A sample of current routing vulnerabilities and How we may hack to live with them

Cristel Pelsser
UCLouvain

October 7, 2022

Focus on the inter-domain routing protocol
BGP

BGP

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

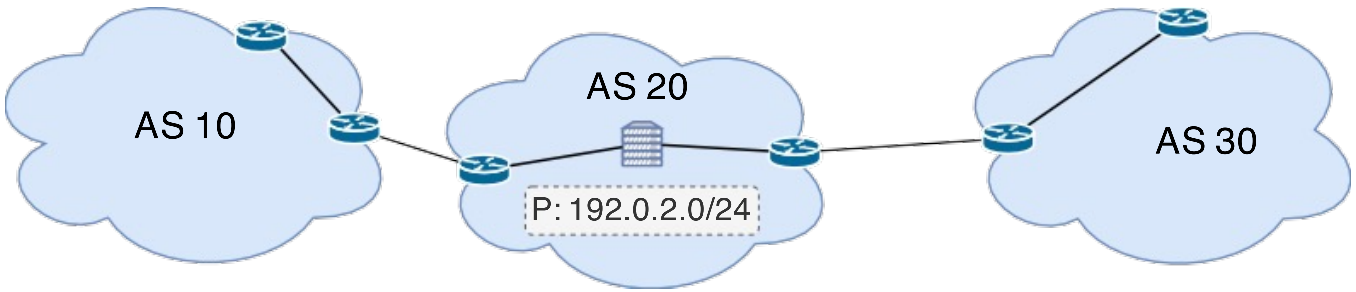

BGP

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.

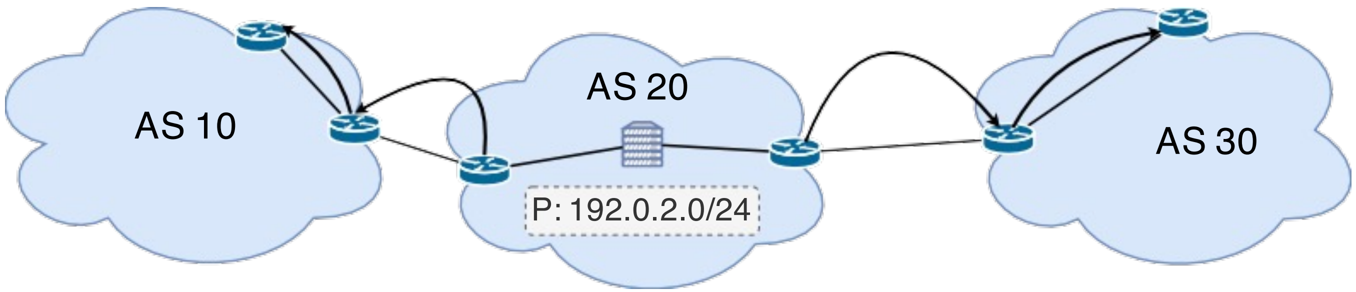

The Internet is composed of **Autonomous Systems** (AS): one or more networks under the control of a single entity.

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

Traffic flows in the reverse direction.

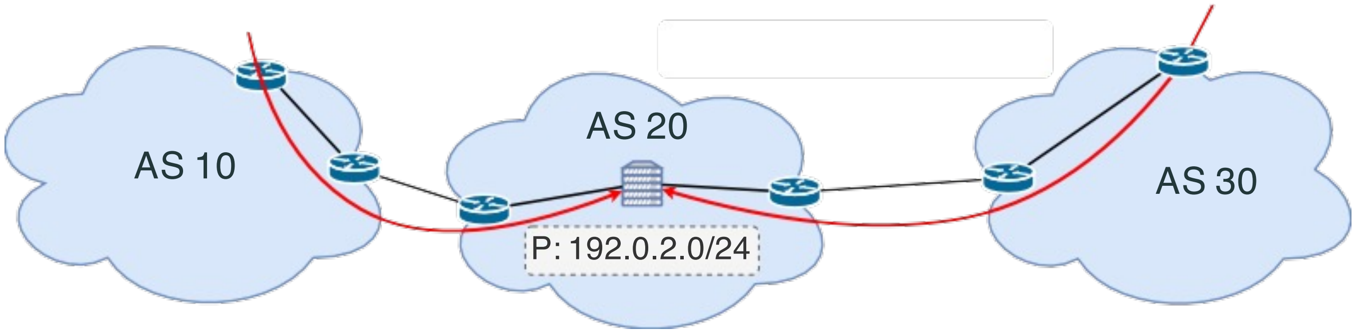

The Internet is a complex ecosystem

There are 73,501 AS advertised as of Oct 2, 2022.

