Profiling BGP Serial Hijackers: Capturing Persistent Misbehavior in the Global Routing Table

ACM Internet Measurement Conference 2019

Cecilia Testart
MIT

Philipp Richter
MIT

Alistair King
CAIDA, UC San Diego

Alberto Dainotti
CAIDA, UC San Diego

David Clark
MIT

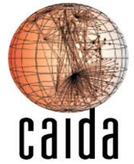

BGP hijacking is pervasive in the Internet

BGP hijacking is pervasive in the Internet

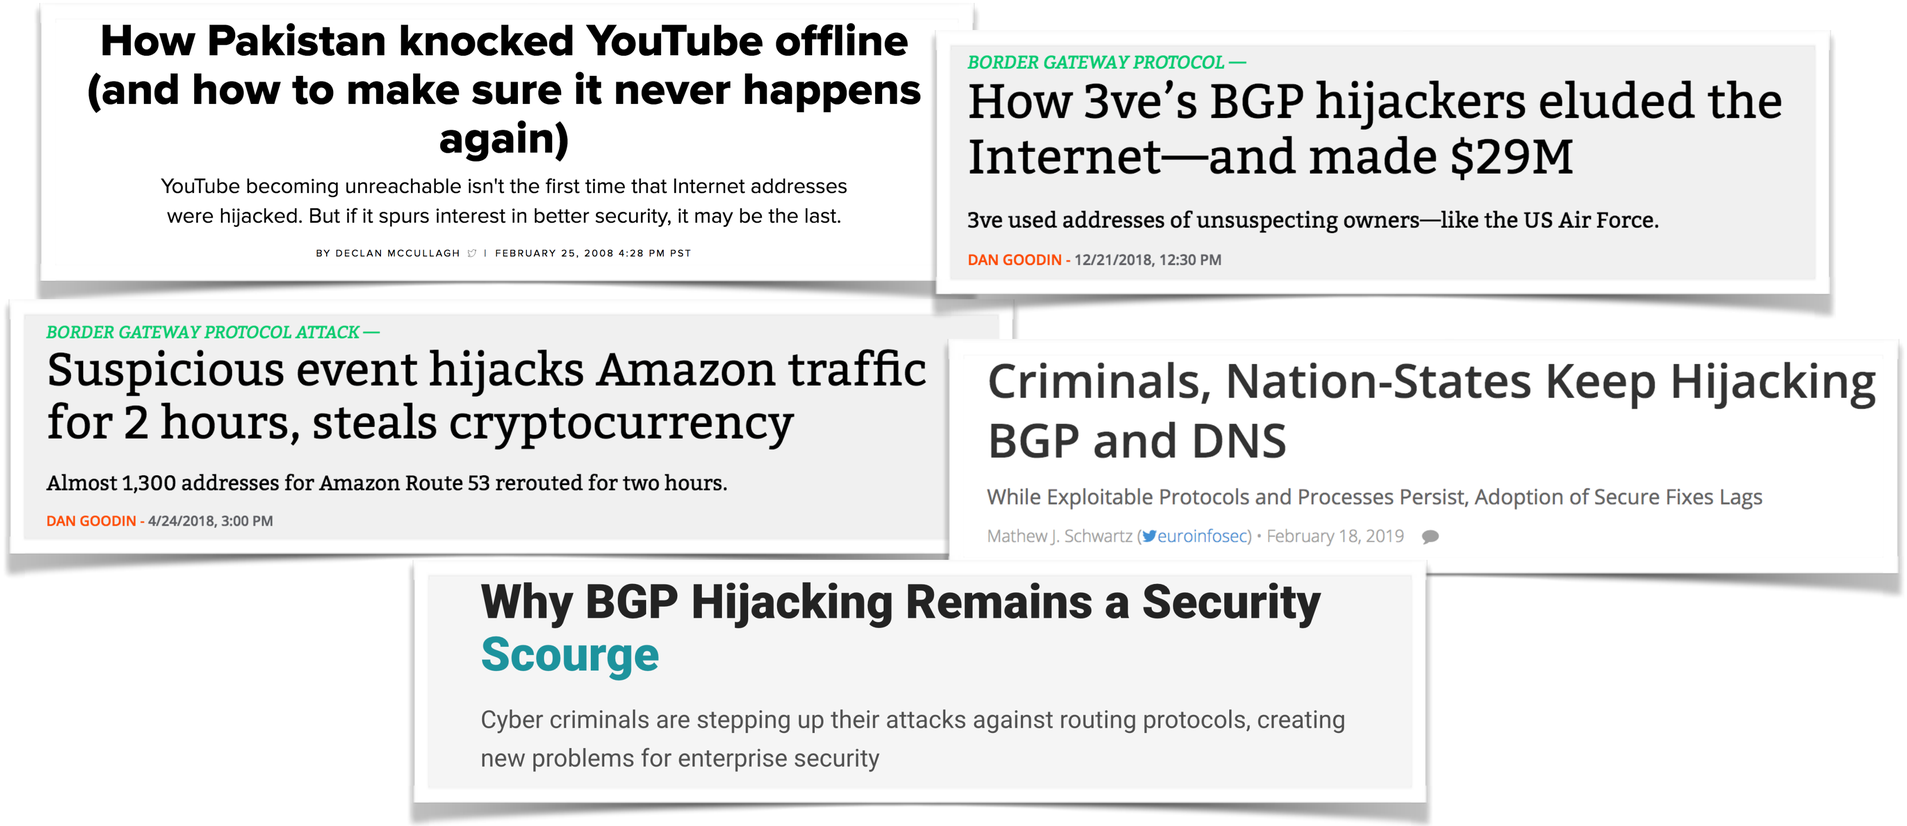

- The problem of BGP hijacking is **still** far from solved.

Hijack disclosure in mailing lists

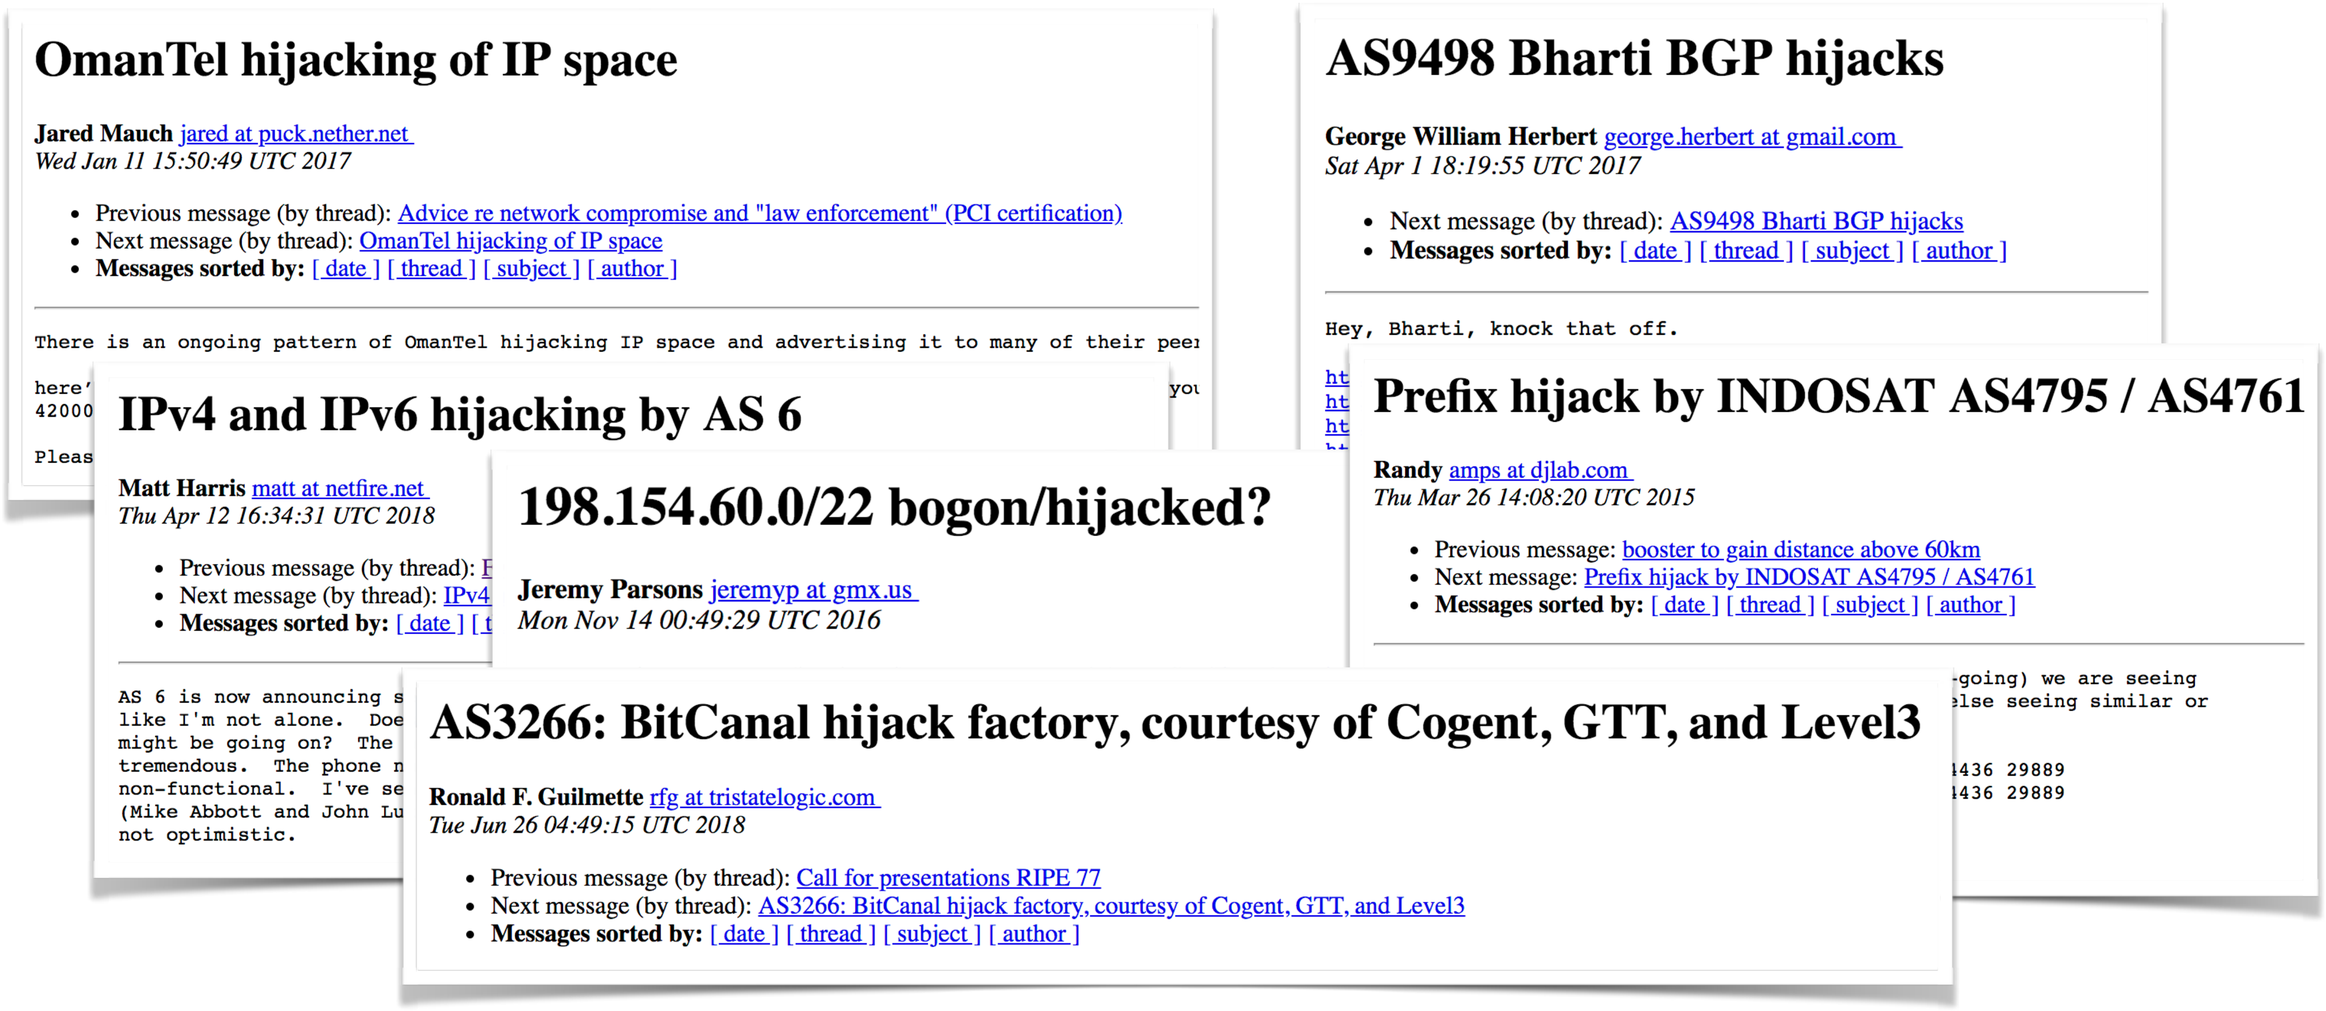

OmanTel hijacking of IP space

Jared Mauch jared at puck.nether.net
Wed Jan 11 15:50:49 UTC 2017

- Previous message (by thread): Advice re network compromise and "law enforcement" (PCI certification)
- Next message (by thread): OmanTel hijacking of IP space
- **Messages sorted by:** [date] [thread] [subject] [author]

There is an ongoing pattern of OmanTel hijacking IP space and advertising it to many of their peer

AS9498 Bharti BGP hijacks

George William Herbert george.herbert at gmail.com
Sat Apr 1 18:19:55 UTC 2017

- Next message (by thread): AS9498 Bharti BGP hijacks
- **Messages sorted by:** [date] [thread] [subject] [author]

Hey, Bharti, knock that off.

