



The impact of an RPKI validator in Bangladesh and Lessons Learned

Md Abdul Awal
awal@nsrc.org

Presented by:
Philip Smith
philip@nsrc.org

#SANOG35

Image source: Internet

Starting with some routing incidents...

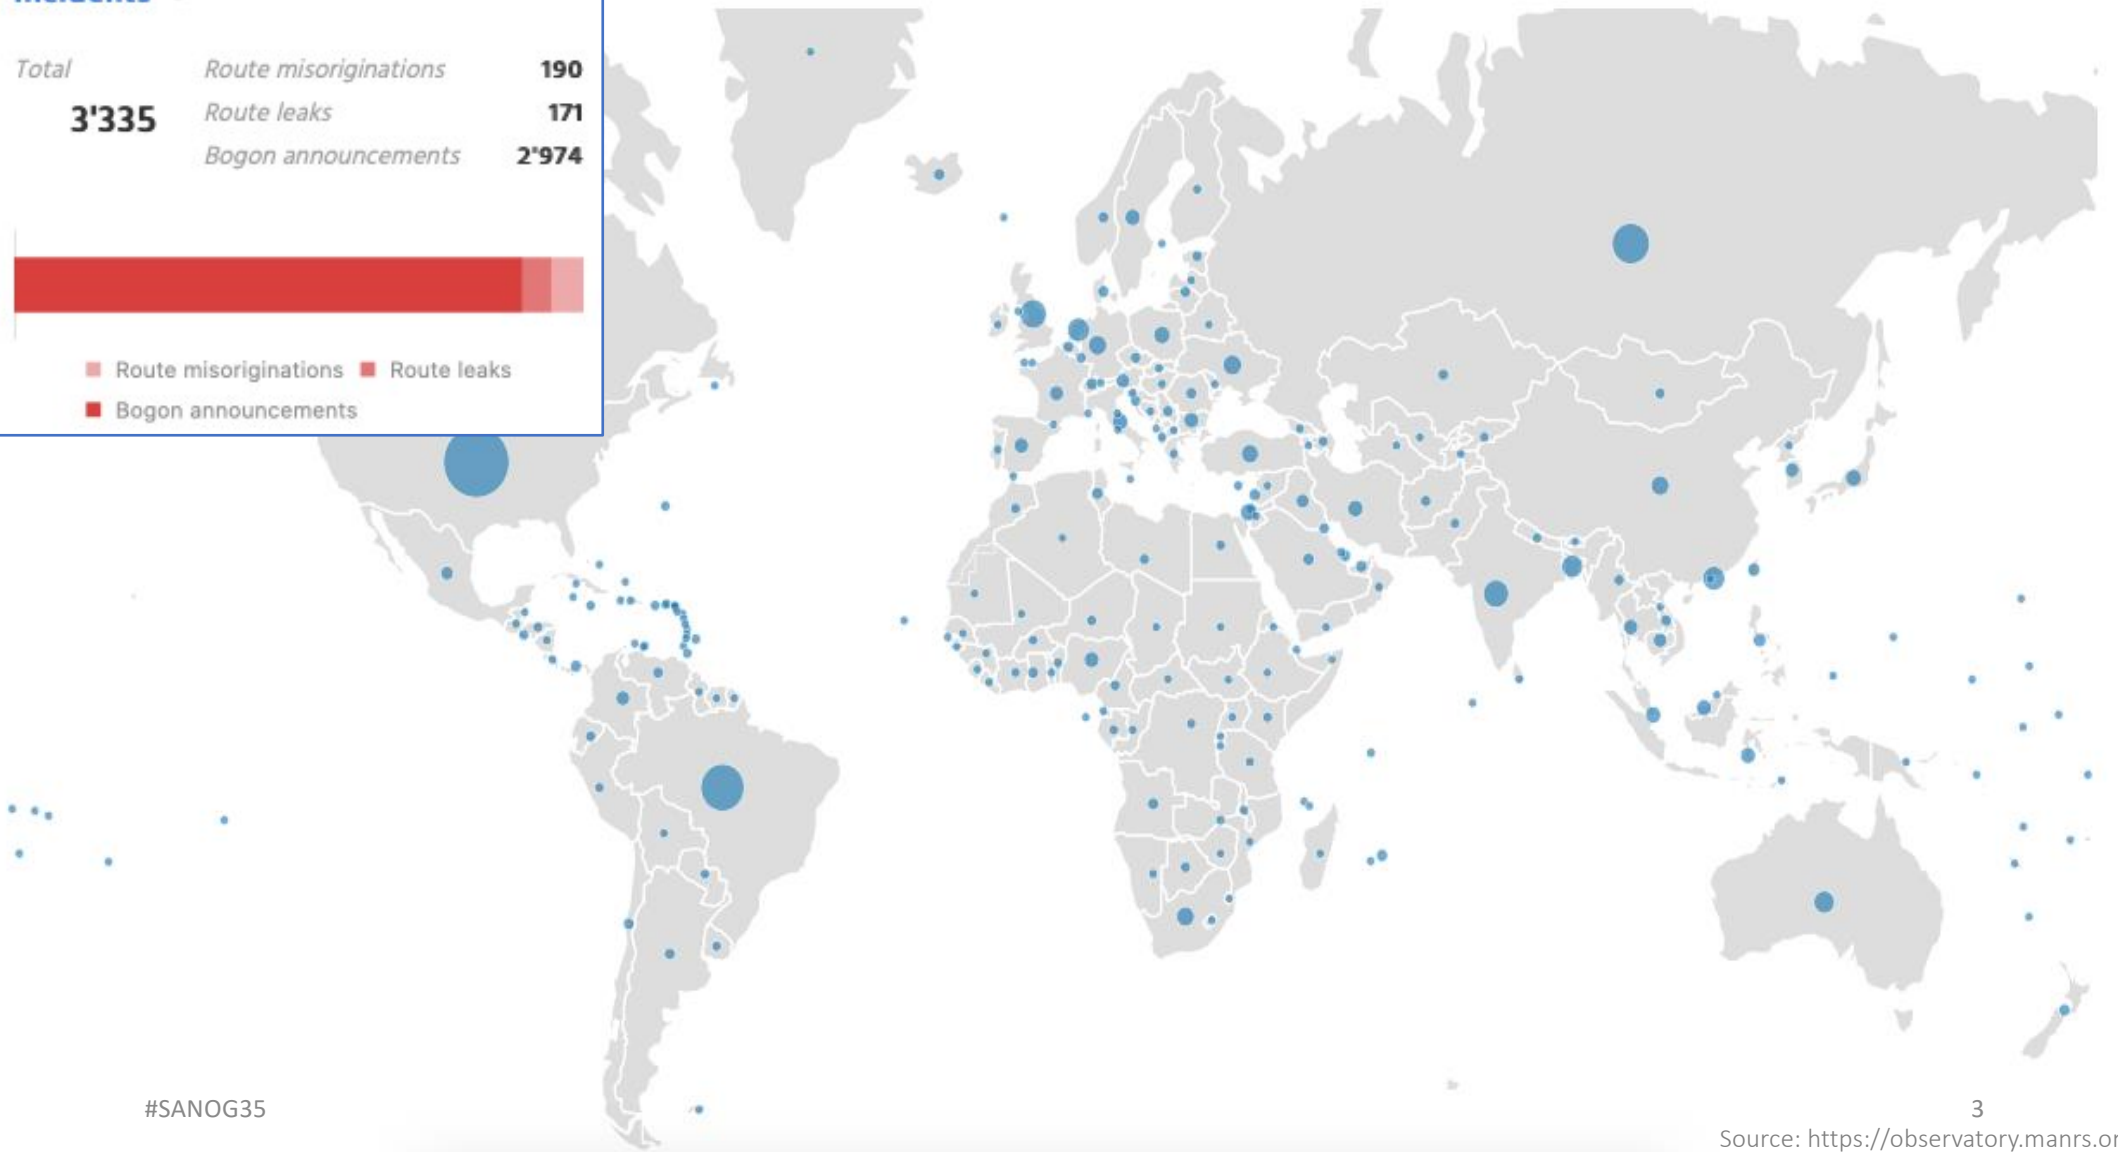
Incidents i

Total	Route misoriginations	190
3'335	Route leaks	171
	Bogon announcements	2'974



- Route misoriginations
- Route leaks
- Bogon announcements

Size: Count | **Incidents** | Culprits | Region: **Country** | UN Regions | UN Sub-Regions | RIR Regions



#SANOG35



- A Prefix is announced by both AS 134599 and AS 133957 (might be unintentional)
- Either AS cloud be closest to different geographic locations
- Legitimate traffic might get blackholed

Announced By		
Origin AS	Announcement	
<u>AS134599</u>	<u>45.118.70.0/24</u>	✓
<u>AS133957</u>	<u>45.118.70.0/24</u>	✓

```
route-views>show ip bgp 45.118.70.0/24 | i 133957
3277 3267 174 58601 138197 138197 138197 138197 138197 133957
24441 3491 3491 3257 58601 138197 138197 138197 138197 138197 133957
3561 209 174 58601 138197 138197 138197 138197 138197 133957
20912 174 58601 138197 138197 138197 138197 138197 138197 133957
852 174 58601 138197 138197 138197 138197 138197 138197 133957
101 101 174 58601 138197 138197 138197 138197 138197 133957
3267 174 58601 138197 138197 138197 138197 138197 138197 133957
3303 3257 58601 138197 138197 138197 138197 138197 138197 133957
54728 20130 23352 3257 58601 138197 138197 138197 138197 138197 133957
```

```
route-views>show ip bgp 45.118.70.0/24 | i 134599
8283 6762 132602 58717 135341 135341 135341 135341 134599 134599 134599 134599 134599 134599 134599 134599 134599
6939 58717 135341 135341 135341 135341 134599 134599 134599 134599 134599 134599 134599 134599 134599 134599
7018 2914 132602 58717 135341 135341 135341 135341 134599 134599 134599 134599 134599 134599 134599 134599
1403 1299 2914 132602 58717 135341 135341 135341 135341 134599 134599 134599 134599 134599 134599 134599 134599
1403 1299 2914 132602 58717 135341 135341 135341 135341 134599 134599 134599 134599 134599 134599 134599 134599
1351 6939 58717 135341 135341 135341 135341 134599 134599 134599 134599 134599 134599 134599 134599 134599 134599
286 6762 132602 58717 135341 135341 135341 135341 134599 134599 134599 134599 134599 134599 134599 134599 134599
57866 6762 132602 58717 135341 135341 135341 135341 134599 134599 134599 134599 134599 134599 134599 134599 134599
4826 58717 135341 135341 135341 135341 134599 134599 134599 134599 134599 134599 134599 134599 134599 134599 134599
```

- Issue informed to the IP owner.
- They removed one announcement
- Created ROA for valid ASN
- Valid announcement visible in the global routing table

Announced By	
Origin AS	Announcement
<u>AS133957</u>	<u>45.118.70.0/24</u>  

```
route-views>show ip bgp 45.118.70.0/24 | i 133957
