access-list 200
ip as-path access-list 200 permit ^10075(_10075)*$
ip as-path access-list 200 permit ^10075(_[0-9]+)*_(7565|7690|9230|9288)$
ip as-path access-list 200 permit ^10075(_[0-9]+)*_(9441|9451|9651|9723)$
ip as-path access-list 200 permit ^10075(_[0-9]+)*_(9825|9832|13335|17469)$
ip as-path access-list 200 permit ^10075(_[0-9]+)*_(17471|17641|17806|17819)$
ip as-path access-list 200 permit ^10075(_[0-9]+)*_(18022|18109|18230|18715)$
ip as-path access-list 200 permit ^10075(_[0-9]+)*_(23456|23688|23893|23923)$
ip as-path access-list 200 permit ^10075(_[0-9]+)*_(23956|23991|24050|24122)$
ip as-path access-list 200 permit ^10075(_[0-9]+)*_(24342|24389|24432|24481)$
ip as-path access-list 200 permit ^10075(_[0-9]+)*_(24556|37972|38011|38017)$
ip as-path access-list 200 permit ^10075(_[0-9]+)*_(38023|38026|38030|38031)$
ip as-path access-list 200 permit ^10075(_[0-9]+)*_(38036|38054|38067|38069)$
ip as-path access-list 200 permit ^10075(_[0-9]+)*_(38071|38137|38138|38192)$
ip as-path access-list 200 permit ^10075(_[0-9]+)*_(38200|38203|38210|38212)$
ip as-path access-list 200 permit ^10075(_[0-9]+)*_(38256|38267|38313|38315)$
```

https://github.com/snar/bgpq3

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

https://team-cymru.com/community-services/bogon-reference/

The Bogon Route Server Project – Team Cymru

Provides bogon tracking and notification through a multihop eBGP peering session.

This can make the automation of filters simple.

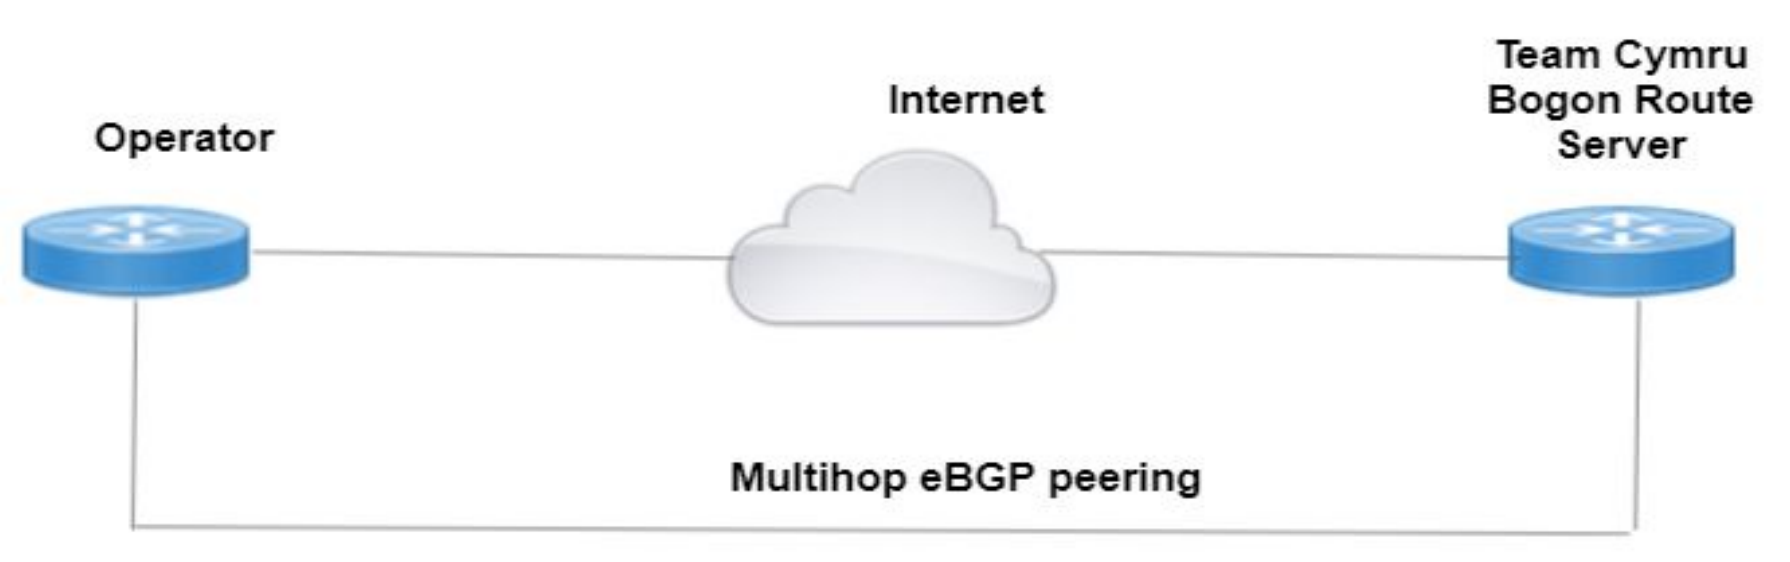

1. Multi Hop BGP
2. Operator advertise Nothing
3. Team Cymru advertise full bogon in both v4 and v6 with community tag.
4. Operator receives the bogon prefixes with tag and point them to null route
5. All the bogon packets are thus dropped.

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 attacks

IP Address Spoofing

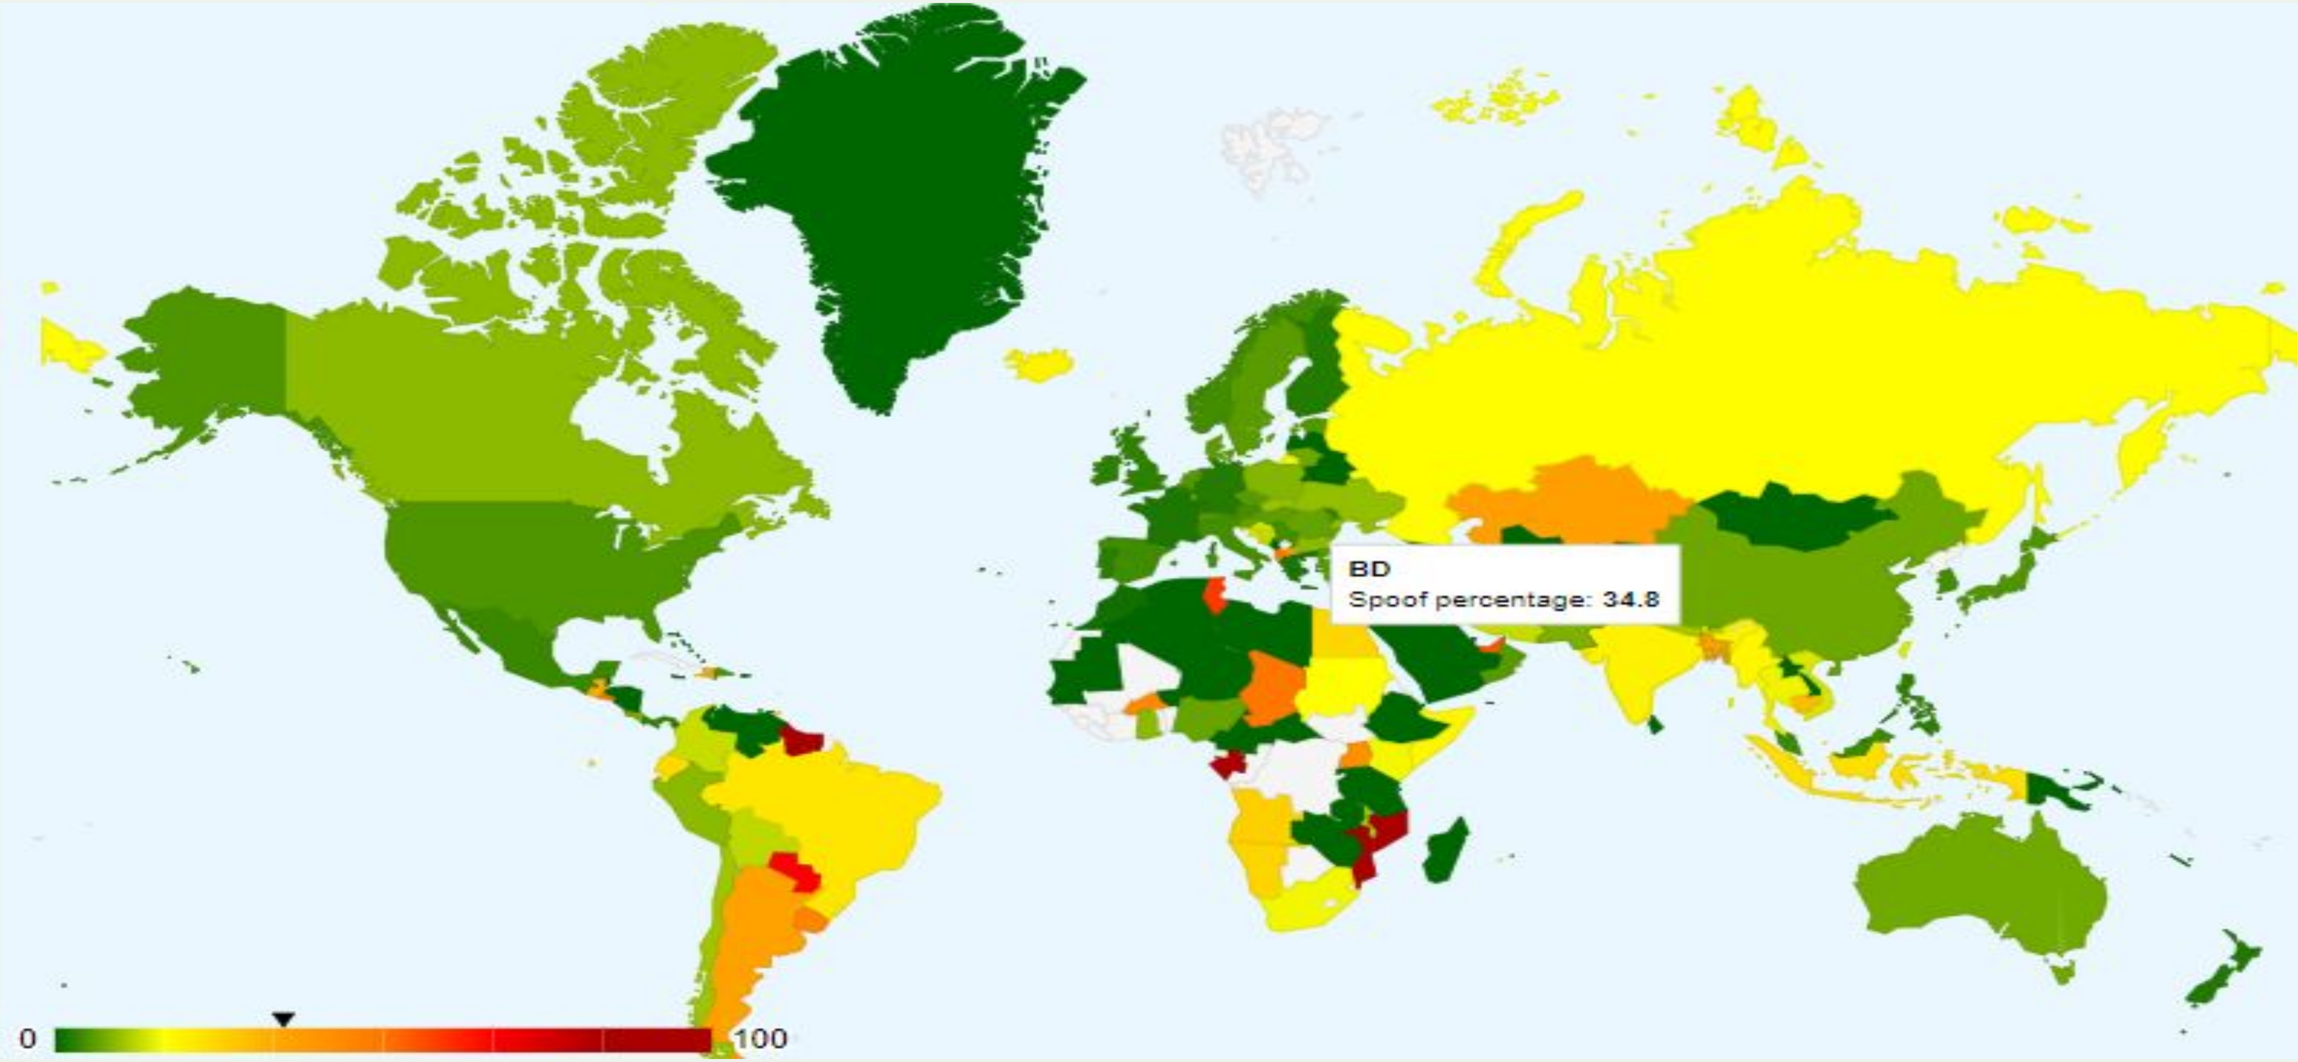

https://spoofer.caida.org/country_stats.php

IP Address Spoofing

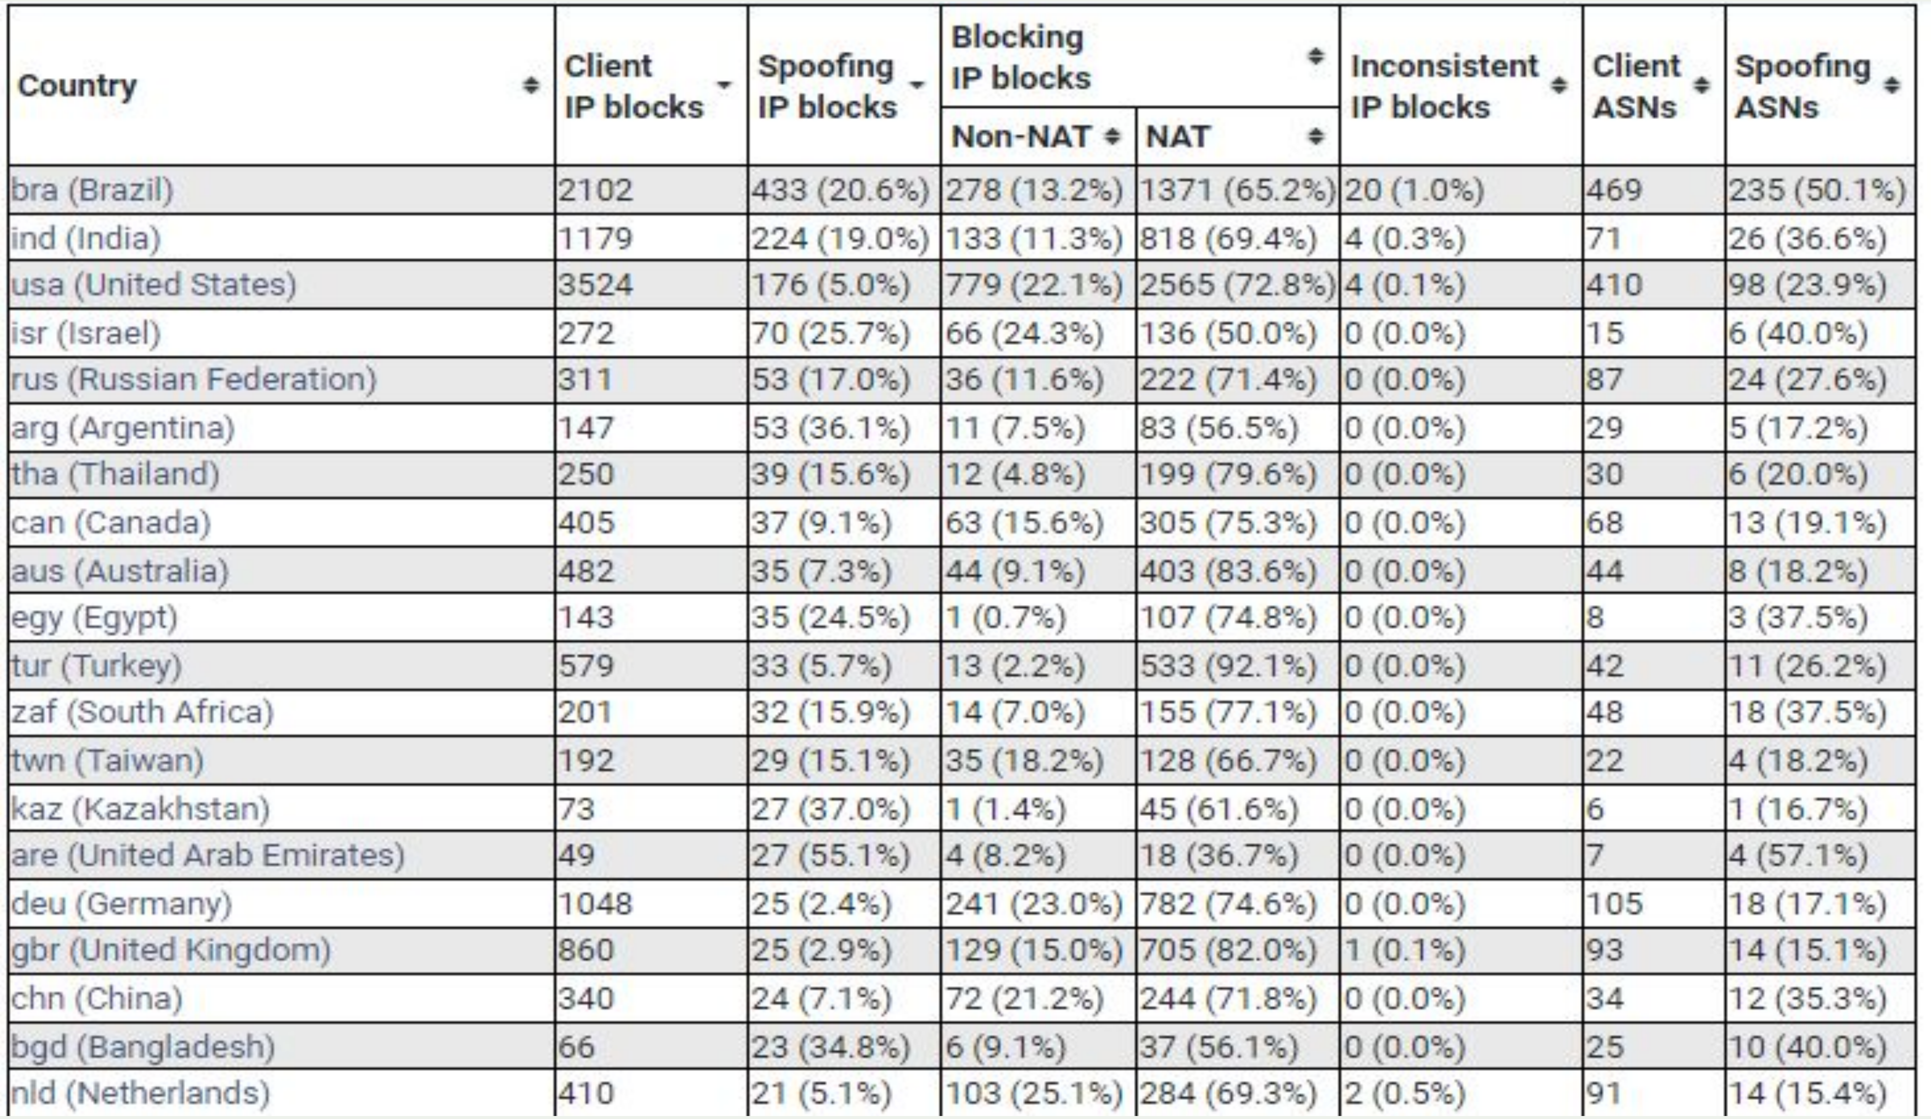

Country	Client IP blocks	Spoofing IP blocks	Blocking IP blocks		Inconsistent IP blocks	Client ASNs	Spoofing ASNs
			Non-NAT	NAT			
bra (Brazil)	2102	433 (20.6%)	278 (13.2%)	1371 (65.2%)	20 (1.0%)	469	235 (50.1%)
ind (India)	1179	224 (19.0%)	133 (11.3%)	818 (69.4%)	4 (0.3%)	71	26 (36.6%)
usa (United States)	3524	176 (5.0%)	779 (22.1%)	2565 (72.8%)	4 (0.1%)	410	98 (23.9%)
isr (Israel)	272	70 (25.7%)	66 (24.3%)	136 (50.0%)	0 (0.0%)	15	6 (40.0%)
rus (Russian Federation)	311	53 (17.0%)	36 (11.6%)	222 (71.4%)	0 (0.0%)	87	24 (27.6%)
arg (Argentina)	147	53 (36.1%)	11 (7.5%)	83 (56.5%)	0 (0.0%)	29	5 (17.2%)
tha (Thailand)	250	39 (15.6%)	12 (4.8%)	199 (79.6%)	0 (0.0%)	30	6 (20.0%)
can (Canada)	405	37 (9.1%)	63 (15.6%)	305 (75.3%)	0 (0.0%)	68	13 (19.1%)
aus (Australia)	482	35 (7.3%)	44 (9.1%)	403 (83.6%)	0 (0.0%)	44	8 (18.2%)
egy (Egypt)	143	35 (24.5%)	1 (0.7%)	107 (74.8%)	0 (0.0%)	8	3 (37.5%)
tur (Turkey)	579	33 (5.7%)	13 (2.2%)	533 (92.1%)	0 (0.0%)	42	11 (26.2%)
zaf (South Africa)	201	32 (15.9%)	14 (7.0%)	155 (77.1%)	0 (0.0%)	48	18 (37.5%)
twn (Taiwan)	192	29 (15.1%)	35 (18.2%)	128 (66.7%)	0 (0.0%)	22	4 (18.2%)
kaz (Kazakhstan)	73	27 (37.0%)	1 (1.4%)	45 (61.6%)	0 (0.0%)	6	1 (16.7%)
are (United Arab Emirates)	49	27 (55.1%)	4 (8.2%)	18 (36.7%)	0 (0.0%)	7	4 (57.1%)
deu (Germany)	1048	25 (2.4%)	241 (23.0%)	782 (74.6%)	0 (0.0%)	105	18 (17.1%)
gbr (United Kingdom)	860	25 (2.9%)	129 (15.0%)	705 (82.0%)	1 (0.1%)	93	14 (15.1%)
chn (China)	340	24 (7.1%)	72 (21.2%)	244 (71.8%)	0 (0.0%)	34	12 (35.3%)
bgd (Bangladesh)	66	23 (34.8%)	6 (9.1%)	37 (56.1%)	0 (0.0%)	25	10 (40.0%)
nld (Netherlands)	410	21 (5.1%)	103 (25.1%)	284 (69.3%)	2 (0.5%)	91	14 (15.4%)

https://spoofer.caida.org/country_stats.php

Source Address Validation

ACL

- 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)

- 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’
- Uses less CPU and RAM, compared to implementing access-lists
- Most preferred mechanism for anti-spoofing technique

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

Unicast Reverse Path Forwarding (uRPF)

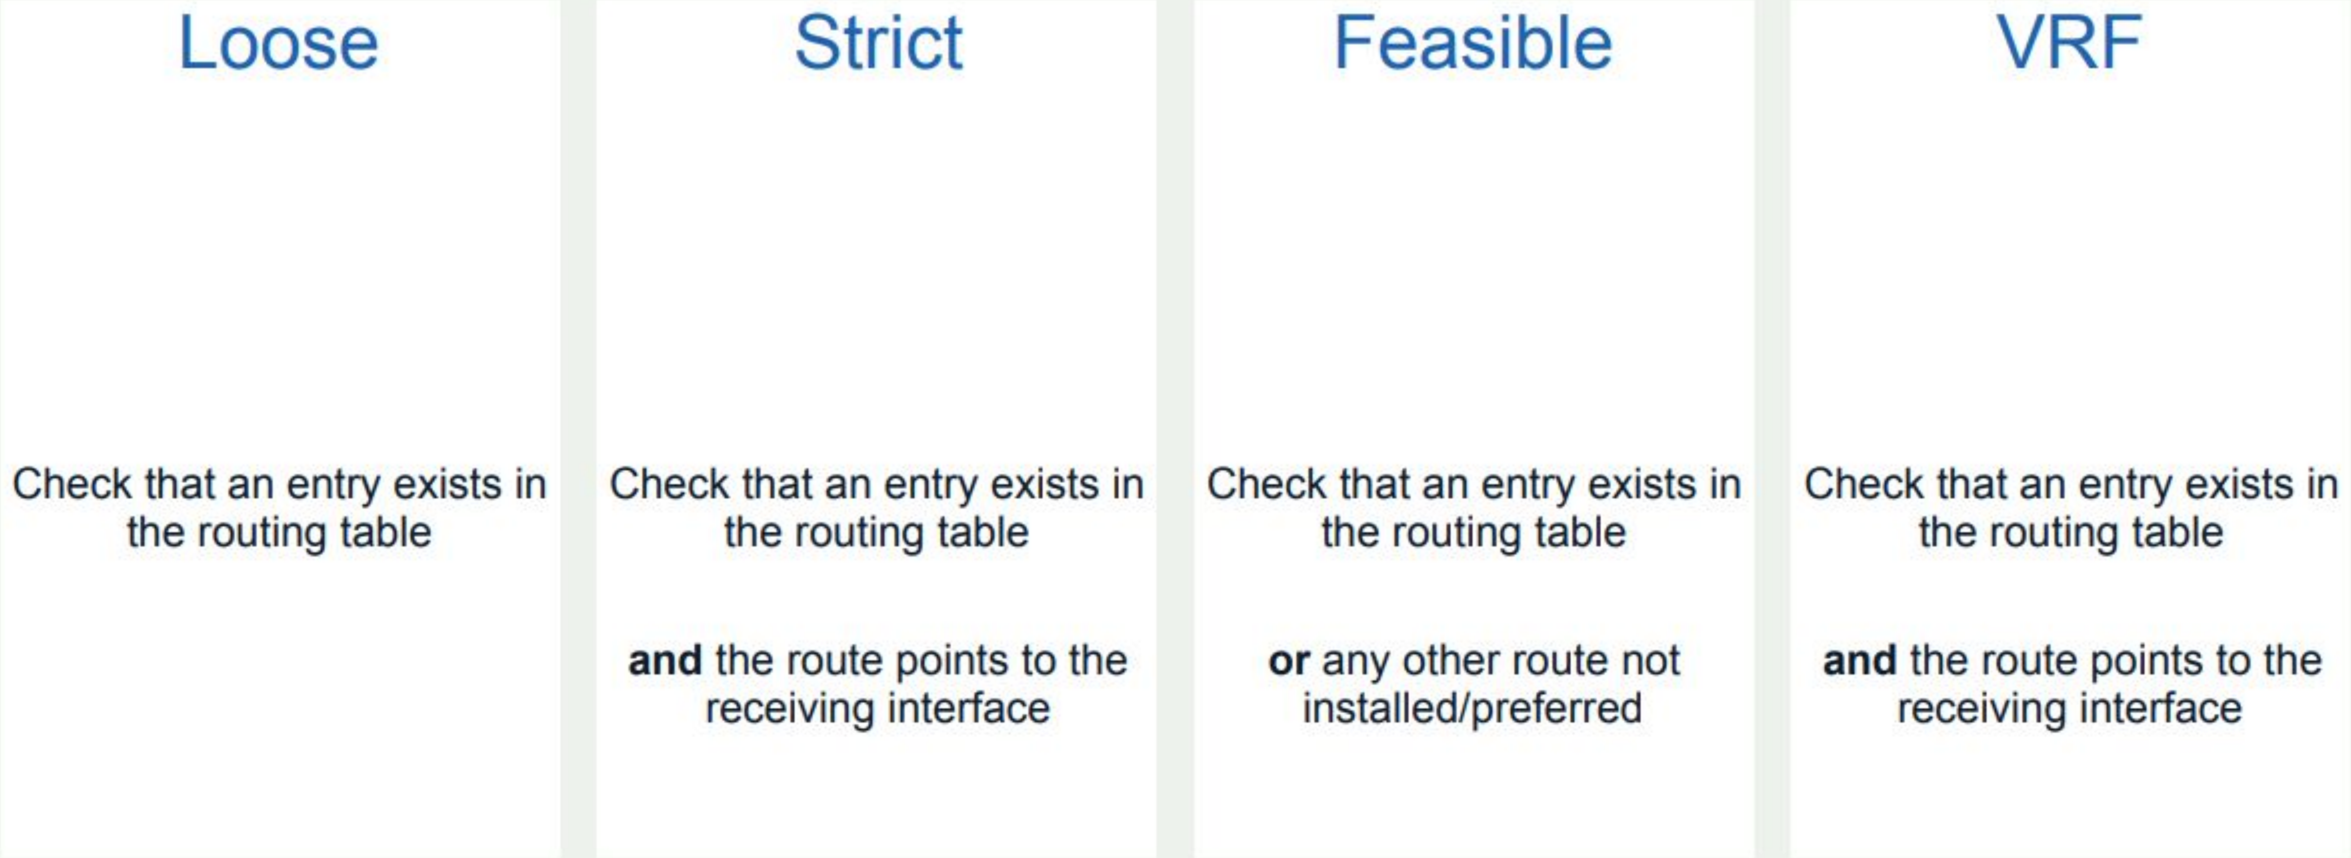

Loose	Strict	Feasible	VRF
Check that an entry exists in the routing table	Check that an entry exists in the routing table	Check that an entry exists in the routing table	Check that an entry exists in the routing table
	and the route points to the receiving interface	**or** any other route not installed/preferred	**and** the route points to the receiving interface

uRPF- Source Address Validation

Configuration in operator’s router

Cisco

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

Juniper

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

Implement uRPF on all single-homed customer facing interfaces

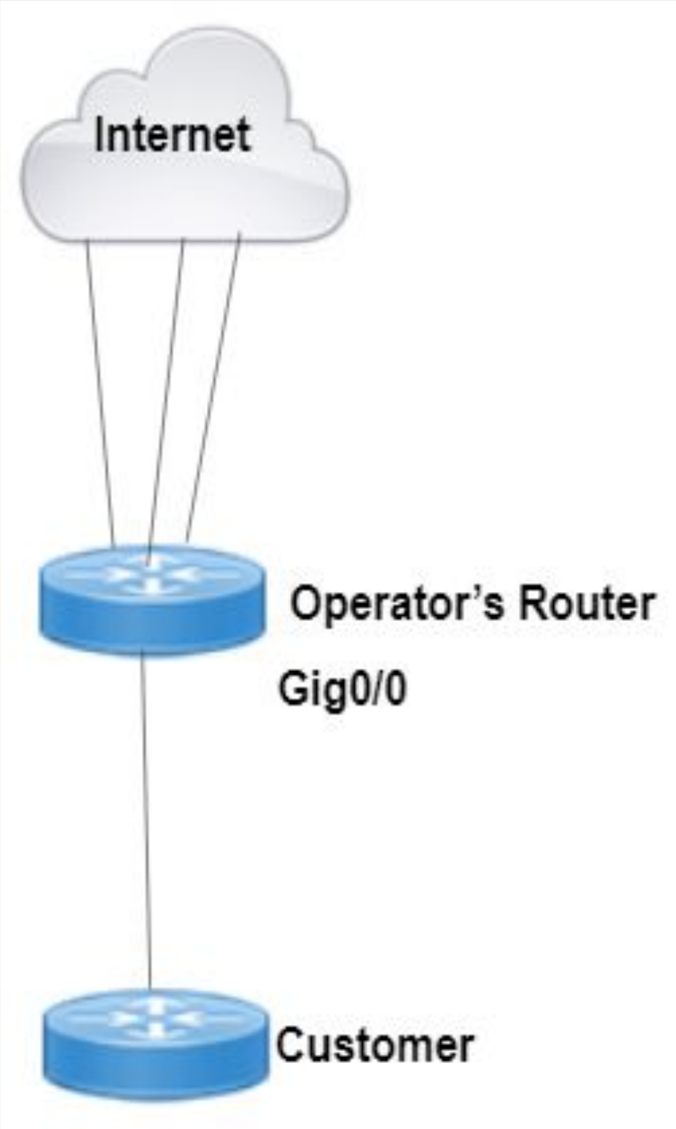

ACL - Source Address Validation

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 Provider Edge / Customer Edge (PE/CE) boundary
- Not a recommended solution, it does not provide complete protection against DoS attacks.

ACL example - Cisco

Configuration in operator’s router

