



<https://About.Me/Terry.Sweetser>

# Routing Security SIG

## *RPKI Data Results for APAC*

# APRICOT 2024

**APNIC 57**

BANGKOK, THAILAND  
21 February - 1 March 2024



# Boilerplate

- All data captured at 0400Z on February 7, 2024.
- It will be Open Source on GITHUB
  - <https://github.com/tcsweetser/bcp185>
- Lots of SQL, you get to see inside the Sausage Factory!

```
bcp185=# \d+
```

List of relations							
Schema	Name	Type	Owner	Persistence	Access method	Size	Description
public	overlapping_signed_routes	materialized view	apnic	permanent	heap	8344 kB	
public	rir_allocations	table	apnic	permanent	heap	26 MB	
public	routes	table	apnic	permanent	heap	2121 MB	
public	rpki_signed_routes	table	apnic	permanent	heap	34 MB	

(4 rows)

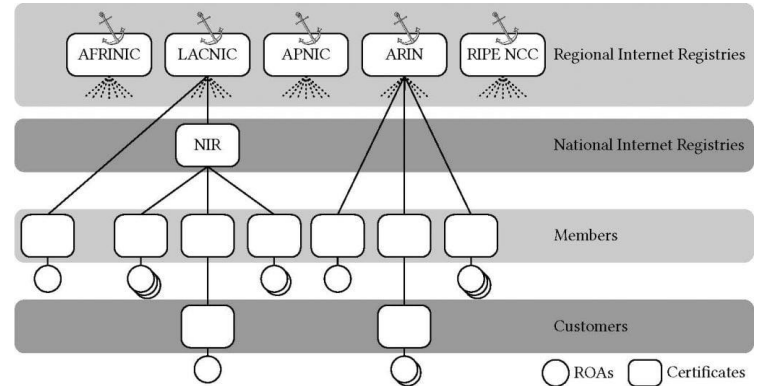
# Data Sources



<https://www.routeviews.org/routeviews/>

<https://ftp.apnic.net/stats/>

<https://blog.apnic.net/2020/01/29/is-rpki-ready-for-the-big-screen/>



# Simple Data Structures

```
bcp185=# \d+ rpki_signed_routes
```

```
Table "public.rpki_signed_routes"
  Column      | Type      | Collation | Nullable | Default | Storage  | Compression | Stats target | Description
-----+-----+-----+-----+-----+-----+-----+-----+-----
ASN           | text      |           |          |         | extended |              |              |
IP Prefix     | cidr      |           |          |         | main     |              |              |
Max Length    | integer   |           |          |         | plain    |              |              |
Trust Anchor  | text      |           |          |         | extended |              |              |
Expires       | bigint    |           |          |         | plain    |              |              |
Indexes:
  "rpki_signed_routes_ASN_idx" btree ("ASN")
  "rpki_signed_routes_IP Prefix_idx" gist ("IP Prefix" inet_ops)
  "rpki_signed_routes_Trust Anchor_idx" btree ("Trust Anchor")
  "rpki_signed_routes_masklen_idx" btree (masklen("IP Prefix"::inet))
Access method: heap
```

```
bcp185=# \d+ routes
```

```
Table "public.routes"
  Column      | Type      | Collation | Nullable | Default | Storage  | Compression | Stats target | Description
-----+-----+-----+-----+-----+-----+-----+-----+-----
dumptype      | text      |           |          |         | extended |              |              |
exported      | text      |           |          |         | extended |              |              |
protocol       | character(1) |           |          |         | extended |              |              |
nexthop       | inet      |           |          |         | main     |              |              |
asn           | bigint    |           |          |         | plain    |              |              |
prefix        | cidr      |           |          |         | main     |              |              |
aspath        | text      |           |          |         | extended |              |              |
origin        | text      |           |          |         | extended |              |              |
Indexes:
  "routes_prefix_idx" gist (prefix inet_ops)
Access method: heap
```

