# Protection against and detection of some routing vulnerabilities

## A measurement approach

**Cristel Pelsser**

**TMA PhD School 2023**

# Highlight

- Intro to BGP and its vulnerabilities
- Some fixes to these vulnerabilities and their impact
  - RPKI time of flight
- Attacks are still possible
- Getting the best of BGP data
  - Most valuable set of Vantage Points (MVP)
- Detecting BGP hijacks
  - Detection of type-1 BGP hijacks (DFOH)

# Focus on the inter-domain routing protocol
# BGP

# The Internet is composed of Autonomous Systems (AS): one or more networks under the control of a single entity.

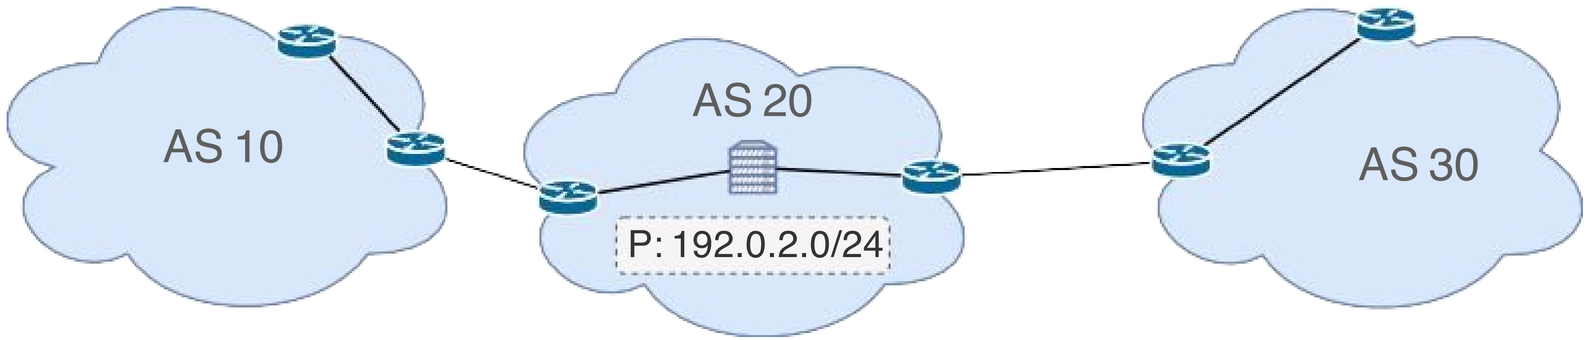

# The Internet is composed of Autonomous Systems (AS): one or more networks under the control of a single entity.

Prefixes of the AS are advertised to the outside using BGP.

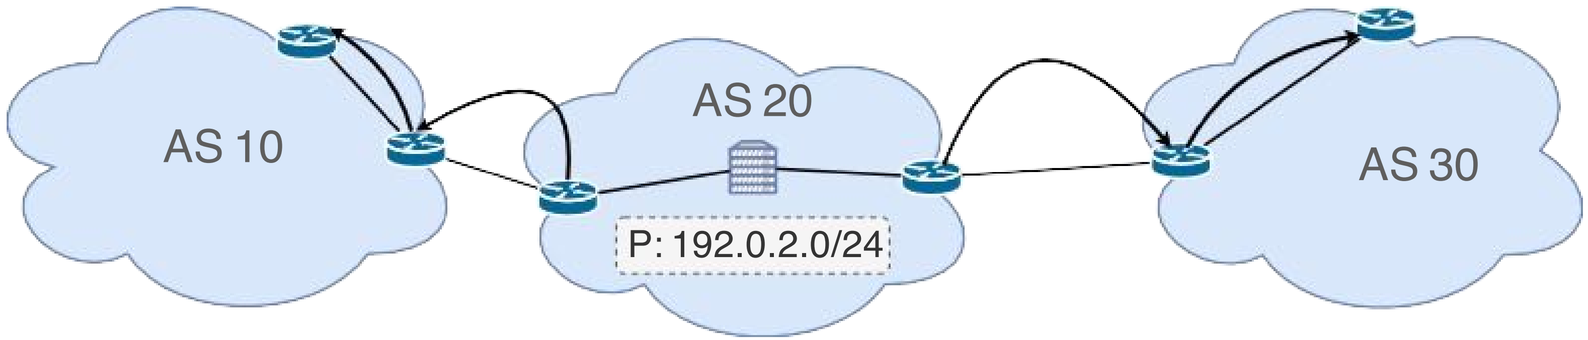

# The Internet is composed of **Autonomous Systems** (AS): one or more networks under the control of a single entity.

Prefixes of the AS are advertised to the outside using BGP.

Traffic flows in the reverse direction.

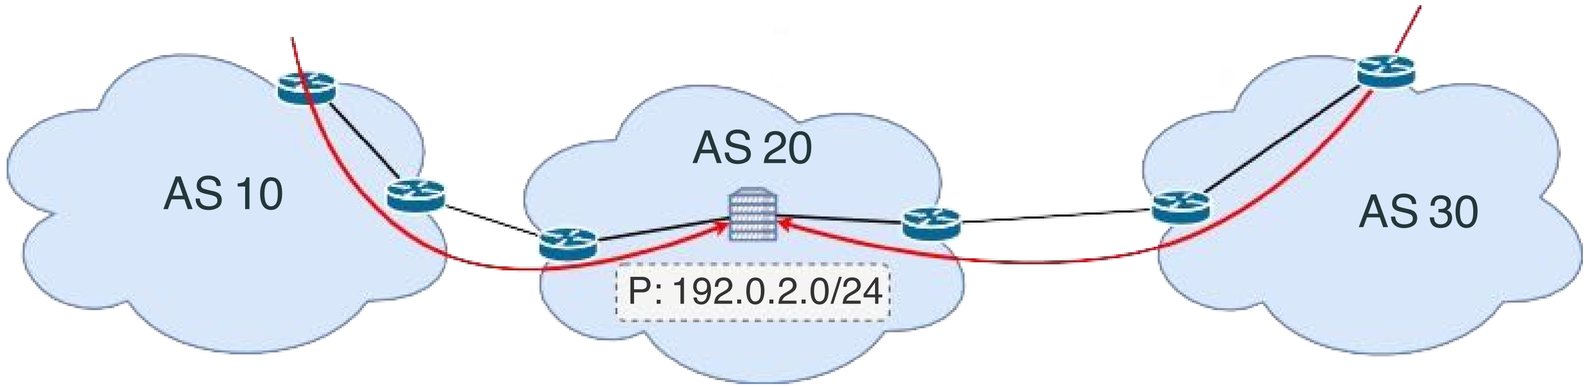

# The Internet is a complex ecosystem

There are 73,803 AS advertised as of May 21, 2023.