IPv4 and IPv6 hijacking by AS 6

Matt Harris matt at netfire.net
Thu Apr 12 16:34:31 UTC 2018

- Previous message (by thread):
- Next message (by thread): IPv4
- **Messages sorted by:** [date]

AS 6 is now announcing s
like I'm not alone. Doe
might be going on? The
tremendous. The phone n
non-functional. I've se
(Mike Abbott and John Lu
not optimistic.

198.154.60.0/22 bogon/hijacked?

Jeremy Parsons jeremyp at gmx.us
Mon Nov 14 00:49:29 UTC 2016

Prefix hijack by INDOSAT AS4795 / AS4761

Randy amps at djlab.com
Thu Mar 26 14:08:20 UTC 2015

- Previous message: booster to gain distance above 60km
- Next message: Prefix hijack by INDOSAT AS4795 / AS4761
- **Messages sorted by:** [date] [thread] [subject] [author]

-going) we are seeing
lse seeing similar or

436 29889
436 29889

AS3266: BitCanal hijack factory, courtesy of Cogent, GTT, and Level3

Ronald F. Guilmette rfg at tristatelogic.com
Tue Jun 26 04:49:15 UTC 2018

- Previous message (by thread): Call for presentations RIPE 77
- Next message (by thread): AS3266: BitCanal hijack factory, courtesy of Cogent, GTT, and Level3
- **Messages sorted by:** [date] [thread] [subject] [author]

Hijack disclosure in mailing lists

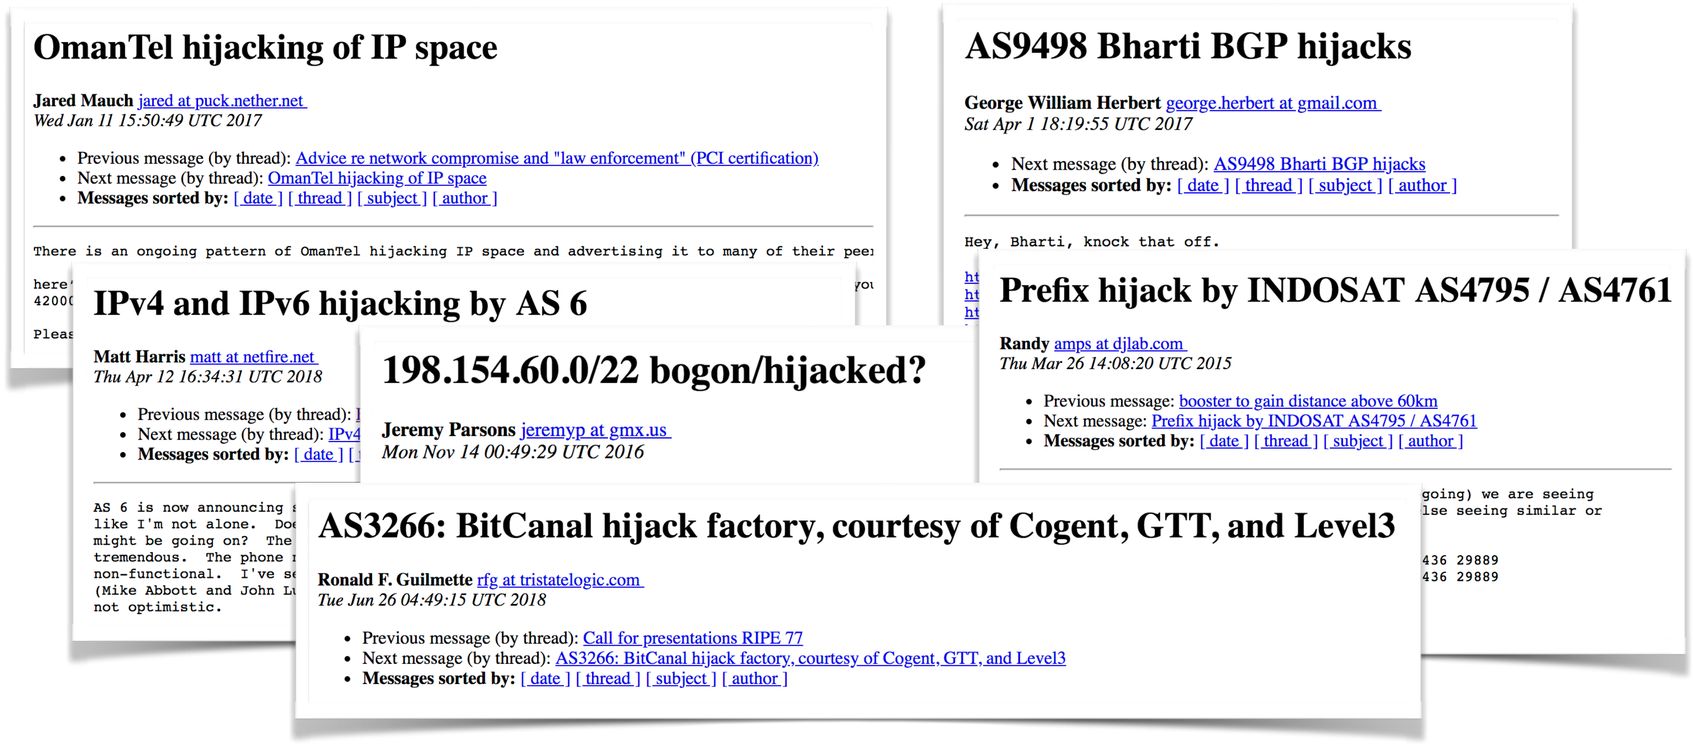

- ***Serial hijackers***: ASes that **repeatedly hijack** over **long periods** of time.

Bitcanal: an infamous serial hijacker

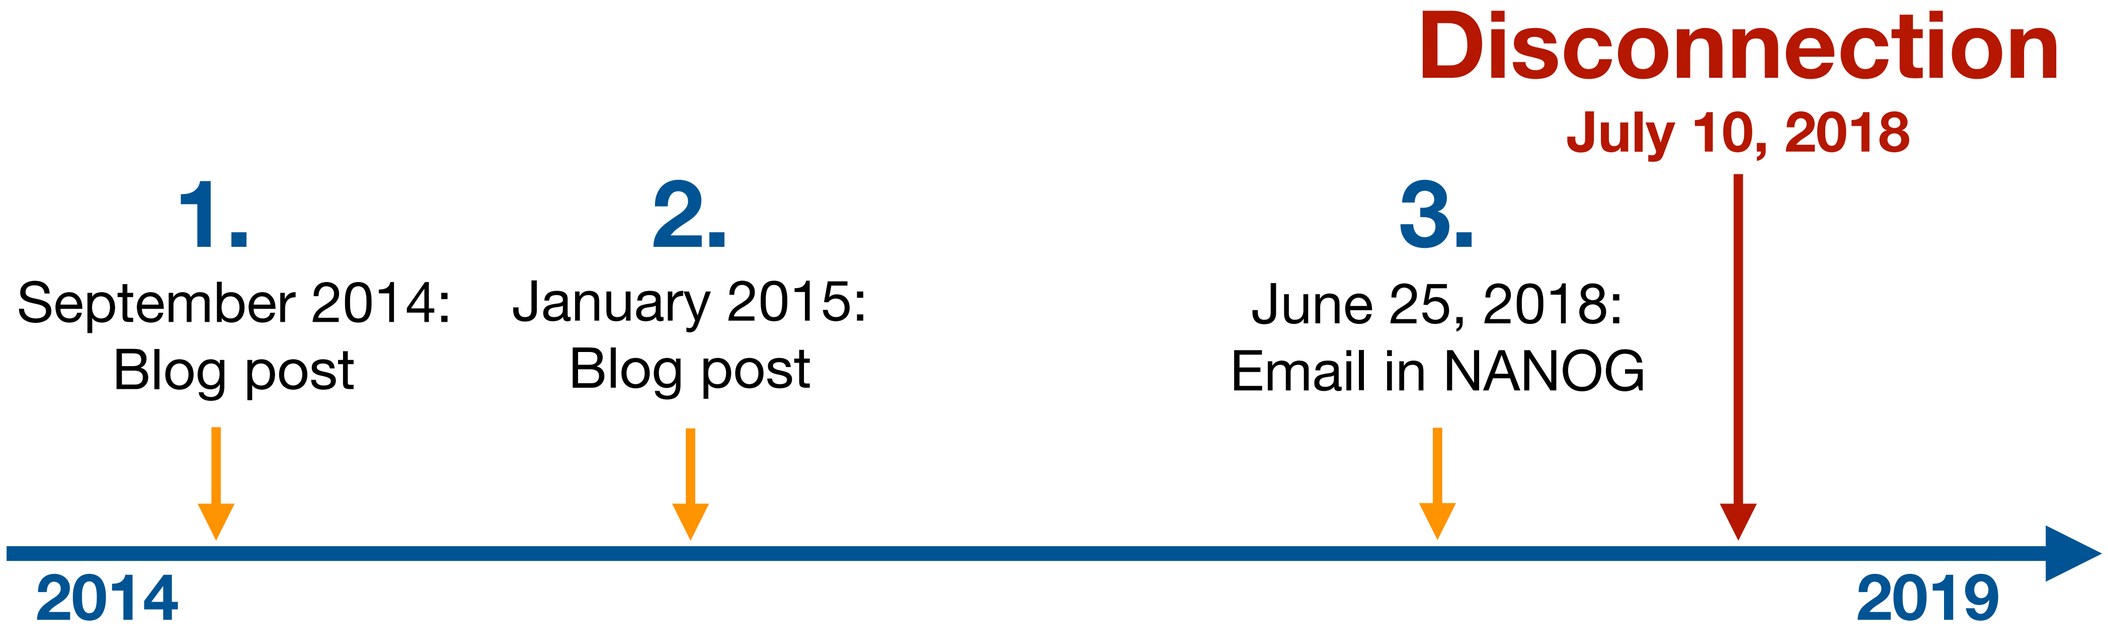

Bitcanal: an infamous serial hijacker

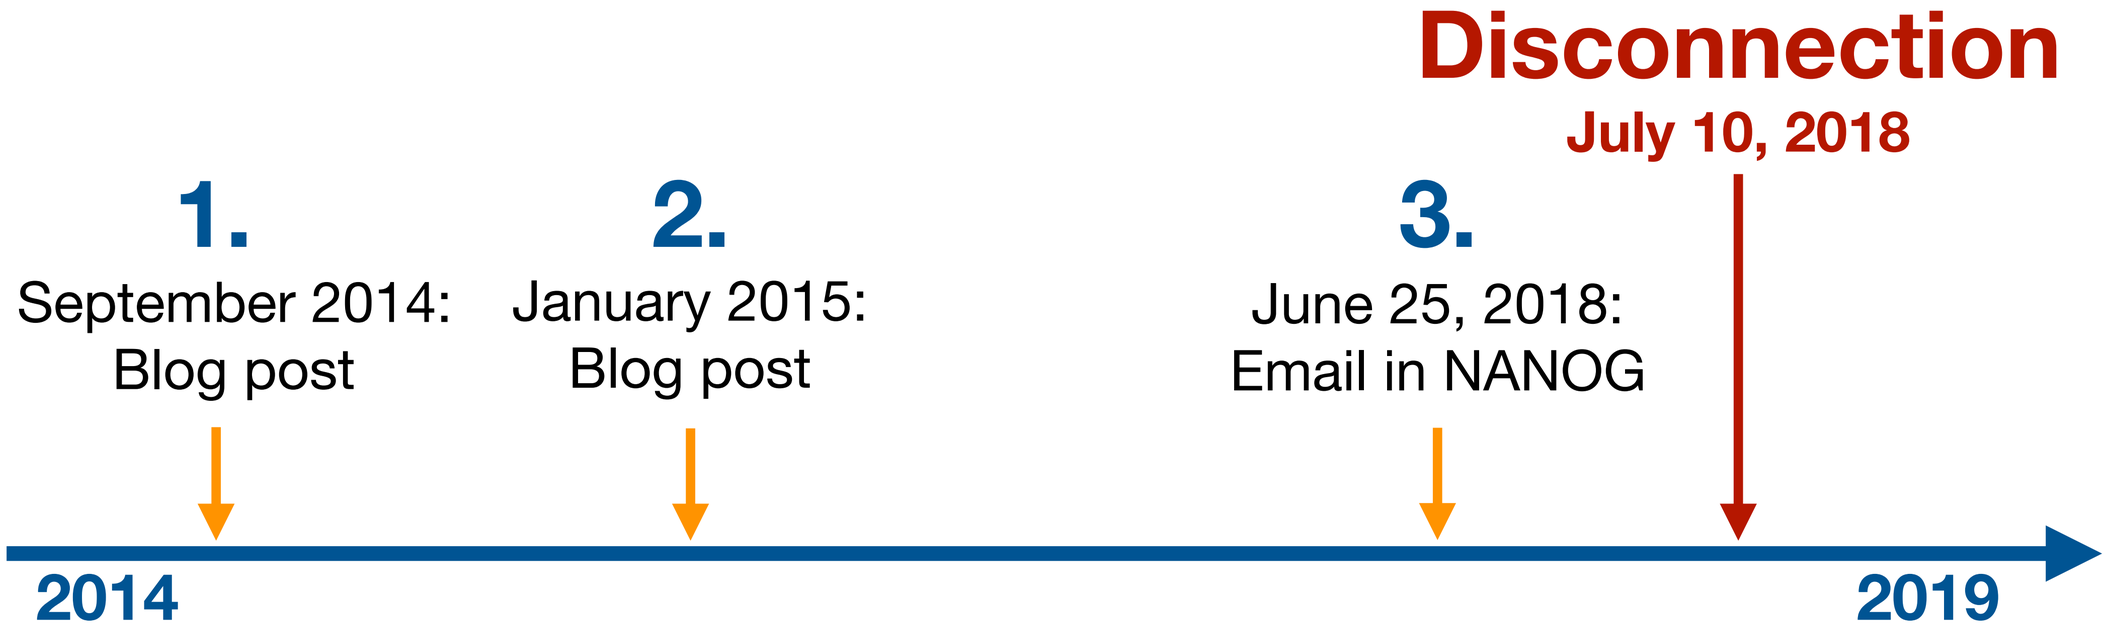

- It took **4 years** to disconnect this serial hijacker.

