

# The Scourge of Excessive AS-Sets

Doug Madory (Kentik)



# AS-SETs vs AS\_SETs

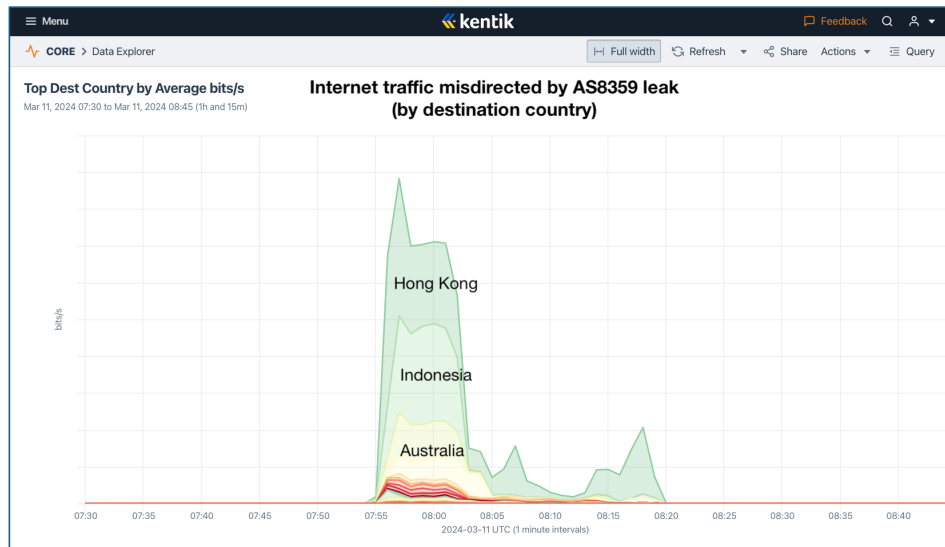
- NOTE: This talk discusses the IRR AS-SET object type!
  - A record in the IRR database that defines a group of ASNs used to simplify the management of routing policies by grouping multiple ASNs together.
- Not BGP AS\_SET construct, which has slated for deprecation.
  - See *Deprecation of AS\_SET and AS\_CONFED\_SET in BGP* (BCP 172, RFC 6472)
  - Aggregate AS\_SETs appear in the AS\_PATH of a BGP announcement as one or more ASNs surrounded by curly brackets.
  - Ex: 300 {200, 100}. This set indicates that the aggregate summarizes routes that have passed through AS200 and AS100.

**Let's begin with a BGP leak**



# AS-SETs and BGP Leaks

- On 07:56 UTC on March 11, 2024, Russian mobile operator MTS (AS8359) mistakenly propagated over 30,000 routes learned from the Hong Kong Internet Exchange (HKIX, AS4635) to its transit providers Lumen (AS3356) and Arelion (AS1299).



**Radars by Qrator**  
@Qrator\_Radar

AS8359 (MTS) leaked 4065 prefixes learned from AS4635 (HKIX-RS1) towards Tier1 AS3356 (LEVEL3), creating 4065 conflicts with 329 ASNs in 28 countries. Asian prefixes were mostly affected.

Max propagation: 39%  
Start: 2024-03-11 07:56 UTC, duration >25 min

**Incident Type** Created Leaks

**Key ASN** AS8359 - MTS - RU

**Overall Info**  
Conflicts count all: 4065  
ASNs affected: 329  
Countries affected: 28

**Prefixes created: 4065**

**Accepted ASNs during incident**

Conflict	Conflicts count
AS4635 - HKIX-RS1 - [RU] → AS8359 - MTS - [RU] → AS3356 - LEVEL3 - [US]	4065 (100.0%)