```
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
!
```

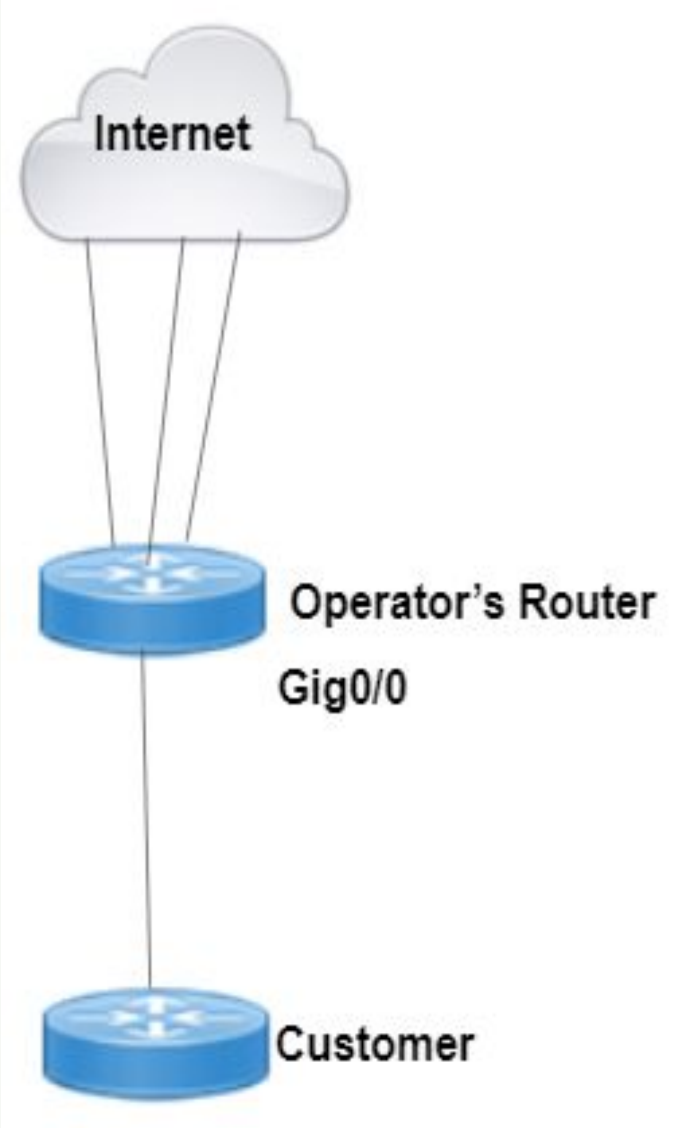

ACL example - Juniper

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;
}
```

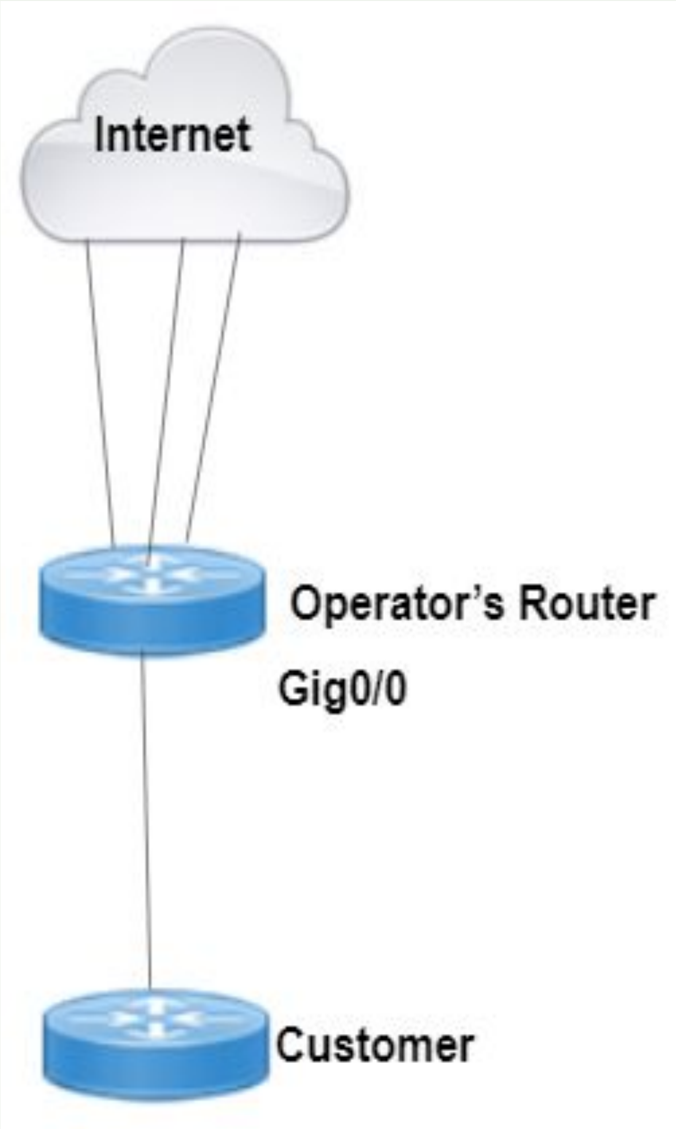

Recommendation

- Test your configuration
  - CAIDA Spoofer Client Software

    https://www.caida.org/projects/spoofer/#download-client-software

- Obtaining a peering session
  - Team Cymru https://www.team-cymru.com/bogon-reference.html
  - Remote Triggered Black Hole Filtering with uRPF

    https://tools.ietf.org/html/rfc5635

Action 3: Coordination

Facilitating global operational communication and coordination between network operators

Coordination

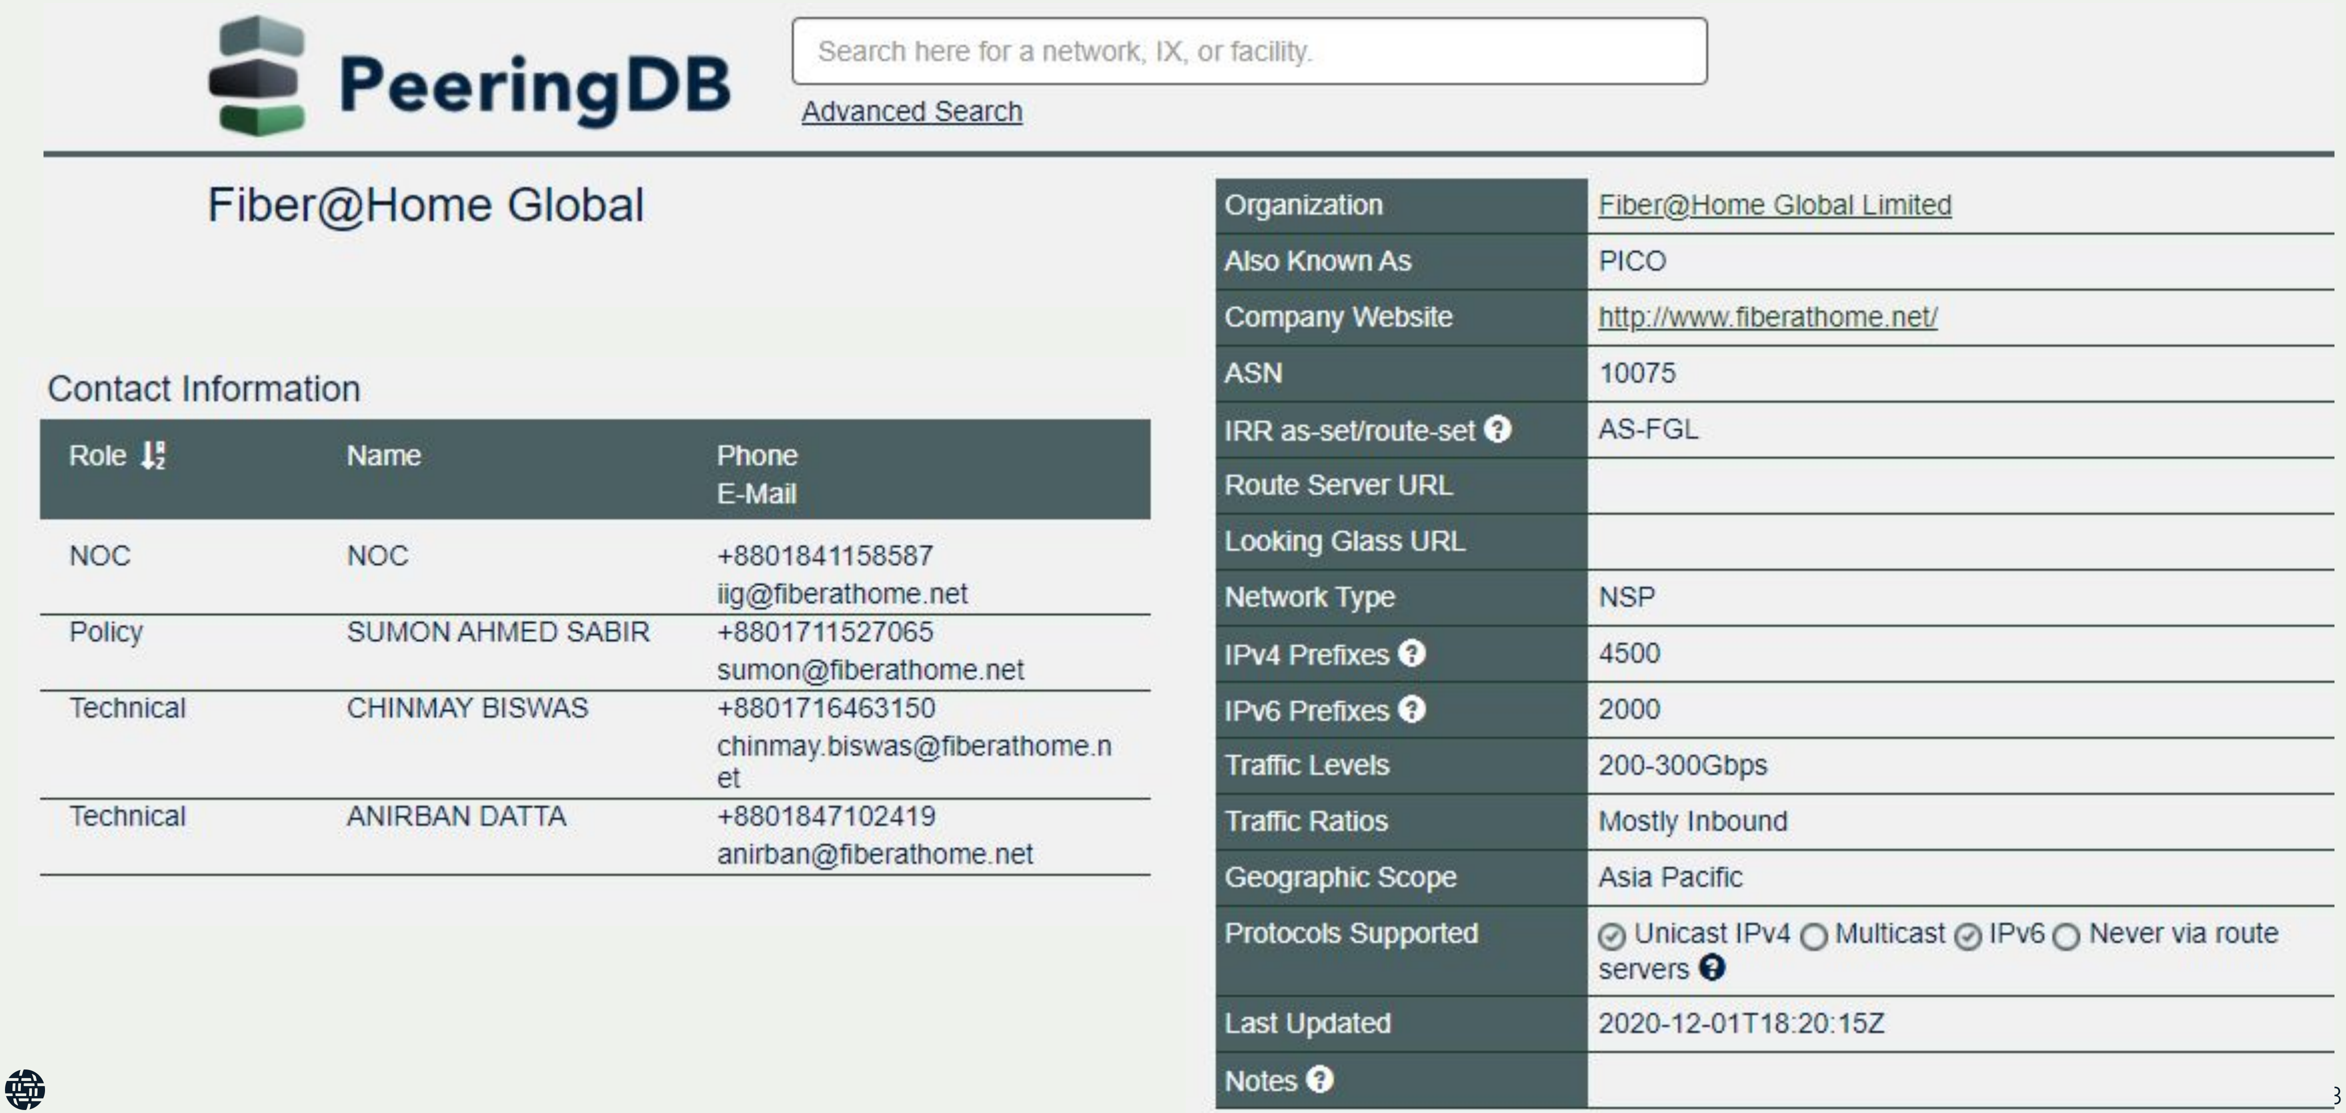

Fiber@Home Global

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 Regional Internet Registries (RIRs): AFRINIC, APNIC, RIPE NCC, LACNIC, ARIN

whois -h whois.apnic.net AS10075

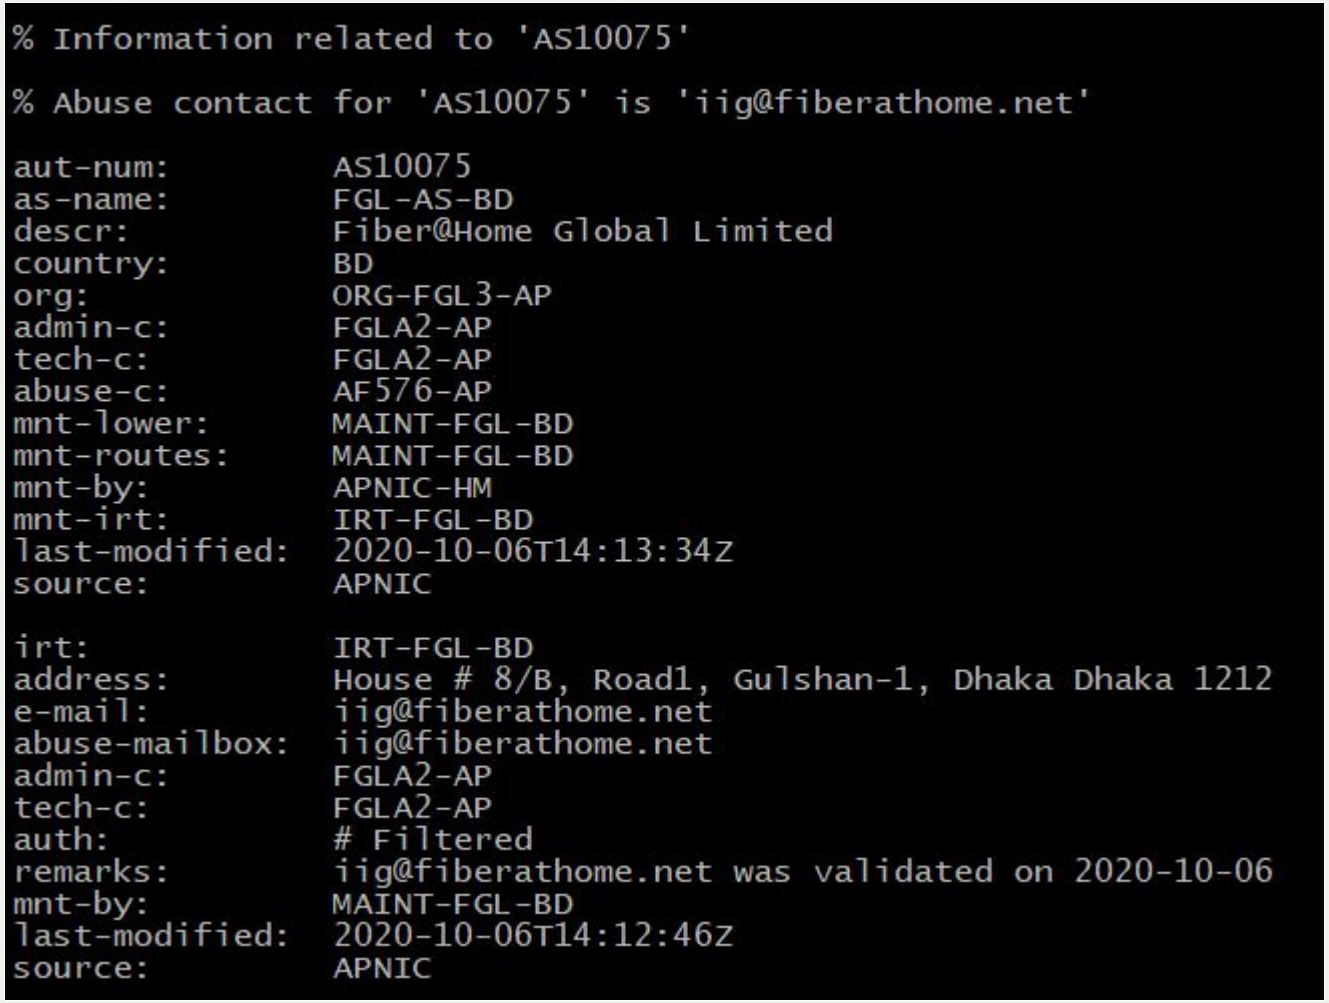

```
% Information related to 'AS10075'

% Abuse contact for 'AS10075' is 'iig@fiberathome.net'

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 2020-10-06
mnt-by:         MAINT-FGL-BD
last-modified:  2020-10-06T14:12:46Z
source:         APNIC
```

Coordination

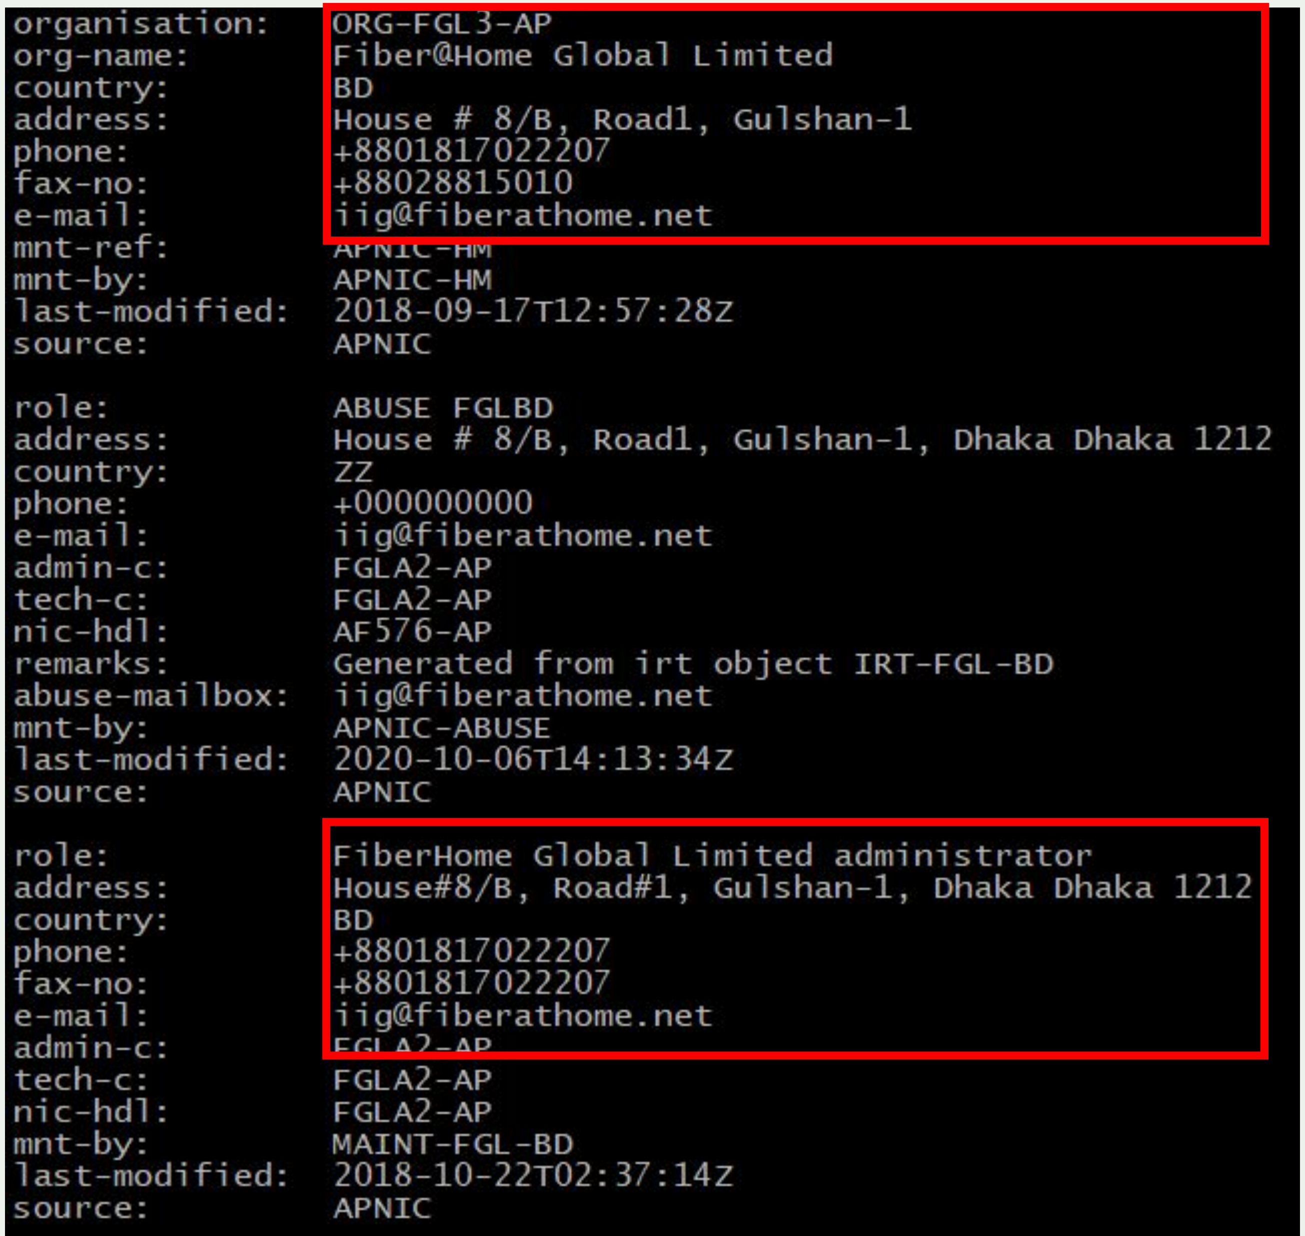

```
organisation:   ORG-FGL3-AP
org-name:       Fiber@Home Global Limited
country:        BD
address:        House # 8/B, Road1, Gulshan-1
phone:          +8801817022207
fax-no:         +88028815010
e-mail:         iig@fiberathome.net
mnt-ref:        APNIC-HM
mnt-by:         APNIC-HM
last-modified:  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
last-modified:  2018-10-22T02:37:14Z
source:         APNIC
```

Action 4: Global Validation

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 your ASN is covered.
- Only the origin ASN is verified, not the full path.
- All Regional Internet Registries (RIR) offers a hosted Resource Certification service.

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)

ROA (Route Origin Authorization)

- 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	103.229.82.0/23
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

RPKI can

- prevent route hijacks/mis-origination/misconfiguration

RPKI Implementation

- Origin validation uses X.509 certificates with extensions specified in RFC 3779
- RPKI Cache server (RPKI Validator) synchronizes its local database with TAs
- RPKI Validator exports a simplified version of ROAs as Route Validation (RV) records
- An RV record is a (prefix, maximum length, origin AS) triple
- Router interacts with cache server to query the RV status for a given route
- Router implements BGP policy based on validation status indicated in the RV record for the route

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)

What is in a ROA ?

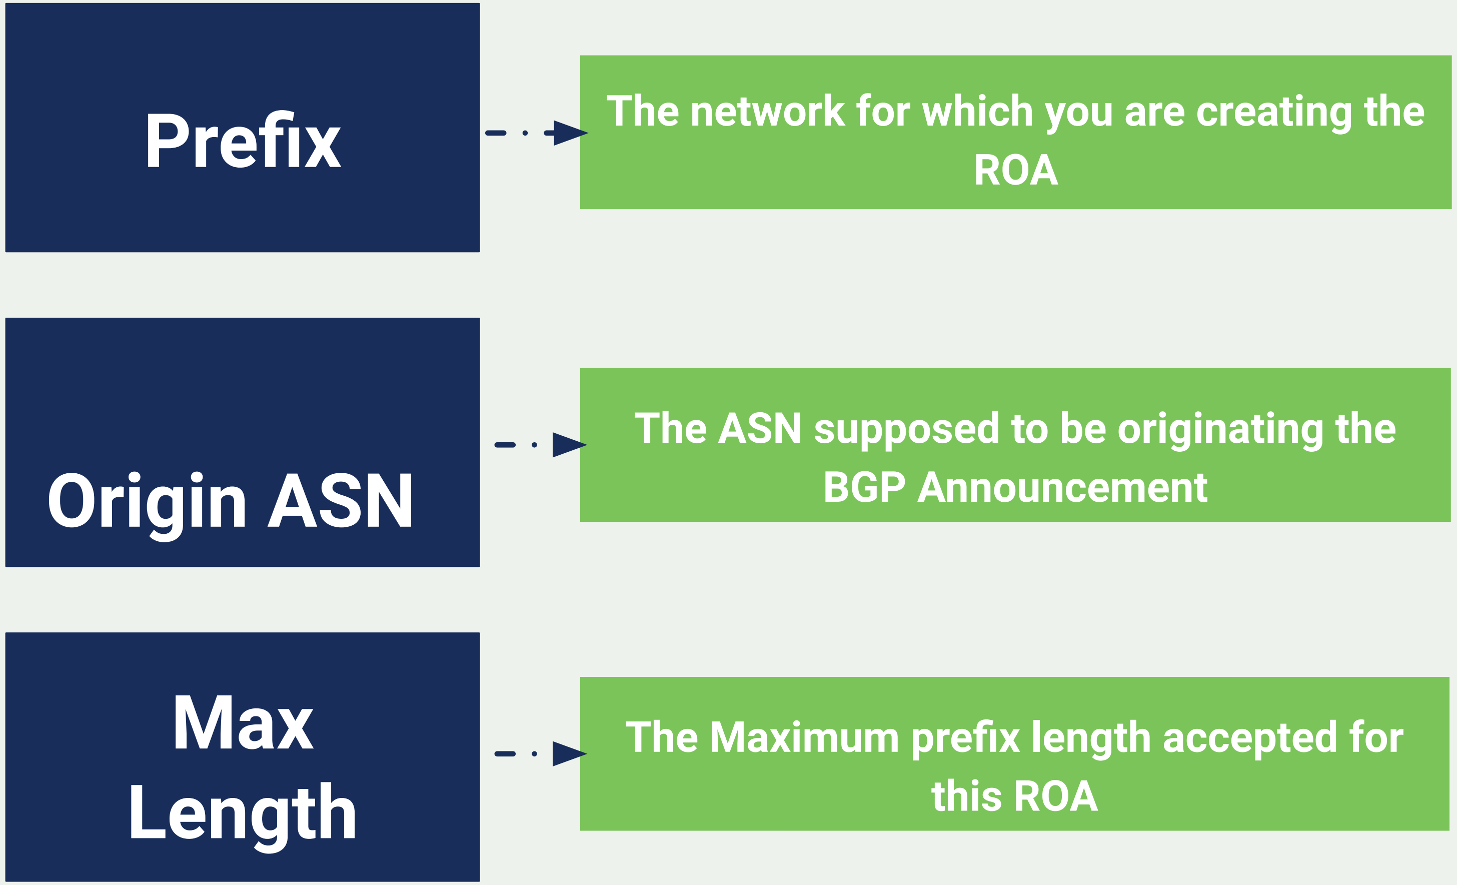

Prefix

The network for which you are creating the ROA

Origin ASN

The ASN supposed to be originating the BGP Announcement

Max Length

The Maximum prefix length accepted for this ROA

RPKI-RTR

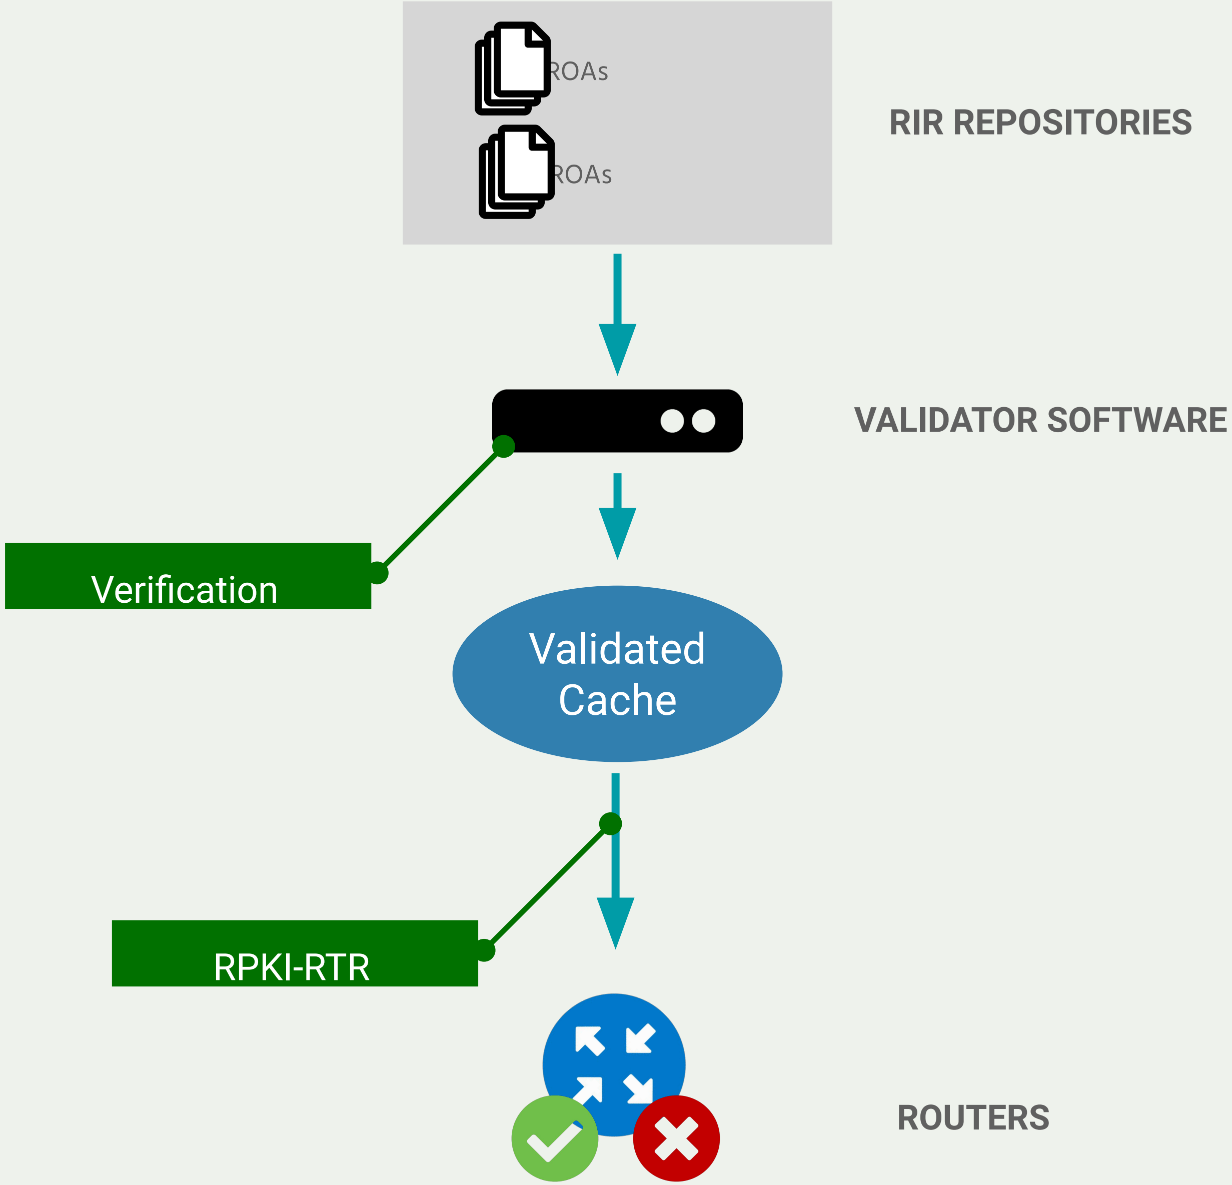
ROAs
ROAs
RIR REPOSITORIES
VALIDATOR SOFTWARE
Verification
Validated Cache
RPKI-RTR
ROUTERS