3277 3267 174 58601 138197 138197 138197 138197 138197 133957
24441 3491 3491 3257 58601 138197 138197 138197 138197 133957
3561 209 174 58601 138197 138197 138197 138197 138197 133957
20912 174 58601 138197 138197 138197 138197 138197 133957
852 174 58601 138197 138197 138197 138197 138197 133957
101 101 174 58601 138197 138197 138197 138197 138197 133957
3267 174 58601 138197 138197 138197 138197 138197 133957
3303 3257 58601 138197 138197 138197 138197 138197 133957
54728 20130 23352 3257 58601 138197 138197 138197 138197 133957
3257 58601 138197 138197 138197 138197 138197 133957
```

```
route-views>sho ip bgp 45.118.70.0/24 | i 134599
route-views>
```

- Client AS didn't update APNIC membership
- Shouldn't have valid prefix allocation with revoked membership
- Transit still announces client's prefixes
- Prefixes marked as BOGON in global routing table

AS Info | Graph v4 | Graph v6 | Prefixes v4 | Prefixes v6 | Peers v4 | Peers v6 | Whois | IRR

AS136555 announces bogons.

Company Website: <http://bkonlinebd.net/>

AS Info | Graph v4 | Graph v6 | Prefixes v4 | Prefixes v6 | Peers v4 | Peers

Prefix	Description
103.67.199.0/24	<input checked="" type="checkbox"/> BK ONLINE
103.92.152.0/24	Asia Pacific Network Information Centre
103.92.153.0/24	Asia Pacific Network Information Centre
103.92.154.0/24	Asia Pacific Network Information Centre
103.92.155.0/24	Asia Pacific Network Information Centre

core3.fmt1.he.net> show ip bgp routes detail 103.92.152.0/24











Status	Network	Next Hop	Metric	LocPrf	Weight	Path	Origin	ROA
BI	103.92.152.0/24	198.32.176.203	15	100	0	9498, 132884, 58717, 136555x5	IGP	?
I	103.92.152.0/24	206.72.210.149	95	100	0	9498, 132884, 58717, 136555x5	IGP	?

Last Update 7h35m37s ago (1 path installed)
Entry cached for another 60 seconds.

2019-10-27 13:18:47 UTC

- The issue has been informed to the transit provider
- Then, they dropped it
- The announcement was removed from global table
- Later on, the AS got membership renewed and has its allocated prefixes back for use

```
route-views>show ip bgp 103.92.152.0/24
% Network not in table
route-views>show ip bgp 103.92.153.0/24
% Network not in table
route-views>show ip bgp 103.92.154.0/24
% Network not in table
route-views>show ip bgp 103.92.155.0/24
% Network not in table
```

Prefix			
<u>103.67.199.0/24</u>			BK ONLINE
<u>103.92.152.0/24</u>			BK ONLINE
<u>103.92.153.0/24</u>			BK ONLINE
<u>103.92.154.0/24</u>			BK ONLINE
<u>103.92.155.0/24</u>			BK ONLINE

- AS 64075 delegated its prefix 103.204.210.0/24 to AS 137842
- AS 137842 announced the prefix
- AS 64075 is also announcing its delegated prefixes as AS 137842 (AS Hijack)
- It's upstream accepting it and further announcing it globally

Announced By		
Origin AS	Announcement	
AS137842	103.204.210.0/24	
AS64075	103.204.210.0/24	