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

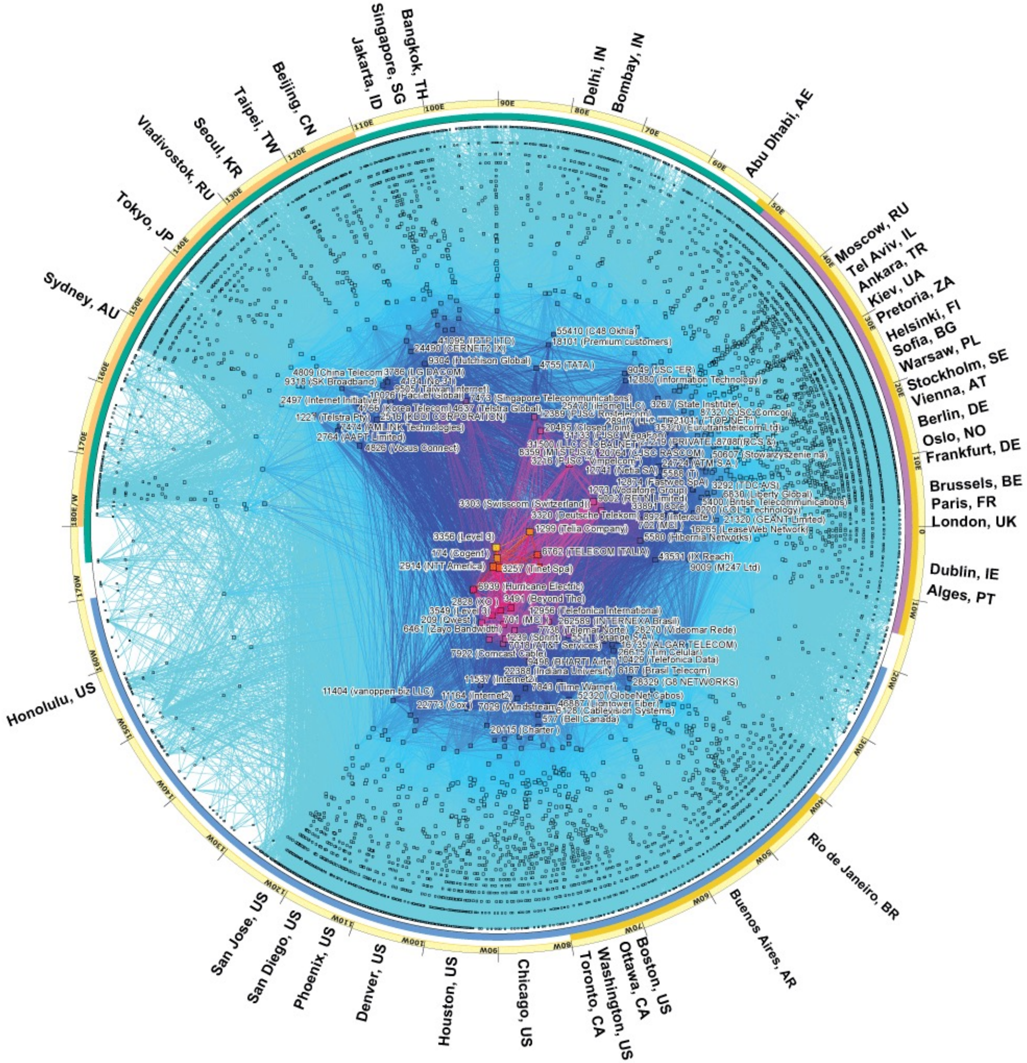

Source: https://www.caida.org/projects/cartography/as-core/2017/

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

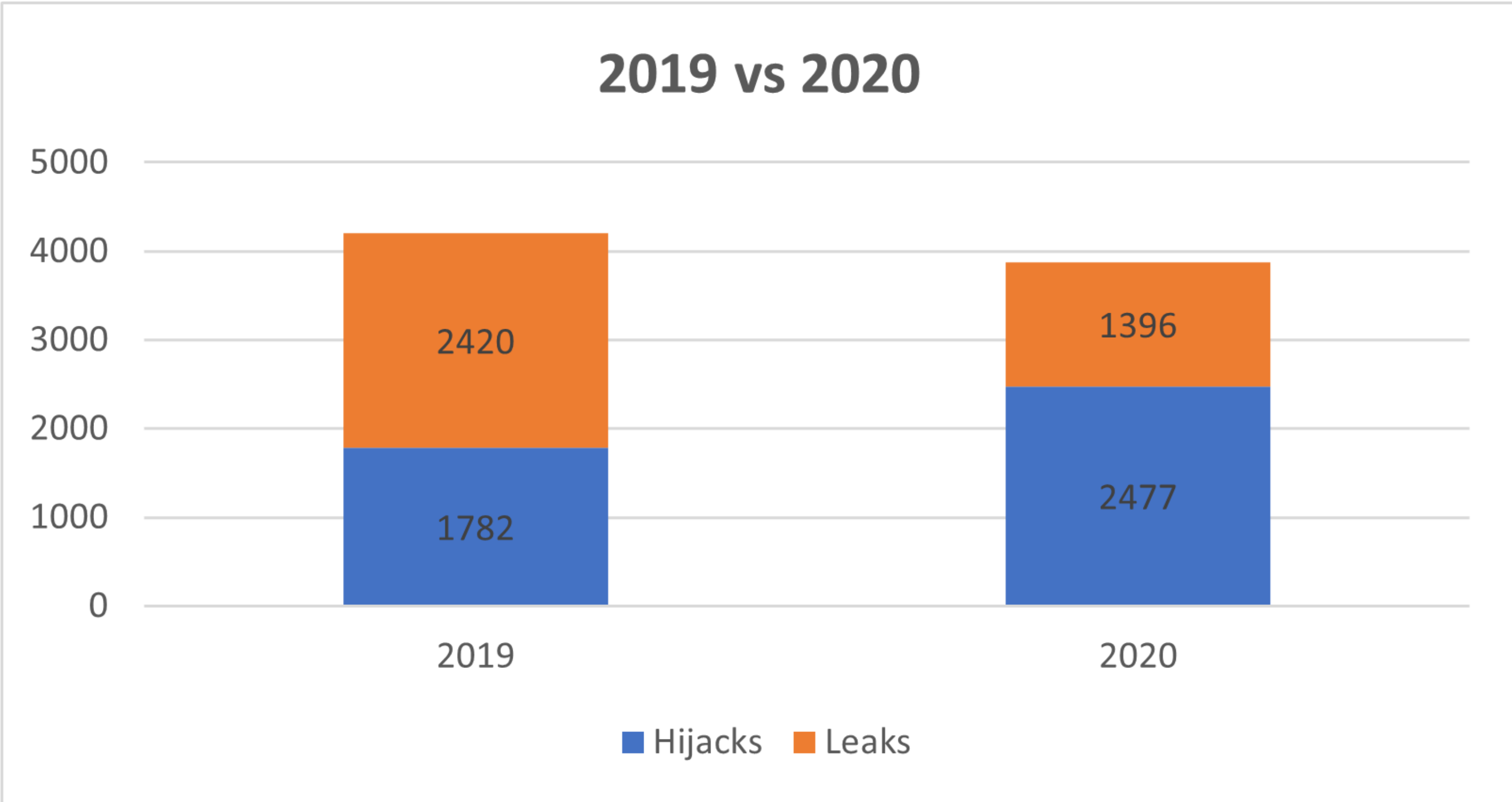

Source: https://www.manrs.org/2021/02/bgp-rpki-and-manrs-2020-in-review/

Some vulnerabilities of BGP

Prefix hijacks

Blackholing

BGP lies

BGP session injection

Hijacks can be used to divert traffic and gain inside knowledge

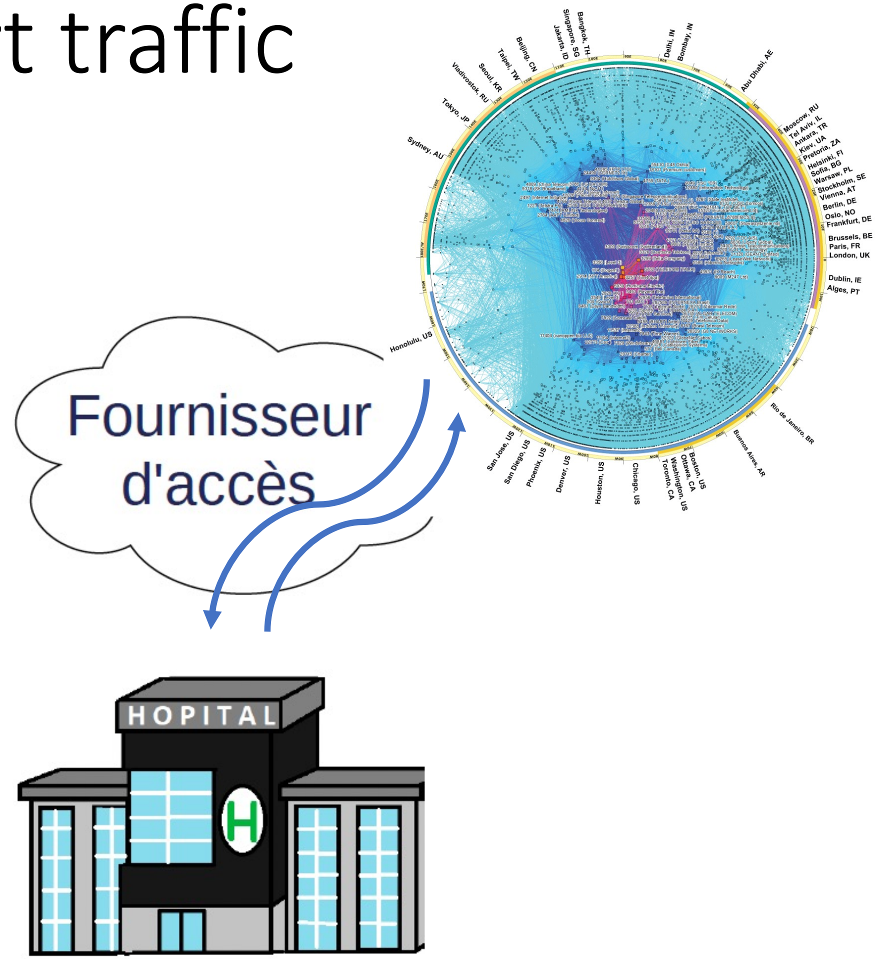

Hijacks can be used to divert traffic and gain inside knowledge

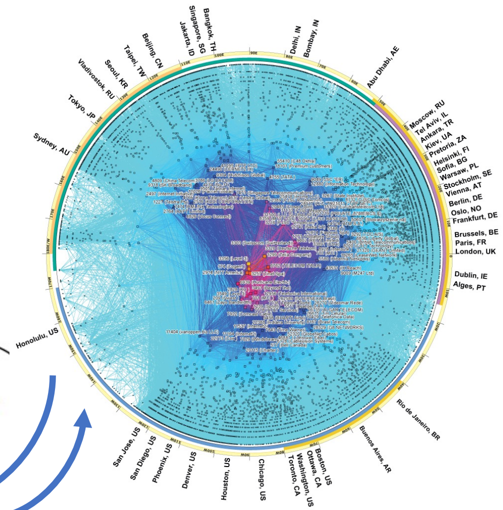

Hijacks can be used to divert traffic and gain inside knowledge

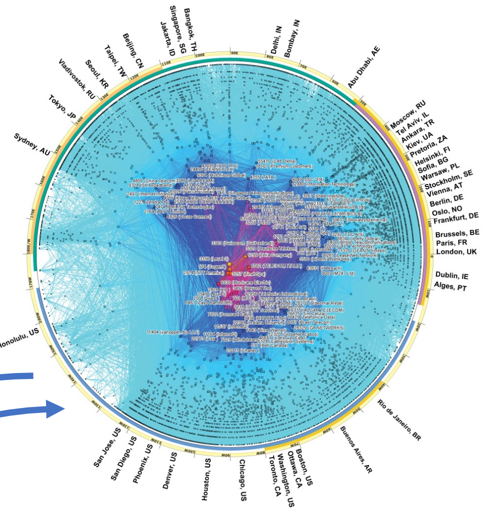

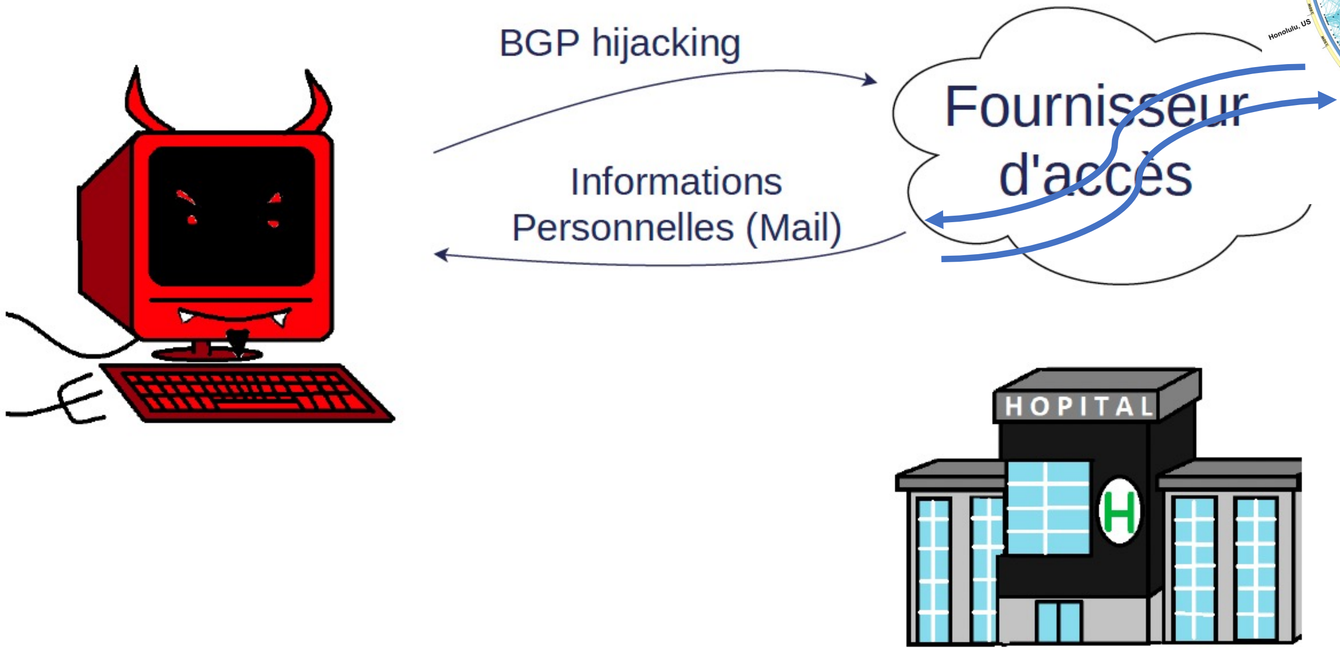

Hijacks can be used to divert traffic and gain inside knowledge

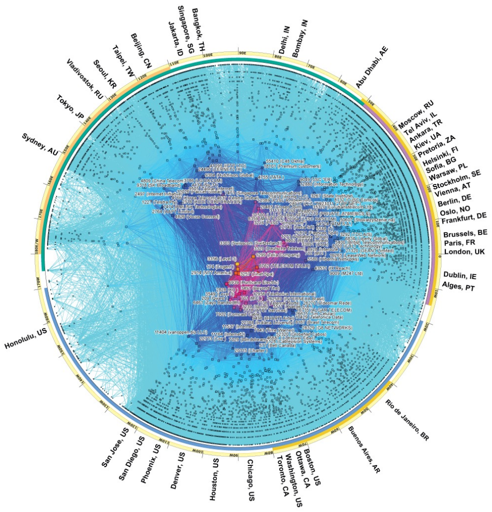

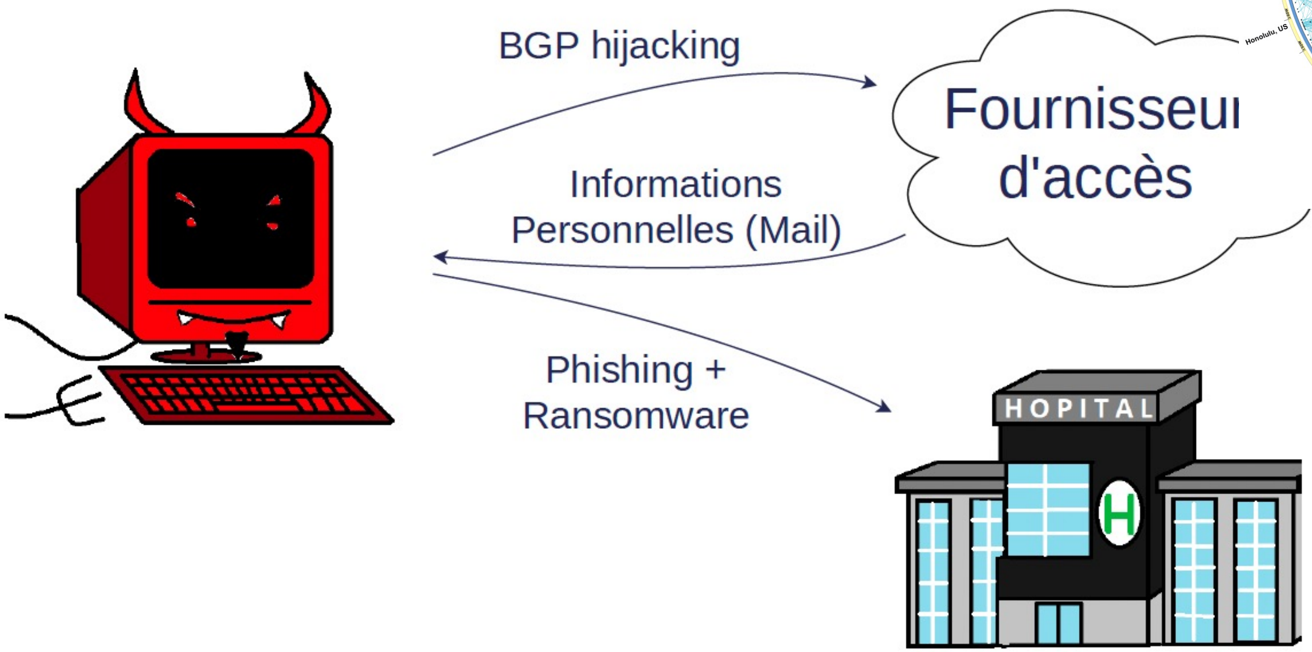

Hijacks can be used to divert traffic and gain inside knowledge

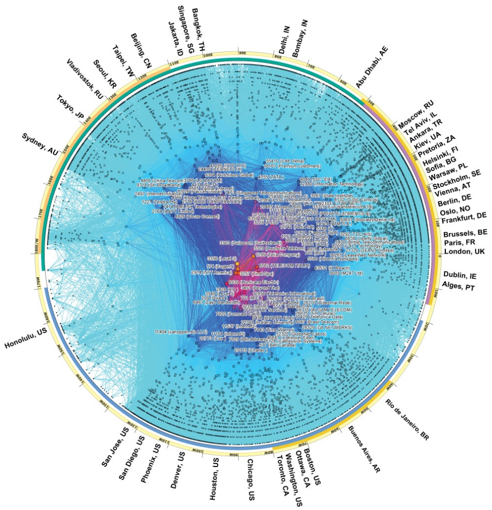

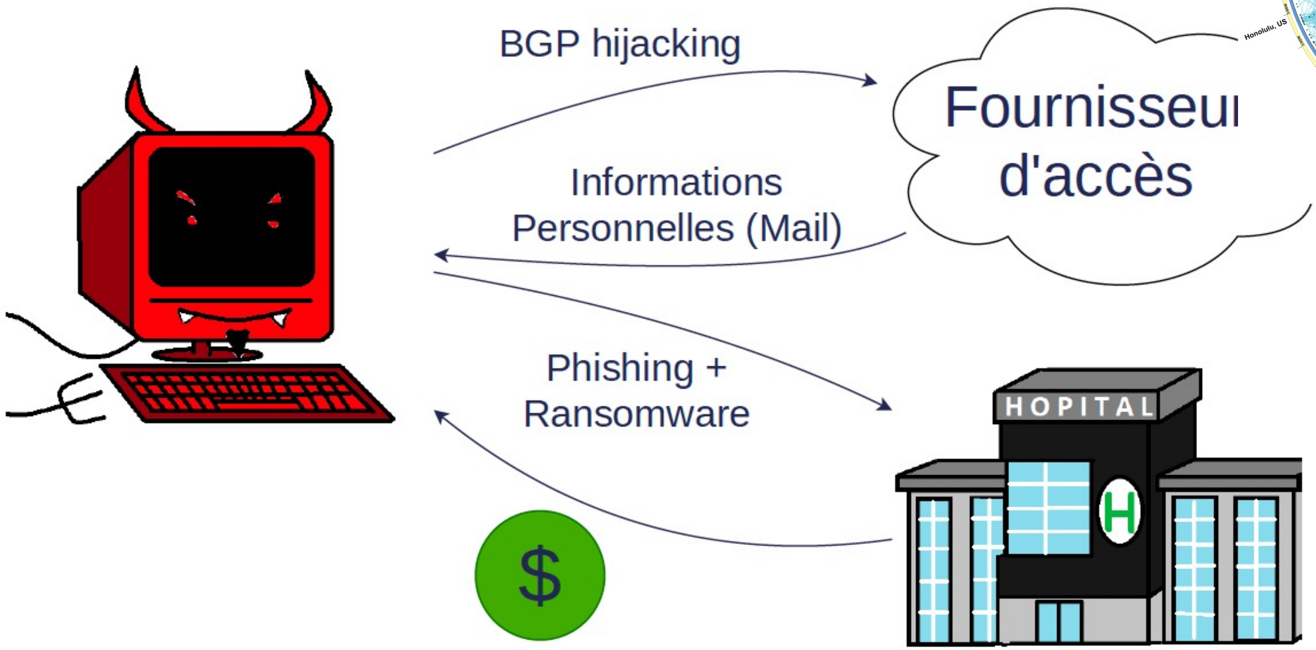

Multiple causes for hijacks

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

Hijacks are frequent

Some vulnerabilities of BGP

Prefix hijacks

Blackholing

BGP lies