# Same Basic Numbers

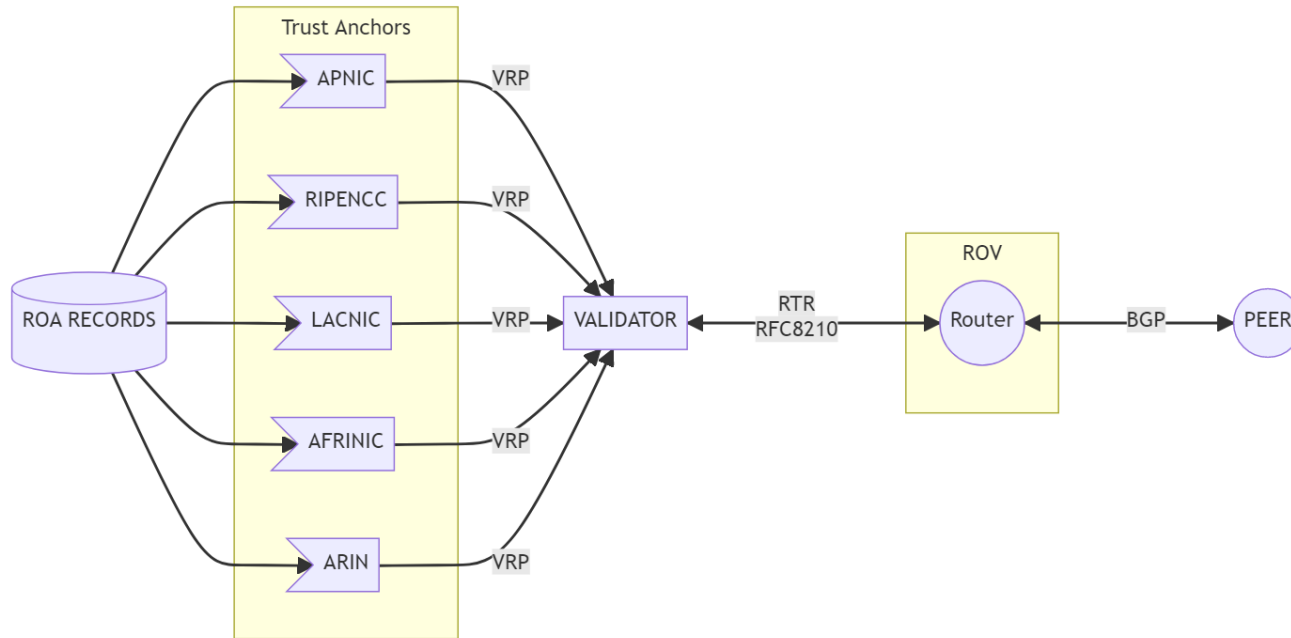
```
bcp185=# select count(*) from rpki_signed_routes;
count
-----
508414
(1 row)
```

```
bcp185=# select count(*) from routes;
count
-----
18535539
(1 row)
```

```
bcp185=# select family("IP Prefix"), "Trust Anchor",
count(distinct "IP Prefix")
from rpki_signed_routes
group by 1,2 order by 2,1;
 family | Trust Anchor | count
-----+-----+-----
4 | afrinic      | 8920
6 | afrinic      | 647
4 | apnic        | 92059
6 | apnic        | 25952
4 | arin         | 74949
6 | arin         | 15991
4 | lacnic       | 21594
6 | lacnic       | 7915
4 | ripe         | 174975
6 | ripe         | 39003
(10 rows)
```

```
bcp185=# select family(prefix),count(distinct prefix) from routes group by 1;
 family | count
-----+-----
4 | 982297
6 | 204965
(2 rows)
```