AS137842 MH ONLINE

[Home](#) | [Report](#) | [Export](#) | [Export](#) | [AS](#) | [t](#) | [Routes](#)

[AS Info](#) | [Graph v4](#) | [Prefixes v4](#) | [Peers v4](#) | [Whois](#) | [IRR](#)

















Rank	Description	
1	BDCOM	
2	CIRCLE NETWORK BANGLADESH	

core1.sin1.he.net> show ip bgp routes detail 202.181.16.0/24















Status	Network	Next Hop	Metric	LocPr#	Weight	Path	Origin	RO#
BEx	202.181.16.0/24	74.82.51.74	0	140	0	132602, 10075, 134371, 137842	IGP	
E	202.181.16.0/24	27.111.228.140	0	100	0	10075, 134371, 137842	IGP	
E	202.181.16.0/24	27.111.228.140	0	100	0	64075, 134371, 137842	IGP	
E	202.181.16.0/24	27.111.228.140	0	100	0	10075, 134371, 137842	IGP	
E	202.181.16.0/24	27.111.228.40	0	100	0	9498, 10075, 134371, 137842	IGP	
E	202.181.16.0/24	27.111.229.76	0	100	0	132602, 10075, 134371, 137842	IGP	
E	202.181.16.0/24	65.49.109.182	0	100	0	9498, 10075, 134371, 137842	IGP	
E	202.181.16.0/24	27.111.228.111	0	100	0	6762, 132602, 10075, 134371, 137842	IGP	
E	202.181.16.0/24	74.82.48.70	0	100	0	3491, 9498, 10075, 134371, 137842	IGP	
E	202.181.16.0/24	216.218.221.142	0	100	0	4637, 9498, 10075, 134371, 137842	IGP	

Last Update 4d5h31m58s ago (1 path installed)
Entry cached for another 60 seconds.

- A /64 IPv6 prefix announced in global routing table
- Most specific announcement in global table is /48
- A /64 should never be in global routing table

Prefix		
<u>2405:1500::/32</u>		
<u>2405:1500::/48</u>		
<u>2405:1500:0:8::/64</u>		
<u>2405:1500:30::/48</u>		
<u>2405:1500:40::/48</u>		
<u>2405:1500:60::/48</u>		
<u>2405:1500:70::/48</u>		
<u>2405:1500:80::/48</u>		

- The issue was informed to the AS
- The announcement has been removed

Prefix		
<u>2405:1500::/32</u>		
<u>2405:1500::/48</u>		
<u>2405:1500:30::/48</u>		
<u>2405:1500:40::/48</u>		
<u>2405:1500:60::/48</u>		
<u>2405:1500:70::/48</u>		
<u>2405:1500:80::/48</u>		

Somebody is announcing non-routable prefixes in the global BGP table.

Prefix	
<u>103.82.232.0/24</u>	✓
<u>103.96.230.0/23</u>	✓
<u>103.96.232.0/24</u>	✓
<u>123.253.228.0/22</u>	✓
<u>172.18.3.0/24</u>	

Prefix	
<u>103.82.232.0/24</u>	🔑 ✓
<u>103.96.230.0/23</u>	🔑 ✓
<u>103.96.232.0/24</u>	🔑 ✓
<u>123.253.228.0/22</u>	🔑 ✓

The announcement has been removed once the issue was informed to them

- Previously, AS 136909 used transit from AS 24342 using static routing.
- AS 24342 announced prefixes of AS 136909 in global BGP table on their behalf.
- Later, AS 136909 stated doing BGP but AS 24342 still didn't stop the announcement.

```

route-views>sh ip bgp 103.98.200.0/24 | i 24342
49788 12552 4637 9498 58601 24342 24342 24342 24342 24342 24342
3303 2914 58601 24342 24342 24342 24342 24342 24342 24342
3561 209 3356 2914 58601 24342 24342 24342 24342 24342 24342
3267 3356 2914 58601 24342 24342 24342 24342 24342 24342
24441 3491 3491 9498 58601 24342 24342 24342 24342 24342
3277 3267 3356 2914 58601 24342 24342 24342 24342 24342
20912 174 9498 58601 24342 24342 24342 24342 24342 24342
852 3491 9498 58601 24342 24342 24342 24342 24342 24342
6939 58601 24342 24342 24342 24342 24342 24342 24342
101 101 11164 7473 9498 58601 24342 24342 24342 24342 24342
1351 6939 58601 24342 24342 24342 24342 24342 24342
54728 20130 6939 58601 24342 24342 24342 24342 24342 24342
3257 7473 9498 58601 24342 24342 24342 24342 24342 24342
3333 1273 2914 58601 24342 24342 24342 24342 24342 24342
6079 4637 9498 58601 24342 24342 24342 24342 24342 24342
701 2914 58601 24342 24342 24342 24342 24342 24342 24342
1239 2914 58601 24342 24342 24342 24342 24342 24342 24342
3549 3356 6453 58601 24342 24342 24342 24342 24342 24342
1221 4637 9498 58601 24342 24342 24342 24342 24342 24342
4901 6079 4637 9498 58601 24342 24342 24342 24342 24342
4826 1221 4637 9498 58601 24342 24342 24342 24342 24342
53767 174 174 9498 58601 24342 24342 24342 24342 24342
2497 2914 58601 24342 24342 24342 24342 24342 24342 24342

```

AS Info	Graph v4	Prefixes v4	Peers v4
Prefix			
		<u>103.98.200.0/22</u>	✓
		<u>103.98.200.0/24</u>	✓
		<u>103.98.201.0/24</u>	✓
		<u>103.98.202.0/24</u>	✓
		<u>103.98.203.0/24</u>	✓

AS Info	Graph v4	Prefixes v4	Peers v4
Prefix			
		<u>103.98.200.0/24</u>	✓
		<u>103.98.201.0/24</u>	✓
		<u>115.127.0.0/17</u>	✓
		<u>115.127.0.0/18</u>	✓
		<u>115.127.0.0/19</u>	✓
		<u>115.127.0.0/20</u>	✓
		<u>115.127.0.0/24</u>	✓

- AS 24342 has been informed to stop announcing client's prefixes
- Client AS 136909 has signed their prefixes
- Issue resolved.

```
route-views>show ip bgp 103.98.200.0/24 | i 24342
7660 2516 6453 58601 24342 24342 24342 136909
286 6453 58601 24342 24342 24342 136909
852 3491 9498 58601 24342 24342 24342 136909
6079 4637 9498 58601 24342 24342 24342 136909
1351 9498 58601 24342 24342 24342 136909
54728 20130 6939 58601 24342 24342 24342 136909
3333 1257 6453 58601 24342 24342 24342 136909
3267 6461 7473 17494 58601 24342 24342 24342 136909
1403 6453 58601 24342 24342 24342 136909
1403 6453 58601 24342 24342 24342 136909
8283 6453 58601 24342 24342 24342 136909
57866 6461 7473 17494 58601 24342 24342 24342 136909
49788 12552 9498 58601 24342 24342 24342 136909
3277 3267 6461 7473 17494 58601 24342 24342 24342 136909
```

AS Info	Graph v4	Prefixes v4	Peers v4
Prefix			
103.98.200.0/22			
103.98.200.0/24			
103.98.201.0/24			
103.98.202.0/24			
103.98.203.0/24			

AS Info	Graph v4	Prefixes v4	Peers v4	Whois
Prefix				
115.127.0.0/17				
115.127.0.0/18				
115.127.0.0/19				
115.127.0.0/20				
115.127.0.0/24				
115.127.1.0/24				
115.127.2.0/24				
115.127.3.0/24				
115.127.4.0/24				

- AS 136901 got an allocation of /22.
- They announce part of its prefix (not the whole), e.g. /23 is announced but the other /23 is not.
- Opportunists can try to use the unannounced /23 for unauthorized activities.