BGP session injection

The purpose of blackholing is to protect against DDoS

DDoS are frequent

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

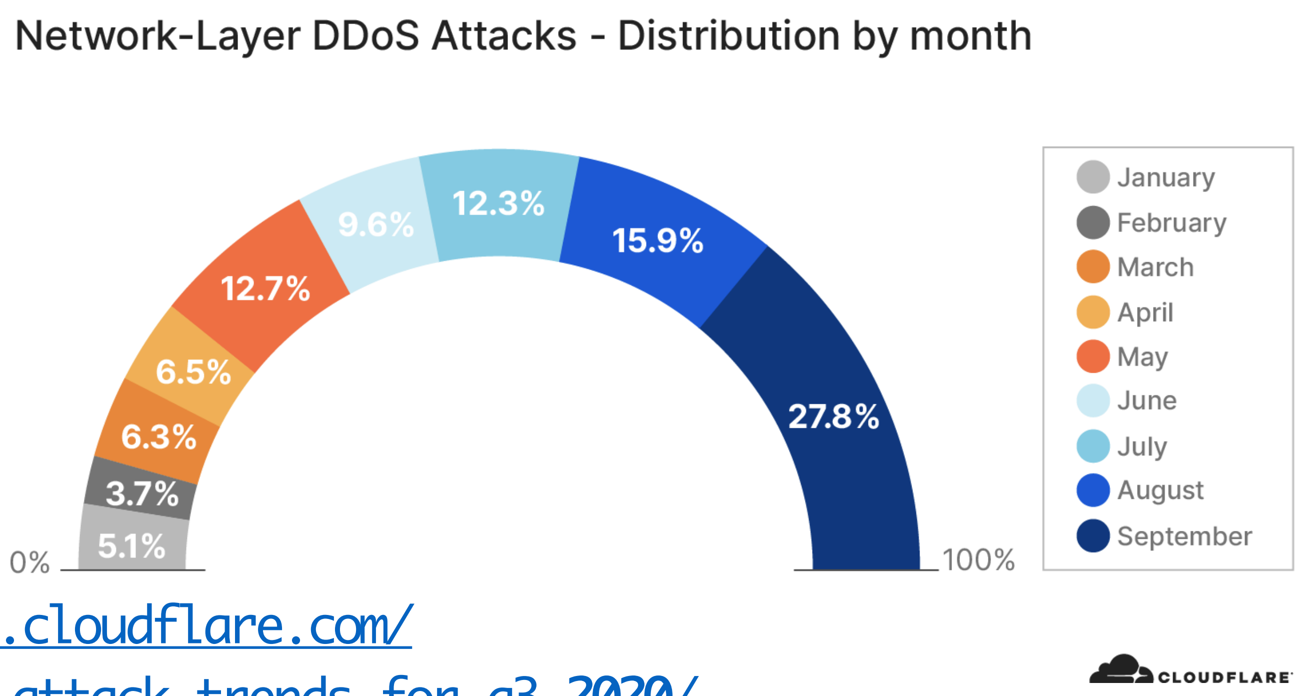

Source: https://blog.cloudflare.com/network-layer-ddos-attack-trends-for-q3-2020/

DDoS

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

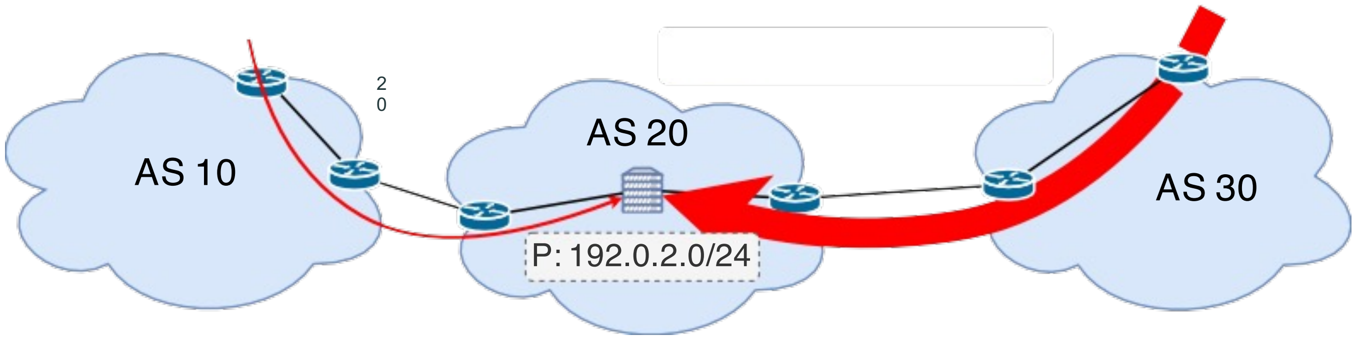

BGP blackholing

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

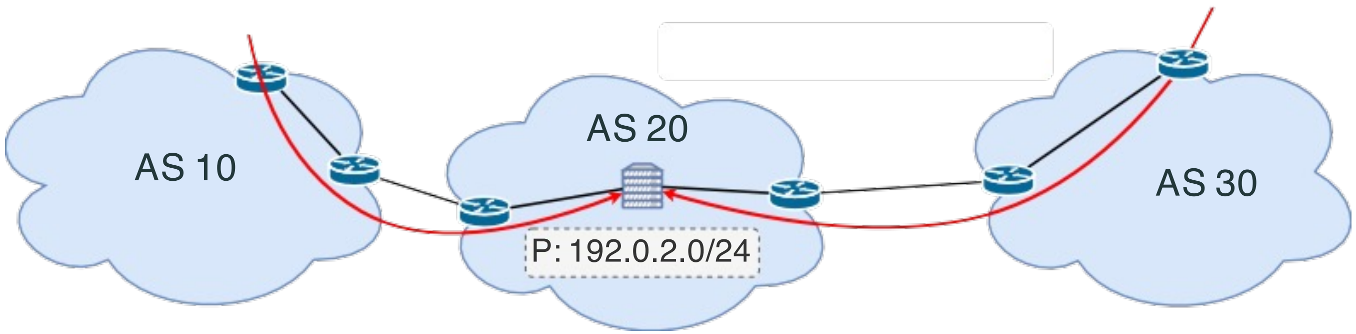

BGP blackholing

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

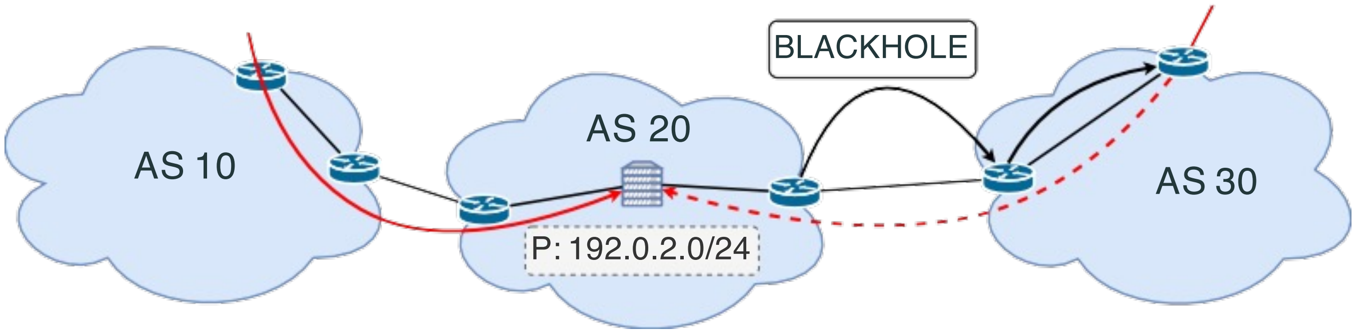

Blackholing has a double-edged sword effect: **all** traffic is dropped.

BGP blackholing

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

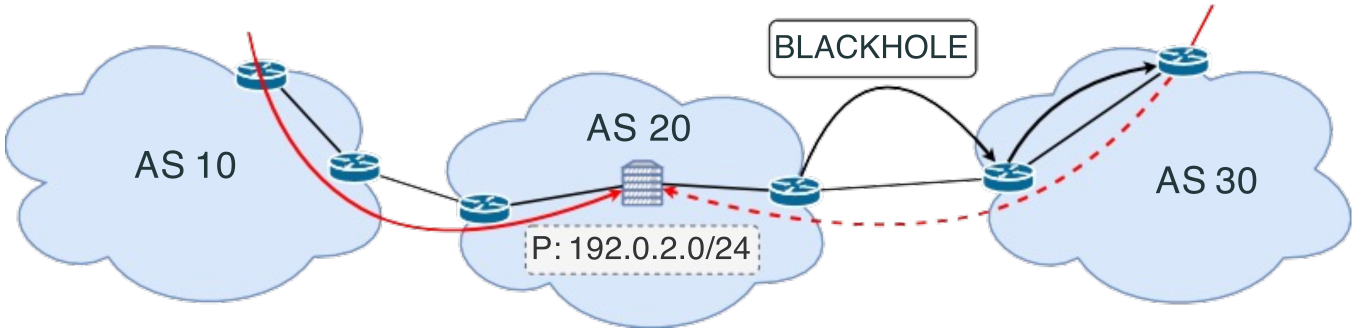

Blackholing is announced via what is called a BGP community.

BGP community usage is increasing

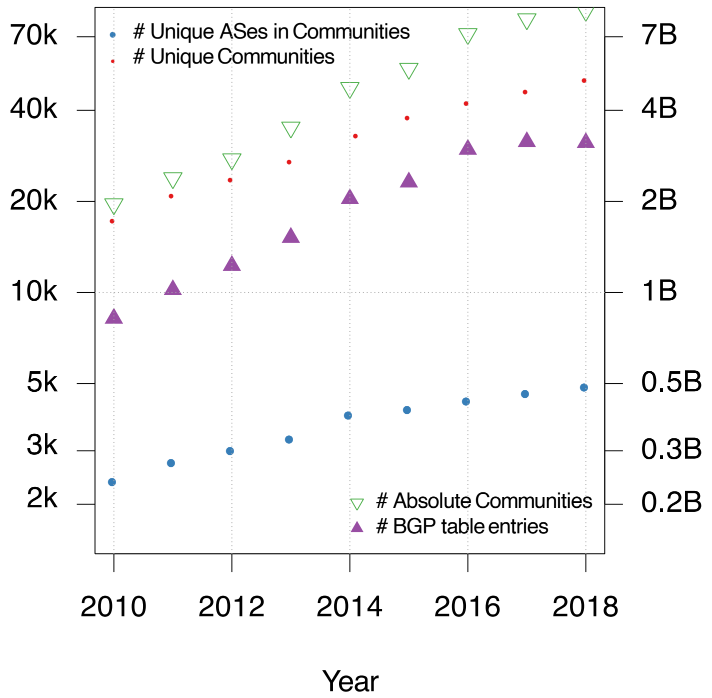

Increasing usage warrants a closer look.

BGP community usage is increasing

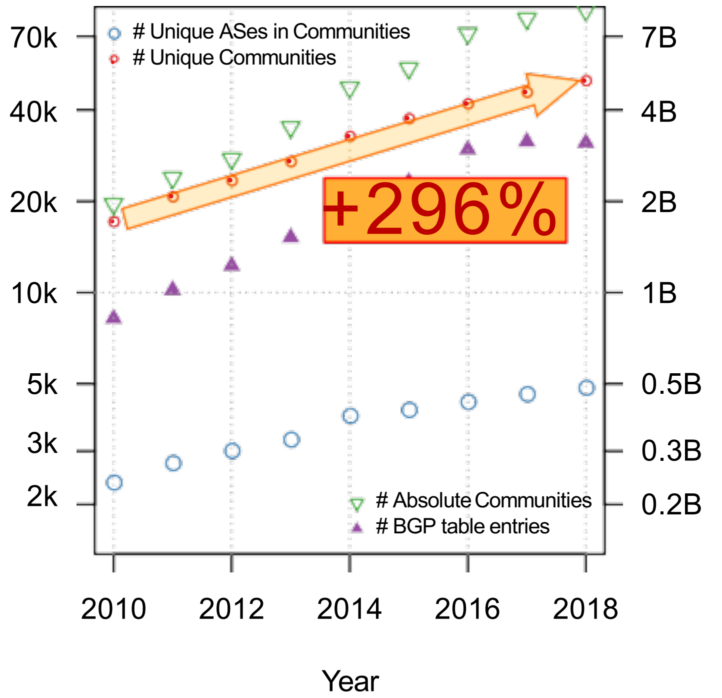

Increasing usage warrants a closer look.

BGP Communities (RFC 1997)

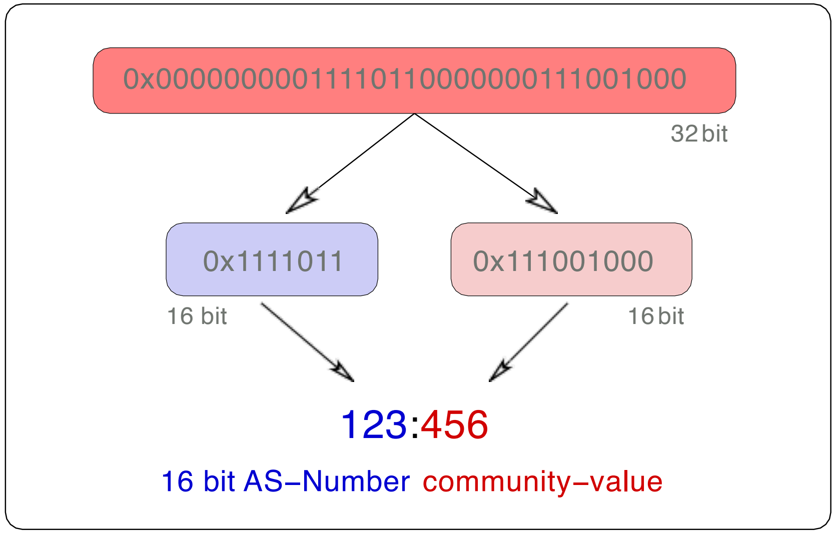

By convention written *ASN:VALUE*

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 enough?