Research goals

Find serial hijackers in the Internet

(i) Identify hijackers distinctive routing characteristics

(ii) Build a machine learning system to flag suspicious ASes

(iii) Evaluate our results

What can we learn about serial hijackers?

Ground truth: serial hijackers

23 serial hijackers:

- 10+ hijacks
- Most have been active over a year
- Up to 30,000 originated prefixes

ASN country and RIR registration

Ground truth: legitimate ASes

230 Legitimate ASes:

- 191 MANRS ASes
- 26 ASes manually selected

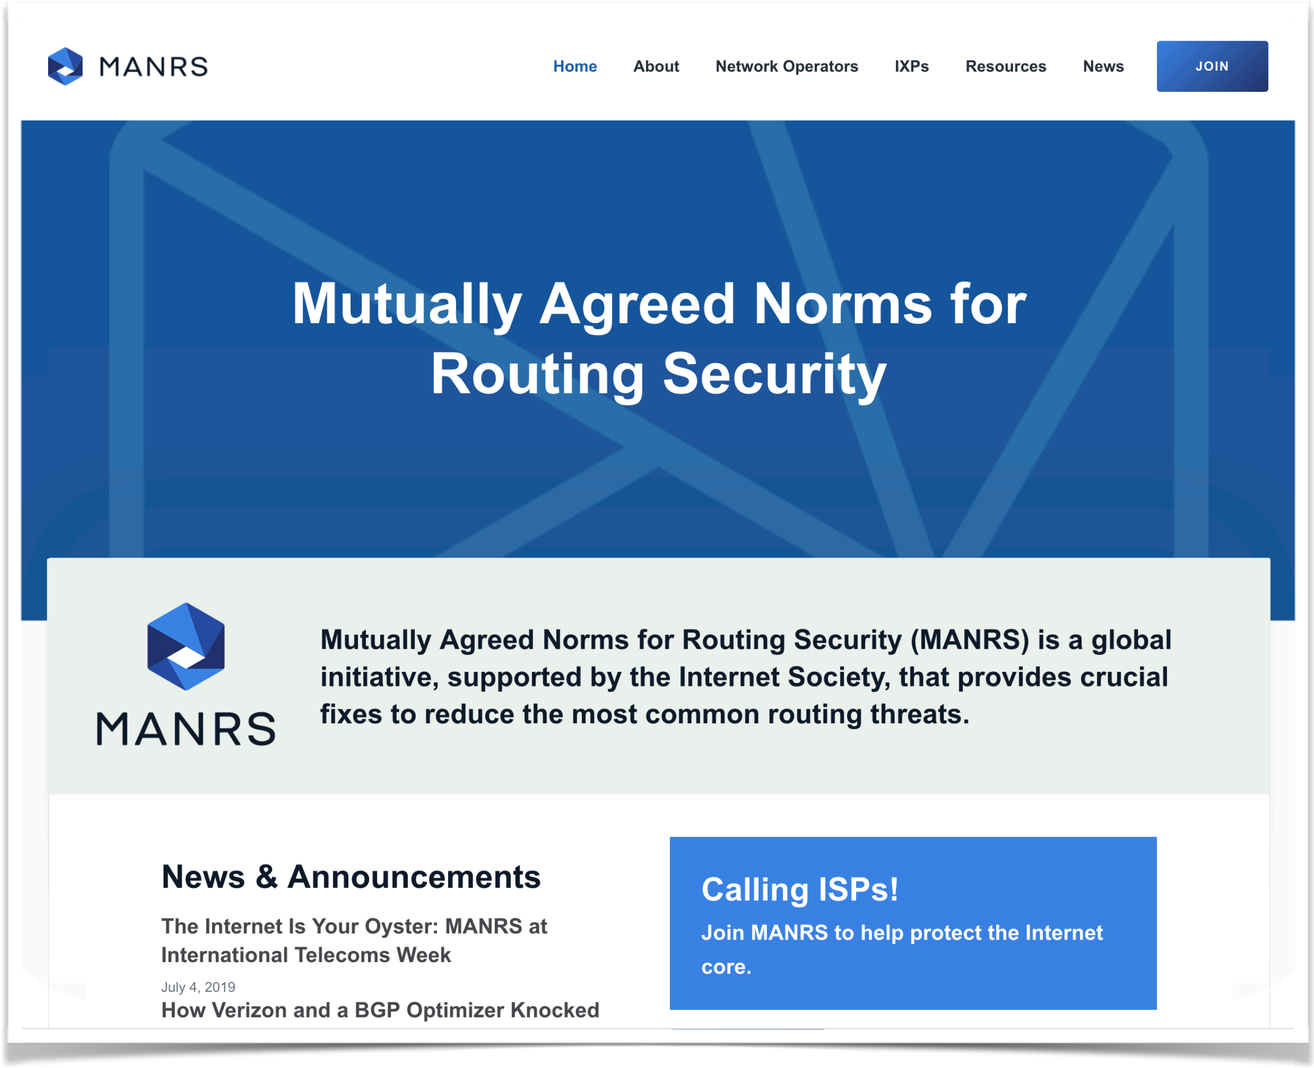

BGP dataset and processing

- RIPE RIS and RouteViews collectors (~40 col., ~1400+ col. peers)
- We process all **BGP updates** to reconstruct peer routing tables
- We extract **(prefix, origin AS)** pairs and the number of peers with each pair in their routing table **(visibility)**
- Data from Jan. 2014 to Dec. 2018

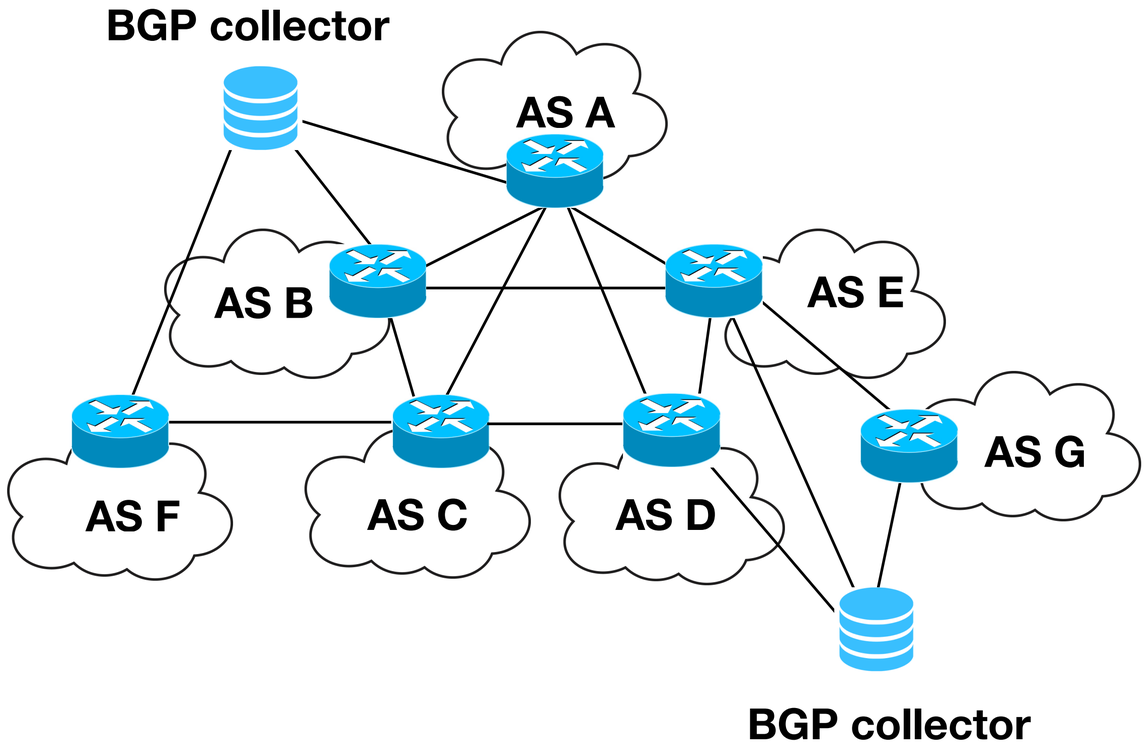

▶ **(prefix, origin AS, visibility, timestamp)** every 5 min.

BGP origination behavior: legitimate vs. serial hijacker

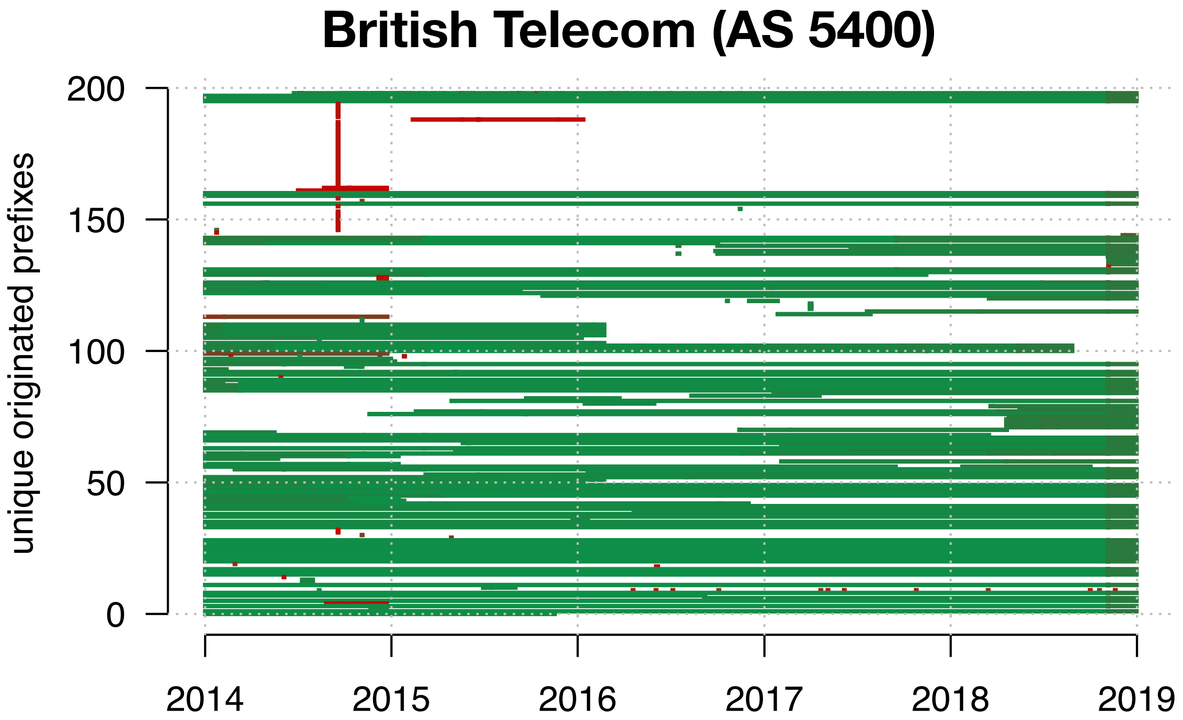

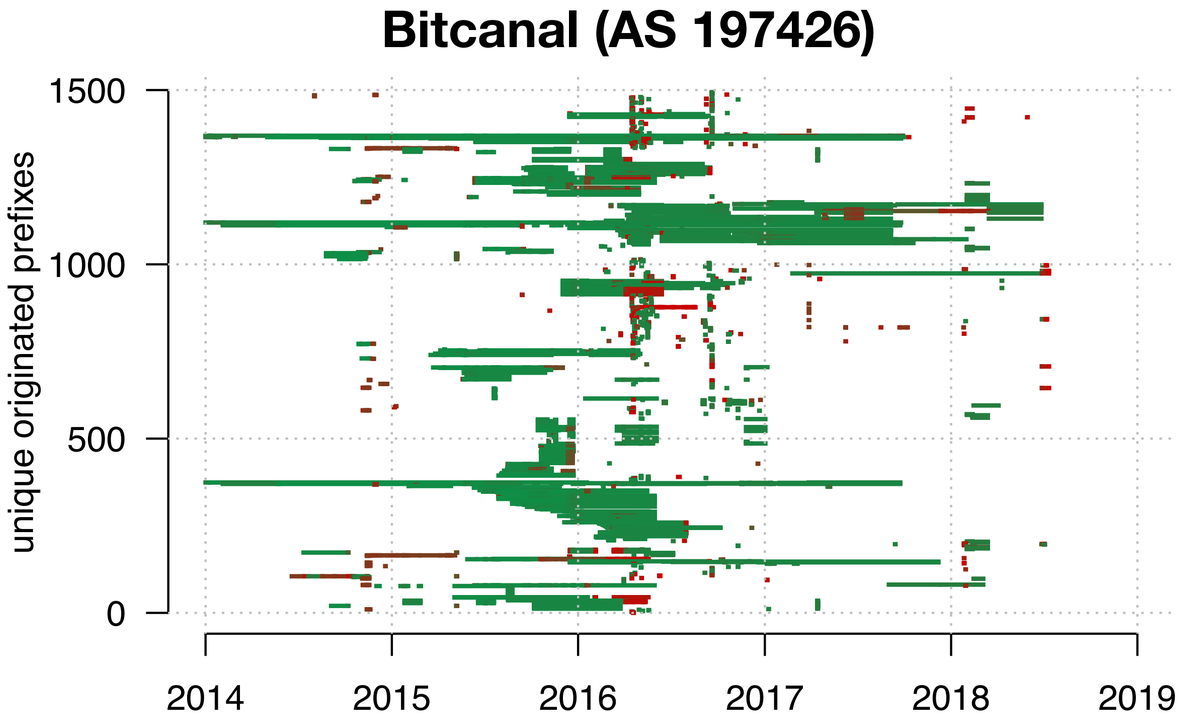

- Legitimate ASes mostly show **stable** BGP behavior.
- Serial hijackers BGP activity is **visually different**.