Relying Party

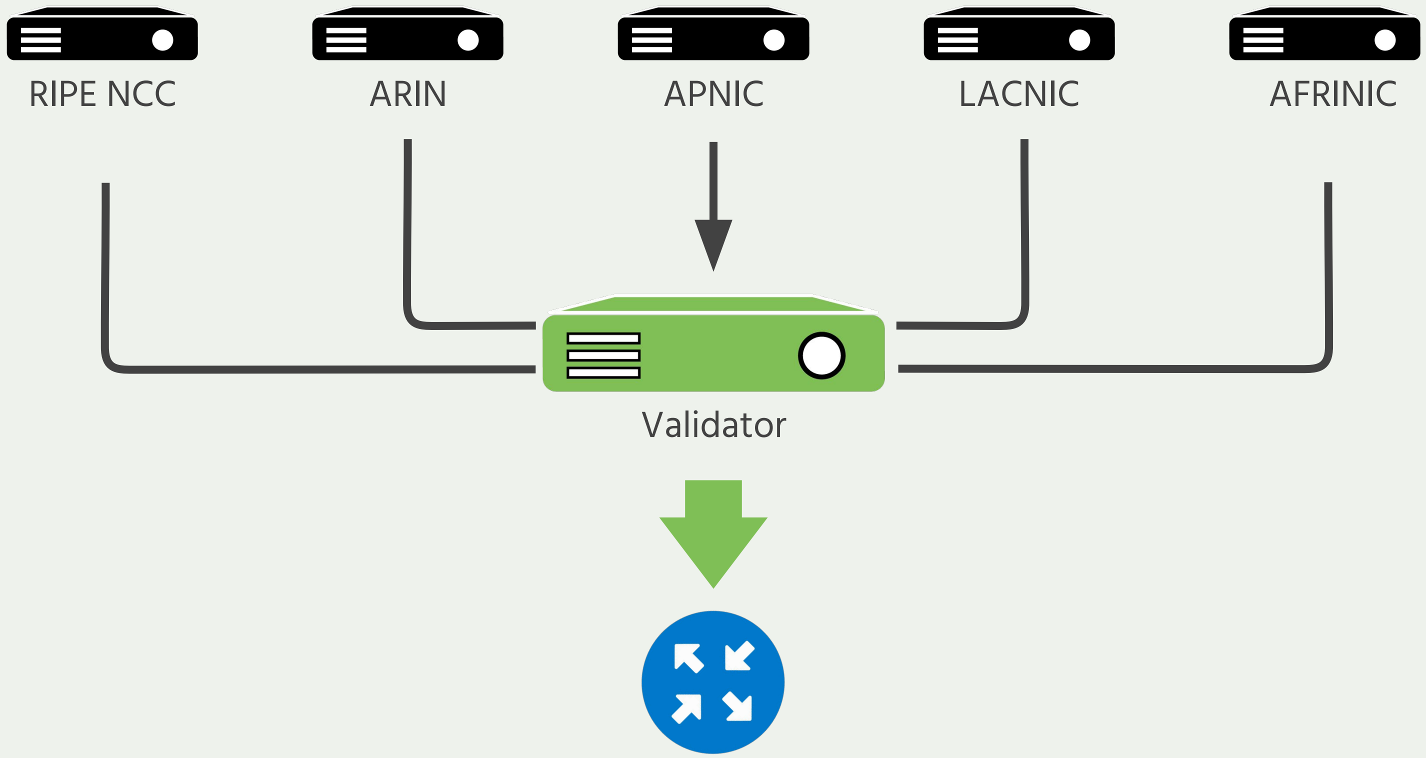

Validator Software

- 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

Validator - Status Check

Routinator status can be checked by having Routinator print a validated ROA payload (VRP) list

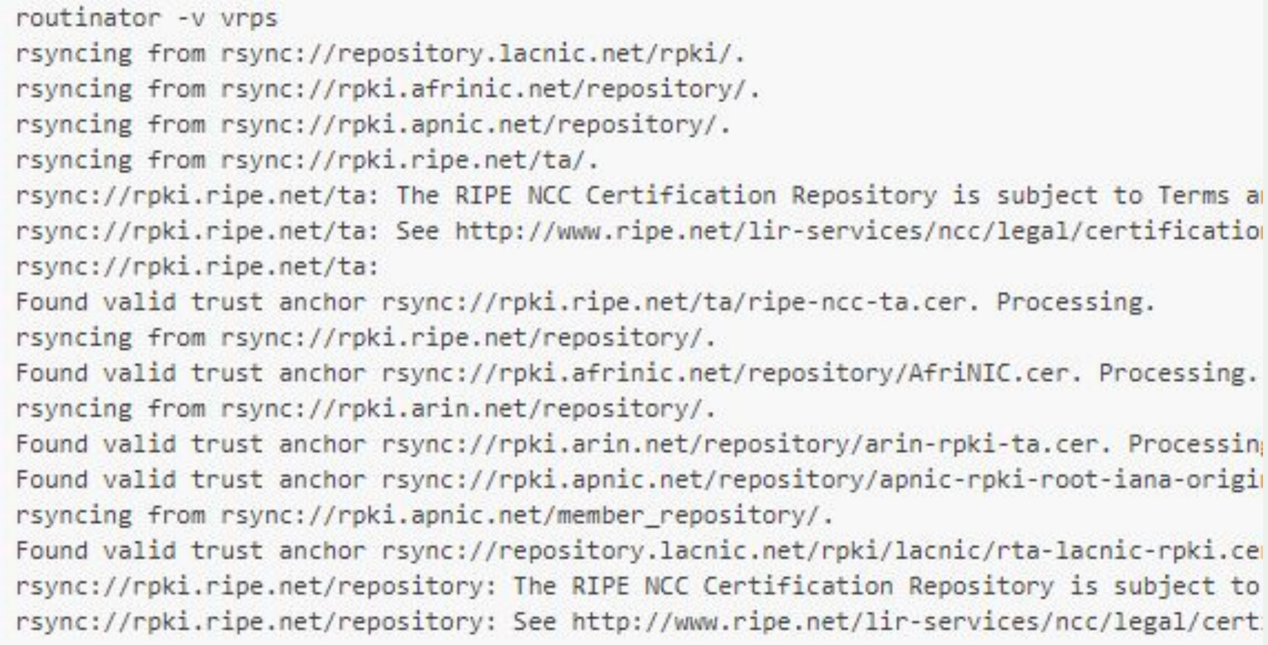

```
routinator -v vrps
rsyncing from rsync://repository.lacnic.net/rpki/.
rsyncing from rsync://rpki.afrinic.net/repository/.
rsyncing from rsync://rpki.apnic.net/repository/.
rsyncing from rsync://rpki.ripe.net/ta/.
rsync://rpki.ripe.net/ta: The RIPE NCC Certification Repository is subject to Terms a
rsync://rpki.ripe.net/ta: See http://www.ripe.net/lir-services/ncc/legal/certificatio
rsync://rpki.ripe.net/ta:
Found valid trust anchor rsync://rpki.ripe.net/ta/ripe-ncc-ta.cer. Processing.
rsyncing from rsync://rpki.ripe.net/repository/.
Found valid trust anchor rsync://rpki.afrinic.net/repository/AfriNIC.cer. Processing.
rsyncing from rsync://rpki.arin.net/repository/.
Found valid trust anchor rsync://rpki.arin.net/repository/arin-rpki-ta.cer. Processin
Found valid trust anchor rsync://rpki.apnic.net/repository/apnic-rpki-root-iana-origi
rsyncing from rsync://rpki.apnic.net/member_repository/.
Found valid trust anchor rsync://repository.lacnic.net/rpki/lacnic/rta-lacnic-rpki.ce
rsync://rpki.ripe.net/repository: The RIPE NCC Certification Repository is subject to
rsync://rpki.ripe.net/repository: See http://www.ripe.net/lir-services/ncc/legal/cert
```

Also server process can be checked

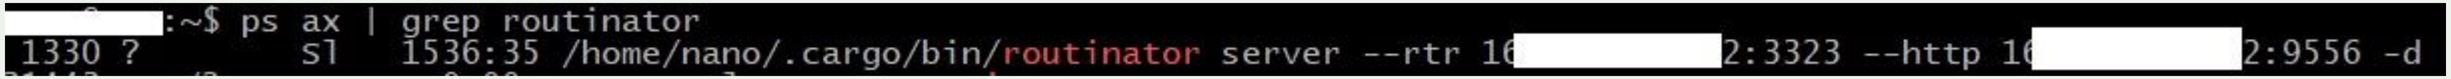

```
[redacted]:~$ ps ax | grep routinator
 1330 ?        Sl   1536:35 /home/nano/.cargo/bin/routinator server --rtr 16[redacted]2:3323 --http 16[redacted]2:9556 -d
```

Origin Validation Configuration (Cisco)

```
(config)# conf t
(config)# router bgp 64500
(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
```

Origin Validation Configuration (Juniper)

```
routing-options {
    autonomous-system 64500;
    validation {
                           group rpki-validator {
            session 100.64.1.1 {
                refresh-time 120;
                hold-time 180;
                port 8282;
                local-address 100.64.1.2;
}}}}
```

Prefix Validation Status

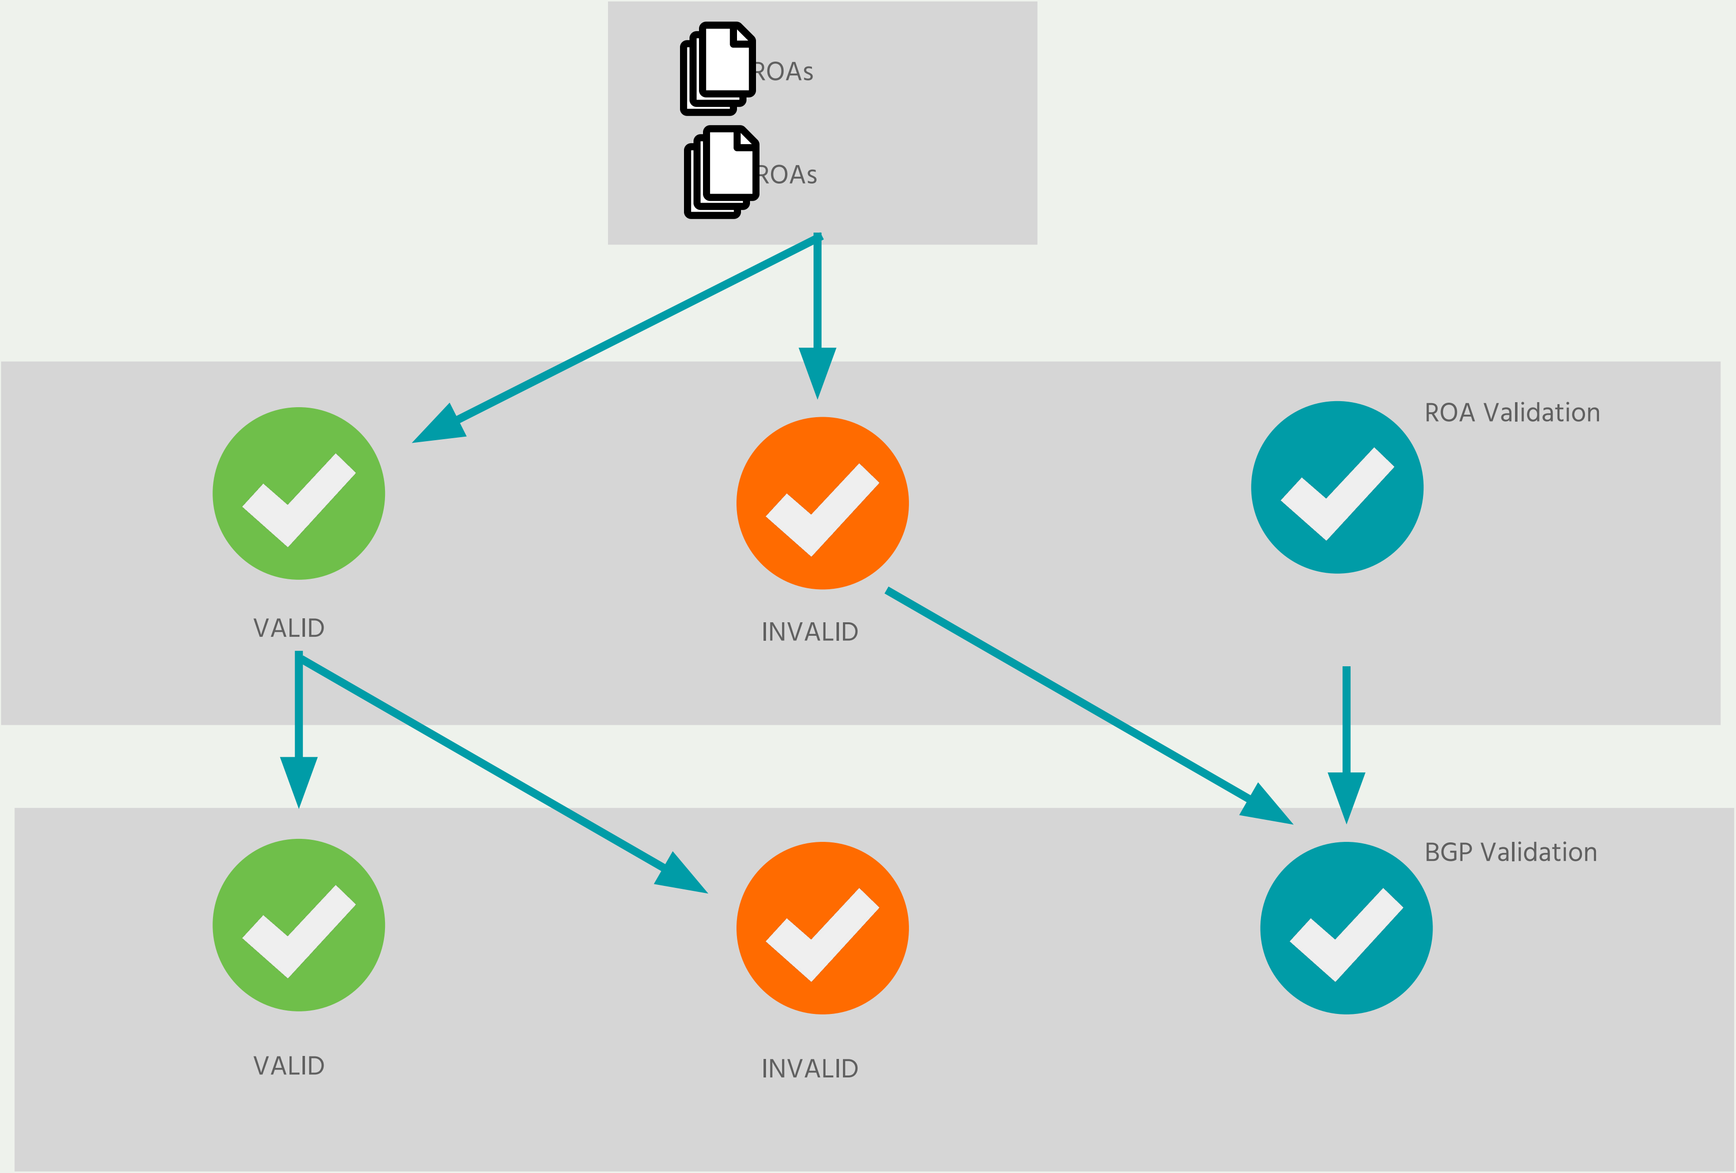

ROAs

ROAs

ROA Validation

VALID

INVALID

BGP Validation

VALID

INVALID

Origin Validation on CISCO (Low pref to invalids)

```
(config)# route-map RPKI-VALIDATION permit 10
(route-map)# match rpki valid
(route-map)# set local-preference 100
!
(route-map)# route-map RPKI-VALIDATION permit 20
(route-map)# match rpki not-found
(route-map)# set local-preference 80
!
(route-map)# route-map RPKI-VALIDATION permit 30
(route-map)# match rpki invalid
(route-map)# set local-preference 60
```

High preference valid ROAs

Medium preference for routes ROAs not found

Very low preference to invalids.

Most operators started dropping invalids

Origin Validation on CISCO (Drop invalids)

```
(config-router)# route-map RPKI-VALIDATION permit 10
(route-map)# match rpki valid
(route-map)# set local-preference 100
!
(route-map)# route-map RPKI-VALIDATION permit 20
(route-map)# match rpki not-found
(route-map)# set local-preference 80
!
(route-map)# route-map RPKI-VALIDATION deny 30
(route-map)# match rpki invalid
```

High preference valid ROAs

Medium preference for routes ROAs not found

Invalids are dropped

Origin Validation Configuration (Cisco)

Applying the route-map to the BGP sessions

```
(config)# router bgp 64511
(config)# address-family ipv4
(config)# neighbor 192.168.1.254 route-map RPKI-VALIDATION in
!
(config)# address-family ipv6
(config)# neighbor 2001:db8:12::2 route-map RPKI-VALIDATION in
```

Origin Validation Configuration (Juniper)

```
policy-statement send-direct {
        from protocol direct;
        then accept;}
policy-statement RPKI-VALIDATION {
        term valid {
                from {
                        protocol bgp;
                        validation-database valid; }
                then {
                        local-preference 100;
                        validation-state valid;
                        community add origin-validation-state-valid;
                        accept;
                        }}
```

High preference valid ROAs

Origin Validation Configuration (Juniper)

```
term unknown {
        from protocol bgp;
        then {
        local-preference 80;
        validation-state unknown;
        community add origin-validation-state-unknown;
        accept;
}}
```

Medium preference for routes ROAs not found

Origin Validation Configuration (Juniper)

```
term invalid {
        from {
                protocol bgp;
                validation-database invalid;}
        then {
                local-preference 60;
                validation-state invalid;
                community add origin-validation-state-invalid;
                accept;
}}
```

Very low preference to invalids.

Origin Validation Configuration (Juniper)

```
term invalid {
        from {
                protocol bgp;
                validation-database invalid;}
        then {
                validation-state invalid;
                reject;
}}
```

Invalids are dropped

RPKI Verification (Server)

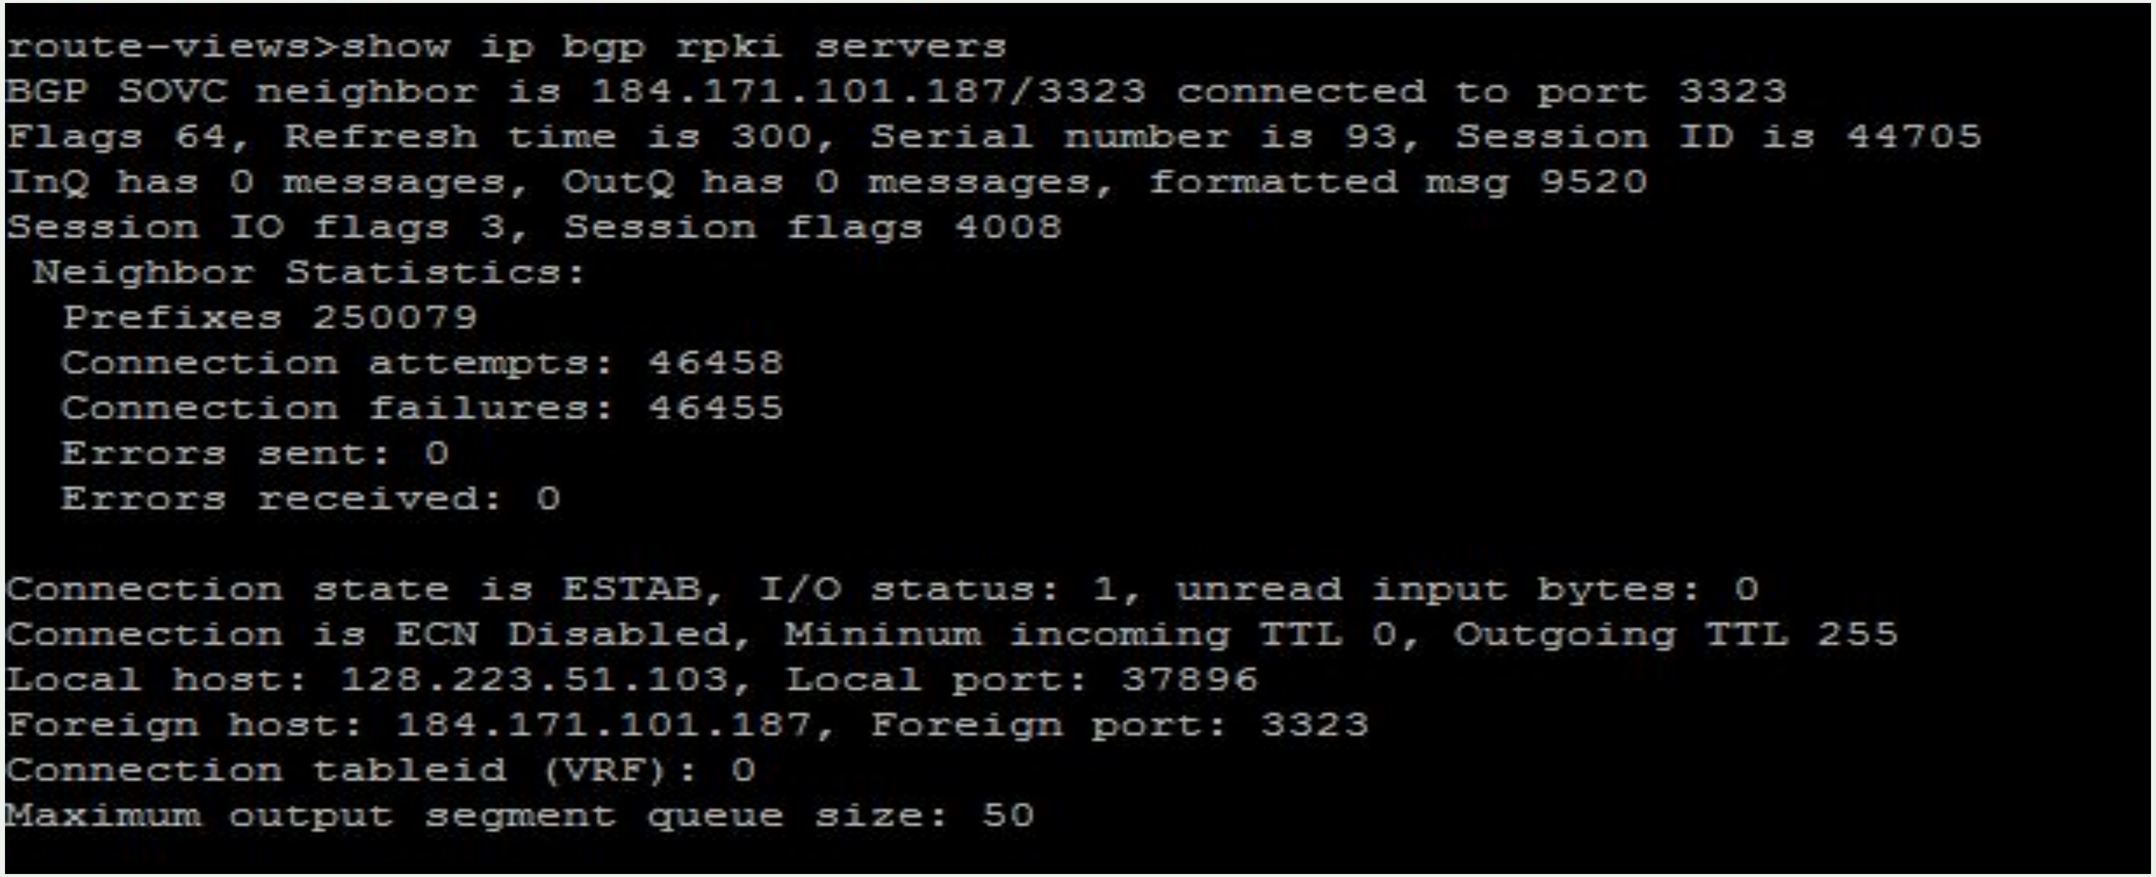

```
route-views>show ip bgp rpki servers
BGP SOVC neighbor is 184.171.101.187/3323 connected to port 3323
Flags 64, Refresh time is 300, Serial number is 93, Session ID is 44705
InQ has 0 messages, OutQ has 0 messages, formatted msg 9520
Session IO flags 3, Session flags 4008
 Neighbor Statistics:
  Prefixes 250079
  Connection attempts: 46458
  Connection failures: 46455
  Errors sent: 0
  Errors received: 0

Connection state is ESTAB, I/O status: 1, unread input bytes: 0
Connection is ECN Disabled, Mininum incoming TTL 0, Outgoing TTL 255
Local host: 128.223.51.103, Local port: 37896
Foreign host: 184.171.101.187, Foreign port: 3323
Connection tableid (VRF): 0
Maximum output segment queue size: 50
```

Courtesy: Routeviews.org

RPKI Verification (Server)

