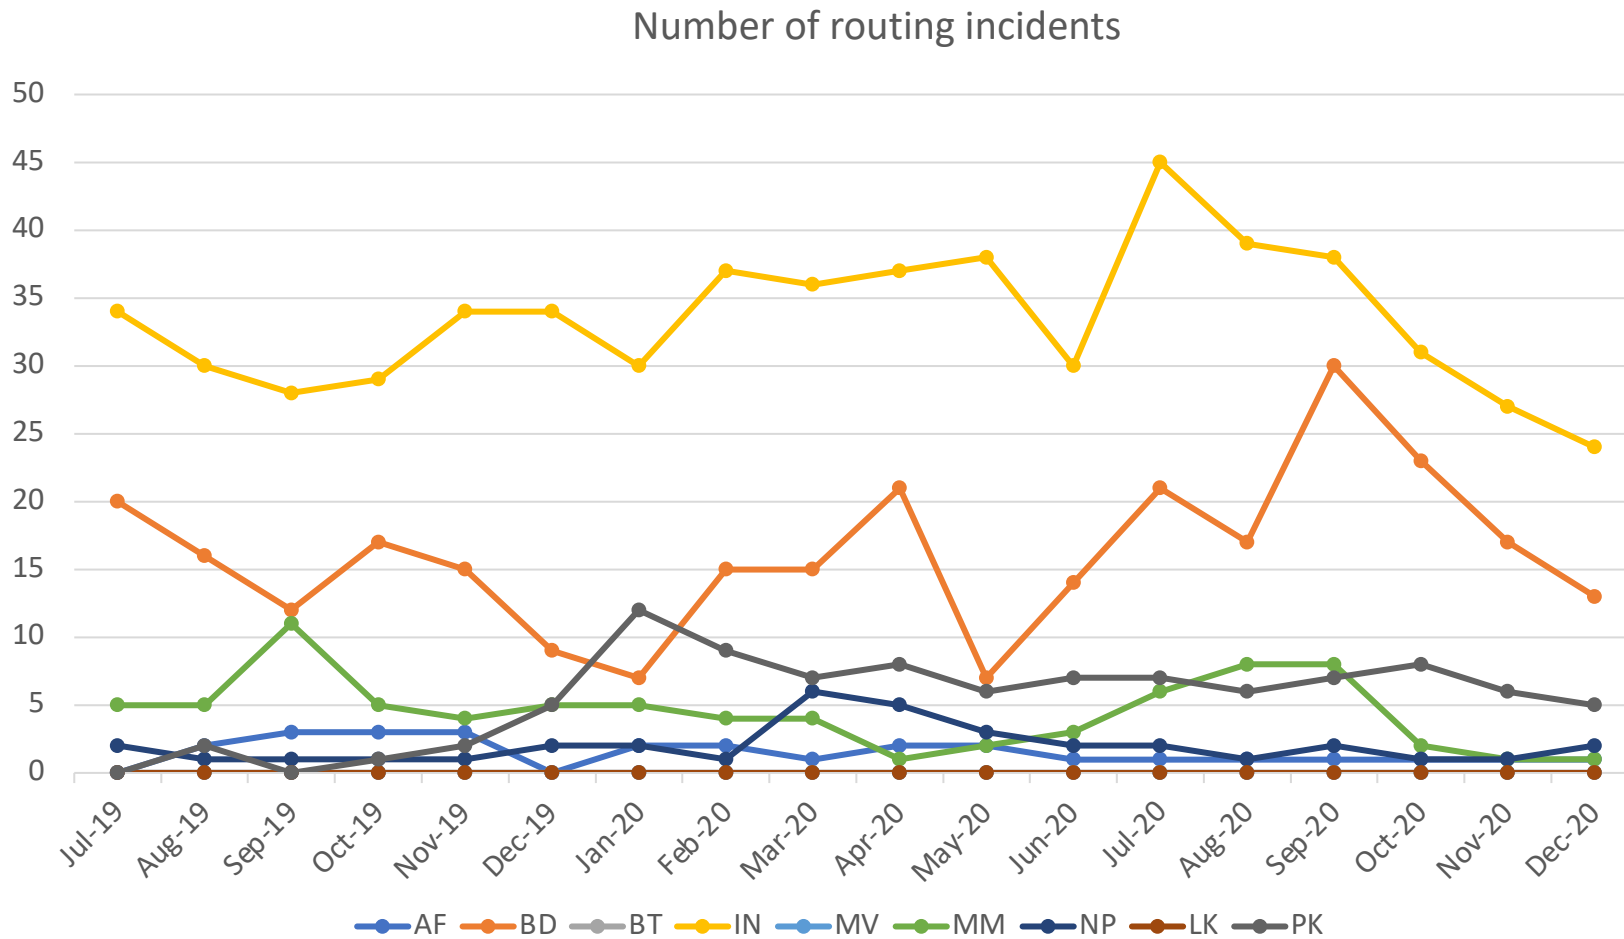


RPKI invalids aren't going away

Md. Abdul Awal

#SANOG36

Routing Incidents in South Asian Countries




