IETF 94 - Technical Plenary
4th Nov 2015, Yokohama, JP

Measuring and Monitoring BGP

Alberto Dainotti,
alberto@caida.org

caida

Center for Applied Internet Data Analysis
University of California, San Diego

MEASURING BGP

Why?

BGP is the central nervous system of the Internet

BGP's design is known to contribute to issues in:

- **Availability**
  - -Labovitz et al. *"Delayed Internet Routing Convergence"*, IEEE/ACM Trans. Netw., 2001.
  - -Varadhan et al. *"Persistent Route Oscillations in Inter-domain Routing"*. Computer Networks, 2000.
  - -Katz-Bassett et al. *"LIFEGUARD: Practical Repair of Persistent Route Failures"*, SIGCOMM, 2012.
- **Performance**
  - -Spring et al. *"The Causes of Path Inflation"*. SIGCOMM, 2003.
- **Security**
  - -Zheng et al. *"A Light-Weight Distributed Scheme for Detecting IP Prefix Hijacks in Realtime"*. SIGCOMM, 2007.

Need to *engineer* protocol evolution!

MEASURING BGP

Why?

Defining problems and make ***protocol engineering*** decisions through realistic evaluations is difficult also because **we know little about the structure and dynamics of the BGP ecosystem!**

- AS-level topology
  - Gregori et al. *"On the incompleteness of the AS-level graph: a novel methodology for BGP route collector placement"*, IMC 2012
- AS relationships
  - Giotsas et al. *"Inferring Complex AS Relationships"*, IMC 2014
- AS interactions: driven by relationships, policies, network conditions, operator updates
  - Anwar et al. *"Investigating Interdomain Routing Policies in the Wild "*, IMC 2015
  - Lychev et al. *"BGP Security in Partial Deployment: Is the Juice Worth the Squeeze?"*, SIGCOMM 2013

MEASURING BGP

two issues - somehow related

1. Literature shows that **we need more/better data**
    - more info from the protocol/routers

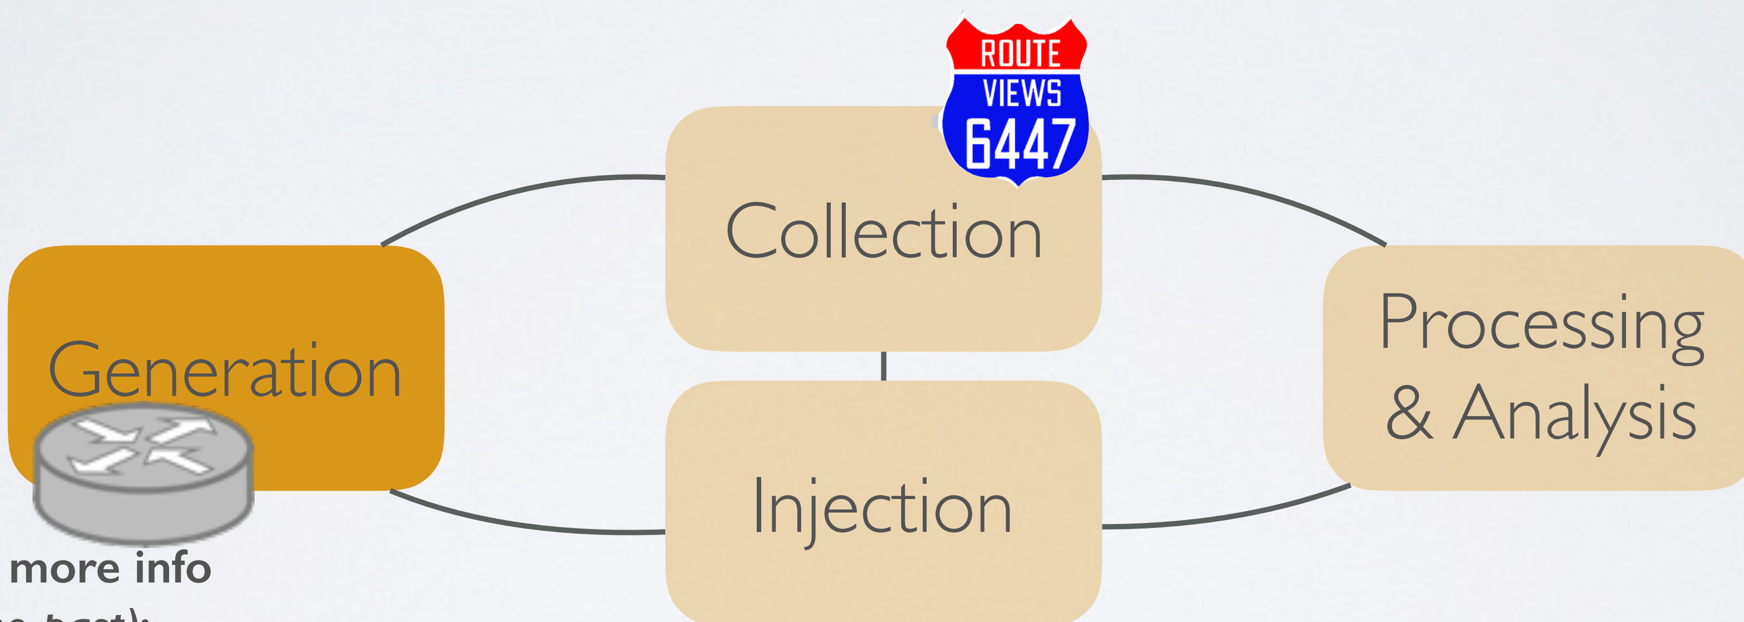

Attempts to generate more info
***(not much traction in the past)*:**

- **`RFC 4384 BGP Communities for Data Collection`**
- **`draft-ymbk-grow-bgp-collector-communities`**

MEASURING BGP

two issues - somehow related

1. Literature shows that **we need more/better data**
    - more info from the protocol/routers, more collectors,

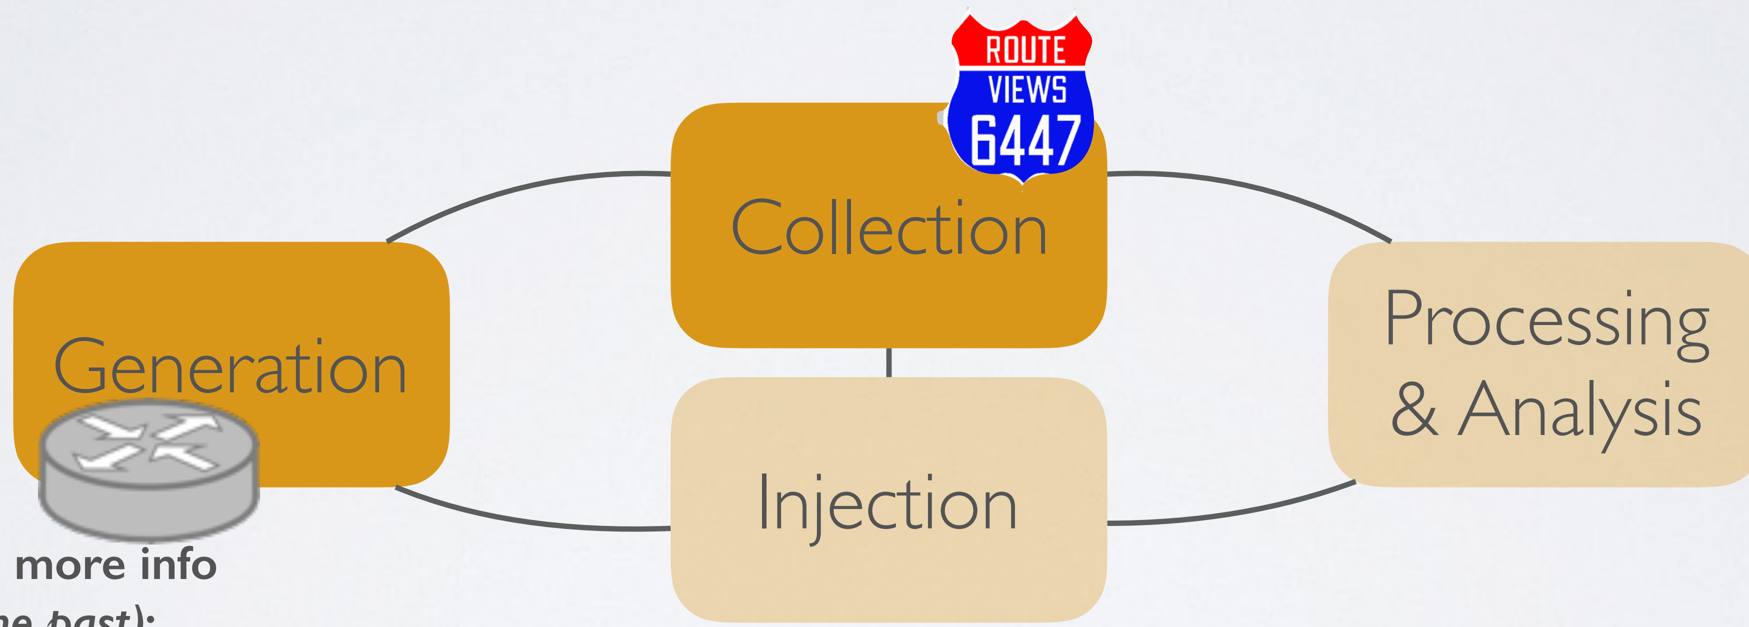

Attempts to generate more info
***(not much traction in the past)*:**

- **`RFC 4384 BGP Communities for Data Collection`**
- **`draft-ymbk-grow-bgp-collector-communities`**

MEASURING BGP

two issues - somehow related

1. Literature shows that **we need more/better data**
   - more info from the protocol/routers, more collectors, more experimental testbeds, …