# RPKI in less than 1 minute ...

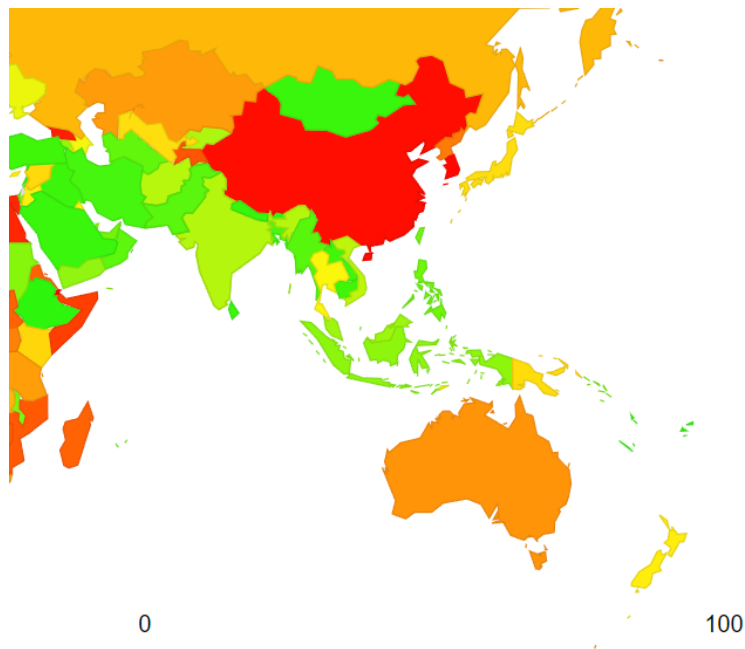


# ROA, are we there yet?

<https://stats.labs.apnic.net/roas>

Lots of Green here, but...

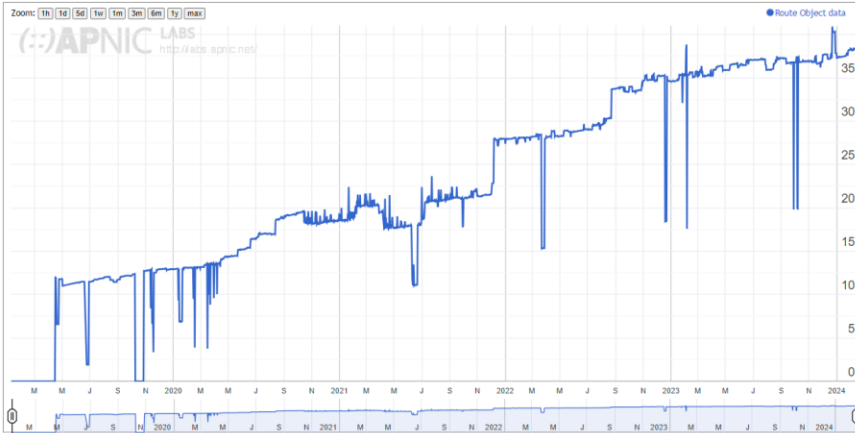
**CN**  
AU  
PG  
JP  
NZ  
TH



# Progress 😊


## Use of Route Object Validation for Asia (XD)

Display: Addresses (Advertised ROA-Valid Advertisements), Total (IPv4 + IPv6), Percent (of Total)



## Use of Route Object Validation for Oceania (XF)

Display: Addresses (Advertised ROA-Valid Advertisements), Total (IPv4 + IPv6), Percent (of Total)