https://www.potaroo.net/tools/asn32/

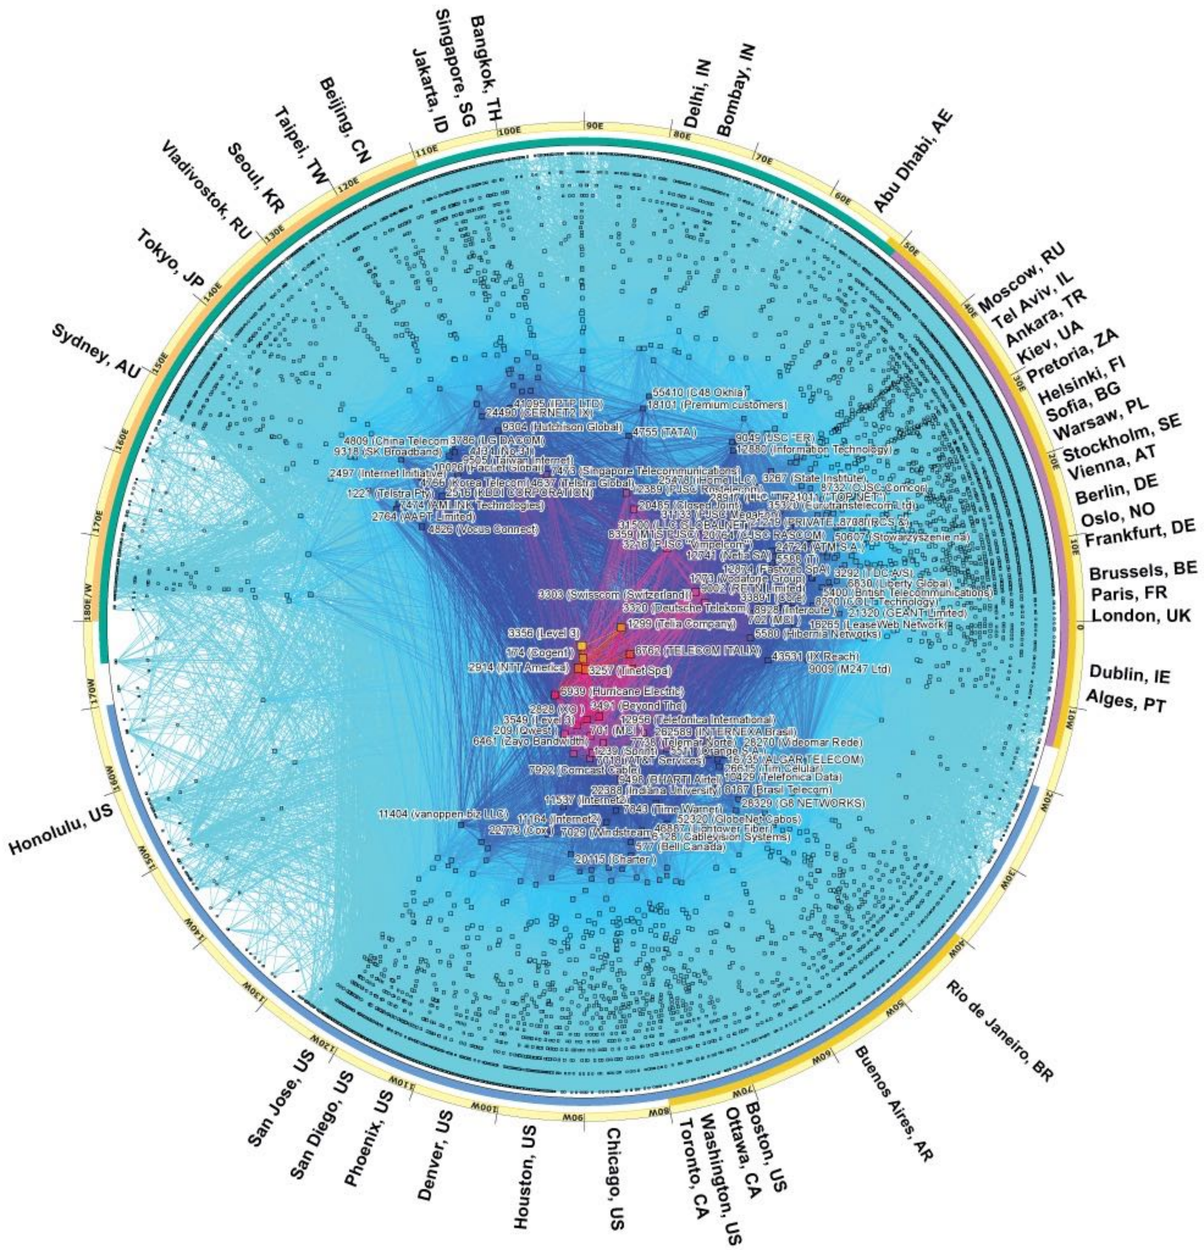


Source: https://www.caida.org/projects/cartography/as-core/2017/

# There is little to no security in the routing protocol used in the Internet

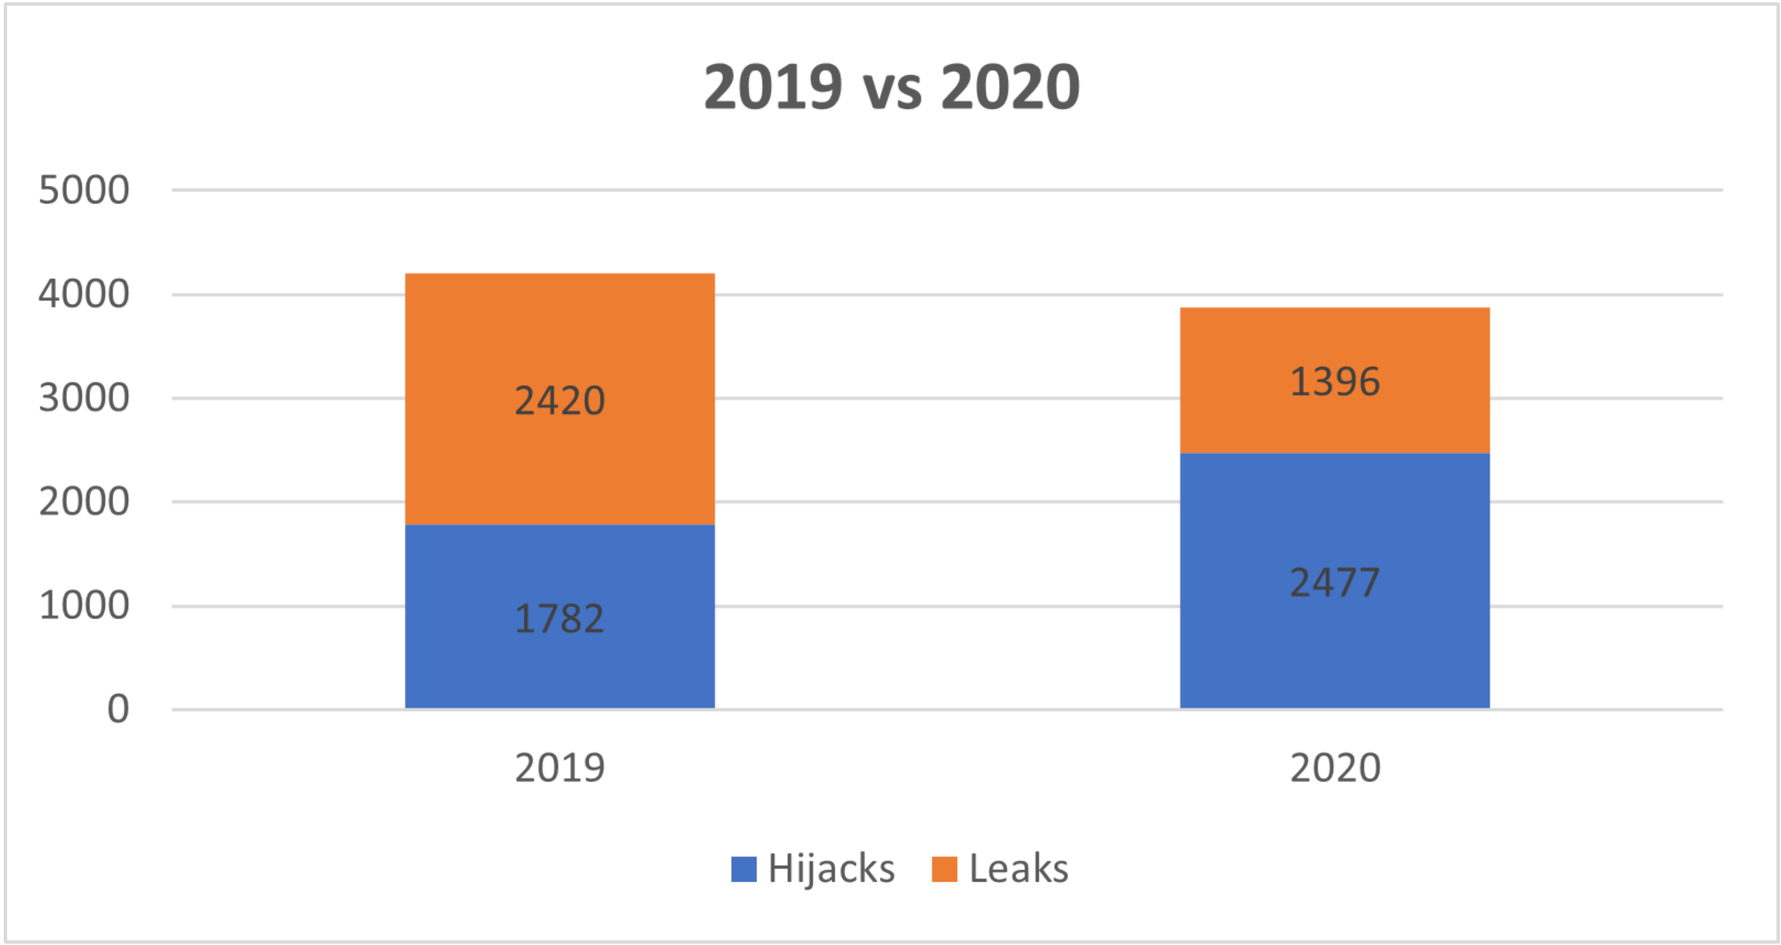


Source: https://www.manrs.org/2021/02/bgp-rpki-and-manrs-2020-in-review/

# Some vulnerabilities of BGP

Prefix hijacks

Blackjack attacks

BGP lies

BGP session injection

# Hijacks can be used to divert traffic and gain inside knowledge

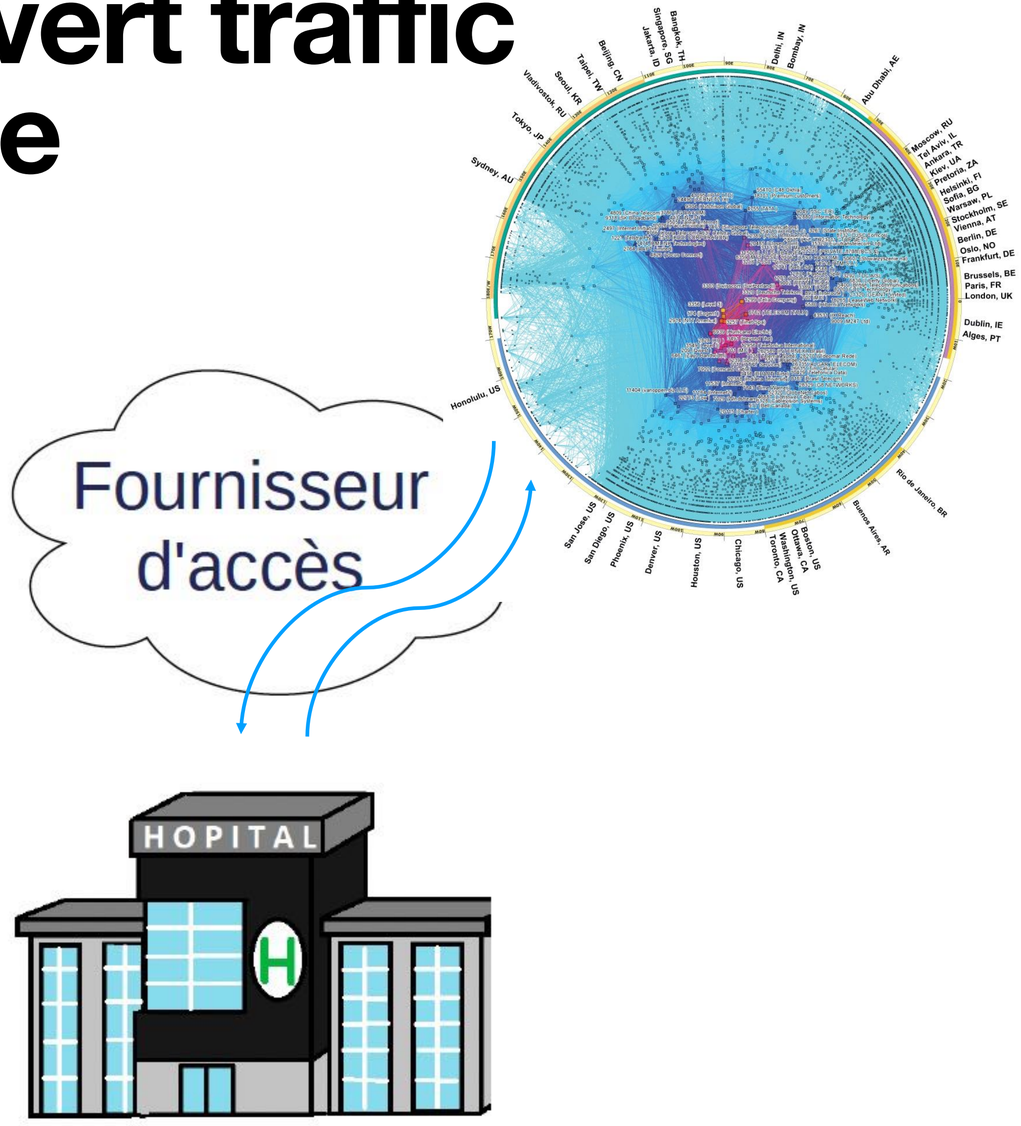

# Hijacks can be used to divert traffic and gain inside knowledge

BGP hijacking

Fournisseur d'accès


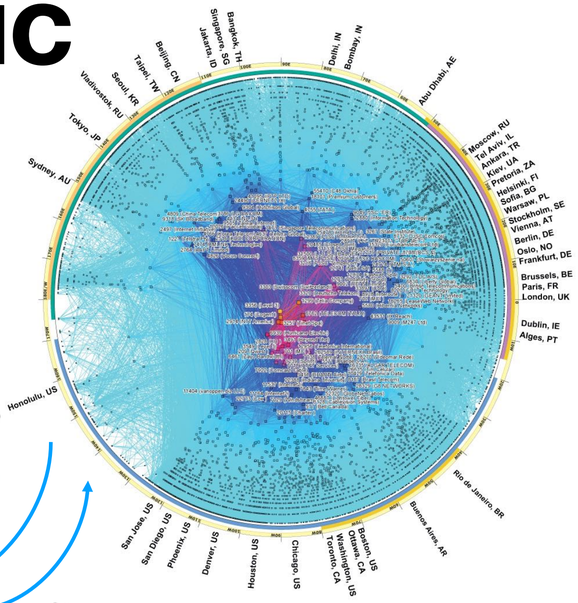

# Hijacks can be used to divert traffic and gain inside knowledge

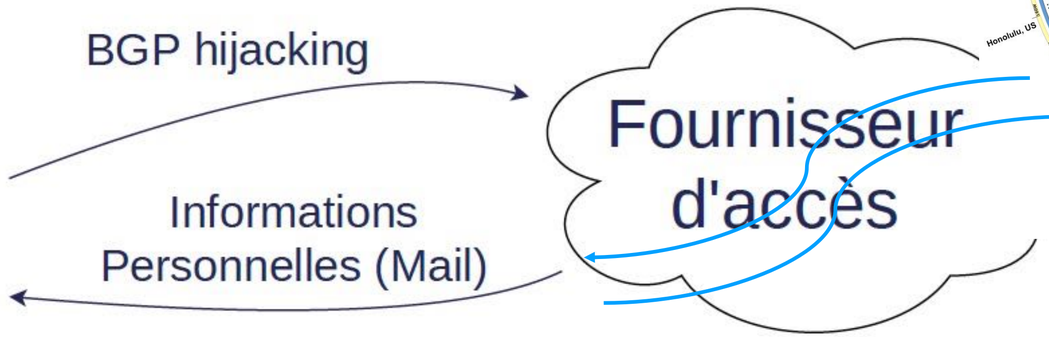


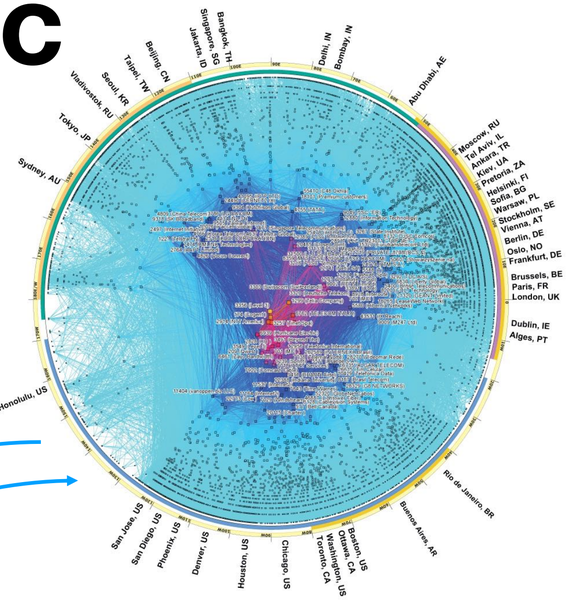


# Hijacks can be used to divert traffic and gain inside knowledge

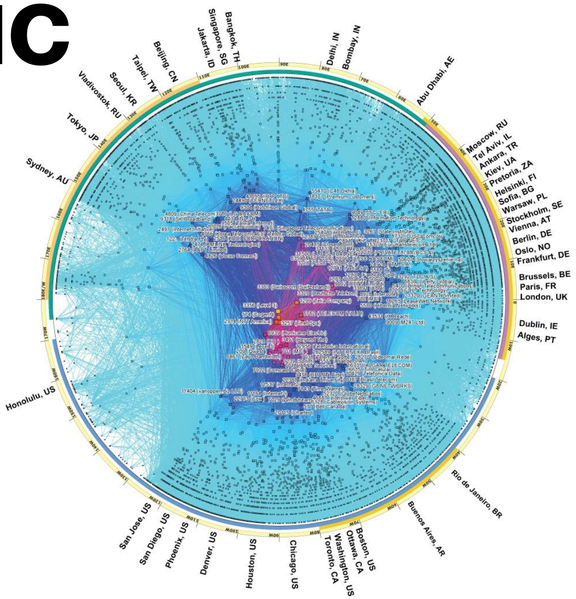


BGP hijacking

Informations Personnelles (Mail)

Phishing + Ransomware

Fournisseur d'accès

# Hijacks can be used to divert traffic and gain inside knowledge

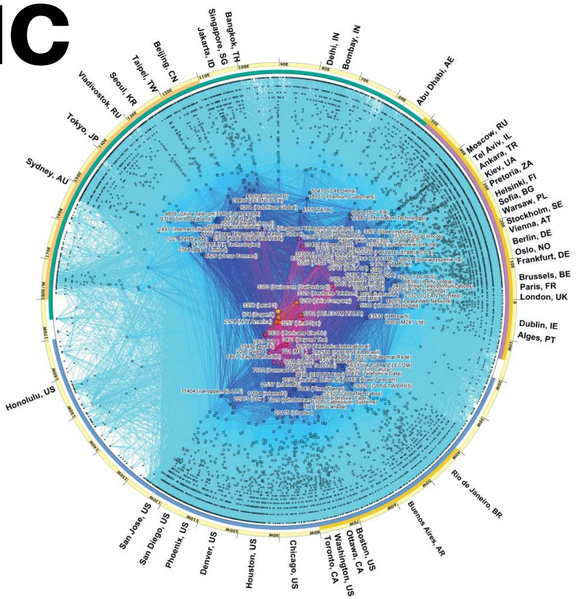


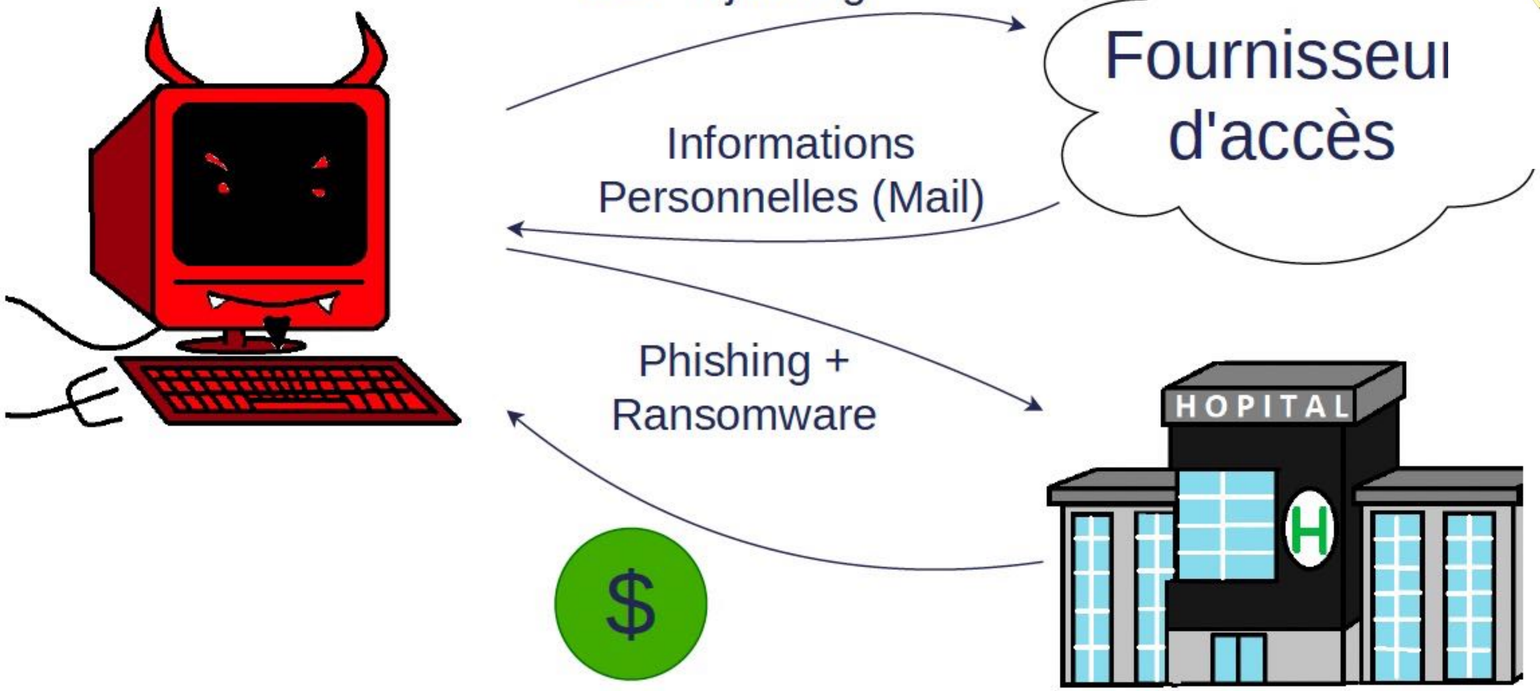

# Multiple causes for hijacks

Hijacks are not always malicious

They can be the result of misconfigurations


https://www.manrs.org/2022/06/configuration-issue-penalizing-single-digit-asns/?utm_source=rss&utm_medium=rss&utm_campaign=configuration-issue-penalizing-single-digit-asns

# Extract from the blog post:

“In recent years, we’ve noticed that single-digit ASNs (ASN1 through ASN9) often appear to be route hijackers. Is this true? We dug into the data and ultimately realized **no, single-digit ASNs are not hijacking address space at an alarming rate**. What’s happening is the result of a misconfiguration issue because of the “AS path prepend” command on Mikrotik routers.”

https://www.manrs.org/2022/06/configuration-issue-penalizing-single-digit-asns/?utm_source=rss&utm_medium=rss&utm_campaign=configuration-issue-penalizing-single-digit-asns

# Some vulnerabilities of BGP

Prefix hijacks

Blackjack attacks

BGP lies

BGP session injection

# A Backjack attack surfs on the blackholing mechanism provided to protect against DDoS

# DDoS are frequent

For examples Cloudflare reports that the number of DDoS quadrupled compared to pre-covid levels

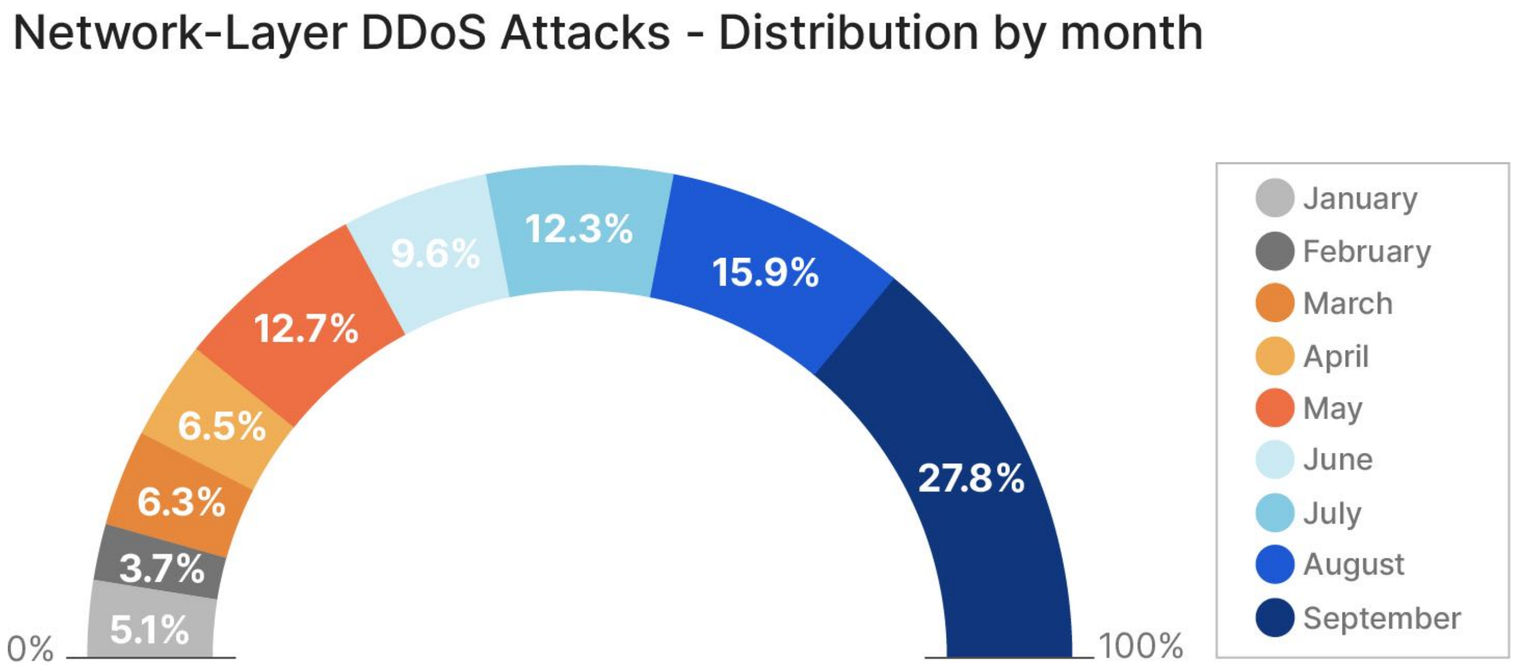


Source: https://blog.cloudflare.com/network-layer-ddos-attack-trends-for-q3-2020/

CLOUDFLARE

# DDoS

In a denial of service attack, the infractucture may be congested.

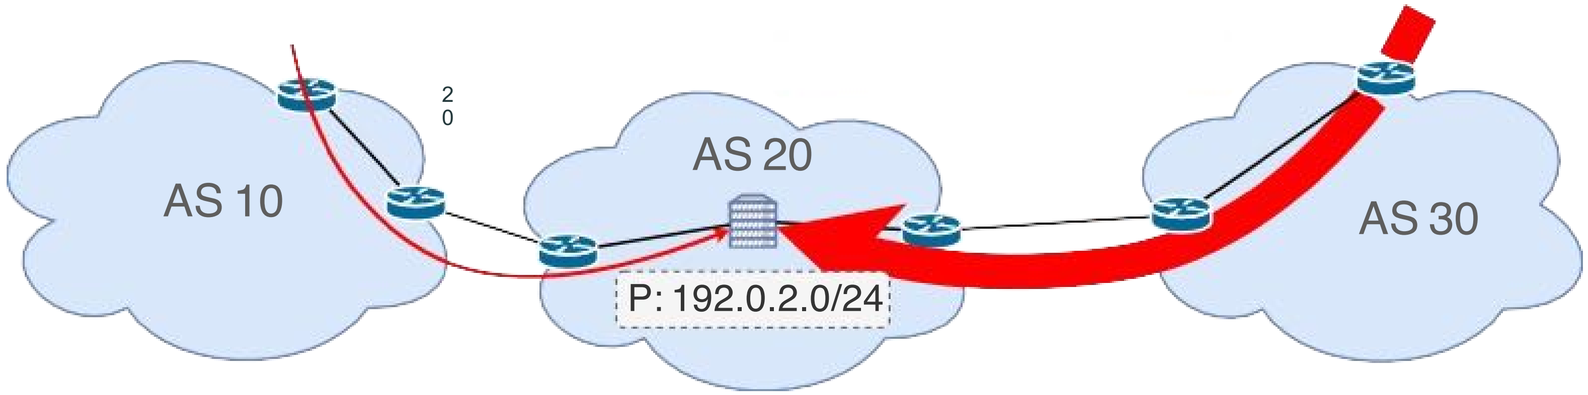

# BGP blackholing

**Blackholing** is a **DDoS mitigation** technique signaled via **BGP** using a community.

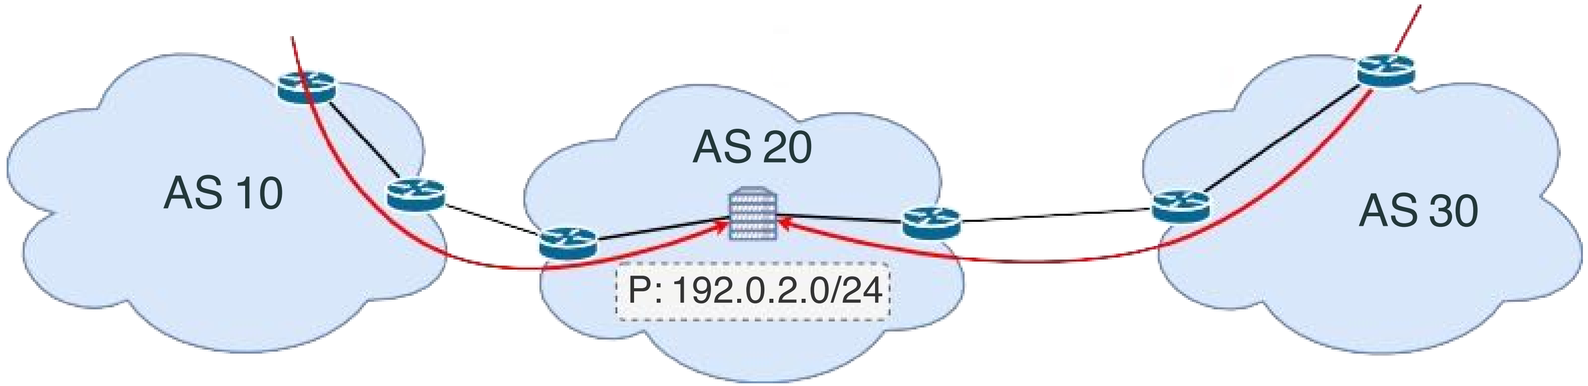

# BGP blackholing

**Blackholing** is a **DDoS mitigation** technique signaled via **BGP** using a community.

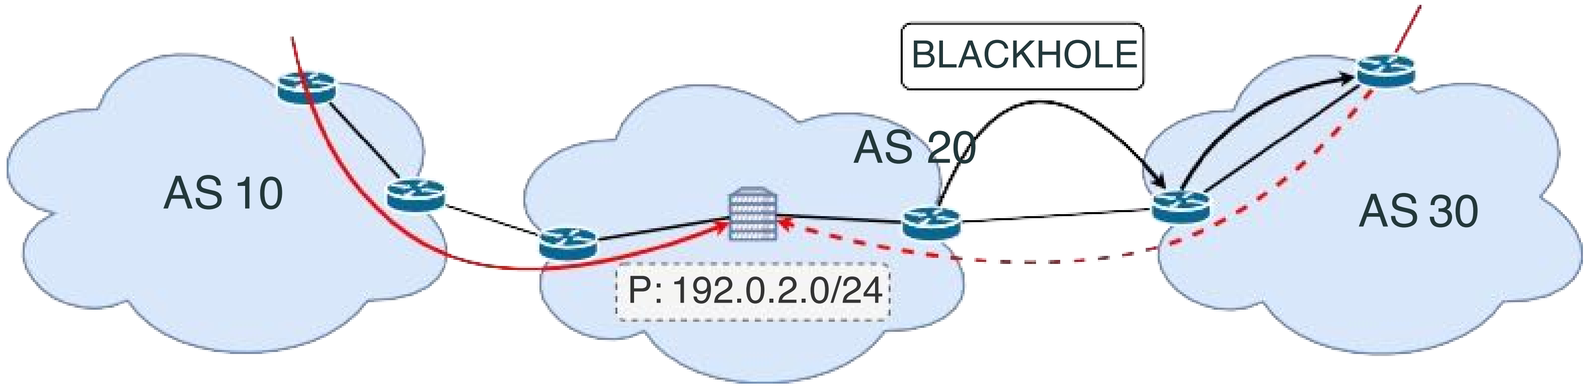


Blackholing has a double-edged sword effect: **all** traffic is dropped.