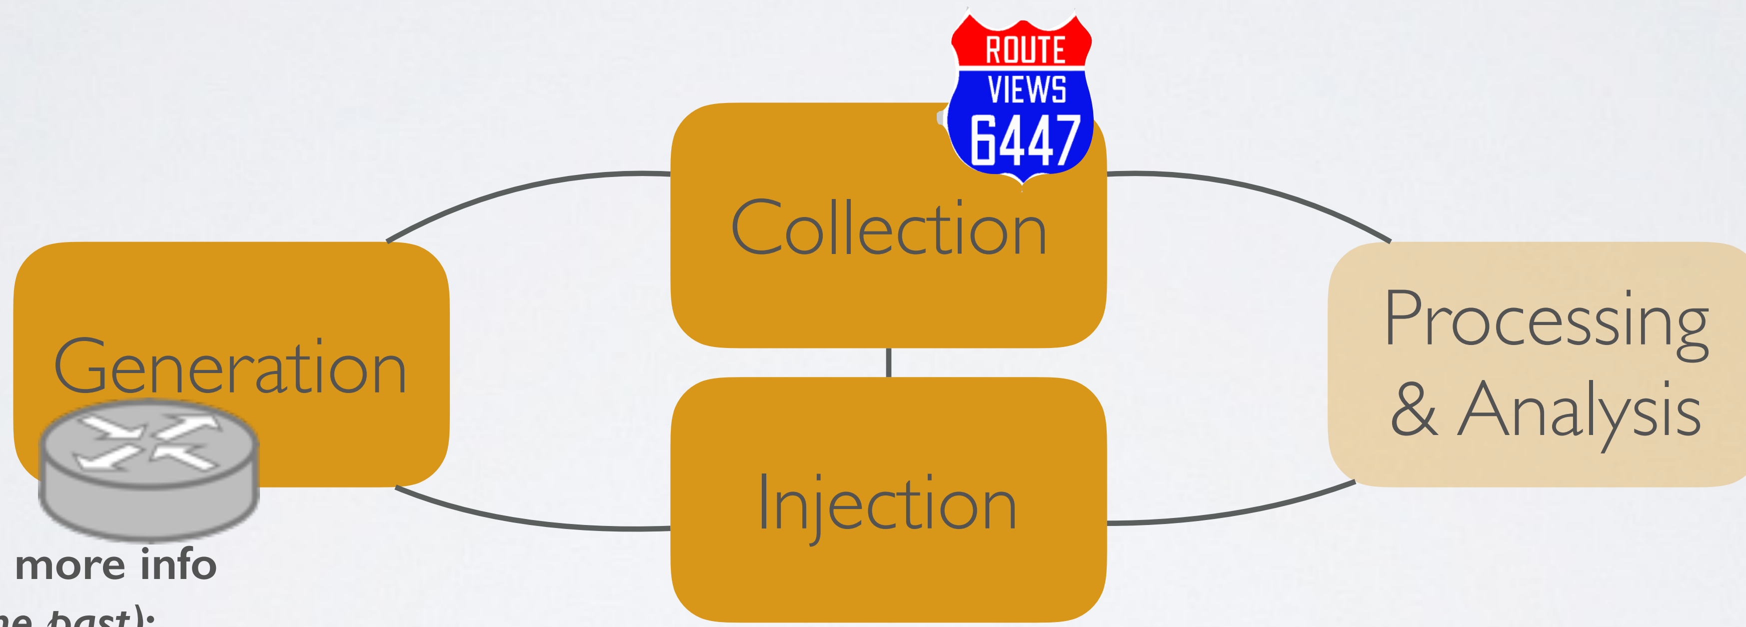

Attempts to generate more info *(not much traction in the past)*:

- **`RFC 4384 BGP Communities for Data Collection`**
- **`draft-ymbk-grow-bgp-collector-communities`**

Inject/Receive Routes & Traffic.
PEERING - http://peering.usc.edu
Schlinker et al. *"PEERING: An AS for Us"*, HotNets 2014

MEASURING BGP

two issues - somehow related

1. Literature shows that **we need more/better data**
   - more info from the protocol/routers, more collectors, more experimental testbeds, …

2. But we also **need better tools to learn from the data**
   - to make data analysis: *easier, faster, able to cope with BIG and heterogeneous data*
   - to monitor BGP in near-realtime
   - tightening data collection, processing, visualization, …

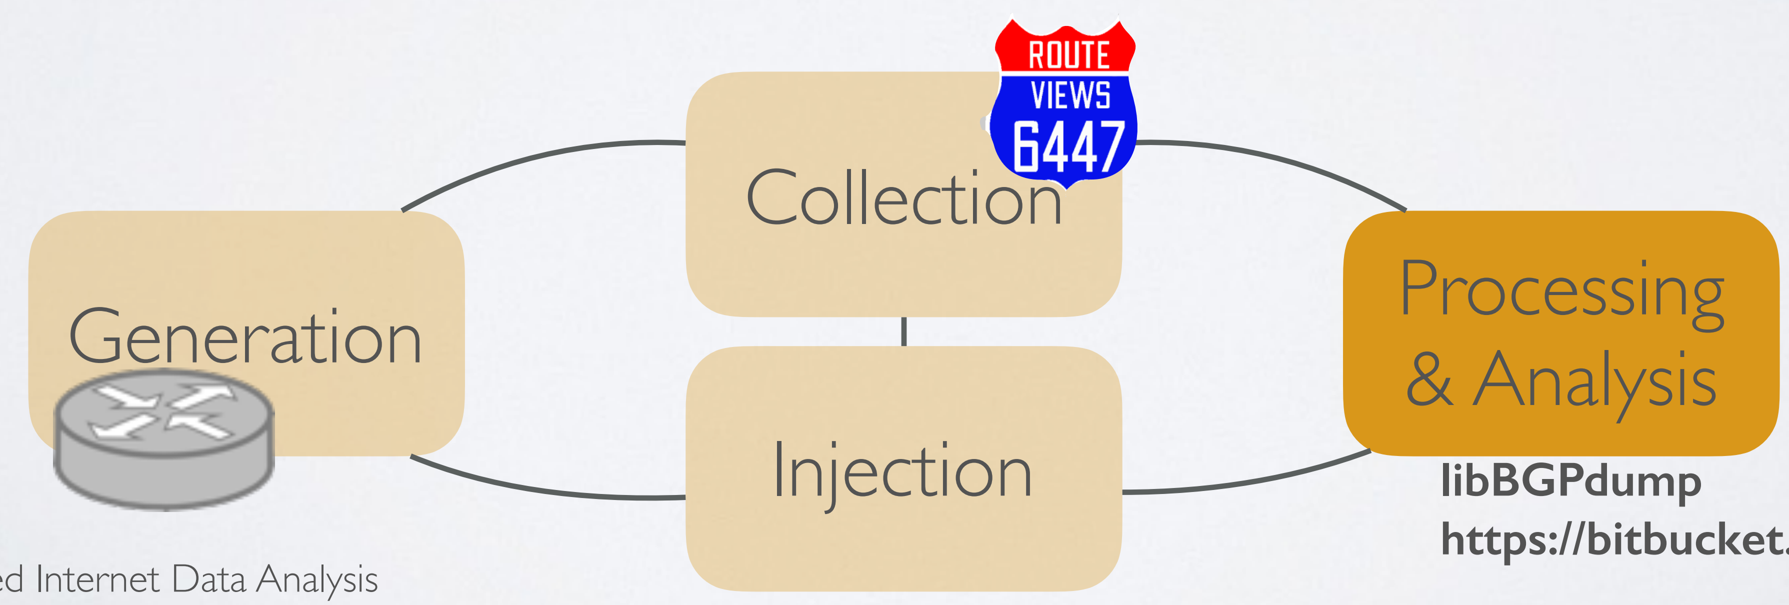

libBGPdump
https://bitbucket.org/ripencc/bgpdump

BGP EVENTS & DYNAMICS

IODA: Detection and Analysis of Internet Outages

- Country-level Internet Blackouts during the Arab Spring

Dainotti et al. "Analysis of Country-wide Internet Outages Caused by Censorship" IMC 2011

EGYPT, JAN 2011
GOVERNMENT ORDERS TO SHUT DOWN THE INTERNET

- Natural disasters affecting the infrastructure

Dainotti et al. "Extracting Benefit from Harm: Using Malware Pollution to Analyze the Impact of Political and Geophysical Events on the Internet" SIGCOMM CCR 2012

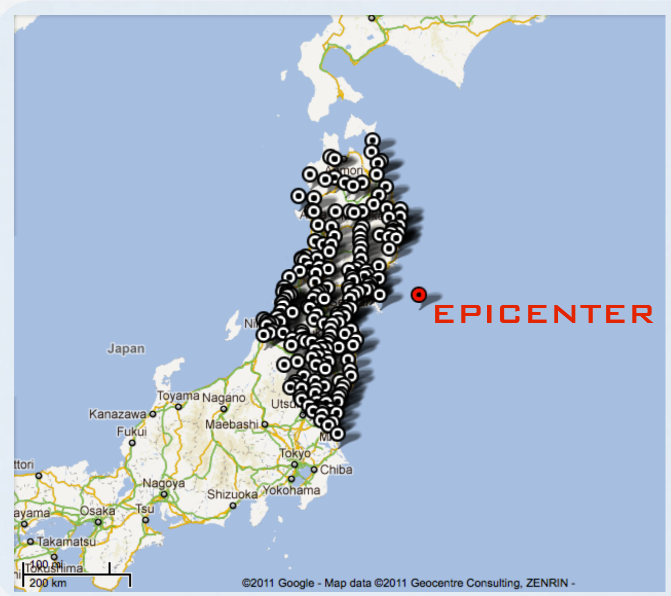

JAPAN, MAR 2011
EARTHQUAKE OF MAGNITUDE 9.0

BGP EVENTS & DYNAMICS

IODA: Detection and Analysis of Internet Outages

Country-wide Internet outages in Iraq that the government ordered in conjunction with the ministerial preparatory exams - Jul 2015