BGP origination behavior: legitimate vs. serial hijacker

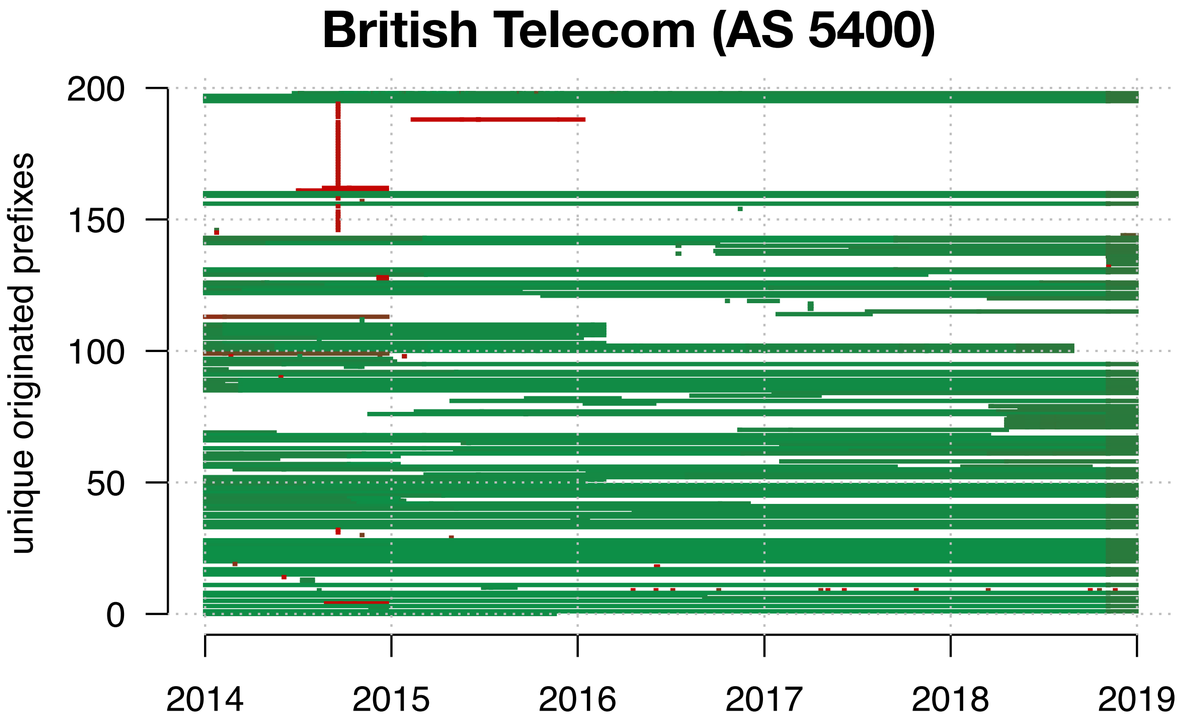

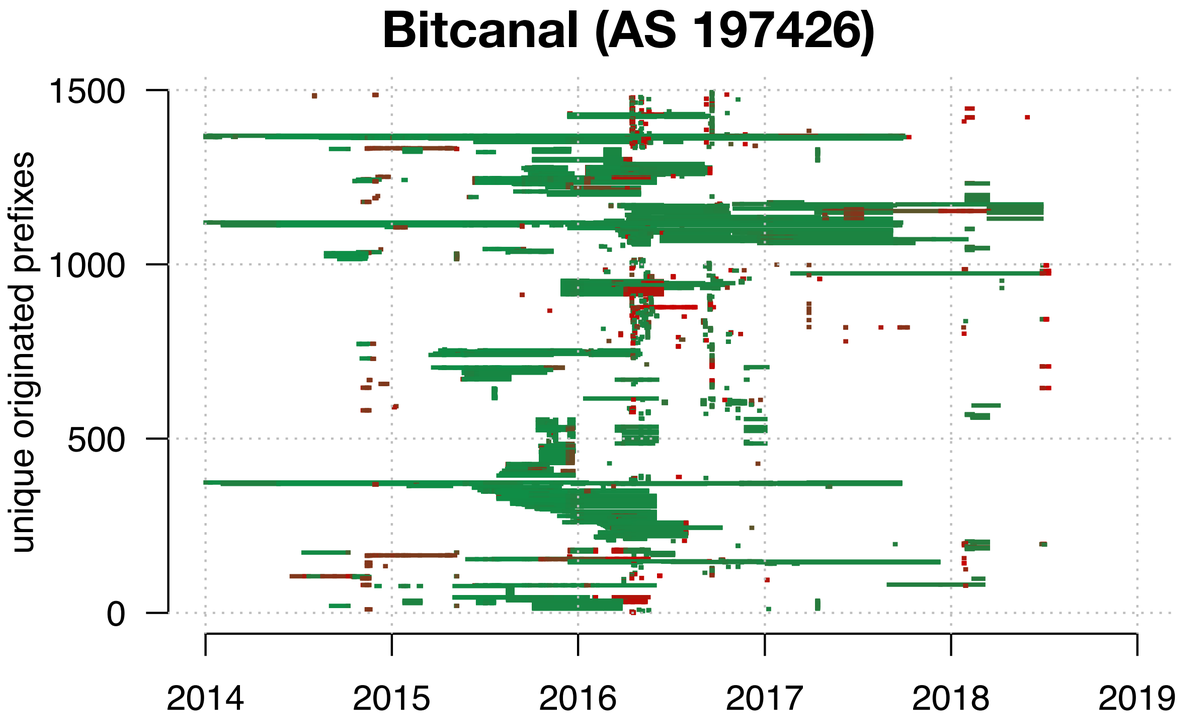

- We need features that **capture** this **behavioral difference**.