# BGP community usage is increasing

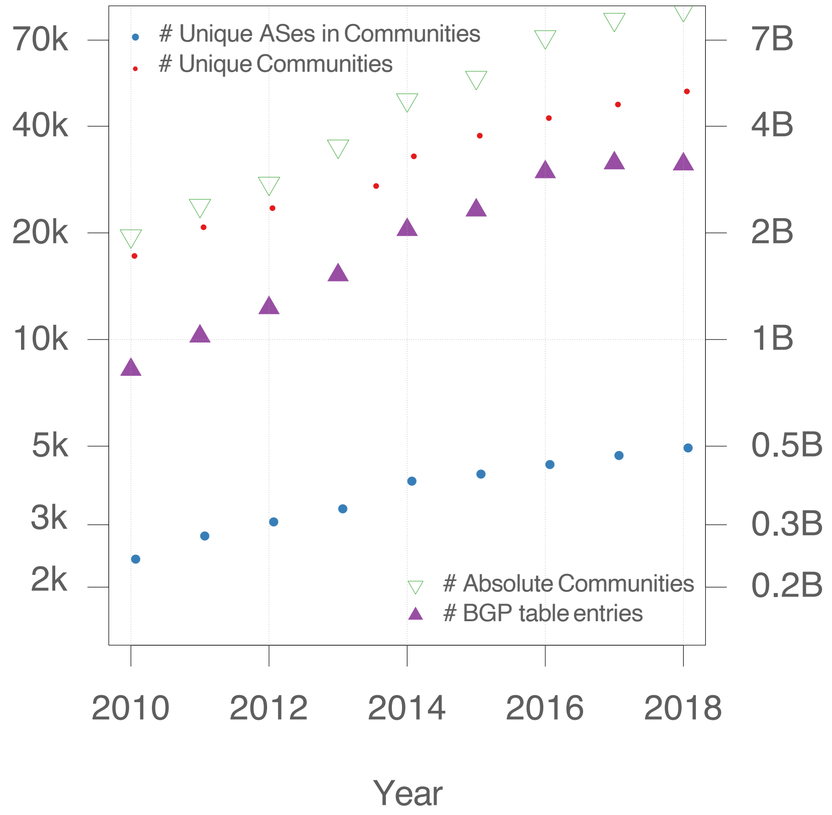


**Increasing usage warrants a closer look.**

# BGP community usage is increasing

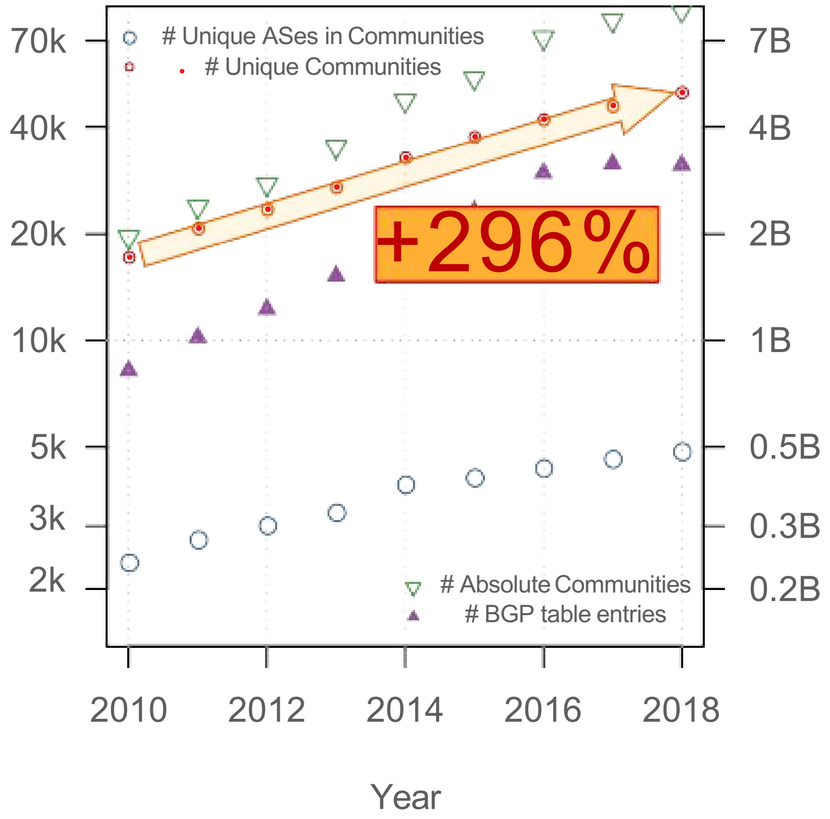


**Increasing usage warrants a closer look.**

# BGP Communities (RFC 1997)

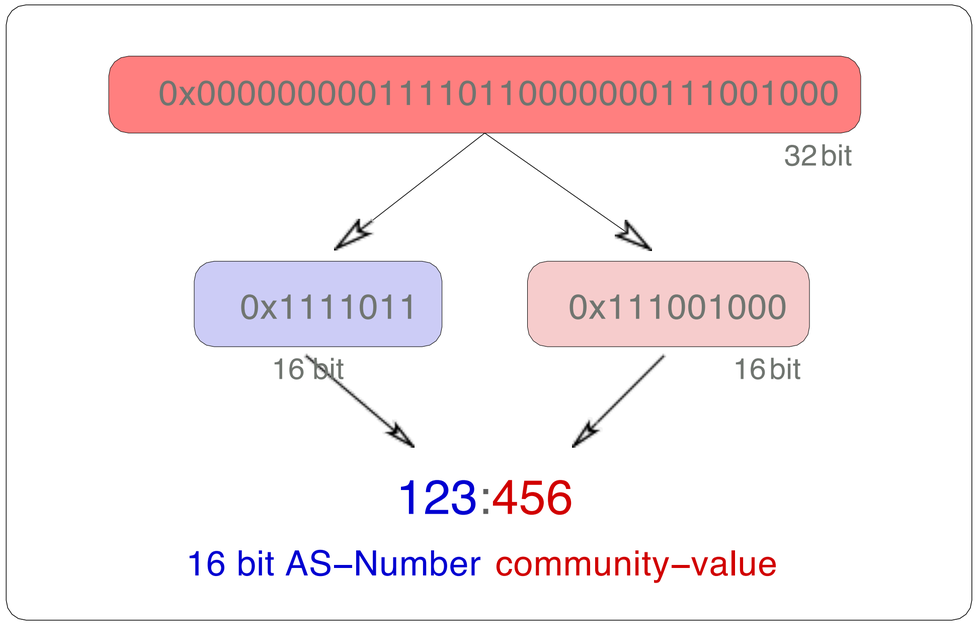


By convention written *ASN:VALUE*

ASN can be both sender or intended 'recipient'

It's up to the peers to agree upon 'values' used

Every network decides on the semantics of values

# BGP Communities: Usage (examples)

| Informational Communities (Passive Semantics) | Action Communities (Active Semantics) |
| --- | --- |
| Location tagging | Remote triggered blackholing |
| RTT tagging | Path prepending |
| | Local pref/MED |
| | Selective announcements |

**Without documentation, you can not tell if a community is active or passive!**

**Blackhole community value is :666** (RFC 7999)

Given the **increasing popularity** of BGP communities
and the ability to **trigger actions** as well as **relay information**,
the first question that comes to the mind of an
Internet measurement researcher is. . .

**What could possibly go wrong?**

Can blackholing be used with malicious intent?

Are there different types of attacks?

Are there any existing and relevant security mechanisms?

Are these mechanisms sufficient?

# Example topology

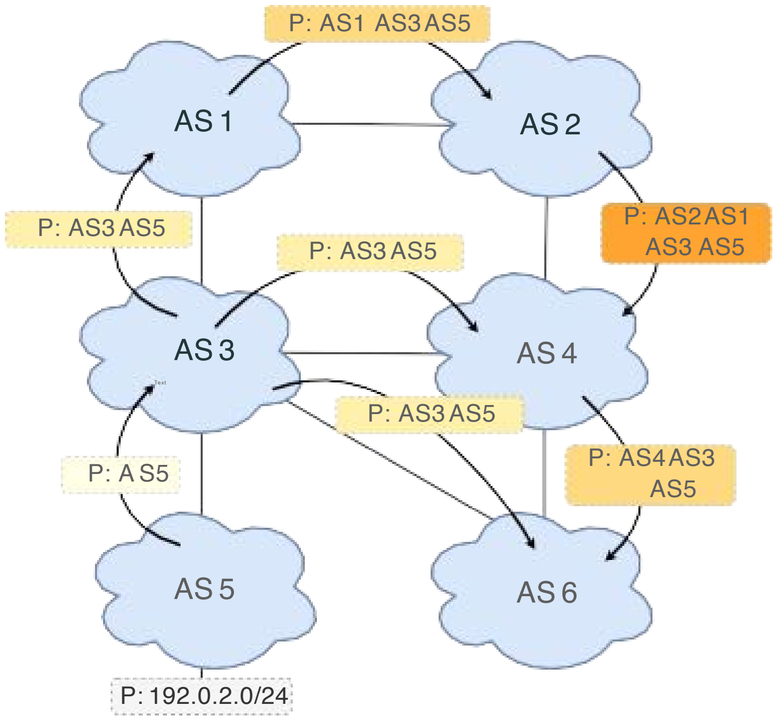


**BGP update propagation**

# Example topology

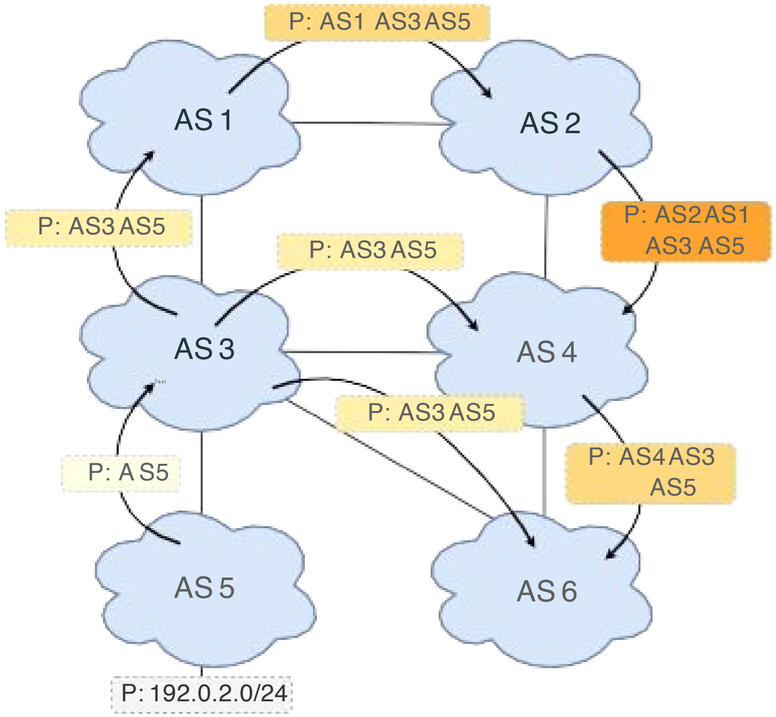


**BGP update propagation**

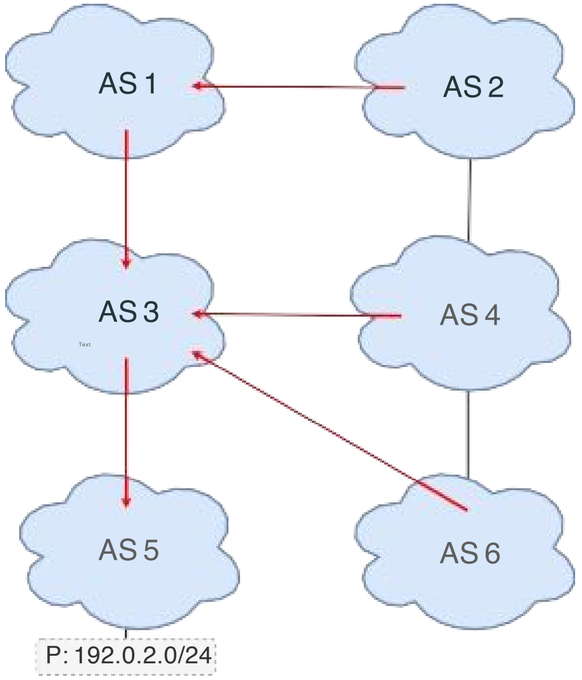


**Traffic flow**

# Example topology

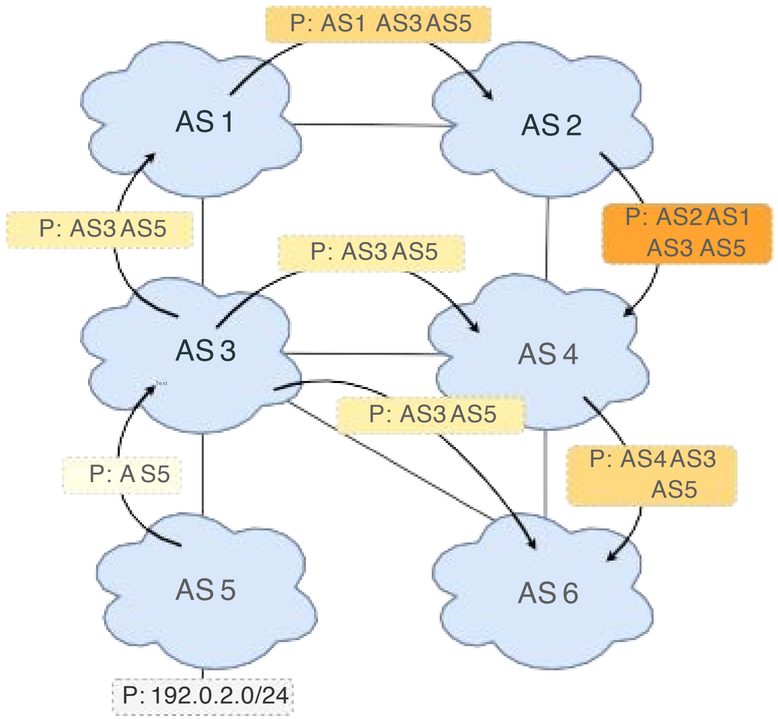


**BGP update propagation**

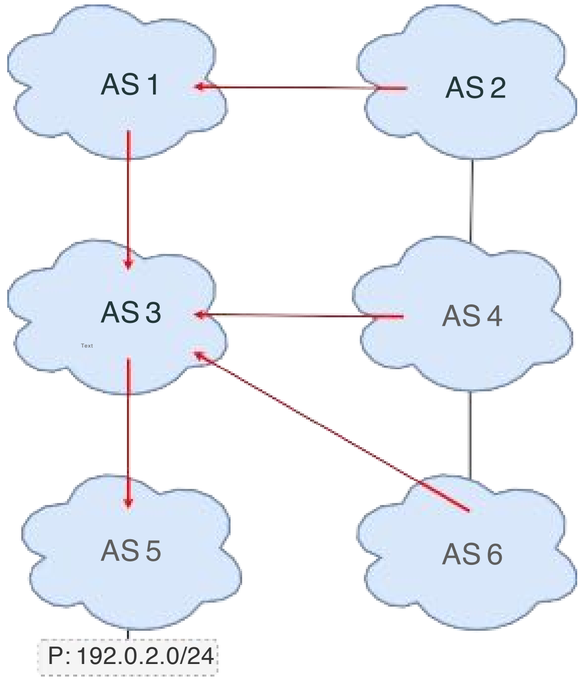


**Traffic flow**

BGP policies make AS2 not learn the path via AS4

# Example topology

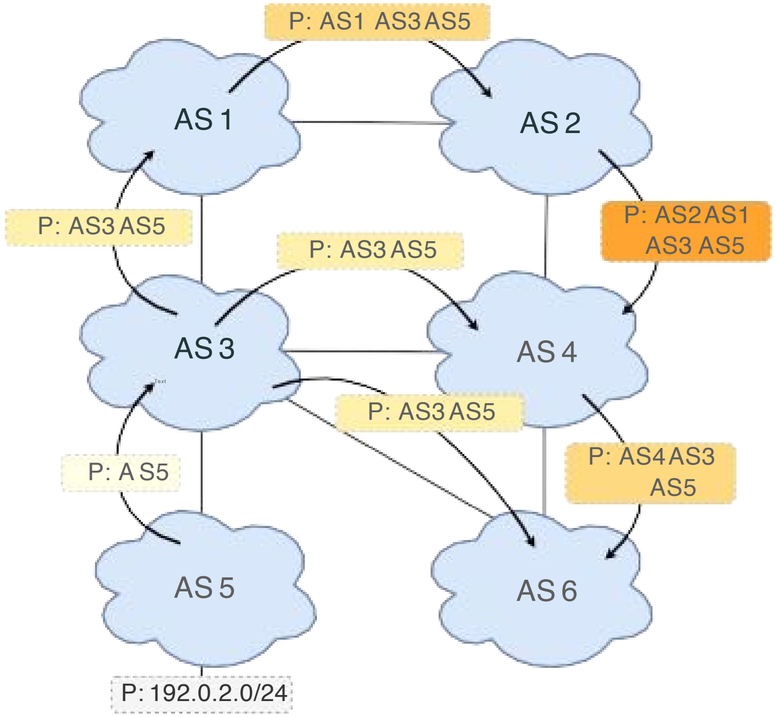


**BGP update propagation**

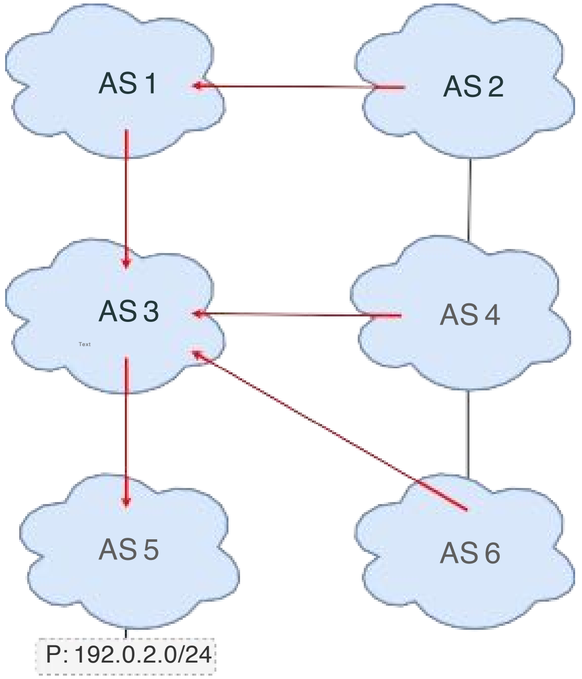


**Traffic flow**

BGP policies make AS2 not learn the path via AS4

BGP policies are distributed in the AS using BGP communities

# Example topology

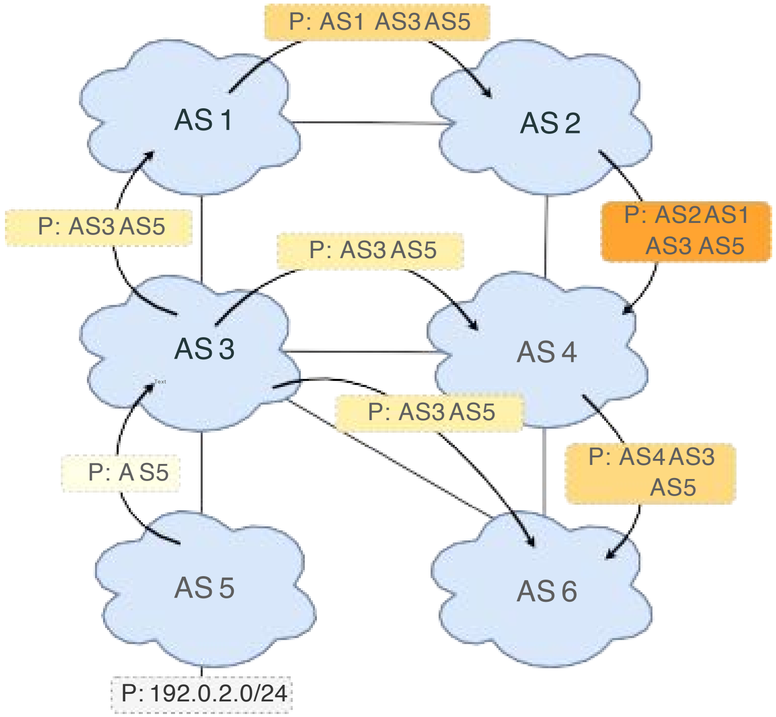


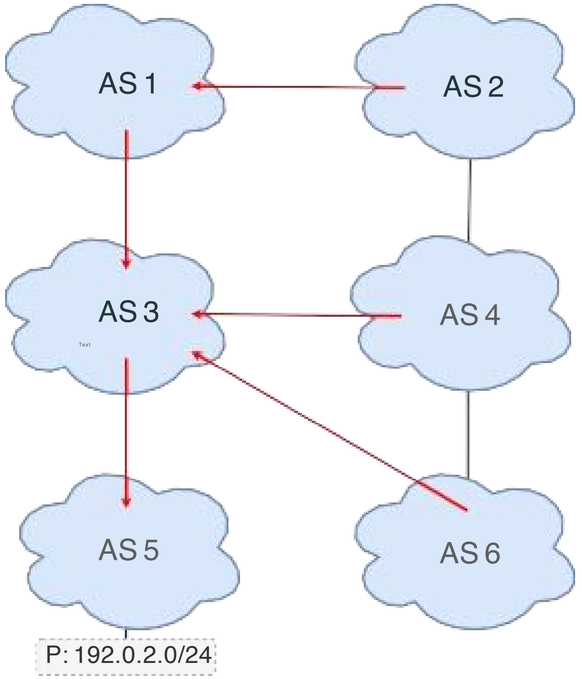


**BGP update propagation**

**Traffic flow**

BGP policies make AS2 not learn the path via AS4

BGP policies are distributed in the AS using BGP communities

**In the next slides AS6 is the attacker**

# Hijack-0 and Blackjack-0

## Sermpezis 2018 (Artemis)

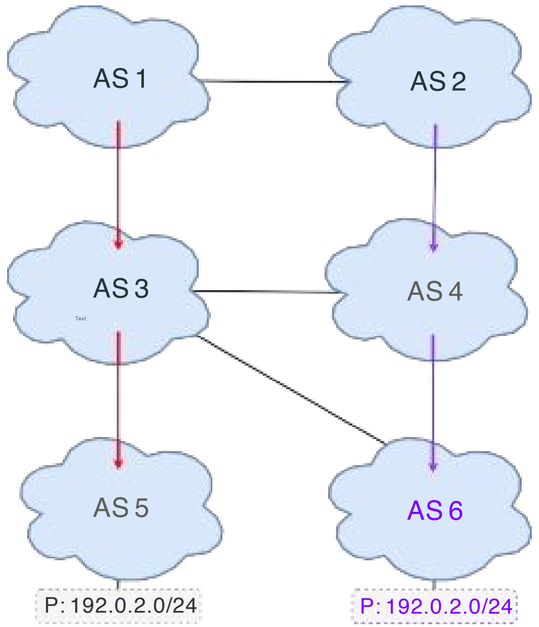


**Hijack type-0**

AS2 and AS4 traffic is de-routed to AS6 because the advertised path is shorter.

# Hijack-0 and Blackjack-0

## Sempezis 2018 (Artemis)

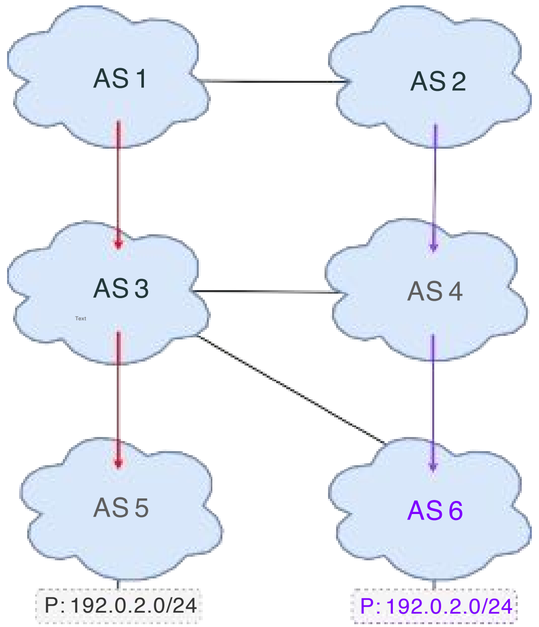


**Hijack type-0**

AS2 and AS4 traffic is de-routed to AS6 because the advertised path is shorter.