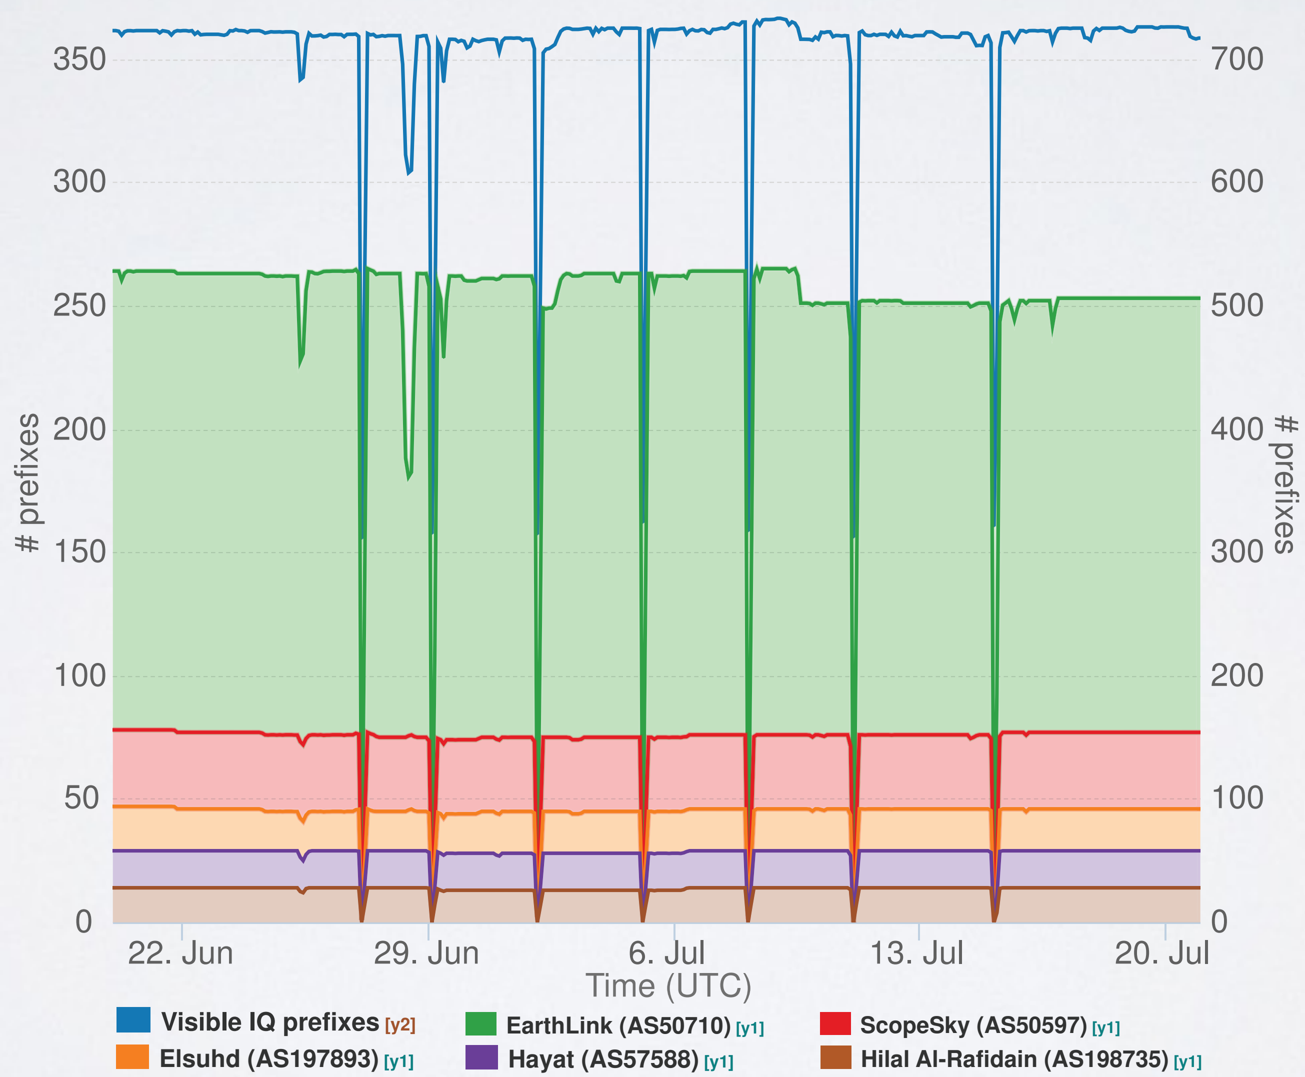

Center for Applied Internet Data Analysis
University of California San Diego

www.caida.org/funding/ioda/

BGP EVENTS & DYNAMICS

IODA: Detection and Analysis of Internet Outages

Outage of AS11351(Time Warner Cable LLC)
September 30, 2015

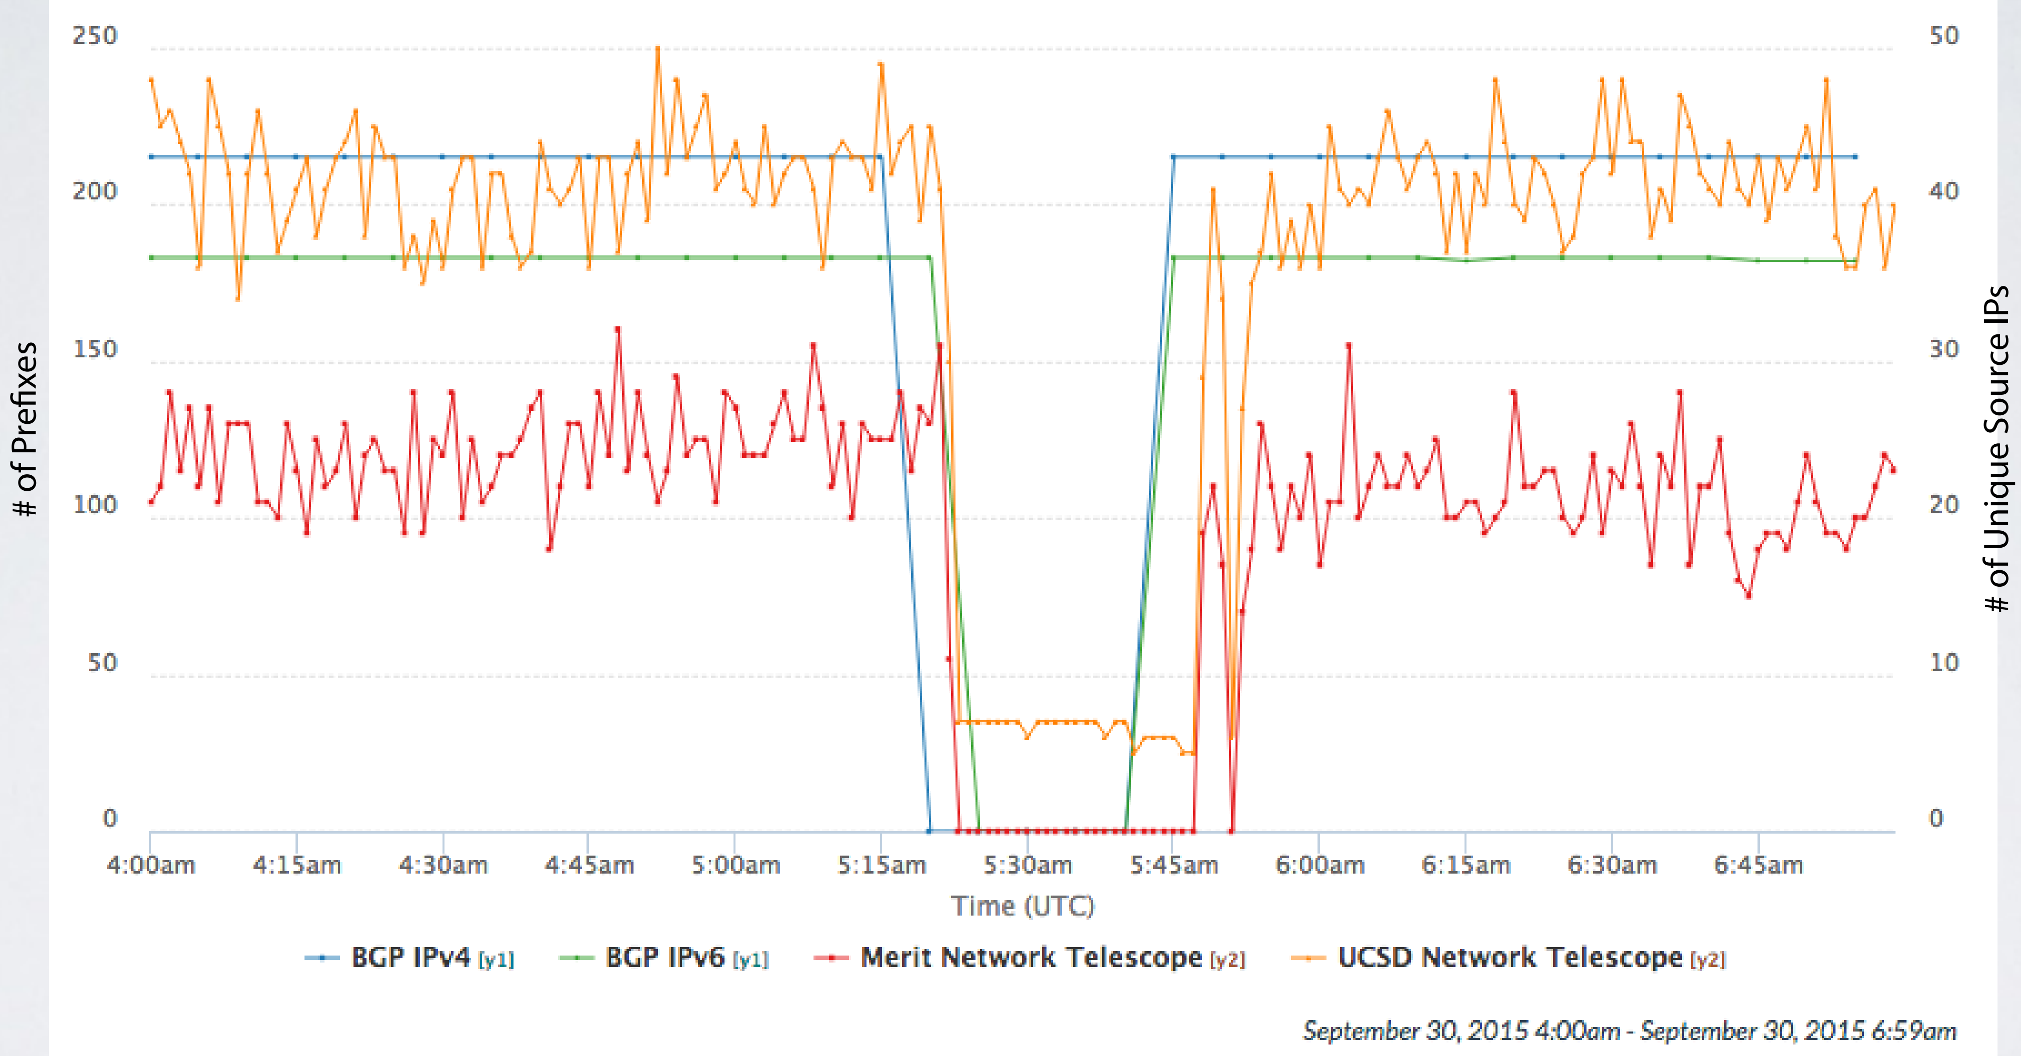

BEFORE IODA

post-event manual analysis

Egypt, Jan 2011
Government orders to shut down the Internet

4 months of work

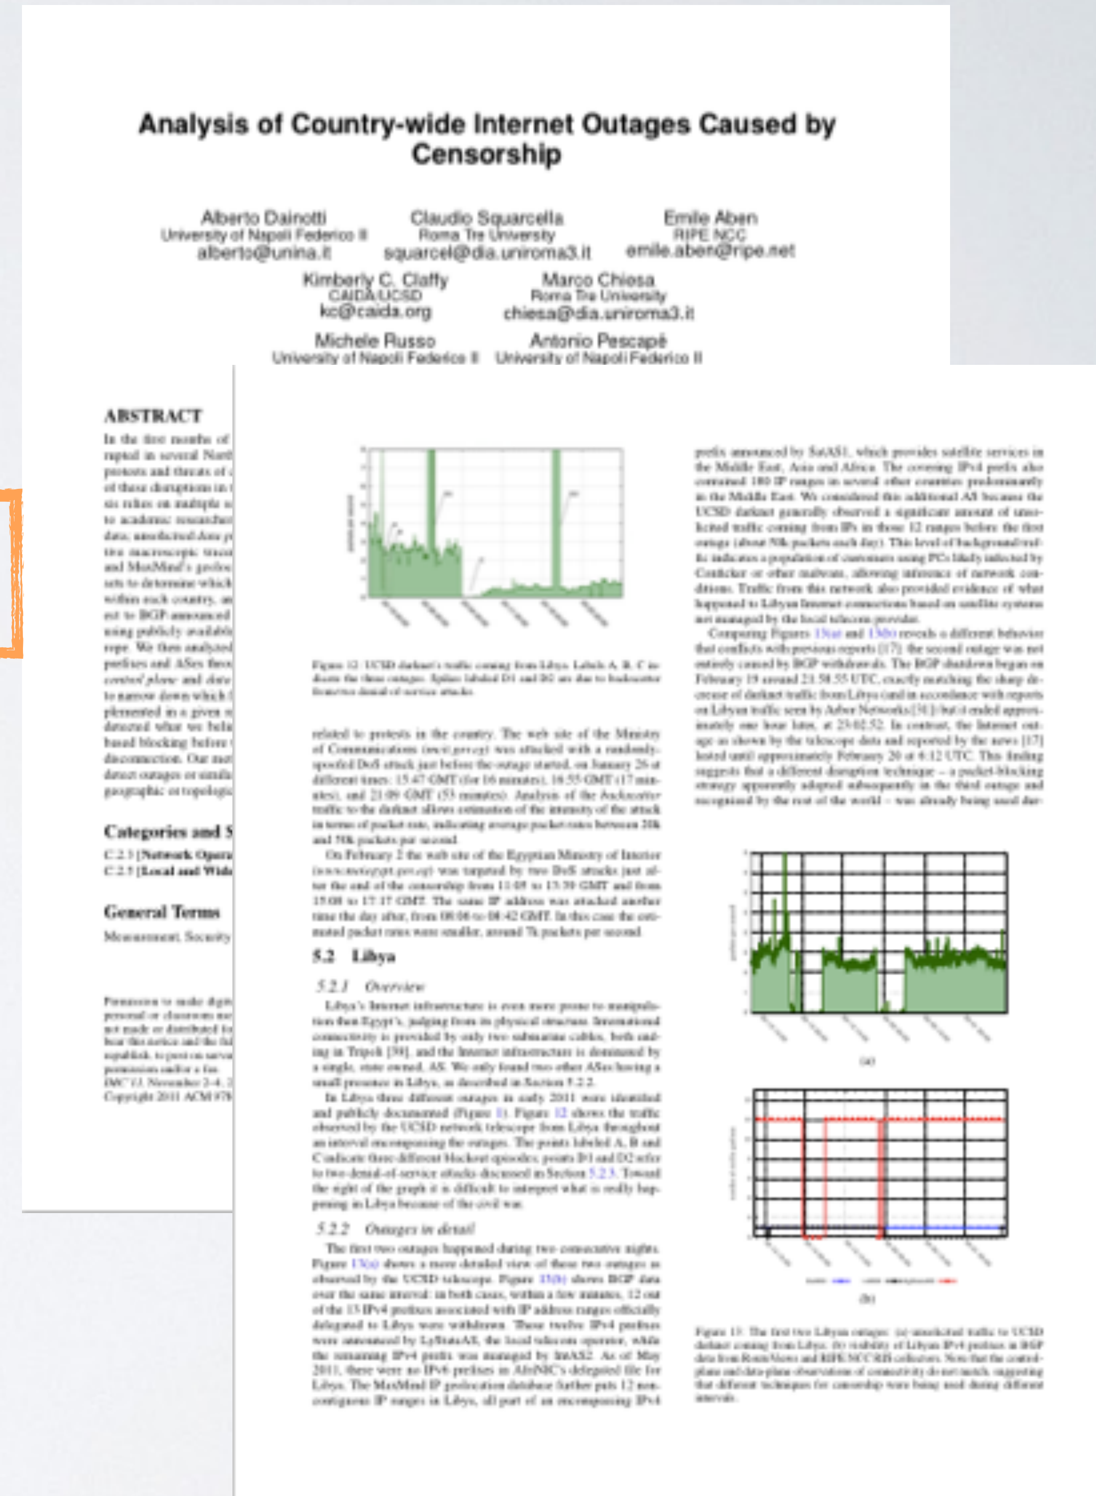

Analysis of Country-wide Internet Outages Caused by Censorship

Alberto Dainotti, University of Napoli Federico II, alberto@unina.it
Claudio Squarcella, Roma Tre University, squarcel@dia.uniroma3.it
Emile Aben, RIPE NCC, emile.aben@ripe.net
Kimberly C. Claffy, CAIDA/UCSD, kc@caida.org
Marco Chiesa, Roma Tre University, chiesa@dia.uniroma3.it
Michele Russo, University of Napoli Federico II