Routing incidents still exist

Stats: observatory.manrs.org

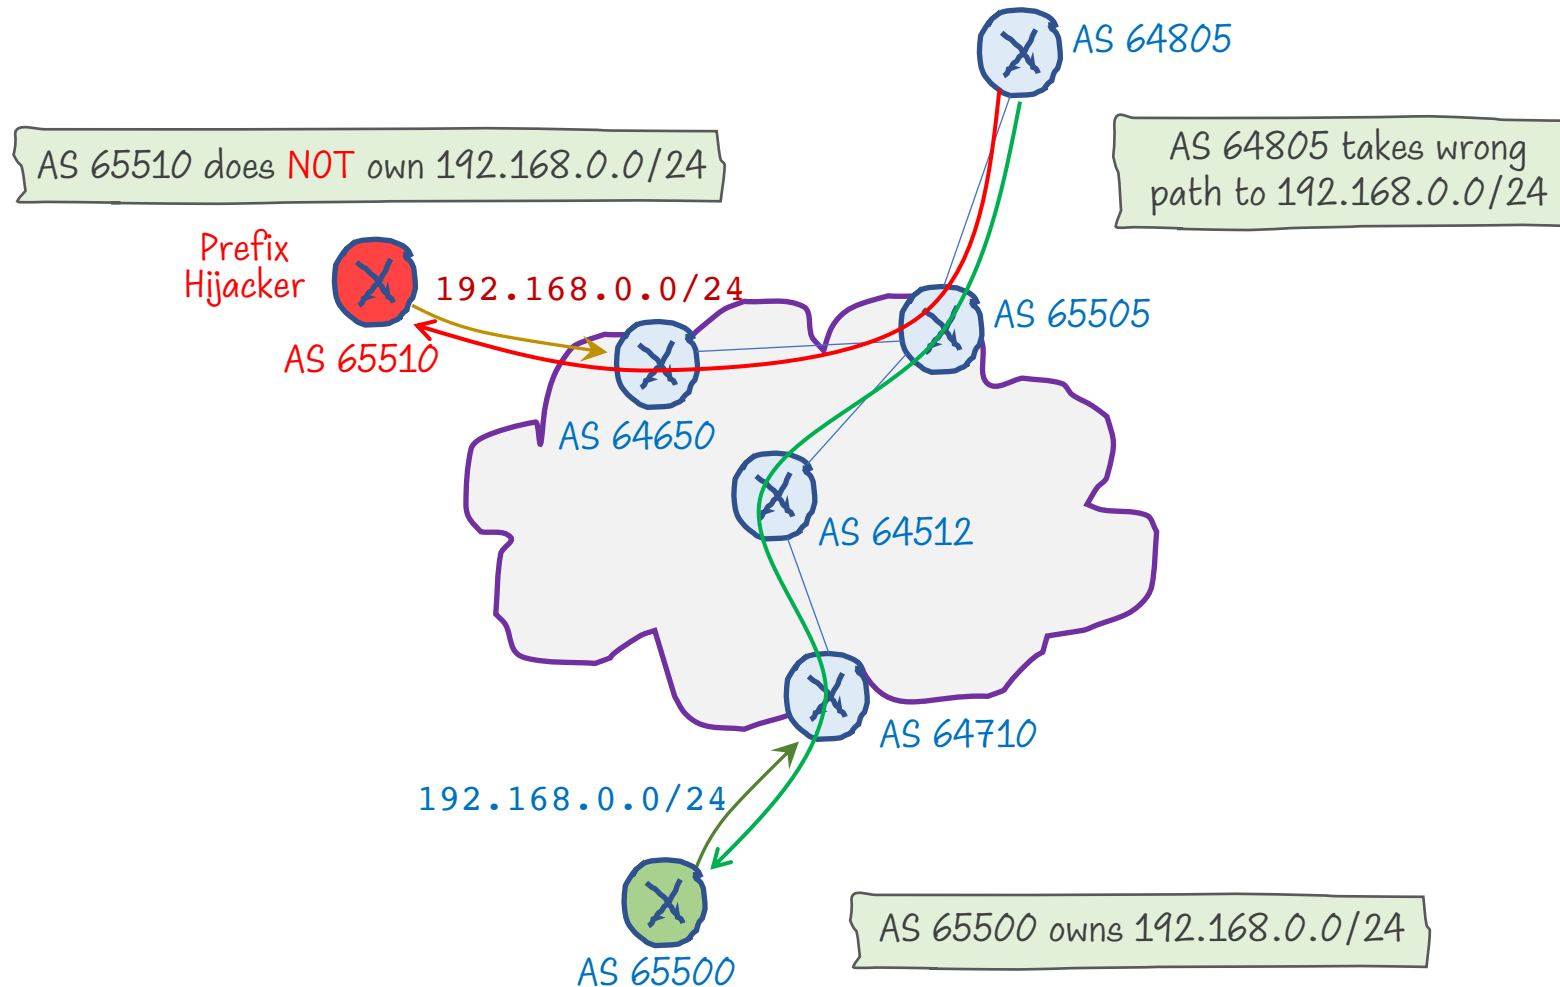
RPKI Status of BGP Prefixes in South Asia

