catchpoint™

Every Experience Matters

One year of BGP (in)security

NANOG 78

February 11th, 2020

San Francisco, USA

Luca Sani

lsani@catchpoint.com

Alessandro Improta

aimprota@catchpoint.com

This year BGP incidents

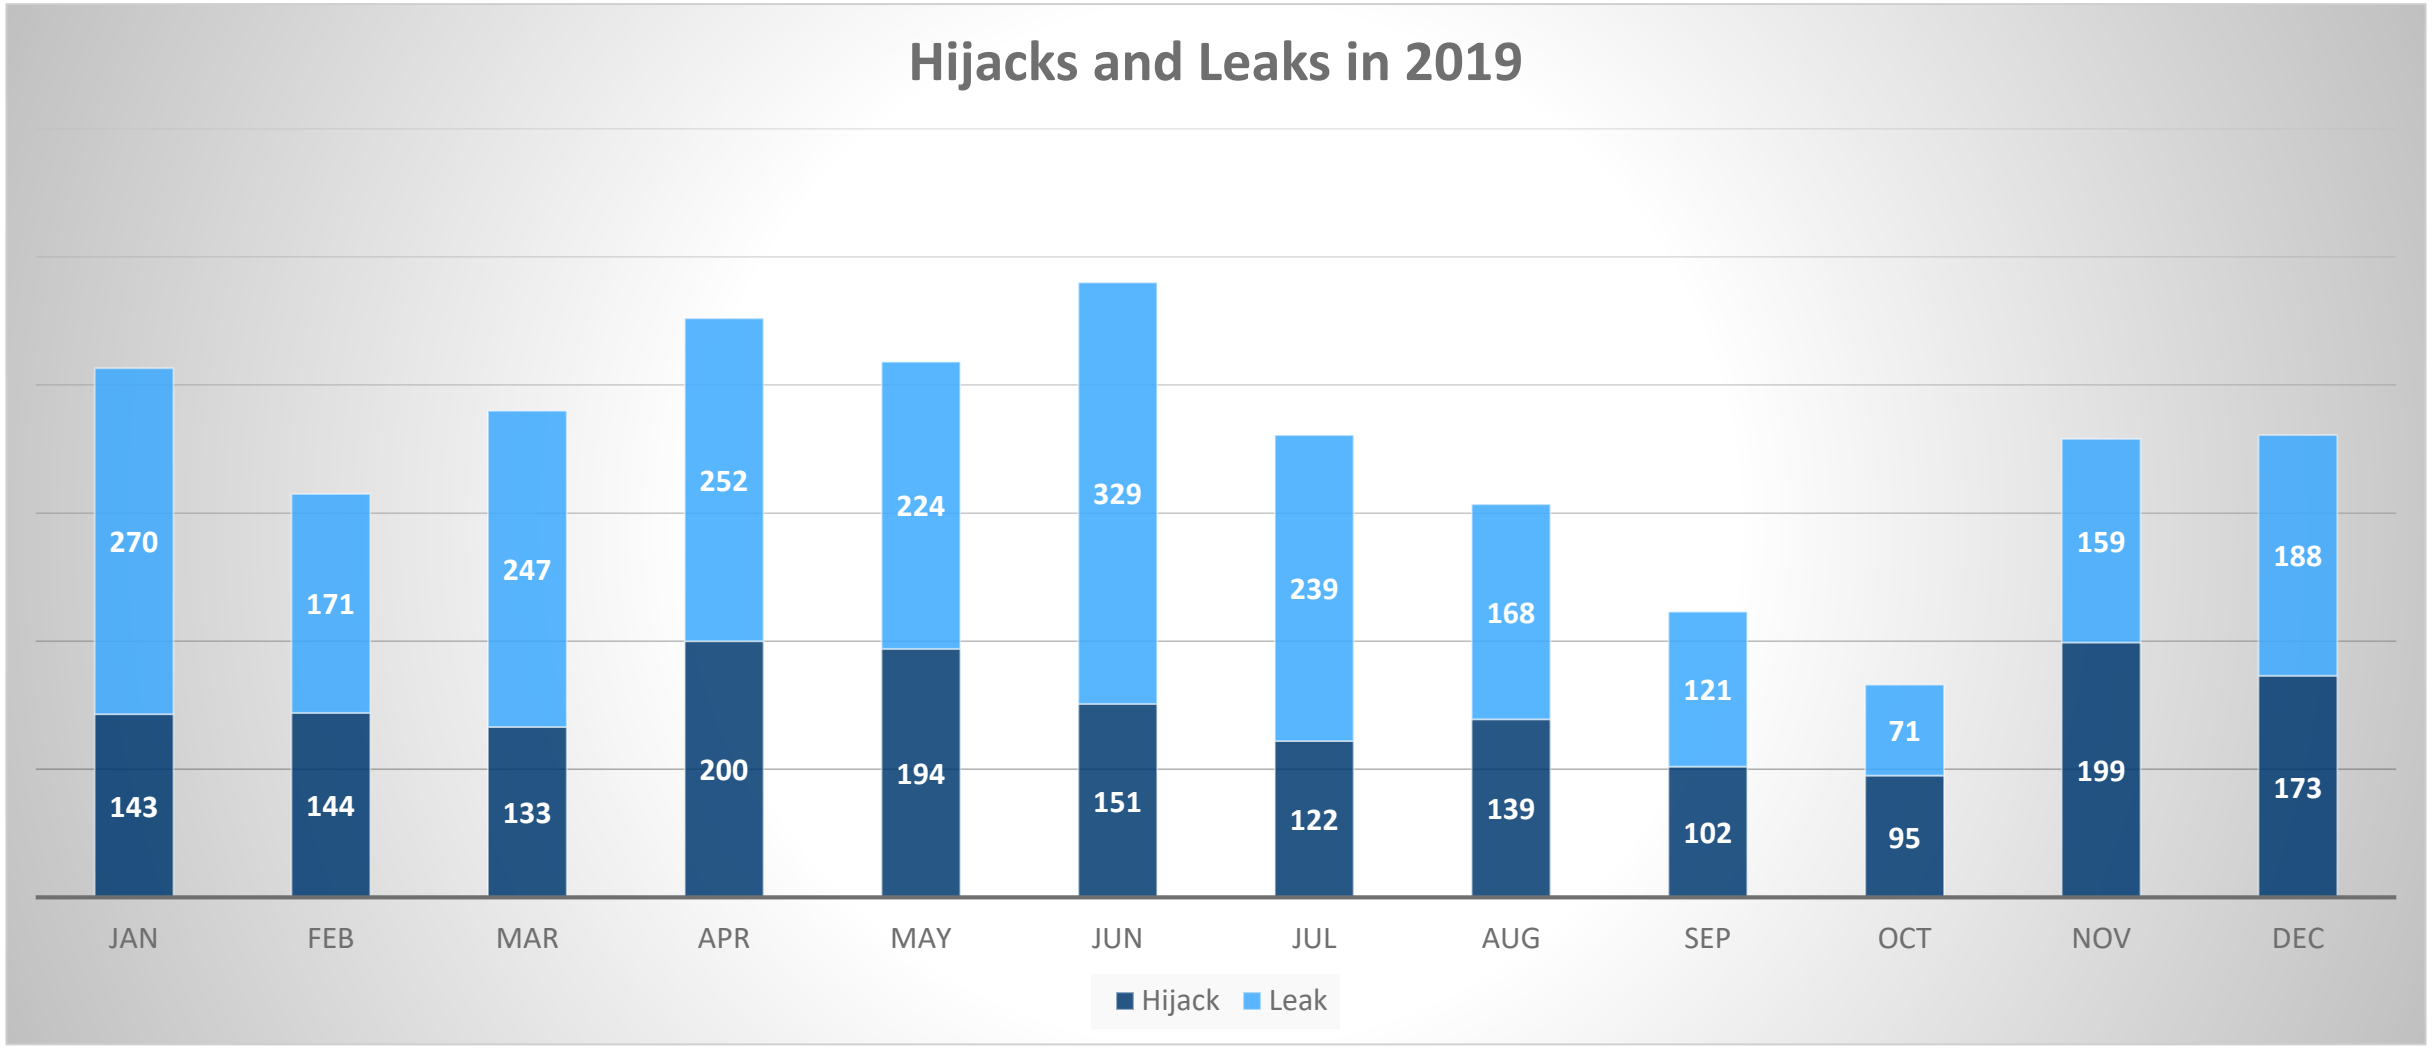

source: https://bgpstream.com/

BGP and security

- BGP turned 30 years old last year!
  - The first version was designed in 1989 by K. Lougheed and Y. Rekhter
  - The current version (four) was standardized in 1994
- BGP was not designed with a focus on security
  - *"In the early days of the Internet, getting stuff to work was the primary goal. There was no concept that people would use this to do malicious things... Security was not a big issue."* (K. Lougheed)
  - Security *"wasn't even on the table"* (Y. Rekhter)
- Therefore it lacks a built-in mechanism to authenticate packets
- BGP is prone to attacks and misconfigurations
  - Prefix hijacks
  - Route leaks

Quotes from: **https://www.washingtonpost.com/sf/business/2015/05/31/net-of-insecurity-part-2**

Prefix hijack

- A prefix hijack happens when an AS originates a prefix that has not been allocated to it
  - Often called mis-origination

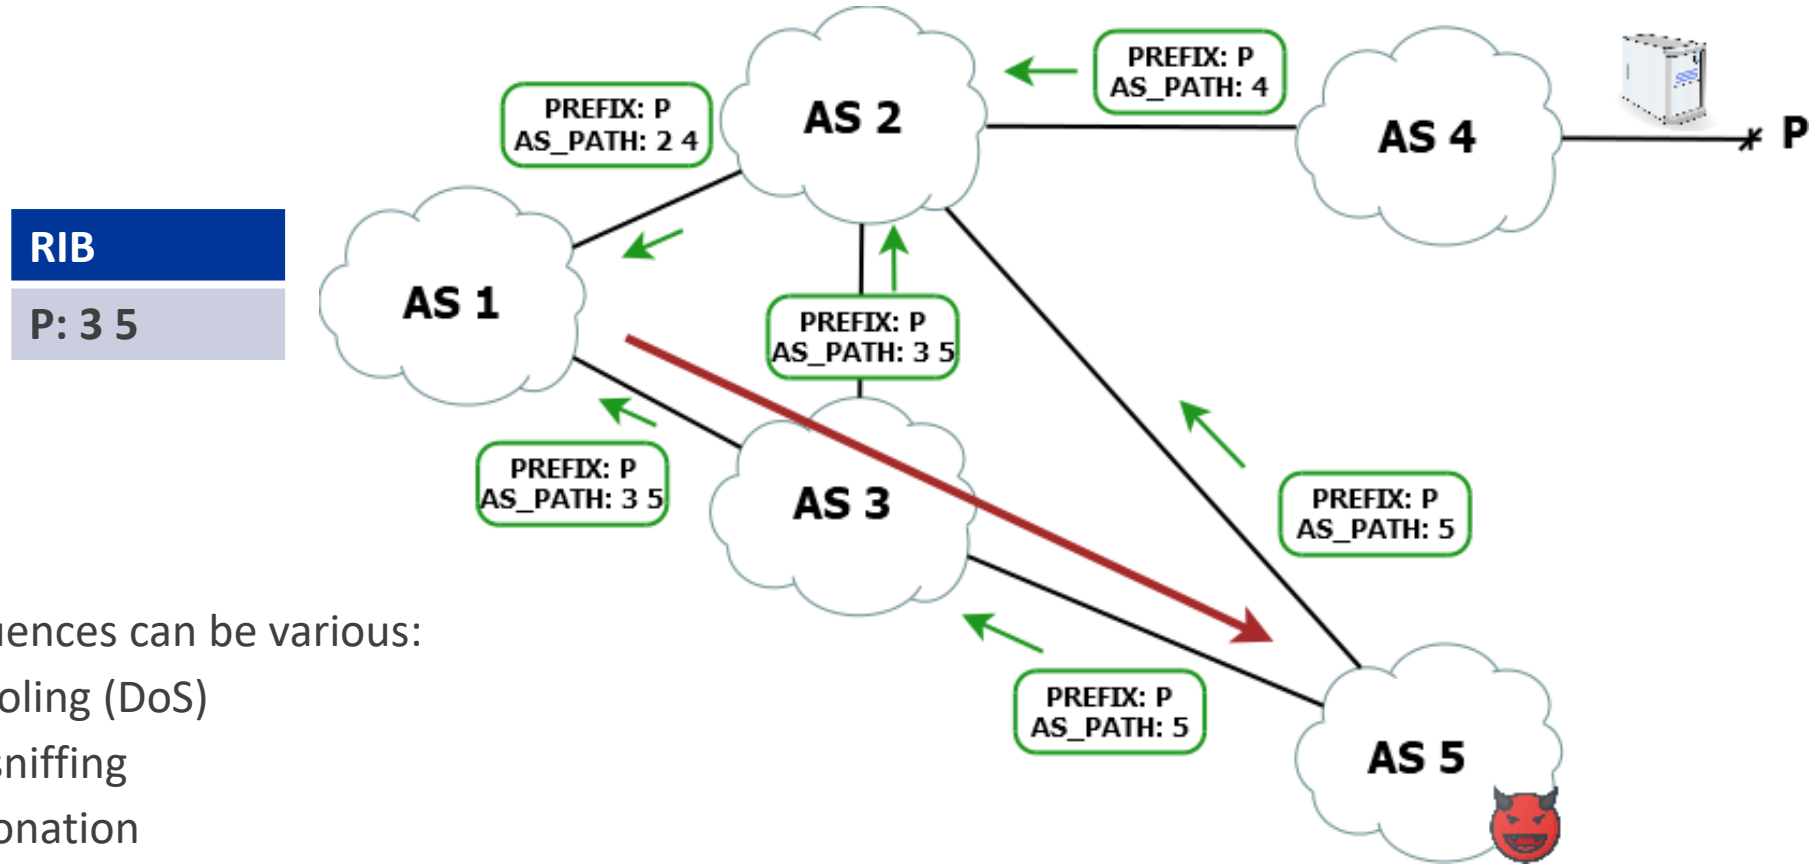

- The consequences can be various:
  - Black-holing (DoS)
  - Traffic sniffing
  - Impersonation

Prefix hijack... not always that easy to detect!