Antonio Pescapè, University of Napoli Federico II

ABSTRACT

Categories and S

General Terms

5.2 Libya

5.2.1 Overview

5.2.2 Outages in detail

Dainotti et al. "Analysis of Country-wide Internet Outages Caused by Censorship" IMC 2011

IODA TODAY

live Internet monitoring

Last Christmas we made it possible for anybody to follow the North Korean disconnection almost live

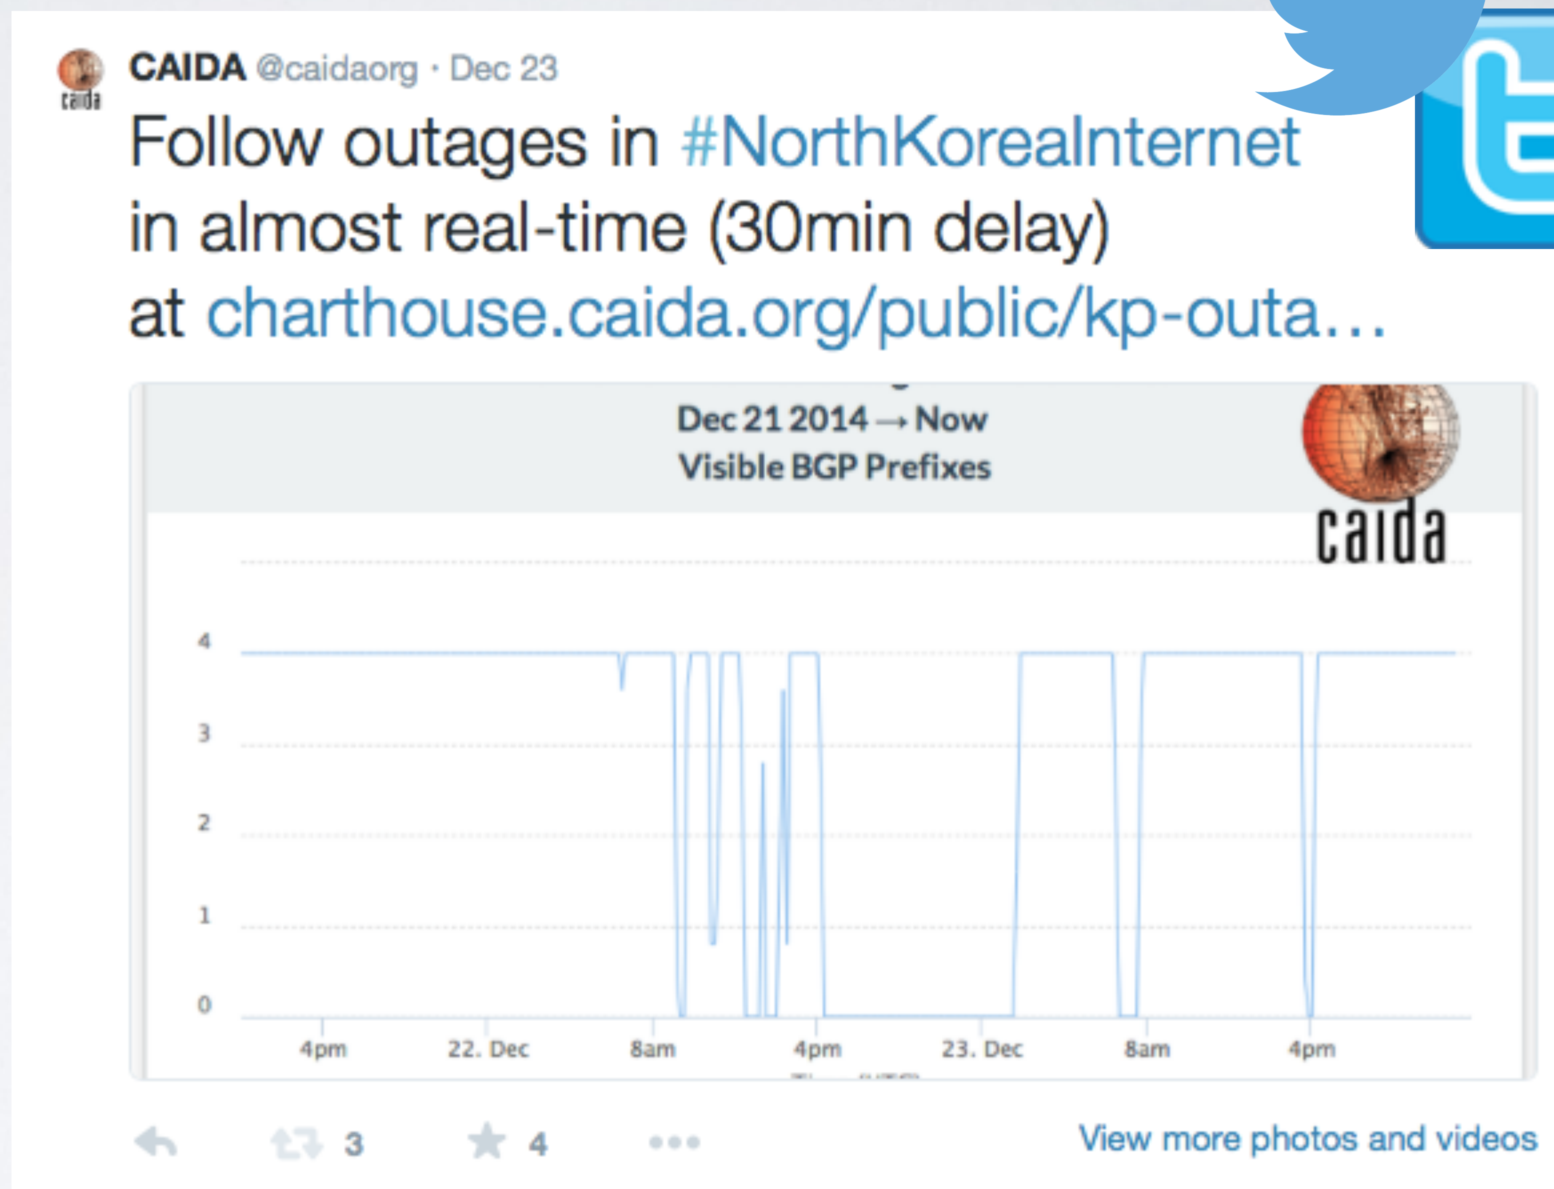

https://charthouse.caida.org/public/kp-outage

MEASURING BGP

two issues - somehow related

1. Literature shows that **we need more/better data**
   - more info from the protocol/routers, more collectors, more experimental testbeds, …

2. But we also **need better tools to learn from the data**
   - to make data analysis: *easier, faster, able to cope with BIG and heterogeneous data*
   - to monitor BGP in near-realtime
   - tightening data collection, processing, visualization, …