Example topology

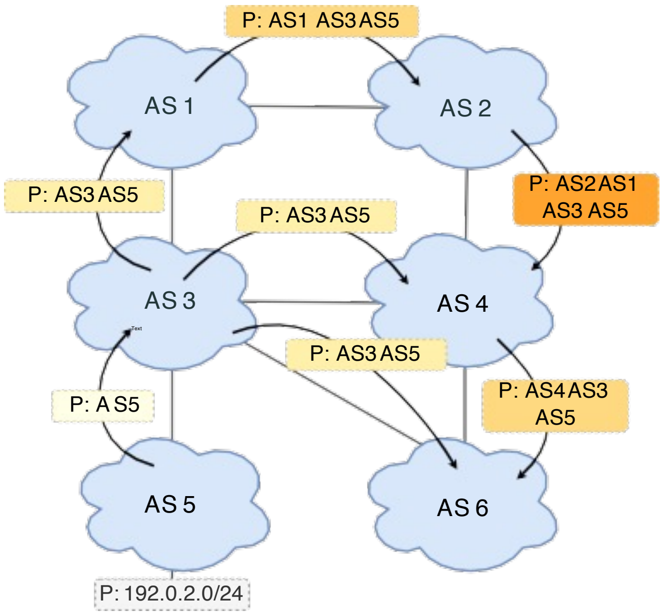

BGP update propagation

Example topology

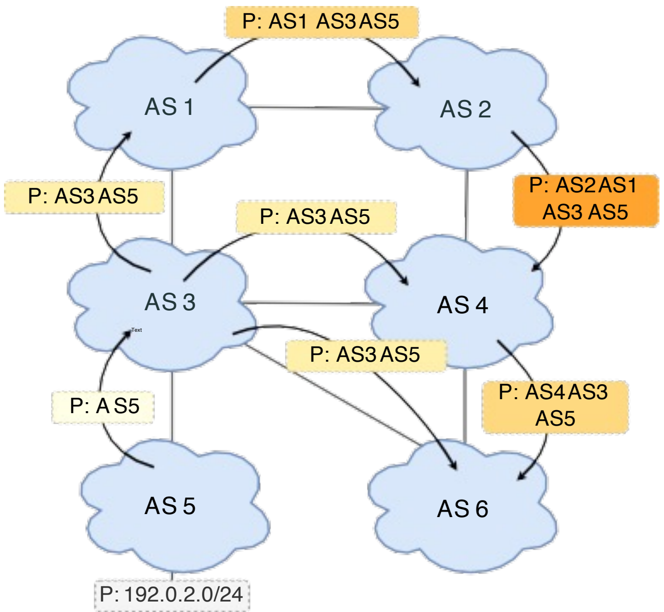

BGP update propagation

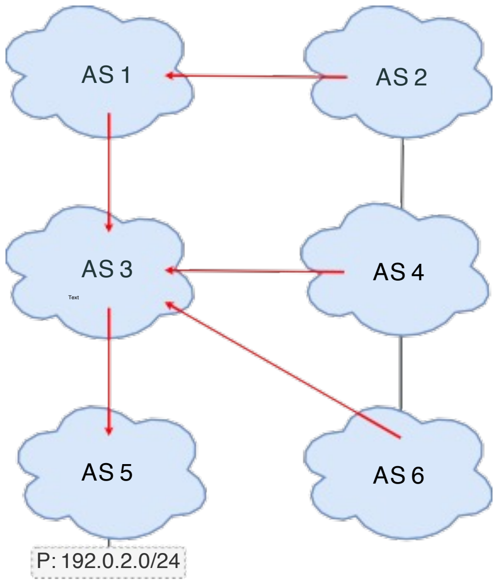

Traffic flow

Example topology

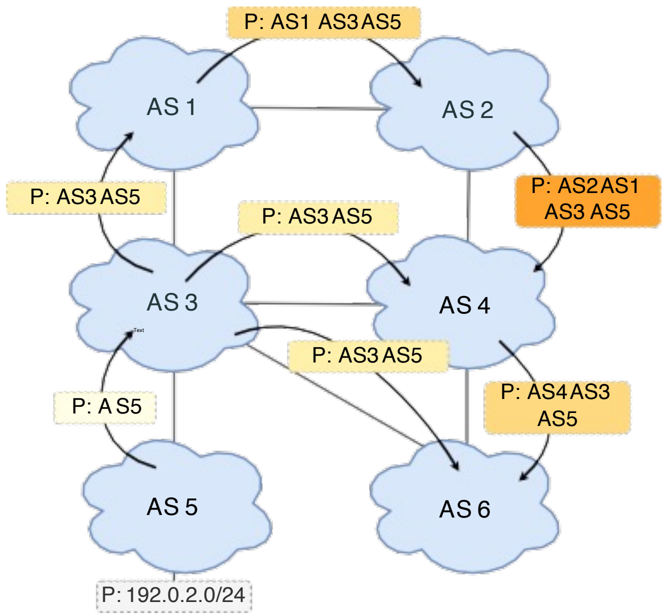

BGP update propagation

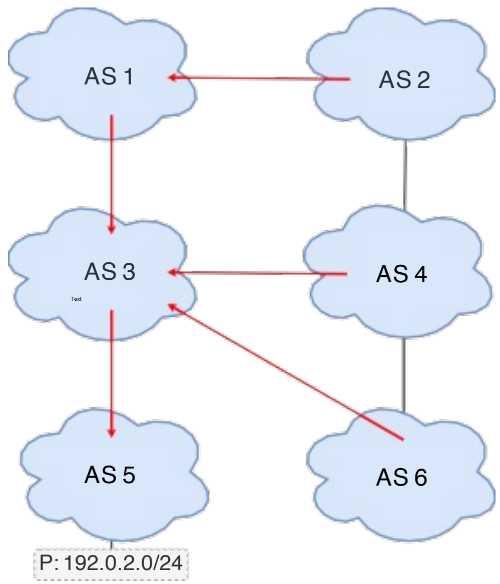

Traffic flow

BGP policies make AS2 not learn the path via AS4

Example topology

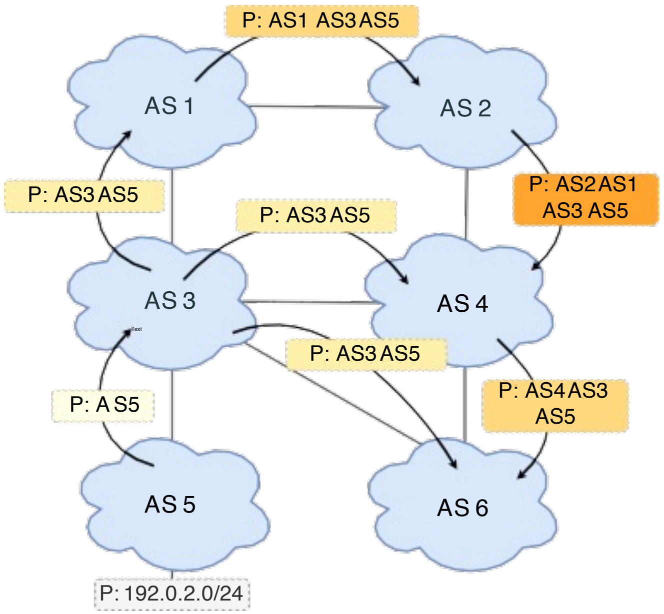

BGP update propagation

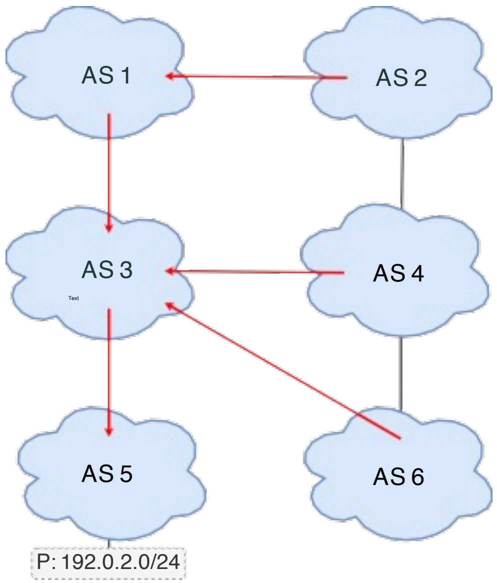

Traffic flow

BGP policies make AS2 not learn the path via AS4

BGP policies are distributed in the AS using BGP communities

Example topology

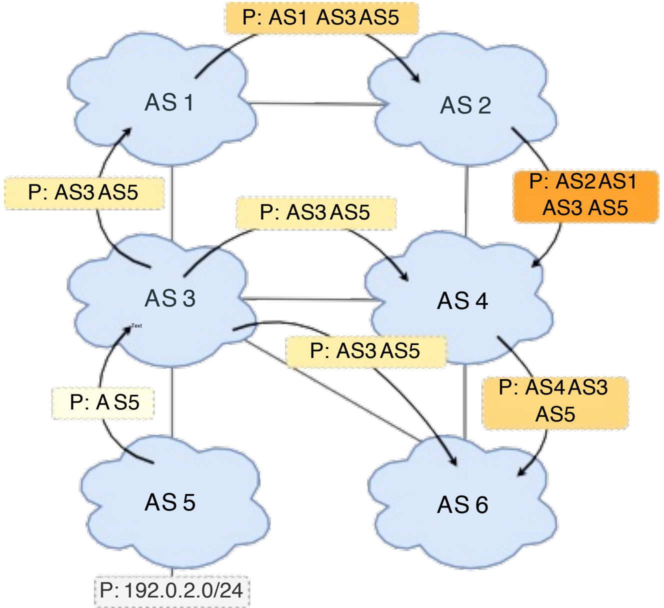

BGP update propagation

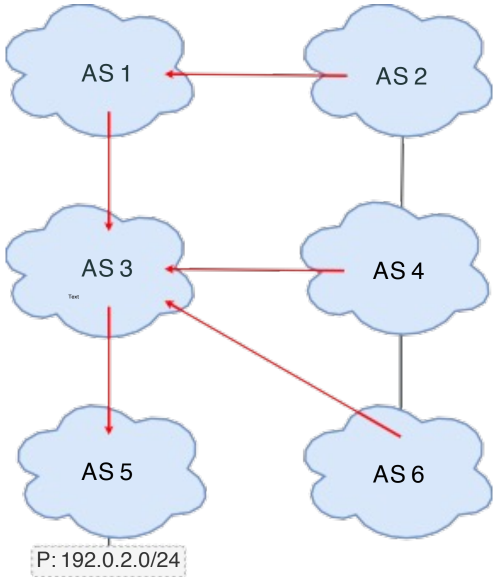

Traffic flow

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)

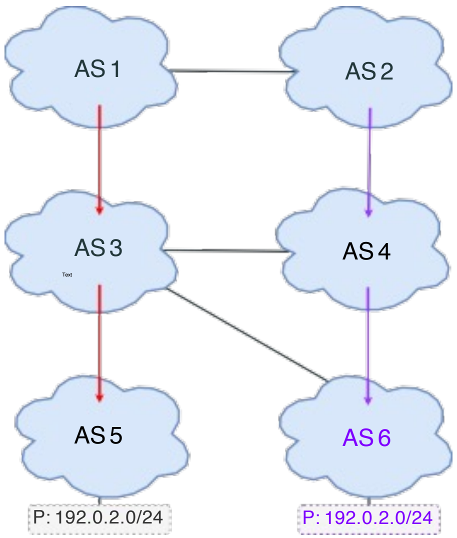

Hijack type-0

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

Hijack-0 and Blackjack-0

Sermpezis 2018 (Artemis)

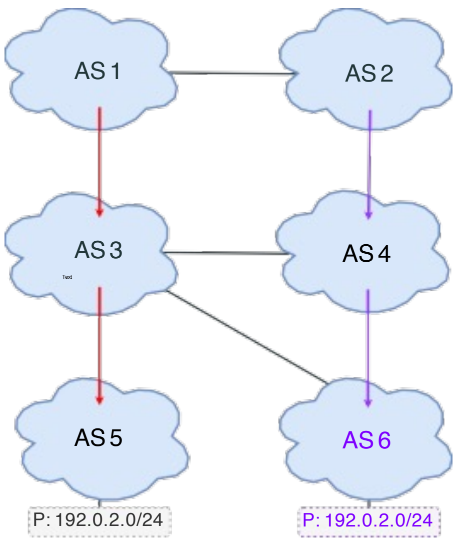

Hijack type-0

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

Miller et Pelsser 2019

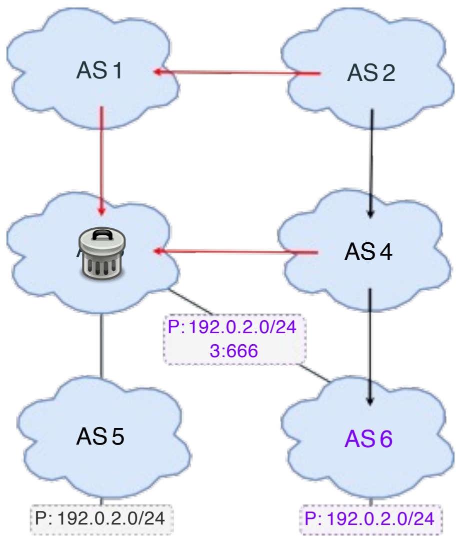

Blackjack type-0

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

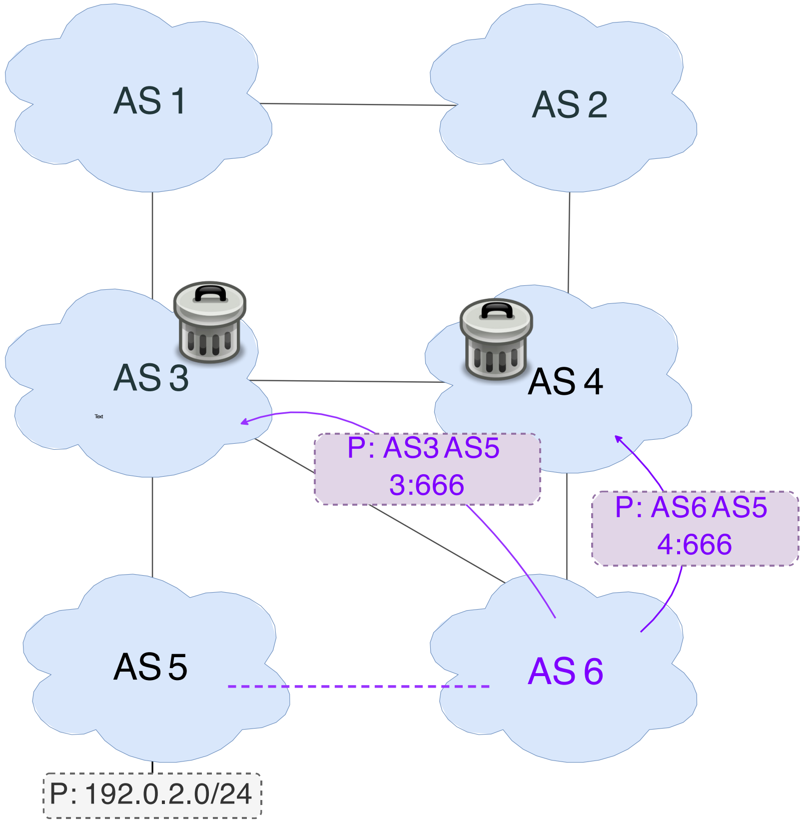

The origin AS is legit. The AS-path is not.