Prefix	
<u>103.98.64.0/24</u>	✓
<u>103.98.65.0/24</u>	✓

```
awal@Awals-MacBook-Air ~> whois -h whois.radb.net 103.98.64.0/22
route:          103.98.64.0/22
origin:         AS136901
```

- AS 137515 announced BCC's prefix 103.48.17.0/24 in a NIX (Prefix hijack)
- Important government services became unavailable to the citizens
- Costs our time to fix.

Prefix	Description
43.229.12.0/22	Bangladesh Computer Council
43.229.15.0/24	Bangladesh Computer Council
103.48.16.0/22	Bangladesh Computer Council
114.130.54.0/23	Bangladesh Computer Council
114.130.116.0/22	Bangladesh Computer Council
180.211.213.0/24	Bangladesh Computer Council

```

sadbm@PXXN-EDGE-3K> show route 103.48.17.0
inet.0: 801403 destinations, 6284521 routes (787268 active, 0 holddown, 3604693 hidden)
+ = Active Route, - = Last Active, * = Both

103.48.17.0/24    *[BGP/170] 00:01:34, localpref 100
                 AS path: 132602 56032 137515 I, validation-state: unverified
                 > to 103.16.155.45 via et-0/0/1.517
                 [BGP/170] 00:01:13, localpref 100
                 AS path: 132602 56032 137515 I, validation-state: unverified
                 > to 103.16.155.41 via et-0/0/2.520
                 [BGP/170] 00:01:34, localpref 100
                 AS path: 132602 56032 137515 I, validation-state: unverified
                 > to 163.47.83.29 via et-0/0/3.73

```

What they reply about it? Funny but they really did

- We mistakenly announced the prefix
- We do not manually check our clients' APNIC membership status
- The client is very close to us and well trustworthy, so we never required to check their announcements
- We don't do prefix or AS filtering for our clients
- Forgot to stop announcing the prefix after it's delegated to another AS



Image source: Internet

How we ensure the routing hygiene manually?

Nightmare...




Image source: Internet



Image source: Internet

RPKI is about 2 things: ROA and ROV

1

Signing prefixes

a.k.a. creating ROAs

Prefix: 2401:ED80::/32

Origin AS: AS63932

MSA: /32

ROA: Enabled

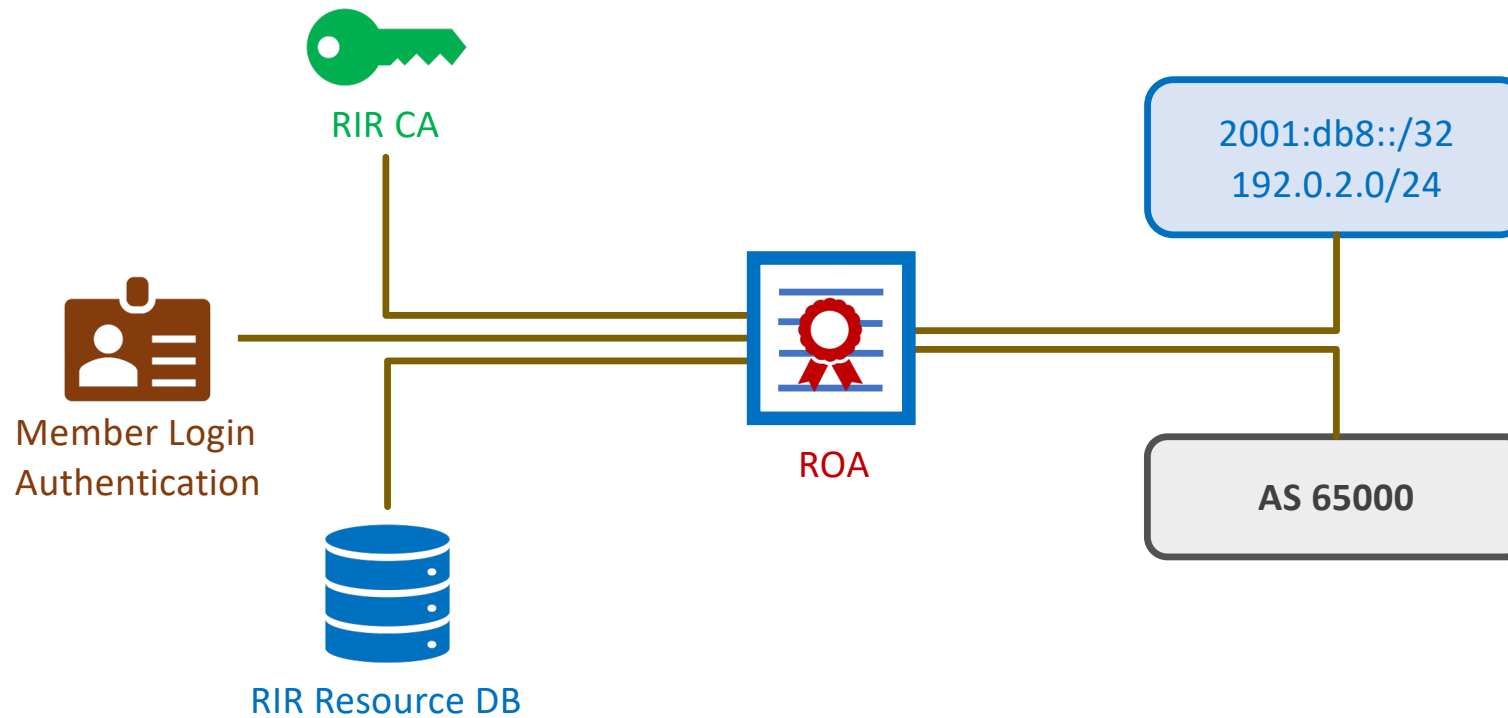
Whois: Enabled

Actions: [Update whois ↗](#)

[Cancel](#) [Submit](#)

Route	Origin AS	ROA status	Whois status
103.48.16.0/22	AS63932	<input checked="" type="checkbox"/>	<input checked="" type="checkbox"/>
2401:ED80::/32	AS63932	<input checked="" type="checkbox"/>	<input checked="" type="checkbox"/>
43.229.12.0/22	AS63932	<input checked="" type="checkbox"/>	<input checked="" type="checkbox"/>
43.229.15.0/24	AS63932	<input checked="" type="checkbox"/>	<input checked="" type="checkbox"/>