Variability of BGP behavior: serial hijackers

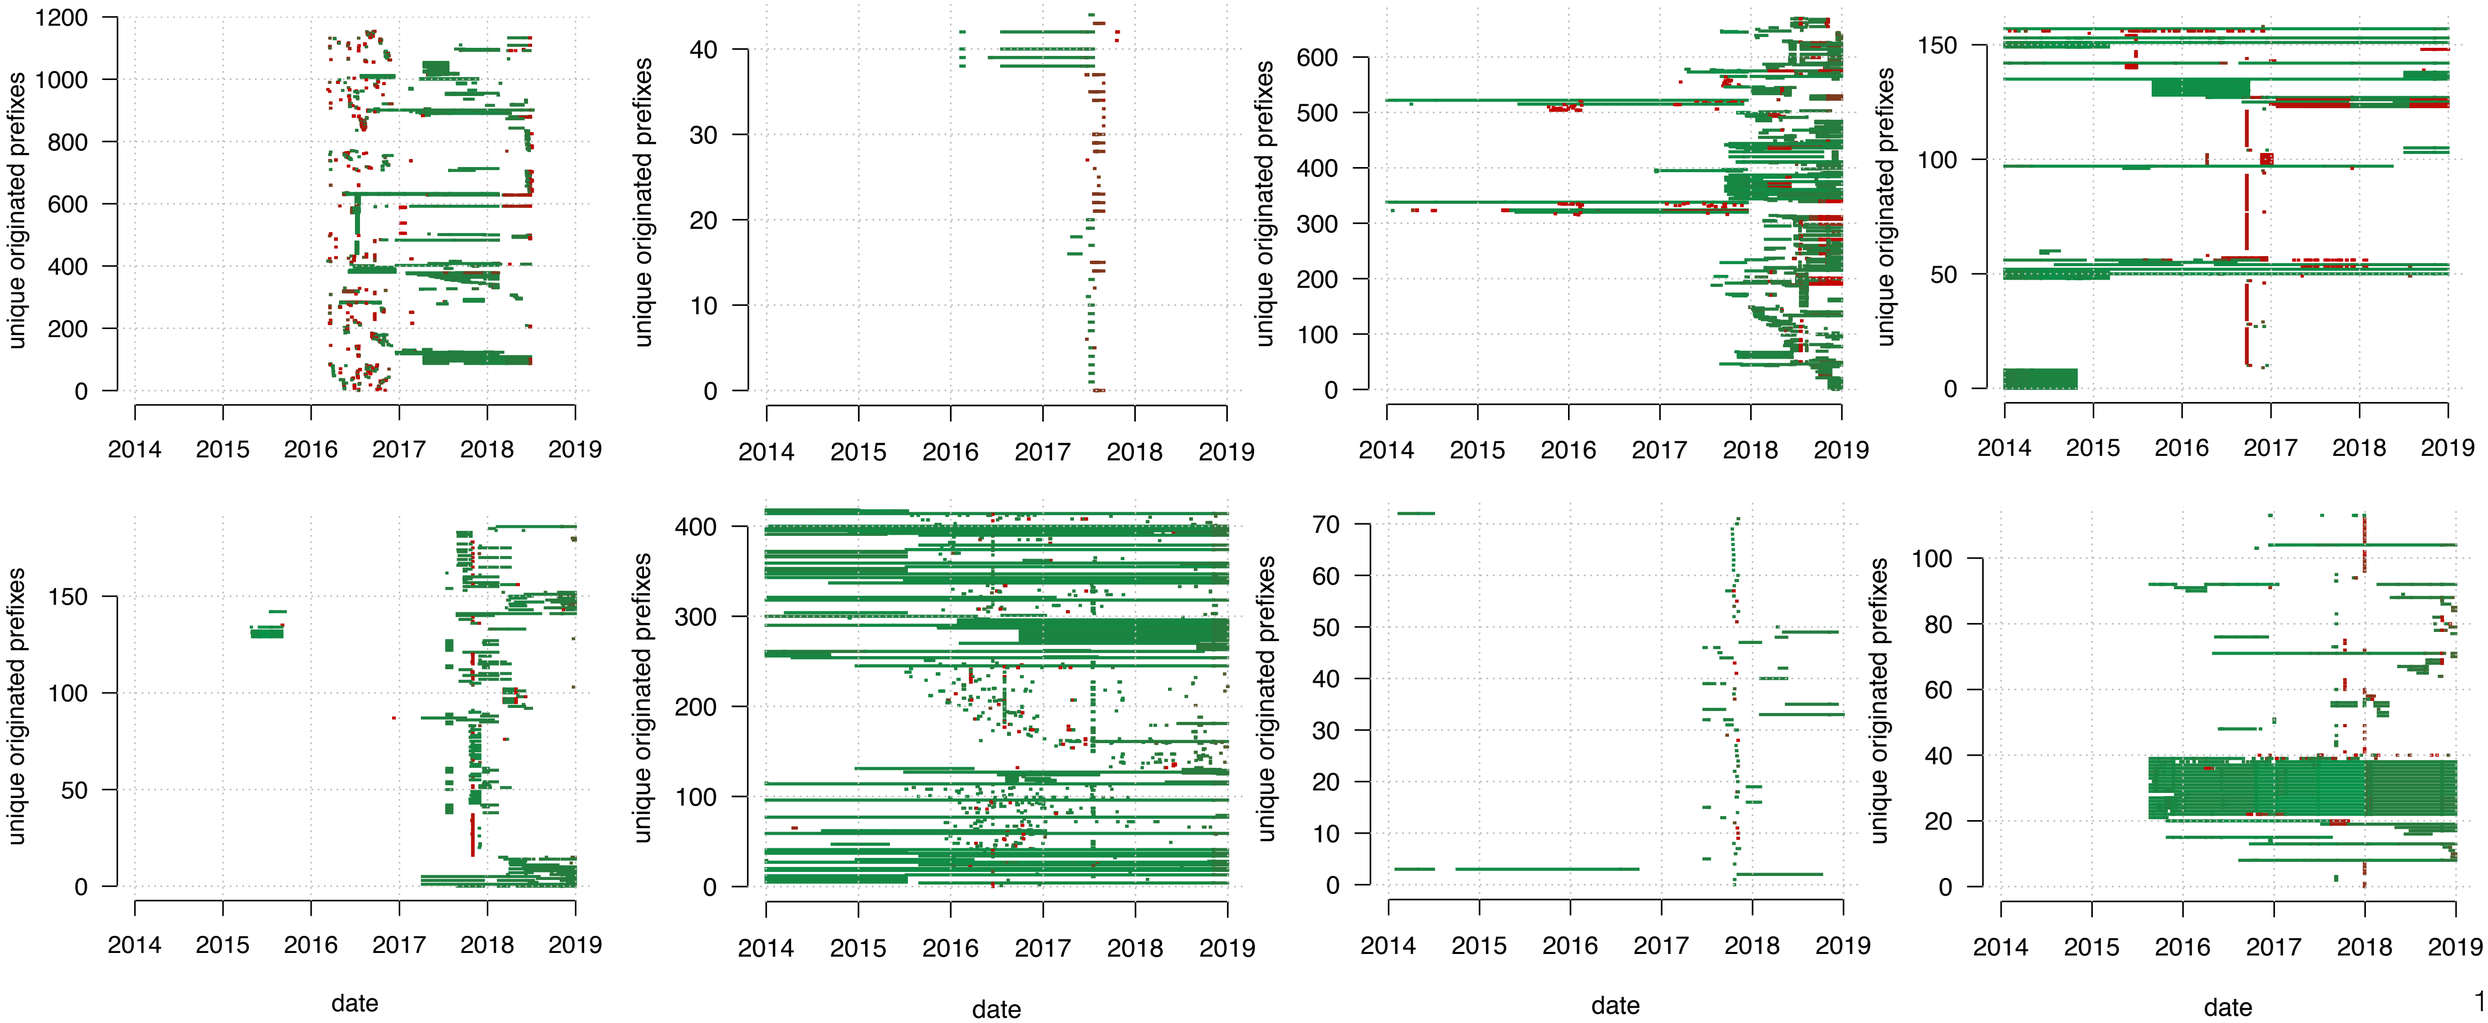

Variability of BGP behavior: legitimate ASes

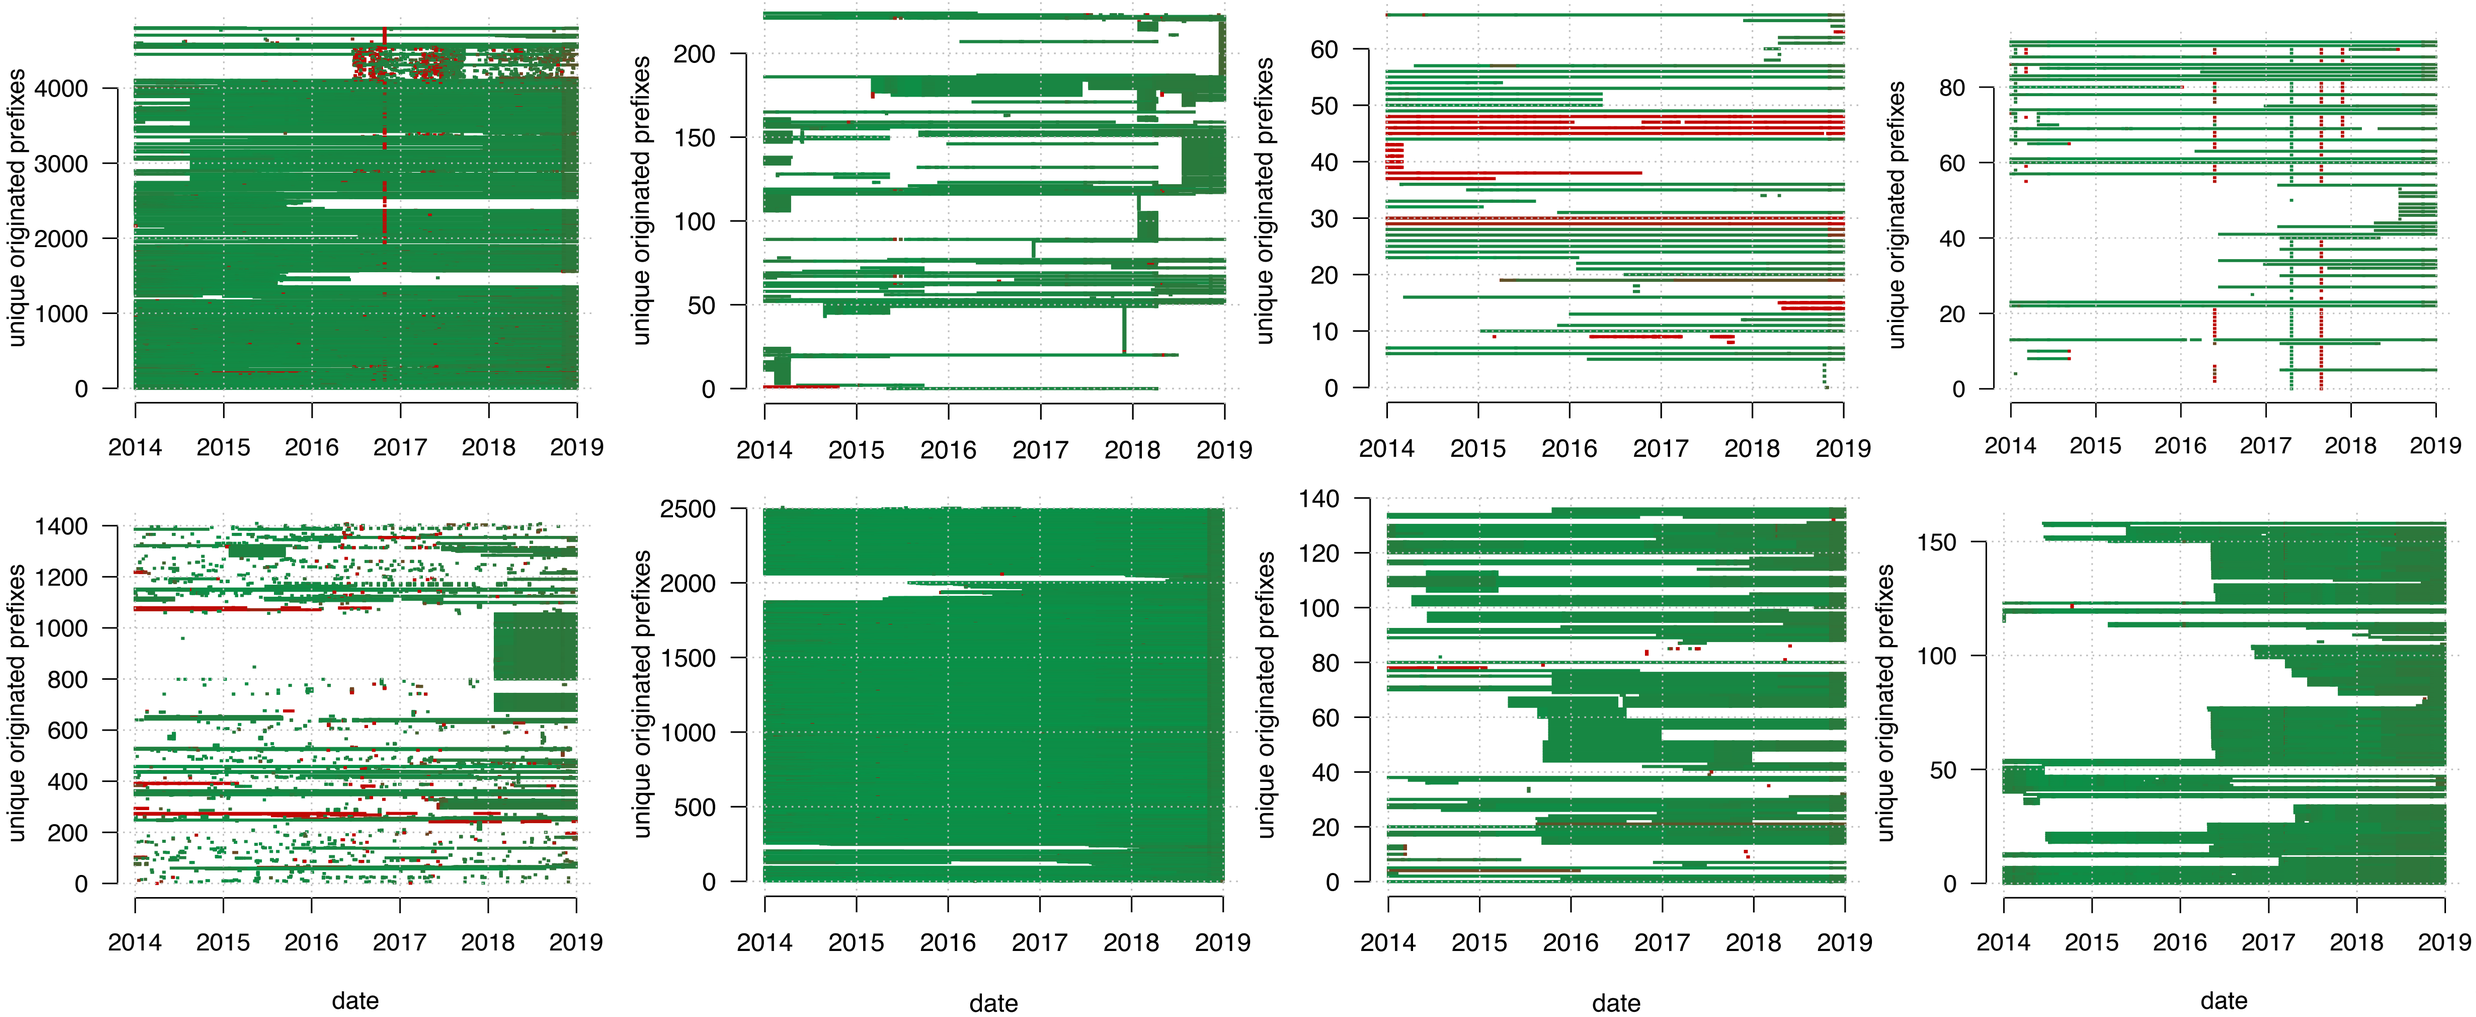

Expected serial hijacker behavior

- Repeated AS absence from the global routing table.

Expected serial hijacker behavior

- Repeated AS absence from the global routing table.
- Short prefix origination times.

Expected serial hijacker behavior

- Repeated AS absence from the global routing table.
- Short prefix origination times.
- More multi-origin conflicts (MOAS).

Expected serial hijacker behavior

- Repeated AS absence from the global routing table.
- Short prefix origination times.
- More multi-origin conflicts (MOAS).
- Volatile count of concurrently advertised prefixes.

Expected serial hijacker behavior

- Repeated AS absence from the global routing table.
- Short prefix origination times.
- More multi-origin conflicts (MOAS).
- Volatile count of concurrently advertised prefixes.

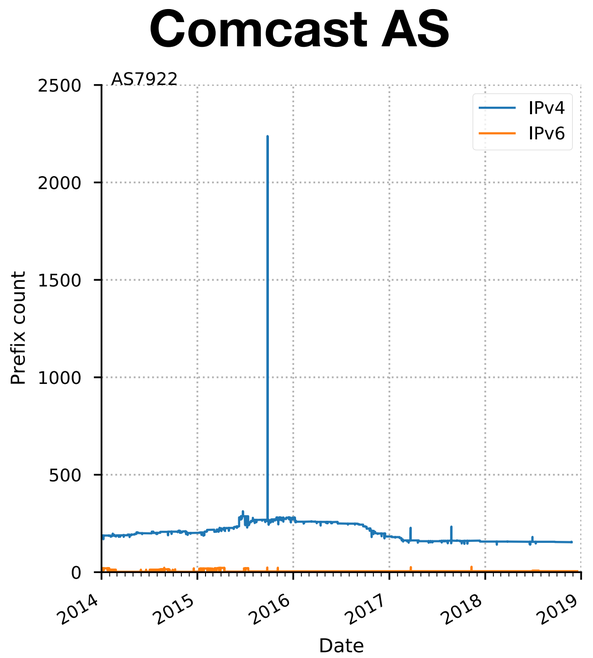

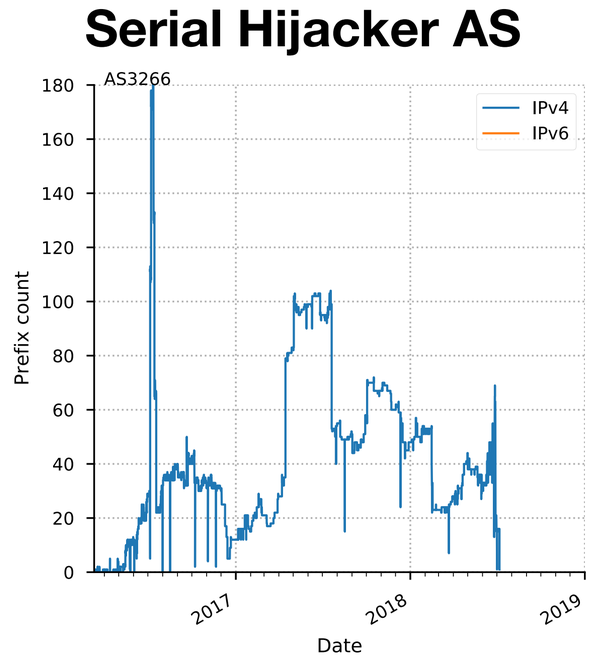

Expected serial hijacker behavior

- Repeated AS absence from the global routing table.
- Short prefix origination times.
- More multi-origin conflicts (MOAS).
- Volatile count of concurrently advertised prefixes.
- Broad geographical distribution of address space originated.