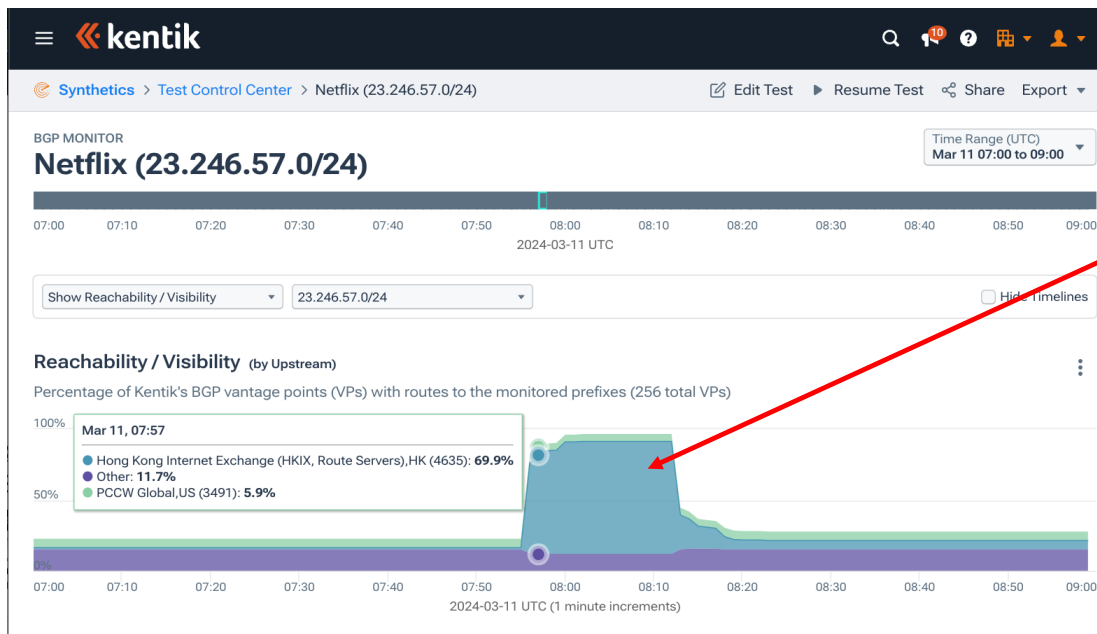
**Unique AS Paths**

Path	Conflicts count
AS7545(1) → AS4635(1) → AS8359(1) → AS3356(1)	1171 (28.8%)
AS7545(2) → AS4635(1) → AS8359(1) → AS3356(1)	281 (6.9%)
AS9340(1) → AS8359(1) → AS136106(1) → AS4635(1) → AS8359(1) → AS3356(1)	178 (4.4%)
AS17444(1) → AS10102(1) → AS4635(1) → AS8359(1) → AS3356(1)	136 (3.3%)

5:09 AM · Mar 11, 2024 · 7,262 Views

# AS-SETs and BGP Leaks

- Propagation of one Netflix's BGP routes announced at HKIX.
  - Normally circulated only regionally.
  - During the leak, the leaked version via AS8359 propagated globally.



The bulge in the middle of the graphic represents the dramatic increase in the number of our BGP sources who had this route in their table (with HKIX as the upstream).

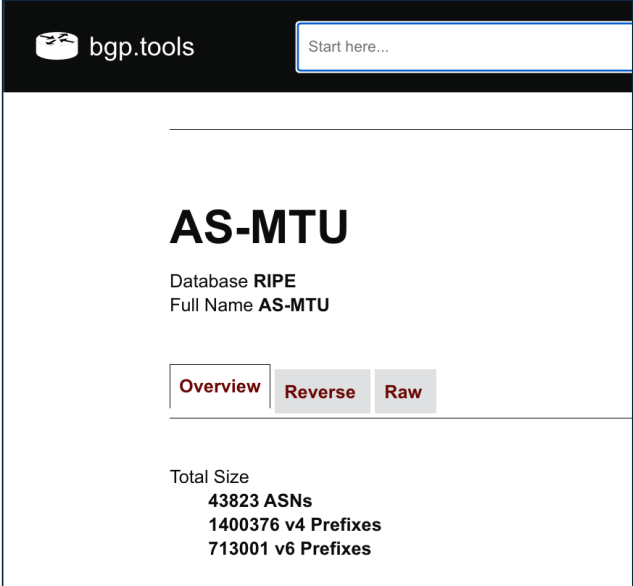
# AS-SETs and BGP Leaks

- Hey, mistakes happen! 🙄
- Since this was an adjacency leak (aka path leak), RPKI ROV can't help.
- At least we have AS-SETs to enable transit providers to programmatically build an appropriate allowlist to prevent the propagation of leaked routes, right? Right?