RPKI status of BGP prefixes in AF, BD, BT, IN, MM, MV, NP, LK and PK



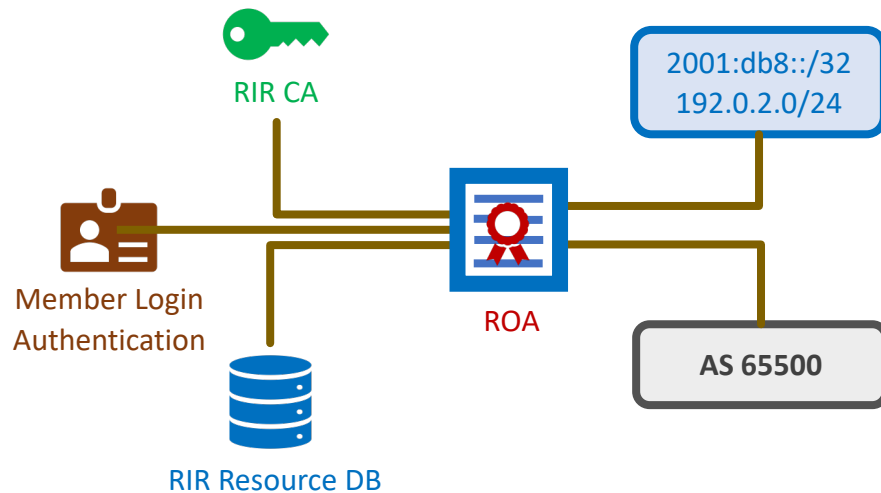
Stats: observatory.manrs.org

Prefix/Route Hijack: The Common Routing Incident




RPKI would be able to solve it

1 Signing prefixes a.k.a. creating ROA

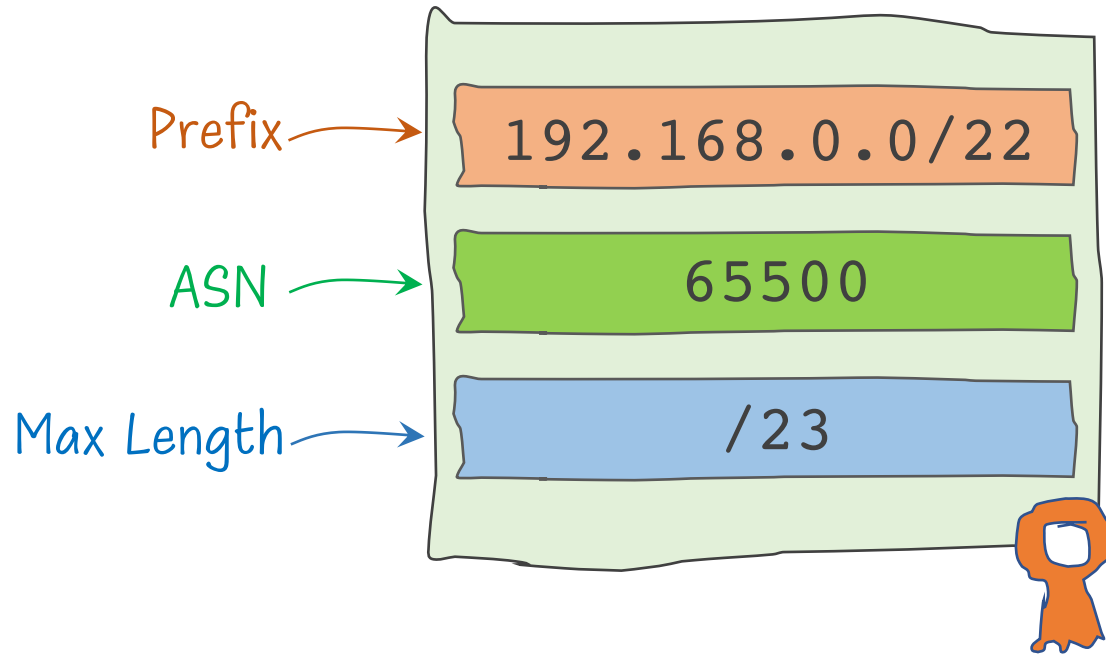


2 Validating ROAs a.k.a doing ROV



What makes a route RPKI Invalid?

