ARTEMIS: Neutralizing BGP Hijacking within a Minute

Alberto Dainotti
alberto@caida.org
Center for Applied Internet Data Analysis
University of California, San Diego

Joint work with:
Pavlos Sermpezis, Vasileios Kotronis, Petros Gigis, Xenofontas Dimitropoulos, Danilo Cicalese, Alistair King

INTERNET ROUTE HIJACKING

a threat to your organization and to critical infrastructure

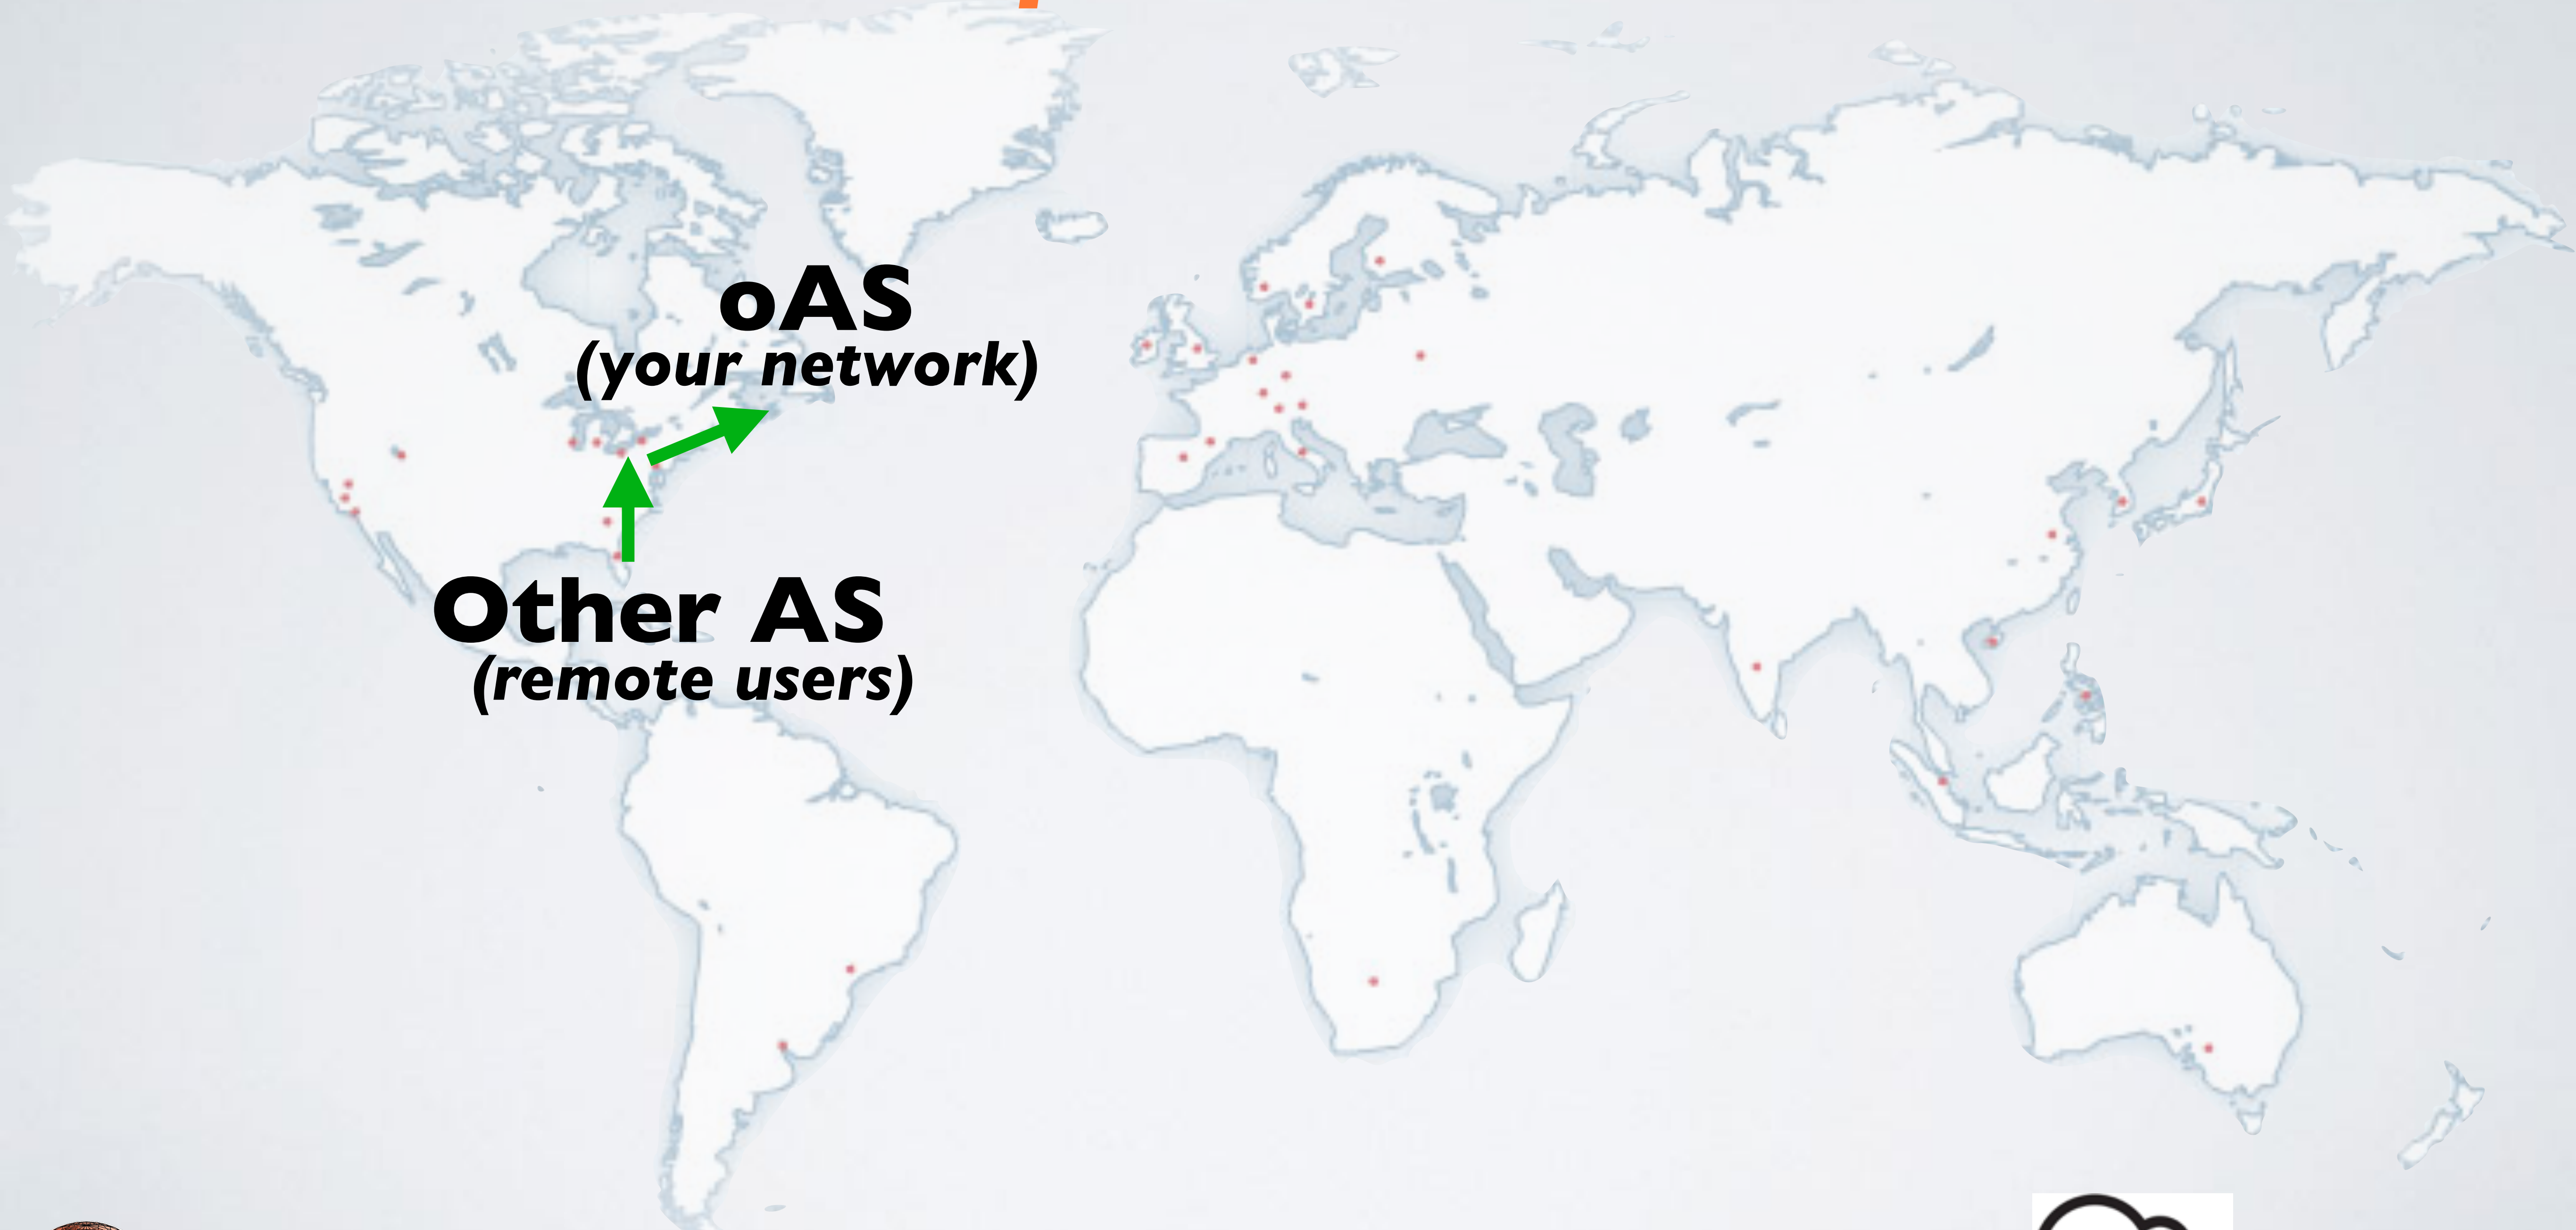