# Excessively Large AS-SETs

- The leaker in the March 11 route leak, uses an AS-SET called AS-MTU.
- Web utility Bgp.tools lists the contents of AS-SETs.
  - Expands AS-MTU to 43,823 ASNs!
  - There are 83,617 ASNs in the global routing table.
  - Any network applying AS-MTU as a filter for an interface with AS8359 is creating an allowlist containing these.
  - Some examples of prefixes allowed by AS-MTU:
    - 6.2.0.0/17 US Department of Defense
    - 8.36.240.0/20 Rural Telephone Service Company, Lenora, Kansas
    - 12.10.219.0/24 American Express, Phoenix, Arizona
    - 23.20.0.0/14 AWS EC2 for us-east-1
    - 41.76.175.0/24 National Government of Kenya



The screenshot shows the bgp.tools website interface. At the top, there is a search bar with the text "Start here...". Below the search bar, the main content area displays the AS-SET "AS-MTU". The database is identified as "RIPE" and the full name is "AS-MTU". There are three tabs: "Overview" (selected), "Reverse", and "Raw". Below the tabs, the "Total Size" is listed as follows:

- Total Size
- 43823 ASNs
- 1400376 v4 Prefixes
- 713001 v6 Prefixes

# Excessively Large AS-SETs

- A popular tool for building BGP filter lists based on IRR data is bgpq4.

- <https://github.com/bgp/bgpq4>

- For AS-MTU, bgpq4 “-J” returns a Junos router configuration that is *almost 1.3 million lines long!*

```
$ bgpq4 -J1 eltel AS-MTU | wc -l
1294200
```

- We can use the -A option to aggregate routes, reducing the lines of configuration to only a third of a million, but *it is still a lot.*

```
$ bgpq4 -A1 eltel AS-MTU | wc -l
271171
```

- The routes contained in this AS-SET represent 1.8 billion unique IPv4 addresses out of a total possible 3 billion addresses currently in the IPv4 routing table.



# Excessively Large AS-SETs

*AS-MTU is not alone, nor anywhere near the worst!*

- So, what are the internet's largest (and most absurd) AS-SETs?
- Ben Cartwright-Cox, creator of Bgp.tools, ran the numbers.
- The biggest AS-SETs contain more ASNs than are in the global routing table (~83k).
- 2,192 AS-SETs expand to over 1,000 ASNs!

RIR	AS-SET	ASNs
RIPE	AS39533:AS-PEERS	102479
RIPE	AS-CLARANETDE-PEERINGS	102335
RADB	AS-ST1-IXPS	102332
RIPE	AS-MERKEL-PEERS	102313
RIPE	as-cloud-ix-pro	102305
RIPE	AS3326:AS-PEERS-DEE	102301
RIPE	AS-DECIX-V6	102300
RIPE	AS12732:AS-UPSTREAMS	102299
RIPE	AS-NFON-DECIX-PEERS-v4	102298
RIPE	AS-NFON-DECIX-PEERS-v6	102298

# Excessively Large AS-SETs

*Why is this a problem?*

- Our only hope to reduce harm from BGP mishaps is automation.
  - IRR data enables automated generation of allowlists.
- Excessively large AS-SETs defeats the purpose of an allowlist.
- Excessively large AS-SETs also breaks automation!
  - Requires large amount of data to be repeatedly transferred and stored.
  - Generates extremely large (and unusable) router configurations.
- Providers have had to create workarounds to deal with this IRR pollution.

# How can you help?

- 1) Check yourself:  
How many ASes/prefixes are included in your network's AS-SETs?
- 2) Be mindful of what you include in your AS-SET.  
The IRR databases are shared spaces. Not pollute them!
- 3) Only include things that your network is likely to transit.
- 4) Avoid including other AS-SETs in your AS-SET unless you absolutely must.
- 5) Help support newer path leak prevention technologies:
  - 1) Autonomous System Provider Authorization (ASPA)
  - 2) Router roles and Only-To-Customer Attribute (RFC9234)

# Thank you!

Doug Madory  
dmadory@kentik.com

 @DougMadory  
 in/DougMadory