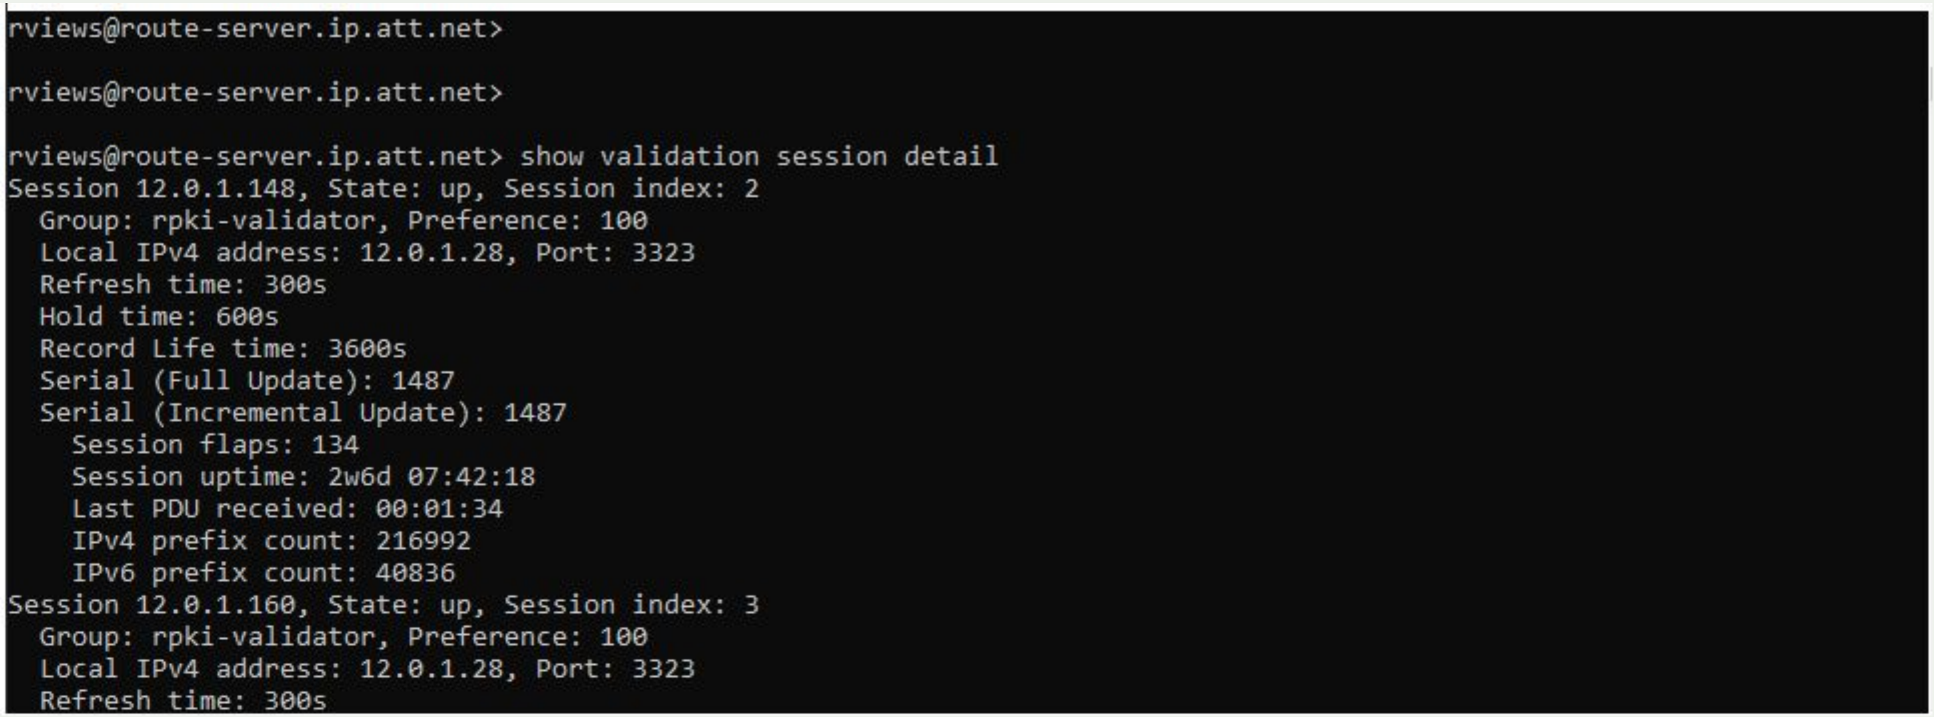

```
rviews@route-server.ip.att.net>

rviews@route-server.ip.att.net>

rviews@route-server.ip.att.net> show validation session detail
Session 12.0.1.148, State: up, Session index: 2
  Group: rpki-validator, Preference: 100
  Local IPv4 address: 12.0.1.28, Port: 3323
  Refresh time: 300s
  Hold time: 600s
  Record Life time: 3600s
  Serial (Full Update): 1487
  Serial (Incremental Update): 1487
    Session flaps: 134
    Session uptime: 2w6d 07:42:18
    Last PDU received: 00:01:34
    IPv4 prefix count: 216992
    IPv6 prefix count: 40836
Session 12.0.1.160, State: up, Session index: 3
  Group: rpki-validator, Preference: 100
  Local IPv4 address: 12.0.1.28, Port: 3323
  Refresh time: 300s
```

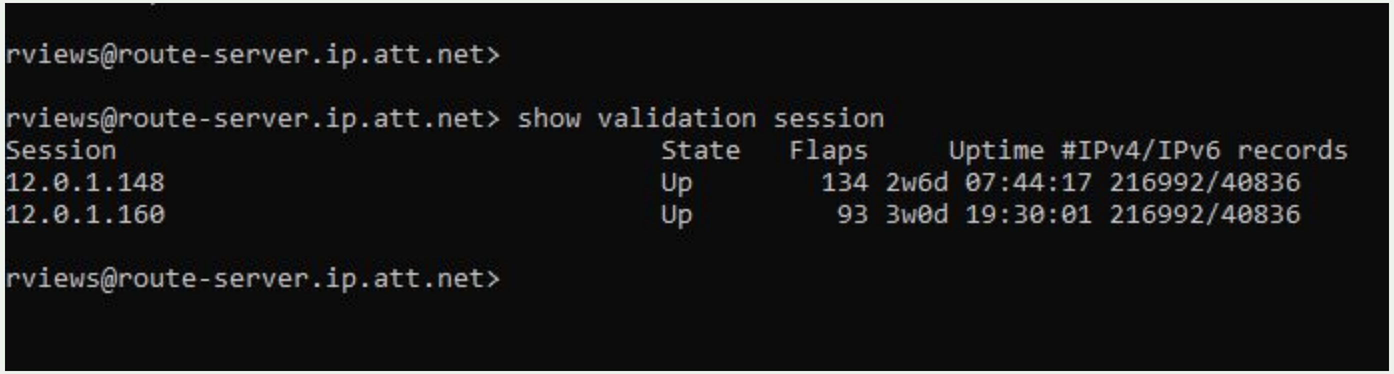

```
rviews@route-server.ip.att.net>

rviews@route-server.ip.att.net> show validation session
Session                                  State   Flaps     Uptime #IPv4/IPv6 records
12.0.1.148                               Up        134 2w6d 07:44:17 216992/40836
12.0.1.160                               Up         93 3w0d 19:30:01 216992/40836

rviews@route-server.ip.att.net>
```

Courtesy: routeserver@ATT AS7018

RPKI Table (IPv4)

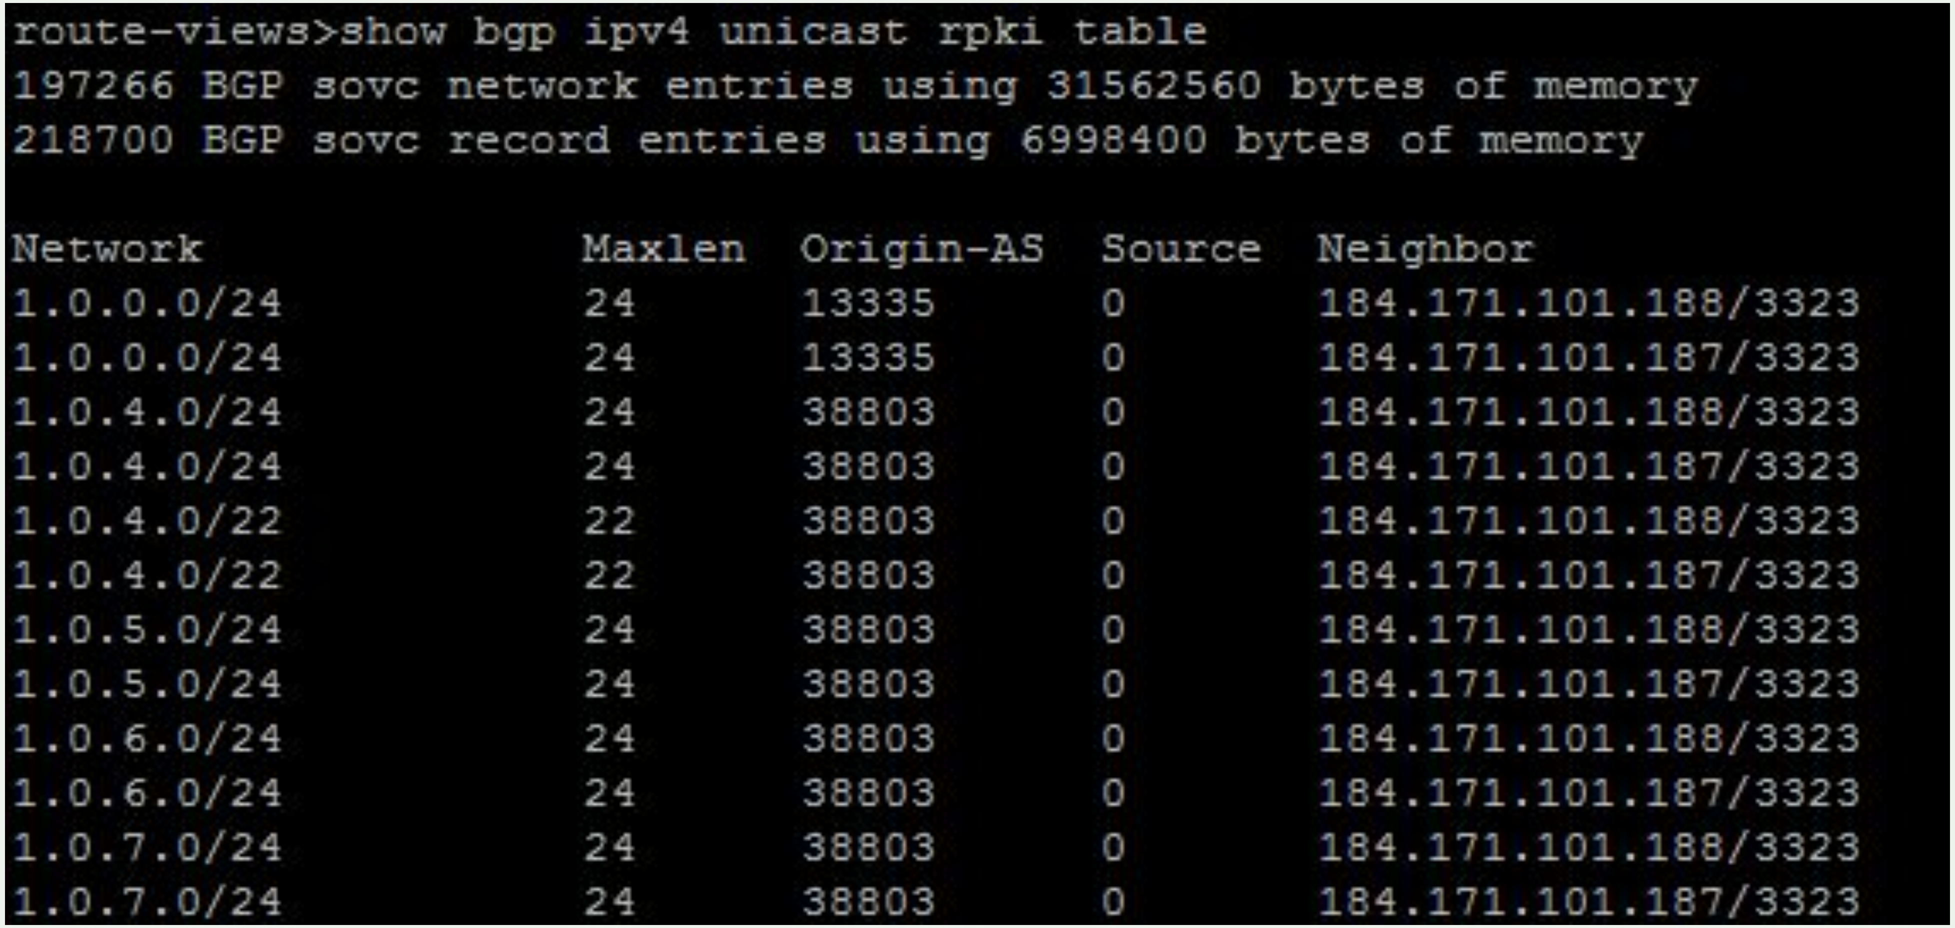

```
route-views>show bgp ipv4 unicast rpki table
197266 BGP sovc network entries using 31562560 bytes of memory
218700 BGP sovc record entries using 6998400 bytes of memory

Network                Maxlen  Origin-AS  Source  Neighbor
1.0.0.0/24             24      13335      0       184.171.101.188/3323
1.0.0.0/24             24      13335      0       184.171.101.187/3323
1.0.4.0/24             24      38803      0       184.171.101.188/3323
1.0.4.0/24             24      38803      0       184.171.101.187/3323
1.0.4.0/22             22      38803      0       184.171.101.188/3323
1.0.4.0/22             22      38803      0       184.171.101.187/3323
1.0.5.0/24             24      38803      0       184.171.101.188/3323
1.0.5.0/24             24      38803      0       184.171.101.187/3323
1.0.6.0/24             24      38803      0       184.171.101.188/3323
1.0.6.0/24             24      38803      0       184.171.101.187/3323
1.0.7.0/24             24      38803      0       184.171.101.188/3323
1.0.7.0/24             24      38803      0       184.171.101.187/3323
```

Courtesy: Routeviews.org

RPKI Table (IPv6)

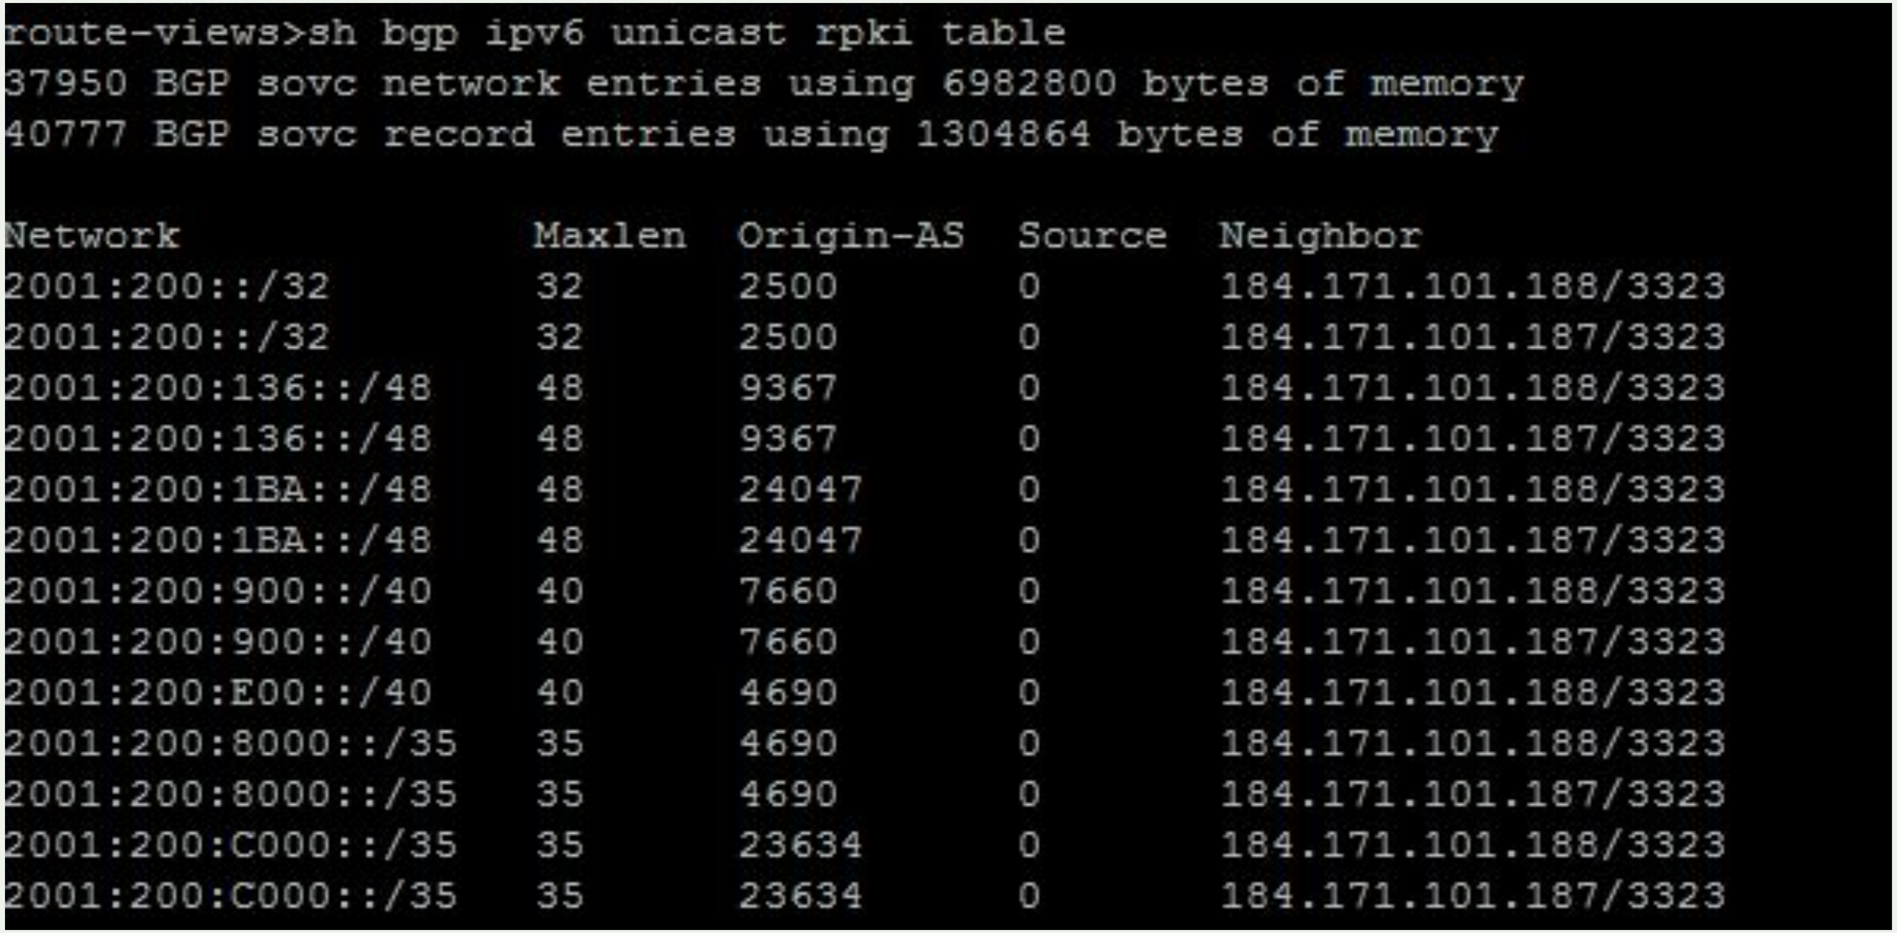

```
route-views>sh bgp ipv6 unicast rpki table
37950 BGP sovc network entries using 6982800 bytes of memory
40777 BGP sovc record entries using 1304864 bytes of memory

Network                Maxlen   Origin-AS   Source   Neighbor
2001:200::/32          32       2500        0        184.171.101.188/3323
2001:200::/32          32       2500        0        184.171.101.187/3323
2001:200:136::/48      48       9367        0        184.171.101.188/3323
2001:200:136::/48      48       9367        0        184.171.101.187/3323
2001:200:1BA::/48      48       24047       0        184.171.101.188/3323
2001:200:1BA::/48      48       24047       0        184.171.101.187/3323
2001:200:900::/40      40       7660        0        184.171.101.188/3323
2001:200:900::/40      40       7660        0        184.171.101.187/3323
2001:200:E00::/40      40       4690        0        184.171.101.188/3323
2001:200:8000::/35     35       4690        0        184.171.101.188/3323
2001:200:8000::/35     35       4690        0        184.171.101.187/3323
2001:200:C000::/35     35       23634       0        184.171.101.188/3323
2001:200:C000::/35     35       23634       0        184.171.101.187/3323
```

Courtesy: Routeviews.org

RPKI State (Valid)

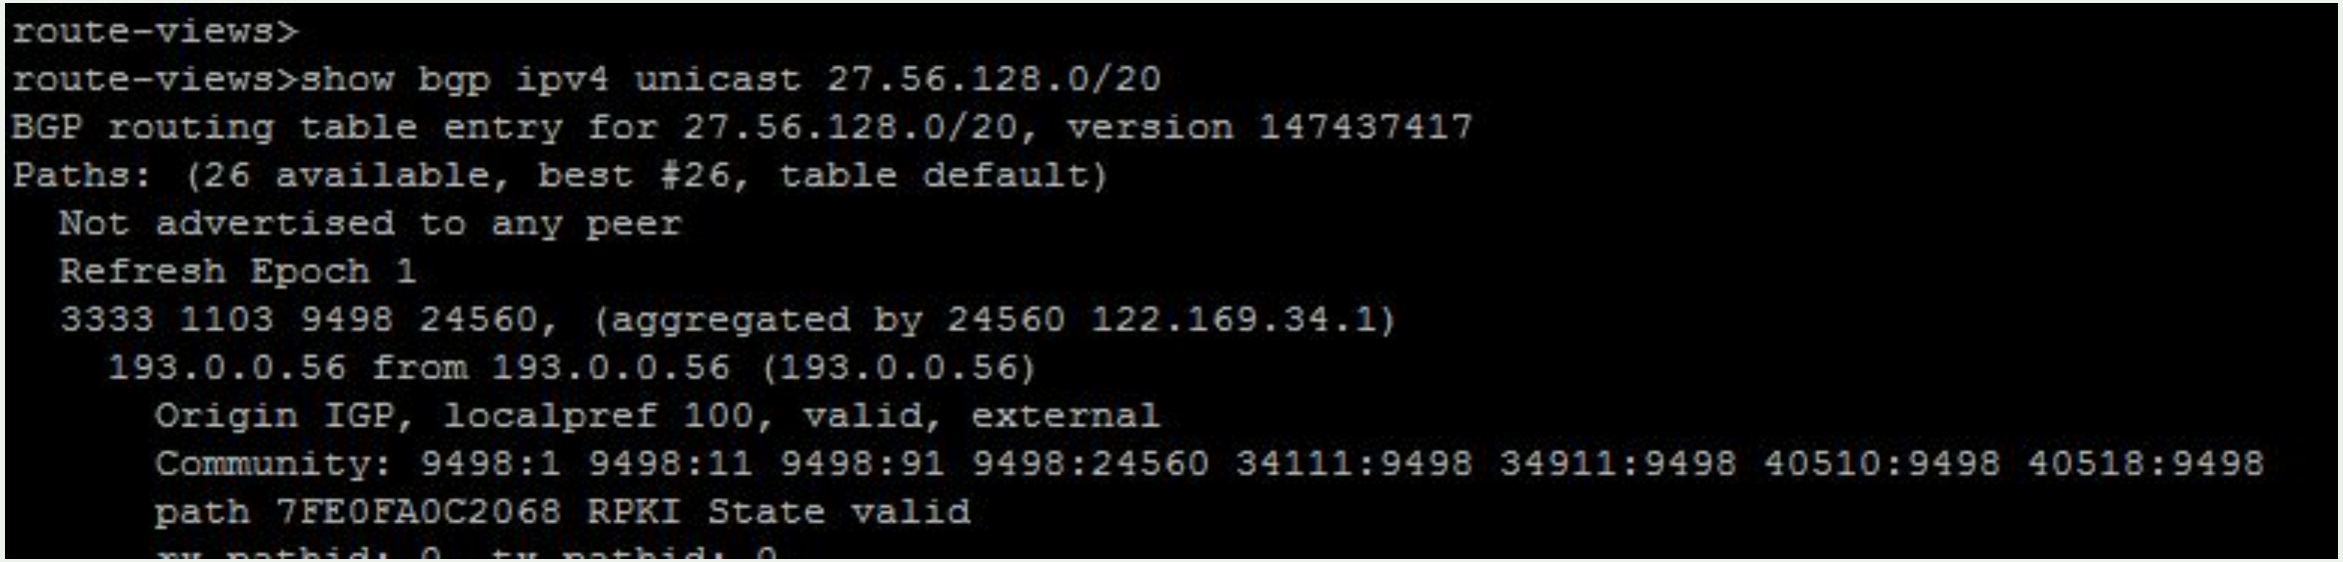

```
route-views>
route-views>show bgp ipv4 unicast 27.56.128.0/20
BGP routing table entry for 27.56.128.0/20, version 147437417
Paths: (26 available, best #26, table default)
  Not advertised to any peer
  Refresh Epoch 1
  3333 1103 9498 24560, (aggregated by 24560 122.169.34.1)
    193.0.0.56 from 193.0.0.56 (193.0.0.56)
      Origin IGP, localpref 100, valid, external
      Community: 9498:1 9498:11 9498:91 9498:24560 34111:9498 34911:9498 40510:9498 40518:9498
      path 7FE0FA0C2068 RPKI State valid
      rx pathid: 0, tx pathid: 0
```

Courtesy: Routeviews.org

RPKI State (Invalid)

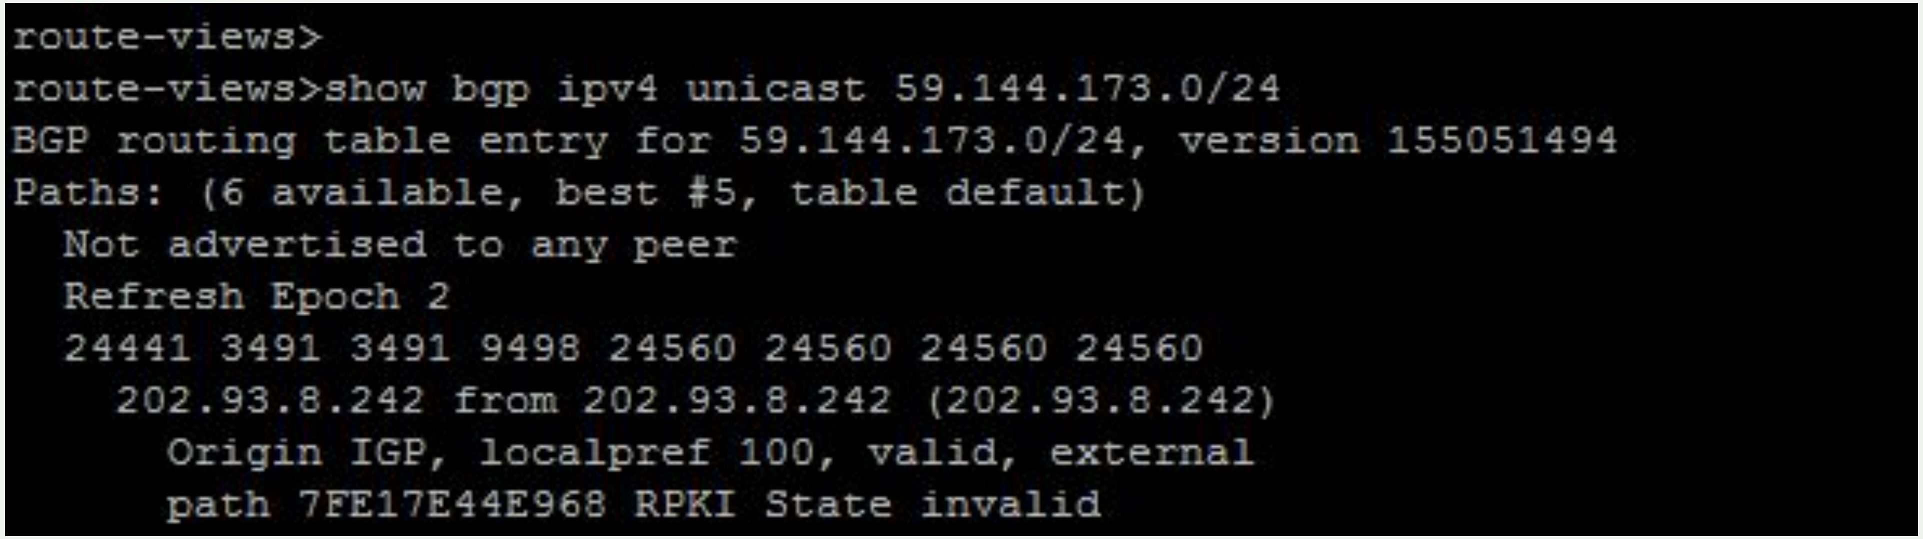

```
route-views>
route-views>show bgp ipv4 unicast 59.144.173.0/24
BGP routing table entry for 59.144.173.0/24, version 155051494
Paths: (6 available, best #5, table default)
  Not advertised to any peer
  Refresh Epoch 2
  24441 3491 3491 9498 24560 24560 24560 24560
    202.93.8.242 from 202.93.8.242 (202.93.8.242)
      Origin IGP, localpref 100, valid, external
      path 7FE17E44E968 RPKI State invalid
```

Courtesy: Routeviews.org

RPKI State (Not Found)

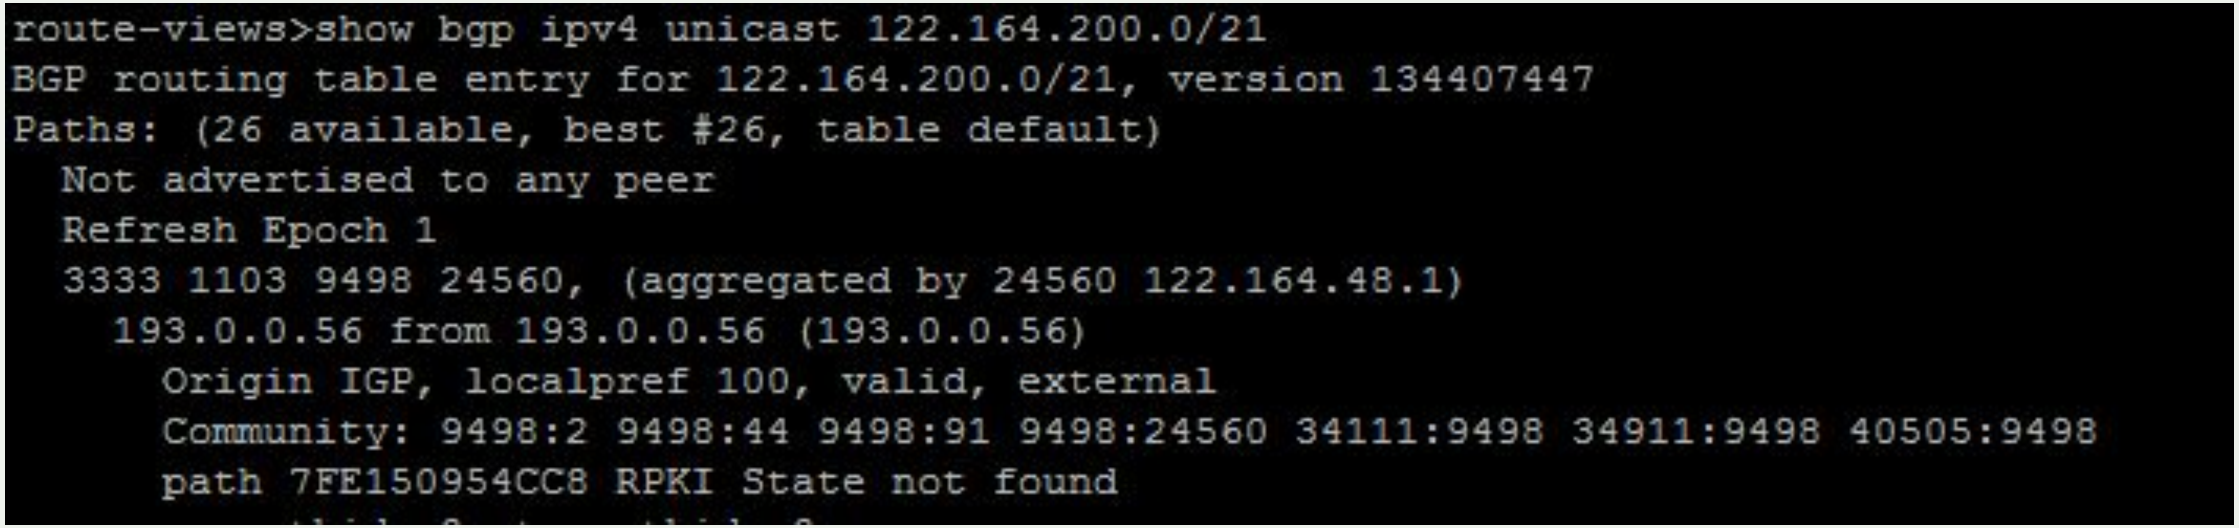

```
route-views>show bgp ipv4 unicast 122.164.200.0/21
BGP routing table entry for 122.164.200.0/21, version 134407447
Paths: (26 available, best #26, table default)
  Not advertised to any peer
  Refresh Epoch 1
  3333 1103 9498 24560, (aggregated by 24560 122.164.48.1)
    193.0.0.56 from 193.0.0.56 (193.0.0.56)
      Origin IGP, localpref 100, valid, external
      Community: 9498:2 9498:44 9498:91 9498:24560 34111:9498 34911:9498 40505:9498
      path 7FE150954CC8 RPKI State not found
```