- The attacker forges the AS_PATH on order to include the expected origin (AS_PATH forgery hijack)

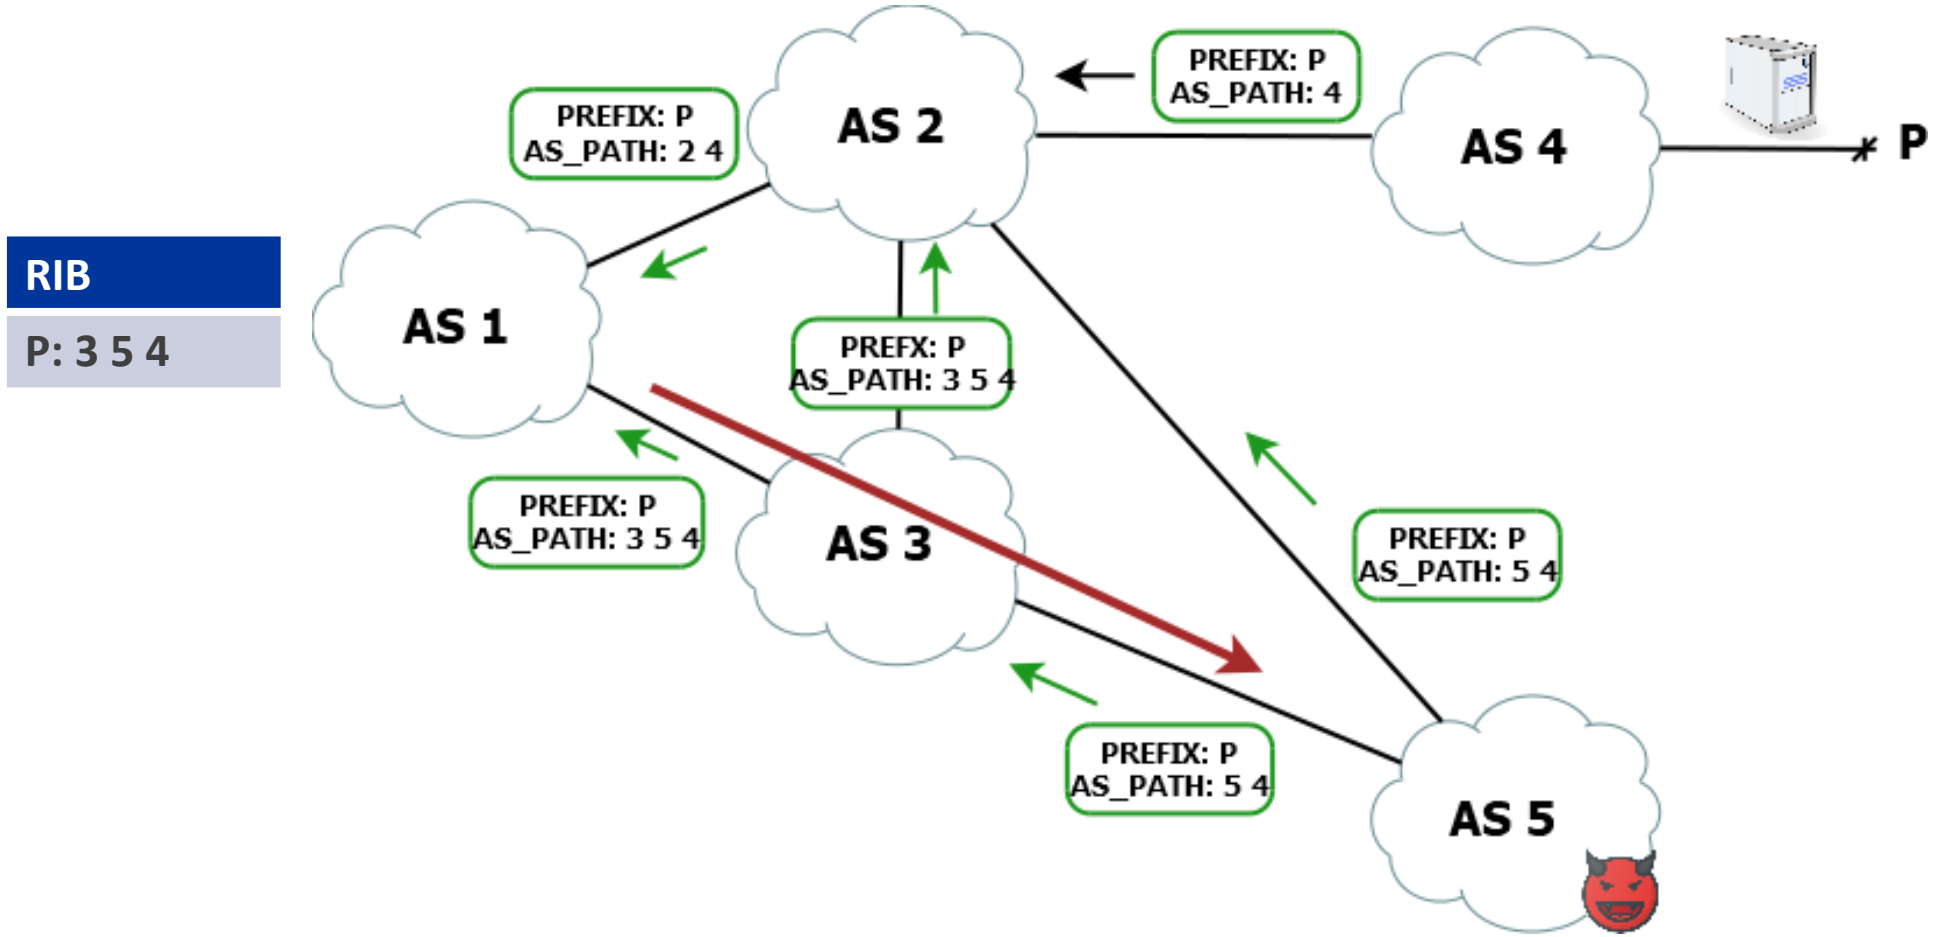

Route leaks

- A route leak is the *propagation of a BGP announcement(s) beyond their intended scope* [RFC 7908]

- BGP is governed by commercial agreements between ASes:
  - **customer-to-provider (c2p):** one of the two ASes (the provider) is providing transit to the whole Internet for the other AS (the customer). Usually the customer pays the provider
  - **peer-to-peer (p2p):** the two ASes decide to announce each other the networks which each AS can reach without using any transit connection or any other p2p relationship. Usually it is a settlement-free agreement

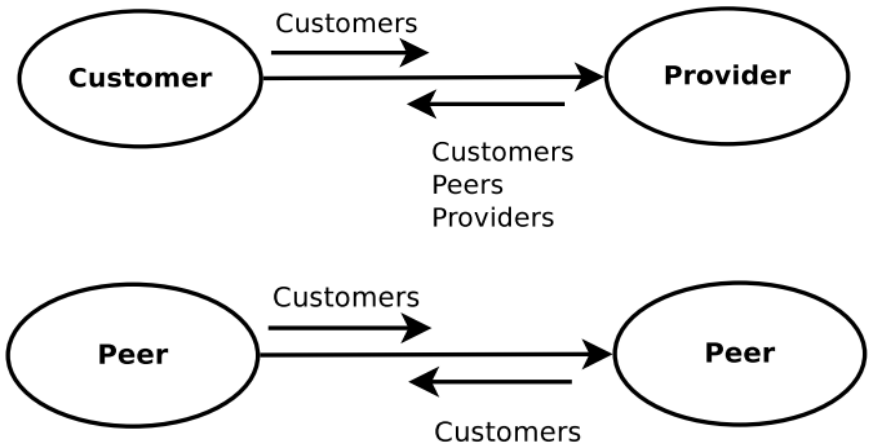

Route leaks

- A customer should not transit traffic between two providers (or peers)!
  - It is not getting paid
  - Its network may be under-provisioned to handle this traffic (performance degradation, DoS)
  - **This** is a leak!

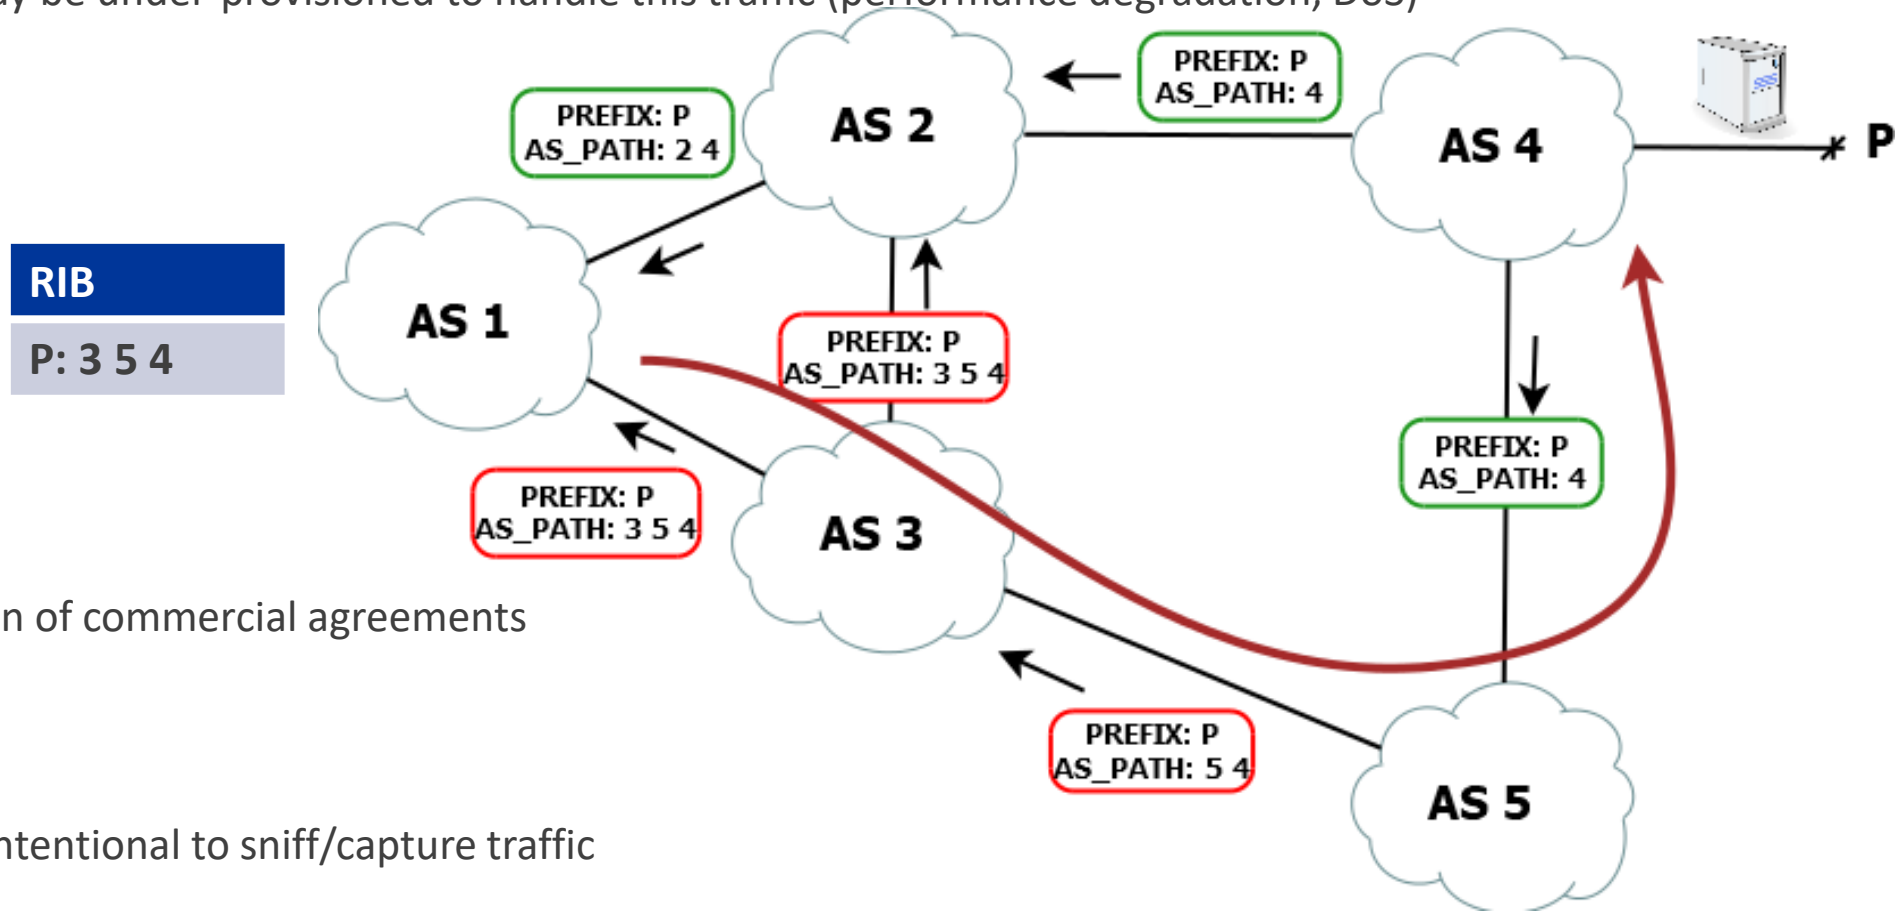

- Unintended violation of commercial agreements
  - Fat finger?
  - Bad filters?

- Also, this could be intentional to sniff/capture traffic

Big trouble in little Switzerland

In the news

GIZMODO | We come from the future

PRIVACY AND SECURITY

China Telecom Swallows Huge Amount of European Mobile Traffic For Over Two Hours

Dell Cameron
6/07/19 7:17PM • Filed to: BGP ROUTE LEAK

31.0K 22 5

INTERNET INTELLIGENCE | ORACLE Cloud Infrastructure

-FEATURED ARTICLES, DISRUPTIONS

June 6, 2019

Large European Routing Leak Sends Traffic Through China Telecom

Doug Madory

computing

Security

BGP route leak sends European mobile traffic via China

Yet another BGP hijack by China Telecom routes internet traffic of several European mobile operators via China

ZDNet

For two hours, a large chunk of European mobile traffic was rerouted through China

It was China Telecom, again. The same ISP accused last year of "hijacking the vital internet backbone of western countries."