Expected serial hijacker behavior

- Repeated AS absence from the global routing table.
- Short prefix origination times.
- More multi-origin conflicts (MOAS).
- Volatile count of concurrently advertised prefixes.
- Broad geographical distribution of address space originated.

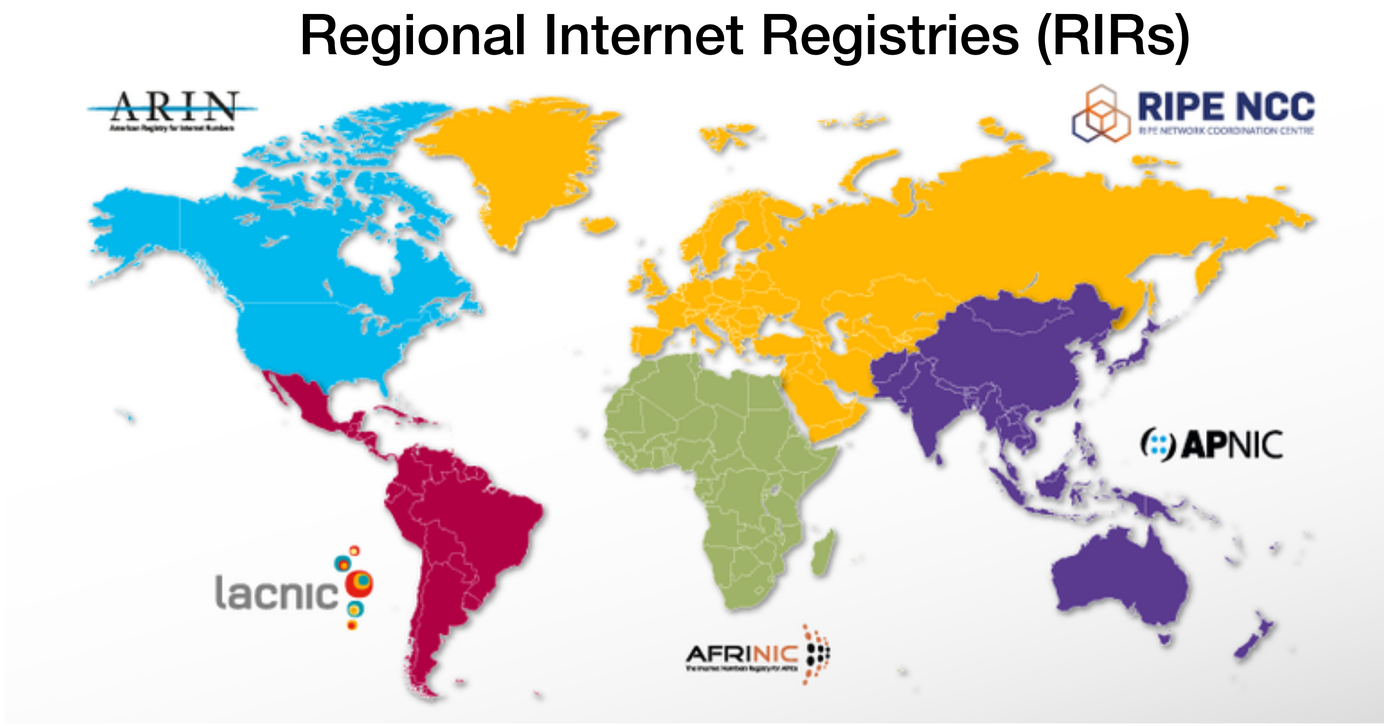

Expected serial hijacker behavior

- Repeated AS absence from the global routing table.
- Short prefix origination times.
- More multi-origin conflicts (MOAS).
- Volatile count of concurrently advertised prefixes.
- Broad geographical distribution of address space originated.

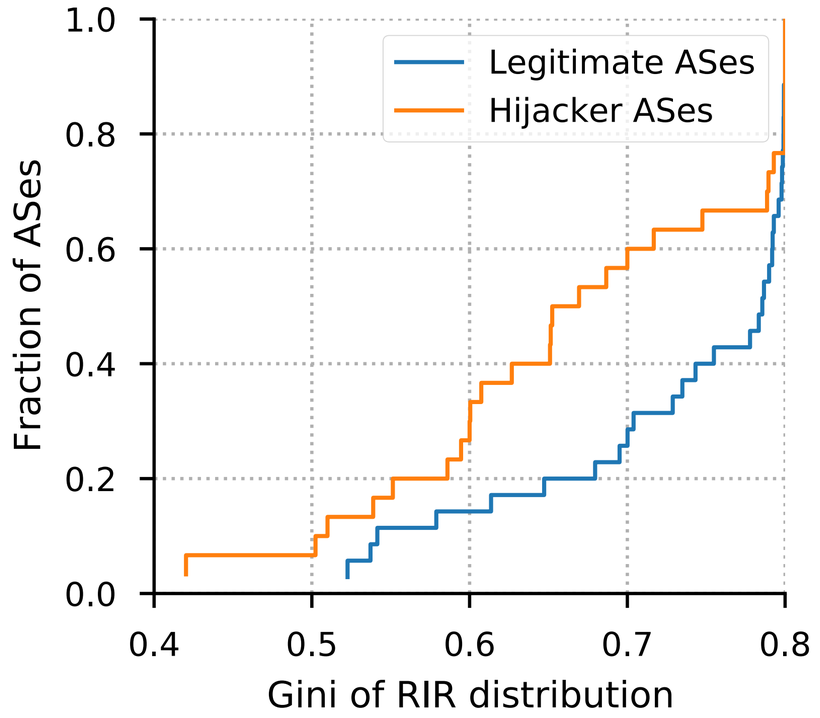

Expected serial hijacker behavior

- Repeated AS absence from the global routing table.
- Short prefix origination times.
- More multi-origin conflicts (MOAS).
- Volatile count of concurrently advertised prefixes.
- Broad geographical distribution of address space originated.

Expected serial hijacker behavior

- Repeated AS absence from the global routing table.
- Short prefix origination times.
- More multi-origin conflicts (MOAS).
- Volatile count of concurrently advertised prefixes.
- Broad geographical distribution of address space originated.

▶ We derived **52 features** to capture differences.

Challenges of applying ML to find more potential serial hijackers

Challenges of applying ML to find more potential serial hijackers

- Heavy-tailed and skewed data:
  Monthly prefix changes [0,2600], Gini in [0,0.8]

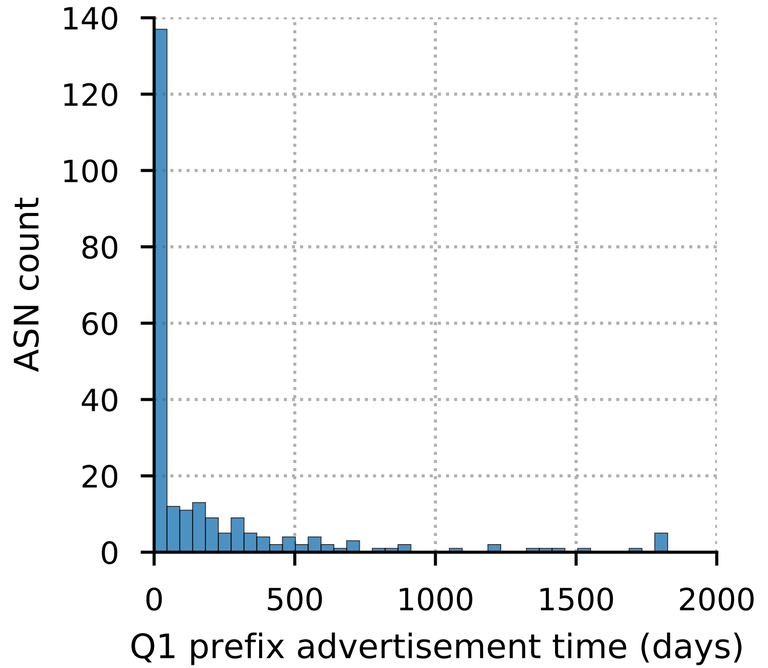

Challenges of applying ML to find more potential serial hijackers

- Heavy-tailed and skewed data:
  Monthly prefix changes [0,2600], Gini in [0,0.8]

- Very small ground truth:
  240 AS for 19,000 ASes

Challenges of applying ML to find more potential serial hijackers

- Heavy-tailed and skewed data:
  Monthly prefix changes [0,2600], Gini in [0,0.8]

- Very small ground truth:
  240 AS for 19,000 ASes

- Class Imbalance:
  23 serial hijacker vs. 217 legitimate networks

Our ML approach

- Tree based classifier.
- Voting ensemble of extremely randomized forests.
- 3 over-sampling techniques.
- All 52 features with positive median drop column importance.

Our ML approach

- Tree based classifier.
- Voting ensemble of extremely randomized forests.
- 3 over-sampling techniques.
- All 52 features with positive median drop column importance.

► **79% precision** and **100% recall**
(in ground-truth using out-of-bag score)

Putting our classifier to work

- **Goal:** Find ASes exhibiting similar BGP behavior to serial hijackers in our ground truth.

Putting our classifier to work

- **Goal:** Find ASes exhibiting similar BGP behavior to serial hijackers in our ground truth.

Putting our classifier to work

- **Goal:** Find ASes exhibiting similar BGP behavior to serial hijackers in our ground truth.

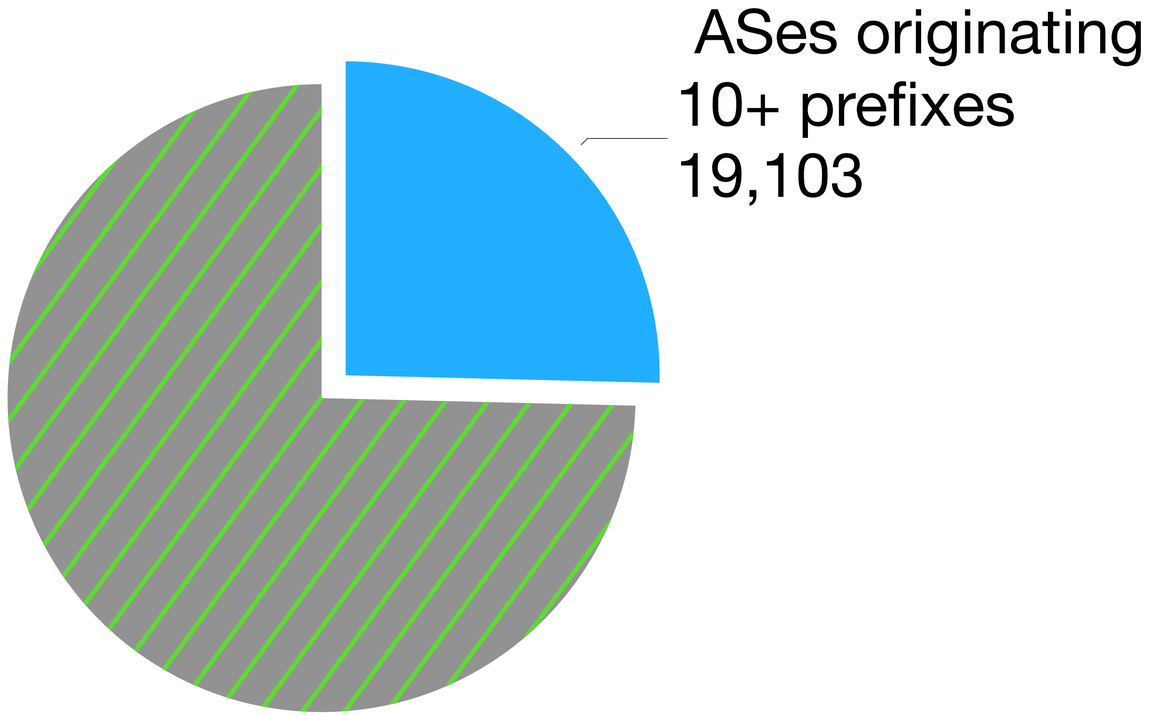

Putting our classifier to work

- **Goal:** Find ASes exhibiting similar BGP behavior to serial hijackers in our ground truth.