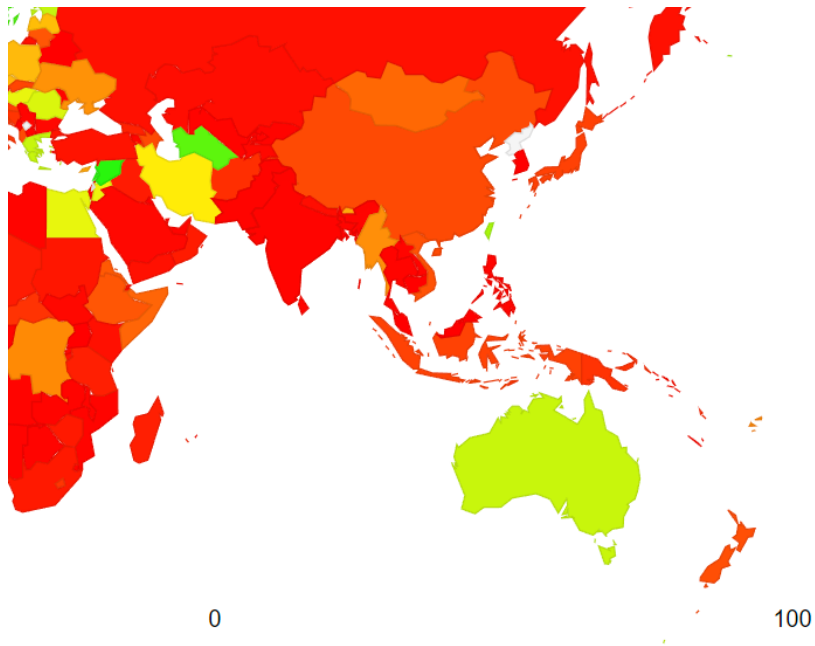


# ROV, how is that going?

<https://stats.labs.apnic.net/rpki>

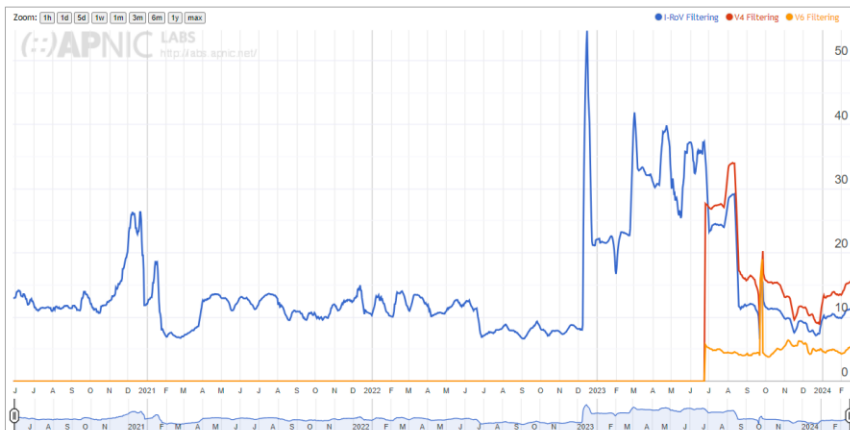
Great job by ...

TW  
AU  
...  
...



# Validation: A Long Road Ahead

Use of RPKI Validation for Asia (XD)



Use of RPKI Validation for Oceania (XF)



# RFC7115 + RFC9319 = BCP185

For this reason, this document recommends that, whenever possible, operators SHOULD use "minimal ROAs" that authorize only those IP prefixes that are actually originated in BGP, and no other prefixes. Further, it recommends ways to reduce the forged-origin attack surface by prudently limiting the address space that is included in ROAs. One recommendation is to avoid using the maxLength attribute in ROAs except in some specific cases. The recommendations complement and extend those in [RFC7115]. The document also

- Members of All the RIRs continue to “OVERSIGN”
  - Huge MAX LENGTHS with no matching ROUTES.
  - Largest and most connected economies are the WORST cases.
  - Highly questionable ROAs signed by those who should know better.

# The Top 10 ...

```
bcp185=# select * from report_bcp185_by_iso where  
("Global Routes"/"ROA Coverage") < 0.00005  
and economy_iso≠''  
and "Global Routes" > 10  
order by 2 desc;
```

economy_iso	ROA Coverage	Global Routes	BCP185 Compliance
FR	1195134656	10450	000.00%
TW	381642809	12649	000.00%
JP	157727426	6087	000.00%
MA	33835598	265	000.00%
AL	32638327	927	000.00%
IS	4457025	178	000.00%
CD	2621515	69	000.00%
LI	2228351	66	000.00%
GD	1974818	34	000.00%
GP	1180482	25	000.00%

(10 rows)