By Catalin Cimpanu for Zero Day | June 7, 2019 -- 19:41 GMT (20:41 BST) | Topic: Security

ars TECHNICA

THANKS, BGP. —

BGP event sends European mobile traffic through China Telecom for 2 hours

Improper leak to Chinese-government-owned telecom lasts up to two hours.

DAN GOODIN - 6/8/2019, 12:05 PM

The Register®
Biting the hand that feeds IT

Data Centre ▸ Networks

You won't guess where European mobile data was rerouted for two hours. Oh. You can. Yes, it was China Telecom

BGP leaks are common but don't usually take hours to fix...

By Kieren McCarthy in San Francisco 10 Jun 2019 at 20:03 62 SHARE ▼

What happened?

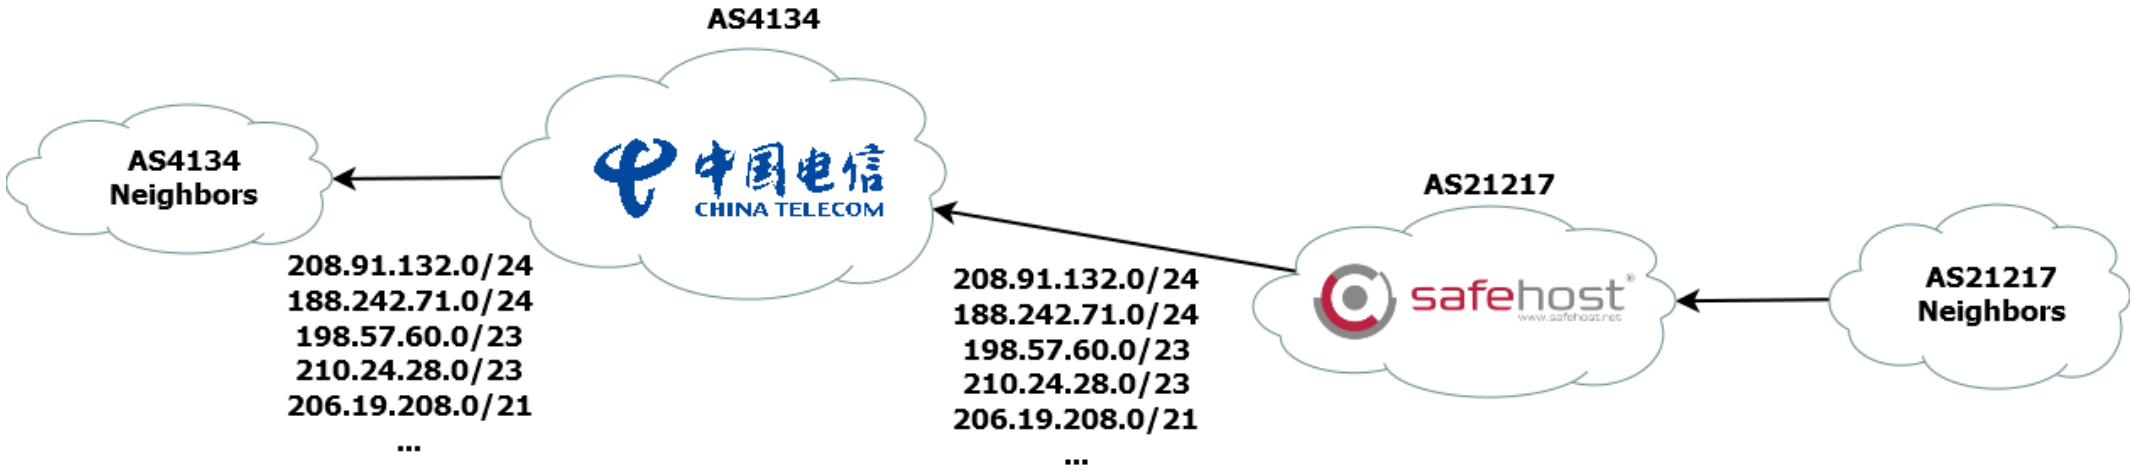

2

China Telecom accepted and propagated those routes to its neighbors (more than 40 neighbors)

1

SafeHost leaked routes regarding more than 40k destinations to China Telecom.
Those destinations were both more and less specific than existing ones

- When: 6 June, 2019 about 09:40am - 1:00pm UTC
- Example of routes seen:
  - `195.209.0.0/19  61832 2914 4134 21217 21217 21217 21217 21217 21217 25091 5568`
  - `129.95.100.0/24 37468 6453 4134 21217 21217 21217 21217 21217 21217 6830 2603 11164 11995`
  - `208.91.132.0/24 7660 2516 4134 21217 21217 21217 21217 21217 21217 3356 15085`

Statistics by collector peer

- Almost every peer sharing a full routing table with Route Views and RIPE NCC RIS detected the leak