RPKI is about 2 things: ROA and ROV



RPKI is about 2 things: ROA and ROV

2

Validating ROAs


a.k.a Route Origin Validation

```
RPKI server is 172.20.22.4, port 3323
RPKI current state: Established, Age: 01m46s
  Local host: 172.19.19.1, Local port: 60462
  Remote host: 172.20.22.4, Remote port: 3323
Refresh time : 900
Aging time : 1800
Session ID : 59791
Serial number : 10
Session Statistics:
  IPv4 record : 98101
  IPv6 record : 16506
```

RPKI validation codes: V - valid, I - invalid, N - not-found

```
Total Number of Routes: 1005452
Network          NextHop          MED          LocPrf    PrefVal Path/Ogn
* V 1.0.0.0/24    103.9.114.161   0            58601 13335i
* V              114.130.31.5    0            17806 17494 6453 13335i
* N 1.0.4.0/22    103.9.114.161   0            58601 4826 38803 56203i
* N 1.0.4.0/24    103.9.114.161   0            58601 4826 38803 56203i
* N 1.0.5.0/24    103.9.114.161   0            58601 4826 38803 56203i
* N 1.0.6.0/24    103.9.114.161   0            58601 4826 38803 56203i
* N 1.0.7.0/24    103.9.114.161   0            58601 4826 38803 56203i
* N 1.0.16.0/24   103.9.114.161   0            58601 2914 2519i
* N 1.0.64.0/18   103.9.114.161   0            58601 6453 2497 7670 18144
```





RPKI is about 2 things: ROA and ROV



Why Create ROA?



To ensure the authenticity of your IP resources and help others verify it if requires

		Announced By	
Origin AS	Announcement		
<u>AS134599</u>	<u>45.118.70.0/24</u>		
<u>AS133957</u>	<u>45.118.70.0/24</u>		

So that your IP resources are not knowingly or unknowingly used or abused by anyone

Why Deploy ROV?



To build and maintain a secure and trustworthy global routing infrastructure



To validate BGP routes and identify the authorized originator of the prefix

RPKI Validation in NDC and subsequent impact

RPKI Validation at National Data Center

- NDC declared to drop invalids since Dec 1, 2019
- Bangladesh has more than 700 active ASN
- BD ROA stats in Sep 2019:
 - Valid – 29%
 - Unknown – 69%
 - Invalid – 2%
- Need to find out ASNs who are getting impact
 - BD ASNs are easy to reach for obvious reason
 - How about the global ASNs?

Awareness Before the ROV

- We helped others to create and/or fix their ROAs
 - LEAs, Police, Special forces
 - Govt. Organizations
 - IXPs
 - Banks and Financial Organizations
 - Transit providers
 - ISPs
 - Data Centers

The impact of awareness campaign

Prefix		
<u>123.49.0.0/18</u>		
<u>123.49.14.0/24</u>		
<u>123.49.16.0/20</u>		
<u>123.49.29.0/24</u>		
<u>123.49.30.0/24</u>		
<u>123.49.31.0/24</u>		
<u>123.49.47.0/24</u>		
<u>180.211.128.0/17</u>		
<u>180.211.201.0/24</u>		
<u>180.211.214.0/24</u>		
<u>180.211.215.0/24</u>		
<u>203.112.192.0/19</u>		
<u>203.112.194.0/24</u>		
<u>209.58.24.0/24</u>		



Prefix		
<u>123.49.0.0/18</u>		
<u>123.49.14.0/24</u>		
<u>123.49.16.0/20</u>		
<u>123.49.29.0/24</u>		
<u>123.49.30.0/24</u>		
<u>123.49.31.0/24</u>		
<u>123.49.47.0/24</u>		
<u>180.211.128.0/17</u>		
<u>180.211.201.0/24</u>		
<u>180.211.214.0/24</u>		
<u>180.211.215.0/24</u>		
<u>203.112.192.0/19</u>		
<u>203.112.194.0/24</u>		

Source: <https://bgp.he.net>

The impact of awareness campaign

AS Info Graph v4 Prefixes v4 Peers v4


Prefix		
<u>43.255.20.0/22</u>		
<u>43.255.20.0/24</u>		
<u>43.255.21.0/24</u>		
<u>43.255.22.0/24</u>		
<u>103.239.4.0/22</u>		
<u>103.239.4.0/24</u>		
<u>103.239.5.0/24</u>		
<u>103.239.6.0/24</u>		
<u>103.239.7.0/24</u>		




AS Info Graph v4 Prefixes v4 Peers v4

Prefix		
<u>43.255.20.0/24</u>		
<u>43.255.21.0/24</u>		
<u>43.255.22.0/24</u>		
<u>103.239.4.0/24</u>		
<u>103.239.5.0/24</u>		
<u>103.239.6.0/24</u>		