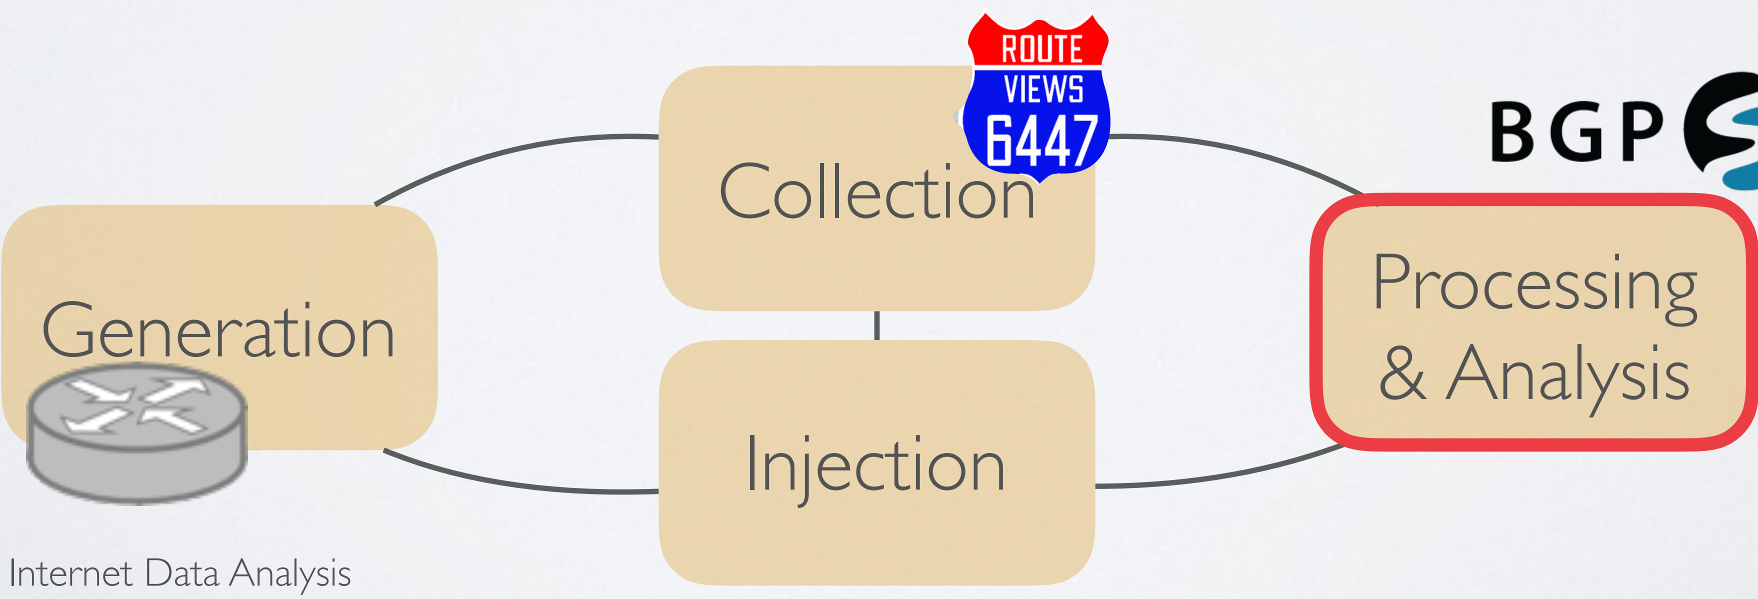

BGP STREAM

BGPSTREAM

bgpstream.caida.org

- A software framework for **historical** and **live** BGP data analysis
- Design goals:
  - -Efficiently deal with large amounts of distributed BGP data
  - -Offer a time-ordered data stream of data from heterogeneous sources
  - -Support near-realtime data processing
  - -Target a broad range of applications and users
  - -Scalable
  - -Easily extensible

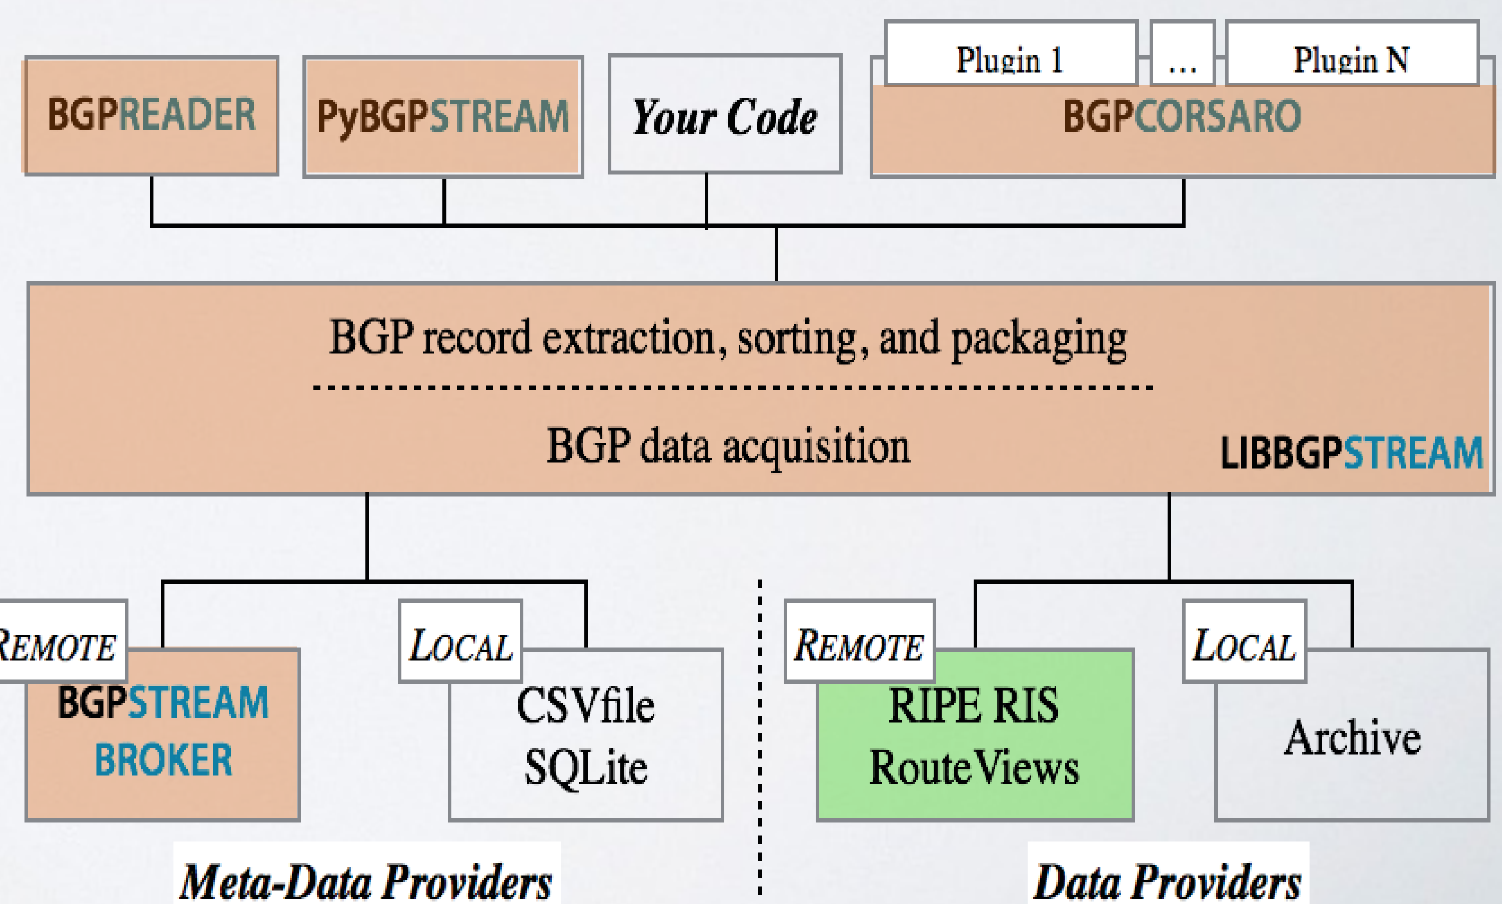

PYBGPSTREAM

Example: studying AS path inflation

How many AS paths are longer than the shortest path between two ASes due to routing policies? (directly correlates to the increase in BGP convergence time)

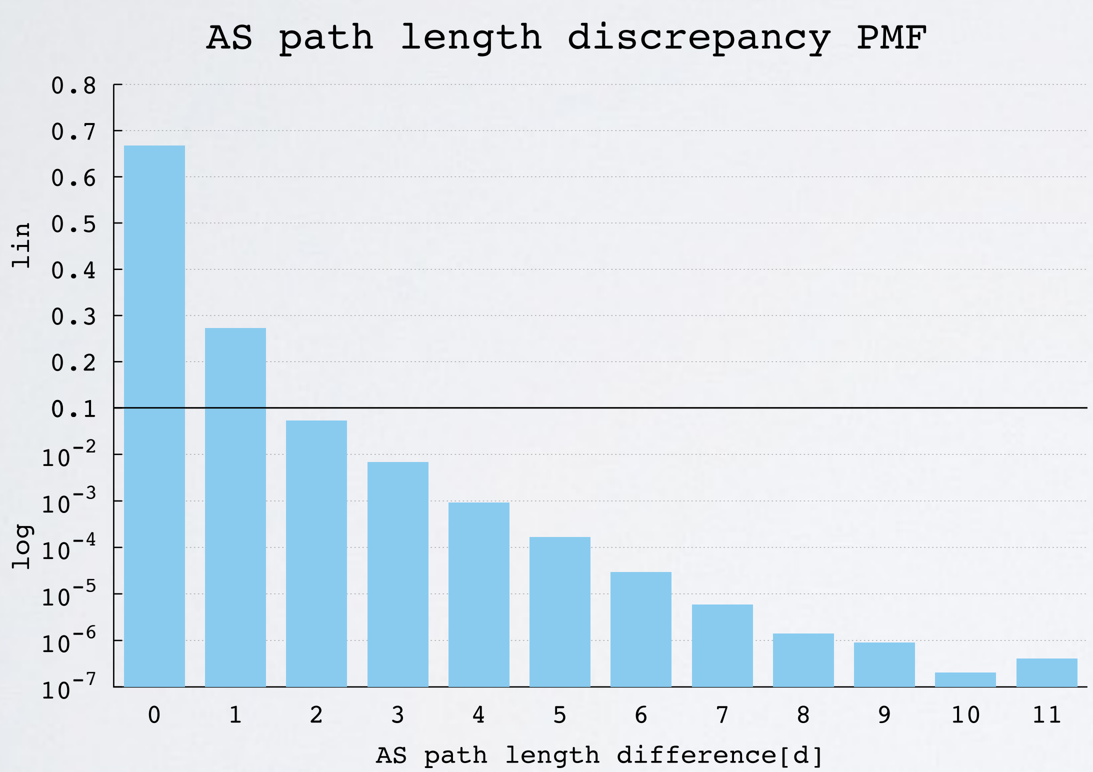


```python
from _pybgpstream import BGPStream, BGPRecord, BGPElem
from collections import defaultdict
from itertools import groupby
import networkx as nx

stream = BGPStream()
as_graph = nx.Graph()
rec = BGPRecord()
bgp_lens = defaultdict(lambda: defaultdict(lambda: None))
stream.add_filter('record-type','ribs')
stream.add_interval_filter(1438415400,1438416600)
stream.start()

while(stream.get_next_record(rec)):
    elem = rec.get_next_elem()
    while(elem):
        monitor = str(elem.peer_asn)
        hops = [k for k, g in groupby(elem.fields['as-path'].split(" "))]
        if len(hops) > 1 and hops[0] == monitor:
            origin = hops[-1]
            for i in range(0,len(hops)-1):
                as_graph.add_edge(hops[i],hops[i+1])
            bgp_lens[monitor][origin] = \
                min(filter(bool,[bgp_lens[monitor][origin],len(hops)]))
        elem = rec.get_next_elem()
for monitor in bgp_lens:
    for origin in bgp_lens[monitor]:
        nxlen = len(nx.shortest_path(as_graph, monitor, origin))
        print monitor, origin, bgp_lens[monitor][origin], nxlen
```