BGPsec[5]

BGPsec allow 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

Blackholing

BGP lies

BGP session injection

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

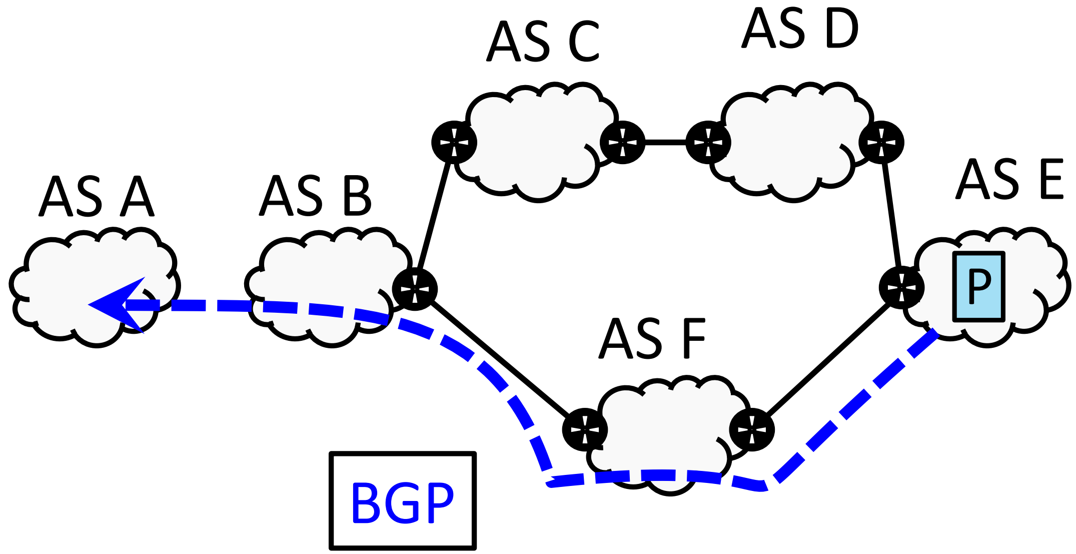

An ISP (AS B) announces a path in BGP but forwards packets along a different path

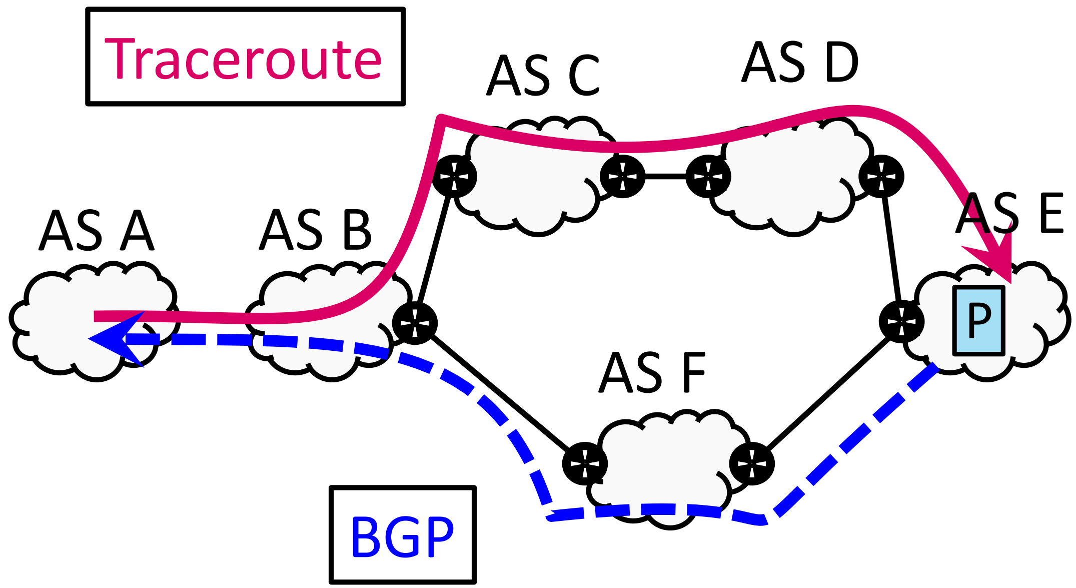

An ISP (AS B) announces a path in BGP but forwards packets along a different path

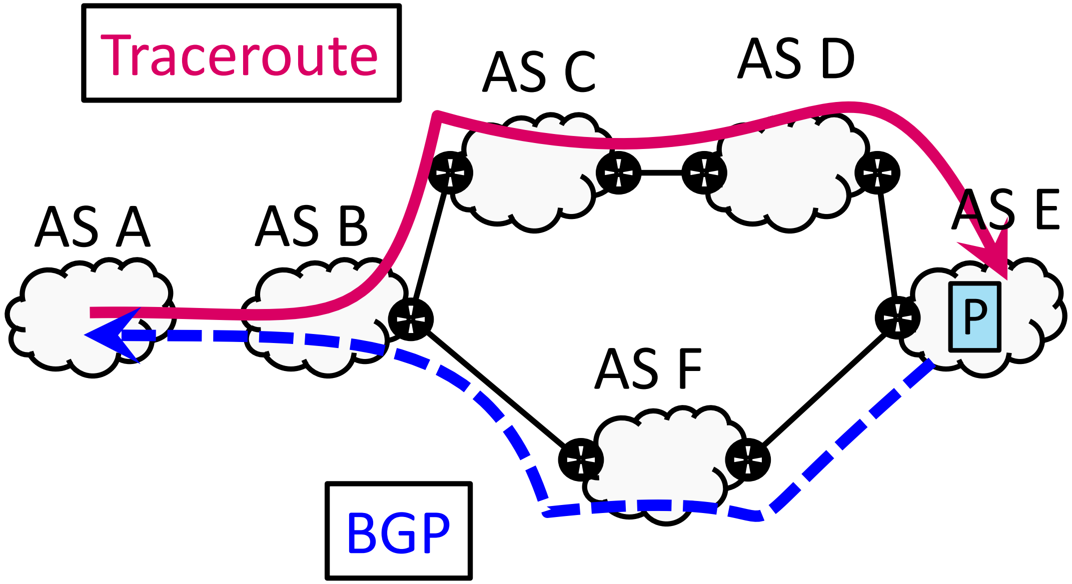

Because the peer C is cheaper

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

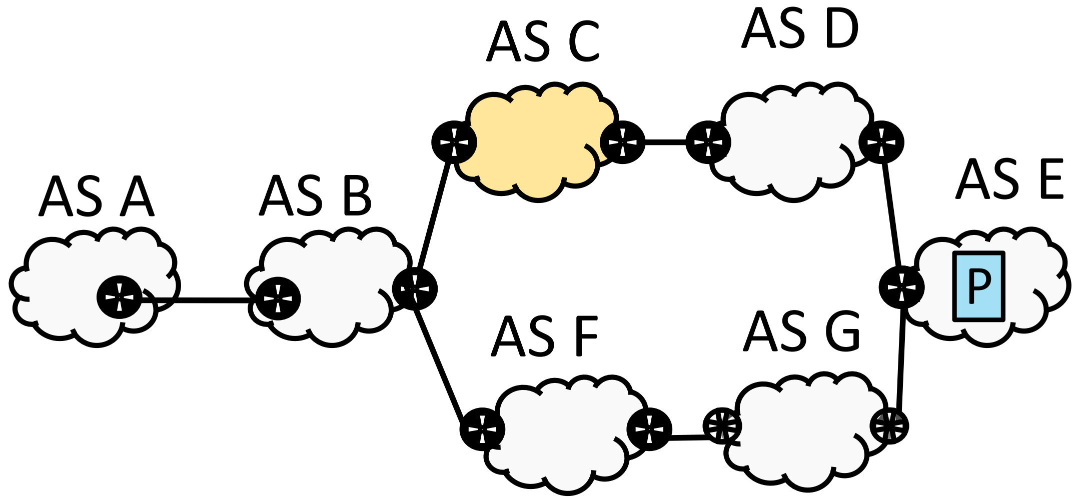

This difference in control and data paths may also be observed in the Kapela-Pilosov BGP monkey-in-the-middle attack

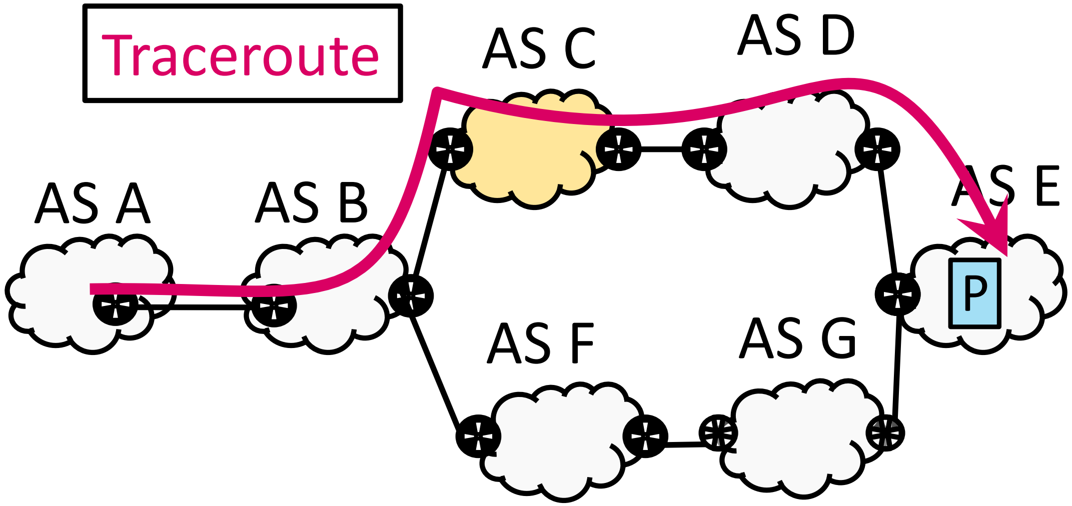

This difference in control and data paths may also be observed in the Kapela-Pilosov BGP monkey-in-the-middle attack

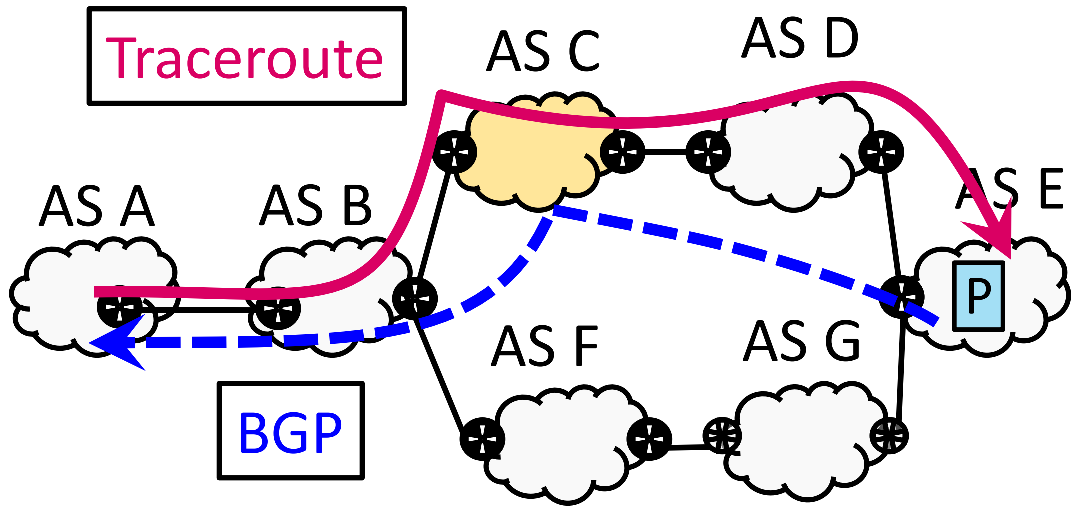

But for packets to follow the traceroute path, the yellow AS faked a direct link to the prefix origin