The impact of awareness campaign



Oct 2019




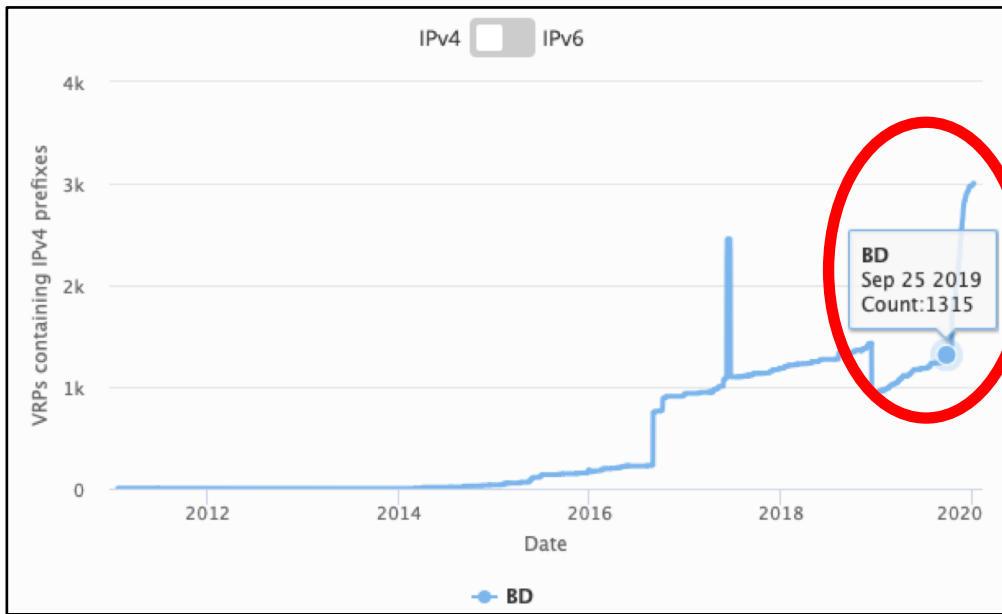
Nov 2019



Dec 2019

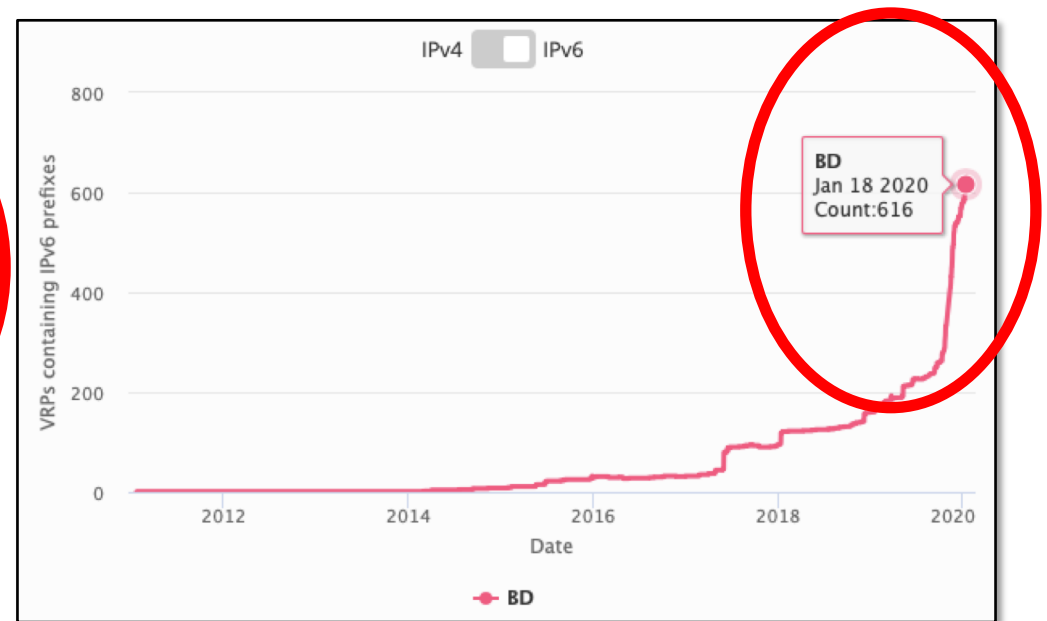
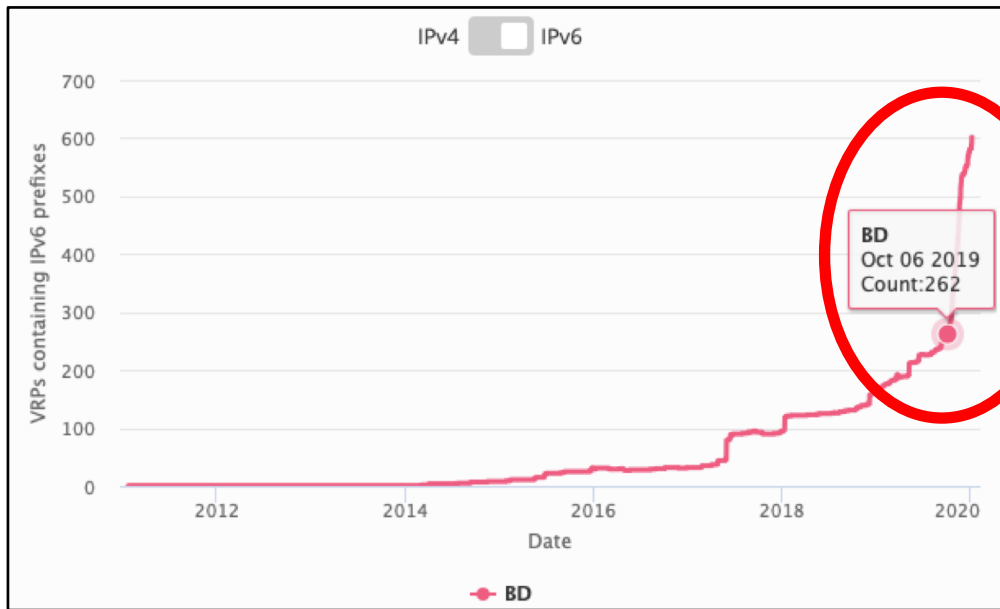
Source: <https://observatory.manrs.org>

The impact of awareness campaign



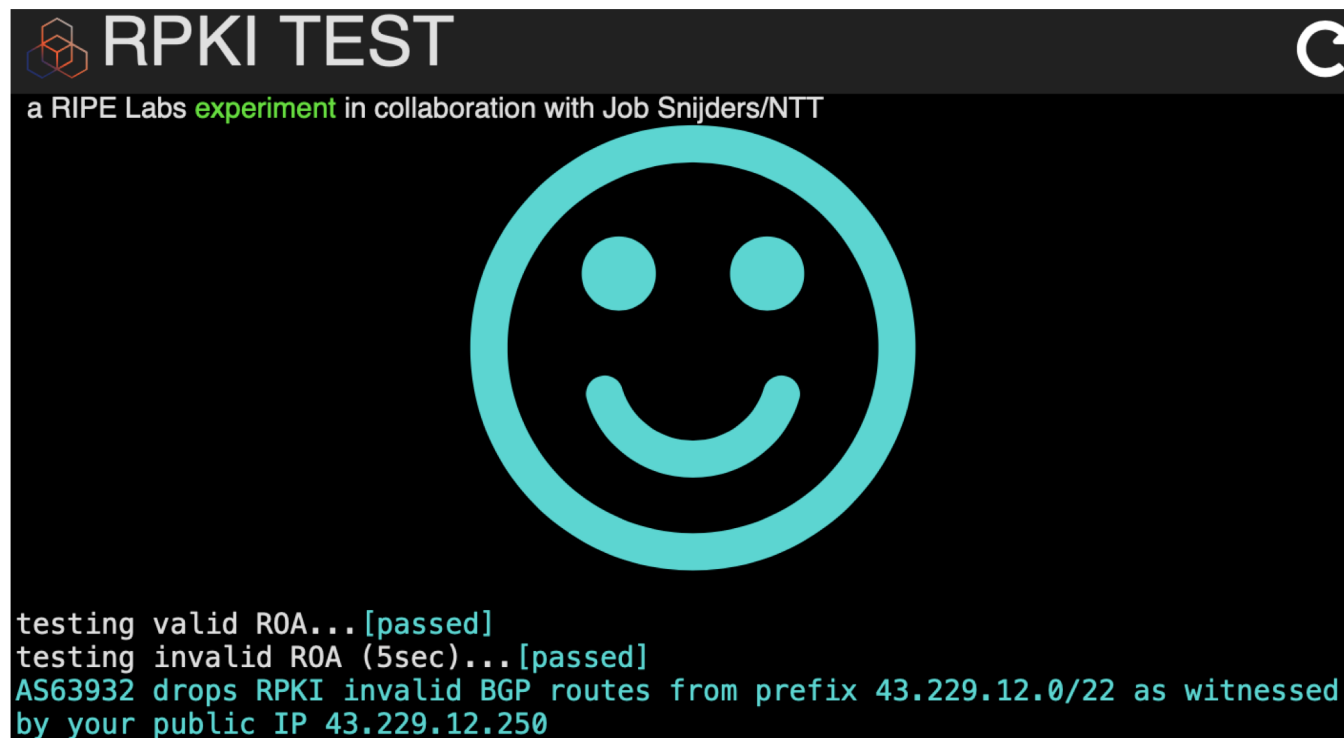
Source: <https://stat.ripe.net/BD#tabId=routing>

The impact of awareness campaign



Source: <https://stat.ripe.net/BD#tabId=routing>

And, finally NDC drops invalids since Dec 1, 2020



The image shows a terminal window titled "RPKI TEST" with a refresh icon in the top right corner. Below the title, it says "a RIPE Labs experiment in collaboration with Job Snijders/NTT". In the center of the terminal is a large cyan smiley face. At the bottom, the terminal output shows the following text:


```
testing valid ROA...[passed]
testing invalid ROA (5sec)...[passed]
AS63932 drops RPKI invalid BGP routes from prefix 43.229.12.0/22 as witnessed
by your public IP 43.229.12.250
```

More organizations to start ROV

- BdREN (AS 63961) - Jan, 2020
- Summit Communications (AS 58717) – Jan 2019
- BD Link (AS 58668) – Jan, 2020
- Link3 Technologies (AS 23688) - TBD
- Cybergate (AS 58599) - TBD


How about South Asia and Pakistan?