Courtesy: Routeviews.org

RPKI-ROV Analysis: Global Analysis

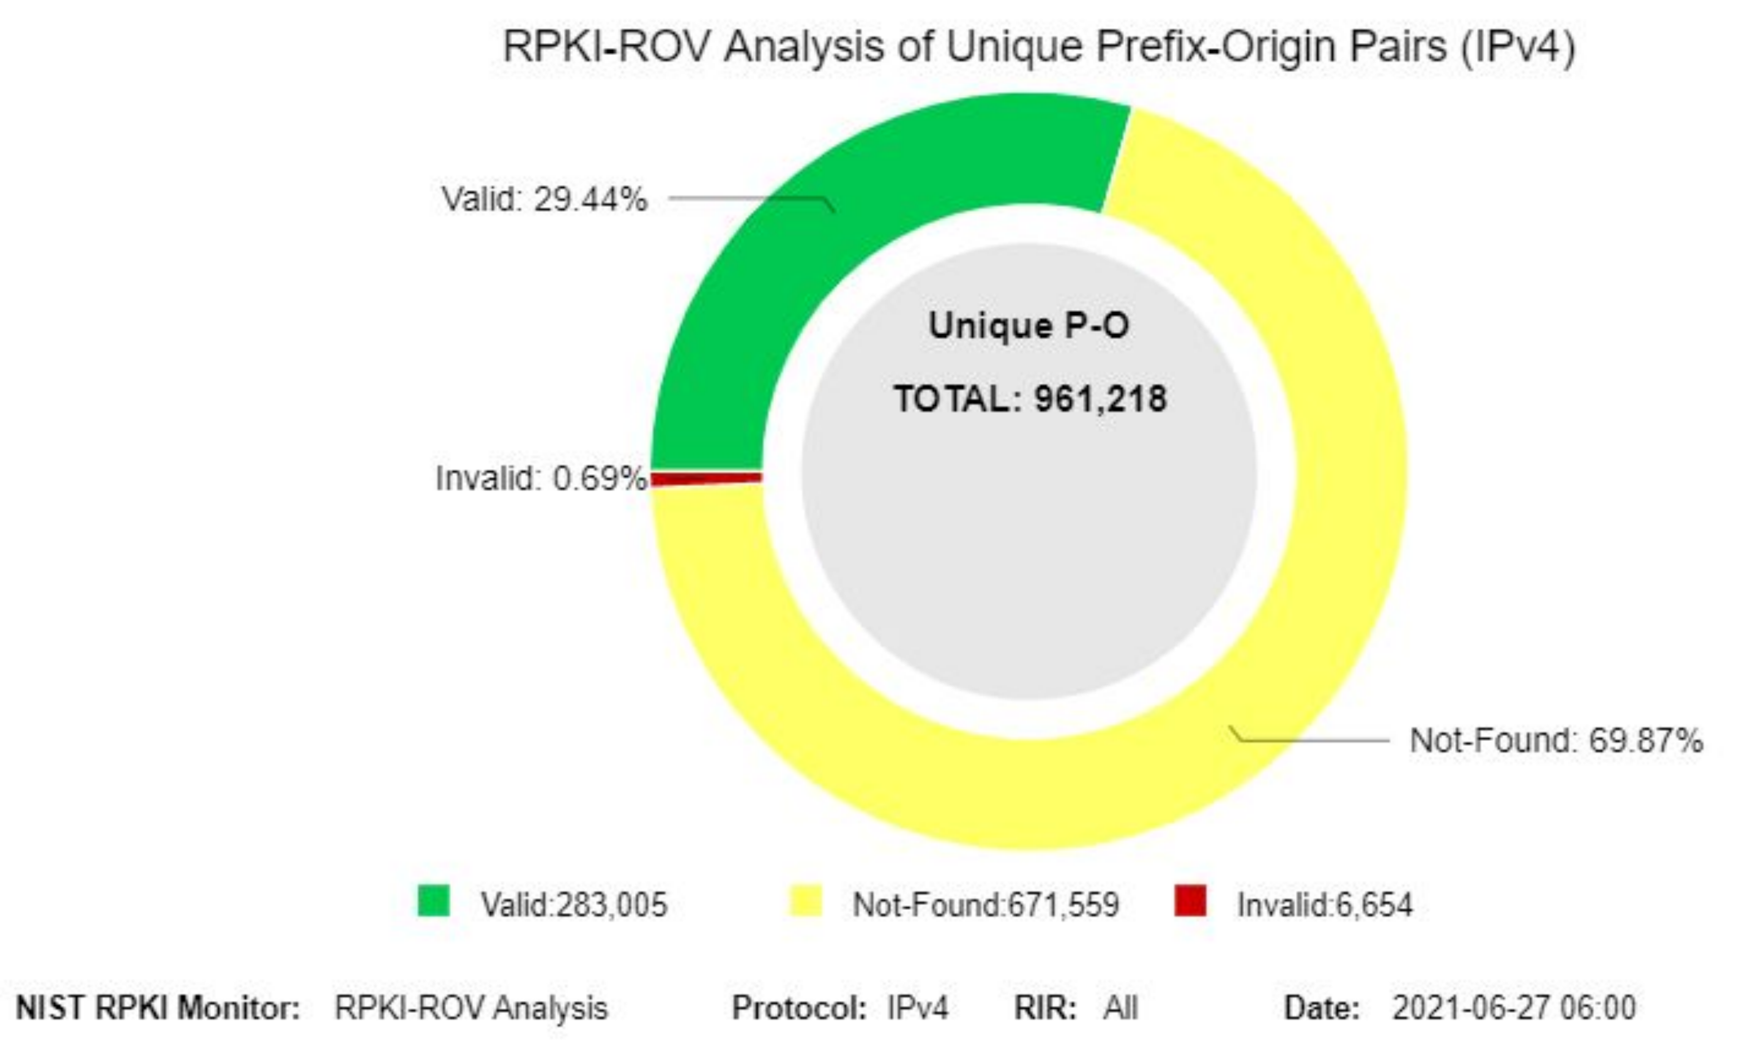

NIST RPKI Monitor: RPKI-ROV Analysis **Protocol:** IPv4 **RIR:** All **Date:** 2021-06-27 06:00

Courtesy: https://rpki-monitor.antd.nist.gov/

RPKI-ROV Analysis: Global Analysis (IPv4 - Comparing RIRs)

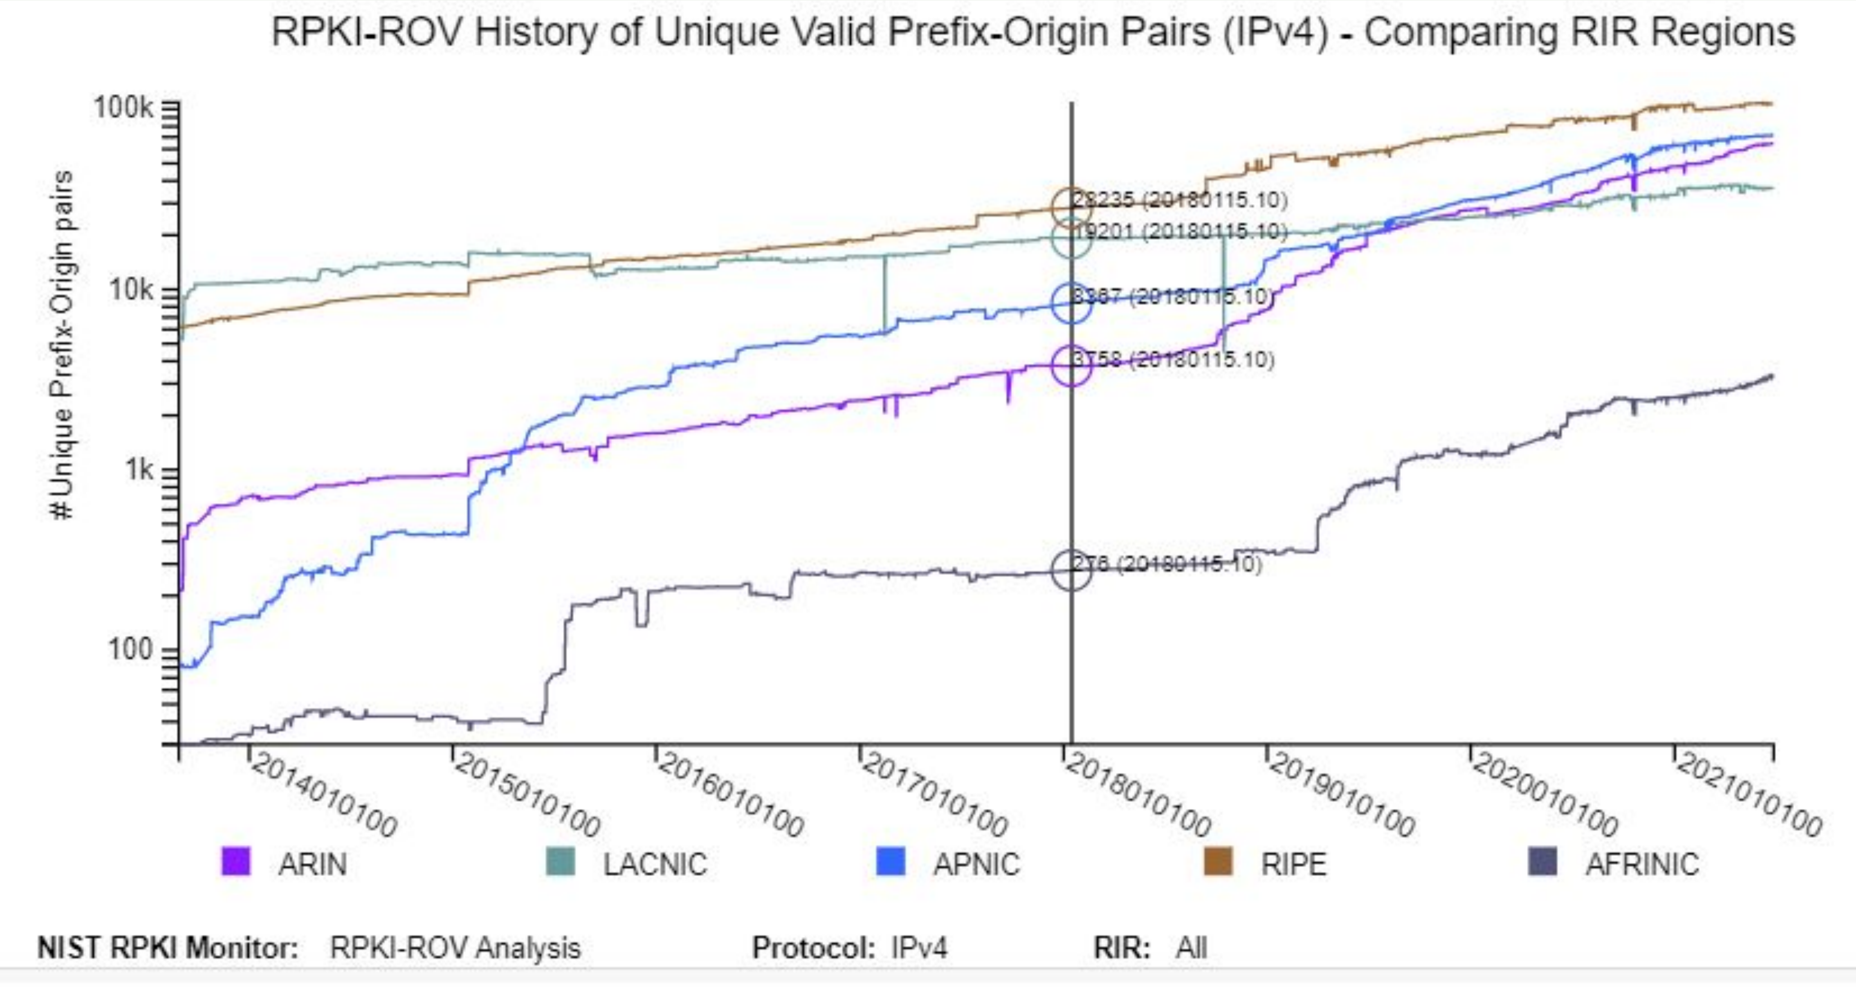

NIST RPKI Monitor: RPKI-ROV Analysis **Protocol:** IPv4 **RIR:** All

Courtesy: https://rpki-monitor.antd.nist.gov/

AS with most BGP Originated Prefixes VALID by RPKI-ROV (IPv4)

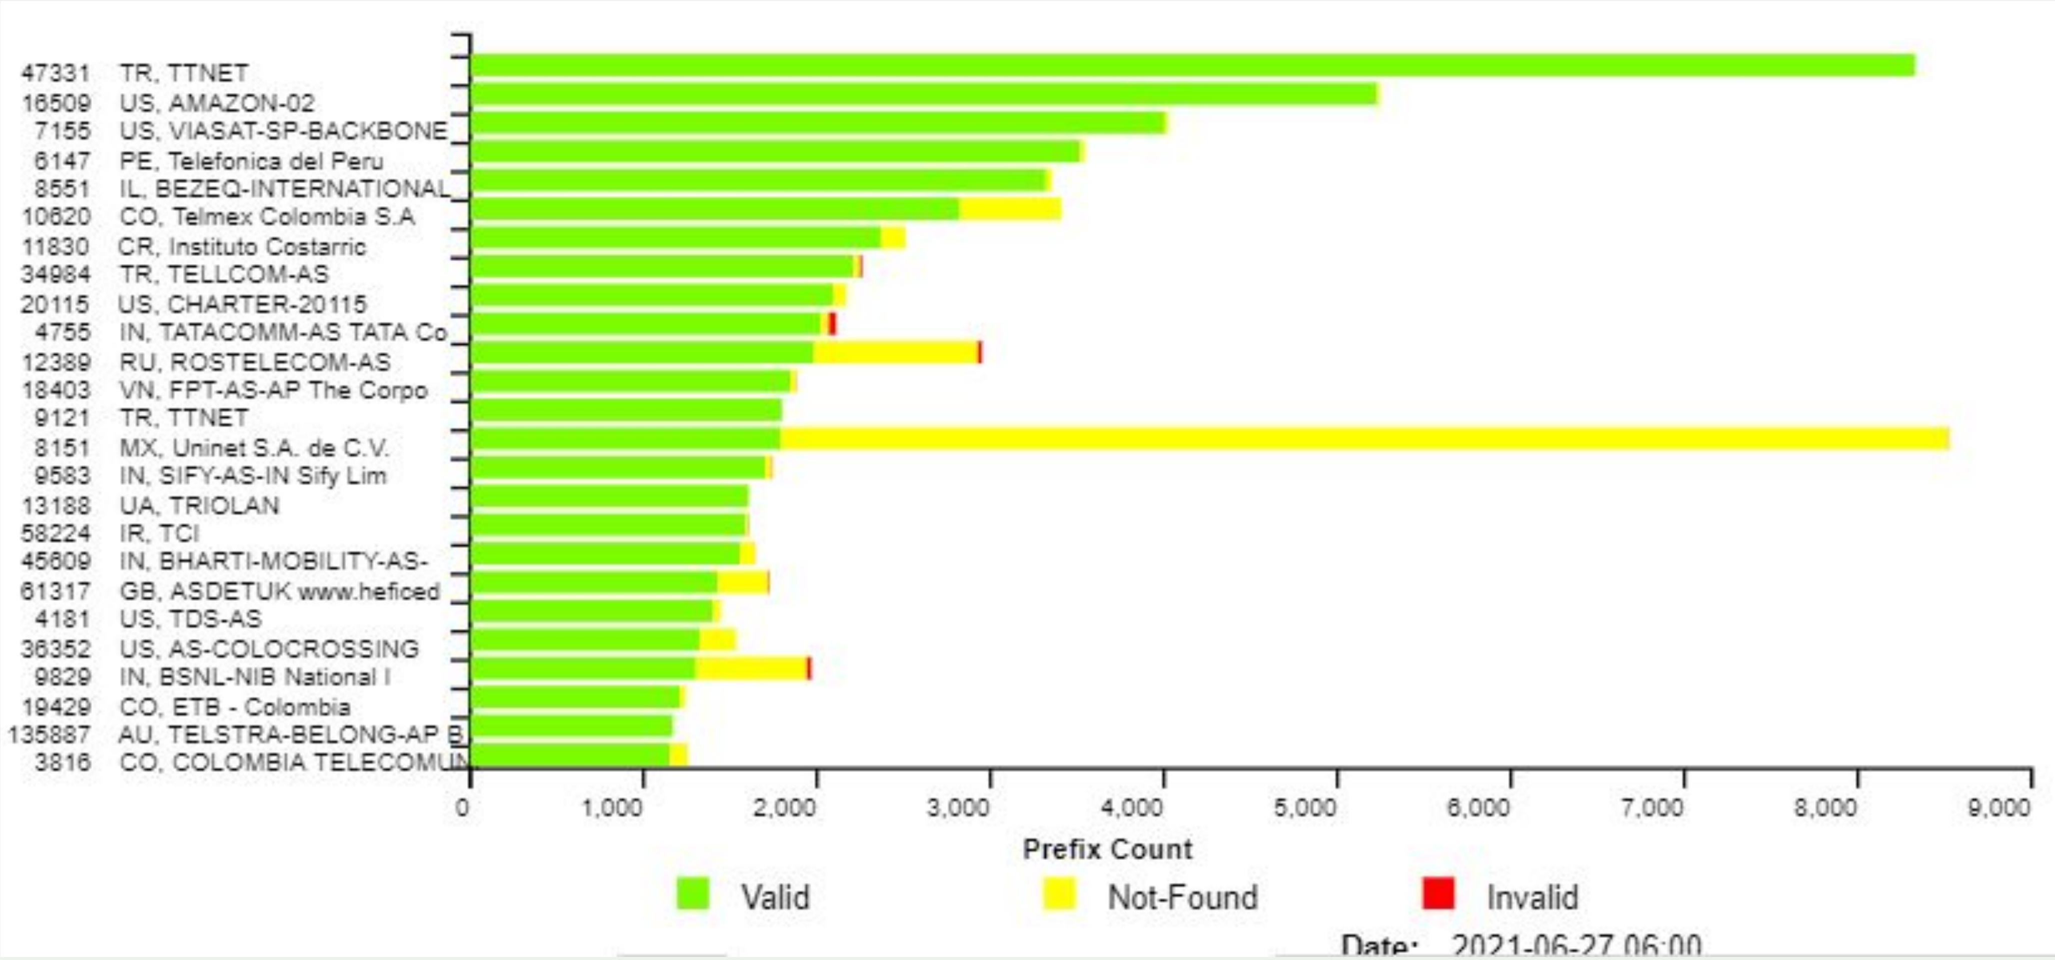

Courtesy: https://rpki-monitor.antd.nist.gov/Val

AS with the most BGP Originated Prefixes INVALID by RPKI-ROV (IPv4)

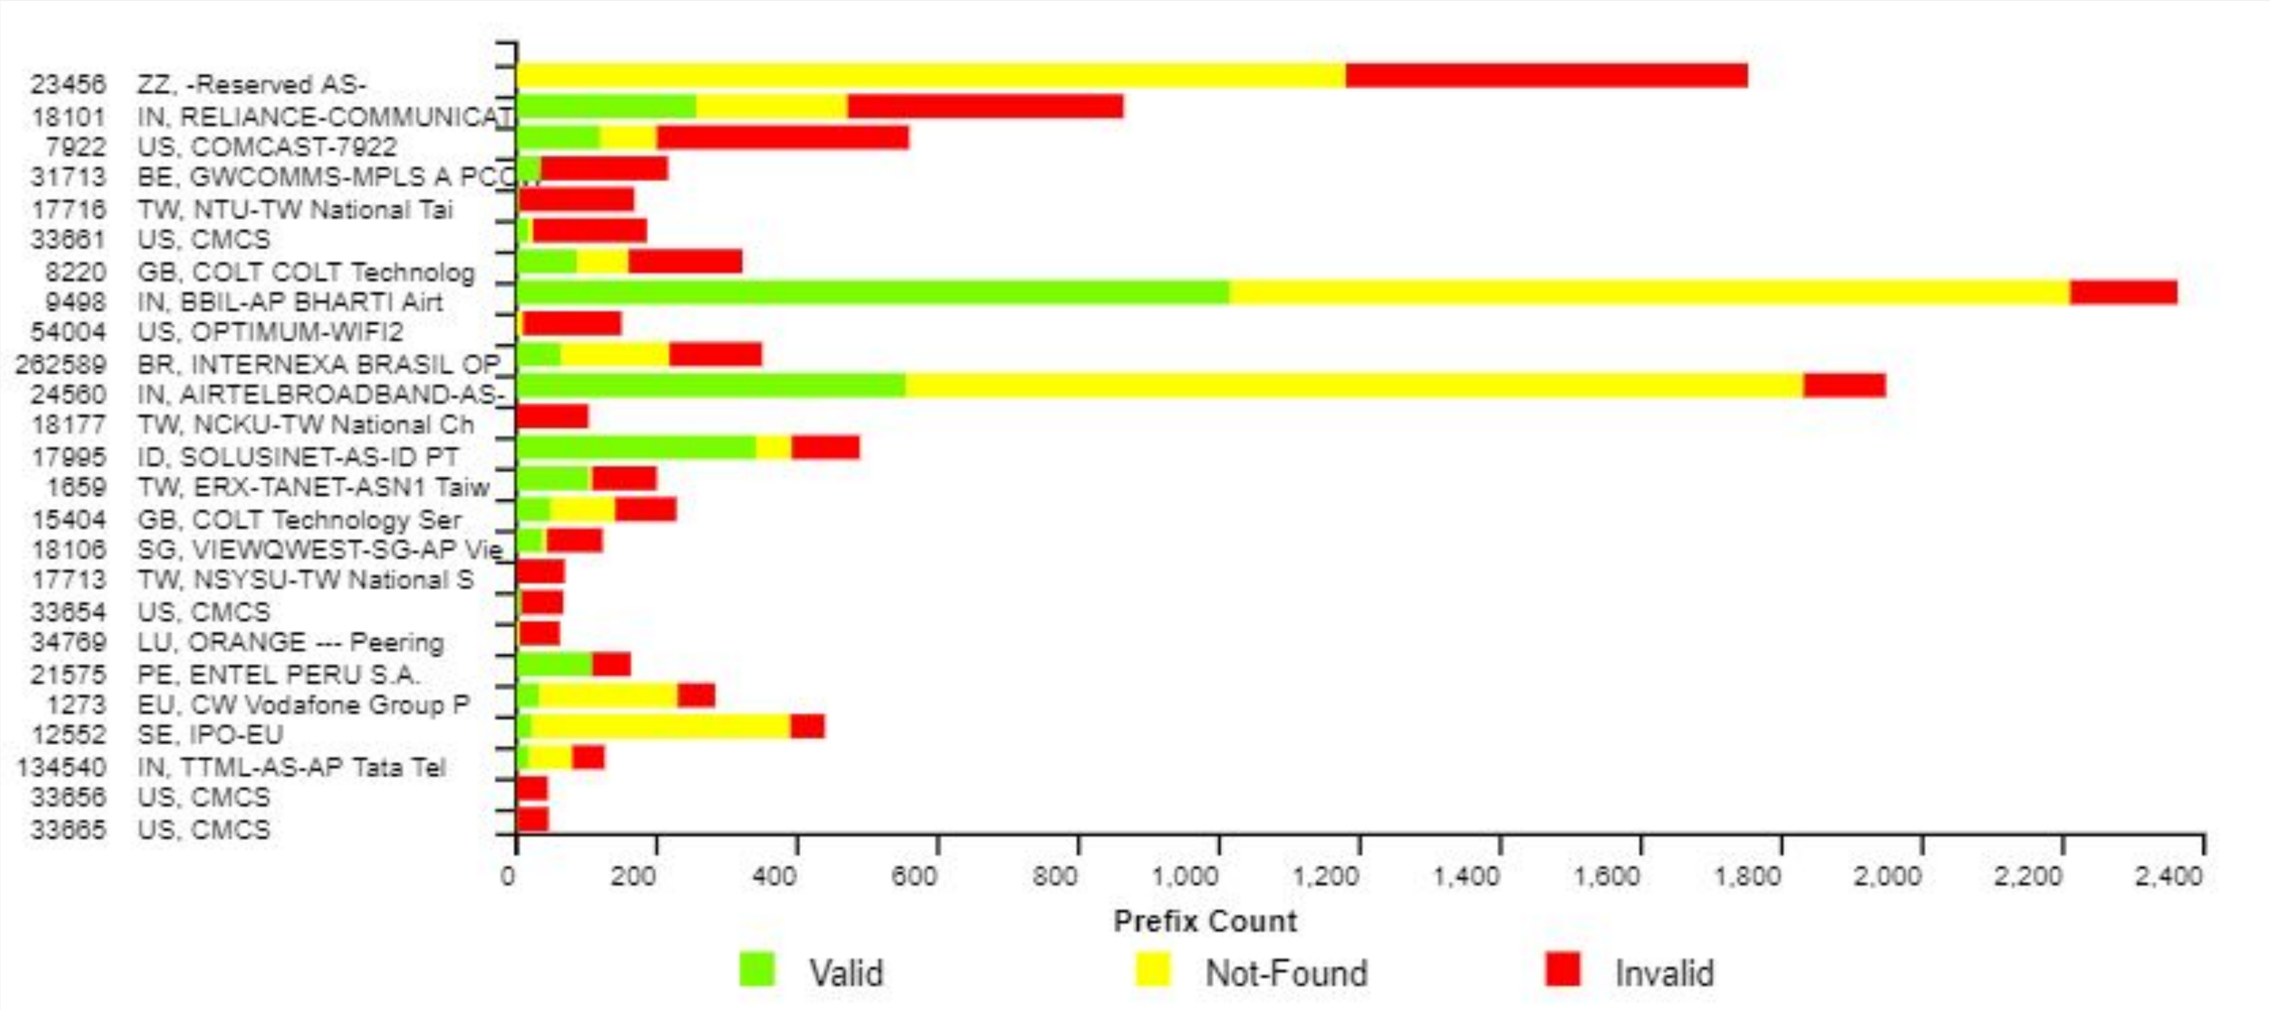

Courtesy: https://rpki-monitor.antd.nist.gov/Inv

Route Origin Validation – AS0

AS0 has been reserved for network operators to identify non-routed networks

“A ROA with a subject of AS0 (AS0 ROA) is an attestation by the holder of a prefix that the prefix described in the ROA, and any more specific prefix, should not be used in a routing context”

- Any prefixes with ROA having AS0 as the origin AS needs to be dropped
    - AS0 ROA has a lower relative preference than any other ROA that has a routable AS as its subject

Route Origin Validation – AS0

Possible use cases of AS0

- IXP LAN that must never appear in the global BGP table.
- Private address space (IPv4) and non-Global Unicast address space (IPv6)
- Unassigned address space.
  - under discussion within the various RIR policy forums
- Special purpose address space (IPv4 and IPv6)

Route Origin Validation – AS0

An address holder can use an AS0 ROA to do two things

- “Disavow” a prefix they have no intention of asserting into BGP
  - Publish the AS0 ROA for the prefix. Do not publish any other ROA
- “turn on” RPKI for a range of resources and selectively announce specific prefixes
  - Publish the AS0 ROA for the covering prefix
  - Publish the more specific (or even identical) prefixes with valid origin-AS
  - Longest match, and “AS0-ROA has lowest priority” combine.

APNIC has now published its AS0 TAL

- Operated separately from the regular TAL

https://www.apnic.net/community/security/resource-certification/tal-archive/

APNIC RPKI Number Resource Management

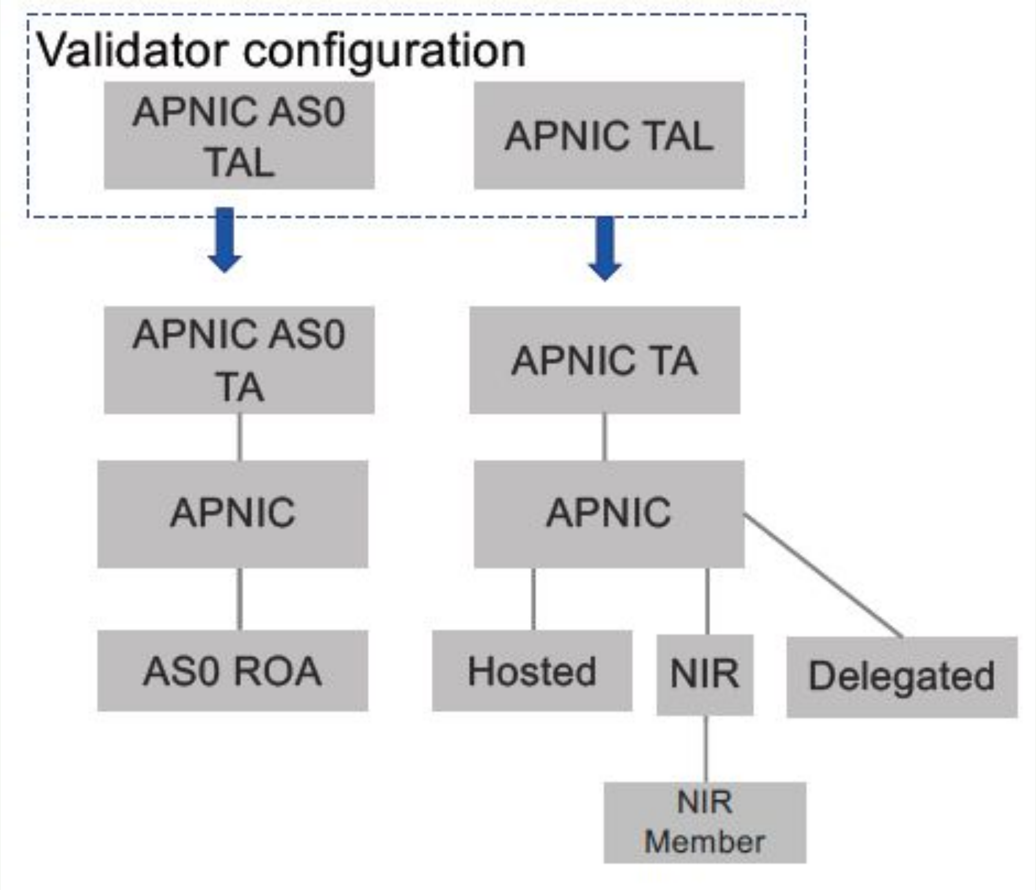

- APNIC runs two independent RPKI systems
- AS0 ROA runs from a fully independent TA
- Unlike mainline RPKI there is no "hosted" or "delegated" or delegated component.
- It's only visible end product is as ROA for AS0

Courtesy: apnic

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 8205)**
- **ASPA (draft)**

  https://datatracker.ietf.org/doc/html/draft-ietf-sidrops-aspa-verification-07

- **AS-Cones (draft)**

  https://datatracker.ietf.org/doc/html/draft-ietf-grow-rpki-as-cones-02

Why join MANRS?