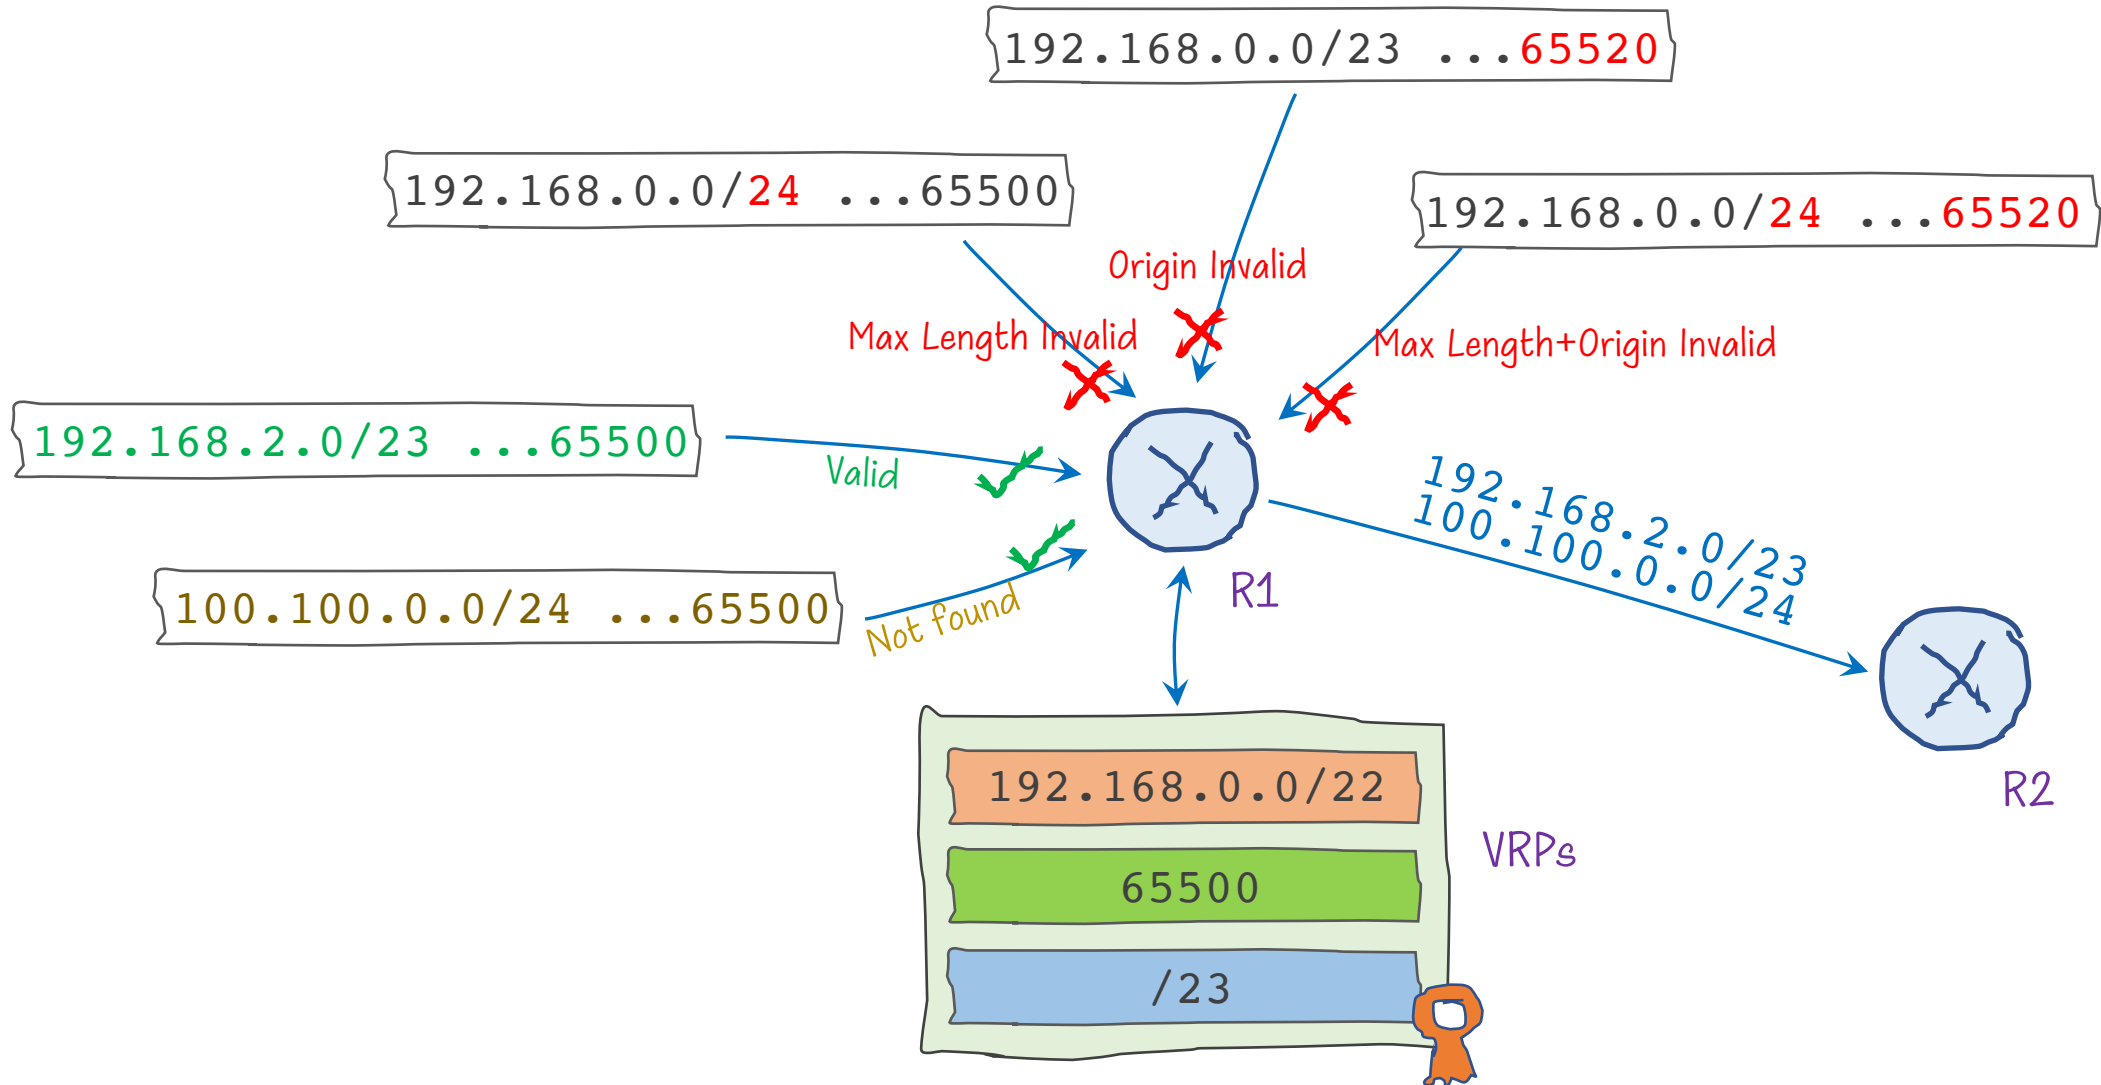
Route Origin Authorization (ROA)