## Miller et Pelsser 2019

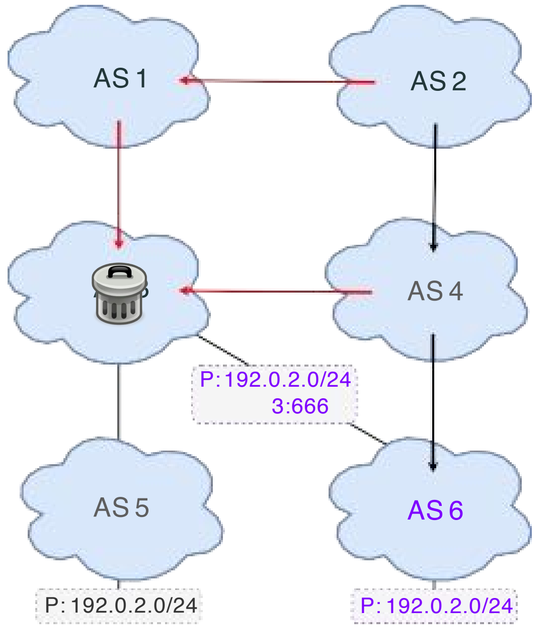


**Blackjack type-0**

All traffic to *P* is blackholed at AS3.

**Hijacking + blackholing**

# Best practices for legitimate blackholing empower blackjacks

**Best Practices for blackholing**[4]

Give a **higher priority** to blackholing.

Do **not propagate** the advertisement across AS borders.

[4]Cisco, **Remotely Triggered Black Hole Filtering - Destination Based and Source Based**.

# Best practices for legitimate blackholing empower blackjacks

**Best Practices for blackholing**[4]

Give a **higher priority** to blackholing.

Do **not propagate** the advertisement across AS borders.

**Consequences**

**Reach**: Precedence over AS path length. Even ASes far away are vulnerable.

**Stealth**: The attacker is not dropping traffic himself.

[4]Cisco, **Remotely Triggered Black Hole Filtering - Destination Based and Source Based**.

# Best practices for legitimate blackholing empower blackjacks

**ROA** Route Origin Authorizations are digitally signed objects attesting that a given AS is **authorized to originate** routes for a set of prefixes.

**ROV** With Route Origin validation, an AS **validates the origin** of the BGP updates with regard to the content of the RPKI Objects.

**But other attacks are possible.**

# BGP Blackjacks - Type-N

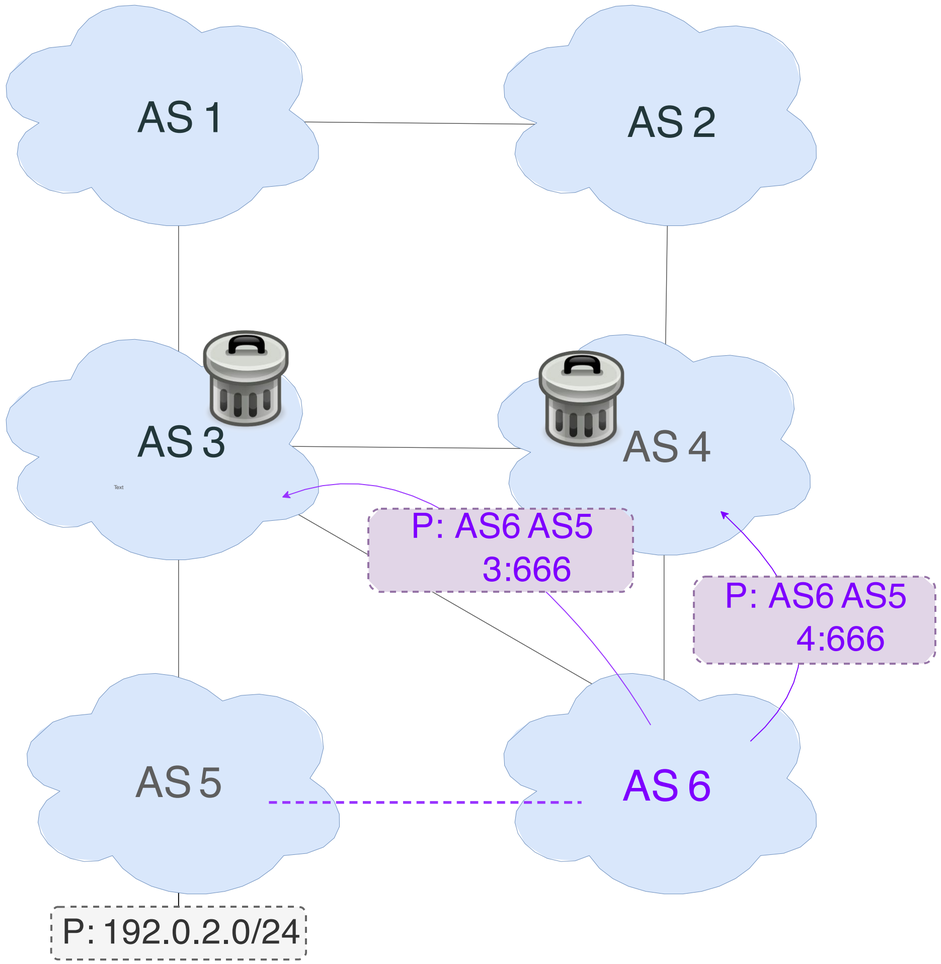


The origin AS is legit. The AS-path is not.

# BGPsec[5]

BGPsec allows ASes to **sign** advertisements.

This guarantees the AS path reflects the **actual path** the advertisement went through.

**But on-paths attacks are still possible.**

[5]Lepinski and Sriram, **BGPsec Protocol Specification**.

# Related publications

**Taxonomy of Attacks using BGP Blackholing.**

Loic Miller (U. Strasbourg), Cristel Pelsser (U. Strasbourg). ESORICS 2019.

**BGP Communities: Even more Worms in the Routing Can.**

Florian Streibelt (MPI[1]), Franziska Lichtblau (MPI), Robert Beverly (NPS[2]), Anja Feldmann (MPI), Cristel Pelsser (U. Strasbourg), Georgios Smaragdakis (TU Berlin), Randy Bush (IIJ[3]). ACM IMC 2018.

[1]Max Planck Institute for Informatics
[2]Naval Postgraduate School
[3]Internet Initiative Japan

# Some vulnerabilities of BGP

Prefix hijacks

Blackjack attacks

BGP lies

BGP session injection

# An ISP (AS B) announces a path in BGP but forwards packets along a different path

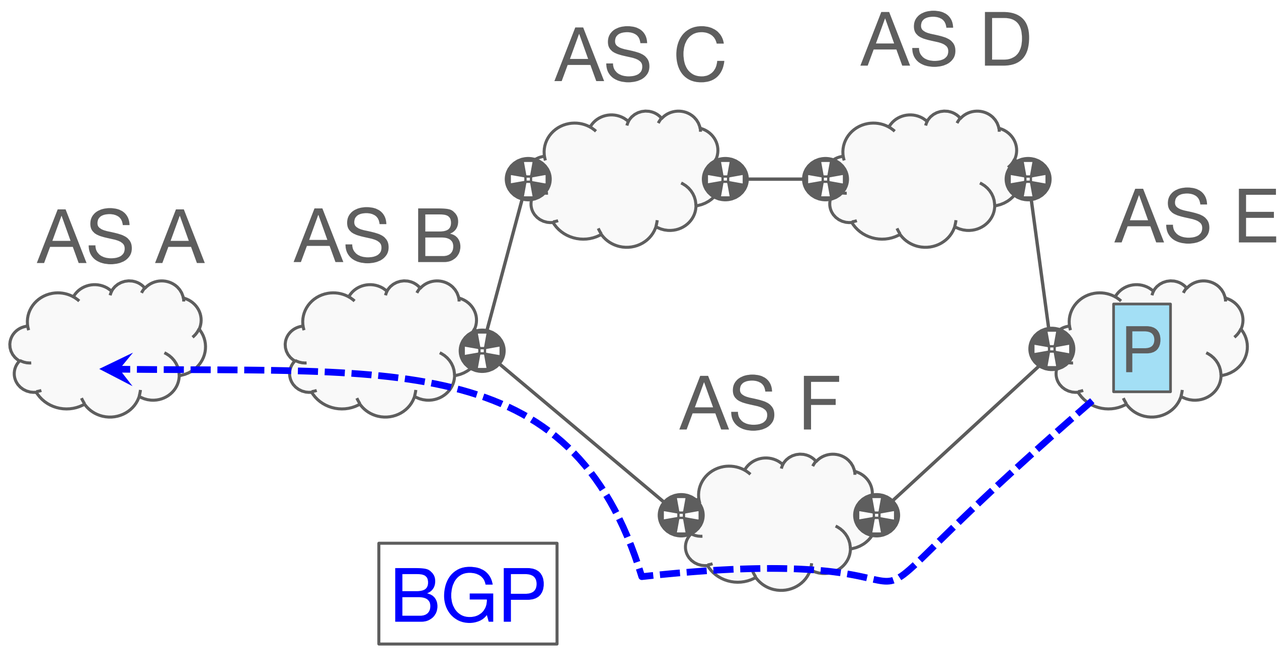

# An ISP (AS B) announces a path in BGP but forwards packets along a different path

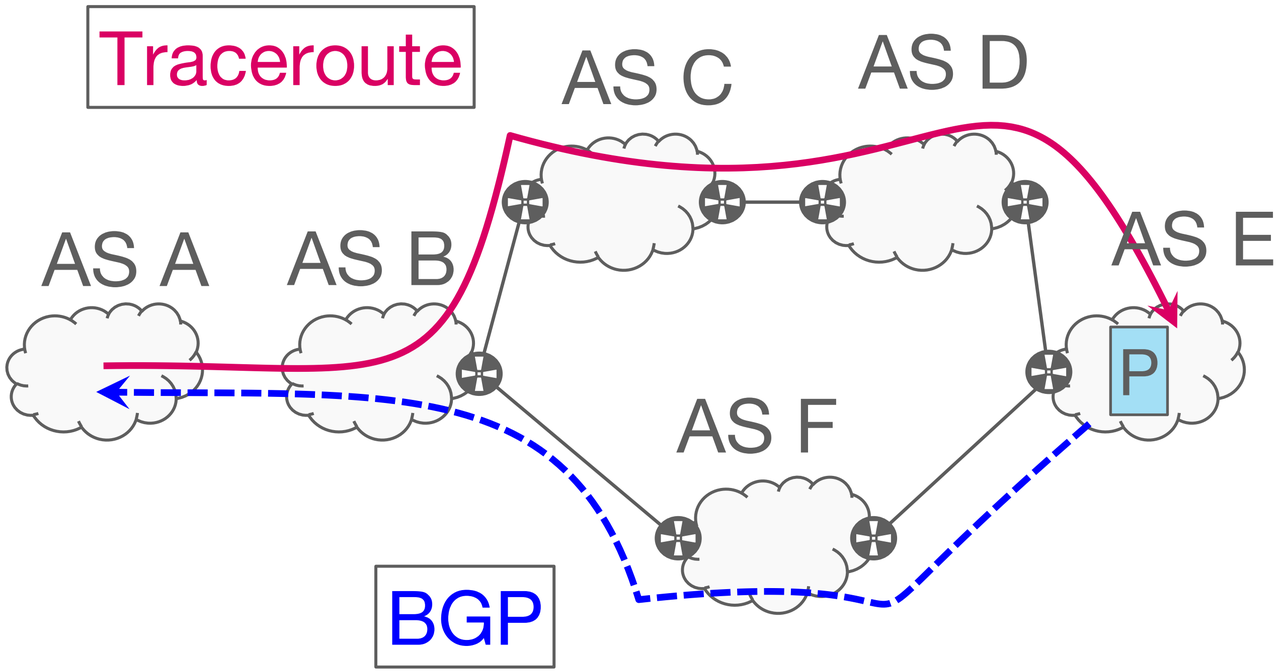

# An ISP (AS B) announces a path in BGP but forwards packets along a different path

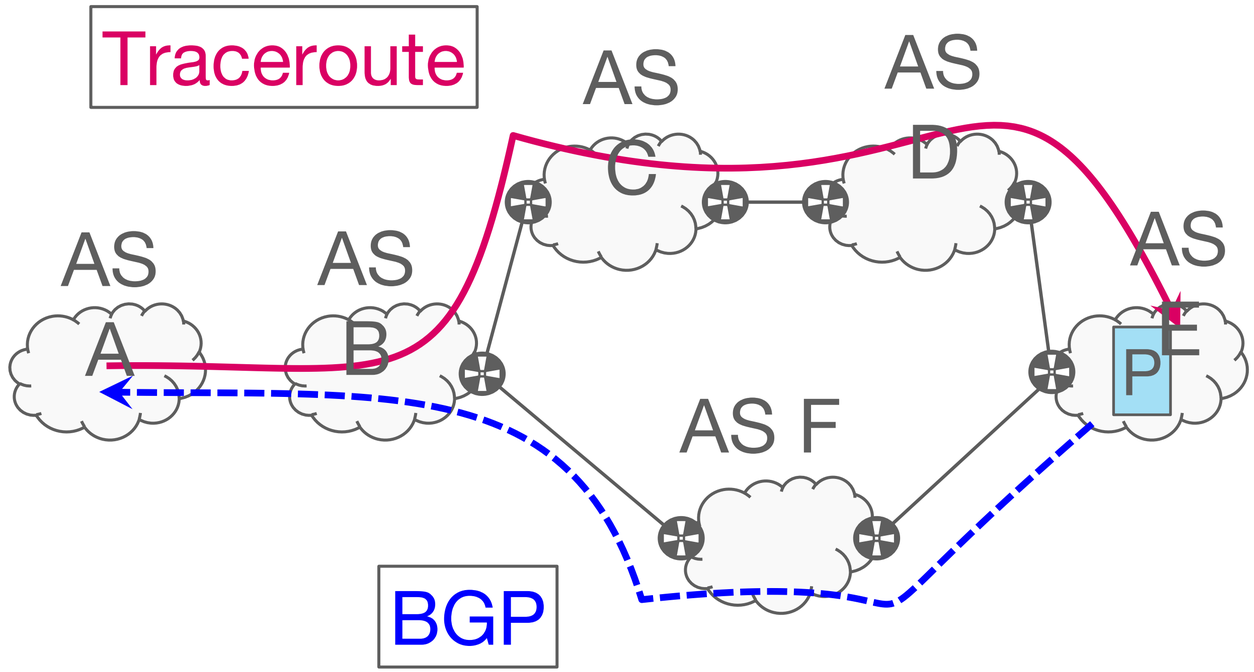


**Because** the peer C is cheaper

**Or** peer C pays B to access traffic data from AS A

**Or** …

# This difference in control and data paths may also be observed in the Kapela-Pilosov BGP monkey-in-the-middle attack

The topology

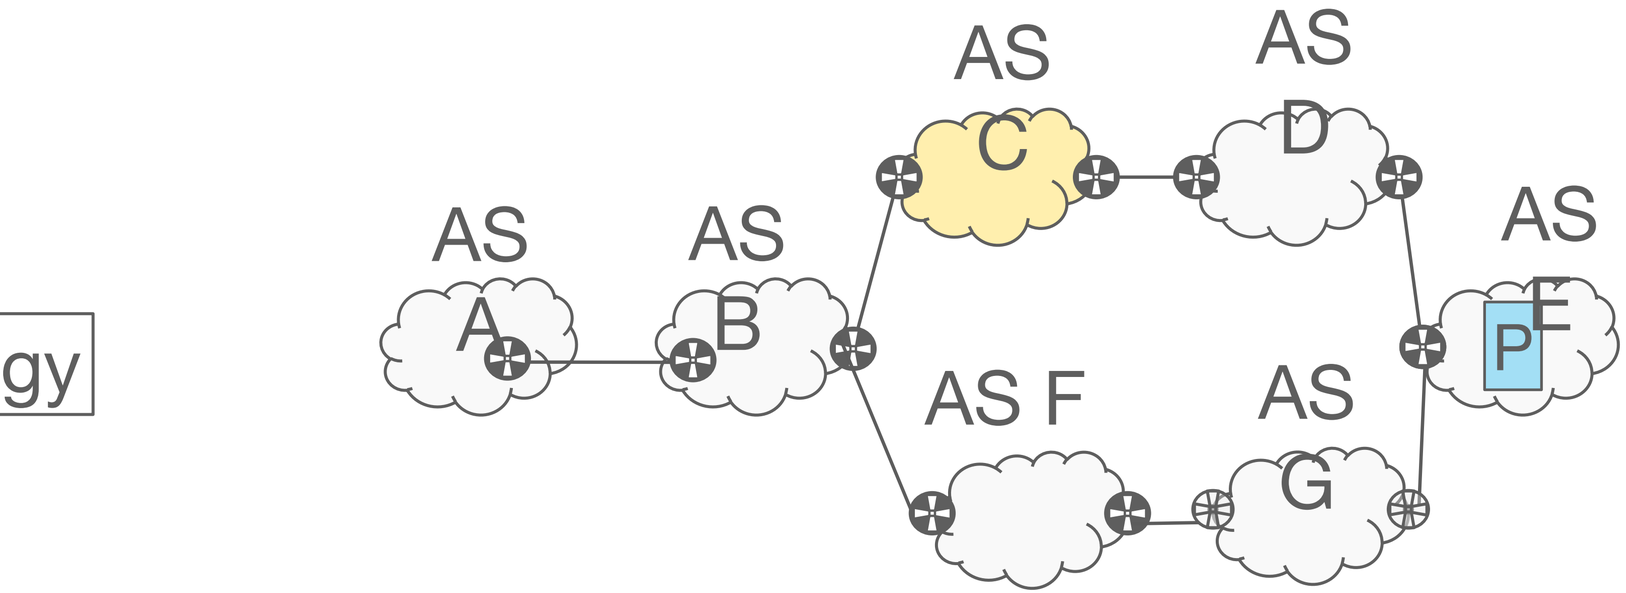

# This difference in control and data paths may also be observed in the Kapela-Pilosov BGP monkey-in-the-middle attack

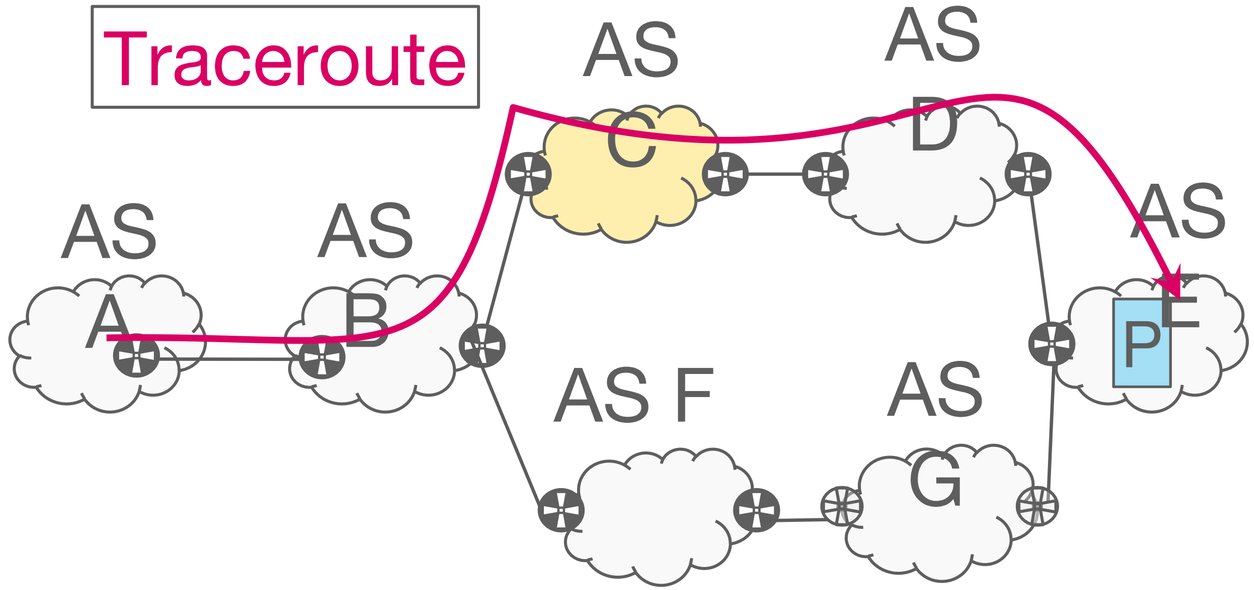

# This difference in control and data paths may also be observed in the Kapela-Pilosov BGP monkey-in-the-middle attack

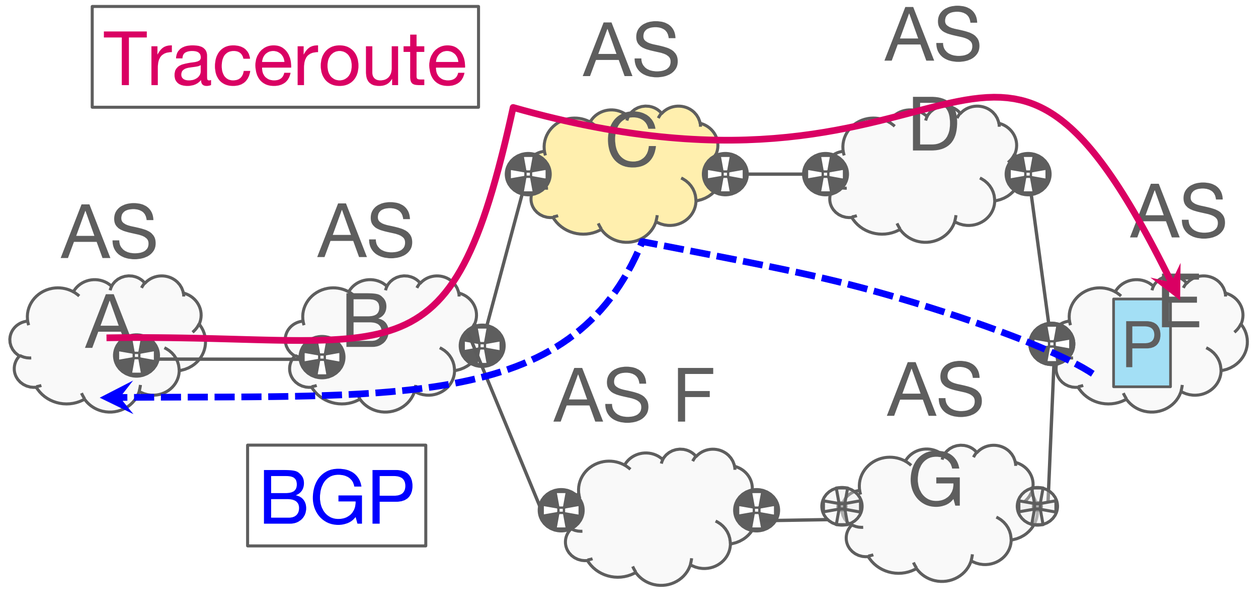


But for packets to follow the traceroute path, the yellow AS faked a direct link to the prefix origin

# The general assumption is that

For each external prefix P...