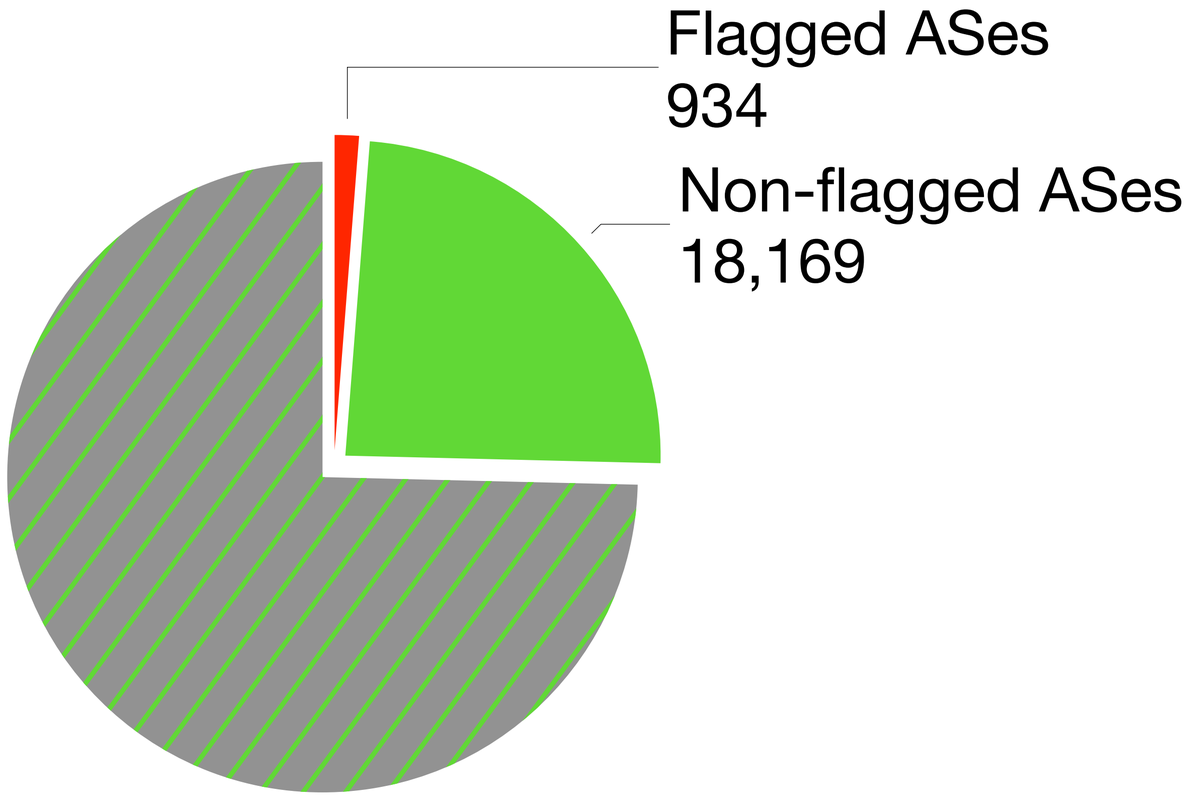

Putting our classifier to work

- **Goal:** Find ASes exhibiting similar BGP behavior to serial hijackers in our ground truth.

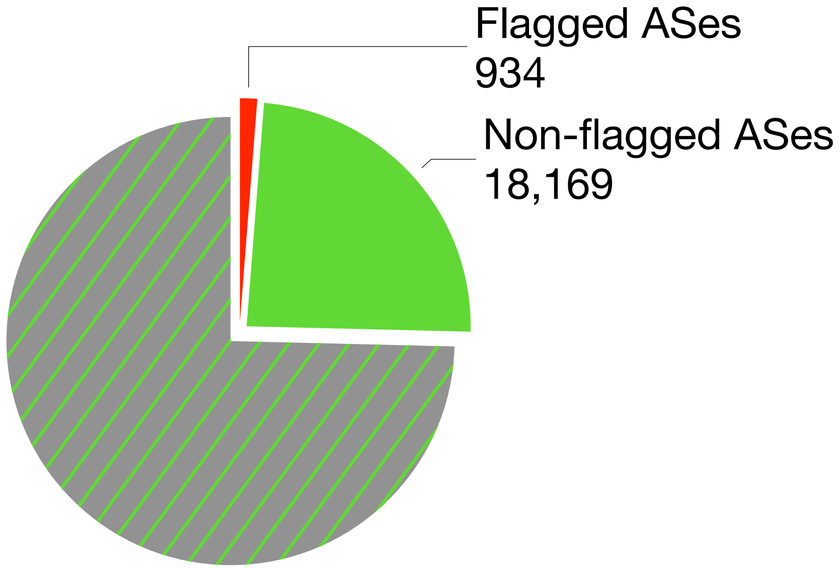

▶ Flagged ASes are:

- **4.9%** of ASes originating 10+ prefixes
- **1.2%** of all ASes.

BGP behavior of flagged ASes

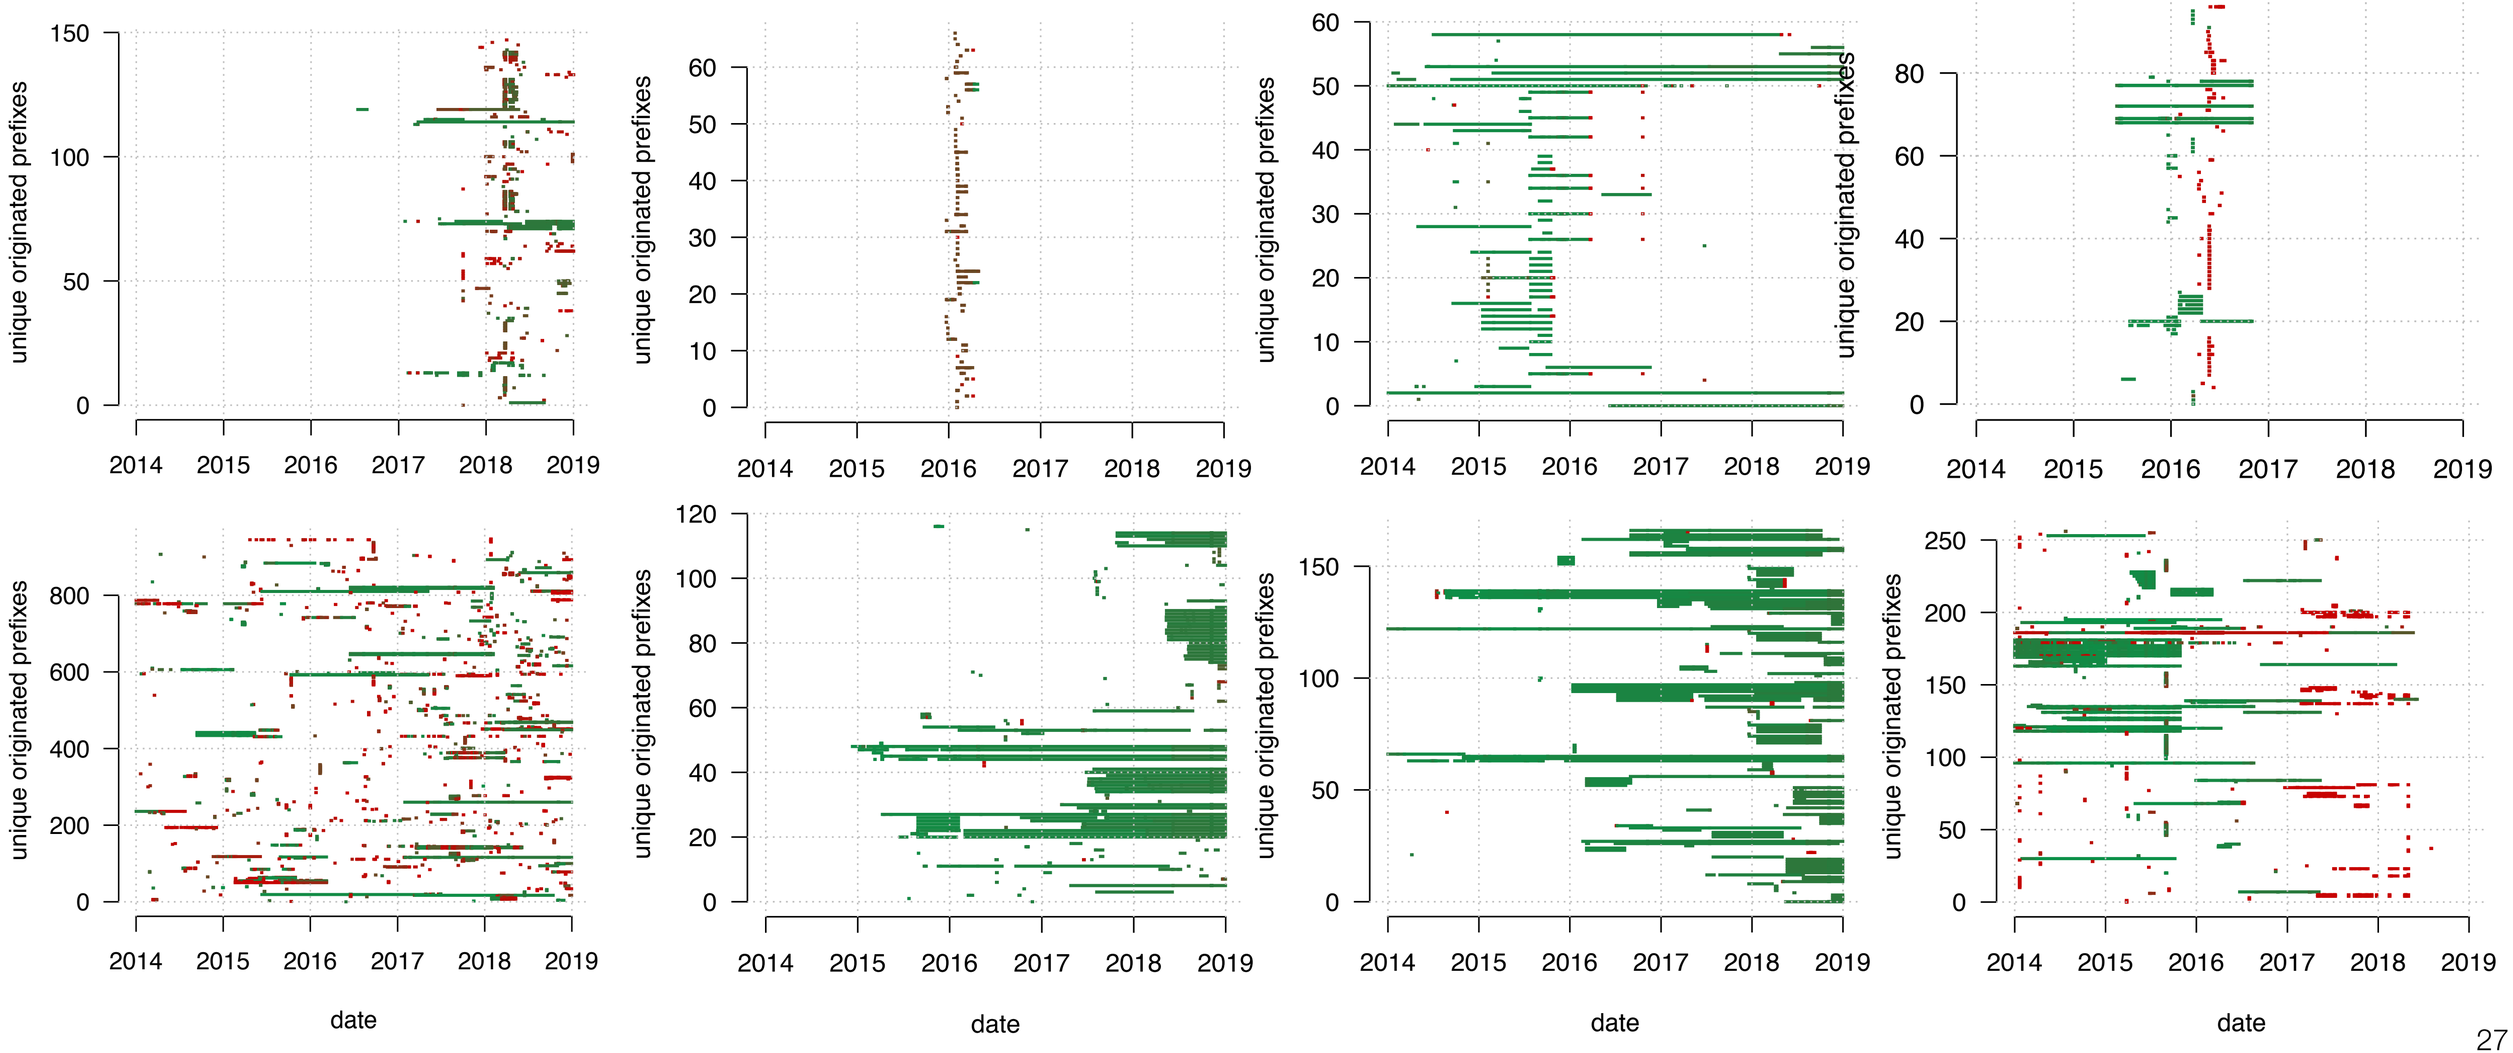

What are ASes flagged by our classifier?

- **Indication of malicious behavior**
  - Blacklisted ASNs:

What are ASes flagged by our classifier?

- **Indication of malicious behavior**
  - Blacklisted ASNs: **84/290** ASes in *Spamhaus ASN DROP list*

What are ASes flagged by our classifier?

- **Indication of malicious behavior**
  - Blacklisted ASNs: **84/290** ASes in *Spamhaus ASN DROP list*

➡ Flagged ASes are **10x** more likely to be blacklisted

What are ASes flagged by our classifier?

- **Indication of malicious behavior**
  - Blacklisted ASNs: **84/290** ASes in *Spamhaus ASN DROP list*
  - Spammer ASNs: **33%** ASes have a prefix in UCE-PROTECT level 2 spam blacklist

What are ASes flagged by our classifier?