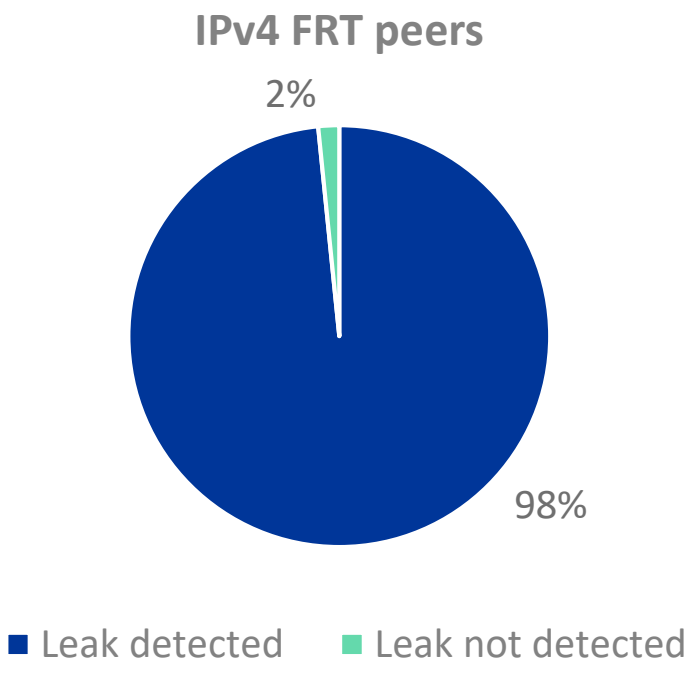

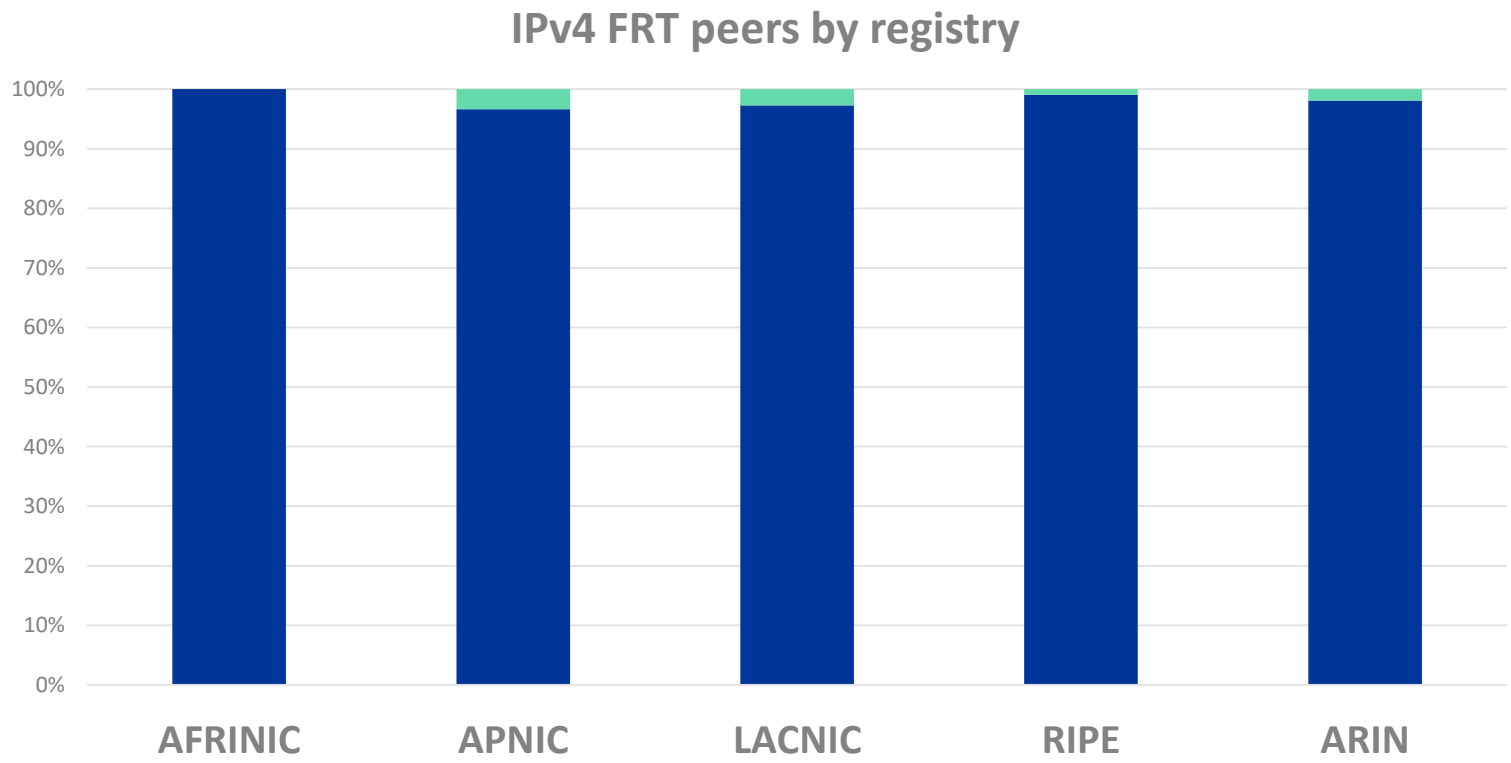

Registry mapping thanks to Team Cymru

Statistics about involved parties

Only IPv4 networks were affected

CCDF number of origin AS's involved per FRT peer

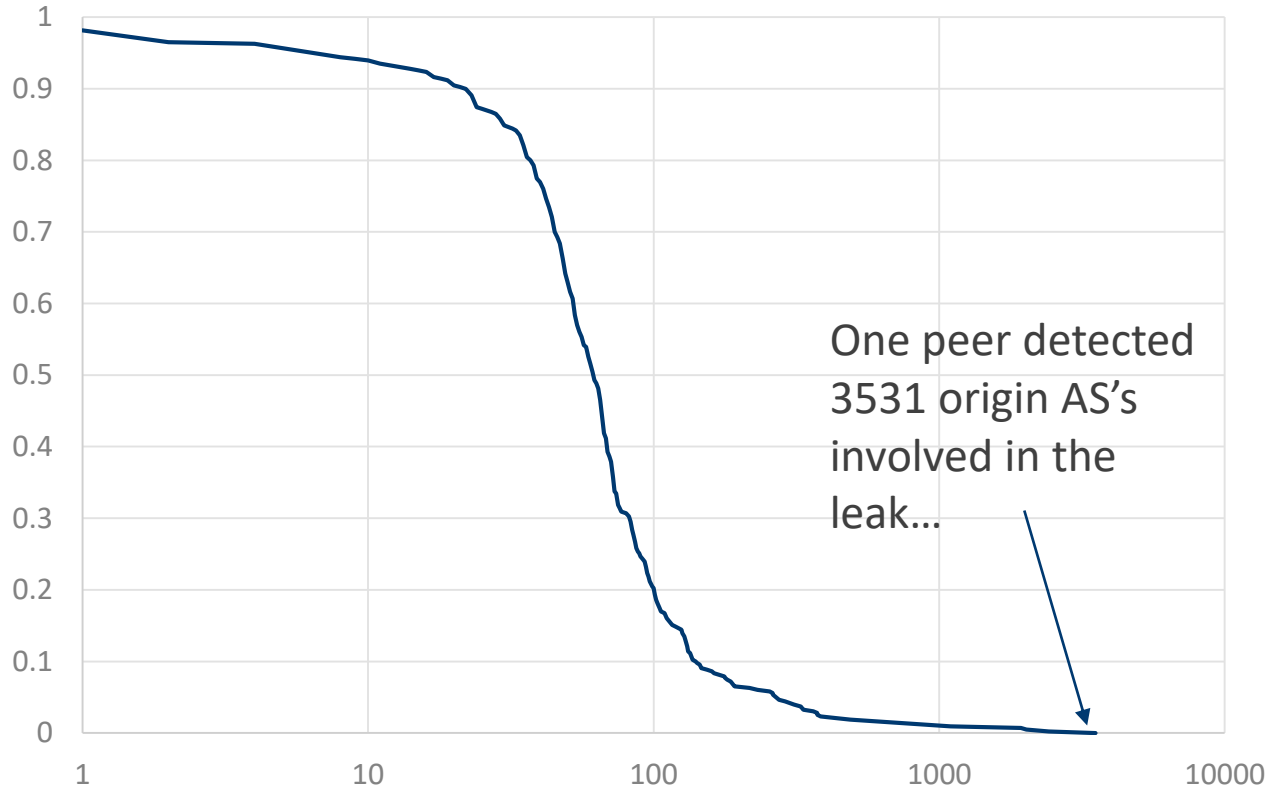

CCDF number of IPv4 subnets involved per FRT peer

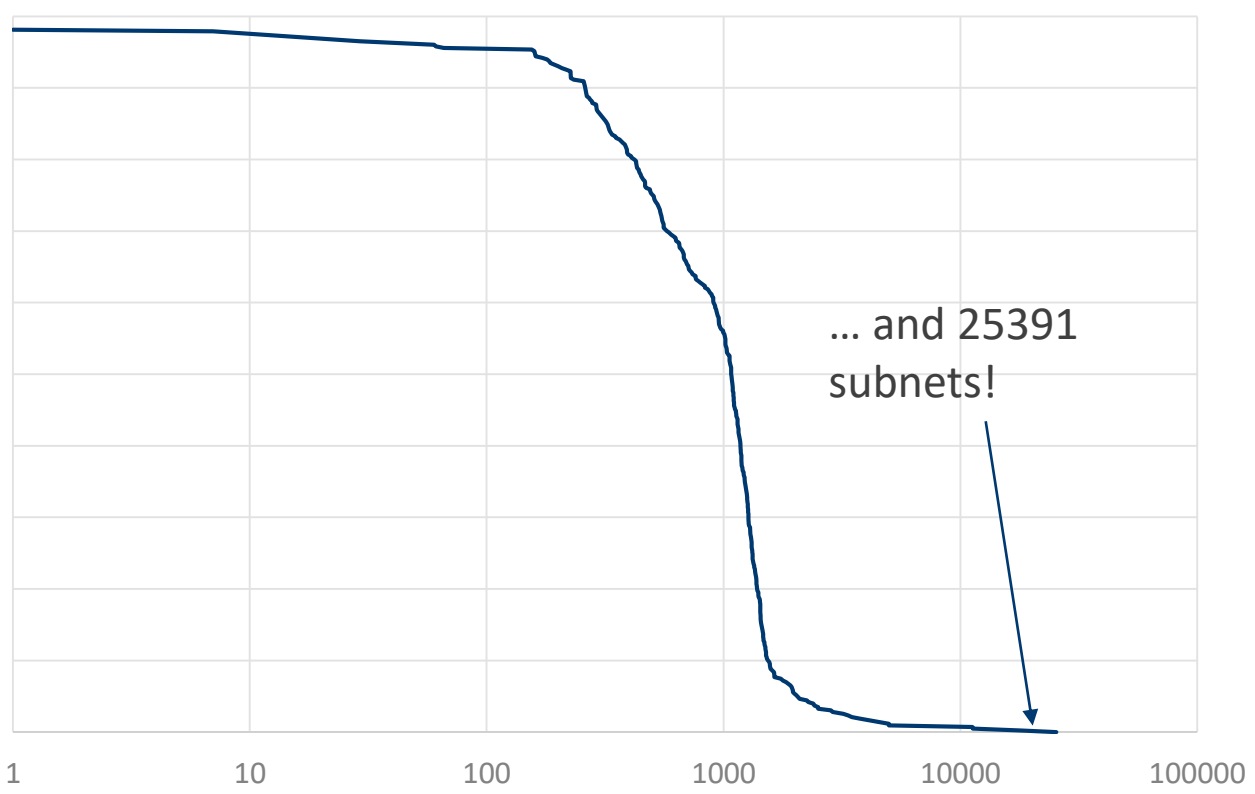

Peers directly connected to China Telecom are seeing the highest number of leaked routes

Origin AS's affected

- More than 6000 different origin AS's involved
- Popular services affected
  - WhatsApp
  - Microsoft
  - …
- Hosting providers
- Transit providers
- Banks
- …

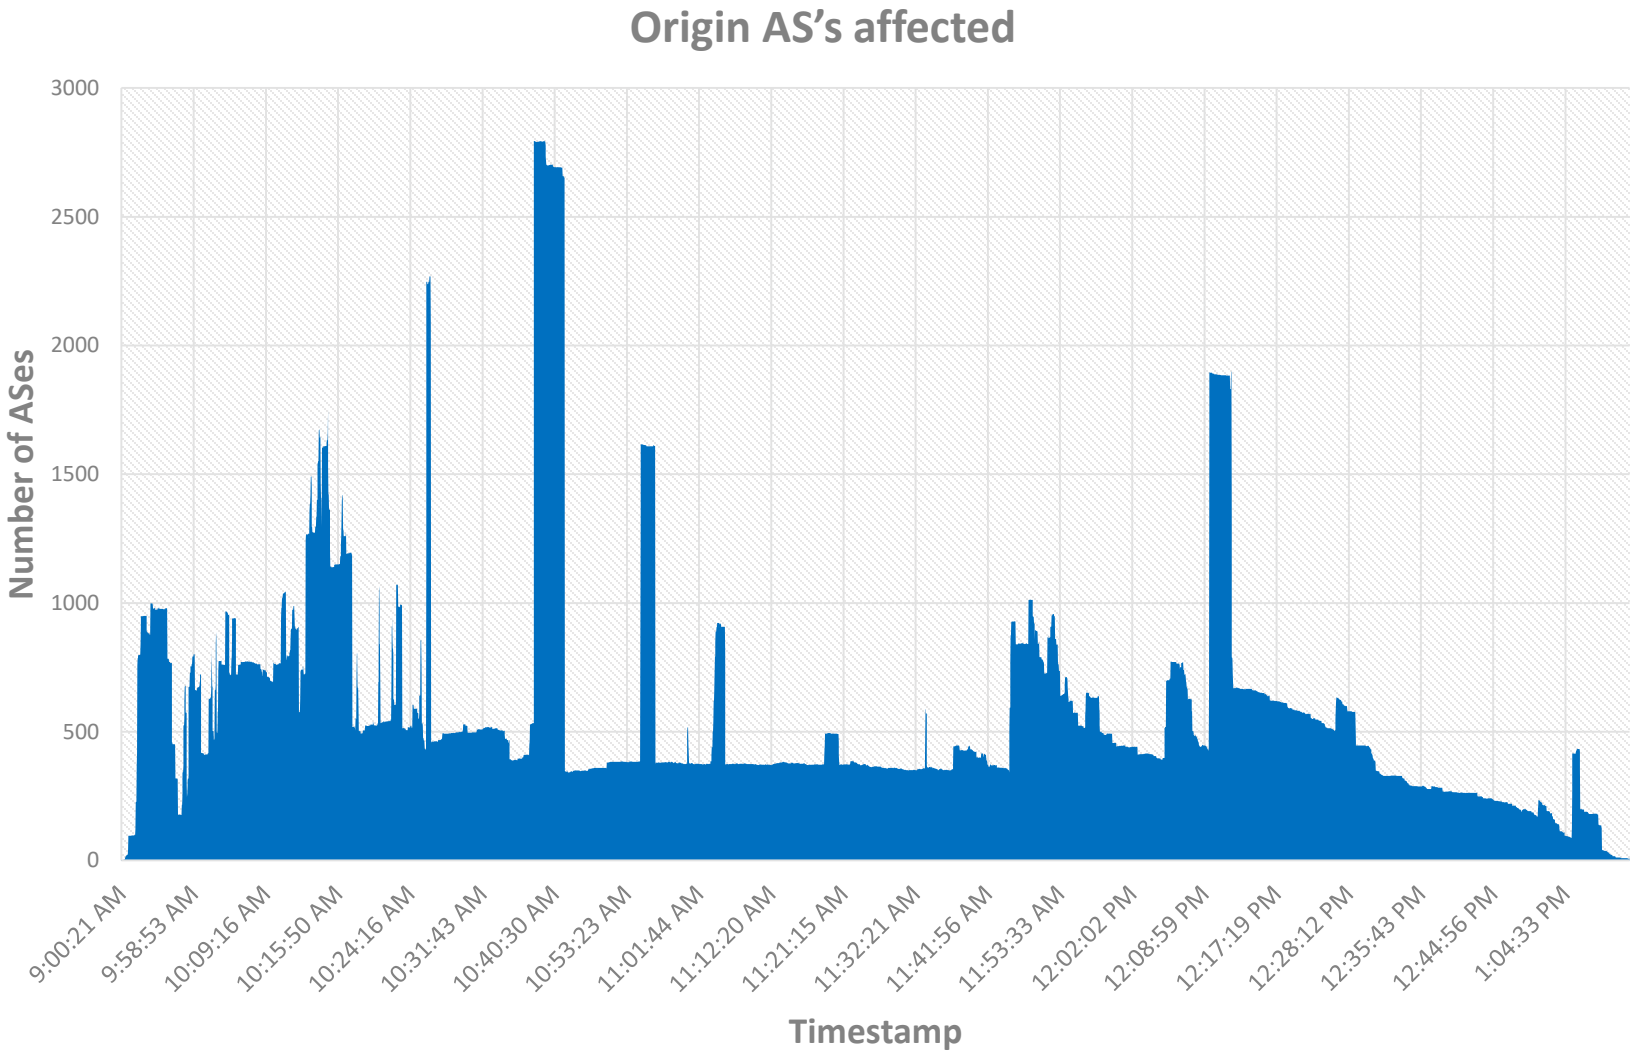

Cutthroat island… just for a few minutes

In the news


Public DNS in Taiwan the latest victim to BGP hijack
May 15, 2019 by Aftab Siddiqui Leave a Comment

What happened?

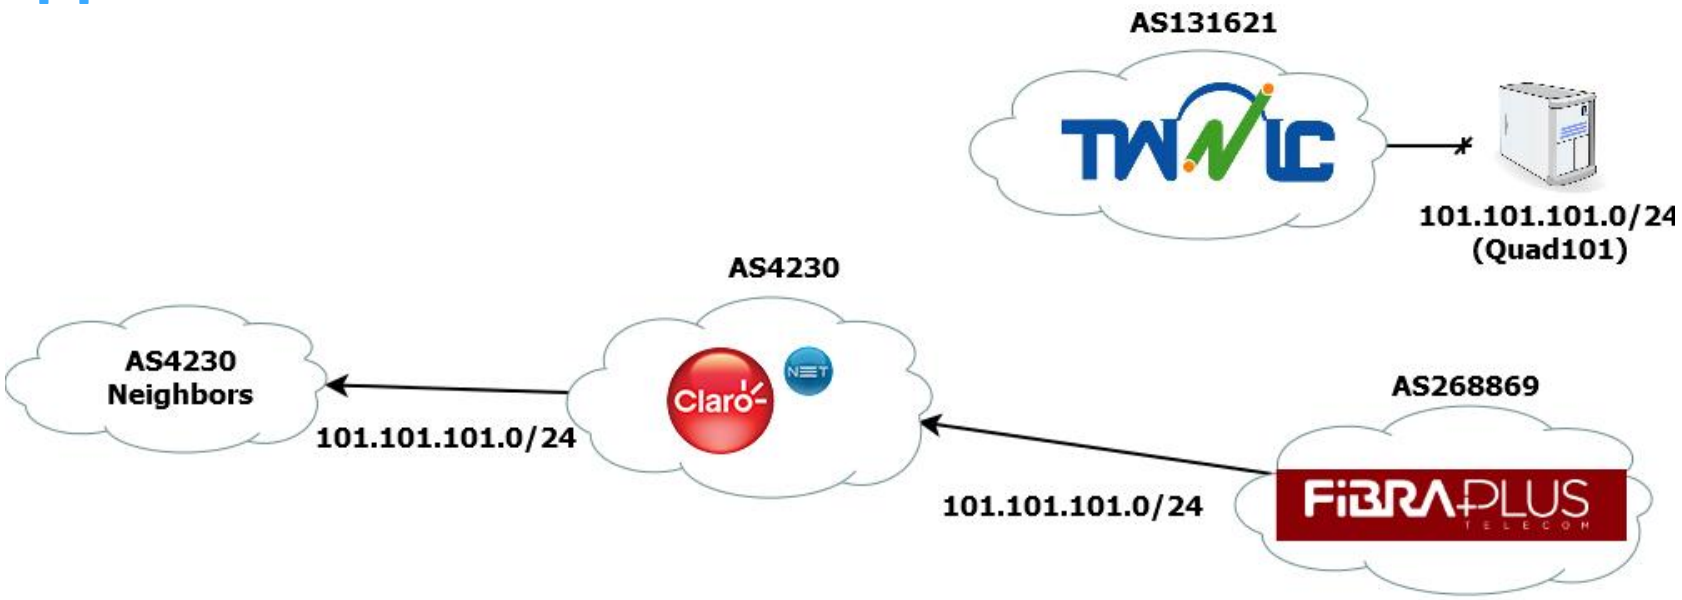

- When: May 8, 2019 about 15:08 UTC to 15:11 UTC
- Example of routes seen:
  - `101.101.101.0/24 8492 9002 4230 268869`
  - `101.101.101.0/24 20912 1267 3356 2828 4230 268869`
  - `101.101.101.0/24 6939 2828 4230 268869`