30 LINES OF PYTHON CODE

BGPCORSARO

Example: monitor your own address space on BGP

*The "**prefix-monitor**" plugin (distributed with source) monitors a set of IP ranges as they are seen from BGP monitors distributed worldwide:*

- *how many prefixes reachable*
- *how many origin ASes*
- *generates detailed logs*

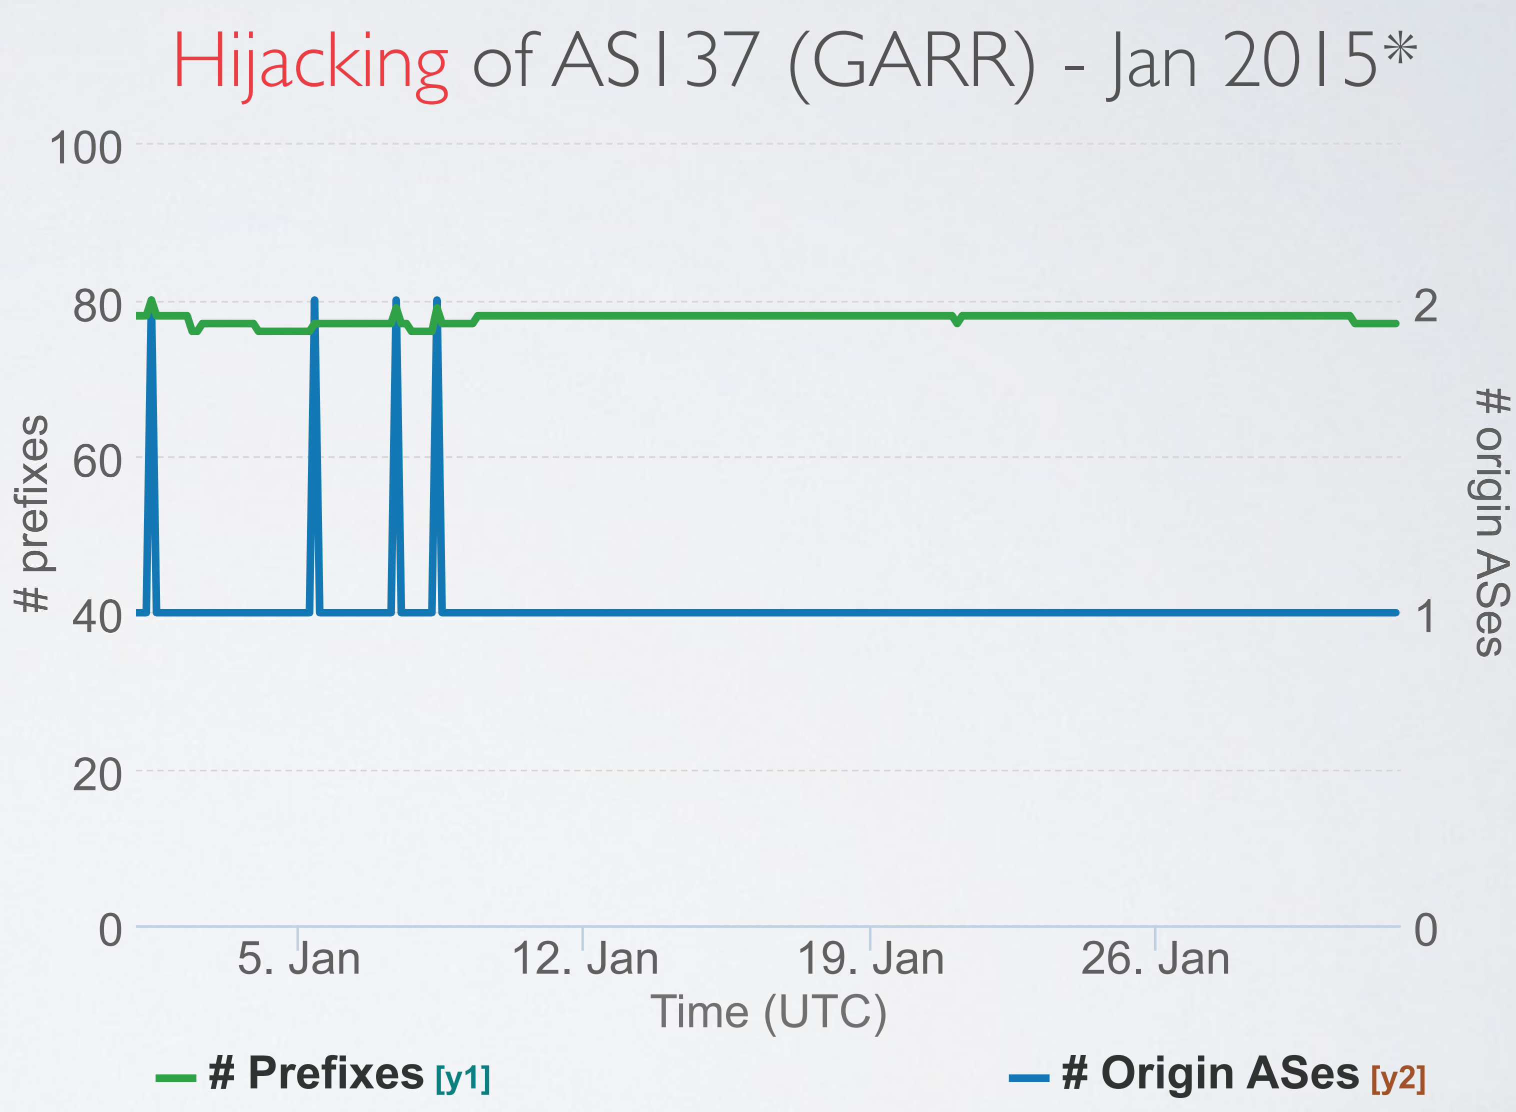

**Originally discovered by Dyn: http://research.dyn.com/2015/01/vast-world-of-fraudulent-routing/*

ANOTHER SUPPORTED PROJECT

Hijacks: detection of MITM BGP attacks

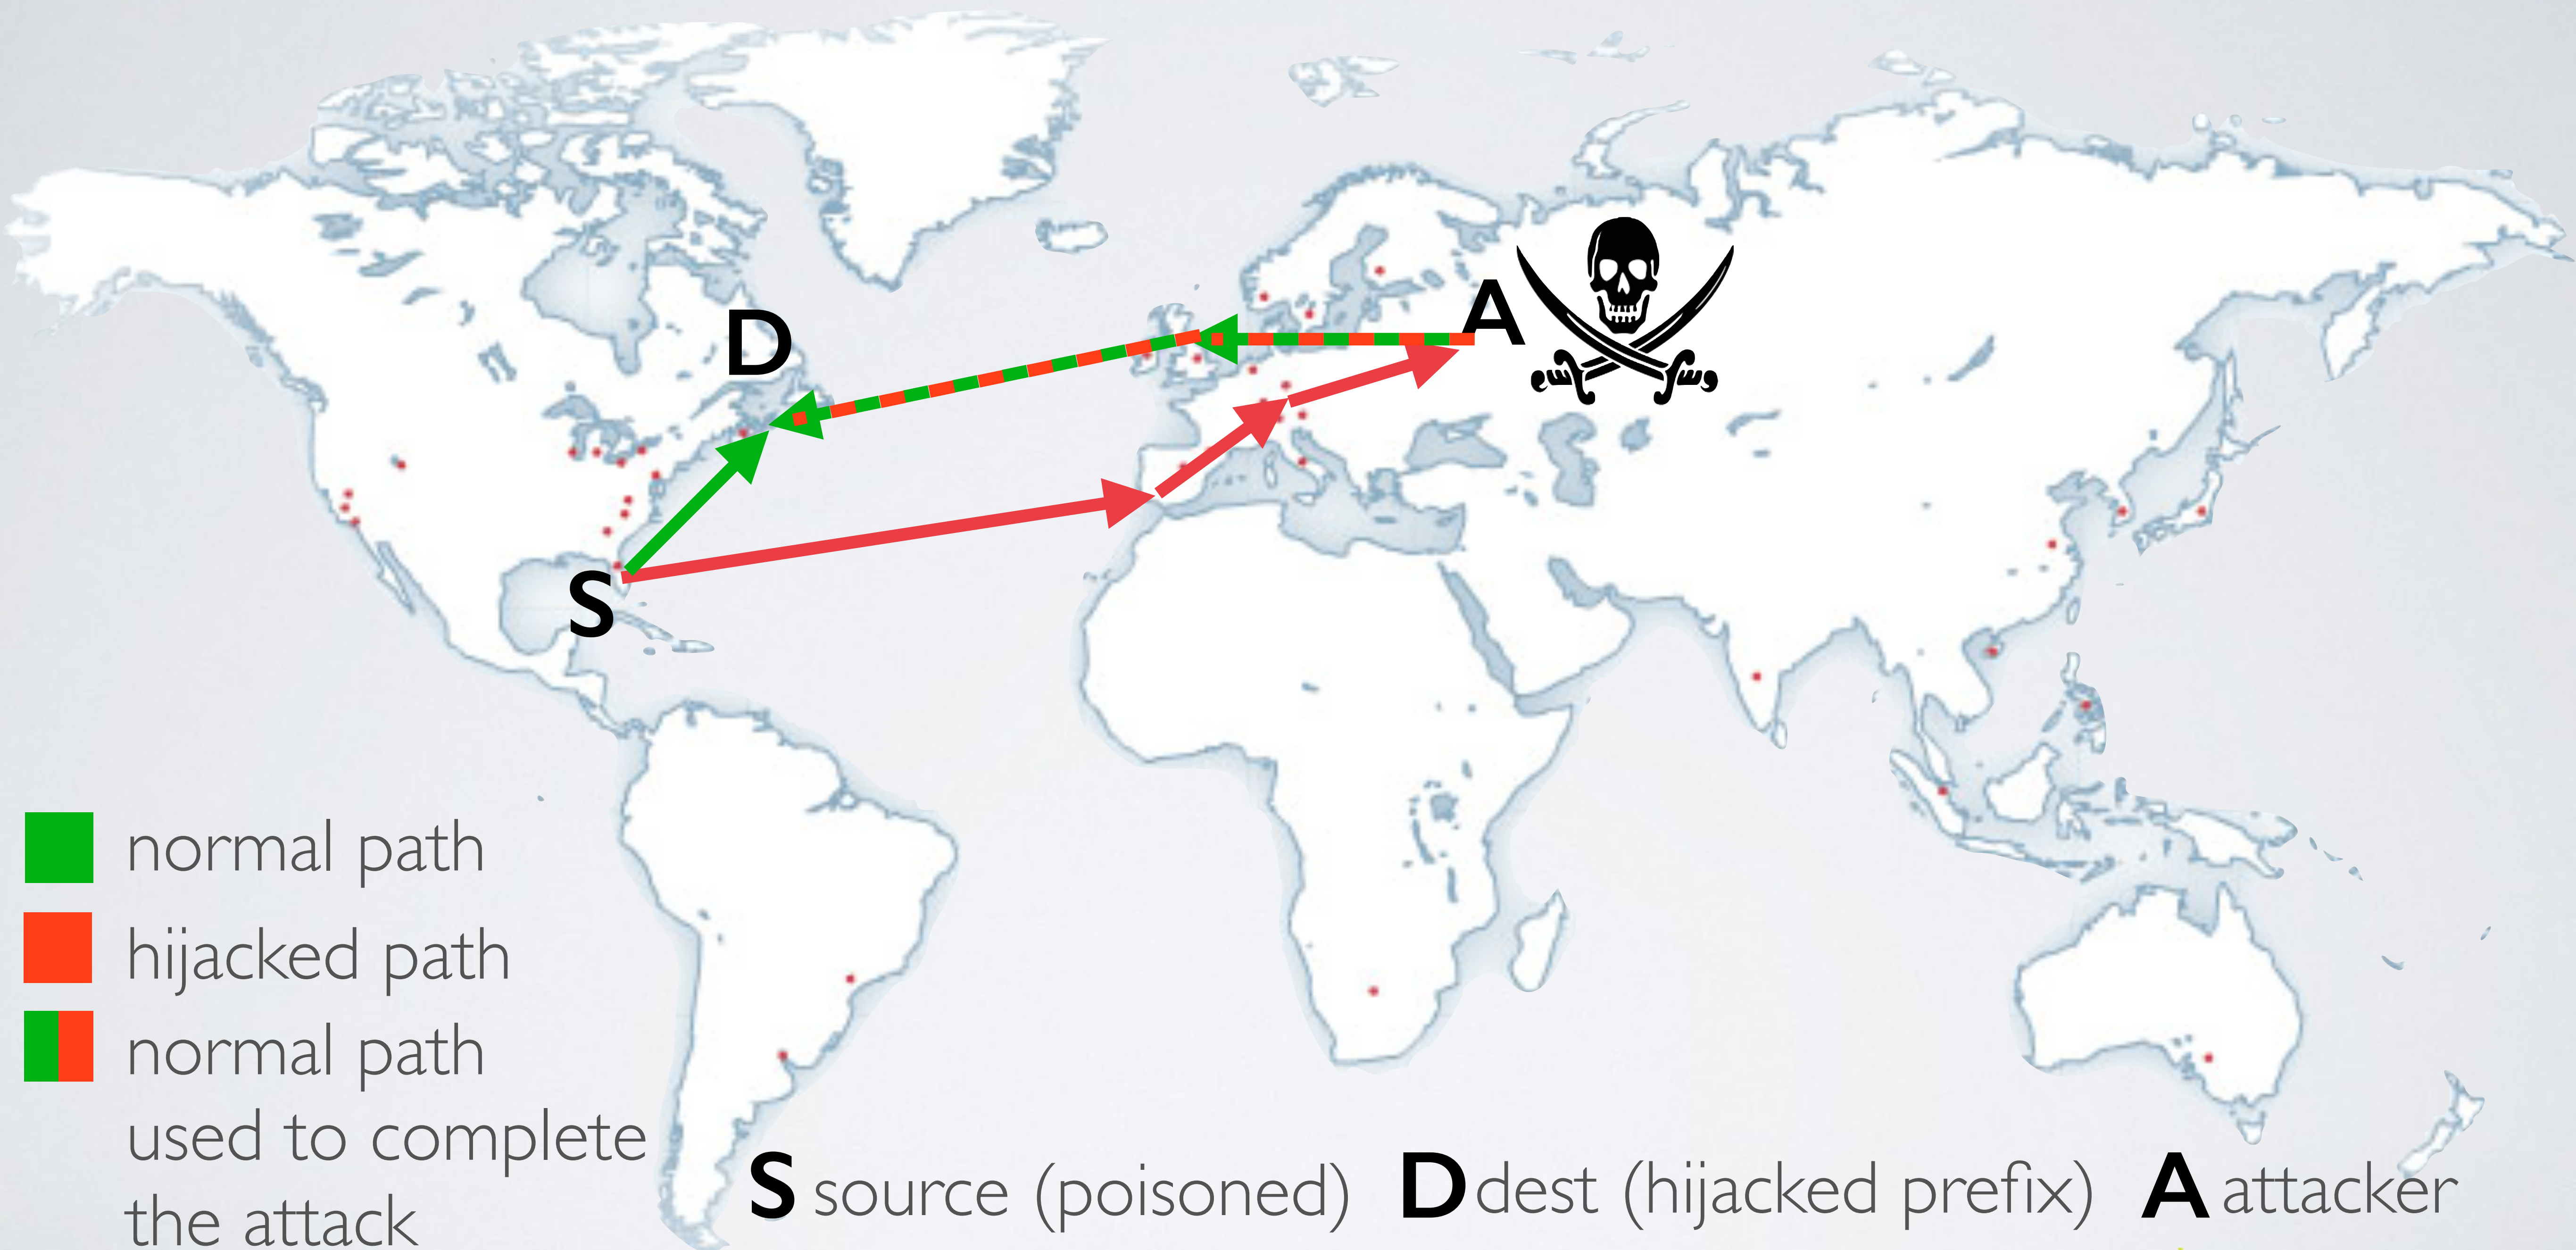

- normal path
- hijacked path
- normal path used to complete the attack

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

