Detecting and Characterizing Internet Traffic Interception based on BGP Hijacking

Alberto Dainotti, Phillipa Gill, Chiara Orsini,
Alistair King, Vasco Asturiano, Matthew Luckie
alberto@caida.org

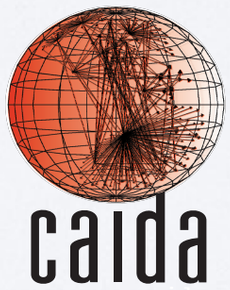

Center for Applied Internet Data Analysis

BGP MITM

BGP-based traffic interception

normal path

hijacked path

normal path used to complete the attack

S source (poisoned) **D** dest (hijacked prefix) **A** attacker

http://research.dyn.com/2013/11/mitm-internet-hijacking/

Center for Applied Internet Data Analysis
University of California San Diego

"HIJACKS" PROJECT

identify BGP-based MiTM

- NSF SaTC, TTP option, started Aug 2014, 3 years
- Collaborative project with Phillipa Gill at Stony Brook University

- Goals:
  - develop methodologies to detect interception
  - live monitoring
  - test/evaluate the system with real hijacks (thanks to the PEERING testbed - *http://peering.usc.edu*)
  - understand/quantify impact of events
  - log events, share data (e.g., through DHS PREDICT)