Implementing MANRS Actions

- **Signals** an organization’s security-forward posture and can eliminate SLA violations that reduce profitability or cost customer relationships.
- **Reduces** routing incidents, helping networks readily identify and address problems with customers or peers.
- **Improves** network’s operations by establishing better and cleaner peering communication pathways, while also providing granular insight for troubleshooting.
- **Addresses** many concerns of security-focused enterprises and other customers.

Everyone Benefits

- Joining MANRS means joining a community of security-minded organizations committed to making the global routing infrastructure more robust and secure.
- Consistent MANRS adoption yields steady improvement, but we need more networks to implement the actions and more customers to demand routing security best practices.
- The more networks apply MANRS actions, the fewer incidents there will be, and the less damage they can do.

MANRS is an important step

Security is a process, not a state. MANRS provides a structure and a consistent approach to solving security issues facing the Internet.

MANRS is the minimum a network should consider, with low risk and cost-effective actions.

MANRS is not a one-stop solution to all the Internet’s routing woes, but it is an important step toward a globally robust and secure routing infrastructure.

MANRS Participants (June 2021)

- 567 Network Operators
- 87 Internet eXchange Points (IXP)
- 17 CDN and Cloud Providers

MANRS Observatory

MANRS Observatory

- The MANRS Observatory is a web-based tool that collates publicly available data sources including BGPStream, the CIDR Report, the CAIDA Spoofer Database, RIR Whois and IRR databases and PeeringDB to view routing incidents on any network (ASN) that is publicly visible on the Internet.
- It is also able to check the general routing health of particular networks, countries and regions, and provide a longer-term overview on whether routing incidents are getting better or worse.

MANRS Observatory

Provide a factual state of security and resilience of the Internet routing system and track it over time

Measurements are:

- Transparent – using publicly accessible data
- Passive – no cooperation from networks required
- Evolving – MANRS community decide what gets measured and how

MANRS Observatory Access

Publicly launched in August 2019

Uses trusted, publicly available third-party data

Anyone may view aggregated data

Only MANRS Participants have access to detailed data about their own network

Caveats:

- There are still some false positives
- Lack of security controls is not always visible

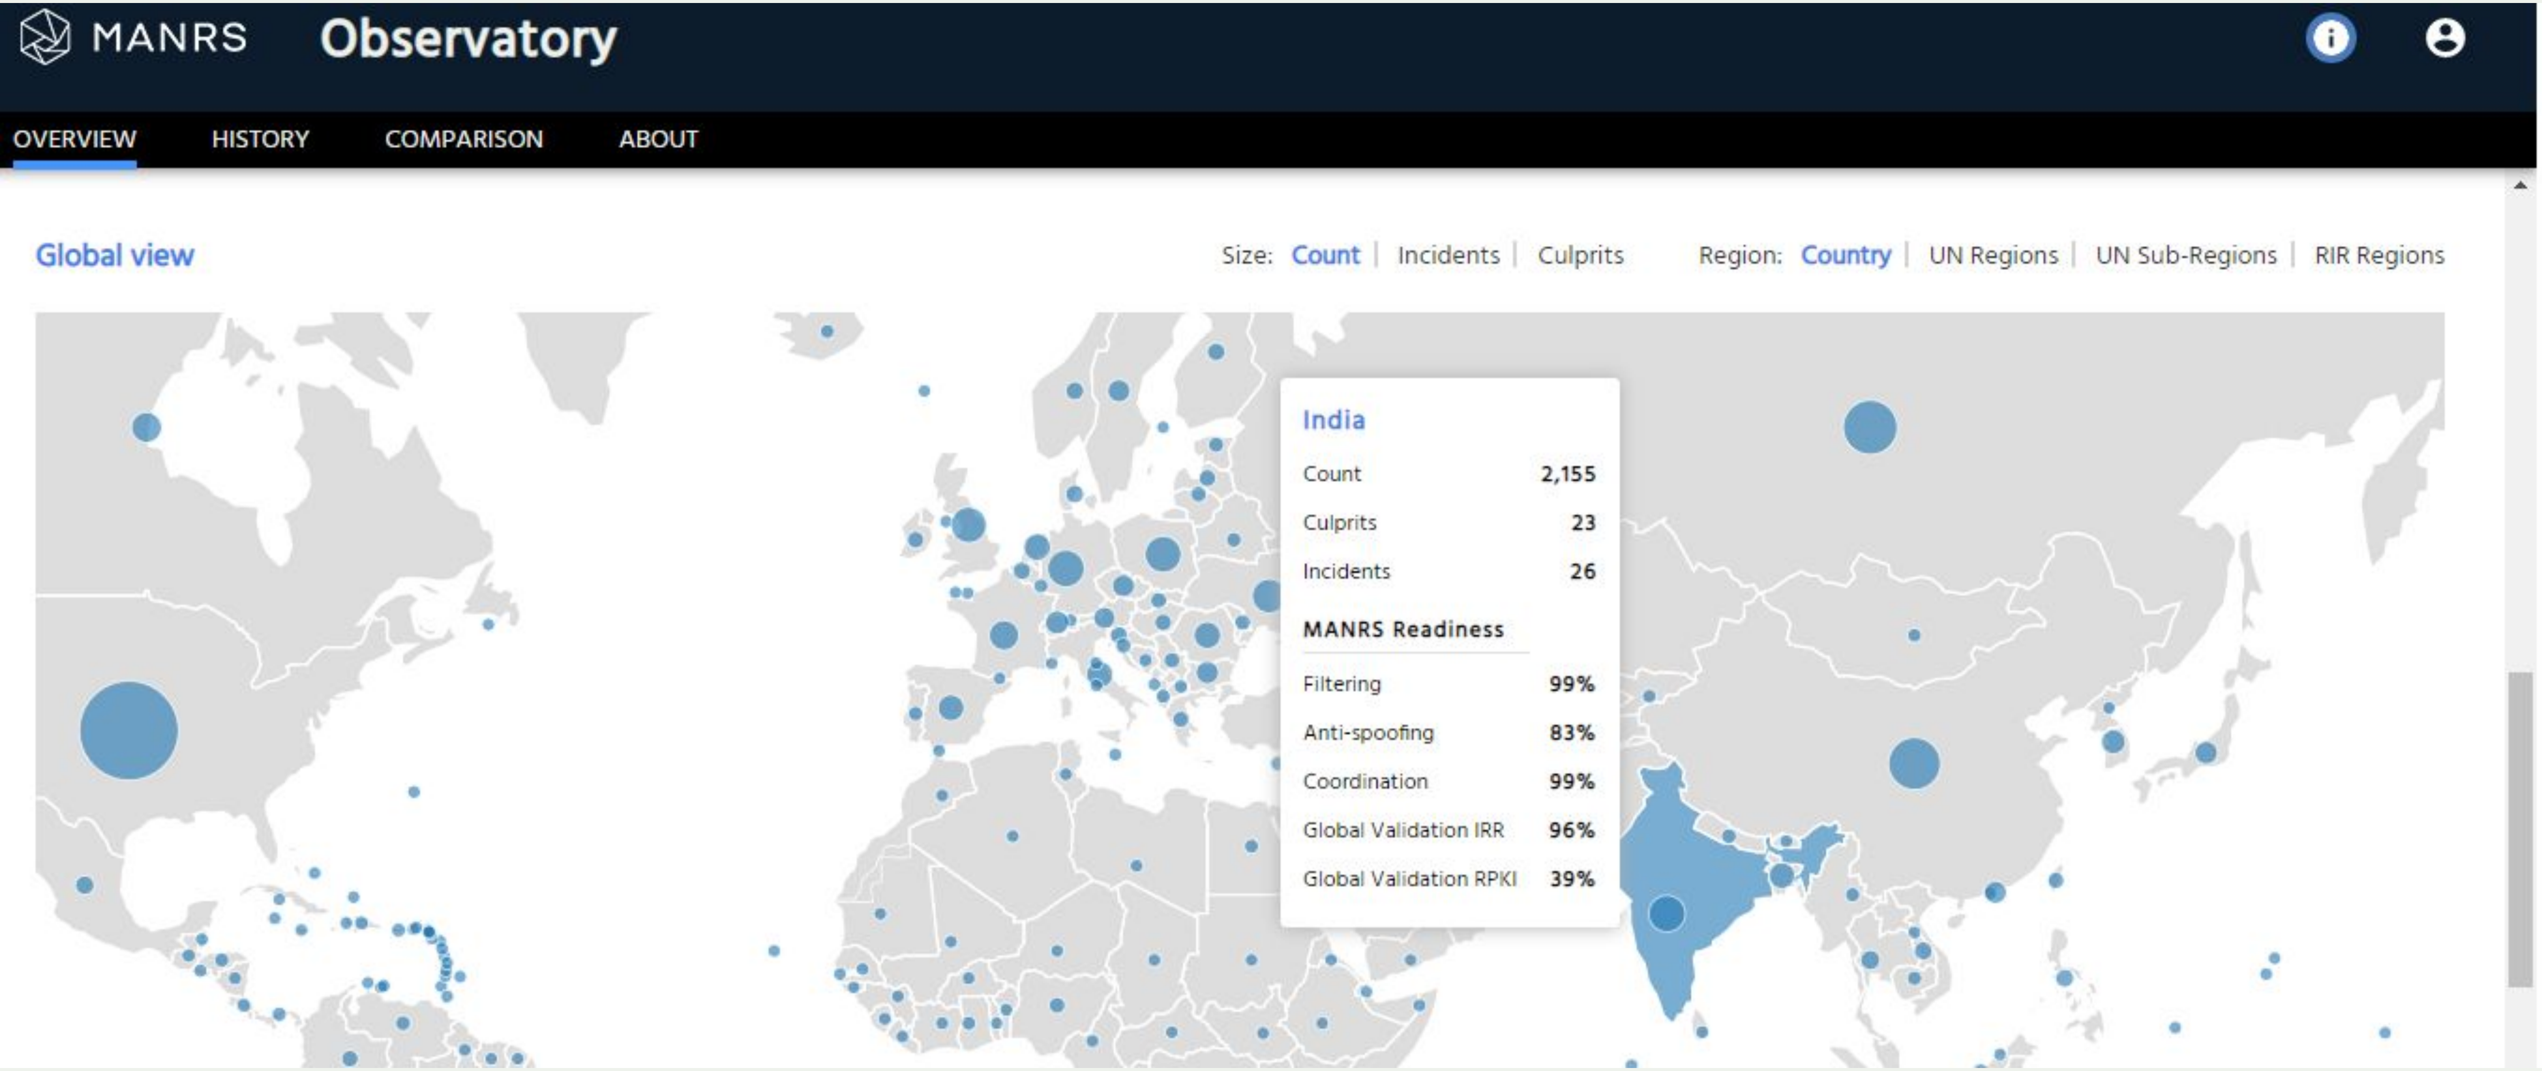
MANRS
Observatory
OVERVIEW
HISTORY
COMPARISON
ABOUT
Global view
Size: Count | Incidents | Culprits
Region: Country | UN Regions | UN Sub-Regions | RIR Regions

India

Count	2,155
Culprits	23
Incidents	26

MANRS Readiness

Filtering	99%
Anti-spoofing	83%
Coordination	99%
Global Validation IRR	96%
Global Validation RPKI	39%

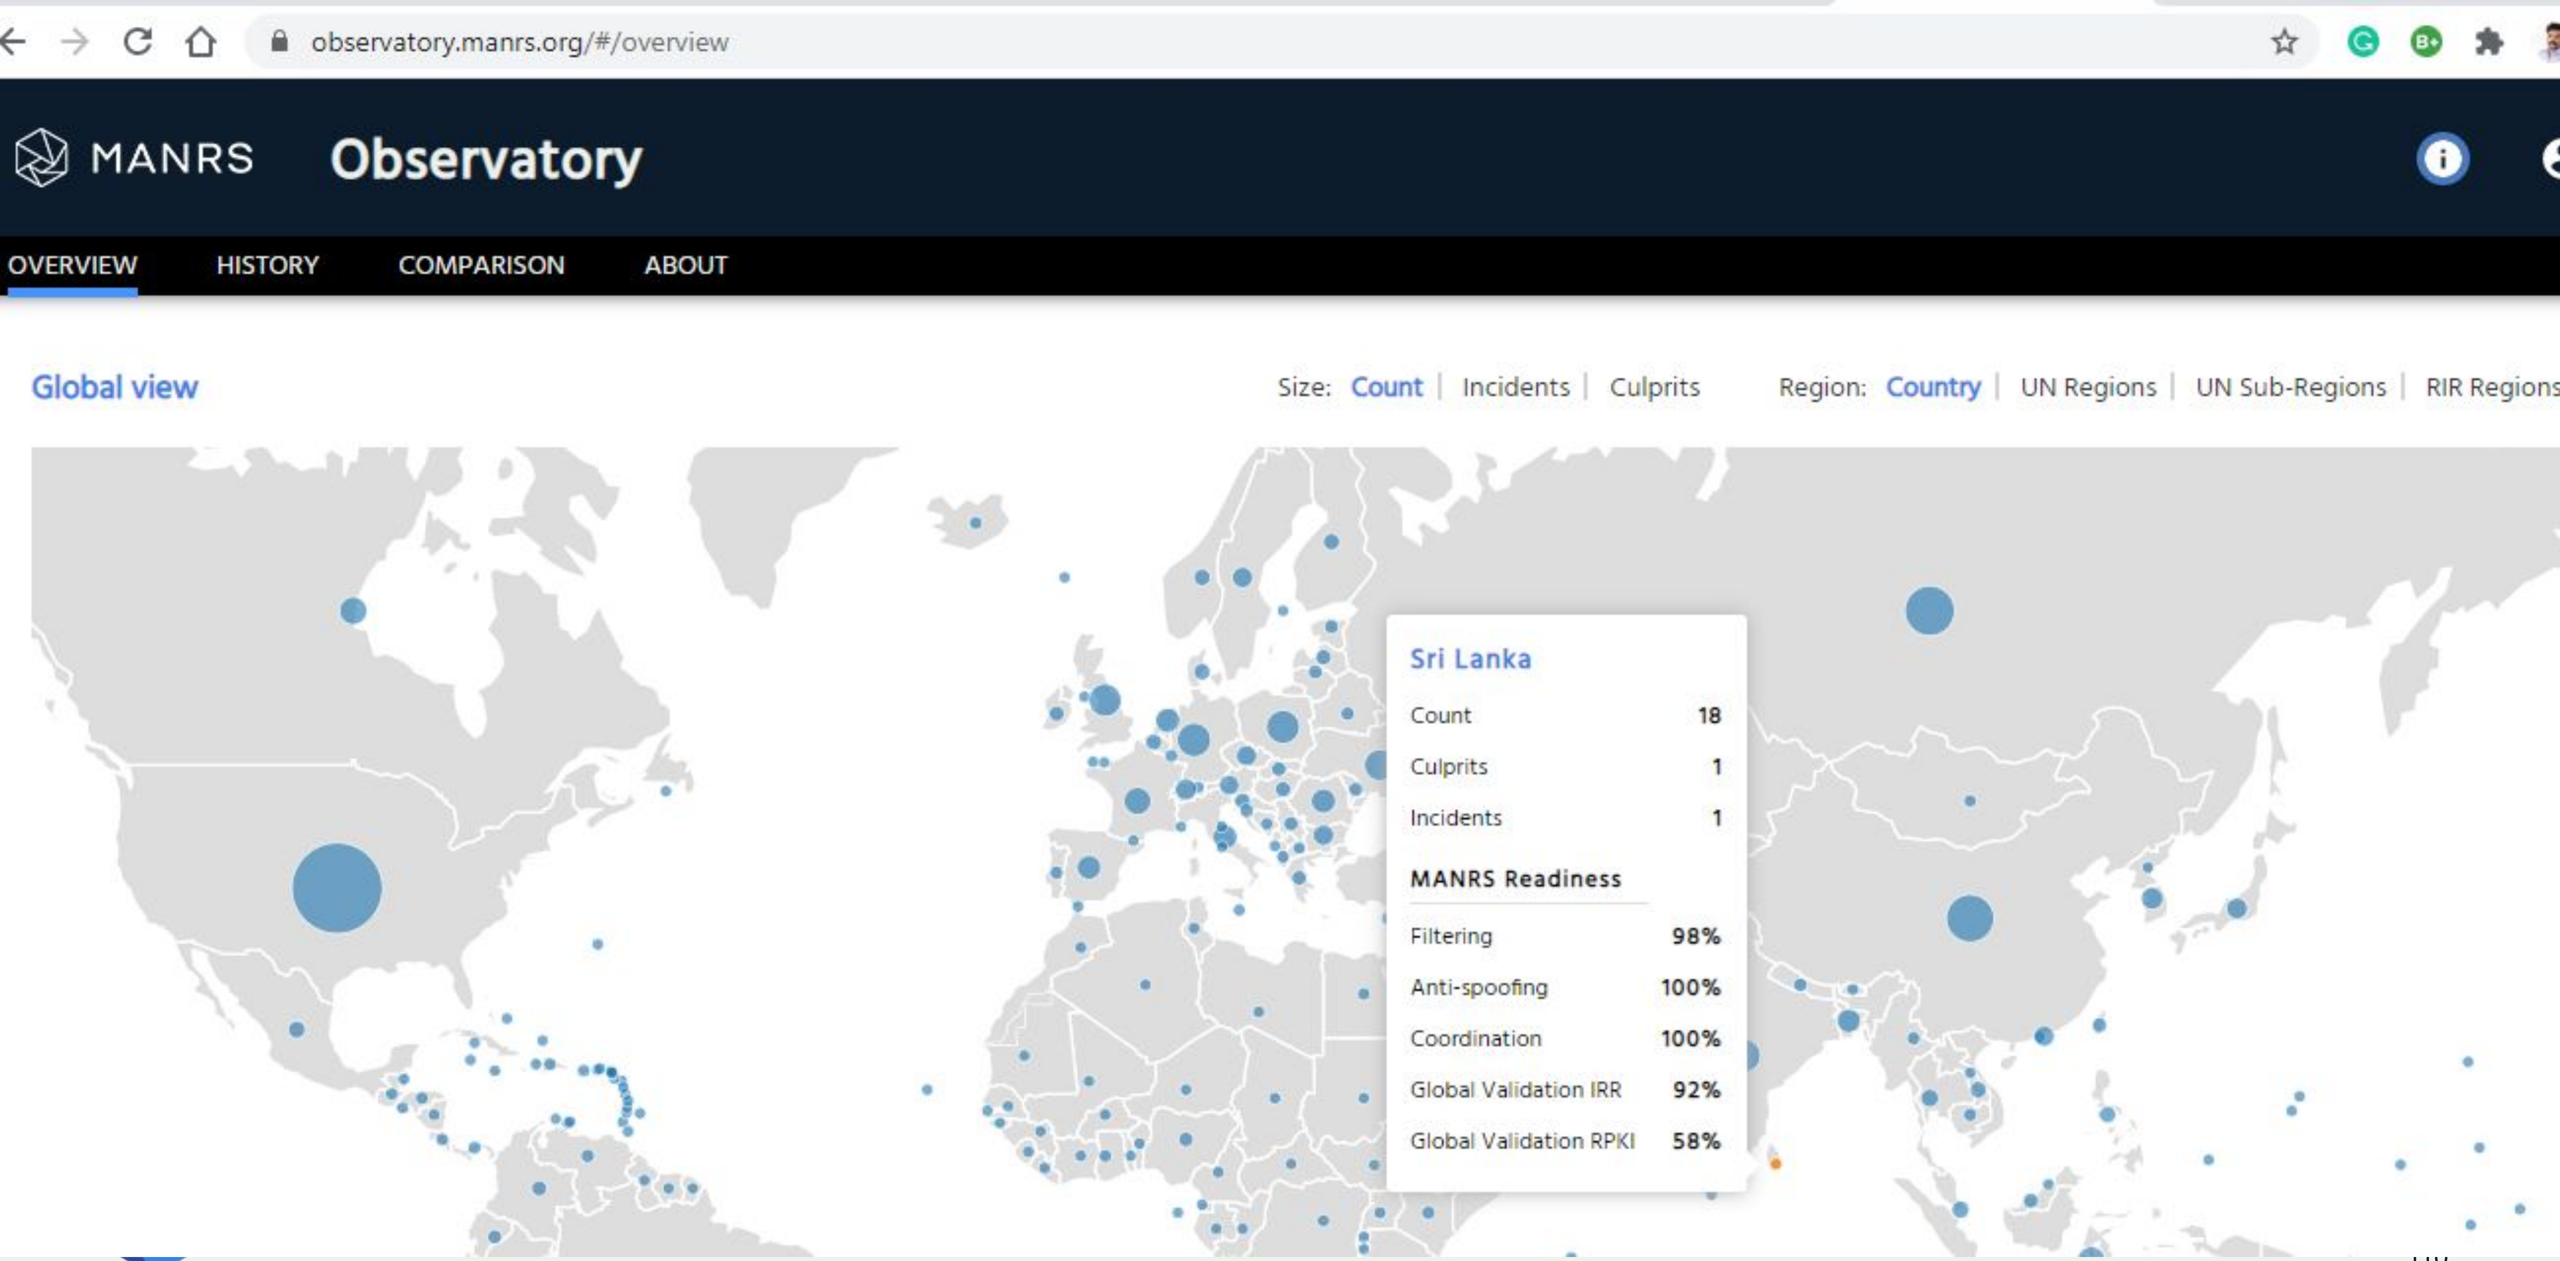

observatory.manrs.org/#/overview
MANRS
OVERVIEW
HISTORY
COMPARISON
ABOUT
Size: Count | Incidents | Culprits
Region: Country | UN Regions | UN Sub-Regions | RIR Regions

Observatory

Global view

Sri Lanka

Count	18
Culprits	1
Incidents	1

MANRS Readiness

Filtering	98%
Anti-spoofing	100%
Coordination	100%
Global Validation IRR	92%
Global Validation RPKI	58%

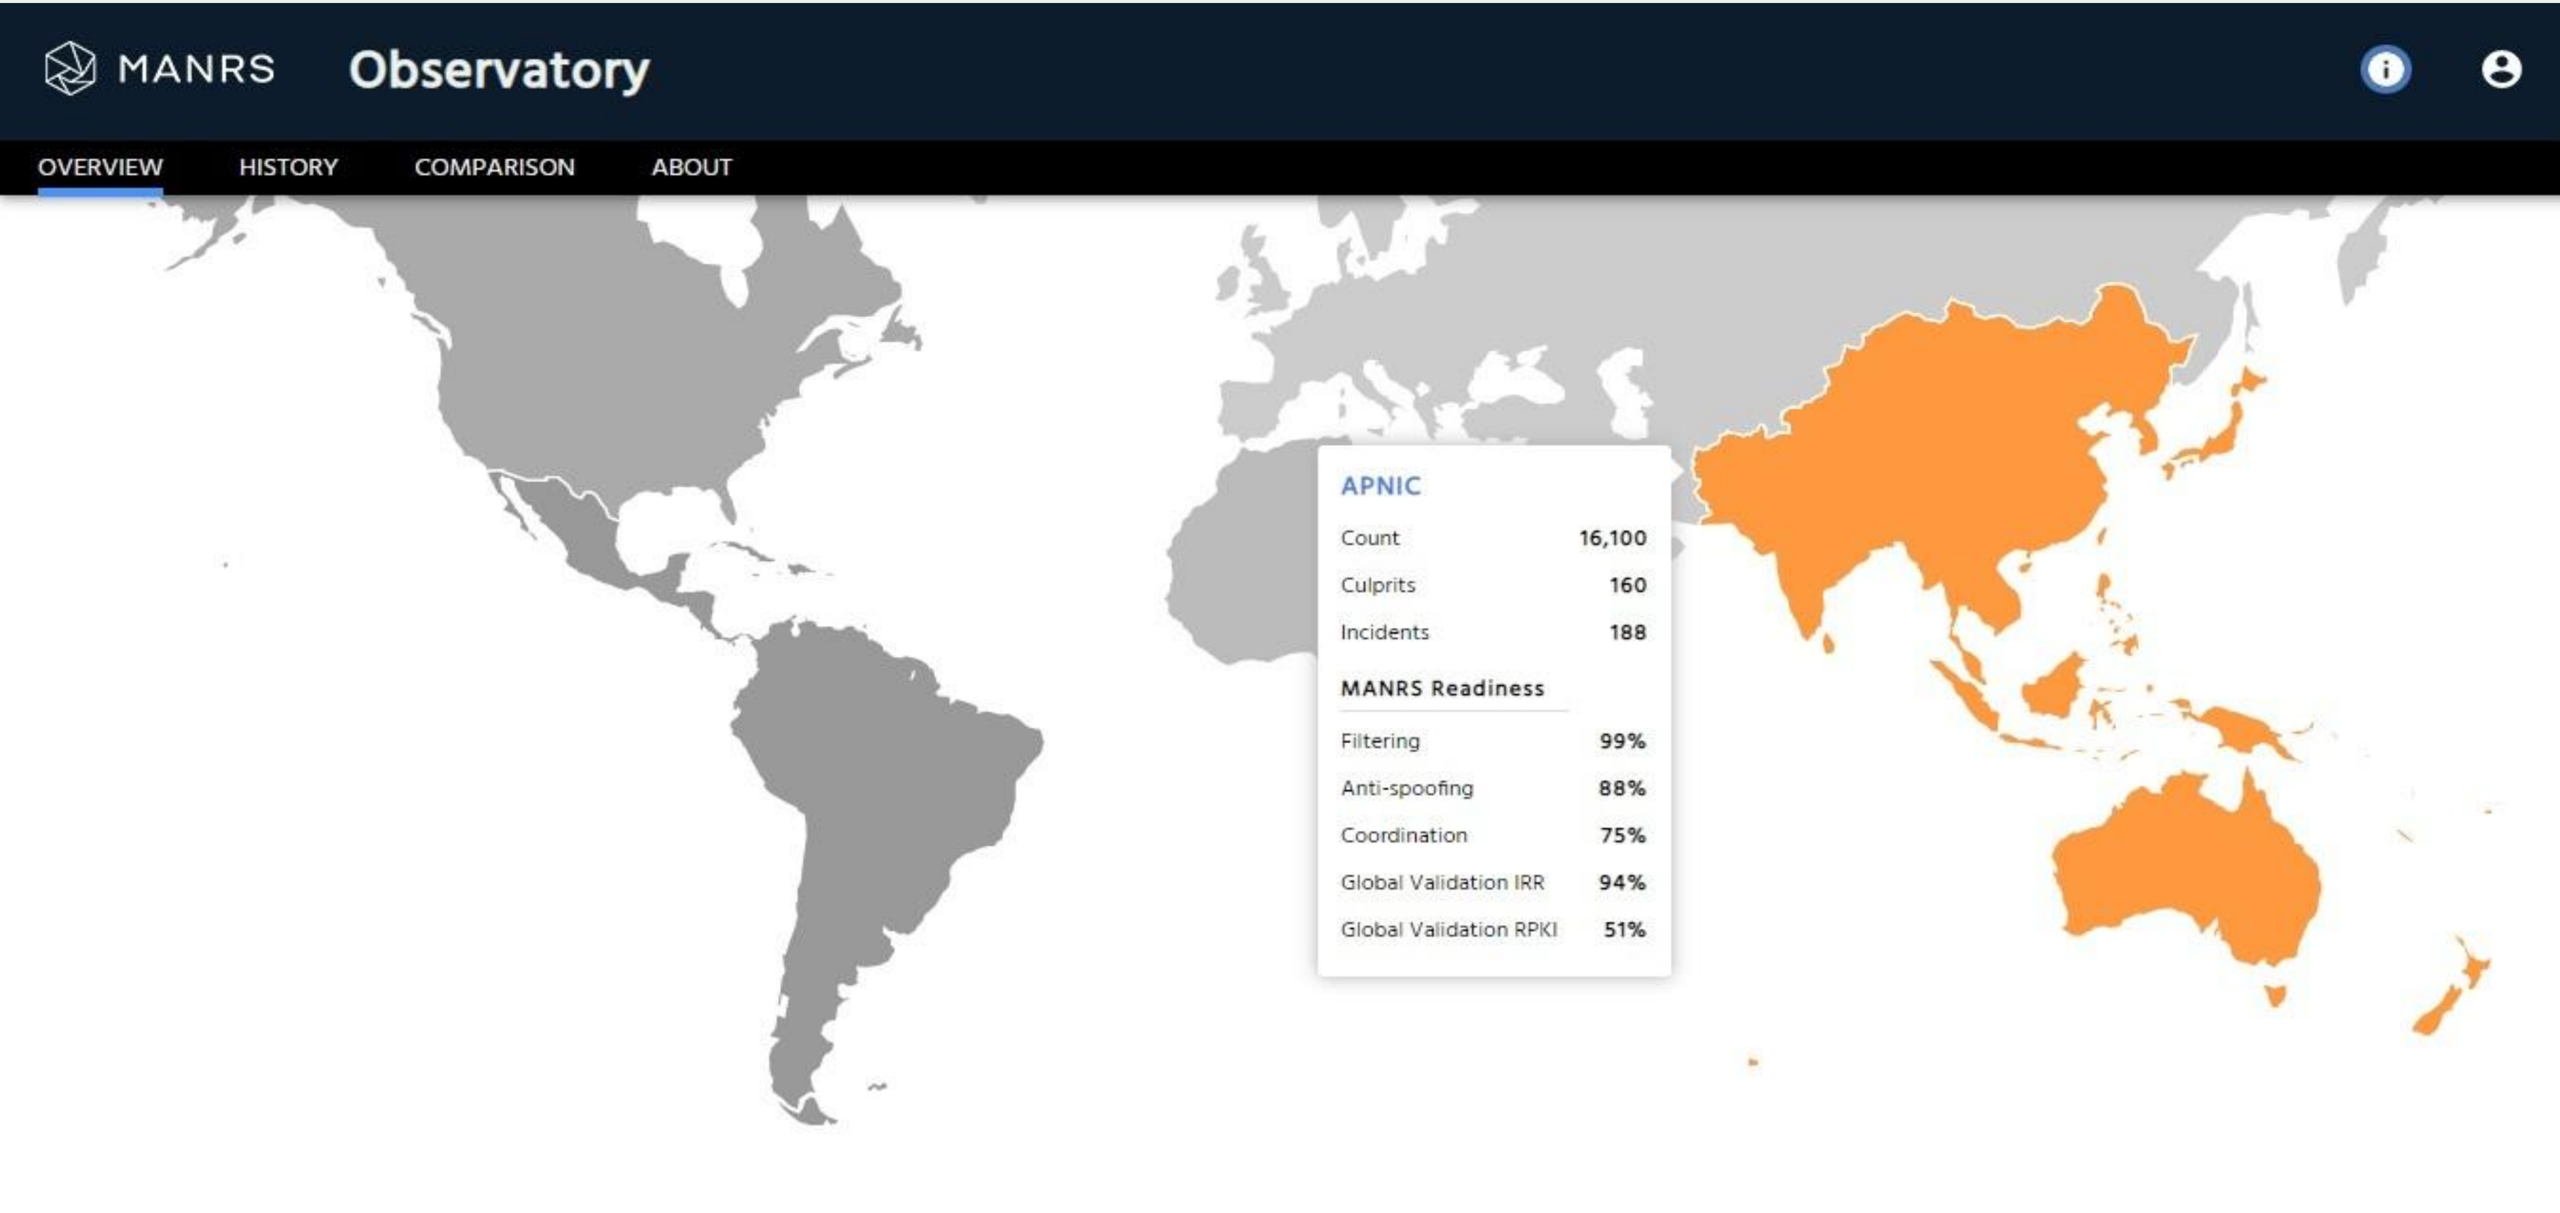

APNIC

Count	16,100
Culprits	160
Incidents	188

MANRS Readiness

Filtering	99%
Anti-spoofing	88%
Coordination	75%
Global Validation IRR	94%
Global Validation RPKI	51%