Stony Brook University

www.caida.org/funding/hijacks/

COMCAST

ANOTHER SUPPORTED PROJECT

Hijacks: detection of MITM BGP attacks

Research informed by (and tested with) ***data in the wild***

Live BGP measurements trigger on-demand dataplane measurements (e.g., traceroutes) ***during*** *a suspicious event.*

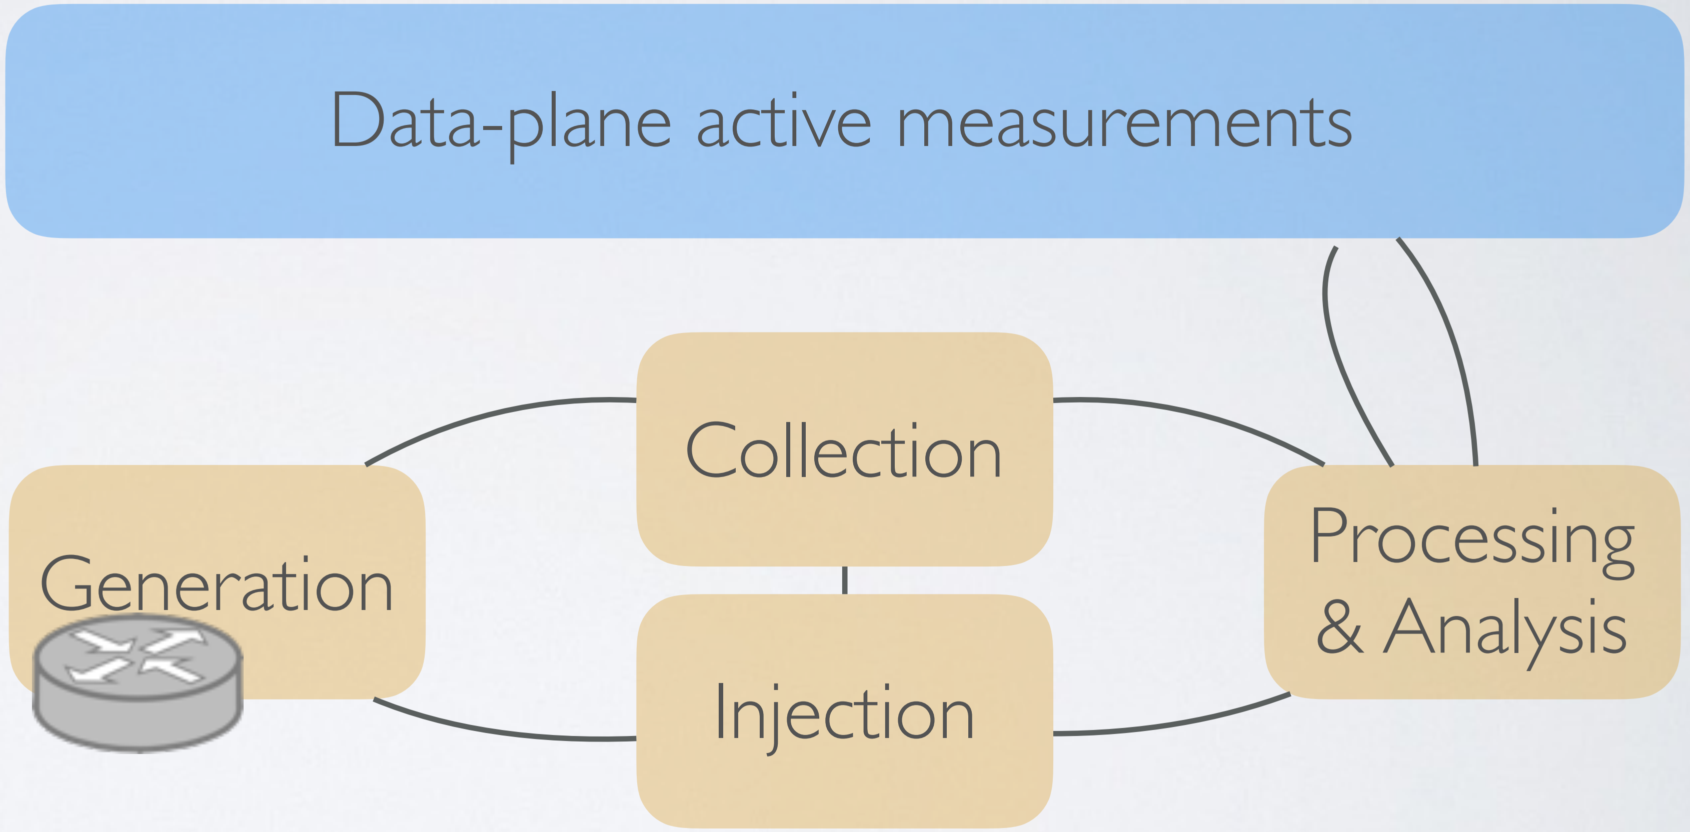

www.caida.org/funding/hijacks/

BGP HACKATHON - FEB 2016

*theme: "**live** BGP measurements & monitoring"*

Improve/Integrate tools to study the BGP eco-system. Target practical problems: topology, hijacks, outages, RPKI deployment, path inflation, circuitous paths, policies, relationships, visualize dynamics, …