- The control path (CP) advertised in BGP
- And the data path (DP) used in **practice** are the same

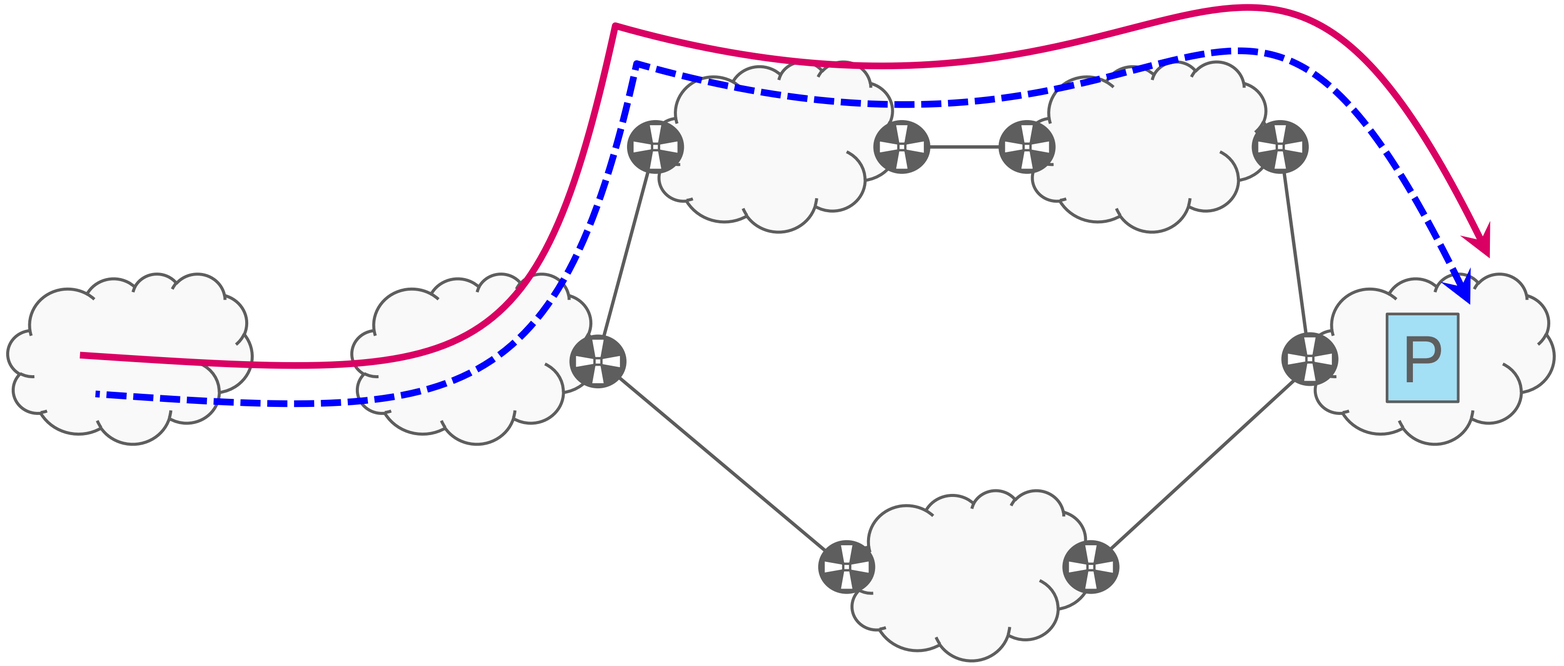

# One form of BGP lie is

when the control path (CP) and data path (DP) for a prefix P do not match

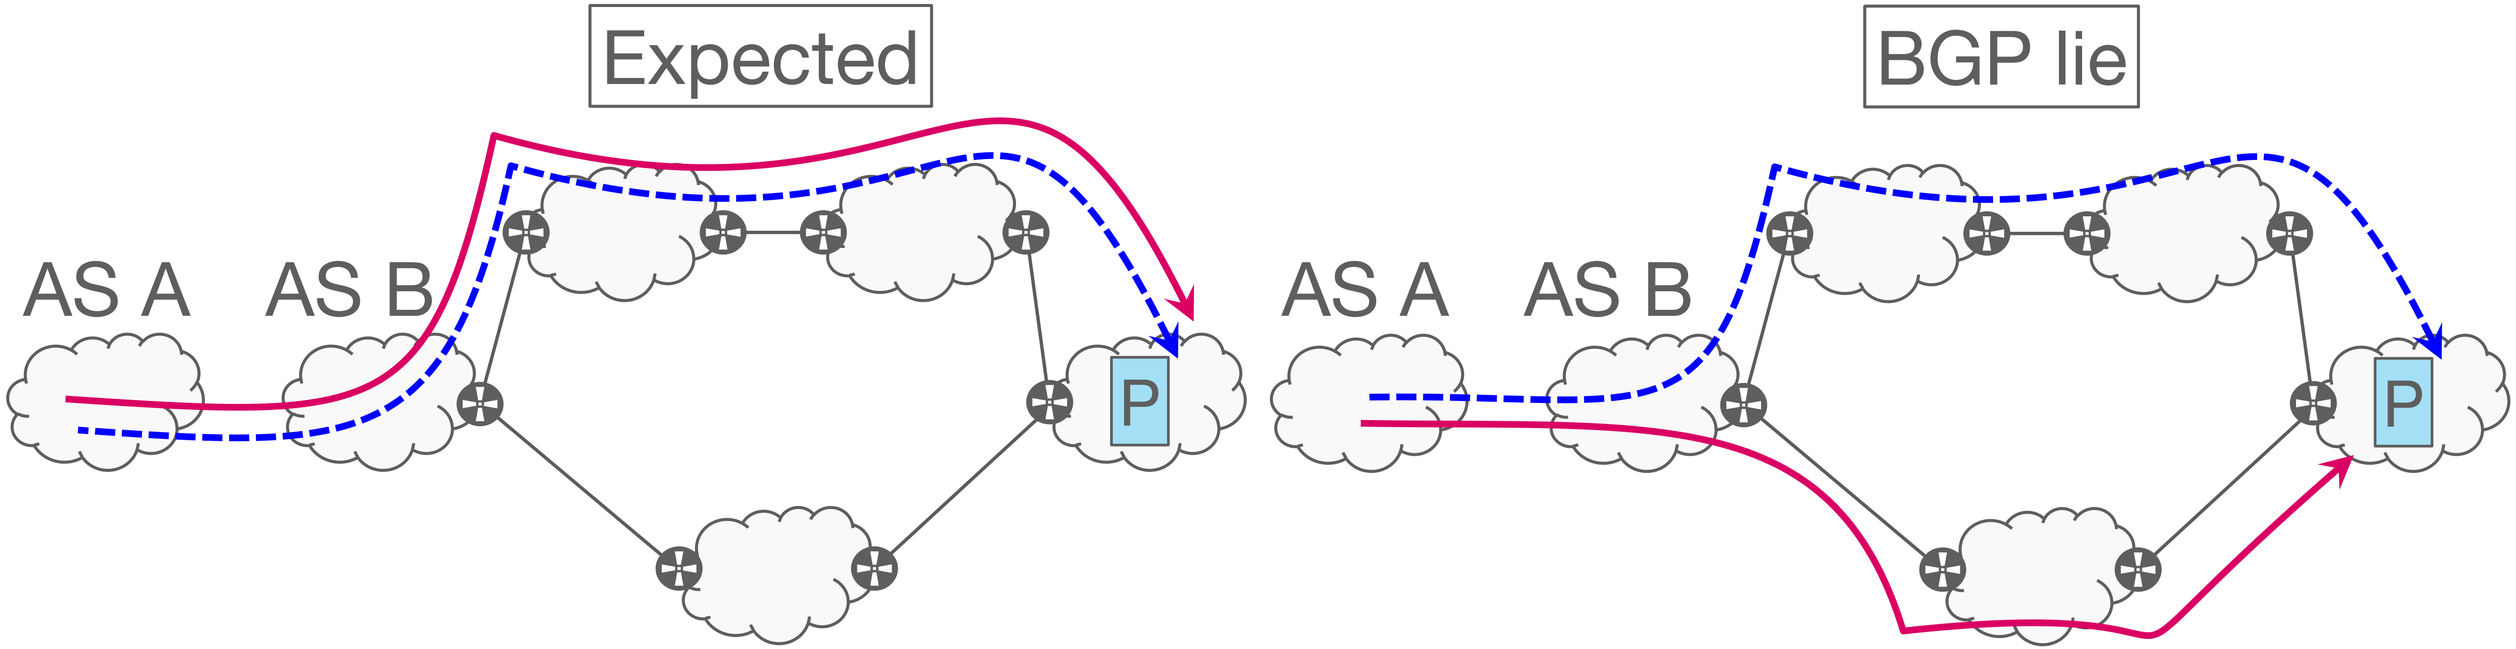

# Related publications

- Julian M. Del Fiore, Pascal Merindol, Valerio Persico, Cristel Pelsser and Antonio Pescape. ***Filtering the Noise to Reveal Inter-Domain Lies***, in 2019 Network Traffic Measurement and Analysis Conference (TMA), pages 17–24, 2019.

- Julian M. Del Fiore, Valerio Persico, Pascal Merindol, Cristel Pelsser and Antonio Pescape. ***The Art of Detecting Forwarding Detours***, in IEEE Transactions on Network and Service Management (IEEE TNSM) 2021.

# Some vulnerabilities of BGP

Prefix hijacks

Blackjack attacks

BGP lies

BGP session injection

# BGP runs on top of TCP

TCP is vulnerable to injection attacks

The attacker

- guesses the next sequence number
- sends a packet  with the sequence number and forged content

The client accepts the content if it arrives before the legit packet

The recommendation is to use MD5 for session authentication.

- But there are tools able to provide payload for a given MD5 digest https://github.com/DavidBuchanan314/monomorph
- The adoption status of TCP Authentication Option (TCP-AO) for BGP is not known

# Related publication

- Routing over QUIC: Bringing transport innovations to routing protocols. Thomas Wirtgen, Nicolas Rybowski, Cristel Pelsser, Olivier Bonaventure (2023). Poster at NSDI 2023.

# Some vulnerabilities of BGP

Prefix hijacks

Blackholing

BGP lies

BGP session injection

$\Rightarrow$BGP designed with no security in mind

Weak authentication

No integrity protection

# How we may hack to live with these vulnerabilities

# Prevention: some fixes

- RPKI ROA and ROV
  - State of deployment
- BGP filters
  - MANRS
- BGPsec

# RPKI ROA

1,078,454 RIB entries covered by ROAs in May 2022 (V4 and V6 together).

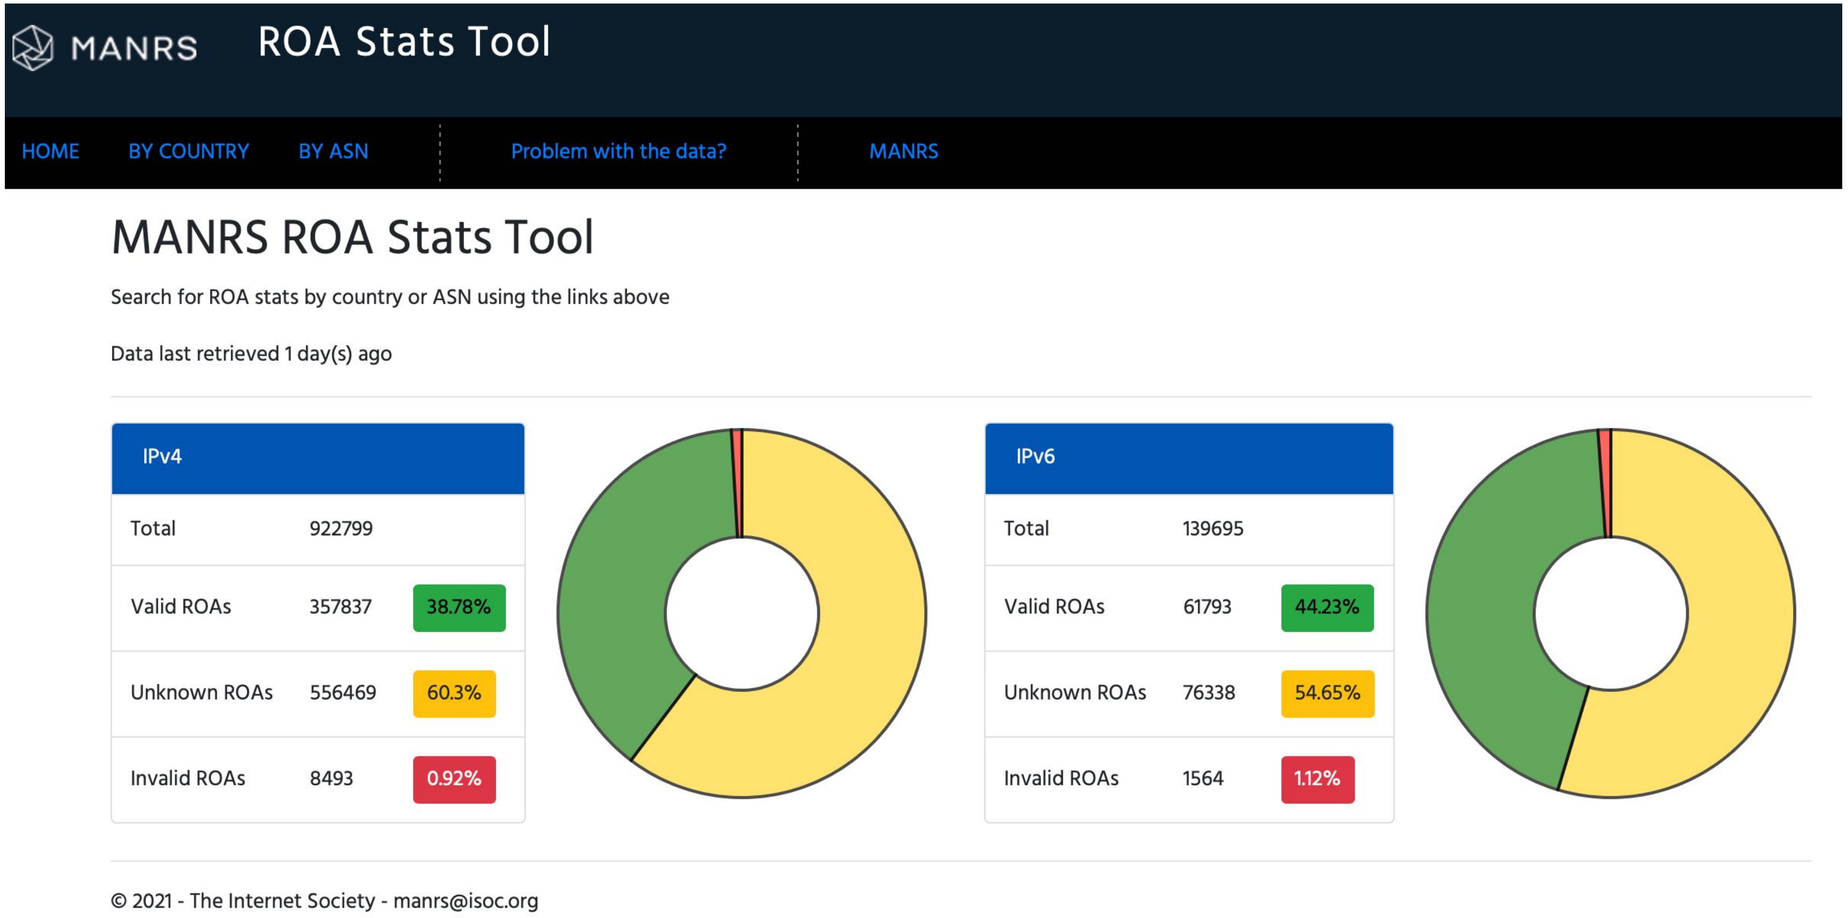


https://roa-stats.manrs.org (October 6, 2022) -> similar results for May 2023

# RPKI ROV

75 ASs deploy ROV (certainty above 0.7) according to **rov.rpki.net** (out of > 73.5k) → Last measurement was on 2020-08-31

Only 5.9 % of user prefixes are protected according to the MANRS Observatory (May 2023)

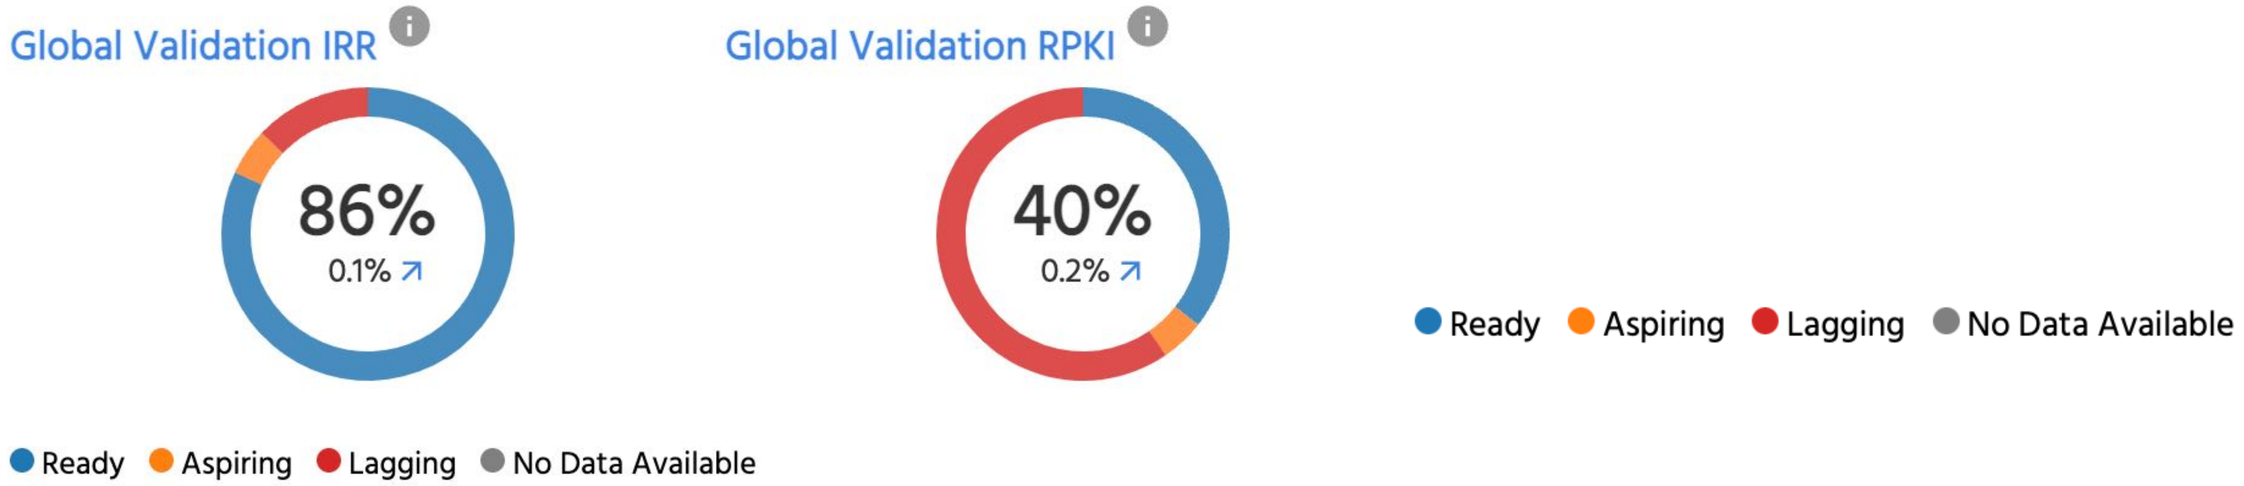


From https://observatory.manrs.org/#/overview (May, 2023)

# BGP filters and MANRS

Mutually Agreed Norms for Routing Security (MANRS)

| Action | Metric | Description | Data source(s) |
|---|---|---|---|
| Filtering | M1 | Route leak by the AS<br>Calculates incidents where the AS was the culprit of BGP leakage events. In the example on Fig 1. if all pink events are route leaks by the AS, M1=3.5 | bgpstream |
| | M2 | Route misorigination by the AS<br>Calculates incidents where the AS was the culprit of BGP misorigination (hijacking) events. | bgpstream or GRIP (*) |
| | M1C | Route leak by a direct customer<br>Calculates incidents where the AS was an accomplice (the misoriginating AS was present in the AS-PATH) to BGP leakage events. Currently only incidents related to adjacent networks are taken into account. | bgpstream |
| | M2C | Route misorigination by a direct customer<br>Calculates incidents where the AS was an accomplice (the leaking AS was present in the AS-PATH) to BGP hijack events. Currently only incidents related to adjacent networks are considered. | bgpstream or GRIP (*) |
| | M3 | Bogon prefixes by the AS.<br>Calculates incidents where the AS originated bogon address space. Note that the duration of each incident is counted per day as the data in the CIDR report is available only on a daily basis. Like with leaks and hijacks all prefixes originated by the AS on a day counted as 1 incident. | CIDR report |
| | M3C | Bogon prefixes propagated by the AS.<br>Calculates incidents where the AS propagated bogon address space announcements received from its peers. | CIDR report |
| | M4 | Bogon ASNs by the AS<br>Calculates incidents where the AS announced bogon ASNs as adjacency.<br>Note that the duration of each incident is counted per day as the data in the CIDR report is available only on a daily basis. | CIDR report |
| | M4C | Bogon ASNs propagated by the AS<br>Calculates incidents where the AS propagated bogon ASNs announcements it received from its peers.<br>Note that the duration of each incident is counted per day as the data in the CIDR report is available only on a daily basis. | CIDR report |

# BGP filters and MANRS

Mutually Agreed Norms for Routing Security (MANRS) rules for filter setting to prevent

- Leaks
- Misorigination
- Bogon prefixes
- Bogon ASs

From the AS itself and from direct customers

# Detecting and localizing who deploys ROV

- **Towards a Rigorous Methodology for Measuring Adoption of RPKI Route Validation and Filtering,** Andreas Reuter et al (2017). Computer Communications Review (CCR).