# The Top 25 in APAC ...

```
bcp185=#
bcp185=# select * from report_bcp185_combined where rir='apnic' order by 3 desc limit 25;
  rir | economy_iso | ROA Coverage | Global Routes | BCP185 Compliance
-----+-----+-----+-----+-----
 apnic | TW          | 380463099   | 12614         | 000.00%
 apnic | JP          | 155498452   | 6027          | 000.00%
 apnic | CN          | 36239128    | 2153          | 000.01%
 apnic | ID          | 29810247    | 17760         | 000.06%
 apnic | IN          | 4275225     | 38941         | 000.91%
 apnic | SG          | 3816797     | 5844          | 000.15%
 apnic | AU          | 3468345     | 7941          | 000.23%
 apnic | VN          | 2285508     | 22620         | 000.99%
 apnic | NZ          | 2105698     | 1802          | 000.09%
 apnic | HK          | 2059318     | 7349          | 000.36%
 apnic | US          | 1105124     | 1612          | 000.15%
 apnic | MY          | 1062999     | 3108          | 000.29%
 apnic | BD          | 958449      | 10300         | 001.07%
 apnic | MN          | 656210      | 514           | 000.08%
 apnic | PK          | 576198      | 7544          | 001.31%
 apnic | PH          | 416962      | 5387          | 001.29%
 apnic | KH          | 268386      | 2165          | 000.81%
 apnic | MM          | 263864      | 1071          | 000.41%
 apnic | MV          | 263226      | 281           | 000.11%
 apnic | NP          | 139692      | 1102          | 000.79%
 apnic | MO          | 132886      | 113           | 000.09%
 apnic | BT          | 131854      | 125           | 000.09%
 apnic | VU          | 131140      | 68            | 000.05%
 apnic | TH          | 117555      | 6822         | 005.80%
 apnic | LK          | 85615       | 743          | 000.87%
(25 rows)
```

# Good News Everyone!

```
bcp185=# select * from report_bcp185_by_iso where  
("Global Routes"/"ROA Coverage") > 0.80 and economy_iso≠''  
order by 3 desc;
```

economy_iso	ROA Coverage	Global Routes	BCP185 Compliance
SY	163	163	100.00%
YE	127	121	095.28%
BW	73	64	087.67%
TJ	46	42	091.30%
GU	26	23	088.46%
GL	24	21	087.50%
TM	20	20	100.00%
TD	24	20	083.33%
PM	18	18	100.00%
SM	17	16	094.12%
MH	18	15	083.33%
AX	16	15	093.75%
NR	15	14	093.33%
FO	14	13	092.86%
IO	12	12	100.00%
CW	13	12	092.31%
DM	7	7	100.00%
TT	7	7	100.00%
AI	5	5	100.00%
GG	6	5	083.33%
CU	5	5	100.00%
NU	5	5	100.00%
BB	4	4	100.00%
VA	2	2	100.00%
VI	1	1	100.00%

(25 rows)



# Overlapping Signed Routes ...

```
bcp185=# select * from overlapping_signed_routes limit 10;
```

asns_subnet	asns_supernet	subnet	subnet_maxlen	supernet	supernet_maxlen
{AS132215}	{AS9583}	1.6.228.0/24	24	1.6.228.0/22	24
{AS4755}	{AS9583}	1.6.229.0/24	24	1.6.228.0/22	24
{AS132215}	{AS9583}	1.7.142.0/24	24	1.7.140.0/22	24
{AS132215}	{AS9583}	1.7.151.0/24	24	1.7.148.0/22	24
{AS132215}	{AS9583}	1.7.161.0/24	24	1.7.160.0/22	24
{AS132215}	{AS9583}	1.7.162.0/24	24	1.7.160.0/22	24
{AS132215}	{AS9583}	1.7.180.0/24	24	1.7.180.0/22	24
{AS396421}	{AS9583}	1.7.225.0/24	24	1.7.224.0/22	24
{AS4755}	{AS9583}	1.7.229.0/24	24	1.7.228.0/22	24
{AS65037}	{AS4788}	1.9.12.0/24	24	1.9.0.0/16	24

(10 rows)

```
bcp185=# select count(*) from overlapping_signed_routes;
```

count
63296

(1 row)

# Please Stop Doing ...

## Max Max Length IPv4

```
bcp185=# select "Trust Anchor",count(distinct "ASN")
from rpki_signed_routes
where family("IP Prefix")=4 and "Max Length" = 32
group by 1 order by 2 desc;
```

Trust Anchor	count
ripe	572
lacnic	281
arin	162
apnic	67
afrinic	20

(5 rows)

## Max Max Length IPv6

```
bcp185=# select "Trust Anchor",count(distinct "ASN")
from rpki_signed_routes
where family("IP Prefix")=6 and "Max Length" = 128
group by 1 order by 2 desc;
```

Trust Anchor	count
lacnic	276
ripe	130
arin	33
apnic	28
afrinic	15

(5 rows)