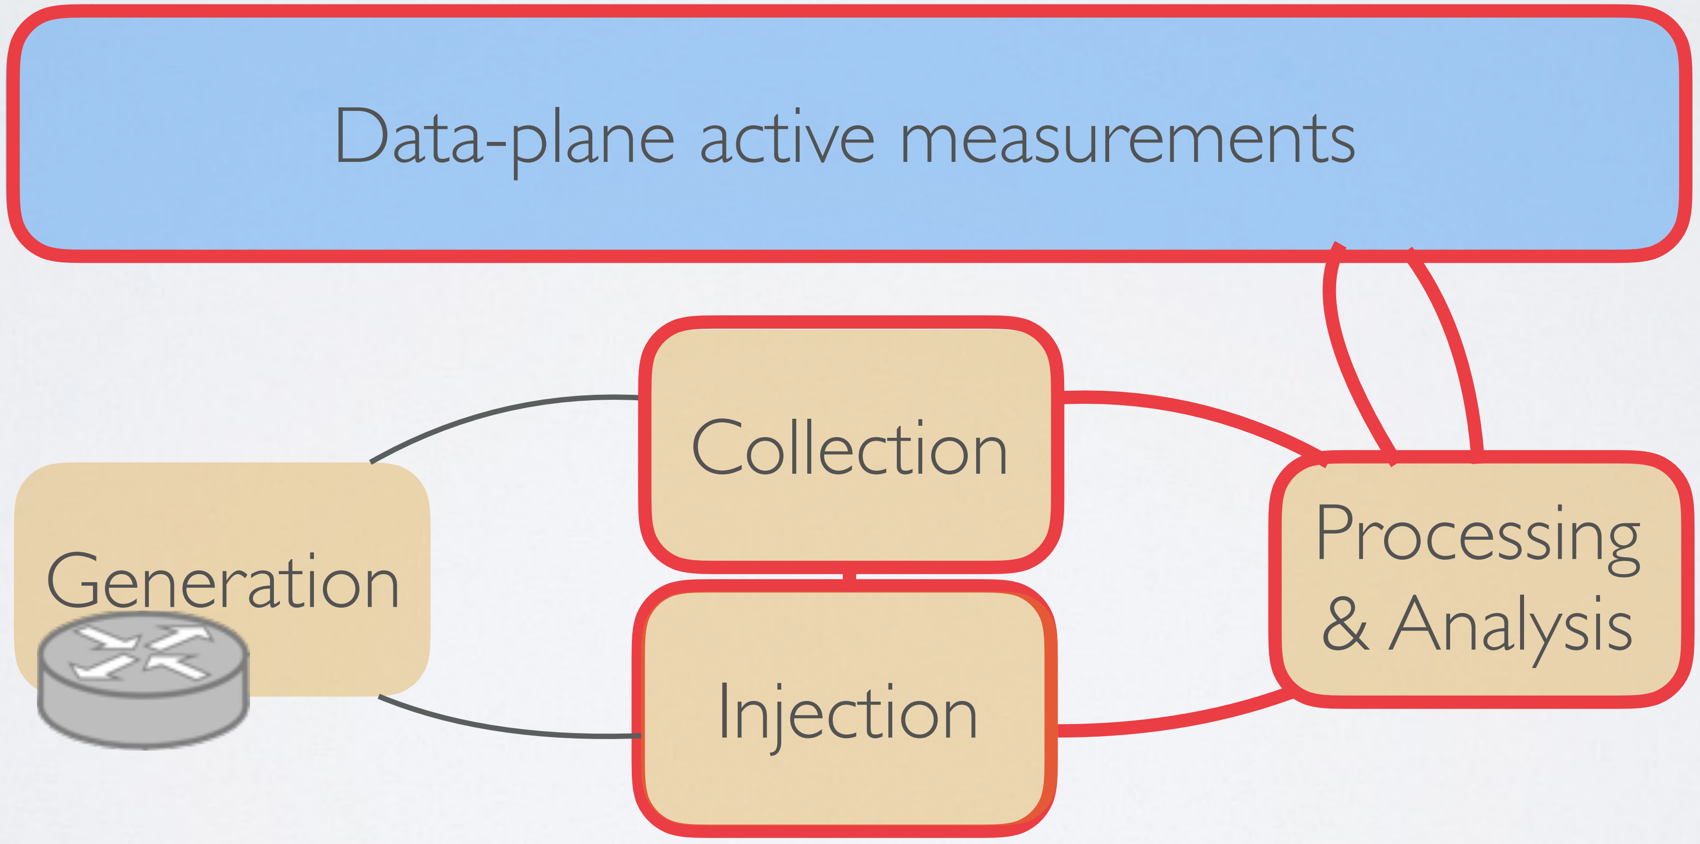

BGP HACKATHON - FEB 2016

*theme: “**live** BGP measurements & monitoring”*

We will provide a rich toolbox and “live” data access:

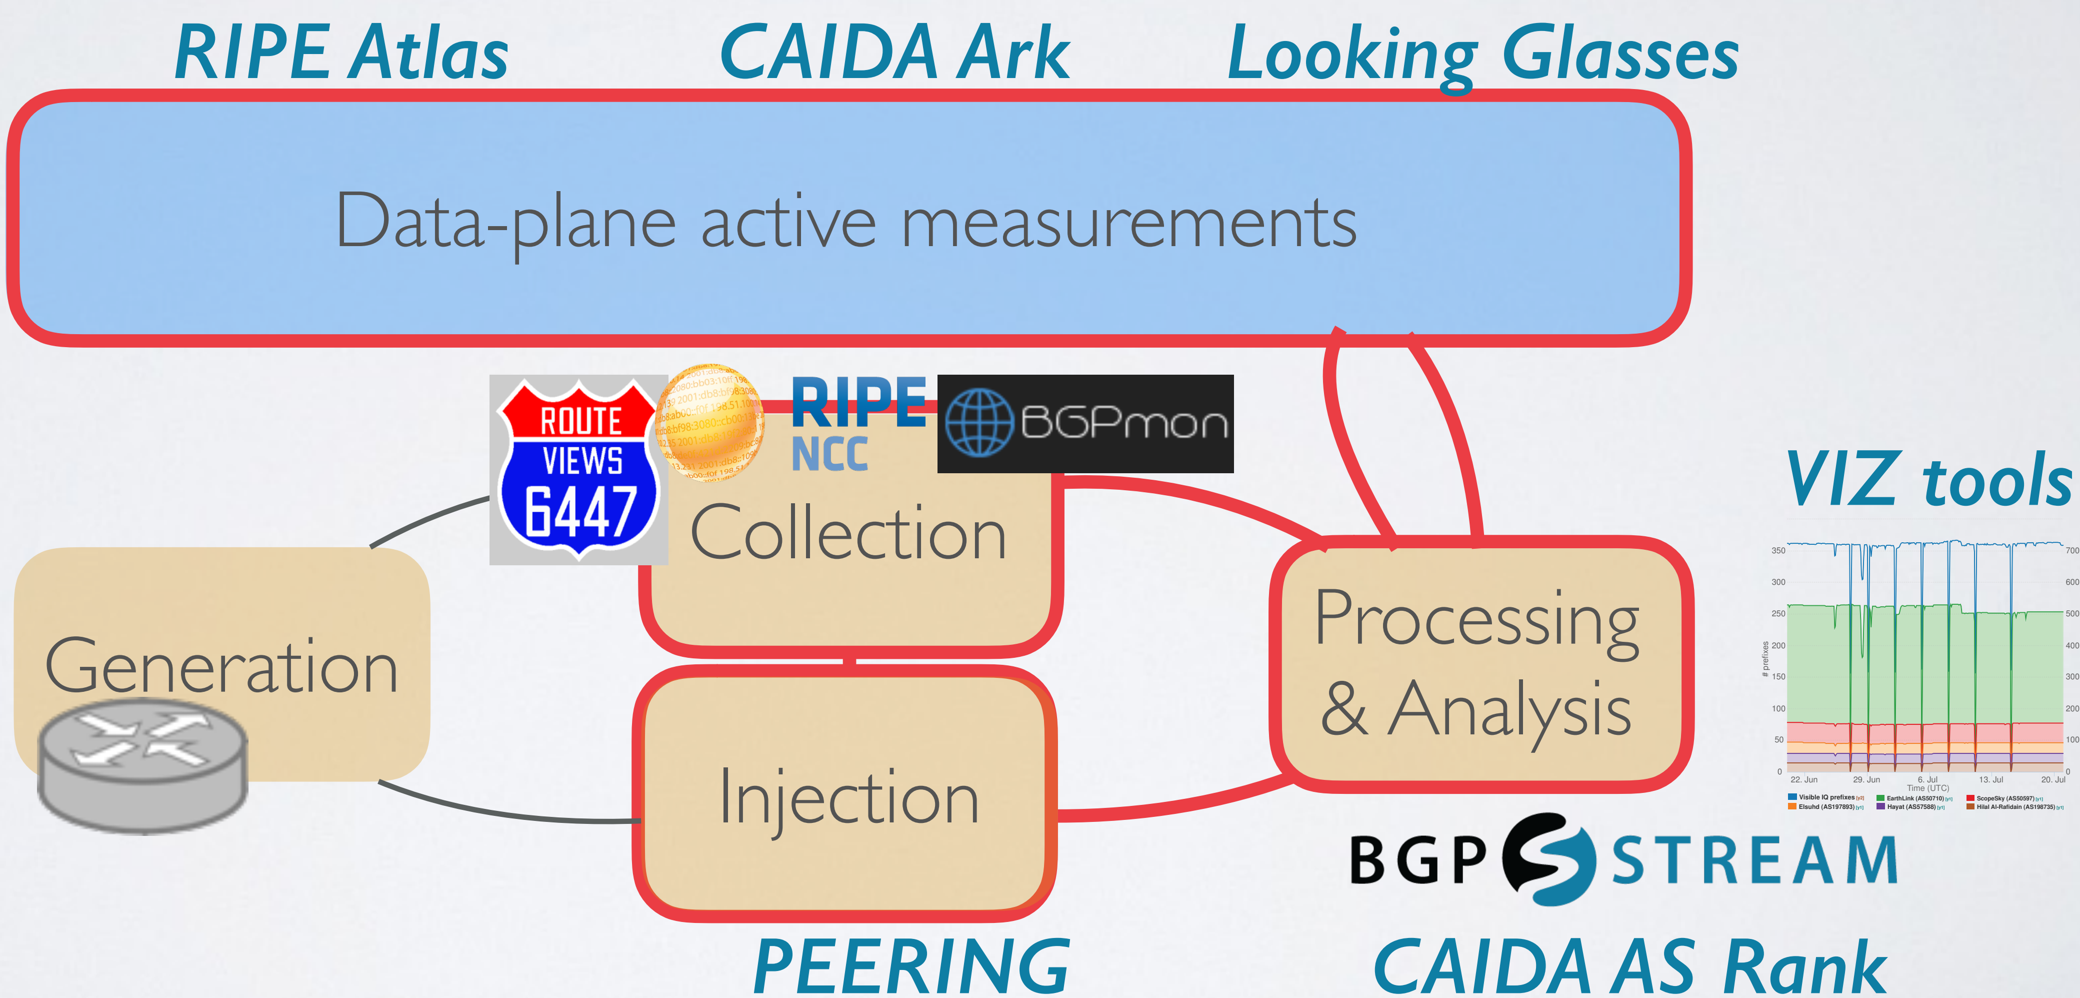

BGP HACKATHON

http://github.com/CAIDA/bgp-hackathon/wiki

- **6-7 February 2016** (weekend before NANOG 66)
- **San Diego** Supercomputer Center, UC San Diego
- **Theme**: **live BGP measurements** and monitoring
- Toolbox: *BGPMon, RIPE RIS, PEERING, BGPStream, RIPE Atlas, CAIDA Archipelago, Route Views, looking glasses, AS relationships, AS Rank, Visualization tools, …*

- How to **contribute**:
  - *join us and come over to hack!*
  - *help teams as a domain expert*
  - *propose projects that hacking teams may pick*
  - *offer to join the jury that will assign awards*

>>> bgp-hackathon-info@caida.org <<<

THANKS

bgpstream.caida.org

github.com/CAIDA/bgp-hackathon/wiki