Statistics by collector peer

- All LACNIC peers detected the leak

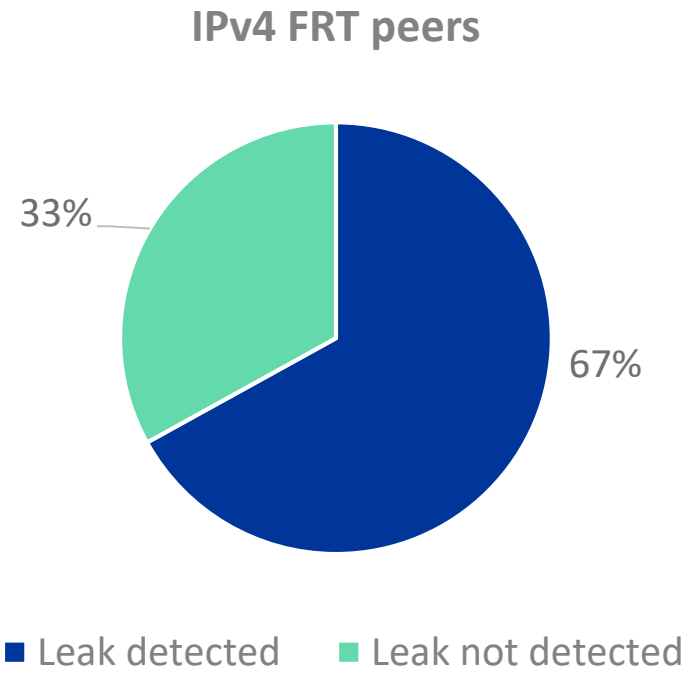

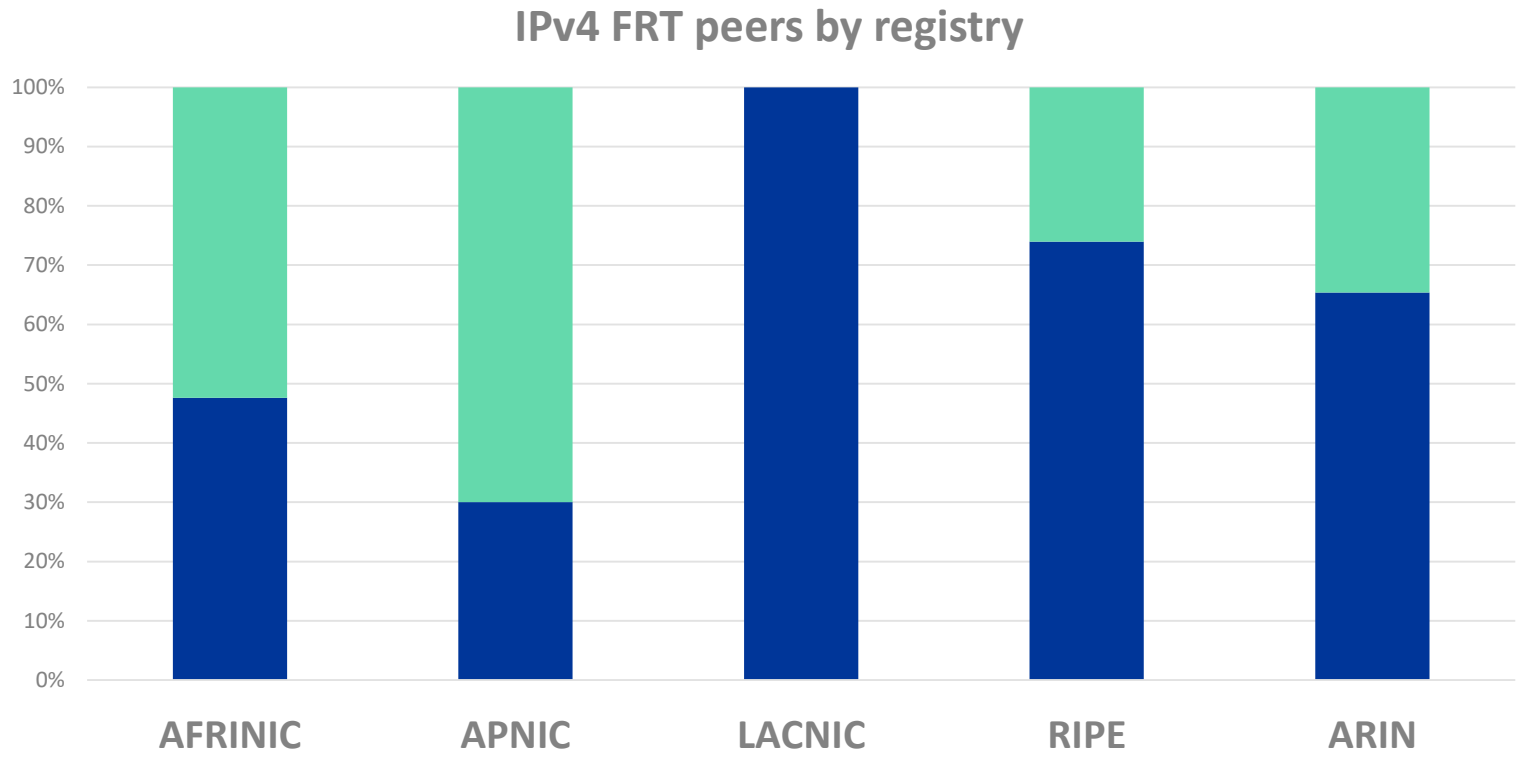

Registry mapping thanks to Team Cymru

The Cloudflare case

In the news


The Register
Biting the hand that feeds IT
Data Centre ▸ Cloud
Cloudflare hits the deck, websites sink from sight after the internet springs yet another BGP leak
Ghost in the machine conspires to ruin CDN biz's 10th birthday, it seems
By Richard Speed 24 Jun 2019 at 13:07
19
SHARE

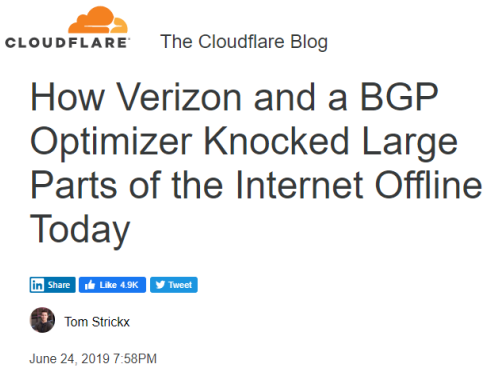
CLOUDFLARE
The Cloudflare Blog
How Verizon and a BGP Optimizer Knocked Large Parts of the Internet Offline Today
Share
Like 4.9K
Tweet
Tom Strickx
June 24, 2019 7:58PM


catchpoint
June 26, 2019
USER EXPERIENCE
BGP Leak Highlights the Fragility of the Internet with Real Consequences


THE VERGE
Discord was down due to Cloudflare and Verizon issues
Cloudflare had to deal with Verizon creating a mess
By Tom Warren | @tomwarren | Jun 24, 2019, 8:29am EDT


The Register
Biting the hand that feeds IT
Data Centre ▸ Networks
BGP super-blunder: How Verizon today sparked a 'cascading catastrophic failure' that knackered Cloudflare, Amazon, etc
'Normally you'd filter it out if some small provider said they own the internet'
By Kieren McCarthy in San Francisco 24 Jun 2019 at 19:01
61
SHARE

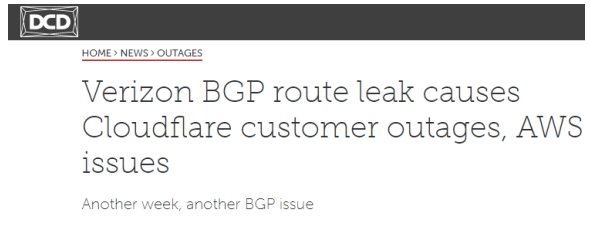
DCD
HOME > NEWS > OUTAGES
Verizon BGP route leak causes Cloudflare customer outages, AWS issues
Another week, another BGP issue
June 24, 2019 By: Sebastian Moss


itnews
Route leak causes internet problems worldwide
By Ry Crozier
Jun 24 2019
10:46PM
Cloudflare, AWS, Google network routes among those impacted.

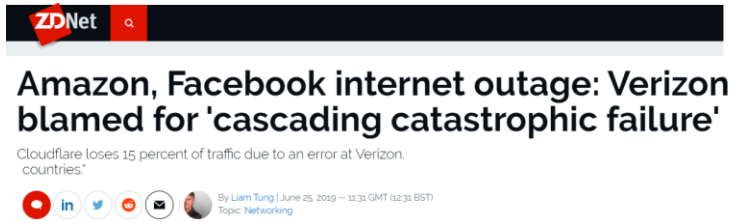
ZDNet
Amazon, Facebook internet outage: Verizon blamed for 'cascading catastrophic failure'
Cloudflare loses 15 percent of traffic due to an error at Verizon.
countries.'
By Liam Tung | June 25, 2019 -- 11:31 GMT (12:31 BST)
Topic: Networking

What happened?

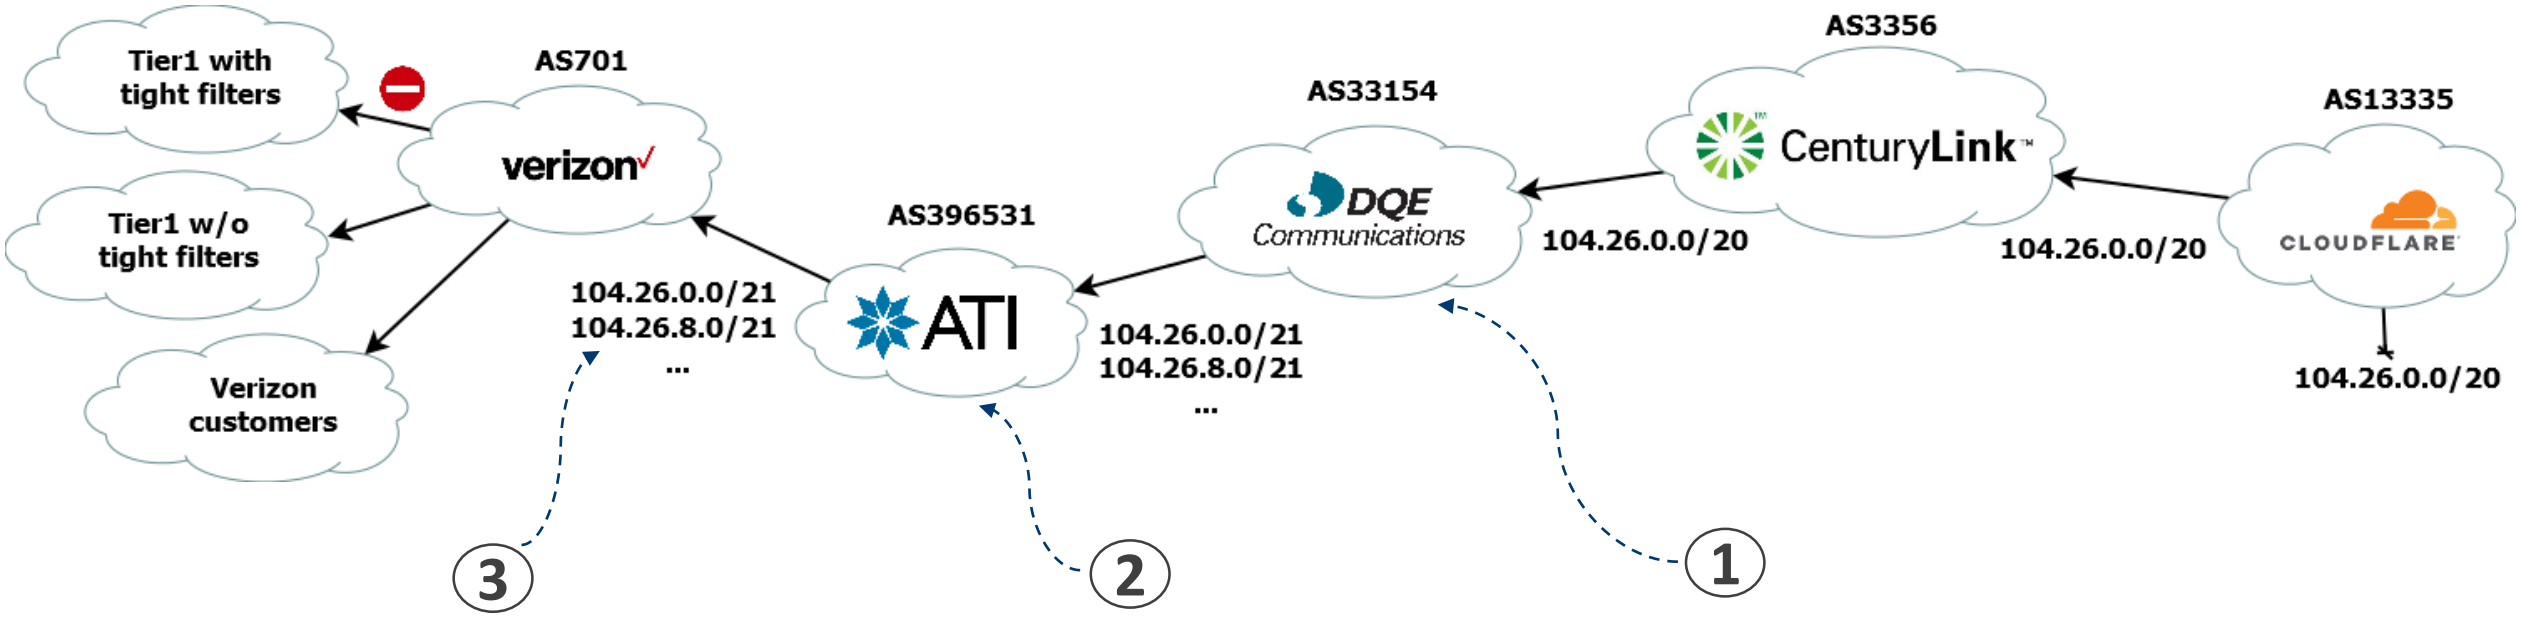

Verizon accepted and propagated these routes to its numerous neighbors

ATI accepted and propagated these routes to Verizon (Provider to Provider leak)

DQE sent to ATI more specific routes to reach Cloudflare network

- When: 24 June, 2019 about 10:30am - 12:30pm UTC
- Routes seen had the form:

```
104.26.0.0/21 … 701 396531 33154 3356 13335
104.26.8.0/21 … 701 396531 33154 3356 13335
```

- Part of the Internet used leaked routes when sending packets to 104.26.0.0/20 (longest match wins)

Statistics about involved subnets

- Almost every peer sharing a full routing table with Route Views and RIPE NCC RIS detected the leak