Prefixes covered
by the ROA

- ✓ 192.168.0.0/22
- ✓ 192.168.0.0/23
- ✗ 192.168.0.0/24
- ✗ 192.168.1.0/24
- ✓ 192.168.2.0/23
- ✗ 192.168.2.0/24
- ✗ 192.168.3.0/24

Route Origin Validation (ROV)



Let's see some examples










Example: RPKI Invalids





```
awal@Awals-MacBook-Air ~> whois -h whois.bgpmon.net " --roa 135076 137.59.180.0/24"  
2 - Not Valid: Invalid Prefix-Length
```





```
awal@Awals-MacBook-Air ~> whois -h whois.bgpmon.net " --roa 132735 103.73.224.0/24"  
2 - Not Valid: Invalid Origin ASN, expected 132530
```

```
awal@Awals-MacBook-Air ~> whois -h whois.bgpmon.net " --roa 138487 103.73.224.0/22"  
2 - Not Valid: Invalid Prefix-Length + Invalid Origin ASN, expected 132530
```





















Example: Invalid Origin

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>	

Announced	
Origin AS	Announcement
<u>AS4613</u>	<u>103.86.56.0/24</u>  
<u>AS136380</u>	<u>103.86.56.0/24</u>  

Announced By	
Origin AS	Announcement
<u>AS3</u>	<u>103.141.3.0/24</u>  
<u>AS139300</u>	<u>103.141.3.0/24</u>  

Example: Invalid Prefix Length

Prefix		
<u>103.208.180.0/22</u>		
<u>103.208.180.0/24</u>		
<u>103.208.181.0/24</u>		
<u>103.208.182.0/24</u>		
<u>103.208.183.0/24</u>		
<u>137.59.180.0/22</u>		
<u>137.59.180.0/24</u>		
<u>137.59.181.0/24</u>		
<u>137.59.182.0/24</u>		
<u>137.59.183.0/24</u>		

```
-----  
ROA Details  
-----  
Origin ASN:      AS135076  
Not valid Before: 2019-12-01 13:03:06  
Not valid After: 2020-05-01 00:00:00 Expires in 100d12h5m8s  
Trust Anchor:    rpki.apnic.net  
Prefixes:        137.59.180.0/22 (max length /22)  
                  103.208.180.0/22 (max length /22)
```