# Please stop doing ...

```
bcp185=# select array_agg("ASN") as asns,"IP Prefix" as subnet, "Max Length" as maxlen,  
split_cidr_count("IP Prefix","Max Length") from rpki_signed_routes  
where family("IP Prefix")=4  
group by 2,3 order by 4 desc limit 10;
```

asns	subnet	maxlen	split_cidr_count
{AS21928}	100.128.0.0/9	24	65535
{AS6253}	48.0.0.0/9	24	65535
{AS8987,AS14618,AS16509}	3.128.0.0/10	24	32767
{AS16509,AS8987,AS14618}	44.192.0.0/10	24	32767
{AS8987,AS14618,AS16509}	3.192.0.0/10	24	32767
{AS14618,AS8987,AS16509}	3.0.0.0/10	24	32767
{AS6253}	48.128.0.0/10	24	32767
{AS14618,AS8987,AS16509}	34.192.0.0/10	24	32767
{AS14618,AS8987,AS16509}	3.64.0.0/10	24	32767
{AS17676}	126.64.0.0/10	24	32767

(10 rows)

# Please explain ...

```
bcp185=#
bcp185=# select * from _route_rpki_subnet_counts where subnet <=<= '100.128.0.0/9'::cidr;
 subnet      | maxlen | expanded_count | routed_count | rir   | economy_iso
-----+-----+-----+-----+-----+-----
 100.128.0.0/9 |    24 |          65535 |             1 | arin  | US
(1 row)

bcp185=# select * from _route_rpki_subnet_counts where subnet <=<= '48.0.0.0/9'::cidr;
 subnet      | maxlen | expanded_count | routed_count | rir   | economy_iso
-----+-----+-----+-----+-----+-----
 48.0.0.0/9 |    24 |          65535 |             0 | arin  | US
(1 row)

bcp185=# select * from _route_rpki_subnet_counts where subnet = '3.0.0.0/10'::cidr;
 subnet      | maxlen | expanded_count | routed_count | rir   | economy_iso
-----+-----+-----+-----+-----+-----
 3.0.0.0/10 |    24 |          32767 |            243 | arin  | US
(1 row)
```

# Wait, there's more ...

```
bcpr185=# select * from rpki_signed_routes where "ASN" = 'AS7473'  
and "IP Prefix" << '2001:c10::/32'::cidr order by 3 limit 25;
```

ASN	IP Prefix	Max Length	Trust Anchor	Expiration
AS7473	2001:c10:8000::/34	34	apnic	170784473
AS7473	2001:c10:4000::/34	34	apnic	170784473
AS7473	2001:c10:c000::/35	35	apnic	170784473
AS7473	2001:c10:2000::/35	35	apnic	170784473
AS7473	2001:c10:e000::/36	36	apnic	170784473
AS7473	2001:c10:1000::/36	36	apnic	170784473
AS7473	2001:c10:f000::/37	37	apnic	170784473
AS7473	2001:c10:800::/37	37	apnic	170784473
AS7473	2001:c10:400::/38	38	apnic	170784473
AS7473	2001:c10:200::/39	39	apnic	170784473
AS7473	2001:c10:fc00::/39	39	apnic	170784473
AS7473	2001:c10:100::/40	40	apnic	170784473
AS7473	2001:c10:fe00::/40	40	apnic	170784473
AS7473	2001:c10:ff80::/41	41	apnic	170784473
AS7473	2001:c10:c0::/42	42	apnic	170784473
AS7473	2001:c10:40::/42	42	apnic	170784473
AS7473	2001:c10:ff40::/42	42	apnic	170784473
AS7473	2001:c10:a0::/43	43	apnic	170784473
AS7473	2001:c10:20::/43	43	apnic	170784473
AS7473	2001:c10:ff20::/43	43	apnic	170784473
AS7473	2001:c10:10::/44	44	apnic	170784473
AS7473	2001:c10:90::/44	44	apnic	170784473
AS7473	2001:c10:ff10::/44	44	apnic	170784473
AS7473	2001:c10:8::/45	45	apnic	170784473
AS7473	2001:c10:88::/45	45	apnic	170784473

(25 rows)

```
85=# select * from rpki_signed_routes where "ASN" = 'AS7473'  
"IP Prefix" << '2001:c10::/32'::cidr order by 3 limit 25 offset
```