- *Happy to identify BGP hijacks in general*

HIJACKS

main components

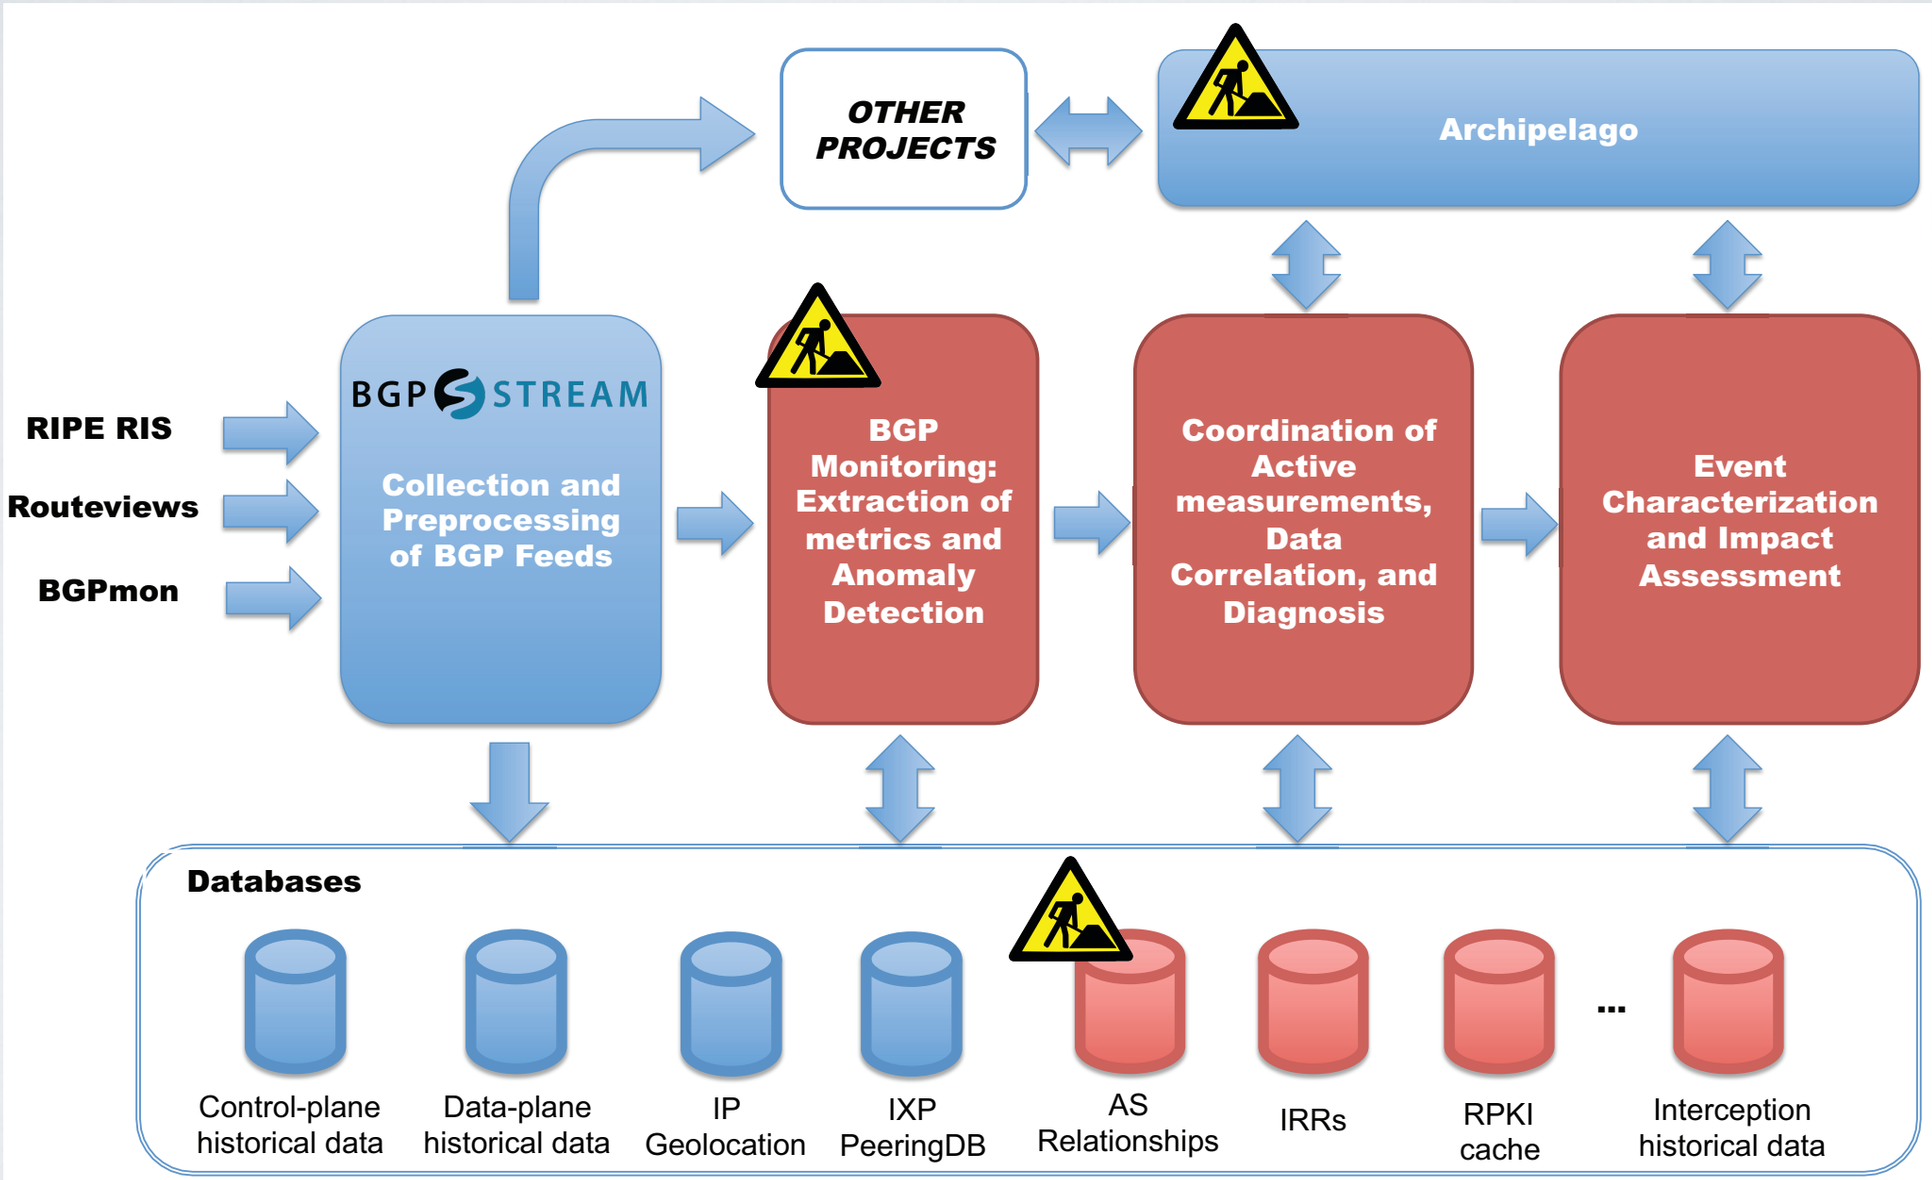

TWO MONITORING PHASES

BGP events are further analyzed through traceroutes

- Detect suspicious events using criteria based on **BGP data**
  - MOAS
  - valley free violations
  - new edges
  - inconsistent prepending
  - …
- We analyze these cases with on-demand **traceroutes** from **Archipelago** probes

ARK

two probing schemes

- Ark's Topo **on Demand** to do traceroutes from *all* Ark probes towards prefixes associated with suspicious events

- **Daily continuous** traceroutes towards all prefixes
  - -target prefix list: updated every day. 1 week sliding window
  - -Purpose:
    - comparison against ad hoc traceroutes
    - infer additional AS relationships
    - historical data analysis

ARK

researchy topics

- Exploit co-location with BGP monitors from RouteViews and RIPE RIS
  - Out of 200 ASes providing a full IPv4 routing table, 20 host an ARK vantage point
    - we plan to increase this fraction
    - how would *you* use it?

- Automatically and accurately translate traceroutes to inferred AS paths
    - collaboration with Matthew Luckie

A TYPICAL SCENARIO

AS-A announces prefix-d, normally announced by AS-D

normal path

hijacked path

normal path
used to complete
the attack

A TYPICAL SCENARIO

AS-A announces prefix-d, normally announced by AS-D

- BGP will observe a MOAS

- Traceroutes (translated in AS paths.. let's call them "*IP AS-paths*") will show:

1. all VPs: IP AS-path will **end at AS-D**
2. VPs co-located with BGP monitor + following hijacked path: **BGP AS-path != IP AS-path**. The first portion of the IP-AS path will match the BGP AS-path
3. VPs following hijacked path: **AS-A** is in the middle of the IP AS-path
4. VPs following hijacked path: IP AS-path typically is **longer than the historical** ones

IP TO AS PATHS

Infer AS paths from traceroutes

In the scope of this project, there are some interesting variations to the classic problem:

- *constraint*: we can't use BGP's AS Paths as ground truth

- *pro*: we can tolerate uncertainty on some hops: looking for large mismatches *[cases 2 and 4]*

- *pro*: we may not care too much about consistent errors *[case 4]*