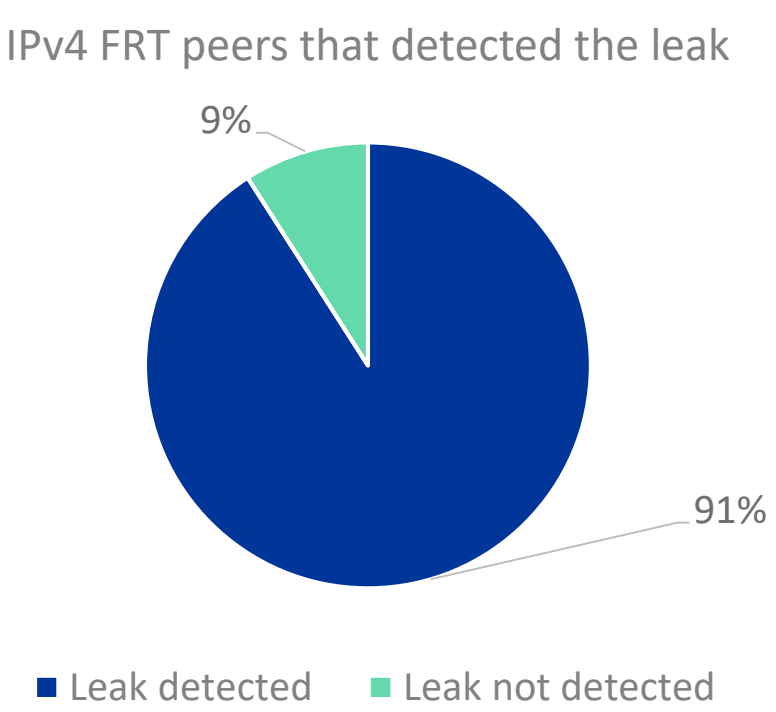

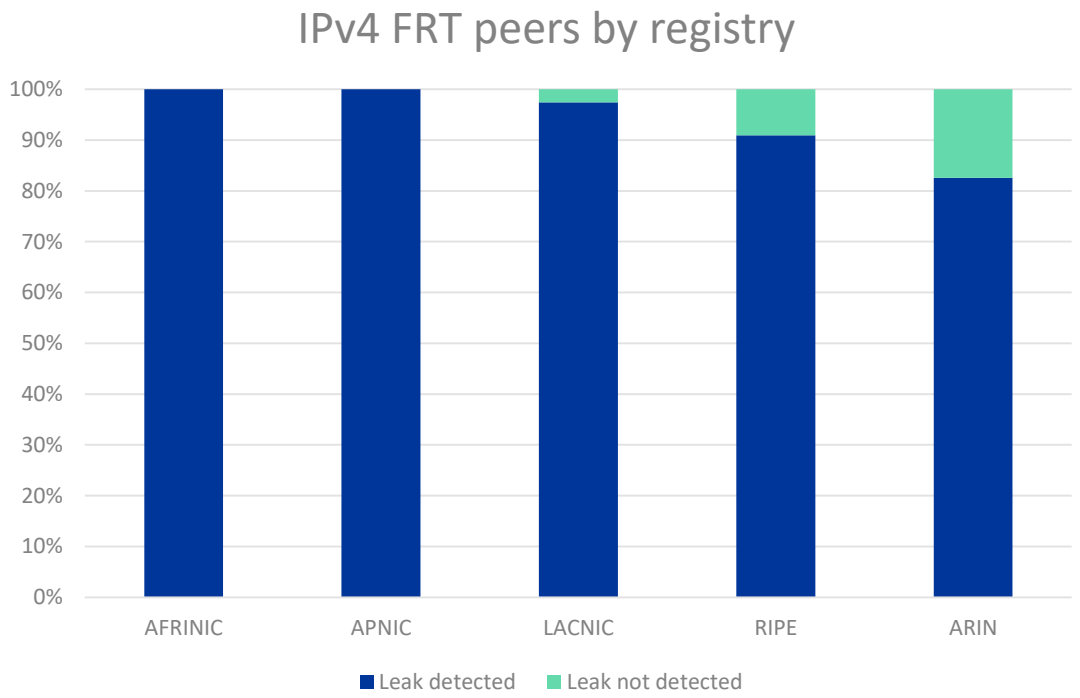

Registry mapping thanks to Team Cymru

Statistics about involved parties

Only IPv4 networks were affected

CCDF number of origin AS's involved per FRT peer

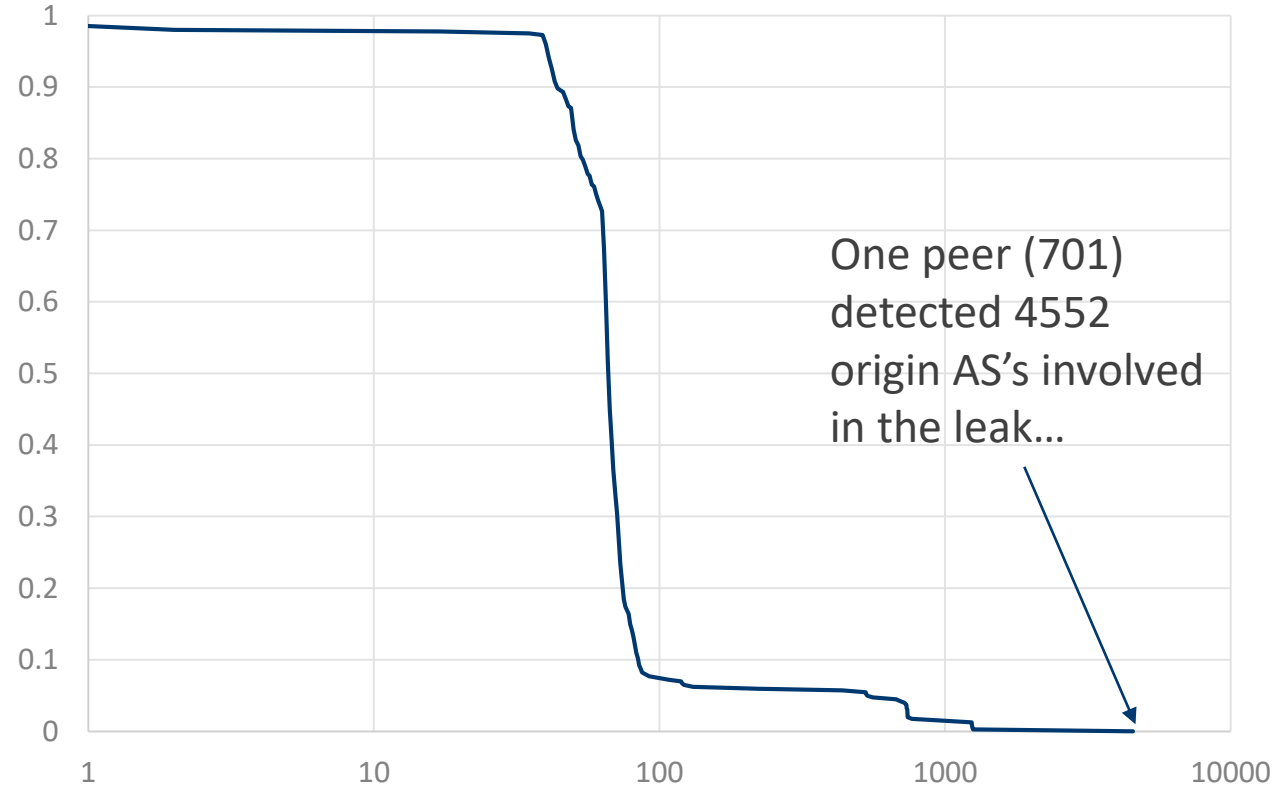

CCDF number of IPv4 subnets involved per FRT peer

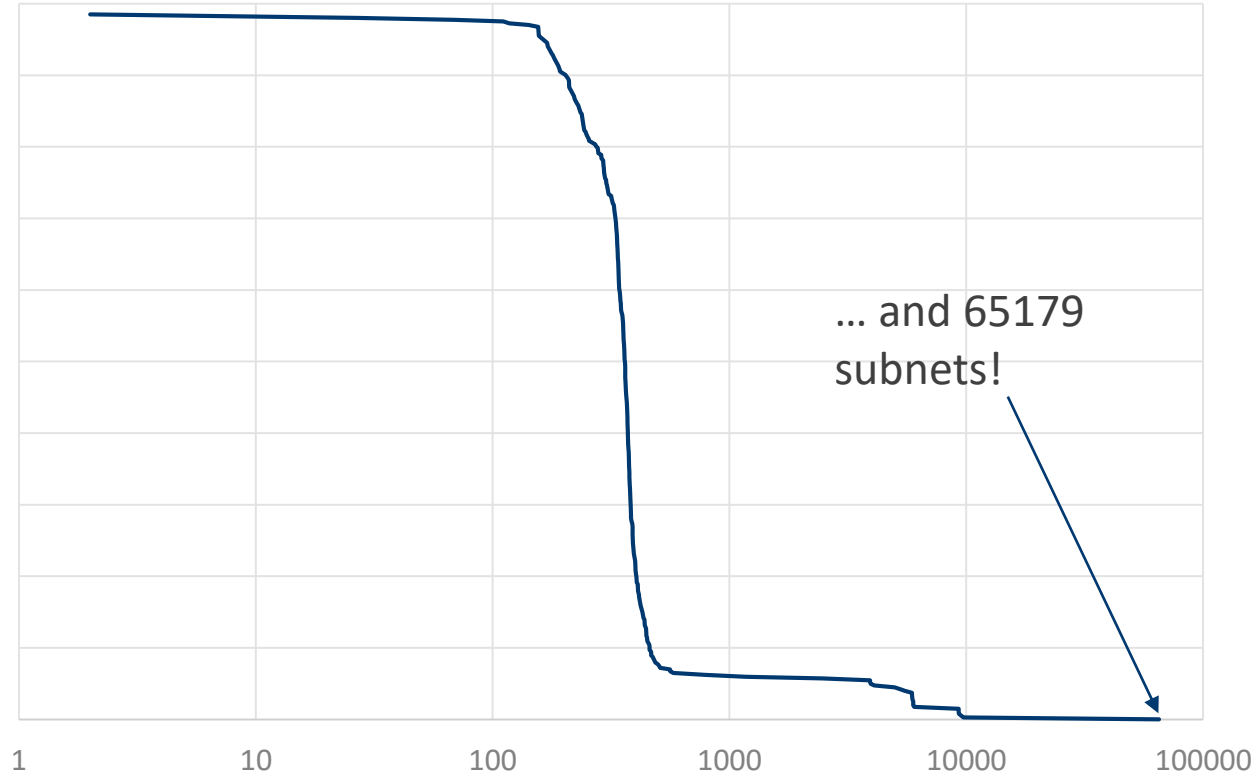

The peer is Verizon (AS701) which is connected to Route Views (route-views2)

Who was affected?

- The leak didn't affect only Cloudflare...
- More than 1200 ASes involved
- Facebook, Comcast, T-Mobile, Bloomberg, ...
- 9 American banks

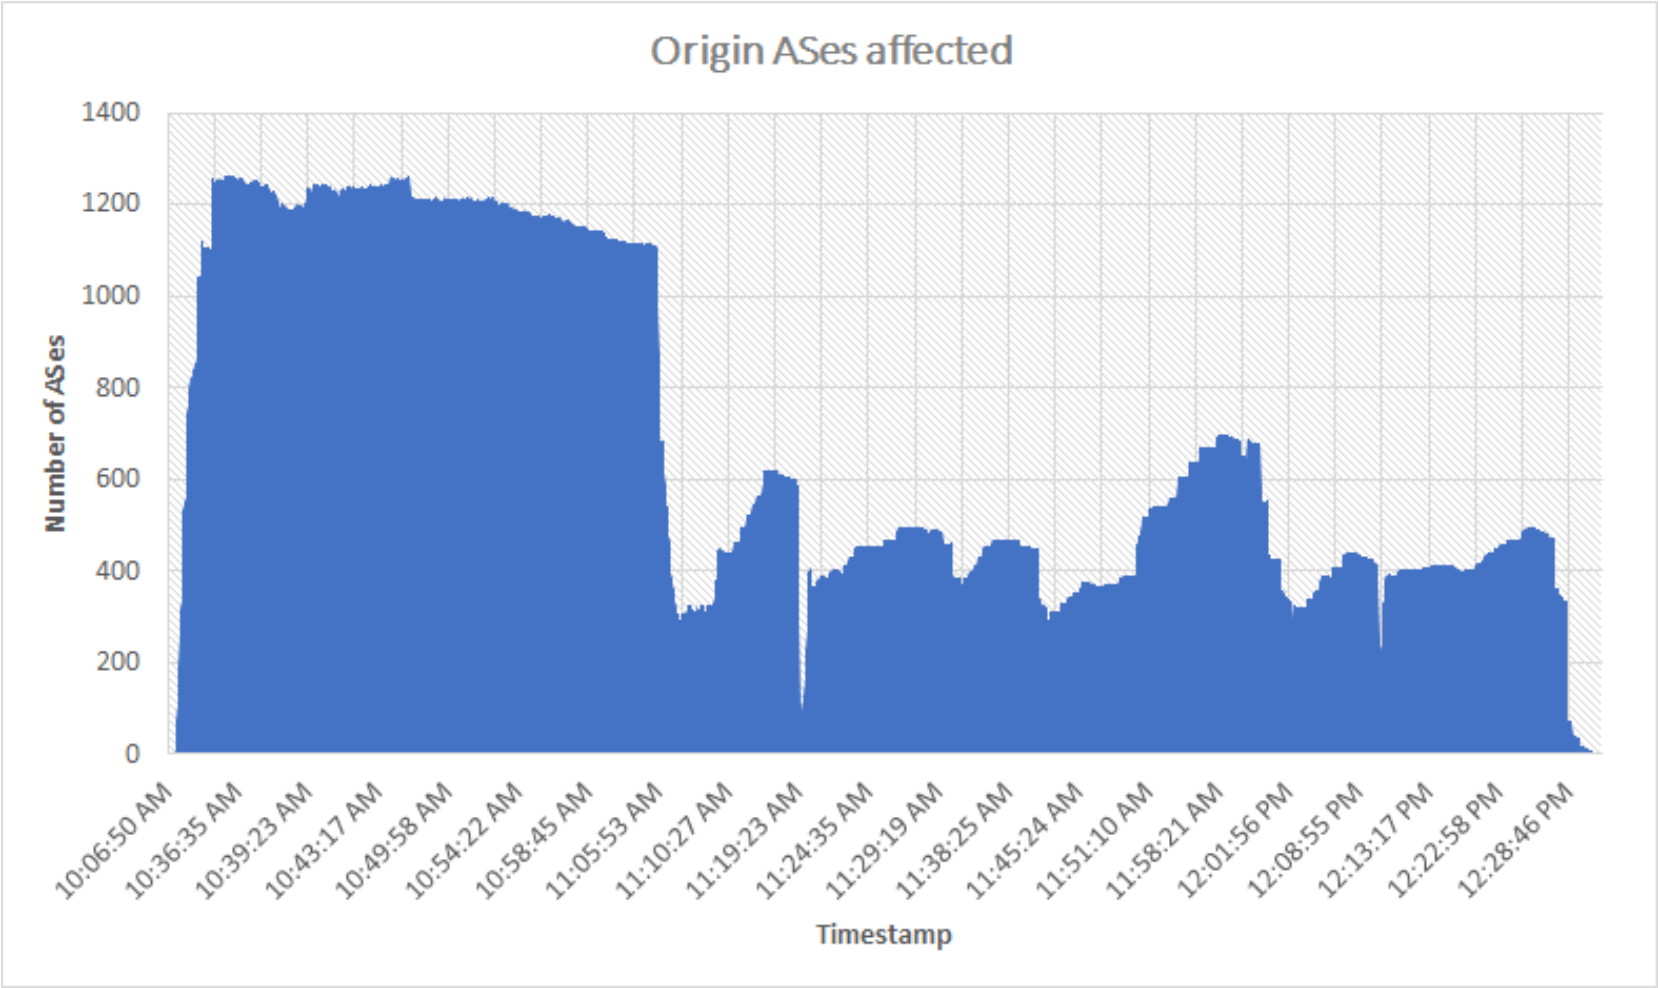

And many more...

Not only famous AS's have been affected!