ASN	IP Prefix	Max Length	Trust Anchor	Expiration
AS7473	2001:c10:80:1::2c0/126	126	apnic	170784473
AS7473	2001:c10:80:1::1f0/126	126	apnic	170784473
AS7473	2001:c10:80:1::16c/126	126	apnic	170784473
AS7473	2001:c10:80:1::164/126	126	apnic	170784473
AS7473	2001:c10:80:1::144/126	126	apnic	170784473
AS7473	2001:c10:80:1::f8/126	126	apnic	170784473
AS7473	2001:c10:80:1::e8/126	126	apnic	170784473
AS7473	2001:c10:80:1::e0/126	126	apnic	170784473
AS7473	2001:c10:80:1::5c/126	126	apnic	170784473
AS7473	2001:c10:80:1::20/126	126	apnic	170784473
AS7473	2001:c10:0:1::64/126	126	apnic	170784473
AS7473	2001:c10:0:1::60/127	127	apnic	170784473
AS7473	2001:c10:0:1::116/127	127	apnic	170784473
AS7473	2001:c10:0:1::136/128	128	apnic	170784473
AS7473	2001:c10:0:1::132/128	128	apnic	170784473
AS7473	2001:c10:0:1::122/128	128	apnic	170784473
AS7473	2001:c10:0:1::115/128	128	apnic	170784473
AS7473	2001:c10:0:1::63/128	128	apnic	170784473
AS7473	2001:c10:0:1::103/128	128	apnic	170784473
AS7473	2001:c10:0:1::62/128	128	apnic	170784473
AS7473	2001:c10:0:1::57/128	128	apnic	170784473
AS7473	2001:c10:0:1::16/128	128	apnic	170784473
AS7473	2001:c10:0:1::60/128	128	apnic	170784473
AS7473	2001:c10:0:1::33/128	128	apnic	170784473
AS7473	2001:c10:0:1::24/128	128	apnic	170784473

(rows)

# One Billion Routes!

```
bcp185=# select array_agg("ASN") as asns, "IP Prefix" as subnet, "Max Length" as maxlen,
split_cidr_count("IP Prefix", "Max Length") from rpki_signed_routes
group by 2,3 order by 4 desc limit 25;
```

asns	subnet	maxlen	split_cidr_count
{AS5511}	2a01:c000::/19	48	1073741823
{AS26810}	2601:9000::/20	48	536870911
{AS7303}	2800:2000::/20	48	536870911
{AS3462}	2001:b000::/21	48	268435455
{AS3462}	2001:b000::/23	64	67108863
{AS17676}	2400:2100::/24	48	33554431
{AS8987, AS14618, AS16509}	240f:8000::/24	48	33554431
{AS17676}	2403:9900::/24	48	33554431
{AS6939}	2600:7000::/24	48	33554431
{AS8987, AS14618, AS16509}	2600:1f00::/24	48	33554431
{AS8987, AS14618, AS16509}	2406:da00::/24	48	33554431
{AS17676}	2400:2200::/24	48	33554431
{AS17676}	2400:2300::/24	48	33554431
{AS7155}	2600:6800::/24	64	33554431
{AS14618, AS8987, AS16509}	2600:f000::/24	48	33554431
{AS36925}	2c0e:6000::/24	48	33554431
{AS14618, AS8987, AS16509}	2a05:d000::/25	48	16777215
{AS22927}	2802:8000::/25	48	16777215
{AS15557, AS198949}	2a02:8400::/25	48	16777215
{AS3215}	2a01:cb00::/26	48	8388607
{AS51964}	2a01:ce80::/26	48	8388607
{AS4780, AS9919}	2401:8000::/26	48	8388607
{AS17709}	2404::/26	48	8388607
{AS4780}	2001:4580::/26	48	8388607
{AS43754}	2a05:1a00::/26	128	8388607

(25 rows)

# What is MY ASK here?



- Sign your routes, but limit that to only what you advertise to the Global Internet.
- Don't use Public Sources of data for bizarre and non-sensical experiments.
- Don't know? Ask!



<https://About.Me/Terry.Sweetser>

# Questions?

THANK YOU  
FOR YOUR ATTENTION