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

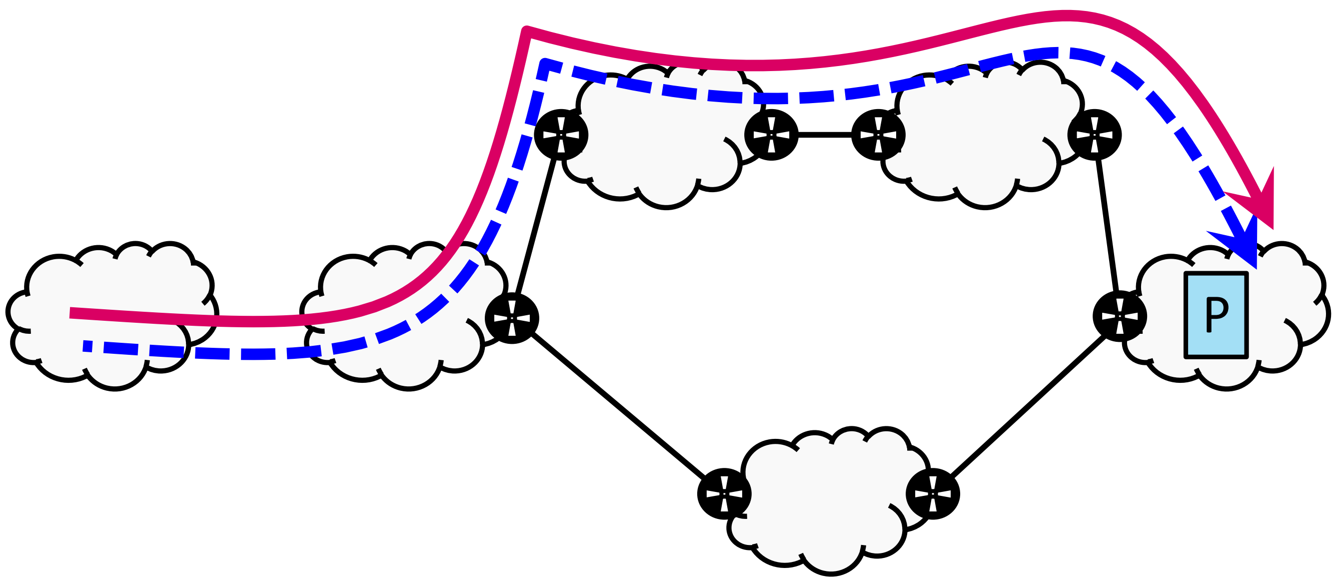

One form of BGP lie is

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

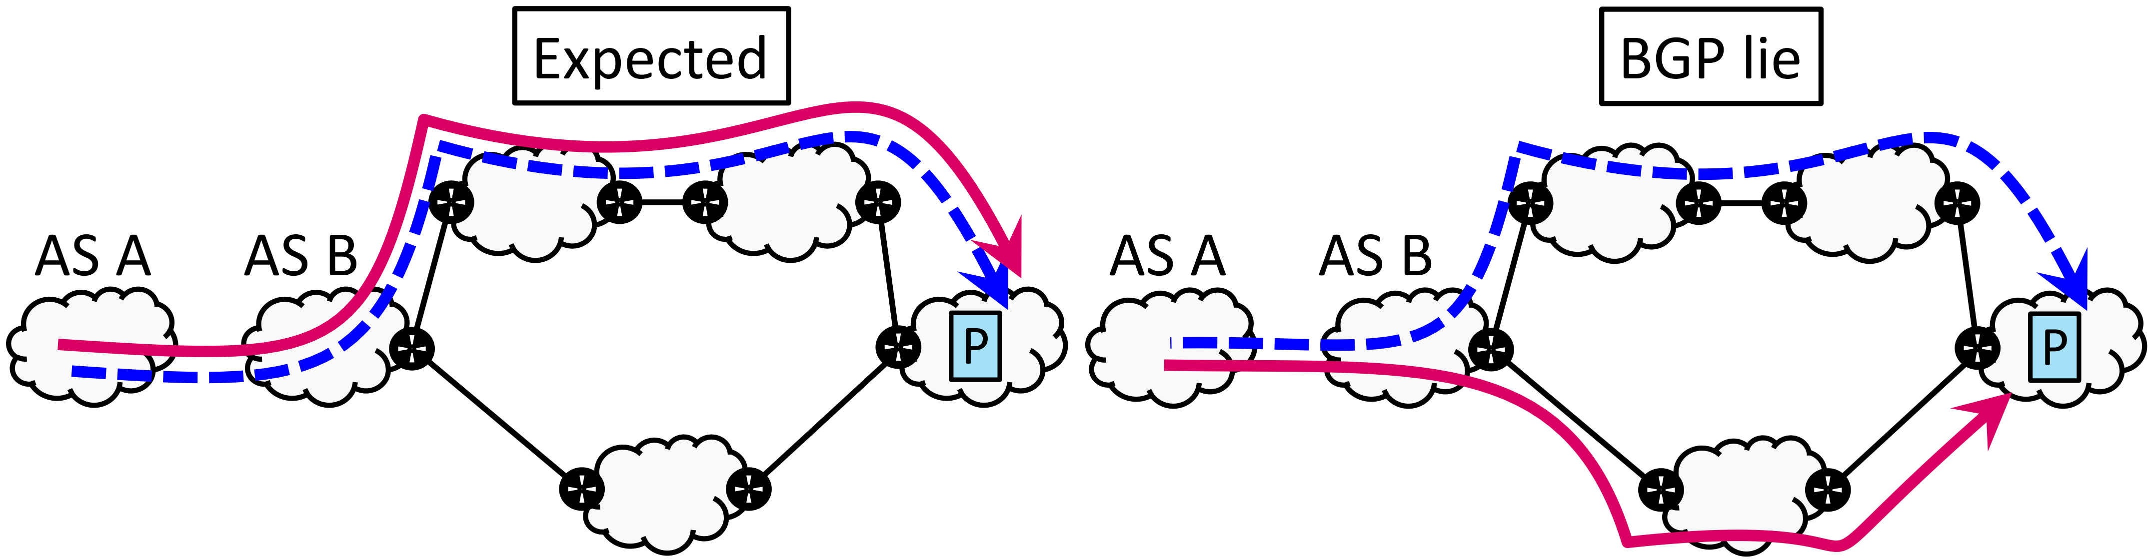

Some vulnerabilities of BGP

Prefix hijacks

Blackholing

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
  - What is the adoption status of TCP Authentication Option (TCP-AO) for BGP?

Some vulnerabilities of BGP

Prefix hijacks

Blackholing

BGP lies

BGP session injection

⇒BGP designed with no security in mind

Weak authentication

No integrity protection

How we may hack to live with these vulnerabilities

Prevention

- 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).

MANRS ROA Stats Tool

HOME BY COUNTRY BY ASN Problem with the data? MANRS

MANRS ROA Stats Tool

Search for ROA stats by country or ASN using the links above

Data last retrieved 1 day(s) ago

IPv4		
Total	922799	
Valid ROAs	357837	38.78%
Unknown ROAs	556469	60.3%
Invalid ROAs	8493	0.92%

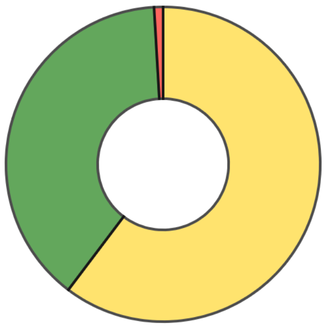

IPv6		
Total	139695	
Valid ROAs	61793	44.23%
Unknown ROAs	76338	54.65%
Invalid ROAs	1564	1.12%

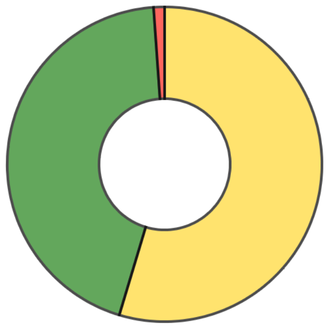

https://roa-stats.manrs.org (October 6, 2022)

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

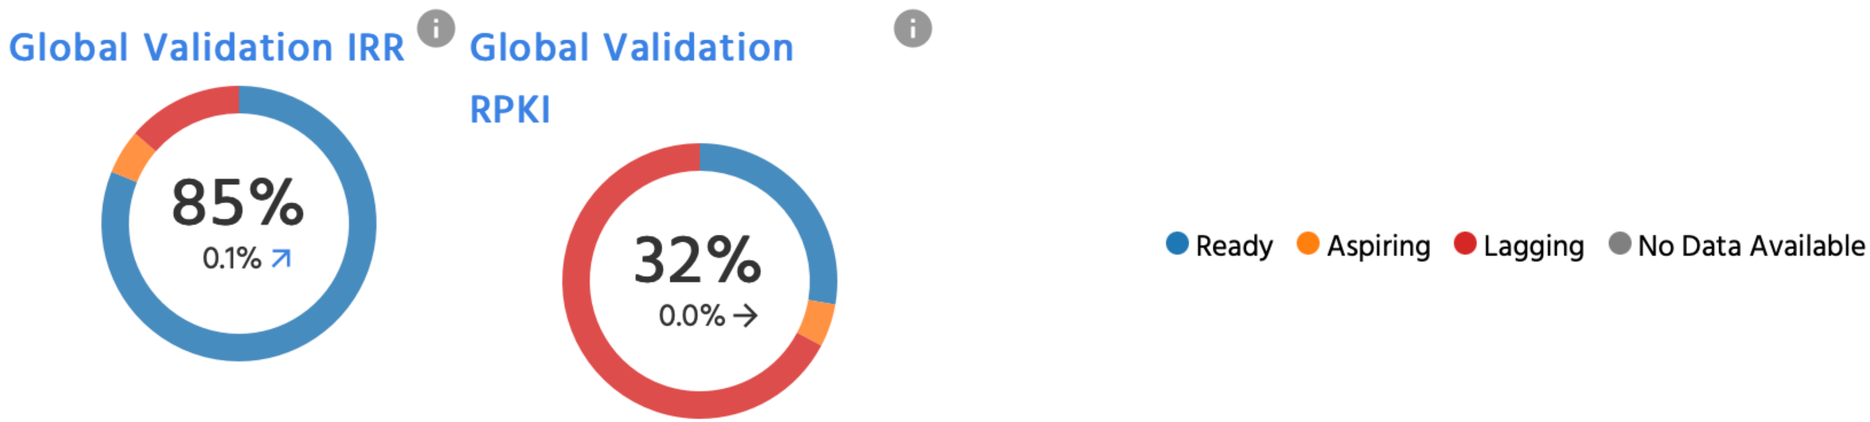

From https://observatory.manrs.org/#/overview (Oct. 6, 2022)

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

Deployement of protection increases but events still occur (NL)

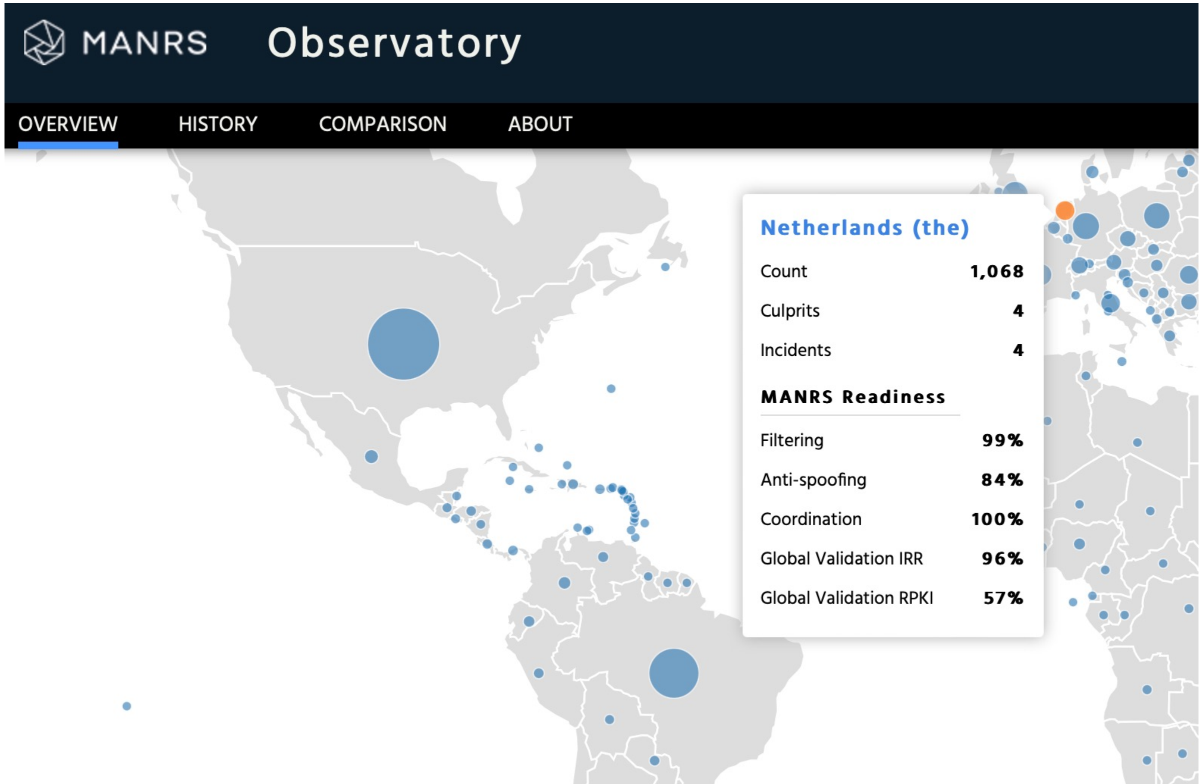

Deployement of protection increases but events still occur (BR)