Possible hijacks recorded in 2019: **911**

AS's causing hijacks: **452** AS's victims of hijacks: **630**

Route leaks recorded in 2019: **1282**

AS's causing a route leak: **294** AS's victims of route leaks: **883**

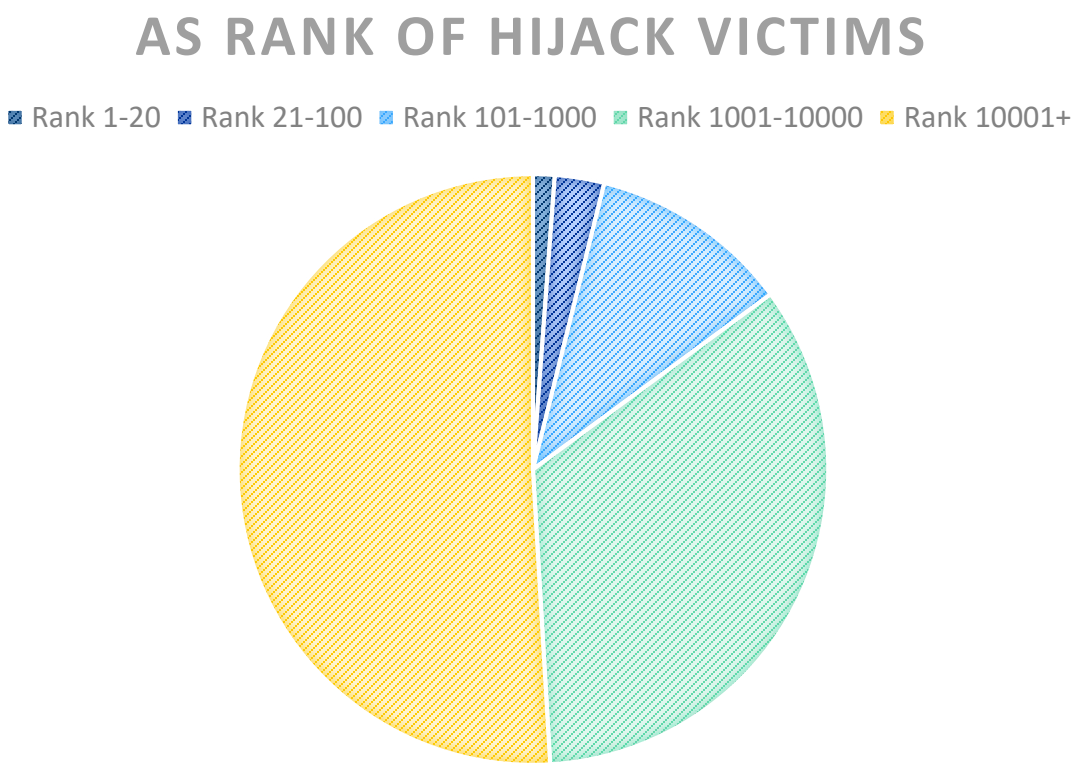

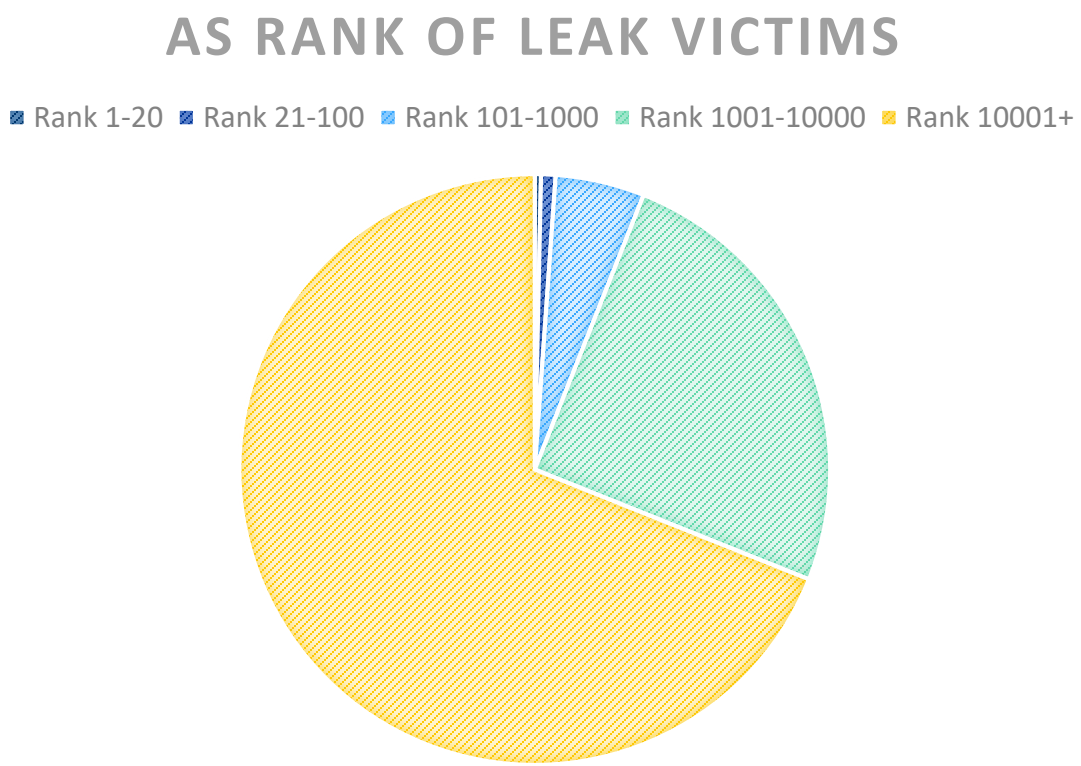

Data courtesy of https://bgpstream.com and https://asrank.caida.org/

What the eyes doesn't see the heart *may* grieve over!

- Leaks and hijacks could remain constrained to a routing region thanks to AS's dropping **RPKI invalid routes**
- This means that if the collectors are not in that routing region, you won't see that
- Assume that E is dropping RPKI invalid routes
- Then, if F starts a hijack/leak
  - The collector will not see it
  - A, B, C and D will
- Even you do not see it, it may still affect you!

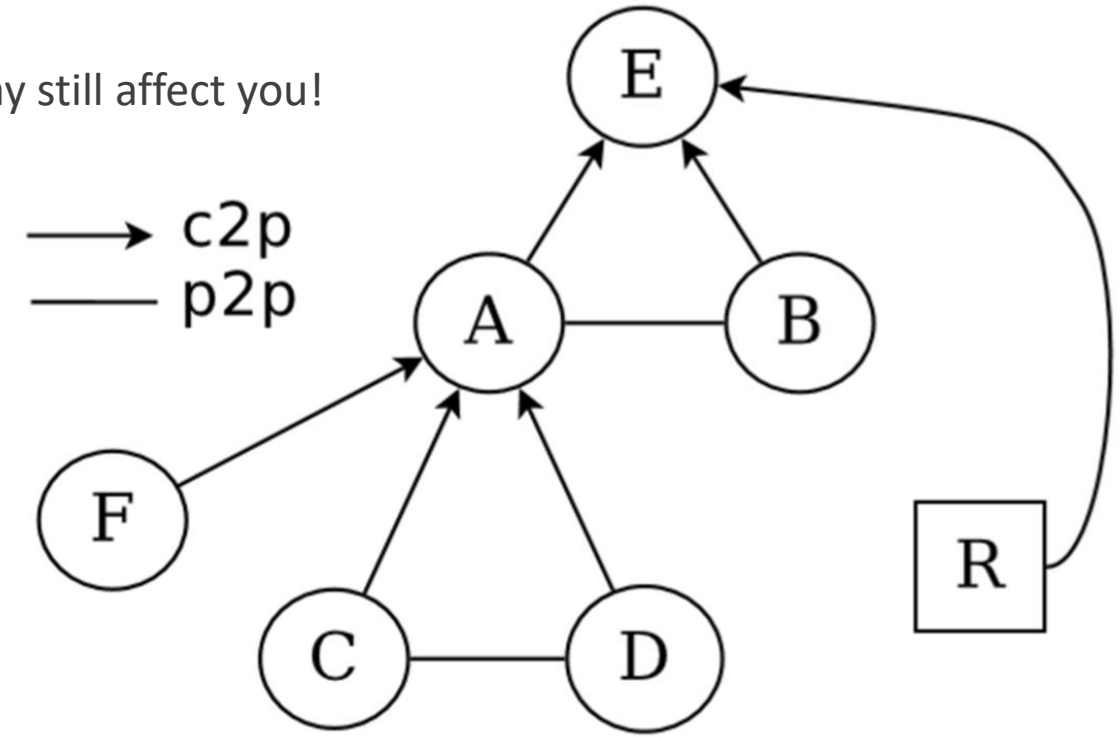

What are the consequences?

Several hijack attempts and route leaks have been seen from a few peers only
Several of them may also remain **unrevealed** from the collectors due to the low number of monitors!

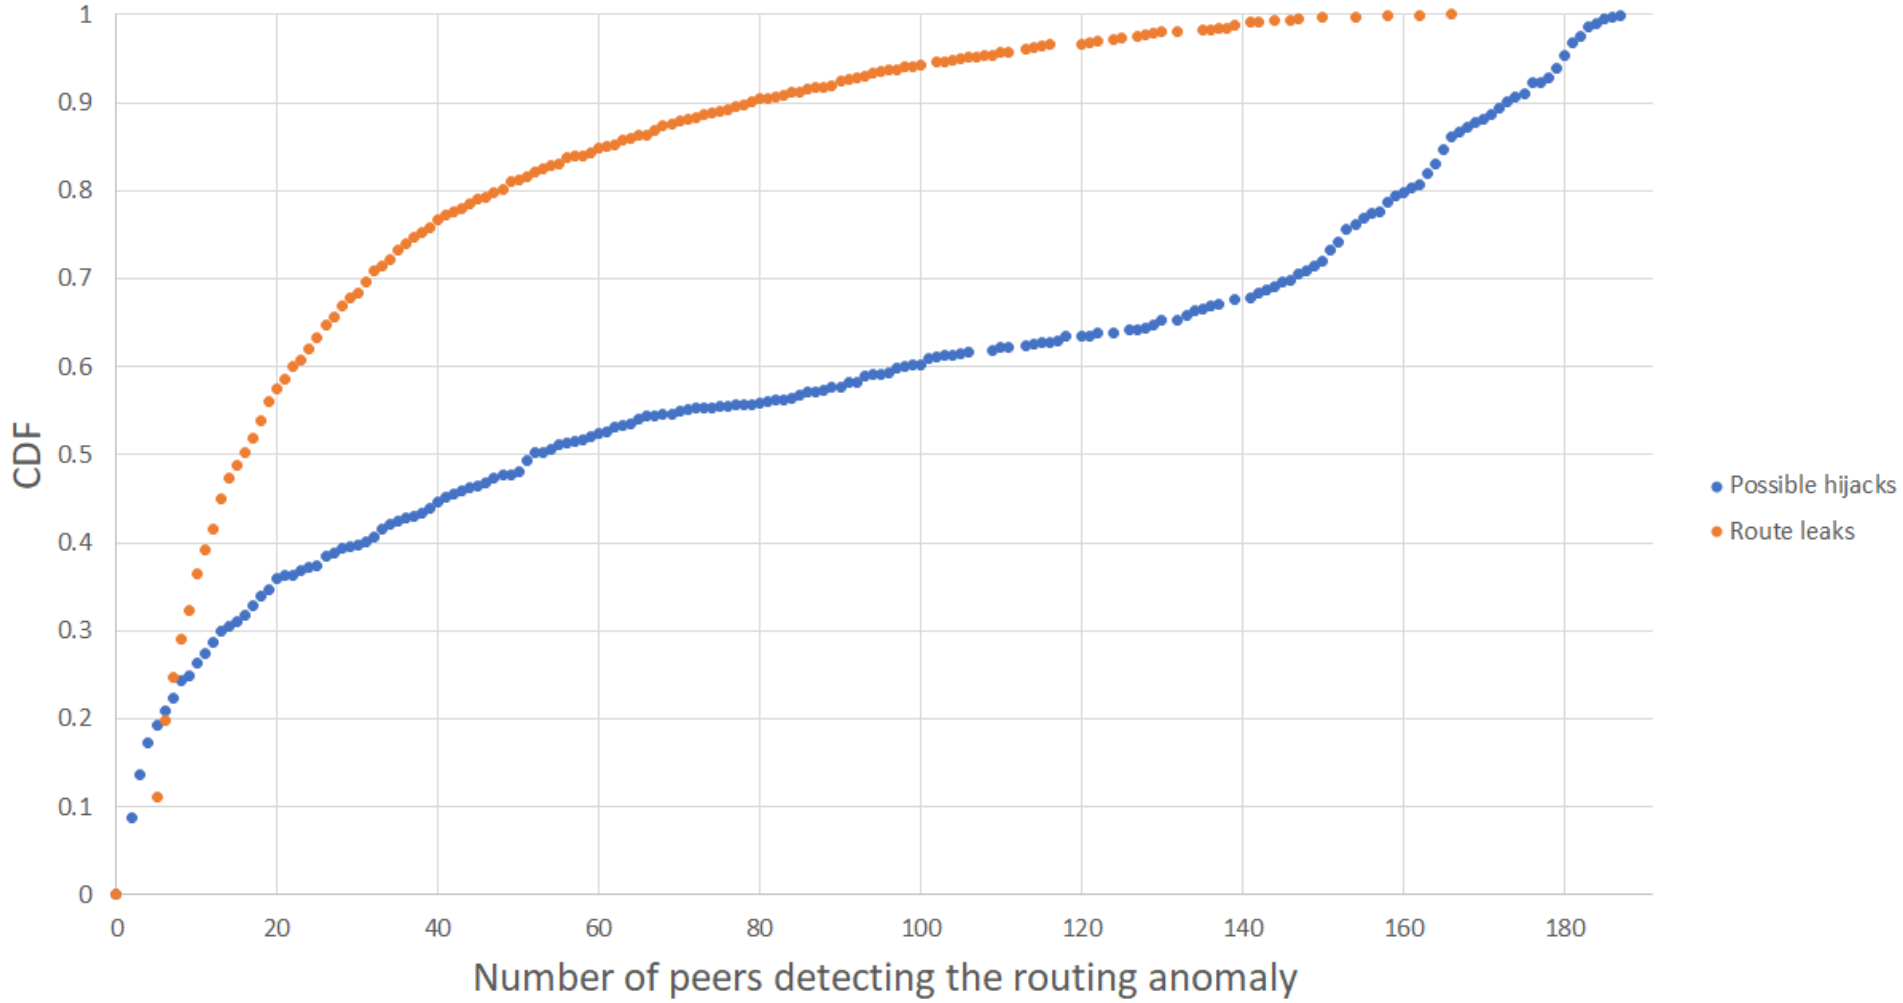

Source: https://bgpstream.com

What can we do about it?

