

The data for these exercises is located in `shared-data/bgp`. Details on the structure can be found here (<https://bgpstream.caida.org/docs/tools/bgpreader>) under *BGP Elem Format*. If you load the data into a DataFrame you can convert a column to time using `pl.to_datetime`. `group_by` works on time as well using a Grouper object.

1. AS Peers

We want to discover peers of our upstream provider.

Tools: polars, pysubnettree¹

Data: RIB from Routeviews route collectors at 4PM on May 13, 2026:

`ribs.routeviews.2026-05-13.csv`.

- (a) First, find out your public IP address. While this should be easy for our servers, just as a thought experiment consider how you could do the same for your home computer, which is likely behind a NAT.
- (b) `mobi8` should be reachable via `141.22.28.18`. Analyze the table dump noted under *Data*. Explain and implement one approach to figure out to which origin AS this IP address belongs to.
- (c) List the autonomous systems that peer with our ISP (based on your data set). Explain why this view is very likely incomplete.

Note that you can resolve the AS numbers via `whois` or `https://www.potaroo.net/bgp/iana/asn-ctl.txt`. A CSV-formatted snapshot of the potaroo list is located in `shared-data/asname/`, the separator is `"|"`.

2. RIPE RIS BGP Beacons Timing

We want to measure the timing behavior of the RIPE RIS BGP beacons².

Tools: polars, matplotlib

Data: Updates from RIPE RIS route collectors RRC00 and RRC15 for May 13, 2026:

`update.ris.rrc{00,15}.2026-05-13.csv.gz`.

- (a) Visualize the update patterns of the BGP beacons `84.205.64.0/24` and `93.175.153.0/24`.
- (b) Do the update patterns comply with the publicly documented announcement and withdrawal schedule?
- (c) Which time offsets do you observe between updates in the BGP dumps compared to the schedule? (Offsets are peer-specific.)

¹<https://github.com/zeek/pysubnettree>

²<https://www.ripe.net/analyse/internet-measurements/routing-information-service-ris/current-ris-routing-beacons>

3. BGP Zombies

When an IP prefix is withdrawn from its origin AS it should disappear from all routing tables—sooner or later.

Tools: polars

Data: RIB from RIPE RIS route collector RRC00 on May 13, 2026 at 08:00 UTC
`rib.ris.rrc00.2026-05-13.csv.gz`.

- (a) Explain the term *BGP zombie*.
- (b) Given the RIBs on the day, how could you check for zombies of the beacon prefix?
- (c) Apply your measurement methodology on the data set. *Warning: In the CSV, RIBs are prefixed with a **B**egin and **E**nd row. These rows have a different column length. Check the *ASCII Output Formats* in the *bgpreader* docs. In Polars, use `read_csv(..., truncate_ragged_lines=True)` (or pre-filter B/E rows).*