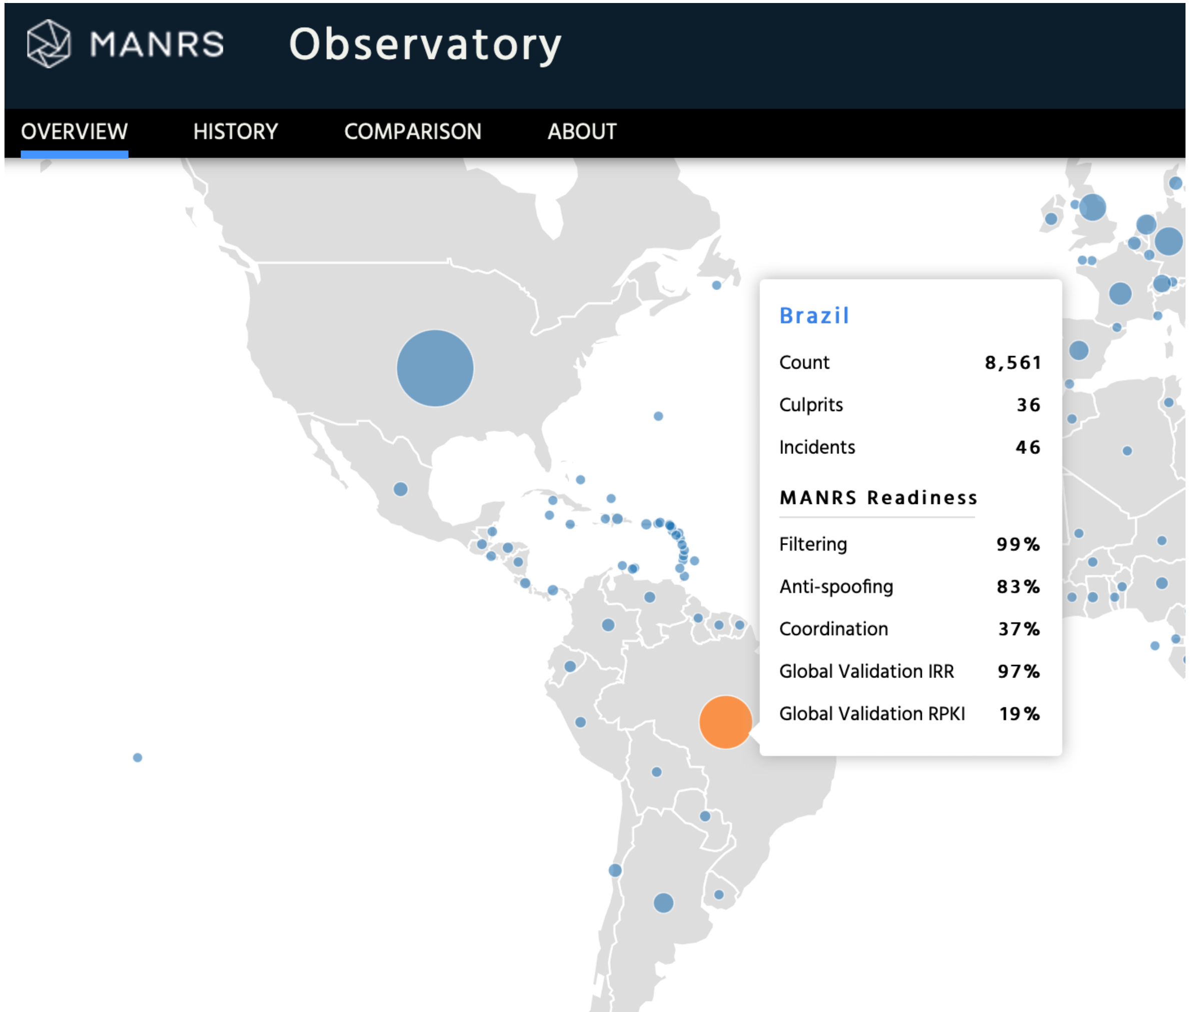

Detection

- BGP lies

Detection of BGP lies

Required data

Control paths

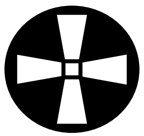

P	CP
P_Y	BCD
P_R	D
P_V	E

Data paths

Vantage Point (VP)

Traceroute per destination

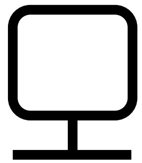

Issues to consider

- Space-synchronization
  - Measurement platform
- Address space and time synchronization
  - Which DP should be compared with which CP
- IP-to-AS mapping
  - CPs come as AS-paths but DPs as IP-paths

Issues to consider

- Space-synchronization
  - Measurement platform
- Address space and time synchronization
  - Which DP should be compared with which CP
- IP-to-AS mapping
  - CPs come as AS-paths but DPs as IP-paths

Space-synchronization

- Control paths are obtained from a given router
- Data paths are gathered from a VP
- To be comparable, DPs need to go through the router that shared the CPs

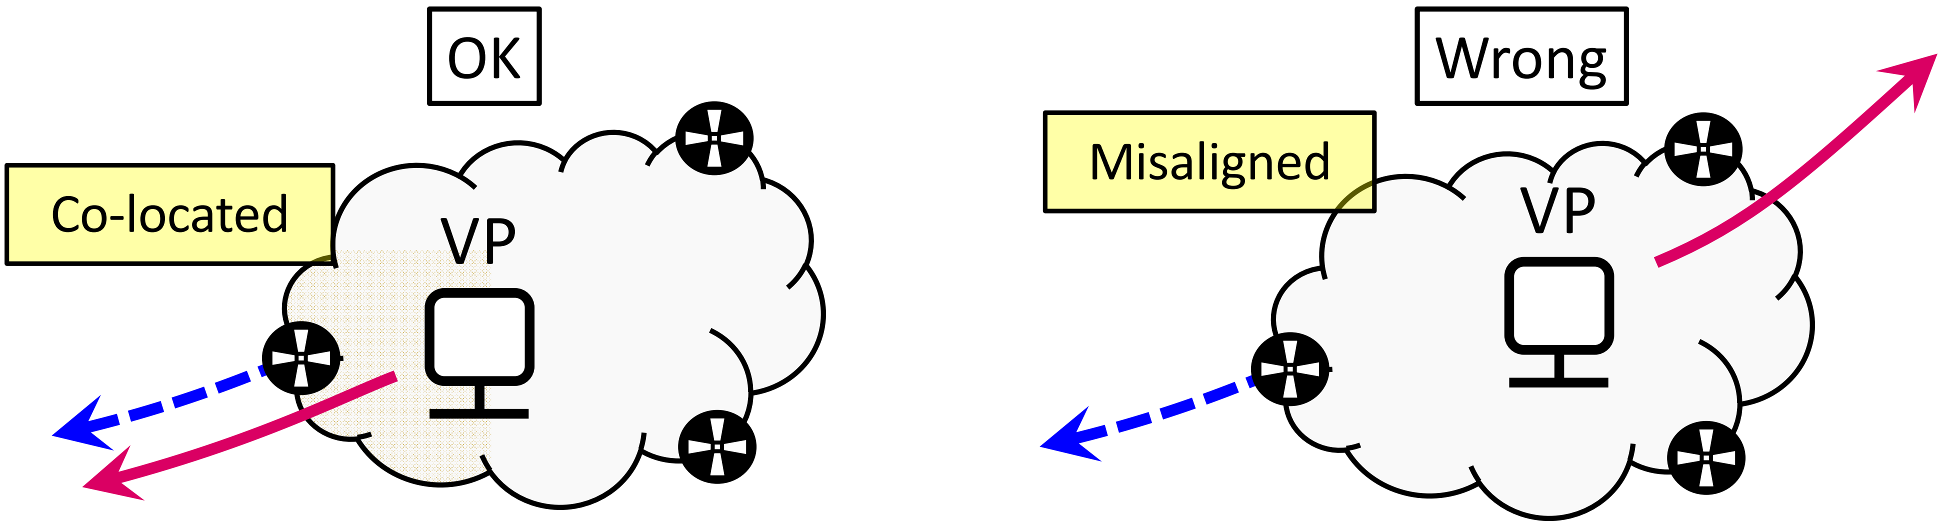

IP-to-AS mapping

- While CPs are AS-paths, DPs are obtained as IP-paths

```
CP: AS A, AS B, AS C…
DP: IP1, IP2, IP3, IP4…
```

To compare them, an IP-to-AS mapping tool is needed !