More Example: Invalid Prefix Length

Prefix		
<u>65.18.112.0/20</u>		
<u>65.18.112.0/21</u>		
<u>65.18.112.0/22</u>		
<u>65.18.112.0/23</u>		
<u>65.18.112.0/24</u>		
<u>65.18.113.0/24</u>		
<u>65.18.114.0/23</u>		
<u>65.18.114.0/24</u>		
<u>65.18.115.0/24</u>		
<u>65.18.116.0/22</u>		
<u>65.18.116.0/23</u>		
<u>65.18.116.0/24</u>		
<u>65.18.117.0/24</u>		
<u>65.18.118.0/23</u>		
<u>65.18.120.0/21</u>		
<u>65.18.120.0/22</u>		
<u>65.18.120.0/23</u>		
<u>65.18.120.0/24</u>		
<u>65.18.121.0/24</u>		
<u>65.18.122.0/23</u>		


Prefix		
<u>119.73.65.0/24</u>		
<u>119.160.0.0/17</u>		
<u>119.160.64.0/18</u>		
<u>119.160.64.0/24</u>		
<u>119.160.65.0/24</u>		
<u>119.160.66.0/24</u>		
<u>119.160.67.0/24</u>		
<u>119.160.68.0/24</u>		
<u>119.160.69.0/24</u>		
<u>119.160.71.0/24</u>		
<u>119.160.80.0/24</u>		
<u>119.160.81.0/24</u>		
<u>119.160.82.0/24</u>		
<u>119.160.83.0/24</u>		
<u>119.160.84.0/24</u>		
<u>119.160.85.0/24</u>		
<u>119.160.86.0/24</u>		
<u>119.160.87.0/24</u>		
<u>119.160.91.0/24</u>		

Prefix		
<u>2400:ac40::/32</u>		
<u>2400:ac40:800::/48</u>		
<u>2400:ac40:801::/48</u>		
<u>2400:ac40:802::/48</u>		
<u>2400:ac40:803::/48</u>		
<u>2400:ac40:804::/48</u>		
<u>2400:ac40:805::/48</u>		
<u>2400:ac40:806::/48</u>		
<u>2400:ac40:807::/48</u>		
<u>2400:ac40:808::/48</u>		
<u>2400:ac40:809::/48</u>		
<u>2400:ac40:80a::/48</u>		
<u>2400:ac40:80b::/48</u>		
<u>2400:ac40:80c::/48</u>		
<u>2400:ac40:80d::/48</u>		
<u>2400:ac40:80e::/48</u>		
<u>2400:ac40:80f::/48</u>		
<u>2400:ac40:811::/48</u>		
<u>2400:ac40:812::/48</u>		
<u>2400:ac40:813::/48</u>		
<u>2400:ac40:814::/48</u>		

RPKI Invalids in South Asia

Country	% of Invalids
IN	55.1
PK	22.2
BD	18.1
LK	2.6
NP	1.2
AF	0.5
MM	0.3
MV	0
BT	0

Invalid Types




Stats: Anurag Bhatia (<https://rpkι.anuragbhatia.com>)

So, why invalids exist in routing atmosphere?


Several Reasons...

- Incorrect ROAs
 - Mostly because of misconfigured Max Length
 - Sometimes because of wrong ASN
 - Lack of awareness?
- Wrong BGP announcements
 - Route advertised without checking its ROA
 - Old habit?
- Most importantly, invalids are not dropped
 - Origin validation is not widely deployed in the region
 - Any reason?