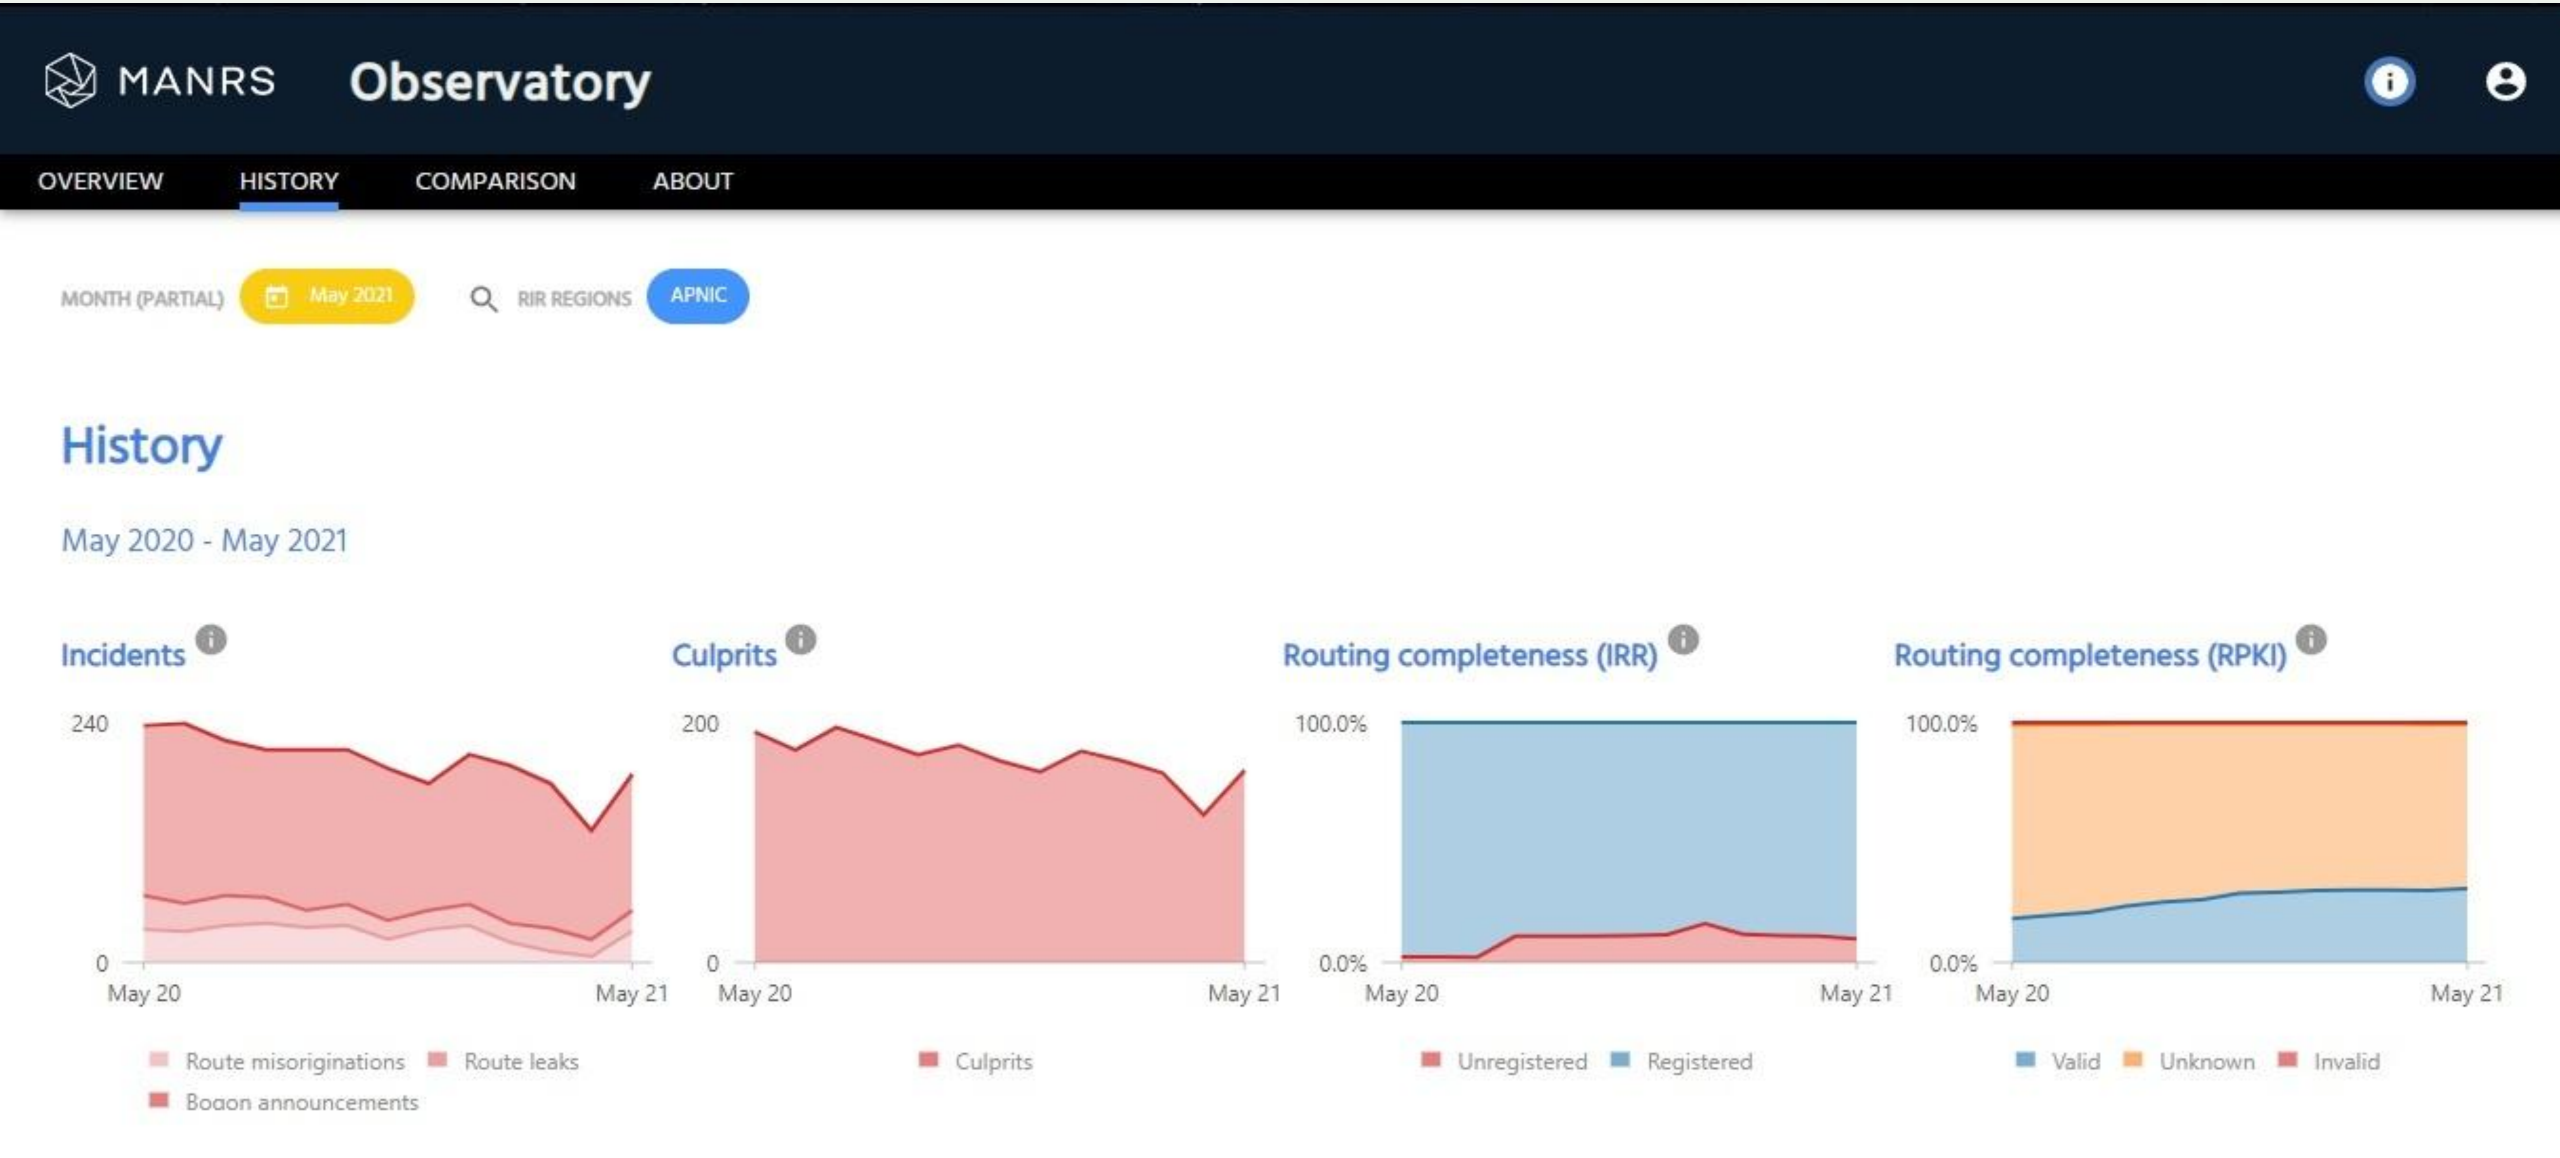

MANRS Observatory

OVERVIEW HISTORY COMPARISON ABOUT

MONTH (PARTIAL) May 2021

RIR REGIONS APNIC

History

May 2020 - May 2021

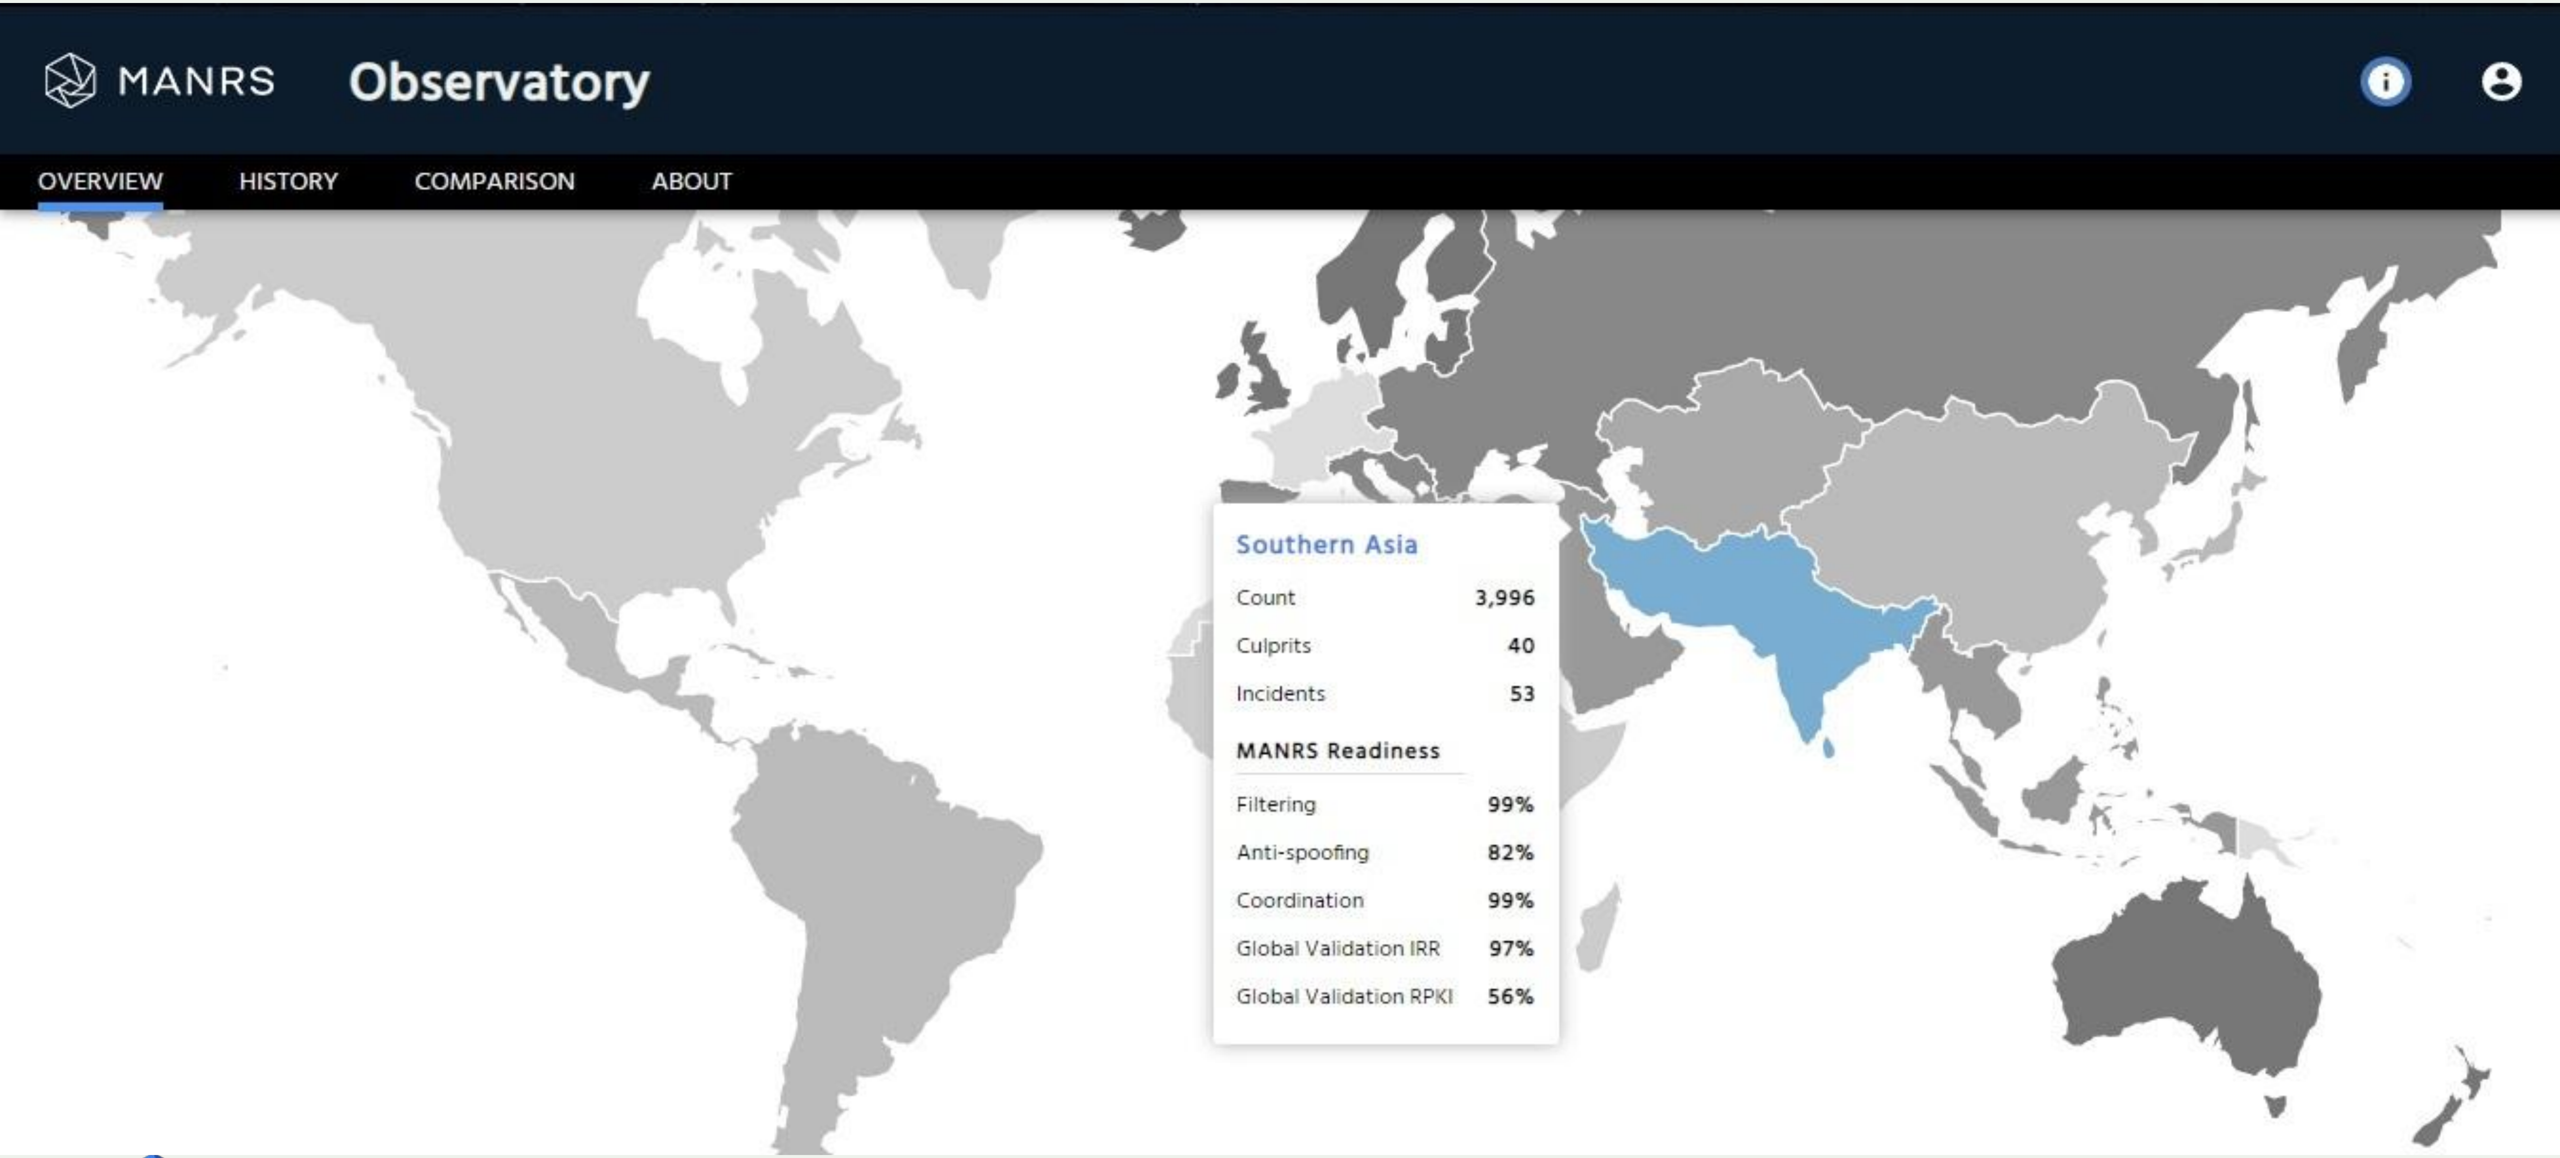

Southern Asia

Count	3,996
Culprits	40
Incidents	53

MANRS Readiness

Filtering	99%
Anti-spoofing	82%
Coordination	99%
Global Validation IRR	97%
Global Validation RPKI	56%