RPKI Status in South Asia



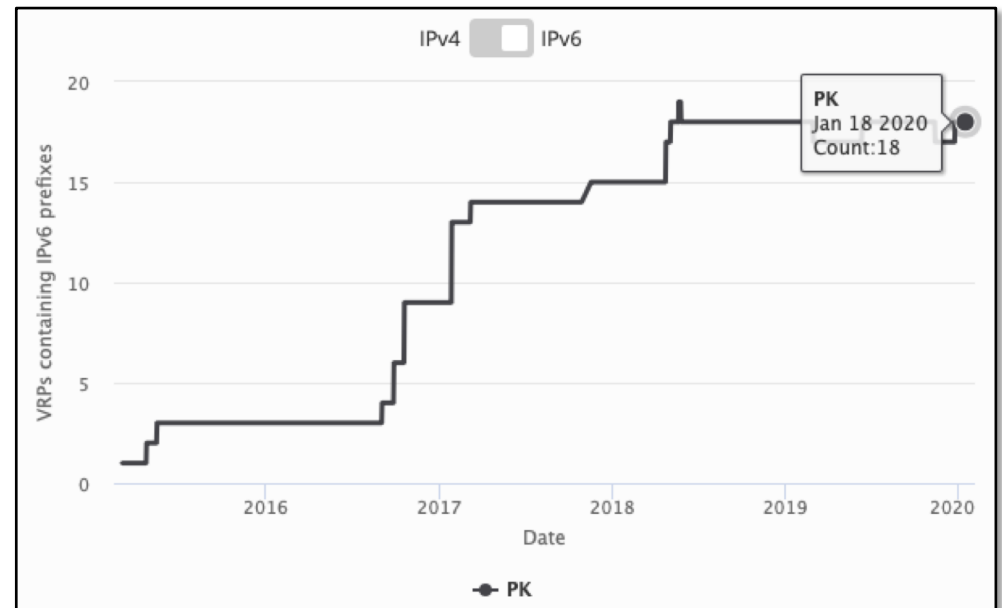

Source: <https://observatory.manrs.org>

RPKI Status in Pakistan



Source: <https://observatory.manrs.org>

RPKI Status in Pakistan



Source: <https://ripe.net/analyse/statistics/ripestat>

- Pakistan has about **1 % INVALID** ROAs
 - We need to eliminate it completely
- **79% UNKNOWN** – That's big
 - Please sign the unsigned!

And, think about doing
Route Origin Validation (ROV)

Considerations about ROA and ROV

Creating ROA

Prefix: 103.48.16.0/22

Origin AS: AS63932

MSA: /22

ROA: Enabled

Whois: Enabled

Actions: [Update whois](#)

VS

Prefix: 103.48.16.0/22

Origin AS: AS63932

MSA: /24

ROA: Enabled

Whois: Enabled

Sub-routes

Show 10 entries Search:

Route	Managed	Actions
103.48.16.0/22	<input type="checkbox"/> Enabled	Update whois
103.48.16.0/23	<input type="checkbox"/> Enabled	Update whois
103.48.16.0/24	<input type="checkbox"/> Enabled	Update whois
103.48.17.0/24	<input type="checkbox"/> Enabled	Update whois
103.48.18.0/23	<input type="checkbox"/> Enabled	Update whois
103.48.18.0/24	<input type="checkbox"/> Enabled	Update whois
103.48.19.0/24	<input type="checkbox"/> Enabled	Update whois

- Not a good idea to create ROAs up to /24 if not announced in BGP
- Better to create ROAs for specific prefixes that are announced in BGP

Creating ROA

Route	Origin AS	ROA status	Whois status
103.48.16.0/22	AS63932	⊙	⊙
2401:ED80::/32	AS63932	⊙	⊙
43.229.12.0/22	AS63932	⊙	⊙
43.229.15.0/24	AS63932	⊙	⊙



Route	Origin AS	ROA status	Whois status
103.48.16.0/22 +	AS63932	⊙	⊙
2401:ED80::/32 +	AS63932	⊙	⊙
43.229.12.0/22 +	AS63932	⊙	⊙

```
awal@Awals-MacBook-Air ~-> whois -h whois.bgpmon.net " --roa 63932 103.48.16.0/22"
0 - Valid
-----
ROA Details
-----
Origin ASN:      AS63932
Not valid Before: 2019-12-06 18:22:35
Not valid After: 2021-01-31 00:00:00 Expires in 1y55d2h39m25.6000000014901s
Trust Anchor:    rpki.apnic.net
Prefixes:        103.48.16.0/22 (max length /22)
                  43.229.12.0/22 (max length /22)
                  2401:ed80::/32 (max length /32)
                  43.229.15.0/24 (max length /24)
```

```
awal@Awals-MacBook-Air ~-> whois -h whois.bgpmon.net " --roa 63932 103.48.16.0/22"
0 - Valid
-----
ROA Details
-----
Origin ASN:      AS63932
Not valid Before: 2019-10-15 10:22:03
Not valid After: 2020-01-31 00:00:00 Expires in 89d6h14m1s
Trust Anchor:    rpki.apnic.net
Prefixes:        2401:ed80::/32 (max length /36)
                  103.48.16.0/22 (max length /24)
                  43.229.12.0/22 (max length /24)
```

Creating ROA

```
awal@Awals-MacBook-Air ~> whois -h whois.bgpmon.net " --roa 17494 103.110.212.0/22"
0 - Valid
-----
ROA Details
-----
Origin ASN:      AS17494
Not valid Before: 2019-11-12 05:41:59
Not valid After:  2020-10-31 00:00:00 Expires in 323d5h27m43s
Trust Anchor:    rpki.apnic.net
Prefixes:        2407:5000:88::/48 (max length /48)
                  203.112.192.0/19 (max length /24)
                  2407:5000::/32 (max length /40)
                  103.110.212.0/22 (max length /24)
                  123.49.0.0/18 (max length /24)
                  180.211.128.0/17 (max length /24)
awal@Awals-MacBook-Air ~> whois -h whois.bgpmon.net " --roa 45588 103.110.212.0/22"
2 - Not Valid: Invalid Origin ASN, expected 17494
```



```
awal@Awals-MacBook-Air ~> whois -h whois.bgpmon.net " --roa 17494 123.49.0.0/18"
0 - Valid
-----
ROA Details
-----
Origin ASN:      AS17494
Not valid Before: 2019-11-12 05:41:59
Not valid After:  2020-10-31 00:00:00 Expires in 323d5h31m22s
Trust Anchor:    rpki.apnic.net
Prefixes:        2407:5000:88::/48 (max length /48)
                  203.112.192.0/19 (max length /24)
                  2407:5000::/32 (max length /40)
                  103.110.212.0/22 (max length /24)
                  123.49.0.0/18 (max length /24)
                  180.211.128.0/17 (max length /24)
awal@Awals-MacBook-Air ~> whois -h whois.bgpmon.net " --roa 45588 123.49.0.0/18"
0 - Valid
-----
ROA Details
-----
Origin ASN:      AS45588
Not valid Before: 2019-11-12 05:42:00
Not valid After:  2020-10-31 00:00:00 Expires in 323d5h31m13s
Trust Anchor:    rpki.apnic.net
Prefixes:        123.49.0.0/18 (max length /24)
                  2407:5000:88::/48 (max length /48)
                  2407:5000::/32 (max length /40)
                  180.211.128.0/17 (max length /24)
                  203.112.192.0/19 (max length /24)
```