- BGP Beacons, Network Tomography, and Bayesian Computation to Locate Route Flap Damping. C. Gray, M. Clemens, R. Bush, C. Pelsser, R. Matthew, T. Schmidt, M. Wählisch (2020). *Internet Measurement Conference (IMC)*.

# Highlight

- Intro to BGP and its vulnerabilities
- Some fixes to these vulnerabilities and **their impact**
  - RPKI time of flight
- Attacks are still possible
- Getting the best of BGP data
  - Most valuable set of Vantage Points (MVP)
- Detecting BGP hijacks
  - Detection of type-1 BGP hijacks (DFOH)

# RPKI time of flight

Tracking delays

- in the management,
- control and
- data plane

**Life cycle of RPKI data?**

**How quickly does it affect Internet routing and reachability?**

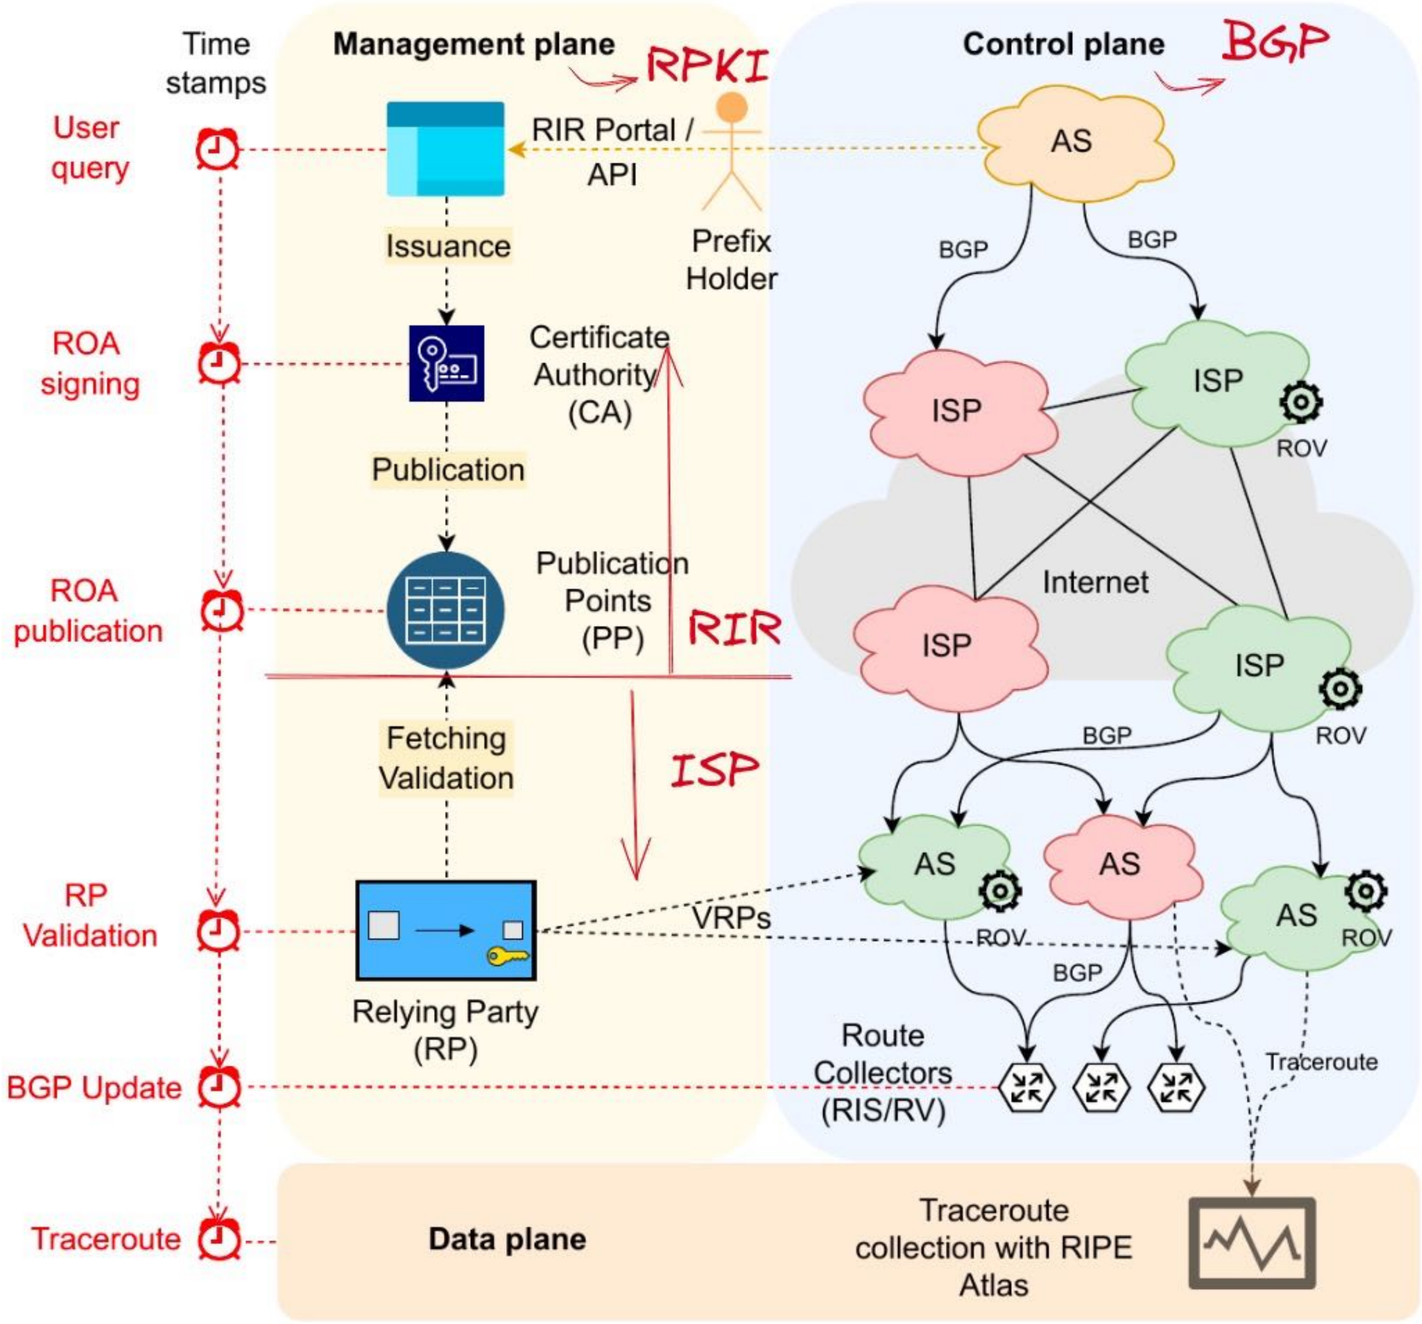

# Experimental setup

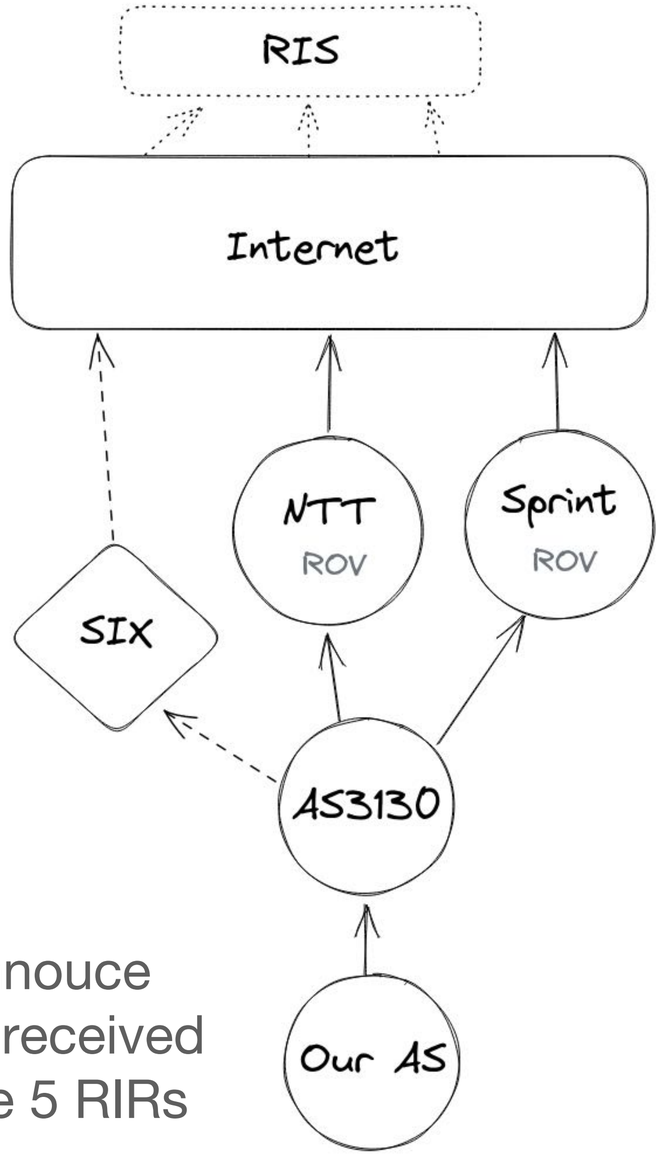


We annouce prefixes received from the 5 RIRs

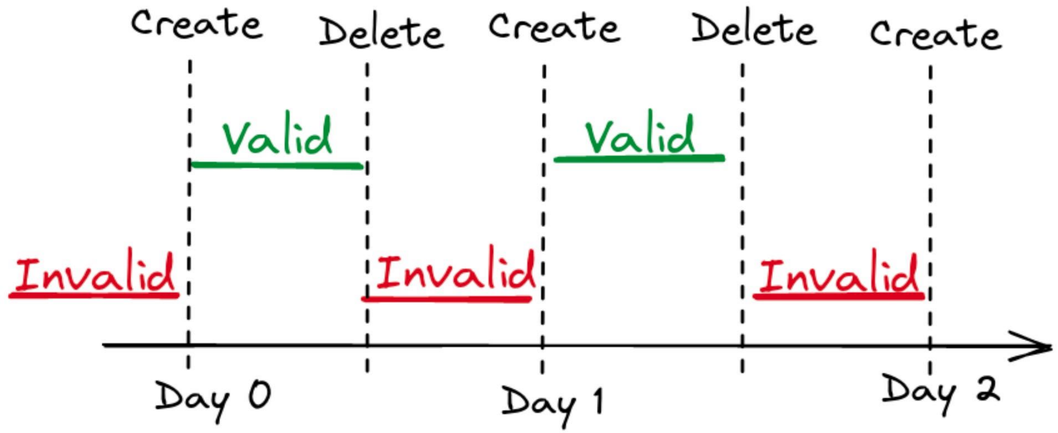


**Initialisation**

Announce a prefix in BGP

Create a ROA to make it **invalid**

**ROA toggling**

1. **Create** a ROA to make it **valid**
2. **Delete** that ROA: **invalid** again
3. Go back to 1)

# ROA creation delay in minutes

| | Sign* | NotBefore* | Publication† | Relying Party† | BGP‡ |
|---|---|---|---|---|---|
| AFRINIC | 0 (0) | 0 (0) | 3 (2) | 14 (13) | 15 (16) |
| APNIC | 10 (13) | 10 (13) | 14 (16) | 34 (38) | 26 (28) |
| ARIN | - (-) | - (-) | 69 (97) | 81 (109) | 95 (143) |
| LACNIC | 0 (0) | - (-) | 54 (32) | 66 (42) | 51 (34) |
| RIPE | 0 (0) | 0 (0) | 4 (4) | 14 (13) | 18 (18) |
| After fix: | | | | | |
| ARIN | - (-) | - (-) | 8 (9) | 21 (22) | 28 (23) |


- APNIC processes requests in batches every 20 minutes
- LACNIC and APNIC had a time zone issue which delayed the publication
  - They fixed the issue after we notified them

# ROA deletion delay in minutes

| | Revocation* | Relying Party† | BGP‡ |
|---|---|---|---|
| AFRINIC | 0 (0) | 13 (14) | 34 (38) |
| APNIC | 10 (12) | 31 (36) | 51 (56) |
| ARIN | 0 (0) | 14 (16) | 45 (51) |
| LACNIC | 0 (0) | 18 (20) | 48 (49) |
| RIPE | 0 (0) | 14 (13) | 41 (50) |


- Effect in BGP twice longer than creation time
  - Because all RPs/ASs have to revoke the ROA
- Batching still present at APNIC

# Lessons learned

- Stuck ROA
- Timezone bug at LACNIC and ARIN

- RPKI is orders of magnitude slower than BGP
- Impact for network operators
  - Time to repair a bad ROA
  - Time to authorize a DDoS mitigator

# Related publication

*RPKI time of flight: Tracking delays in the management, control and data plane.*
Romain Fontugne, Amreesh Phokeer, Cristel Pelsser, Kevin Vermeulen, Randy Bush. PAM 2023

# Highlight

- Intro to BGP and its vulnerabilities
- Some fixes to these vulnerabilities and their impact
  - RPKI time of flight
- **Attacks are still possible**
- Getting the best of BGP data
  - Most valuable set of Vantage Points (MVP)
- Detecting BGP hijacks
  - Detection of type-1 BGP hijacks (DFOH)

# Deployment of protection increases but events still occur (FR)

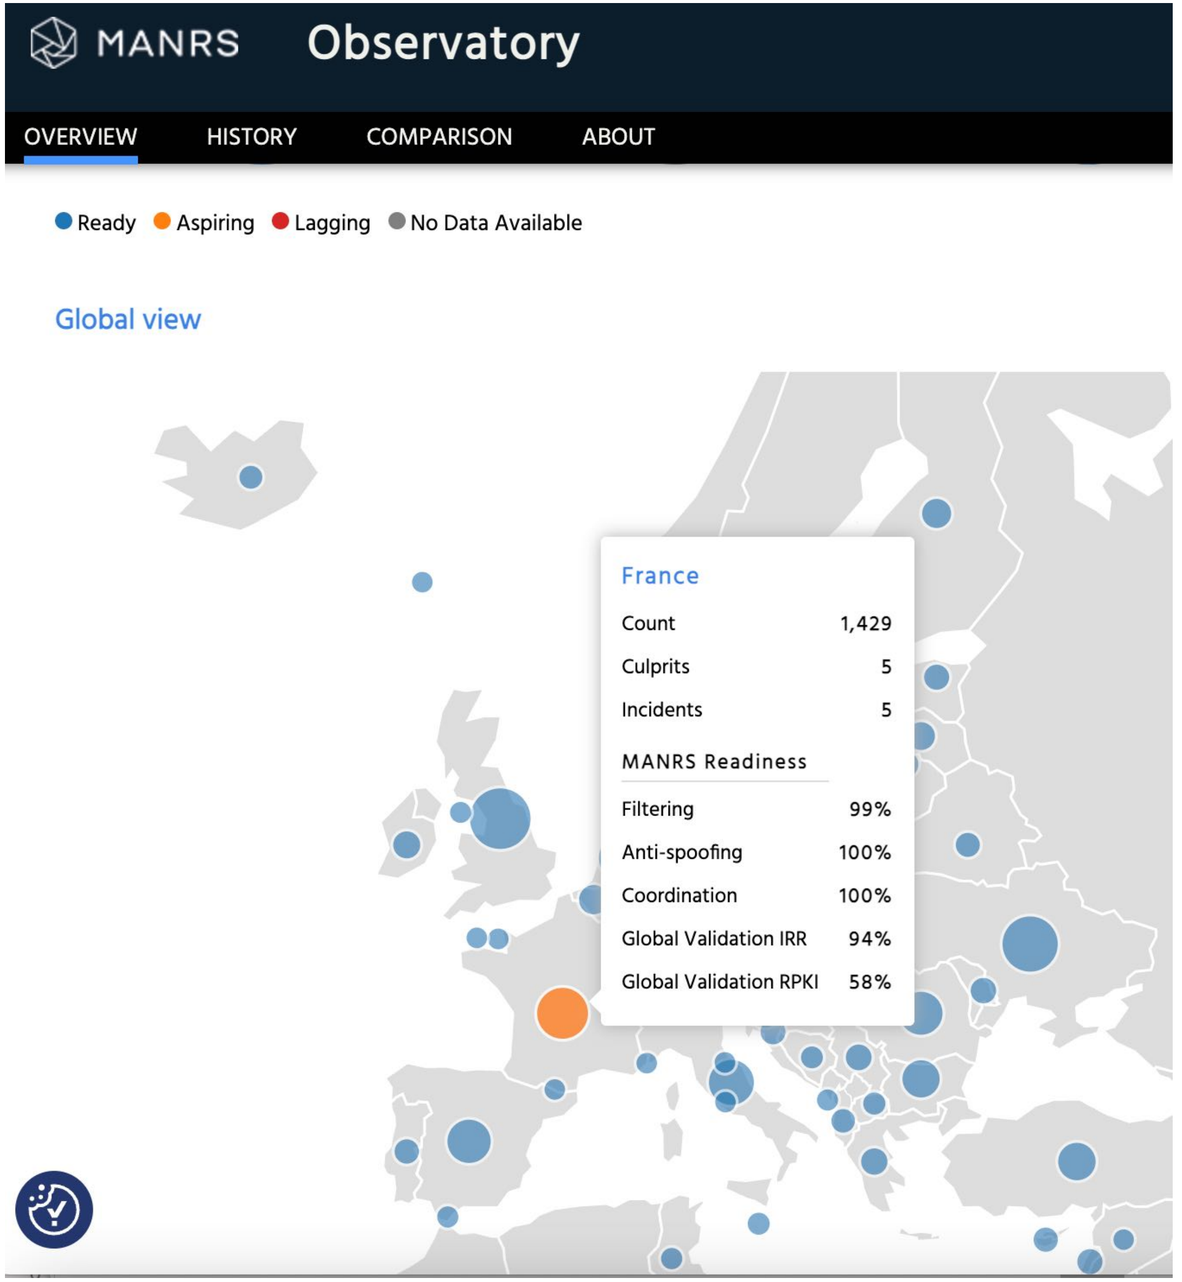

# Deployement of protection increases but events still occur (US)

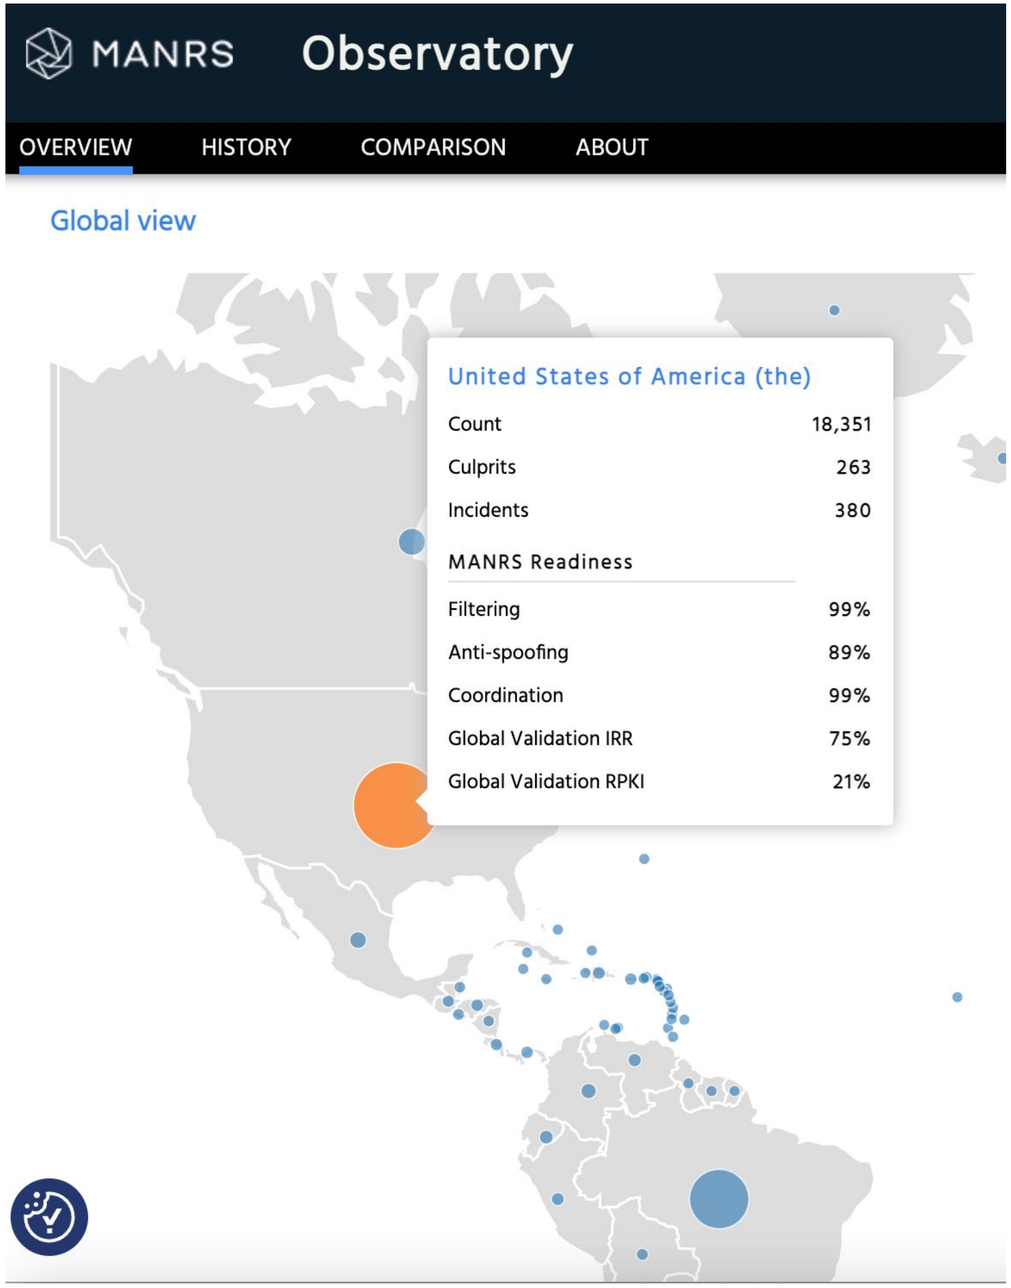

# Highlight

- Intro to BGP and its vulnerabilities
- Some fixes to these vulnerabilities and their impact
  - RPKI time of flight
- Attacks are still possible
- **Getting the best of BGP data**
  - Most valuable set of Vantage Points (MVP)
- Detecting BGP hijacks
  - Detection of type-1 BGP hijacks (DFOH)