- Indication of malicious behavior
- **Indication of misconfigurations**
  - Private ASNs **12%**

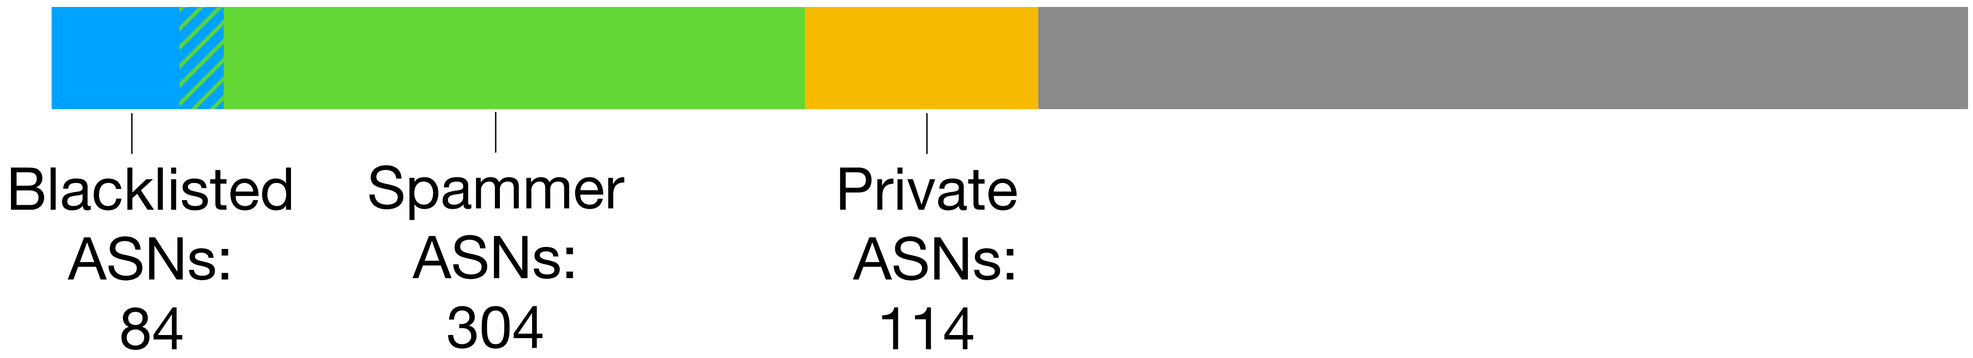

What are ASes flagged by our classifier?

- Indication of malicious behavior
- **Indication of misconfigurations**
  - Private ASNs **12%**
  - Fat-finger error ASNs **1%**

What are ASes flagged by our classifier?

- Indication of malicious behavior
- Indication of misconfigurations
- **Known false positives**
  - DDos protection ASNs **2%**

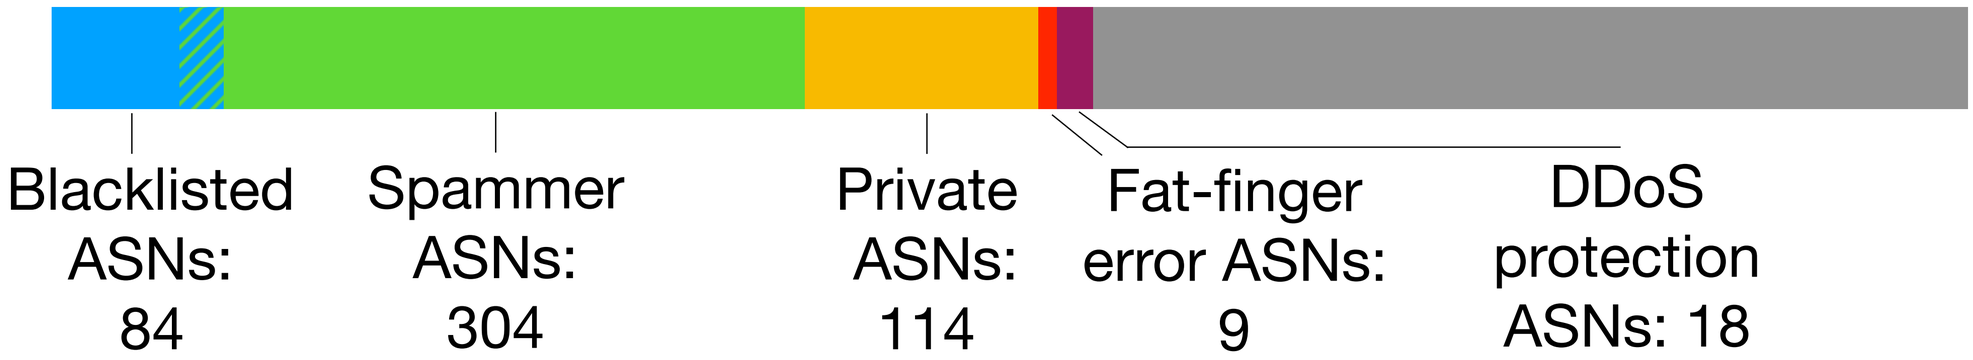

What are ASes flagged by our classifier?

- Indication of malicious behavior
- Indication of misconfigurations
- Known false positives

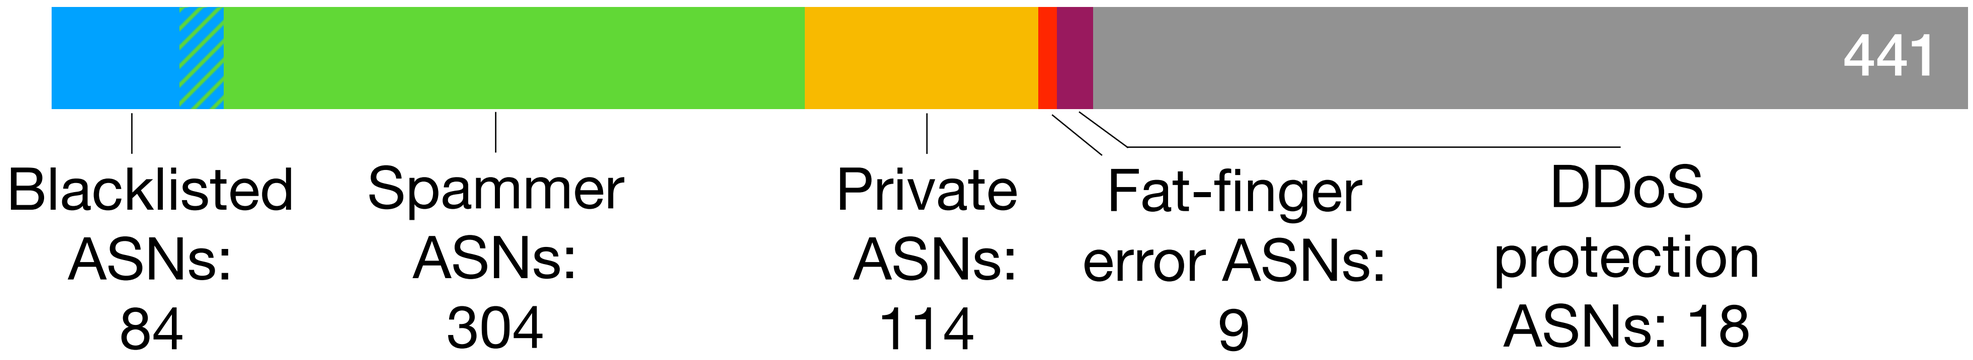

What are ASes flagged by our classifier?

- Indication of malicious behavior
- Indication of misconfigurations
- Known false positives

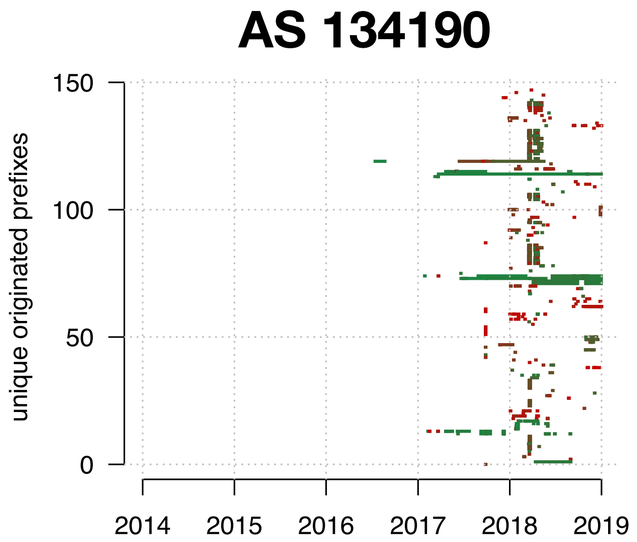

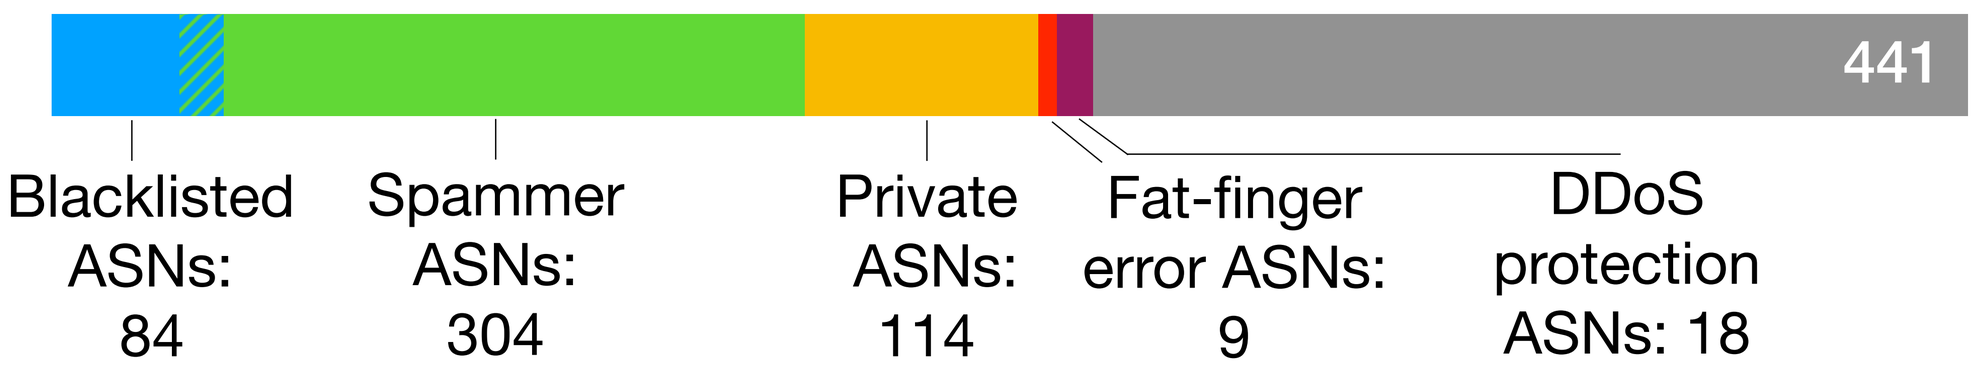

What are ASes flagged by our classifier?

- Indication of malicious behavior
- Indication of misconfigurations
- Known false positives

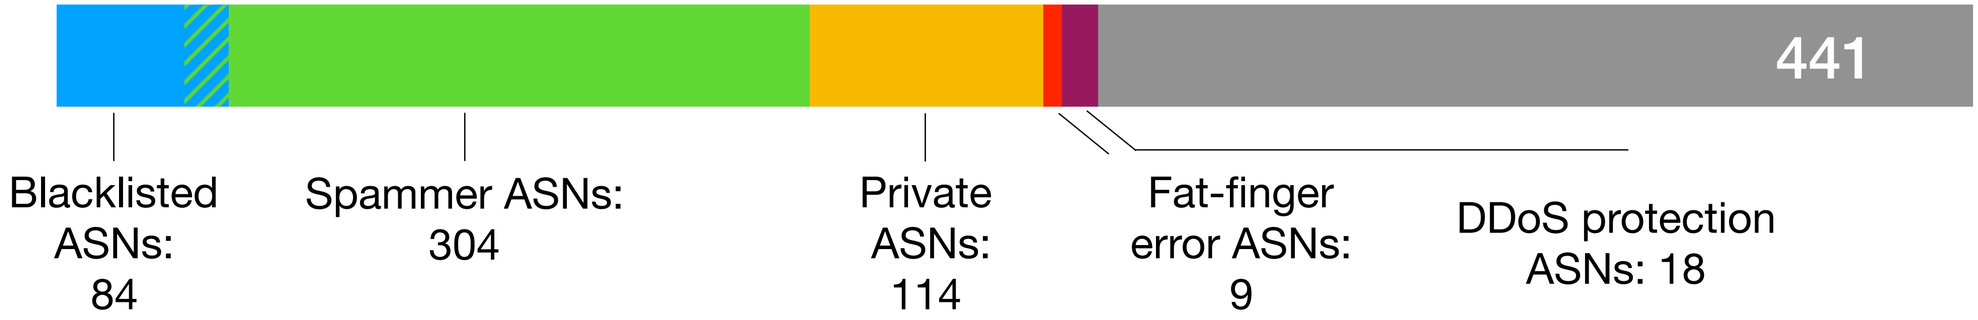

- 53% of flagged ASes are in known categories.
- Many interesting ASes are in the other 47%.

What our classifier is not…

- A bulletproof identifier of malicious ASes.
- A system that exhaustively captures hijackers.

Key takeaways

- **First** longitudinal analysis of **serial hijacker** ASes.
- Features offer **state of affairs** of AS-wide **BGP behavior**.
- Classifier outcome provides **new data for network reputation** scoring systems.
- Effectively **narrows the focus on suspicious networks**, with much future work to be done.

Key takeaways

- **First** longitudinal analysis of serial hijacker ASes.
- Features offer **state of affairs** of AS-wide BGP behavior.
- Classifier outcome provides **new data for network reputation** scoring systems.
- Effectively **narrows the focus on suspicious networks**, with much future work to be done.