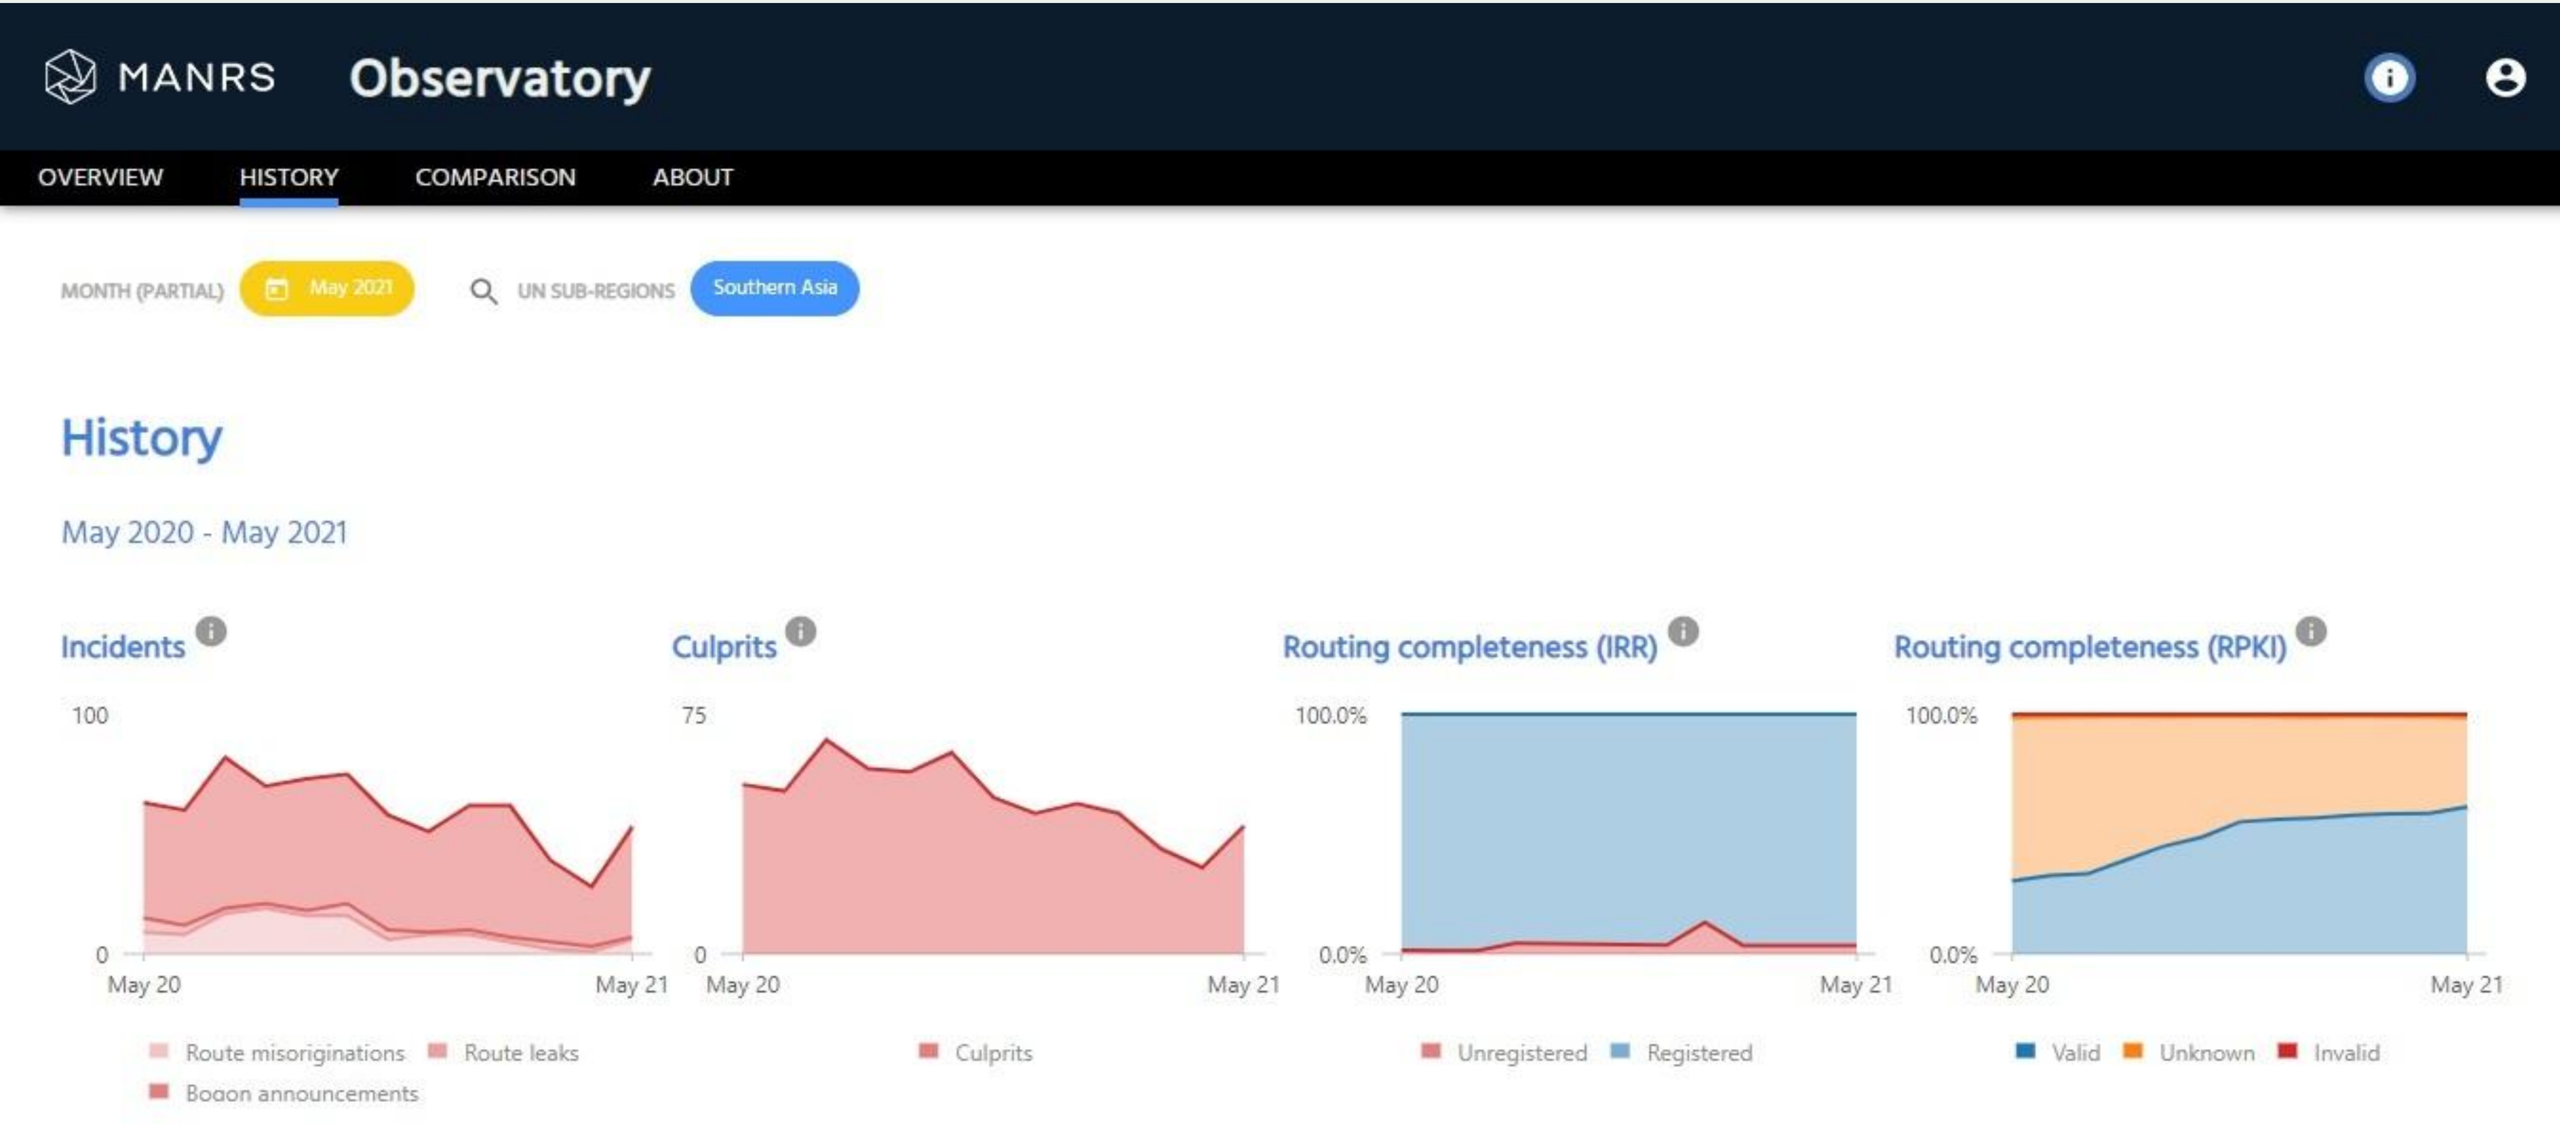

MONTH (PARTIAL): May 2021

UN SUB-REGIONS: Southern Asia

History

May 2020 - May 2021

MANRS - CDN and Cloud Participants

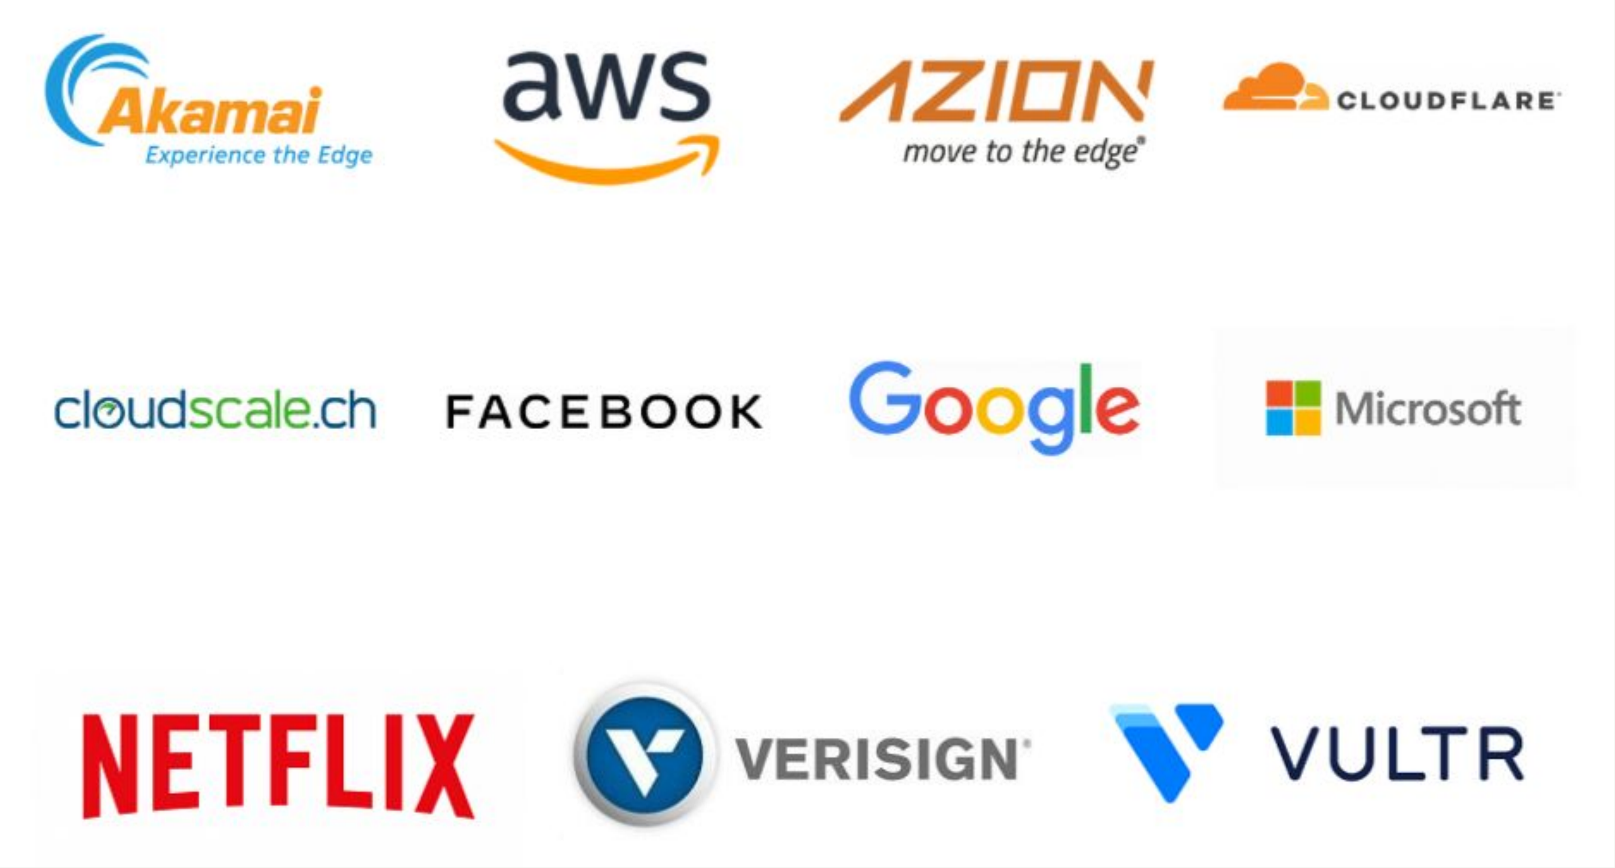

https://www.manrs.org/cdn-cloud-providers/participants/

Join Us

Visit https://www.manrs.org

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

Get Involved in the Community

- Members support the initiative and implement the actions in their own networks
- Members maintain and improve the document and promote MANRS objectives

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

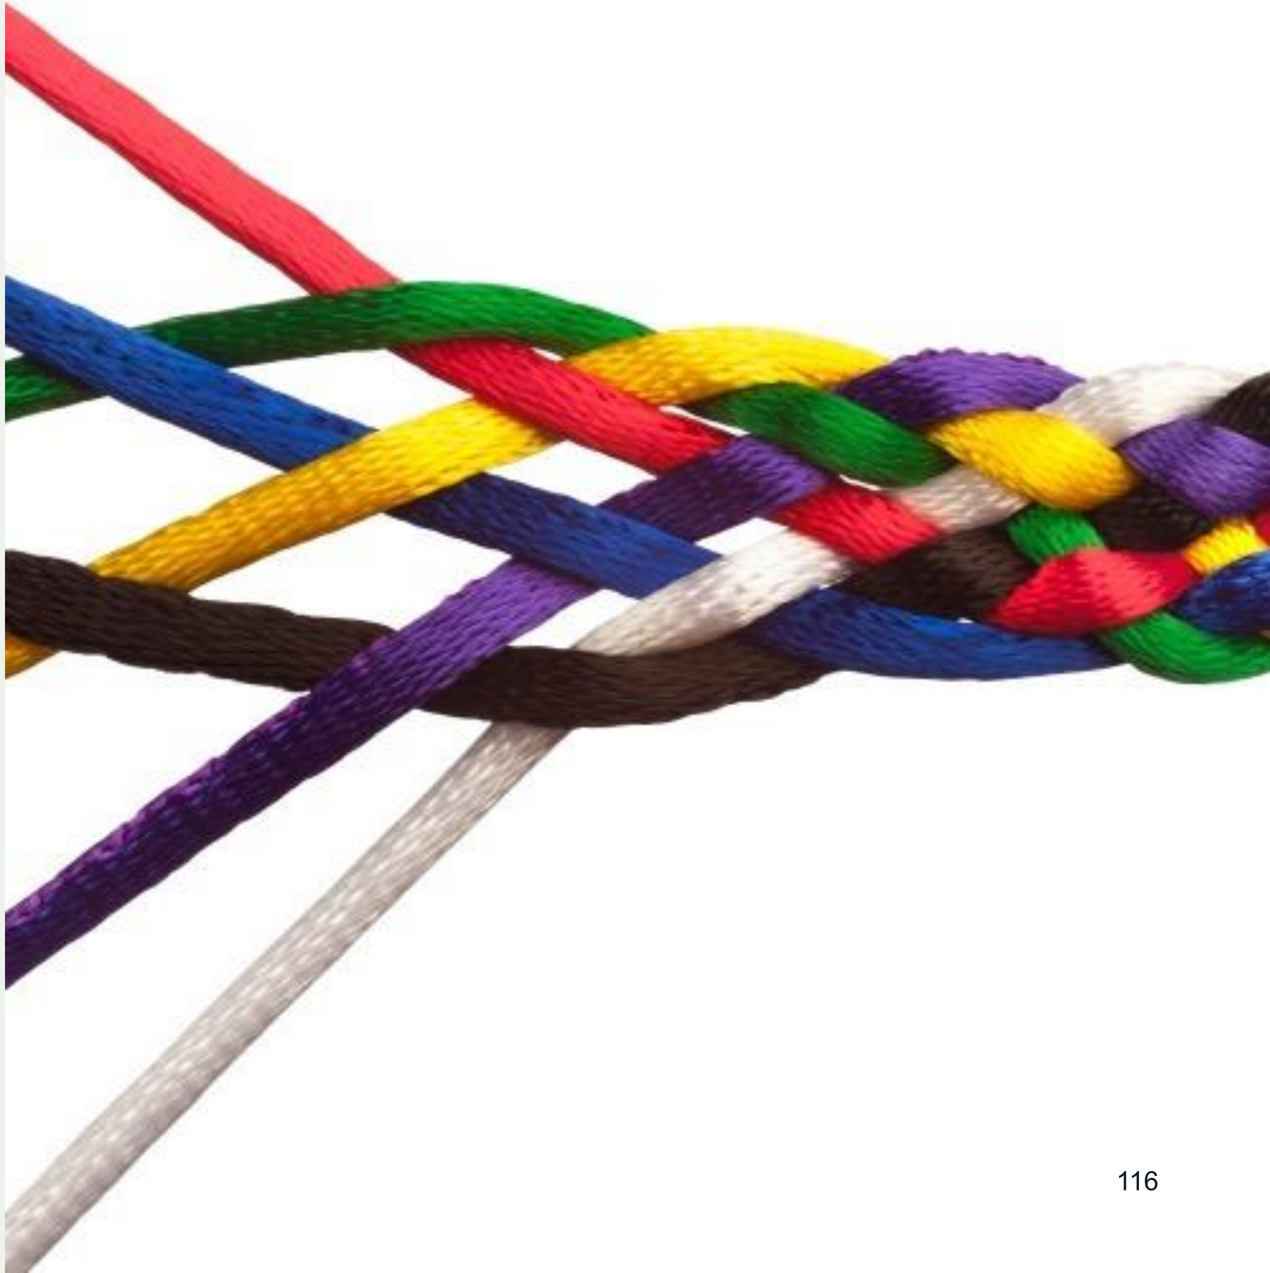

MANRS Implementation Guide

If you're not ready to join yet, implementation guidance is available to help you.

- Based on Best Current Operational Practices deployed by network operators, IXPs, CDNs and Cloud providers around the world

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

  https://www.manrs.org/ixps/

  https://www.manrs.org/cdn-cloud-providers/

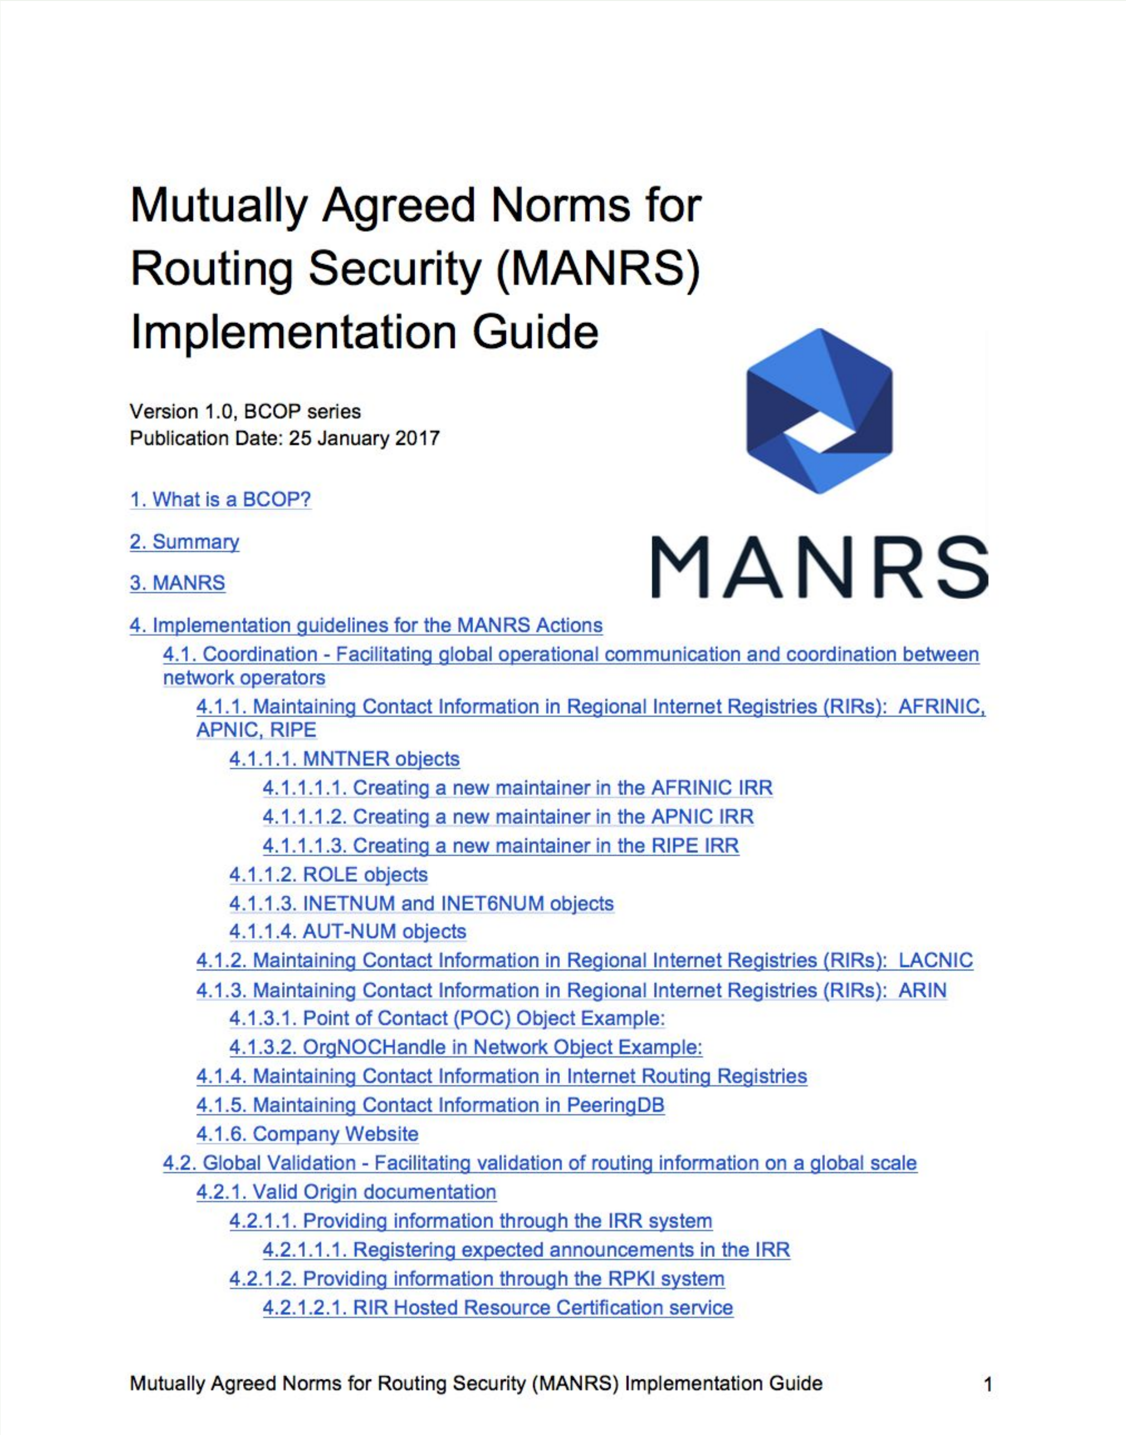

Mutually Agreed Norms for Routing Security (MANRS) Implementation Guide

Version 1.0, BCOP series
Publication Date: 25 January 2017

Mutually Agreed Norms for Routing Security (MANRS) Implementation Guide 1

MANRS Training Modules (Self Paced)

6 training modules based on information in the Implementation Guide.

Walks through the tutorial with a test at the end of each module.

Working with and looking for partners that are interested in integrating it in their curricula.

https://academy.apnic.net/en/course/manrs/

Thanks to APNIC for hosting MANRS Tutorial

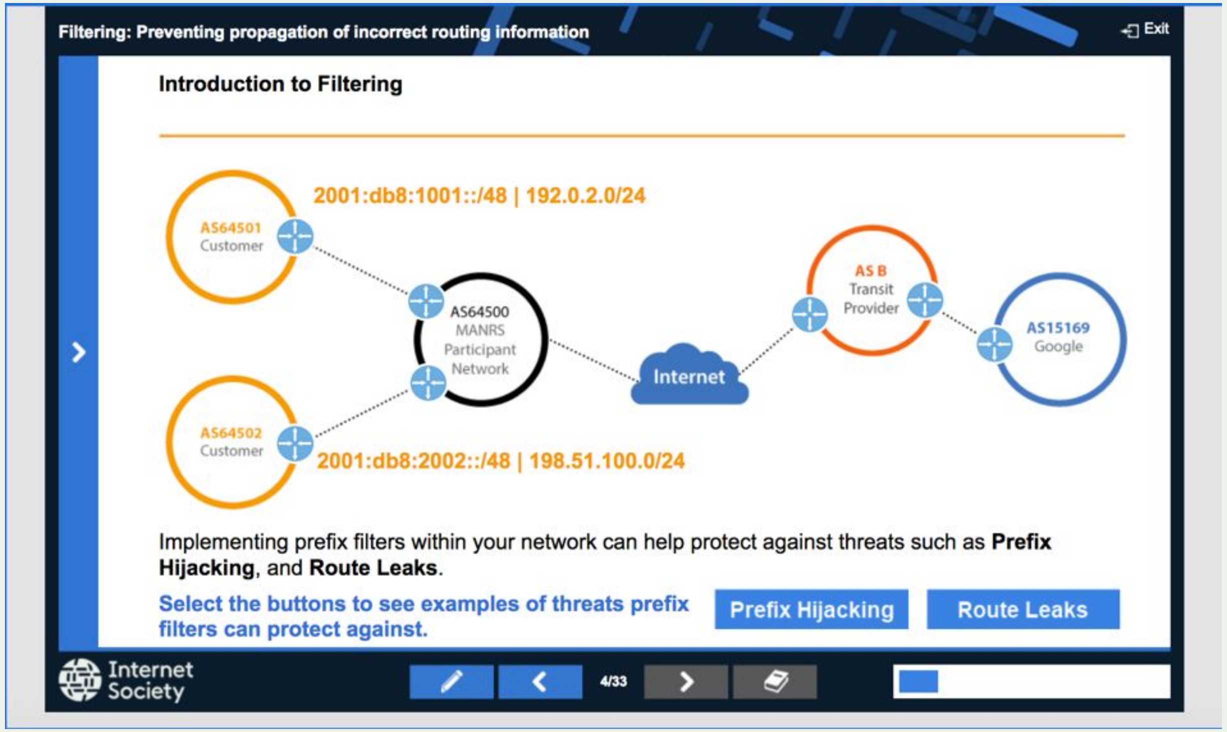

Filtering: Preventing propagation of incorrect routing information

Introduction to Filtering

Implementing prefix filters within your network can help protect against threats such as **Prefix Hijacking**, and **Route Leaks**.

Select the buttons to see examples of threats prefix filters can protect against.

MANRS Hands on Lab (ISOC Hosted)

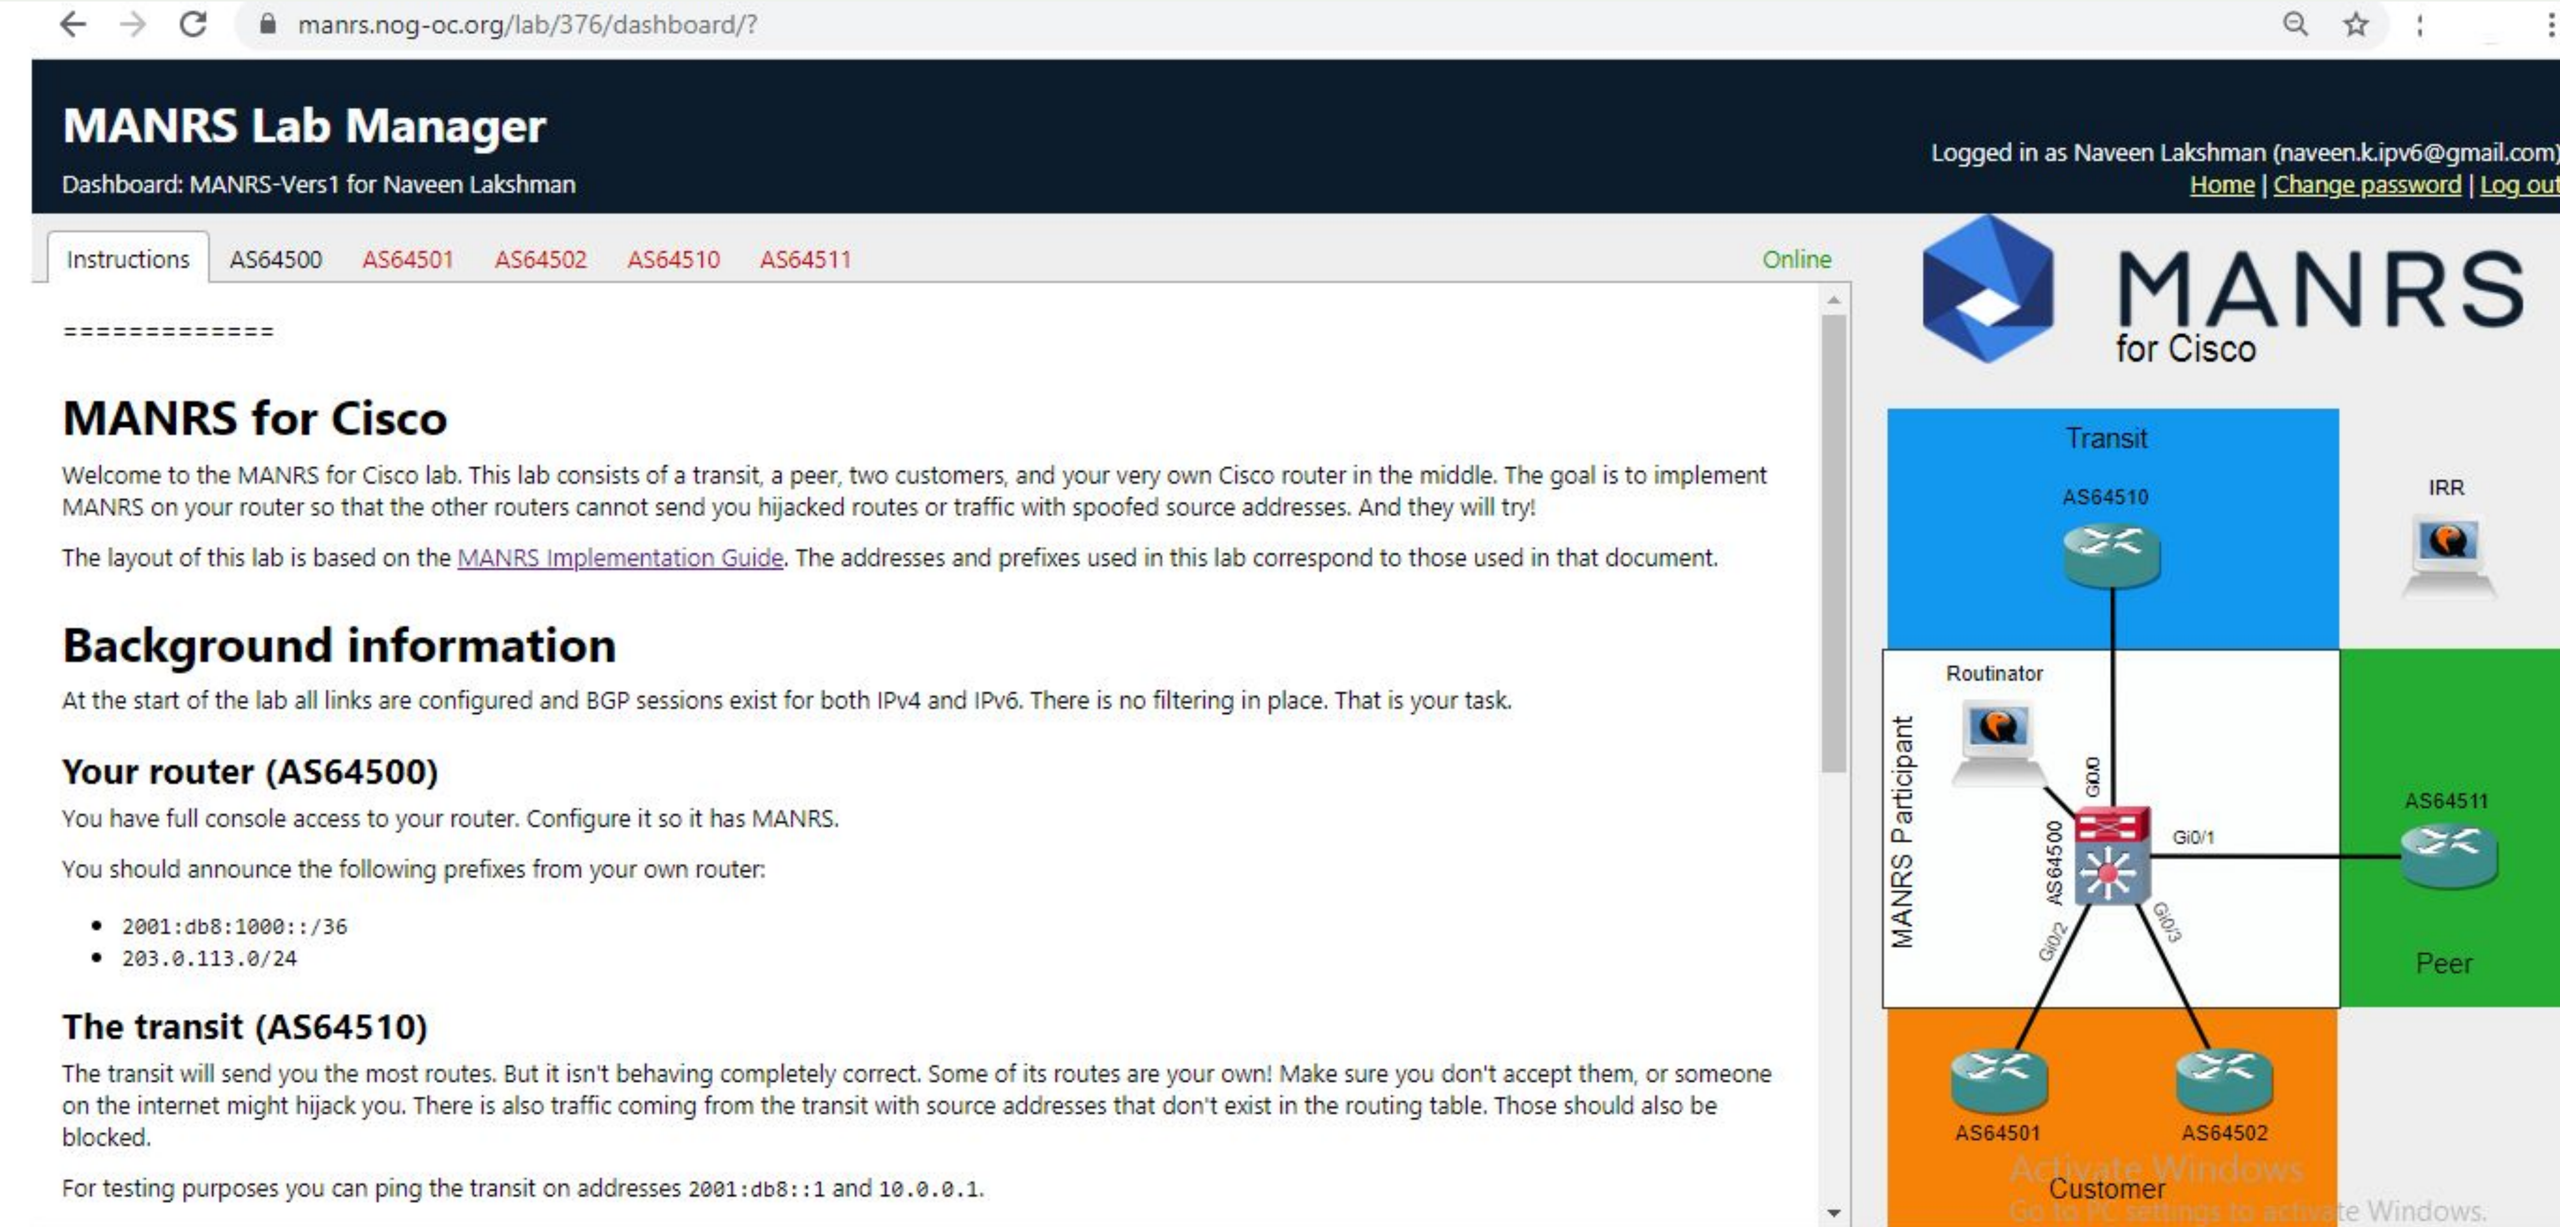

MANRS Lab Manager

Dashboard: MANRS-Vers1 for Naveen Lakshman

Logged in as Naveen Lakshman (naveen.k.ipv6@gmail.com)
Home | Change password | Log out

=============

MANRS for Cisco

Welcome to the MANRS for Cisco lab. This lab consists of a transit, a peer, two customers, and your very own Cisco router in the middle. The goal is to implement MANRS on your router so that the other routers cannot send you hijacked routes or traffic with spoofed source addresses. And they will try!

The layout of this lab is based on the MANRS Implementation Guide. The addresses and prefixes used in this lab correspond to those used in that document.

Background information

At the start of the lab all links are configured and BGP sessions exist for both IPv4 and IPv6. There is no filtering in place. That is your task.

Your router (AS64500)

You have full console access to your router. Configure it so it has MANRS.

You should announce the following prefixes from your own router:

- `2001:db8:1000::/36`
- `203.0.113.0/24`

The transit (AS64510)

The transit will send you the most routes. But it isn't behaving completely correct. Some of its routes are your own! Make sure you don't accept them, or someone on the internet might hijack you. There is also traffic coming from the transit with source addresses that don't exist in the routing table. Those should also be blocked.

For testing purposes you can ping the transit on addresses `2001:db8::1` and `10.0.0.1`.

MANRS Lab Modules

Lab guide is based on https://www.manrs.org/isps/bcop/

Tutorial is a mix of lectures and hands-on lab sessions to deploy MANRS actions based on best current operational practices. Lab runs on dual-stack infrastructure.

MANRS Actions Agenda (Lab Configuration | Cisco IOS Platform)

- Anti-Spoofing (uRPF)
  - BCP38/uRPF Strict Mode
- Filtering (Preventing propagation of incorrect routing information)
  - Specific-prefix outbound filtering of your network to peers and upstreams/transits.
  - Specific-prefix inbound filtering from customers.
  - Specific-prefix Inbound filtering of peers and upstreams to your network.

Lab Guide

https://drive.google.com/file/d/1p5vXJeQSo83PUdv9CsdgSaSPq0SONAoa/view?usp=sharing

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

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

Thank you.

manrs.org