# Quadratic increase of BGP data

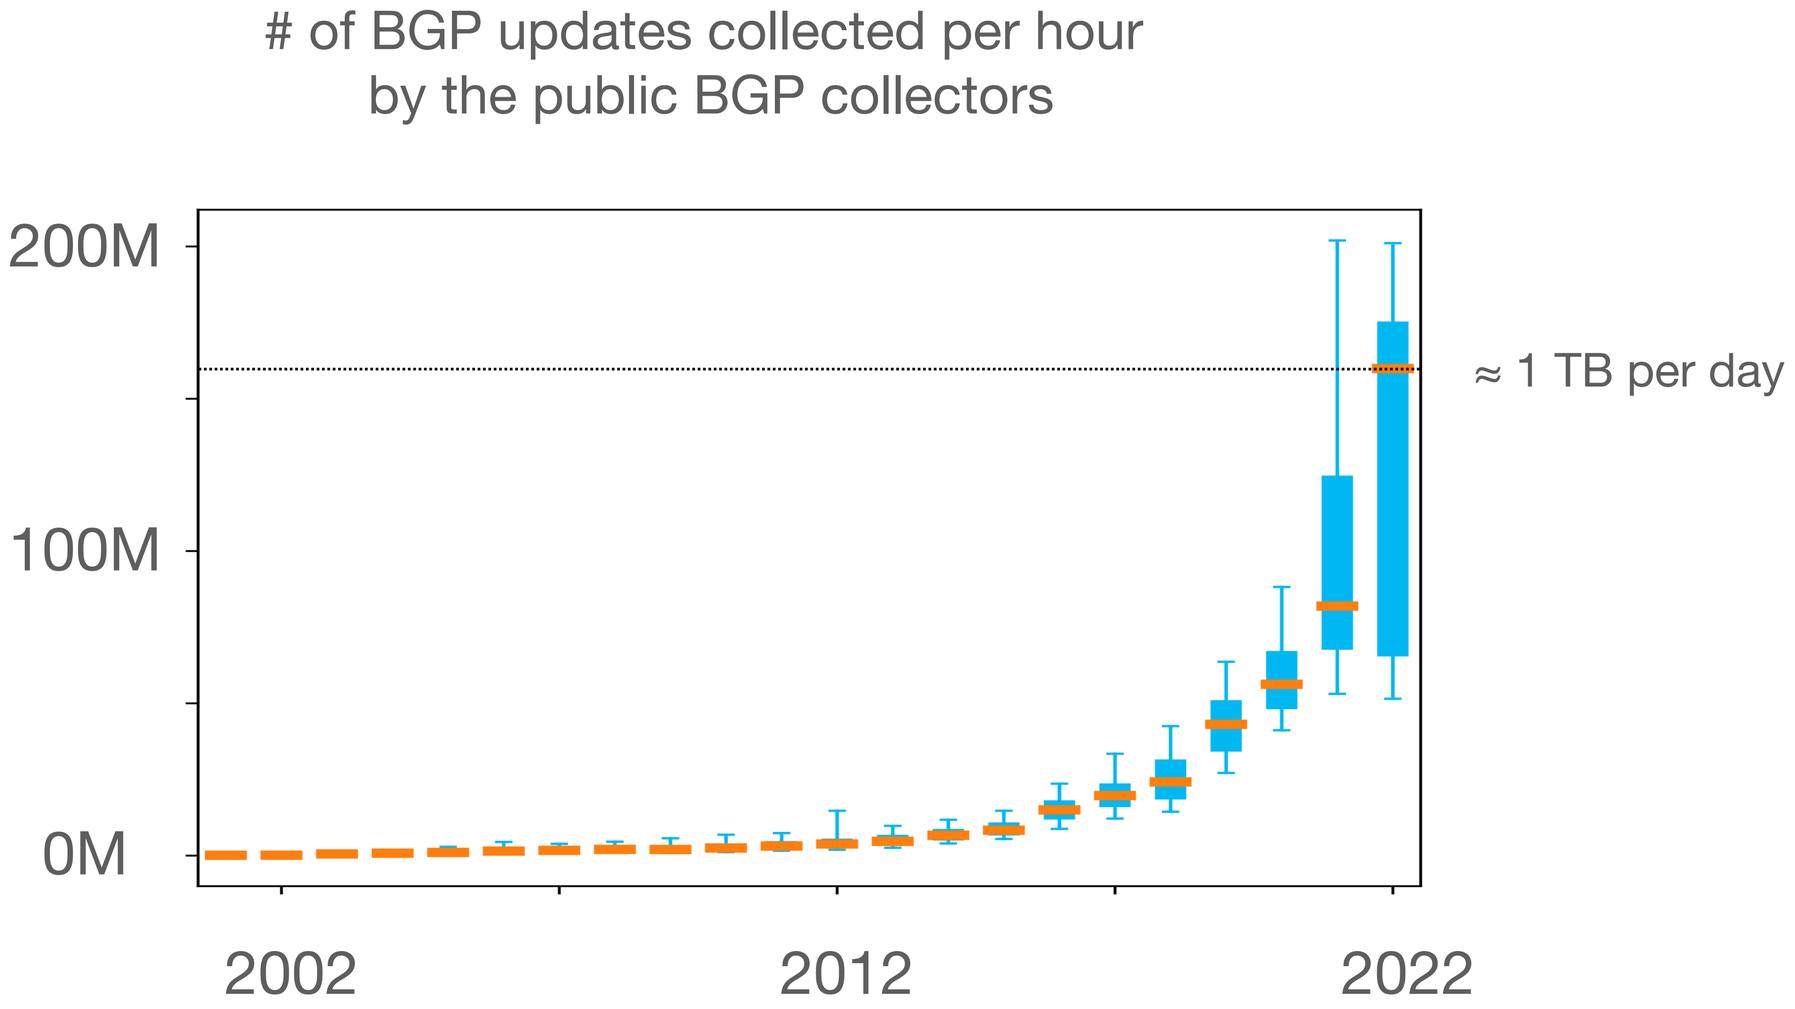

# With their limited processing power, users often arbitrarily focus on a subset of the VPs

But many events are detected by only a few vantage points

Proportion of the BGP Hijacks and mis-originations

Number of vantage points that detected the event

| | **2019** | **2020** |
|---|---|---|
| **1 - 5** | 21% | 22% |
| **6 - 30** | 20% | 25% |
| **> 30** | 60% | 53% |
| **Total #** | 1782 | 2477 |

MANRS blogpost: BGP Security in 2021

# Our goal: Find a set of BGP vantage points that maximises utility and minimizes volume of data

We introduce a redundancy score that uses the distances to evaluate redundancy across all events for a pair of VP

VP1 - VP2 : High average distance → Lowly redundant

VP2 - VP3 : Low average distance → Highly redundant

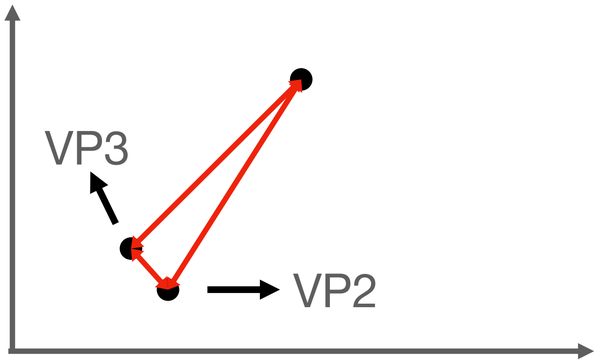


Event 1

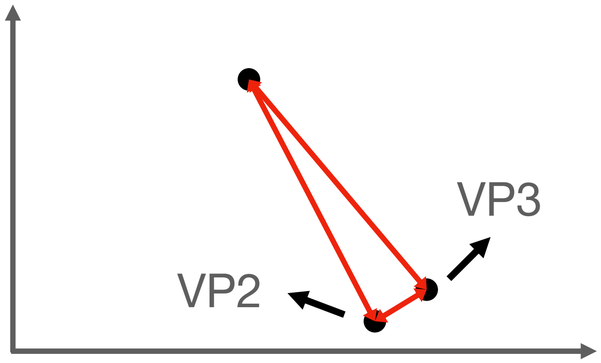


Event 2

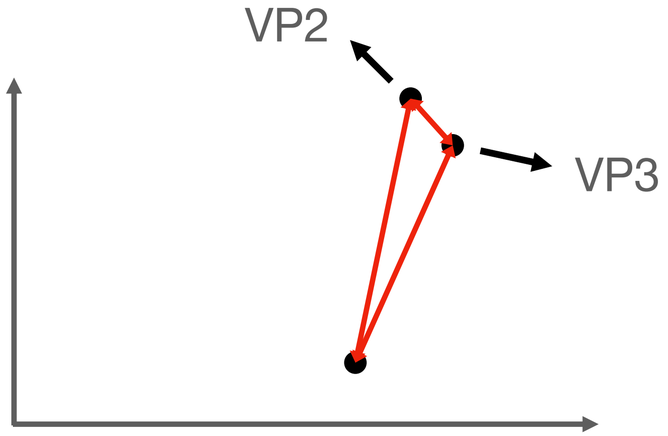


Event 3

MVP discovers 400k AS links with 2.5 times less process messages compared to random selection (95th percentile)

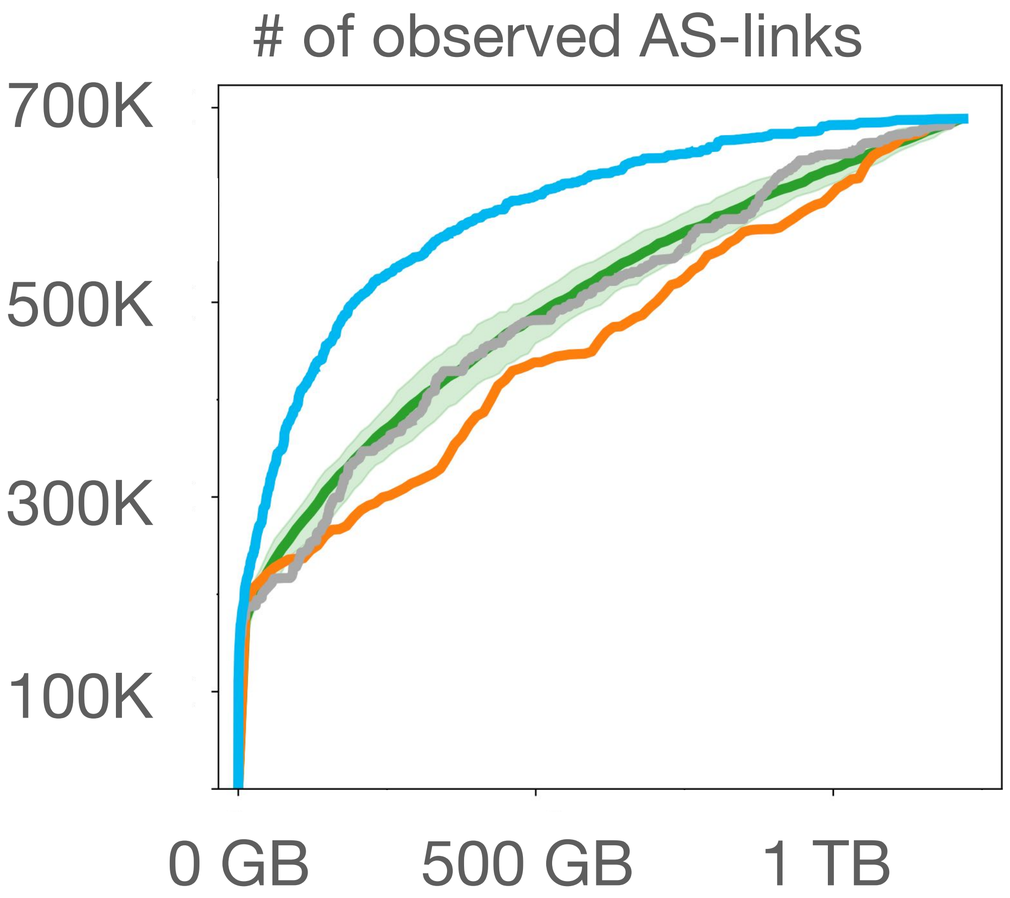


MVP
Unbiased
Distance-Based
Random

MVP detects 60% of the events with 2.2 times less process messages
compared to random selection (95th percentile)

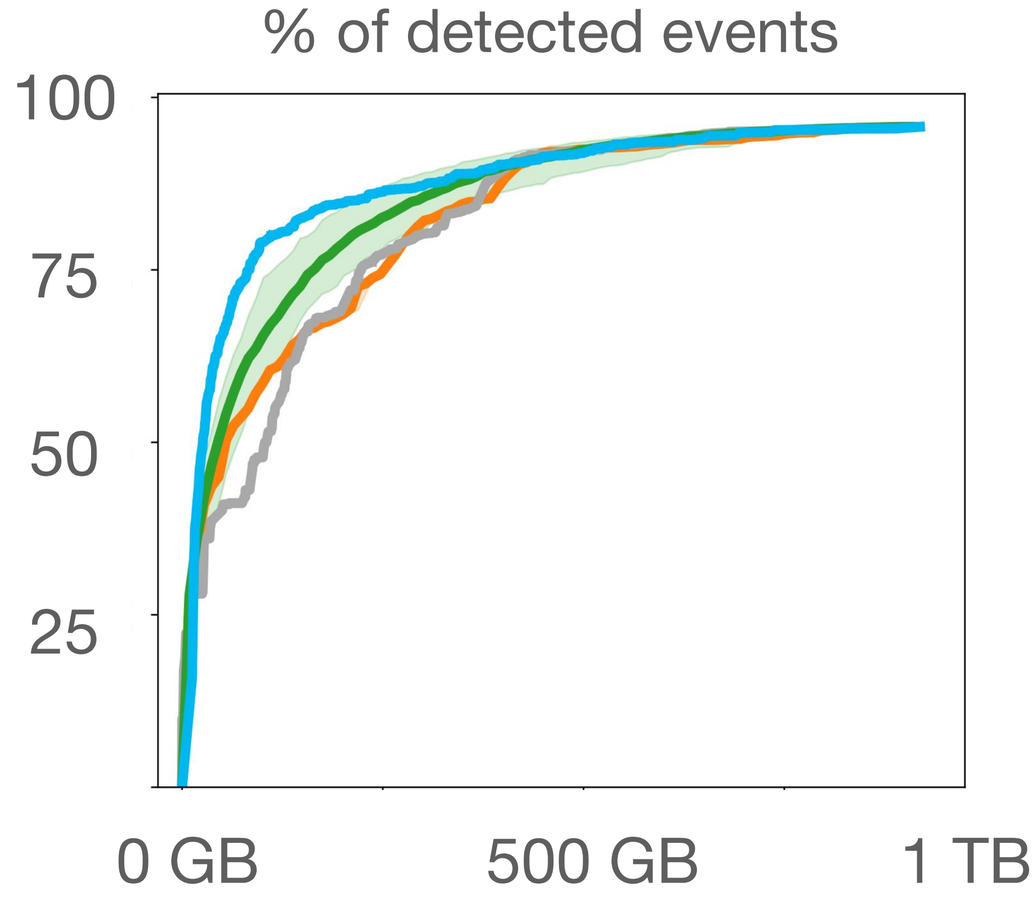

# Related publication

Thomas Alfroy, Thomas Holterbach, **Cristel Pelsser** (2022). MVP: Measuring Internet Routing from the Most Valuable Points. *Poster in the Proceedings of the Internet Measurement Conference (IMC)*.

# MVP has a significant impact on the results of other research works

MVP helps inferring more AS relationships without any cost
in terms of volume or algorithmic change

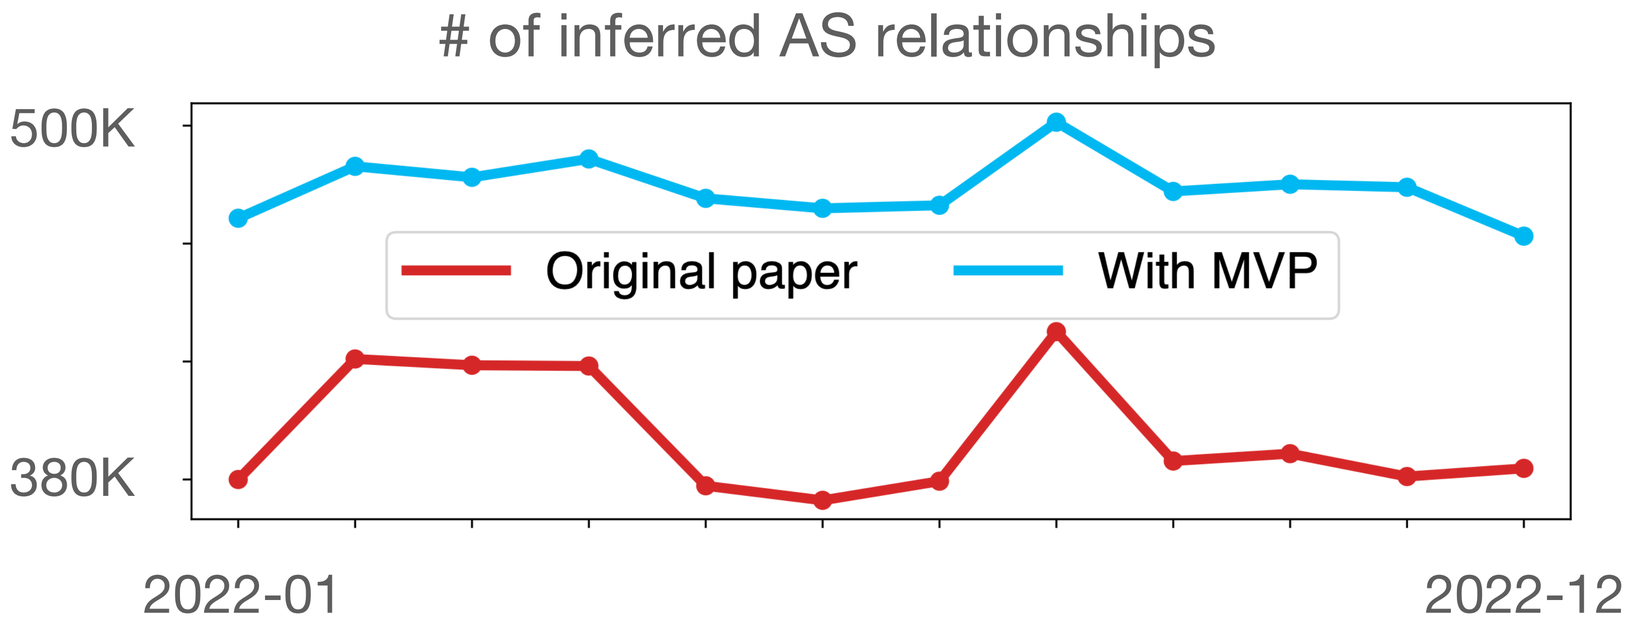

# We’ll use this selection for the detection of attacks

The output of MVP

# Highlight

- Intro to BGP and its vulnerabilities
- Some fixes to these vulnerabilities and their impact
  - RPKI time of flight
- Attacks are still possible
- Getting the best of BGP data
  - Most valuable set of vantage points (MVP)
- **Detecting BGP hijacks**
  - Detection of type-1 BGP hijacks

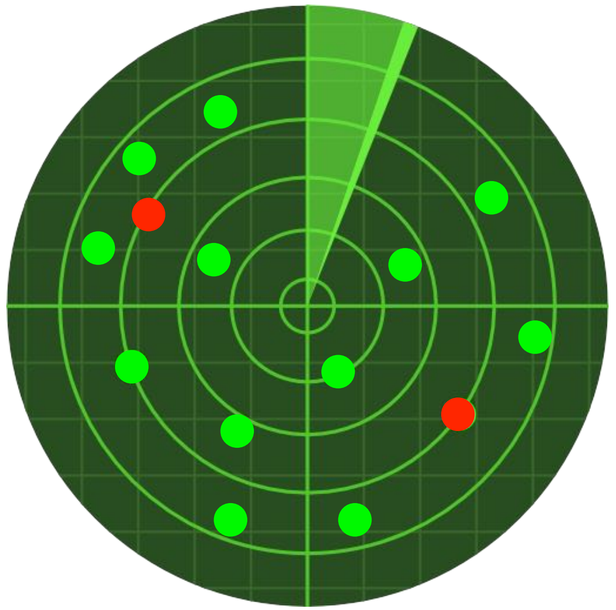

# What is a type-1 hijack ?

The origin AS is legit.

The AS-path is not.