You may sign same prefix with multiple ASNs but do if you really really have to

Doing ROV

	Network	NextHop
*>	I 45.123.40.0/23	103.125.54.54
*>	I 103.63.158.0/23	103.125.54.54
*>	I 103.73.104.0/22	103.125.54.73
*	I 103.73.105.0/24	103.125.54.73
*	I 103.73.106.0/24	103.125.54.73
*	I 103.73.107.0/24	103.125.54.73
*>	I 103.87.249.0/24	103.125.54.54
*	I 103.88.234.0/24	103.125.54.54
*	I 103.114.39.0/24	103.125.54.53
*	I 103.121.60.0/24	103.125.54.54
*>	I 103.123.168.0/24	103.125.54.54
*	I 103.123.169.0/24	103.125.54.54
*>	I 103.126.20.0/22	103.125.54.54
*>	I 103.126.20.0/23	103.125.54.54
*>	I 103.126.22.0/23	103.125.54.54
*	I 103.130.114.0/24	103.125.54.54
*	I 103.133.201.0/24	103.125.54.73
*	I 103.134.56.0/24	103.125.54.54
*	I 103.134.57.0/24	103.125.54.54

VS

	Network	NextHop
I	45.123.40.0/23	103.125.54.54
I	103.63.158.0/23	103.125.54.54
I	103.73.104.0/22	103.125.54.73
I	103.73.105.0/24	103.125.54.73
I	103.73.106.0/24	103.125.54.73
I	103.73.107.0/24	103.125.54.73
I	103.87.249.0/24	103.125.54.54
I	103.88.234.0/24	103.125.54.54
I	103.114.39.0/24	103.125.54.53
I	103.121.60.0/24	103.125.54.54
I	103.123.168.0/24	103.125.54.54
I	103.123.169.0/24	103.125.54.54
I	103.126.20.0/22	103.125.54.54
I	103.126.20.0/23	103.125.54.54
I	103.126.22.0/23	103.125.54.54
I	103.130.114.0/24	103.125.54.54
I	103.133.201.0/24	103.125.54.73
I	103.134.56.0/24	103.125.54.54
I	103.134.57.0/24	103.125.54.54

Validation with allowing invalids for BGP best path selection

Validation without allowing invalids for BGP best path selection

ROA for Small ISPs and Enterprises

- Have own Internet resources?
 - Creating ROA is straightforward using RIR's resource management portal
- Got assignment for LIR?
 - Have public ASN?
 - Ask the LIR to create ROA with your ASN and verify
 - Don't have public ASN?
 - Ask the LIR to create ROA for the assigned prefix and verify

ROV for Small ISPs and Enterprises

- Have BGP with transits and peers?
 - Receive full routes from neighbors?
 - Implementing ROV using validator cache is straightforward
 - Receive partial routes with default from neighbors?
 - Ask transits to do ROV for you
 - And, implement ROV using validator cache to validate peer and IX routes
 - Receive default route only
 - ROV wouldn't fit, however, you may ask transits to do ROV in their routers 😊
- Have static routing with transits?
 - ROV wouldn't fit, however, you may ask transits to do ROV in their routers 😊

Still thinking why we need ROA and ROV?

- Check the issues discussed in first couple slides
- Reduce the opportunity of routing incidents, prefix hijacks, route leaks, DDoS, outages
- You wouldn't want to be a target of those incidents
- Help improve global routing infrastructure security
- Help each other to maintain routing hygiene
- We are engineers working hard to make Internet better, remember?

We all can help improve global routing security

- Create/fix ROAs for your prefixes
- If you are a transit provider, ask your clients to do the same
- If you're receiving BGP full route, implement ROV
- Share this among other colleagues in the community
- Help others fix their ROAs



Image source: Internet

Let's check your own ASN

- Go to <https://bgp.he.net> , search for your AS number and check v4 and v6 prefixes

Or,

- Use whois on unix terminal:

```
whois -h whois.bgpmon.net " --roa ASN Prefix"
```

- If you find issues with ROA, please fix it

<https://blog.apnic.net/2019/09/11/how-to-creating-rpki-roas-in-myapnic/>

HURRICANE ELECTRIC
INTERNET SERVICES

63932 Search

AS63932 Bangladesh Computer Council

AS Info Graph v4 Graph v6 Prefixes v4 Prefixes v6 Peers v4 Peers v6

Prefix	Description
43.229.12.0/22	Bangladesh Computer Council
43.229.15.0/24	Bangladesh Computer Council
103.48.16.0/22	Bangladesh Computer Council
114.130.54.0/23	Bangladesh Computer Council
114.130.116.0/22	Bangladesh Computer Council
180.211.213.0/24	Bangladesh Computer Council

```
awal@Awals-MacBook-Air -> whois -h whois.bgpmon.net " --roa 63932 103.48.16.0/22"
0 - Valid
-----
ROA Details
-----
Origin ASN:      AS63932
Not valid Before: 2019-12-06 18:22:35
Not valid After:  2021-01-31 00:00:00 Expires in 1y55d2h39m25.6000000014901s
Trust Anchor:    rpki.apnic.net
Prefixes:        103.48.16.0/22 (max length /22)
                  43.229.12.0/22 (max length /22)
                  2401:ed80::/32 (max length /32)
                  43.229.15.0/24 (max length /24)
```

References

1. [https://learn.nsrc.org/bgp/MANRS4 RPKI and ROA](https://learn.nsrc.org/bgp/MANRS4_RPKI_and_ROA)
2. <https://nsrc.org/workshops/2019/mnno1/riso/networking/routing-security/en/presentations/BGP-Origin-Validation.pdf>
3. <https://www.manrs.org/manrs>
4. <https://blog.cloudflare.com/rpki-details/>
5. <https://www.apnic.net/get-ip/faqs/rpki/>
6. <https://www.apnic.net/community/security/resource-certification/>
7. <https://blog.apnic.net/2019/10/28/how-to-installing-an-rpki-validator/>
8. <https://blog.apnic.net/2019/09/11/how-to-creating-rpki-roas-in-myapnic/>
9. [https://www.apnic.net/wp-content/uploads/2017/12/ROUTE MANAGEMENT GUIDE.pdf](https://www.apnic.net/wp-content/uploads/2017/12/ROUTE_MANAGEMENT_GUIDE.pdf)
10. <https://www.cirt.gov.bd/জাতীয়-ডাটা-সেন্টারে-আরপি/>

Questions?