Fix it: Who and How



Create appropriate ROAs for your prefixes

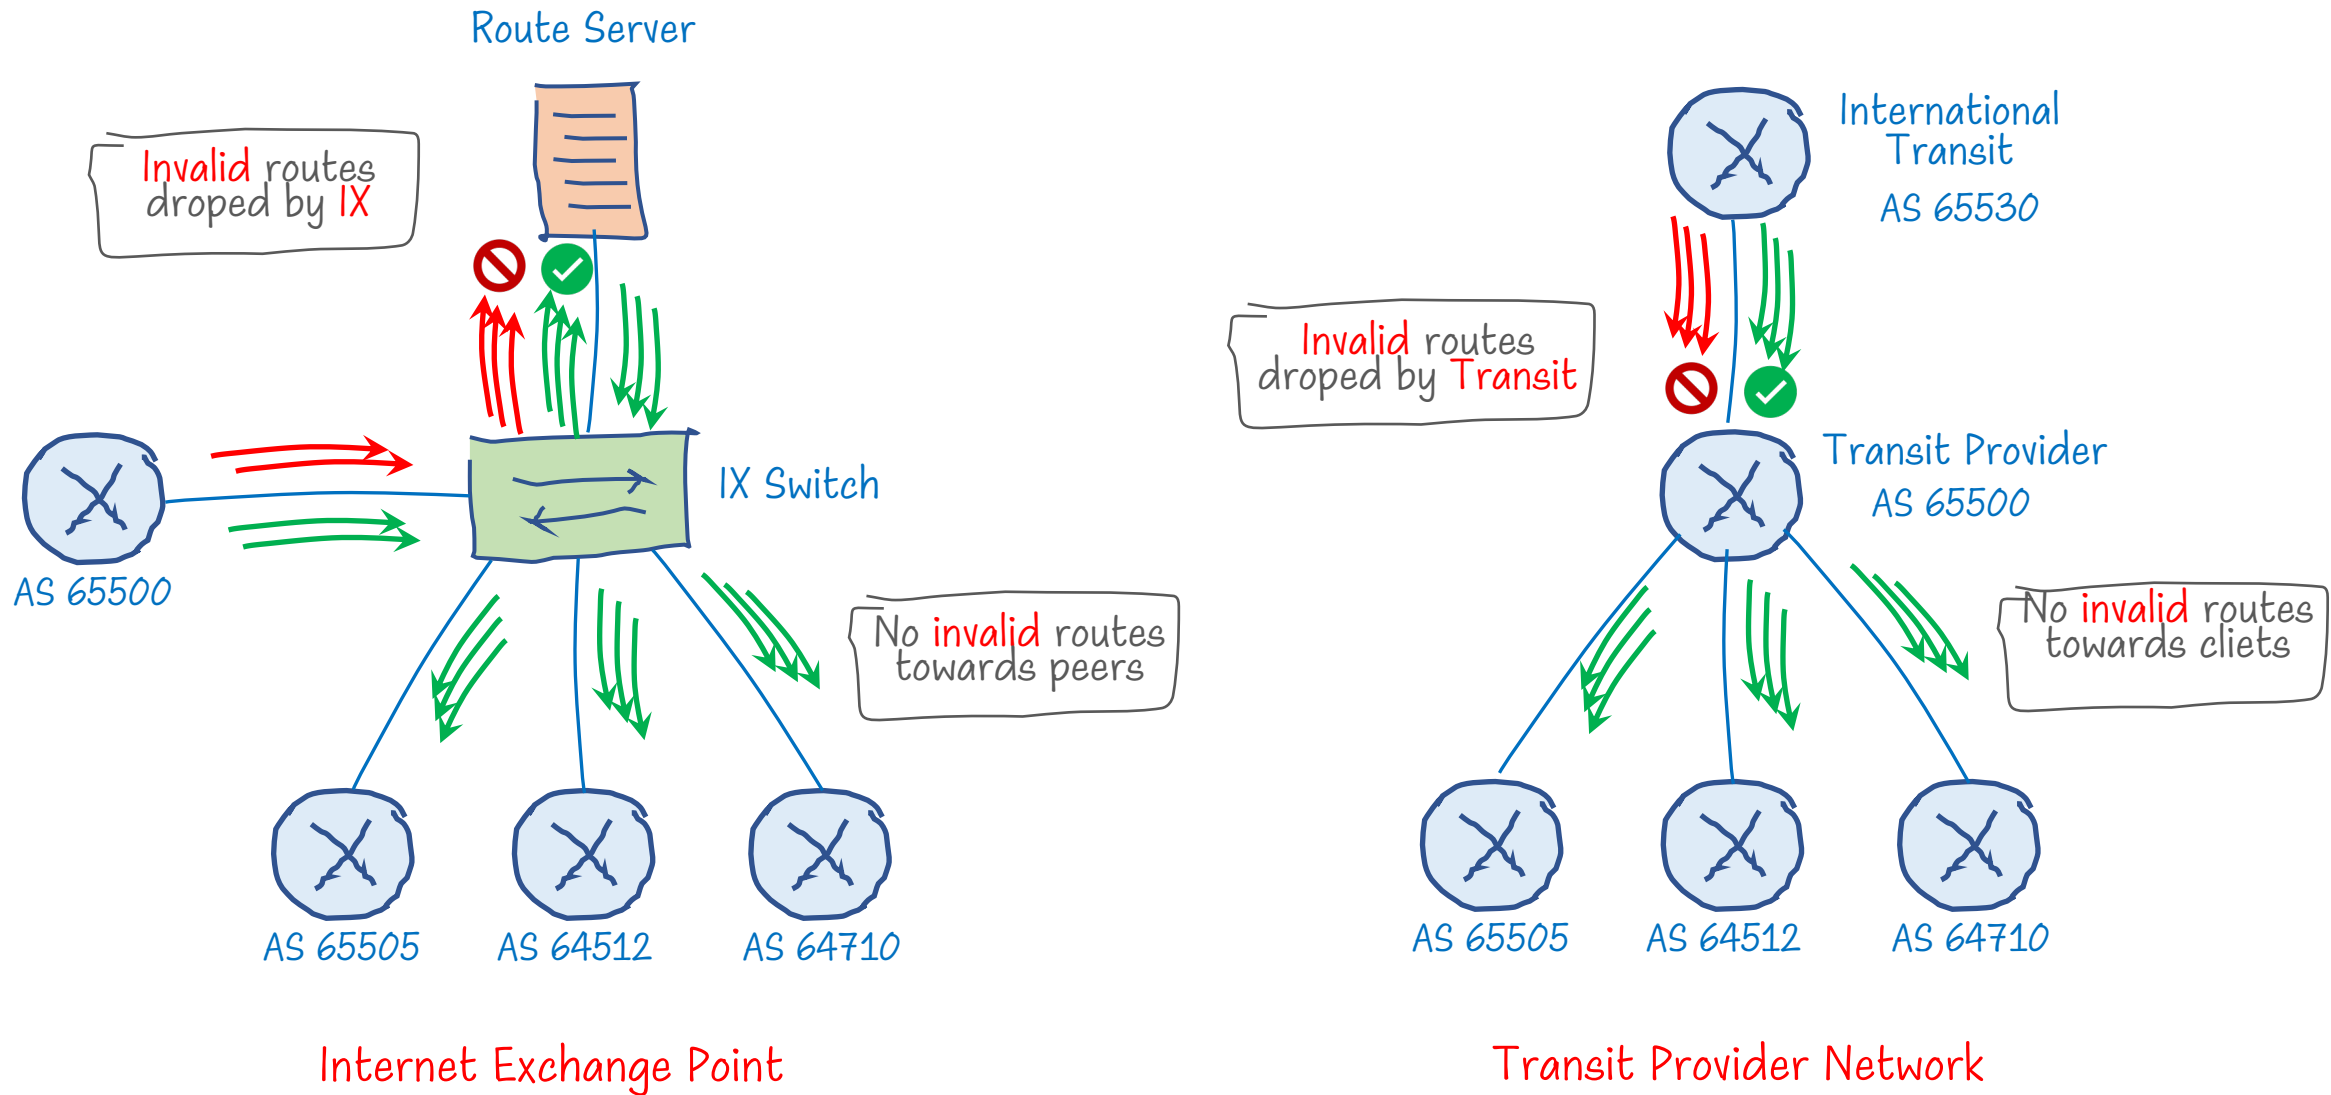


Announce only the correct prefix in BGP



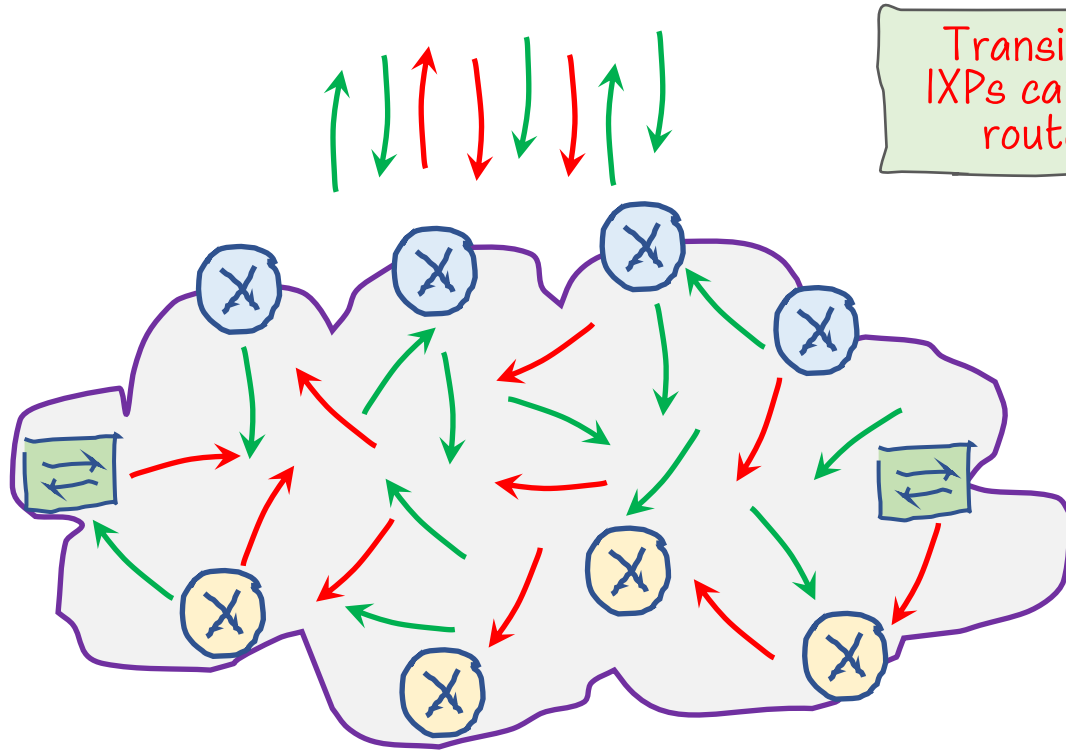
Implement origin validation
i.e. drop RPKI Invalids

Route Origin Validation at IX and Transit



Validation can make our routing table Invalid-free


International Transits



Internet Routing Infrastructure
Without Validation

Transit providers and IXP's can prevent Invalid route propagation

International Transits



Internet Routing Infrastructure
With Validation

Thanks!

Questions?