A new link appears in the Internet topology

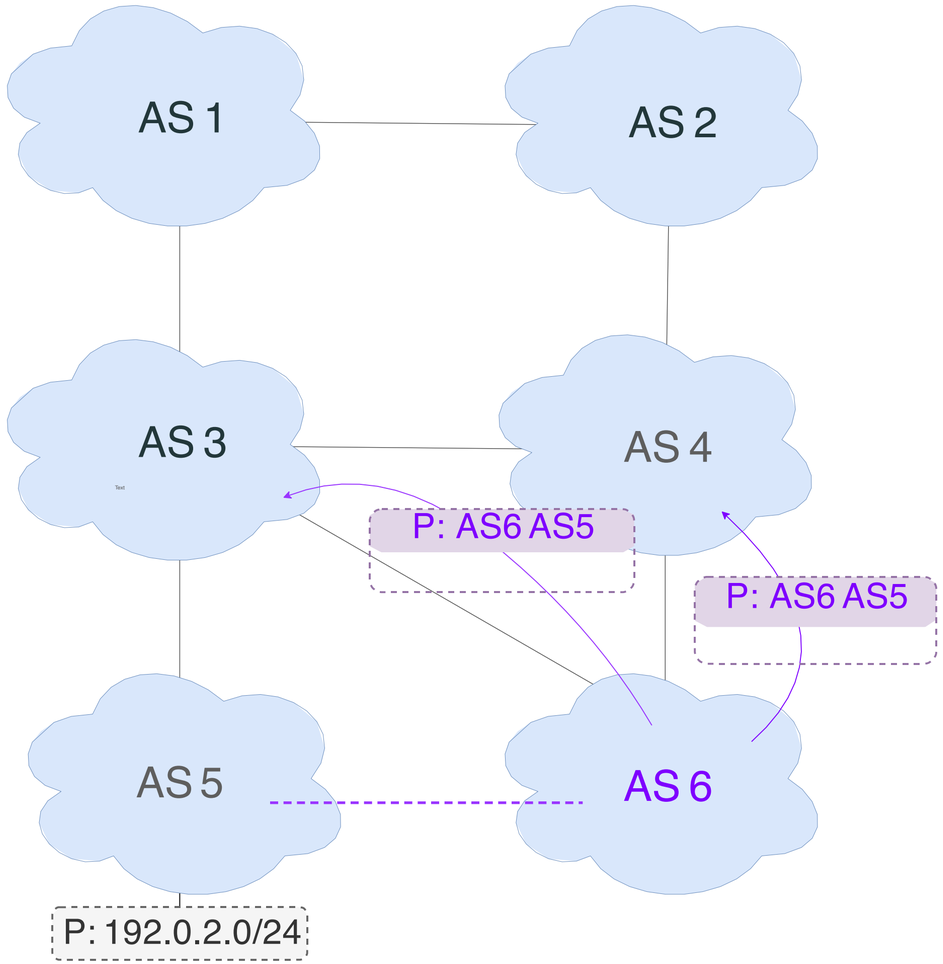

# Type-1 hijacks often introduce a new link in the AS-level topology

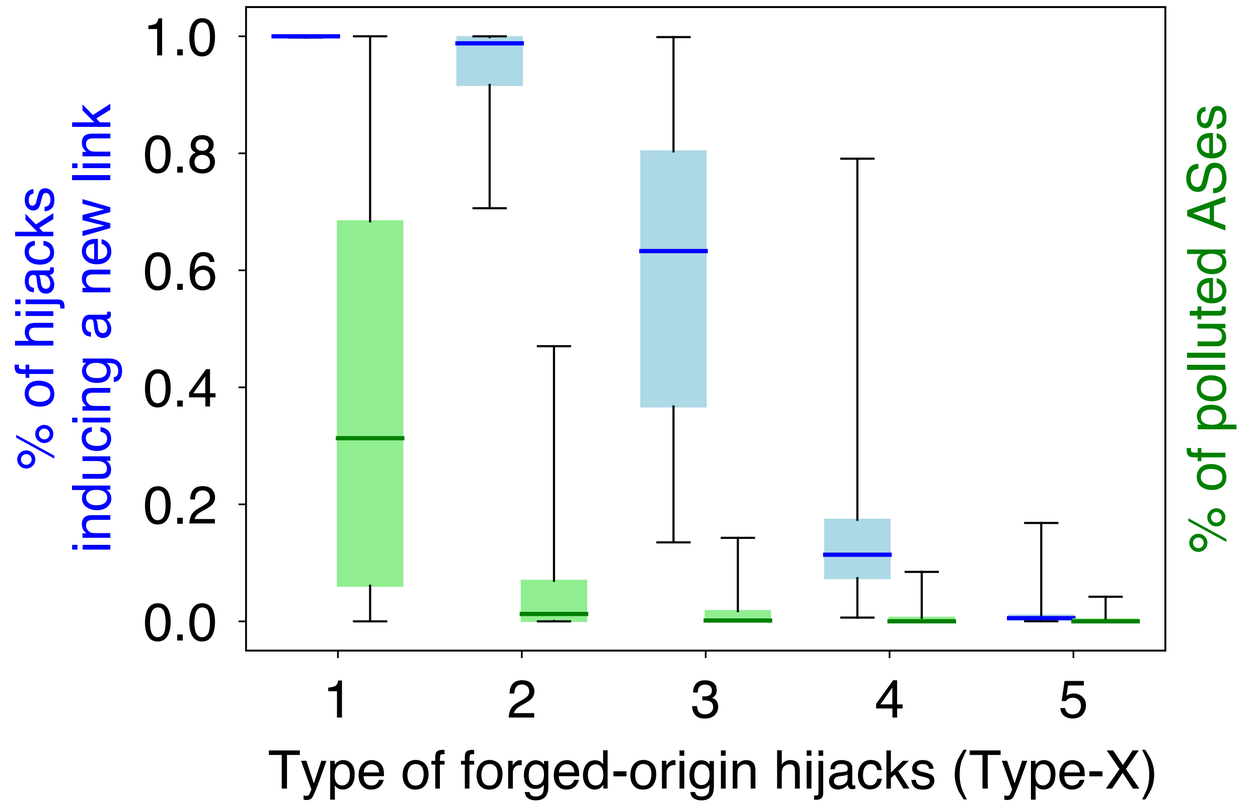

# To detect type-1 to type-X hijacks we aim to determine if new links are legitimate.

166 new AS links
every day (median)

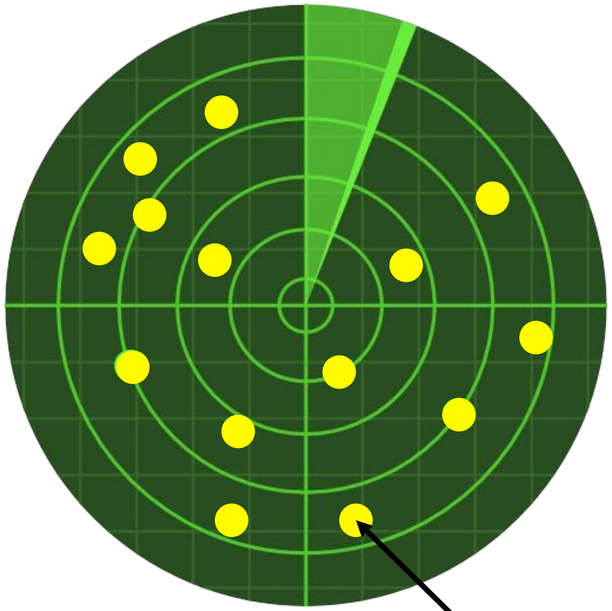

New AS link

***Detector***

Real AS link

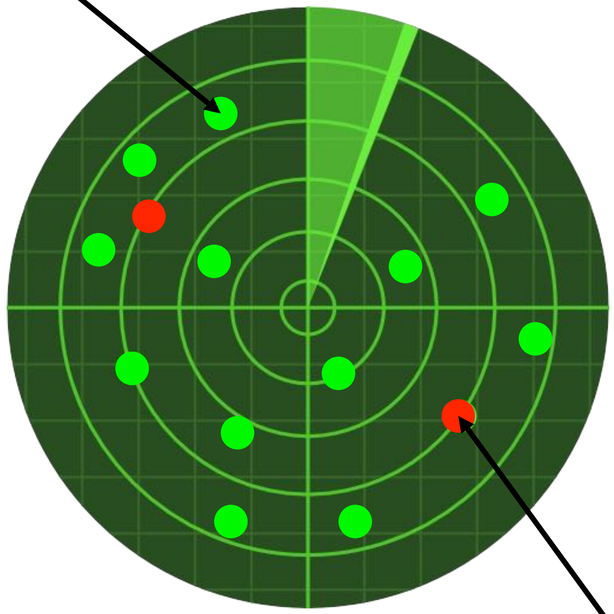

Fake AS link

***DFOH*** runs its own domain-specific inference algorithm
to discriminate fake links from the real ones

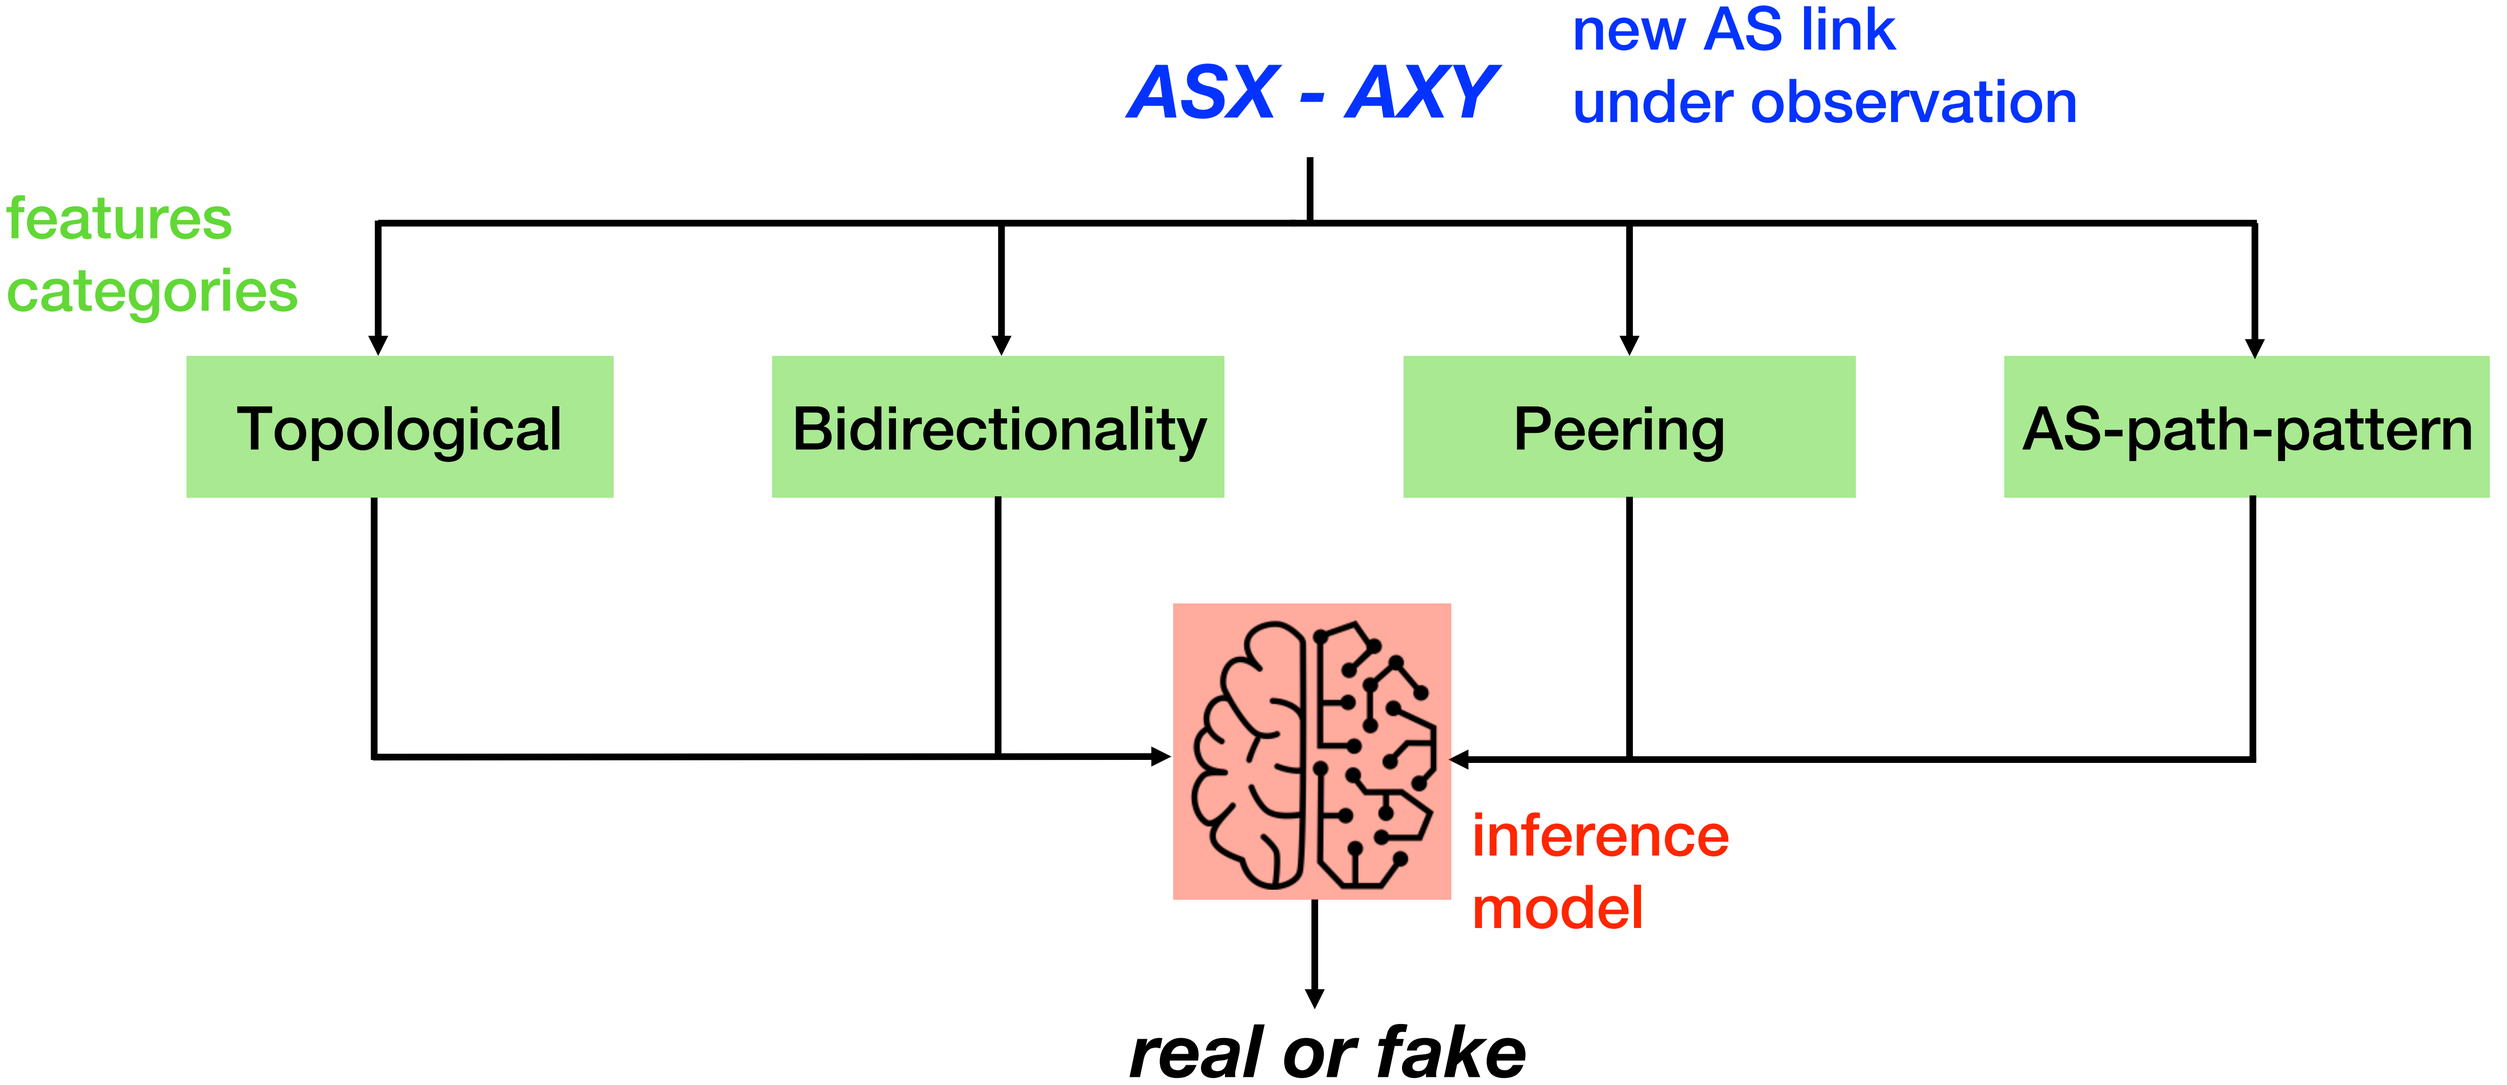

# ***Our detector*** makes the network-wide detection of forged-origin hijacks practical

***Our detector*** greatly limits the number of alarms seeing by every AS

Proportion of ASes seeing at least one alarm every day

Number of ASes involved in at least one case every day

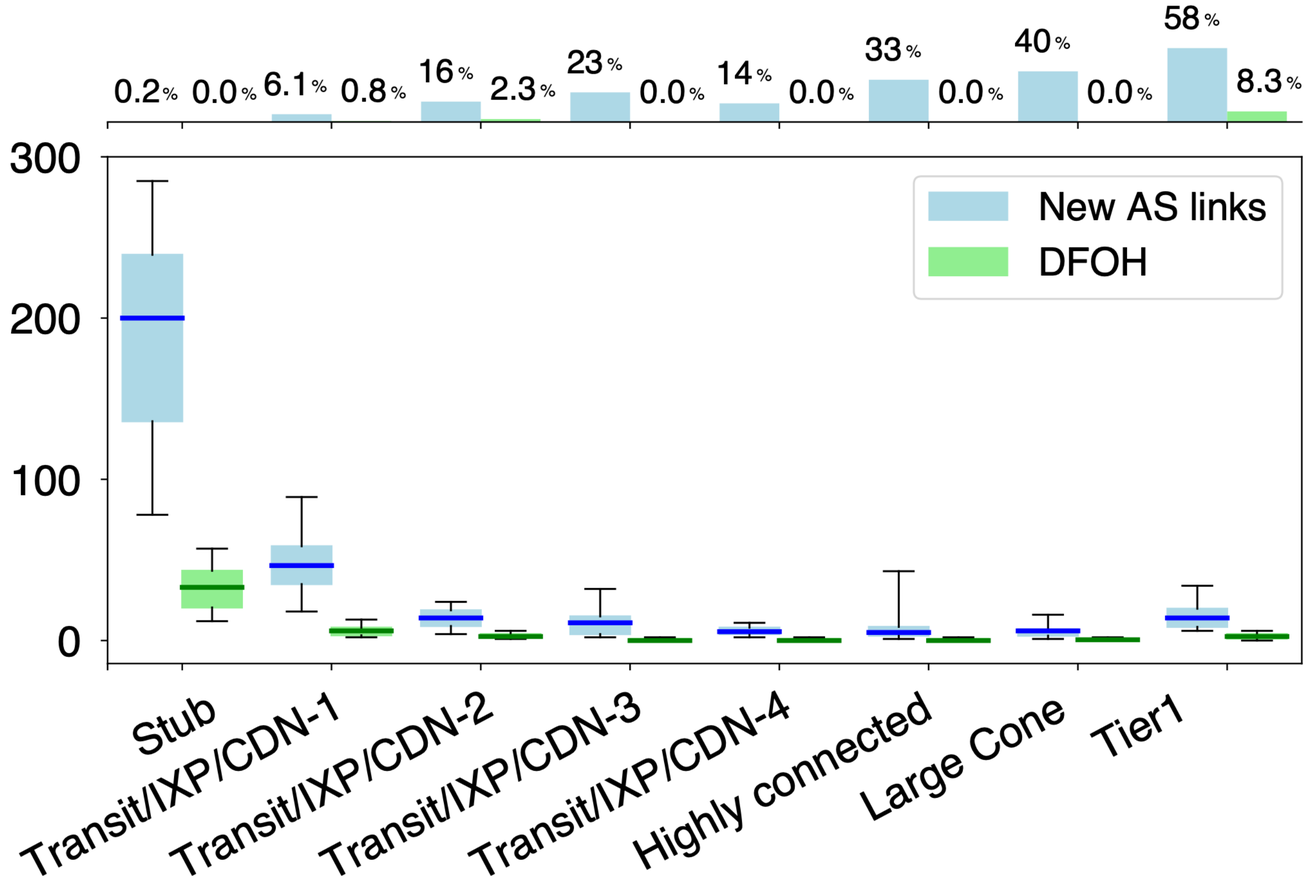

# Detection of type-1 hijacks

Thomas Holterbach, Thomas Alfroy, Cristel Pelsser

# Some of my work on detecting outages

- R. Fontugne , E. Aben , C. Pelsser, R. Bush. *Pinpointing Delay and Forwarding Anomalies Using Large-Scale Traceroute Measurements*, IMC 2017.
- A. Guillot, R. Fontugne , P. Winter , P. Merindol, A. King , A. Dainotti , C. Pelsser. *Chocolatine: Outage Detection for Internet Background Radiation*, TMA 2019.
- Odnan Ref Sanchez , Simone Ferlin , Cristel Pelsser, Randy Bush. *Comparing Machine Learning Algorithms for BGP Anomaly Detection using Graph Features*. Big-DAMA'19: ACM CoNEXT Workshop 2019.
- Anant Shah , Romain Fontugne , Emile Aben , Cristel Pelsser, Randy Bush. *Disco: Fast, Good, and Cheap Outage Detection*. TMA 2017.

# Conclusion

- Today we only have partial fixes to BGP vulnerabilities
- Their deployment can affect current network operations

- Our knowledge of the Internet topology is partial
- Better selecting VP may enable to deploy more VPs and improve our view of the Internet

- We use diverse features and data sets to detect anomalies
- Robustness to attack is important