RPKI – Resource Public Key Infrastructure

- RPKI allows AS administrator to create Route Origin Authorizations (ROAs)
  - ROAs are cryptographically signed objects

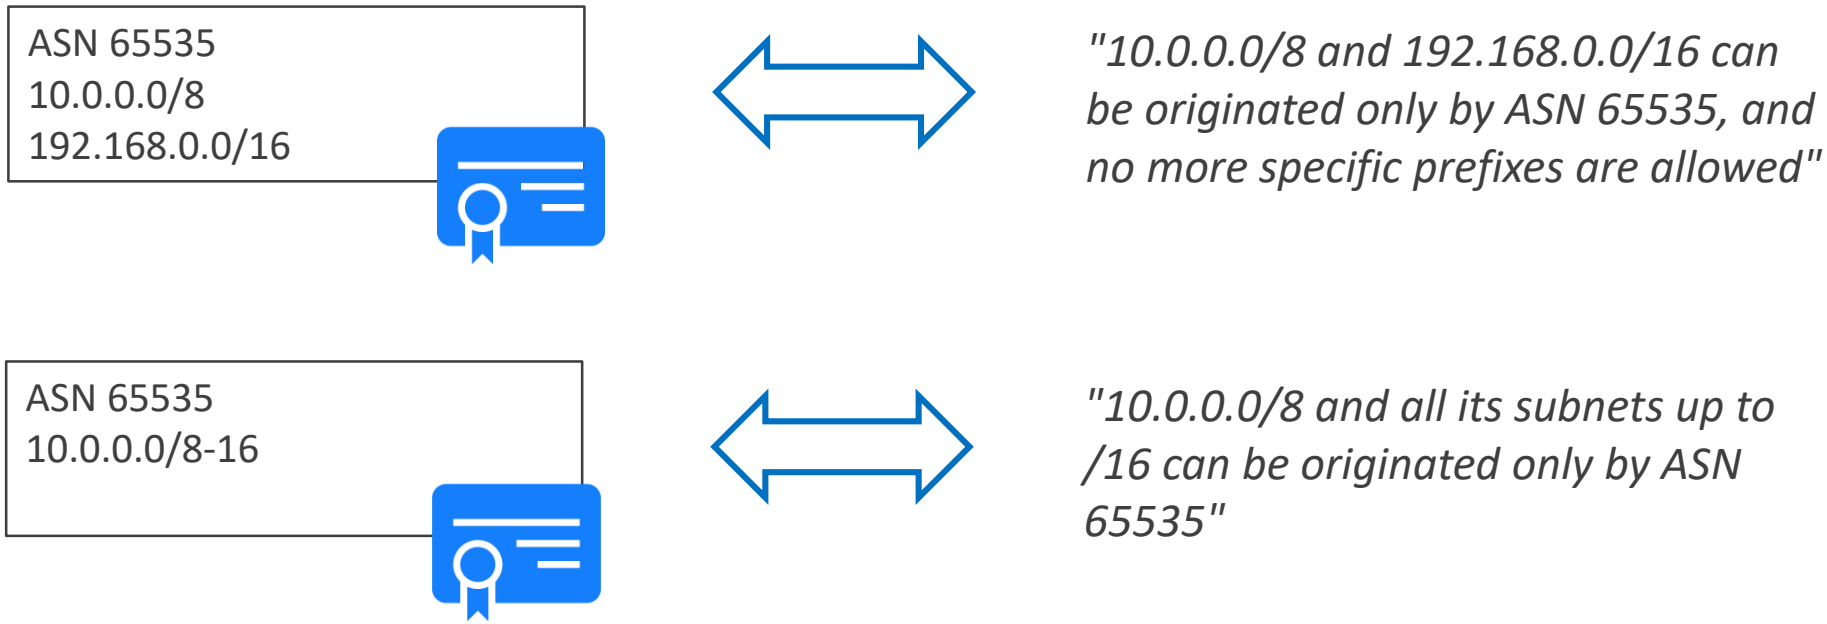

- A BGP router can check each announcement against the RPKI database and the result can be:
  - VALID
  - INVALID (could be dropped, e.g. NTT, AT&T and GTT)
  - NOT FOUND

… but is that enough?

- RPKI is a powerful mechanism to filter invalid announcements and **everyone should sign their prefixes**
- Unfortunately, it is not enough to detect and drop all the invalid announcements
  - BGP leaks (valid prefix origin but unexpected AS-PATH)
  - Intentional attacks (sub-prefix, AS_PATH forgery)
- Also, very little adoption up to date (about 15% of ASes signed at least a ROA)

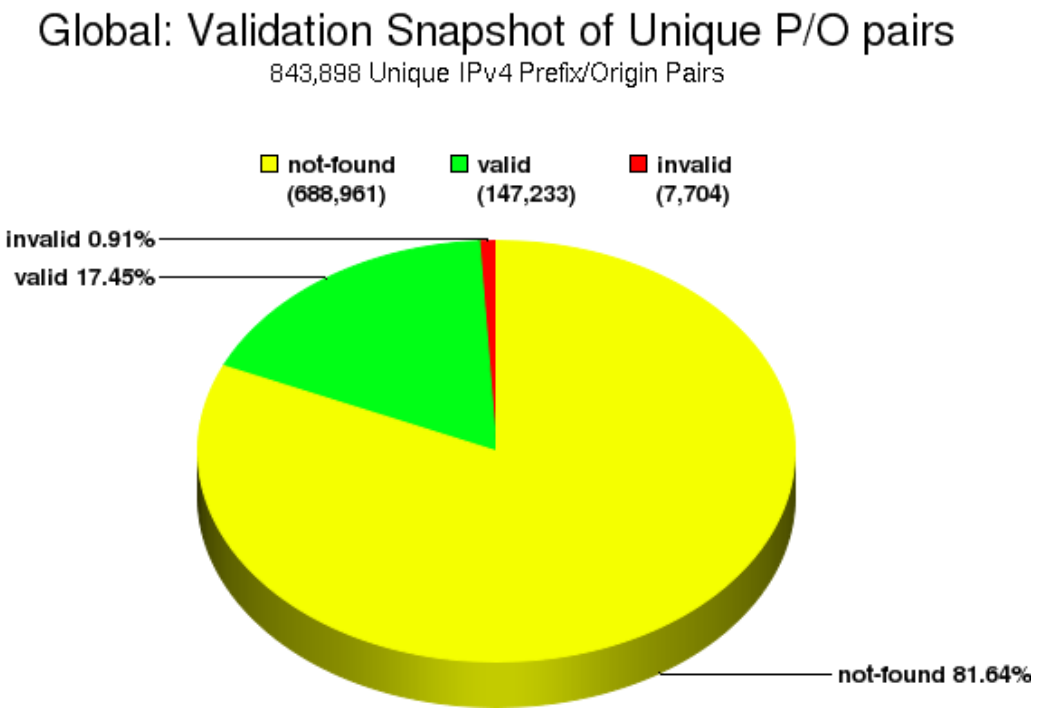

Source: https://rpki-monitor.antd.nist.gov/

Example – Route Leak

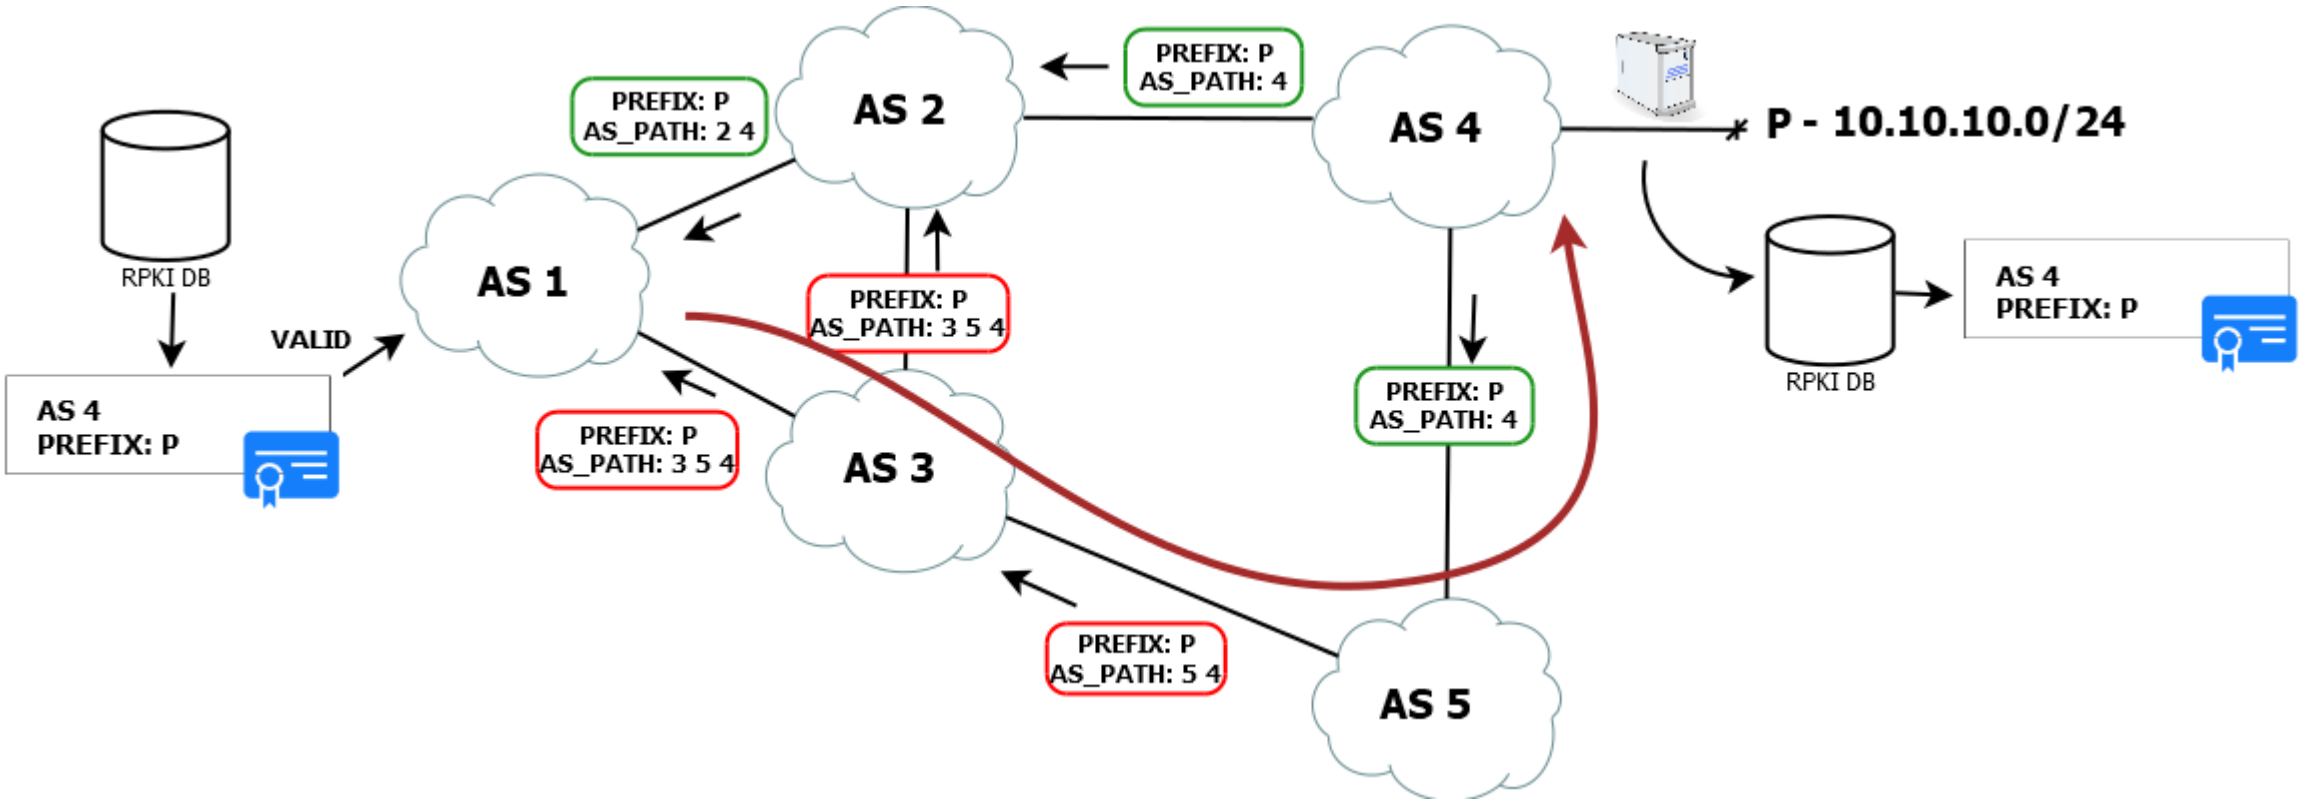

What about the future?

- Future is going to BGPSec, where BGP packets will be cryptographically signed
  - Main challenge: each router incurs a computational overhead due to digital signature/verify of each packet
  - Also, BGPSec will not be the solution to everything, for example BGP leaks

- IETF is discussing about how to detect invalid paths/route leaks
  - LDM [https://tools.ietf.org/html/draft-ietf-grow-route-leak-detection-mitigation-01]
  - ASPA [https://tools.ietf.org/html/draft-azimov-sidrops-aspa-verification-01]
  - Path RPKI: [https://tools.ietf.org/html/draft-van-beijnum-sidrops-pathrpki-00]
  - AS Cones: [https://tools.ietf.org/html/draft-ietf-grow-rpki-as-cones-01]

- Other prevention mechanisms are currently in place (e.g. peer-lock, IRR-based filtering, max-prefixes) but still they are not enough to impede the happening of those events

- In the meanwhile you can rely on BGP monitoring tools and platforms to react as soon as possible!

Questions?

lsani@catchpoint.com

catchpoint™