The problem of IP-to-AS mapping

Noise or sources of errors

- AS siblings

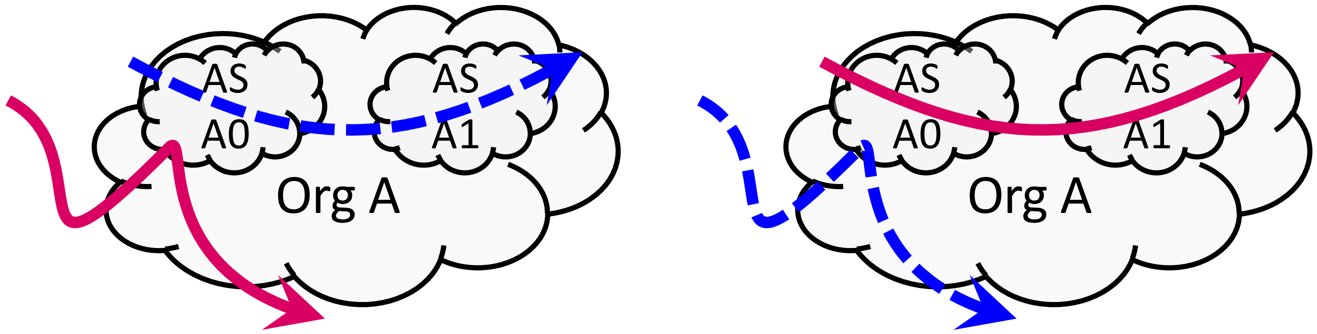

Noise or sources of errors

- AS siblings
- Third-party addresses

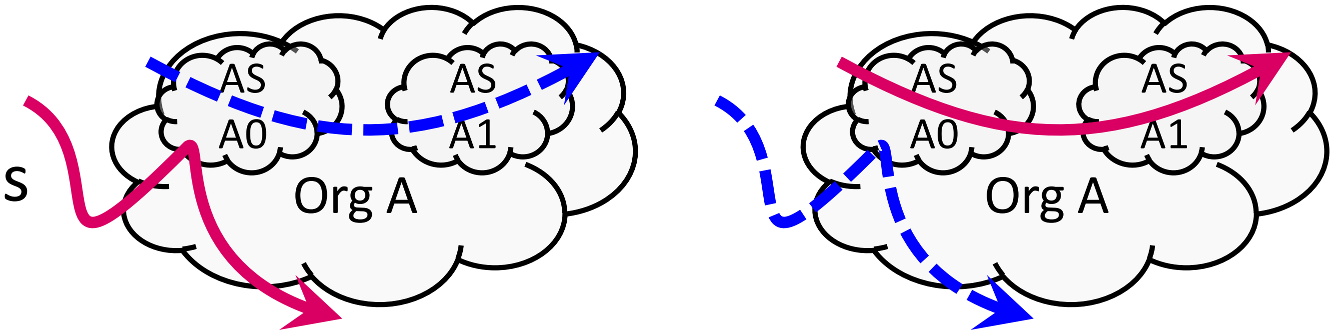

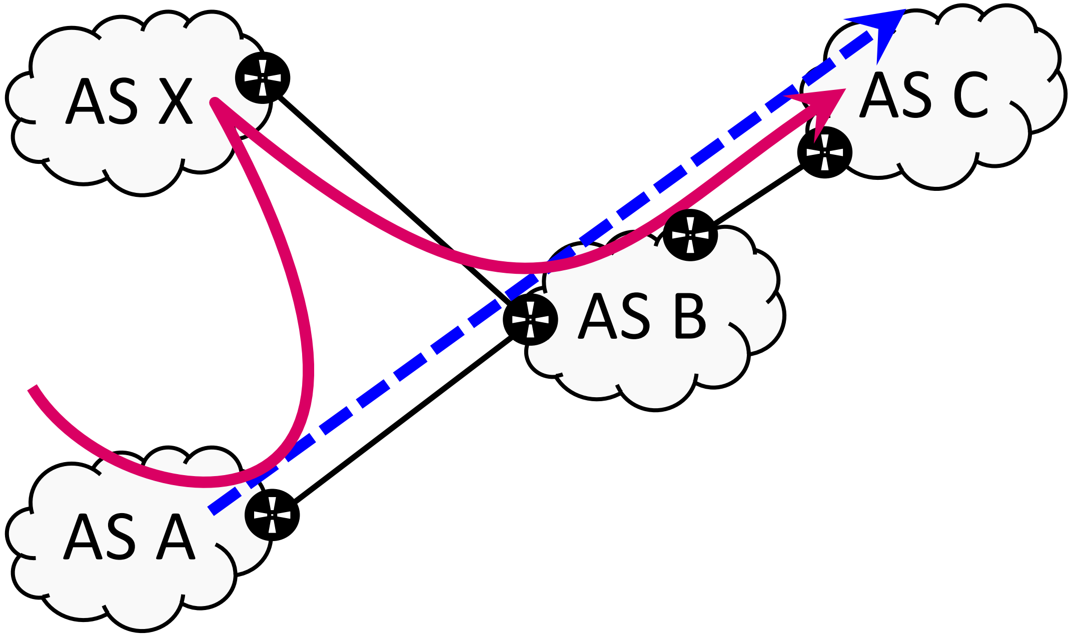

Noise or sources of errors

- AS siblings
- Third-party addresses
- Missing hops

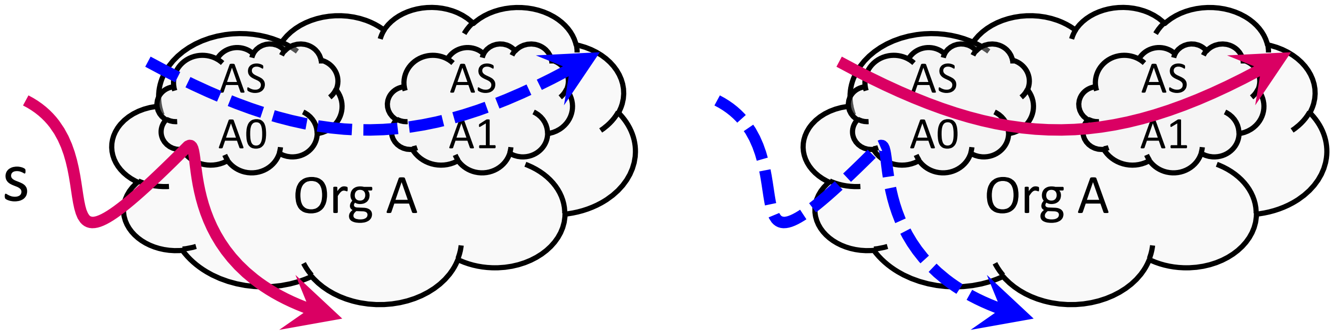

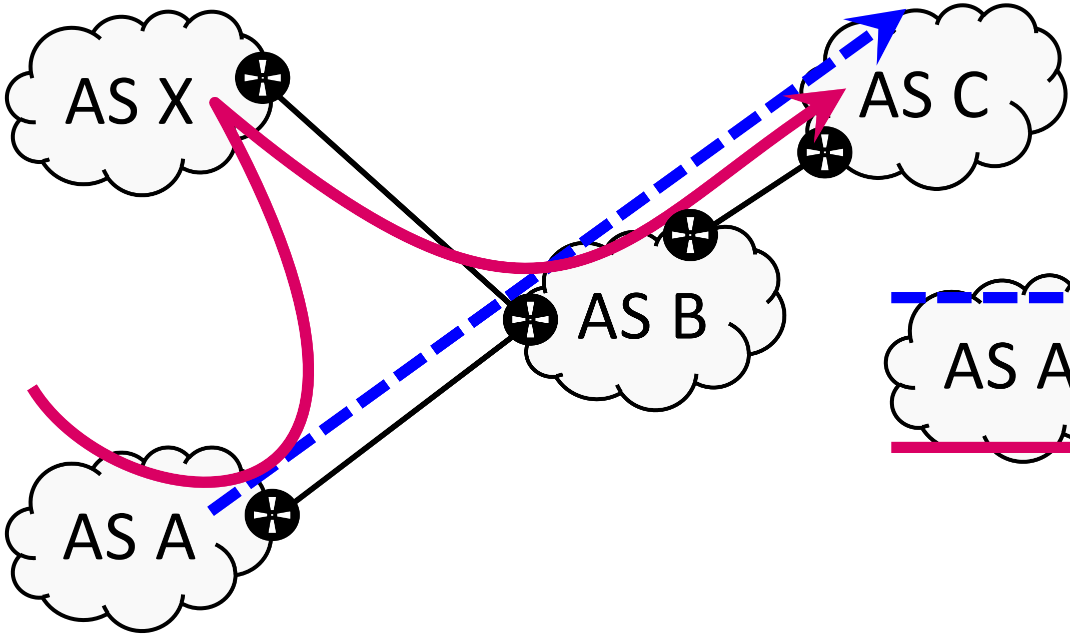

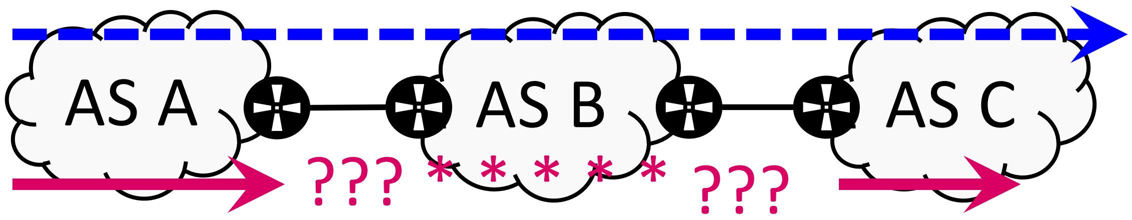

Our solution

A framework to detect BGP lies

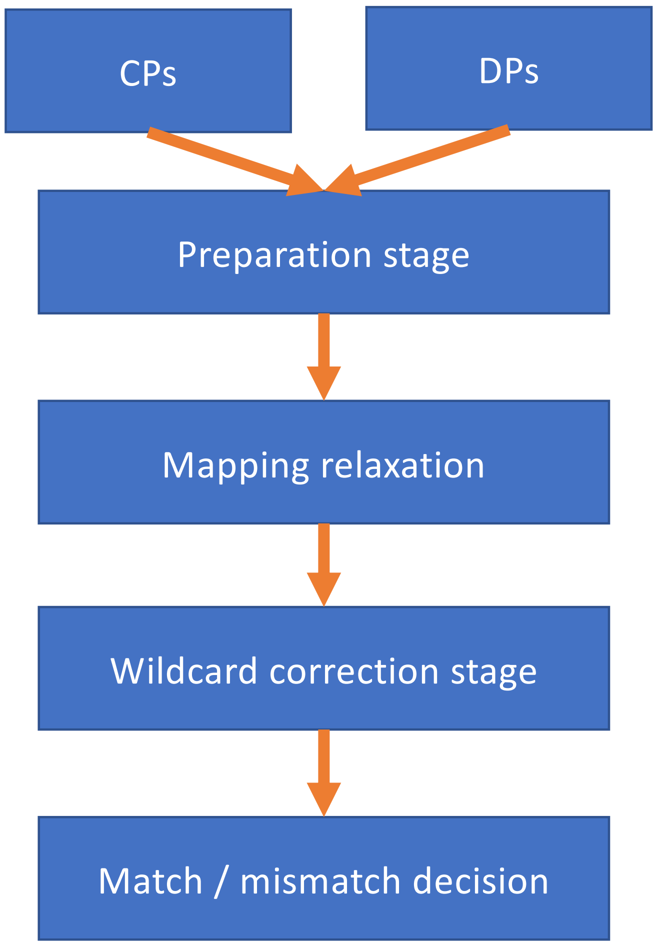

- ✓ **Input:** CPs and DPs from a co-located VP
- ✓ **Output:** rate of BGP lies

A framework to detect BGP lies

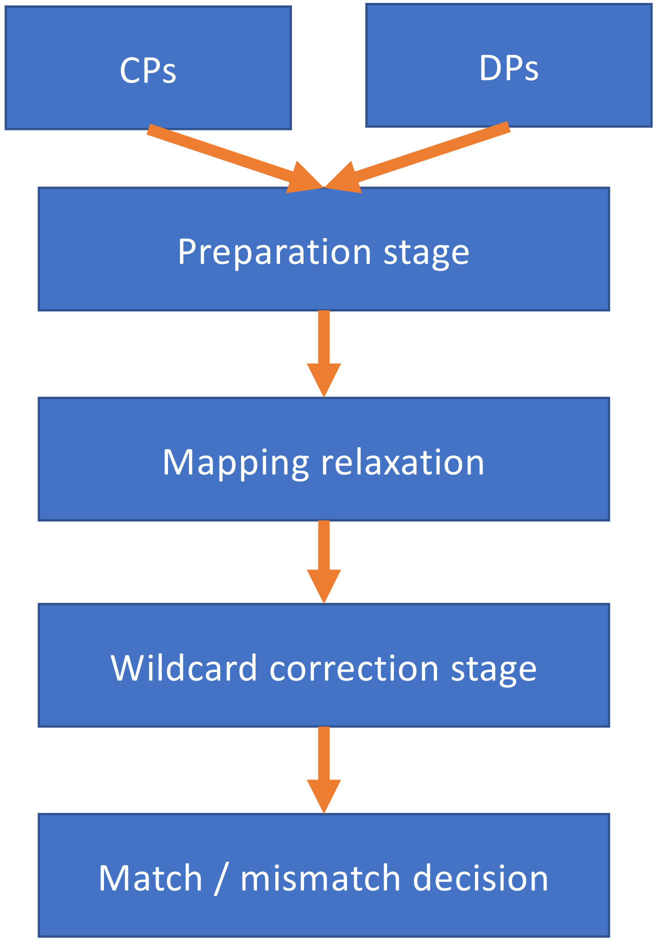

✓ **Input:** CPs and DPs from a co-located VP
✓ **Output:** rate of BGP lies

❑ **Preparation stage:**
- Address space synchronization
- Time synchronization
- Basic IP-to-AS mapping

❑ **Mapping relaxation**
- AS siblings
- Third-party addresses

❑ **Wildcards correction stage**
- Missing hops

A framework to detect BGP lies

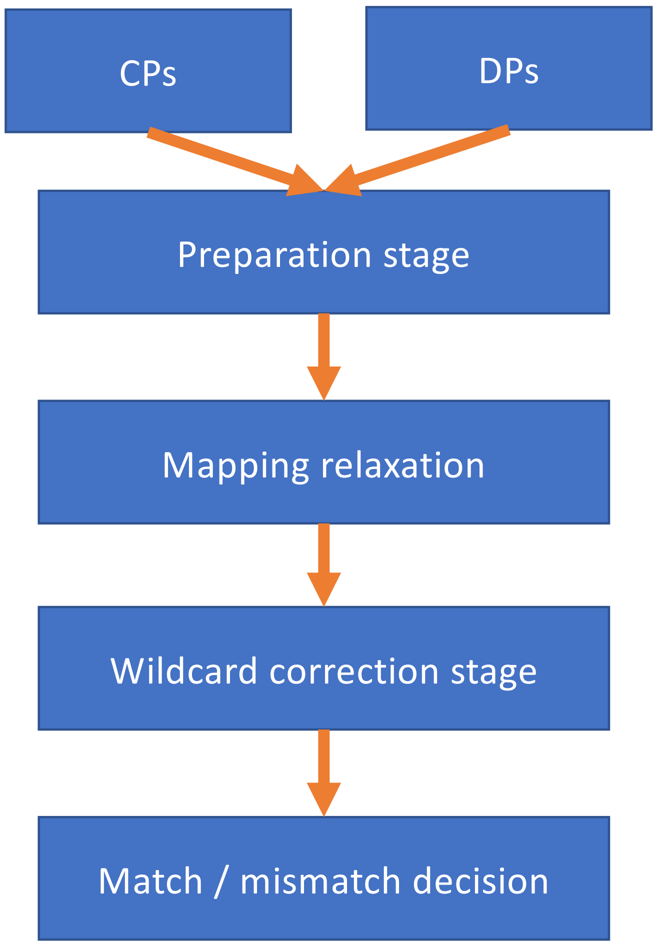

- ✓ **Input:** CPs and DPs from a co-located VP
- ✓ **Output:** rate of BGP lies

- ❑ **Preparation stage:**
  - Address space synchronization
  - Time synchronization
  - Basic IP-to-AS mapping
- ❑ **Mapping relaxation**
  - AS siblings
  - Third-party addresses
- ❑ **Wildcards correction stage**
  - Missing hops

...we are conservative!

Our measurements

Experiment setup

- Deployed 8 co-located VPs
- CPs collected every two hours
- DPs gathered targeting 80K destinations per day
- We run measurements multiple days (at least 13 days)

Low number of mismatches for most vantage points but they exist

At VP 7, the high number of ”lies” is due to partial forwarding tables in the provider AS.

These partial tables also create detours in the provider.

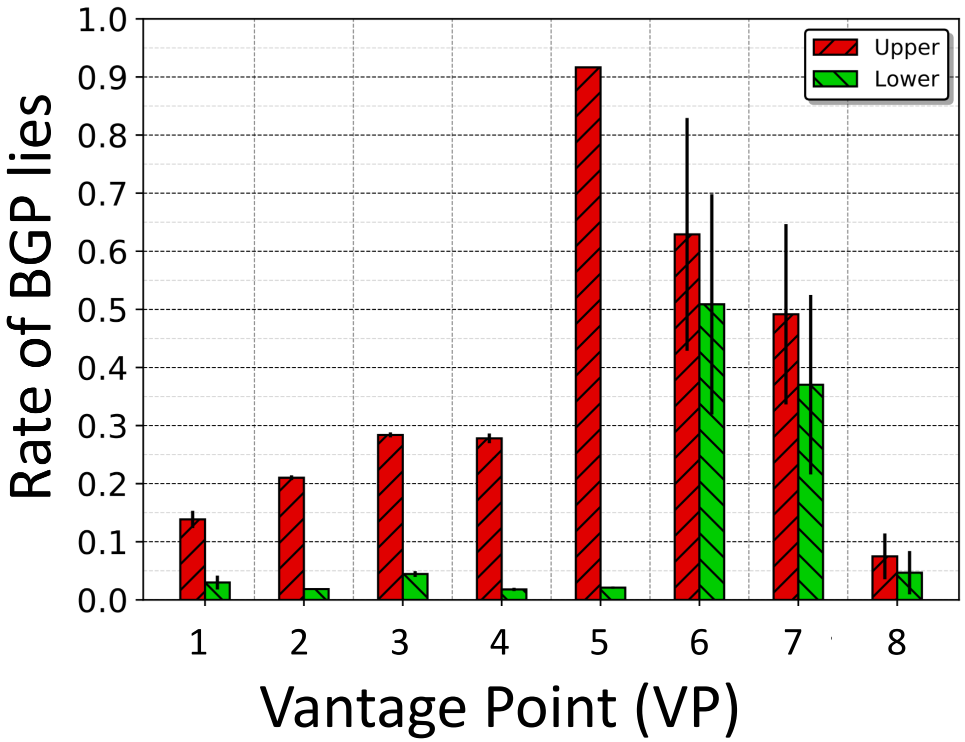

Related publications

- 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.

To conclude

Still a lot progres to be made in detection

- Because IP allocations, AS level, IP level topologies are not fully known we rely on heuristics to determine what is legit.
- We have collectors for BGP data and measurement platforms for traceroutes but the data is biased and redundant.
  - We work on methods to select the collection points for good coverage with reduced redundancy for BGP

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.

And prevention is not a given

- It hardens operations, lengthen the feedback loop

Thank you!

Special thanks to my collaborators