INTERNET ROUTE HIJACKING

a threat to your organization and to critical infrastructure

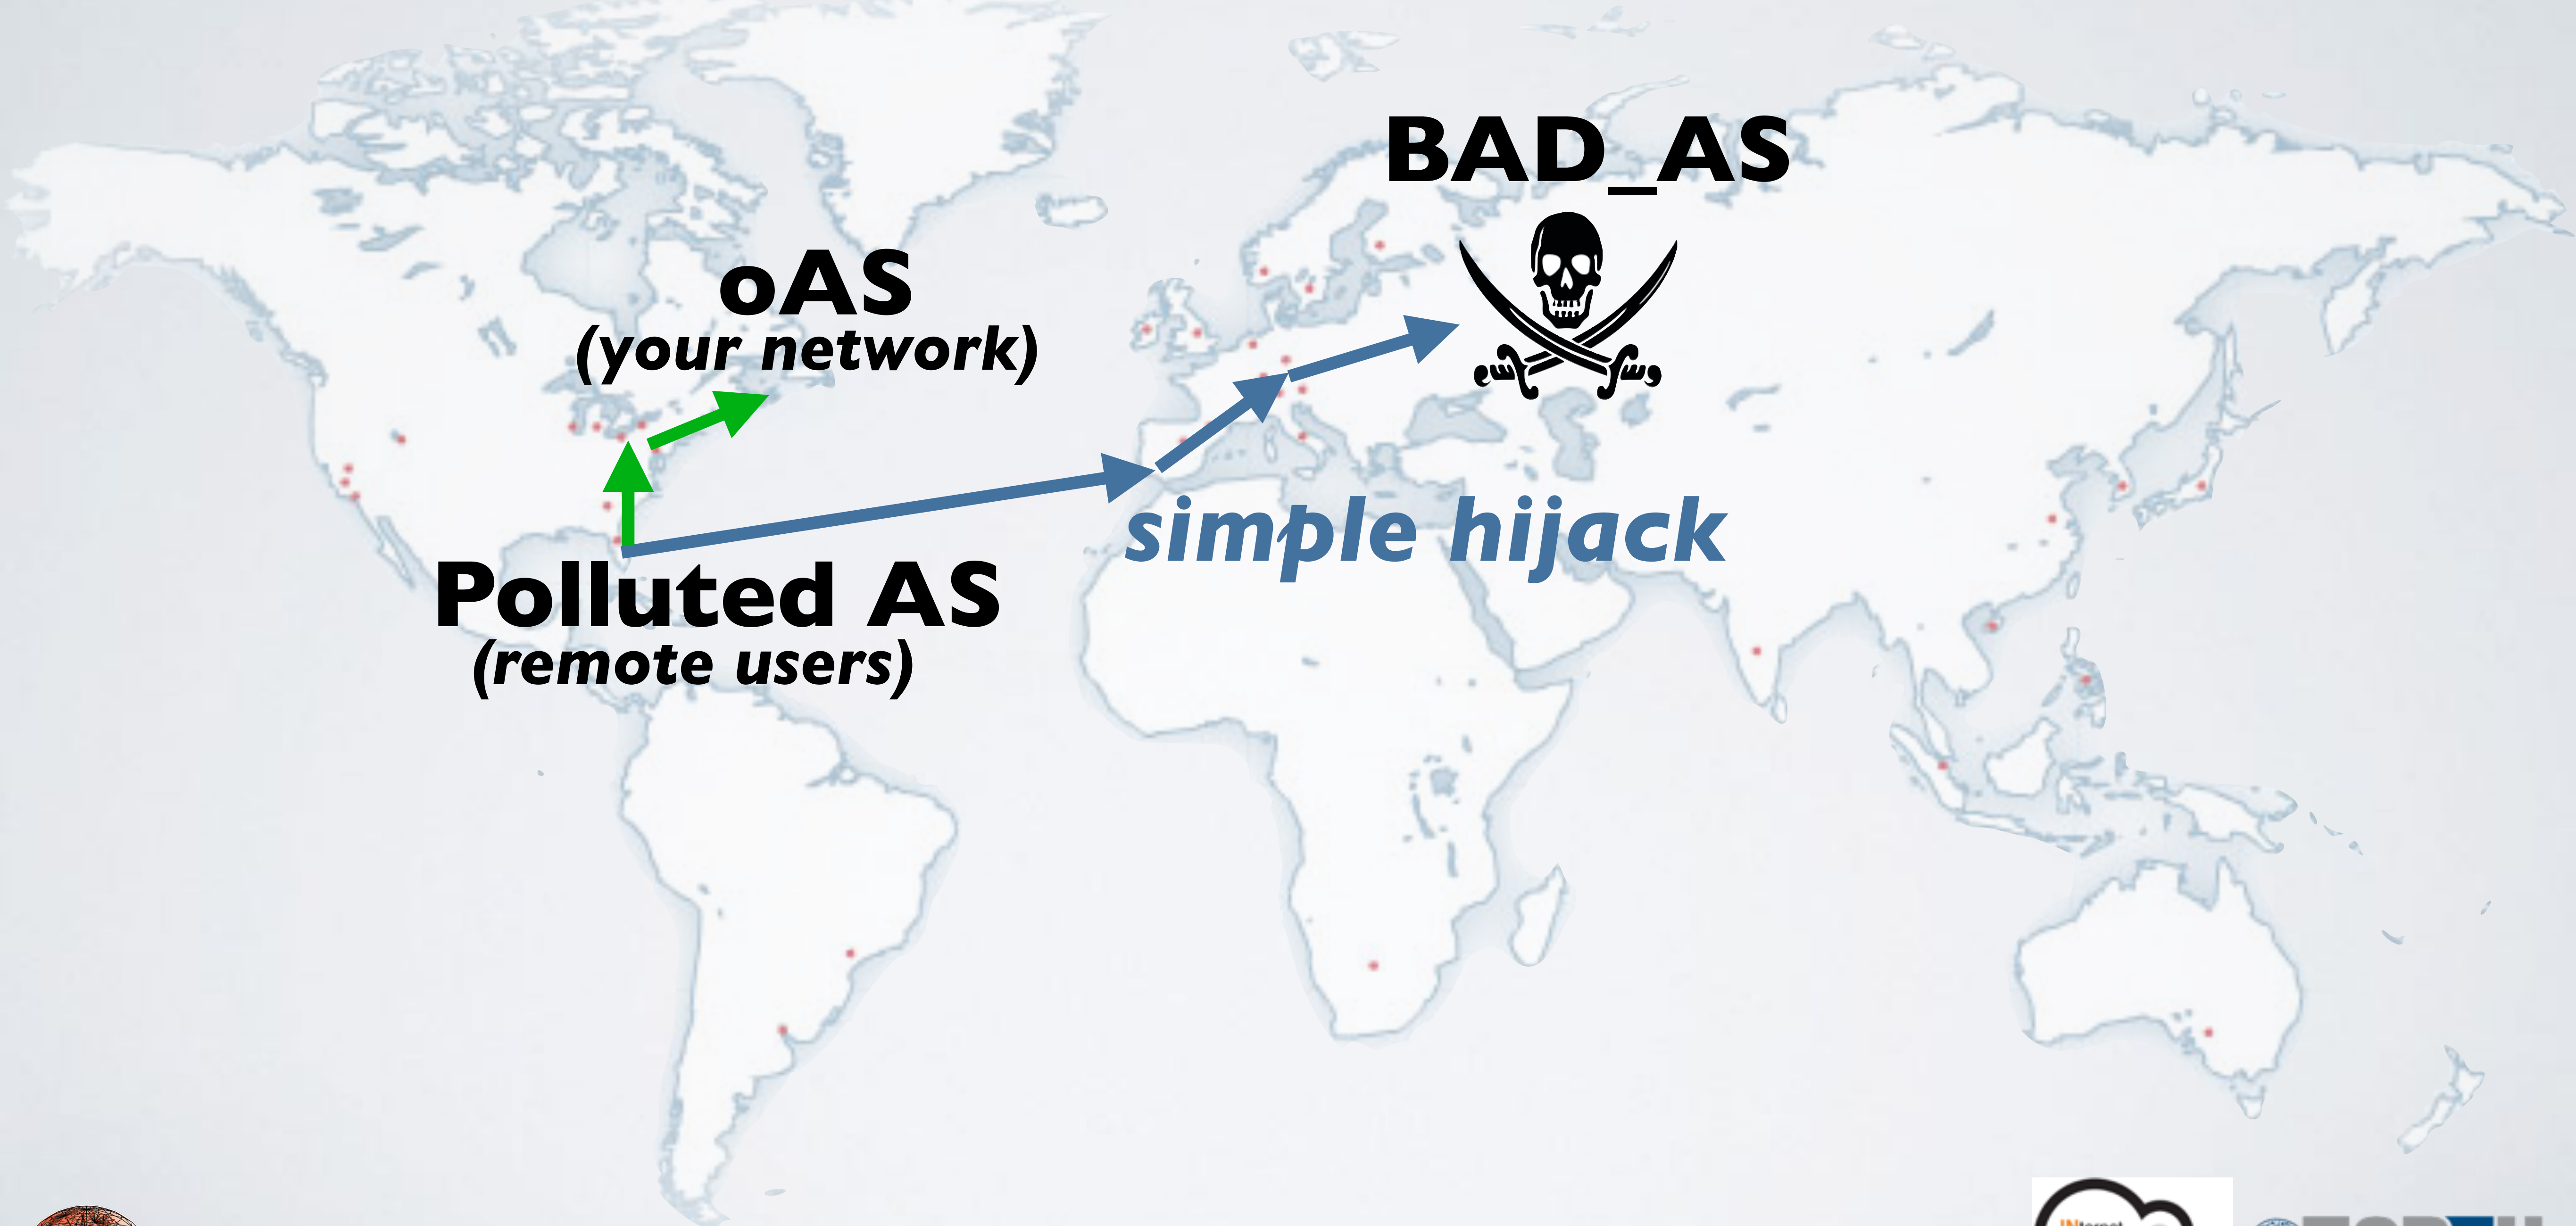

BAD_AS

oAS
(your network)

Polluted AS
(remote users)

simple hijack

INTERNET ROUTE HIJACKING

a threat to your organization and to critical infrastructure

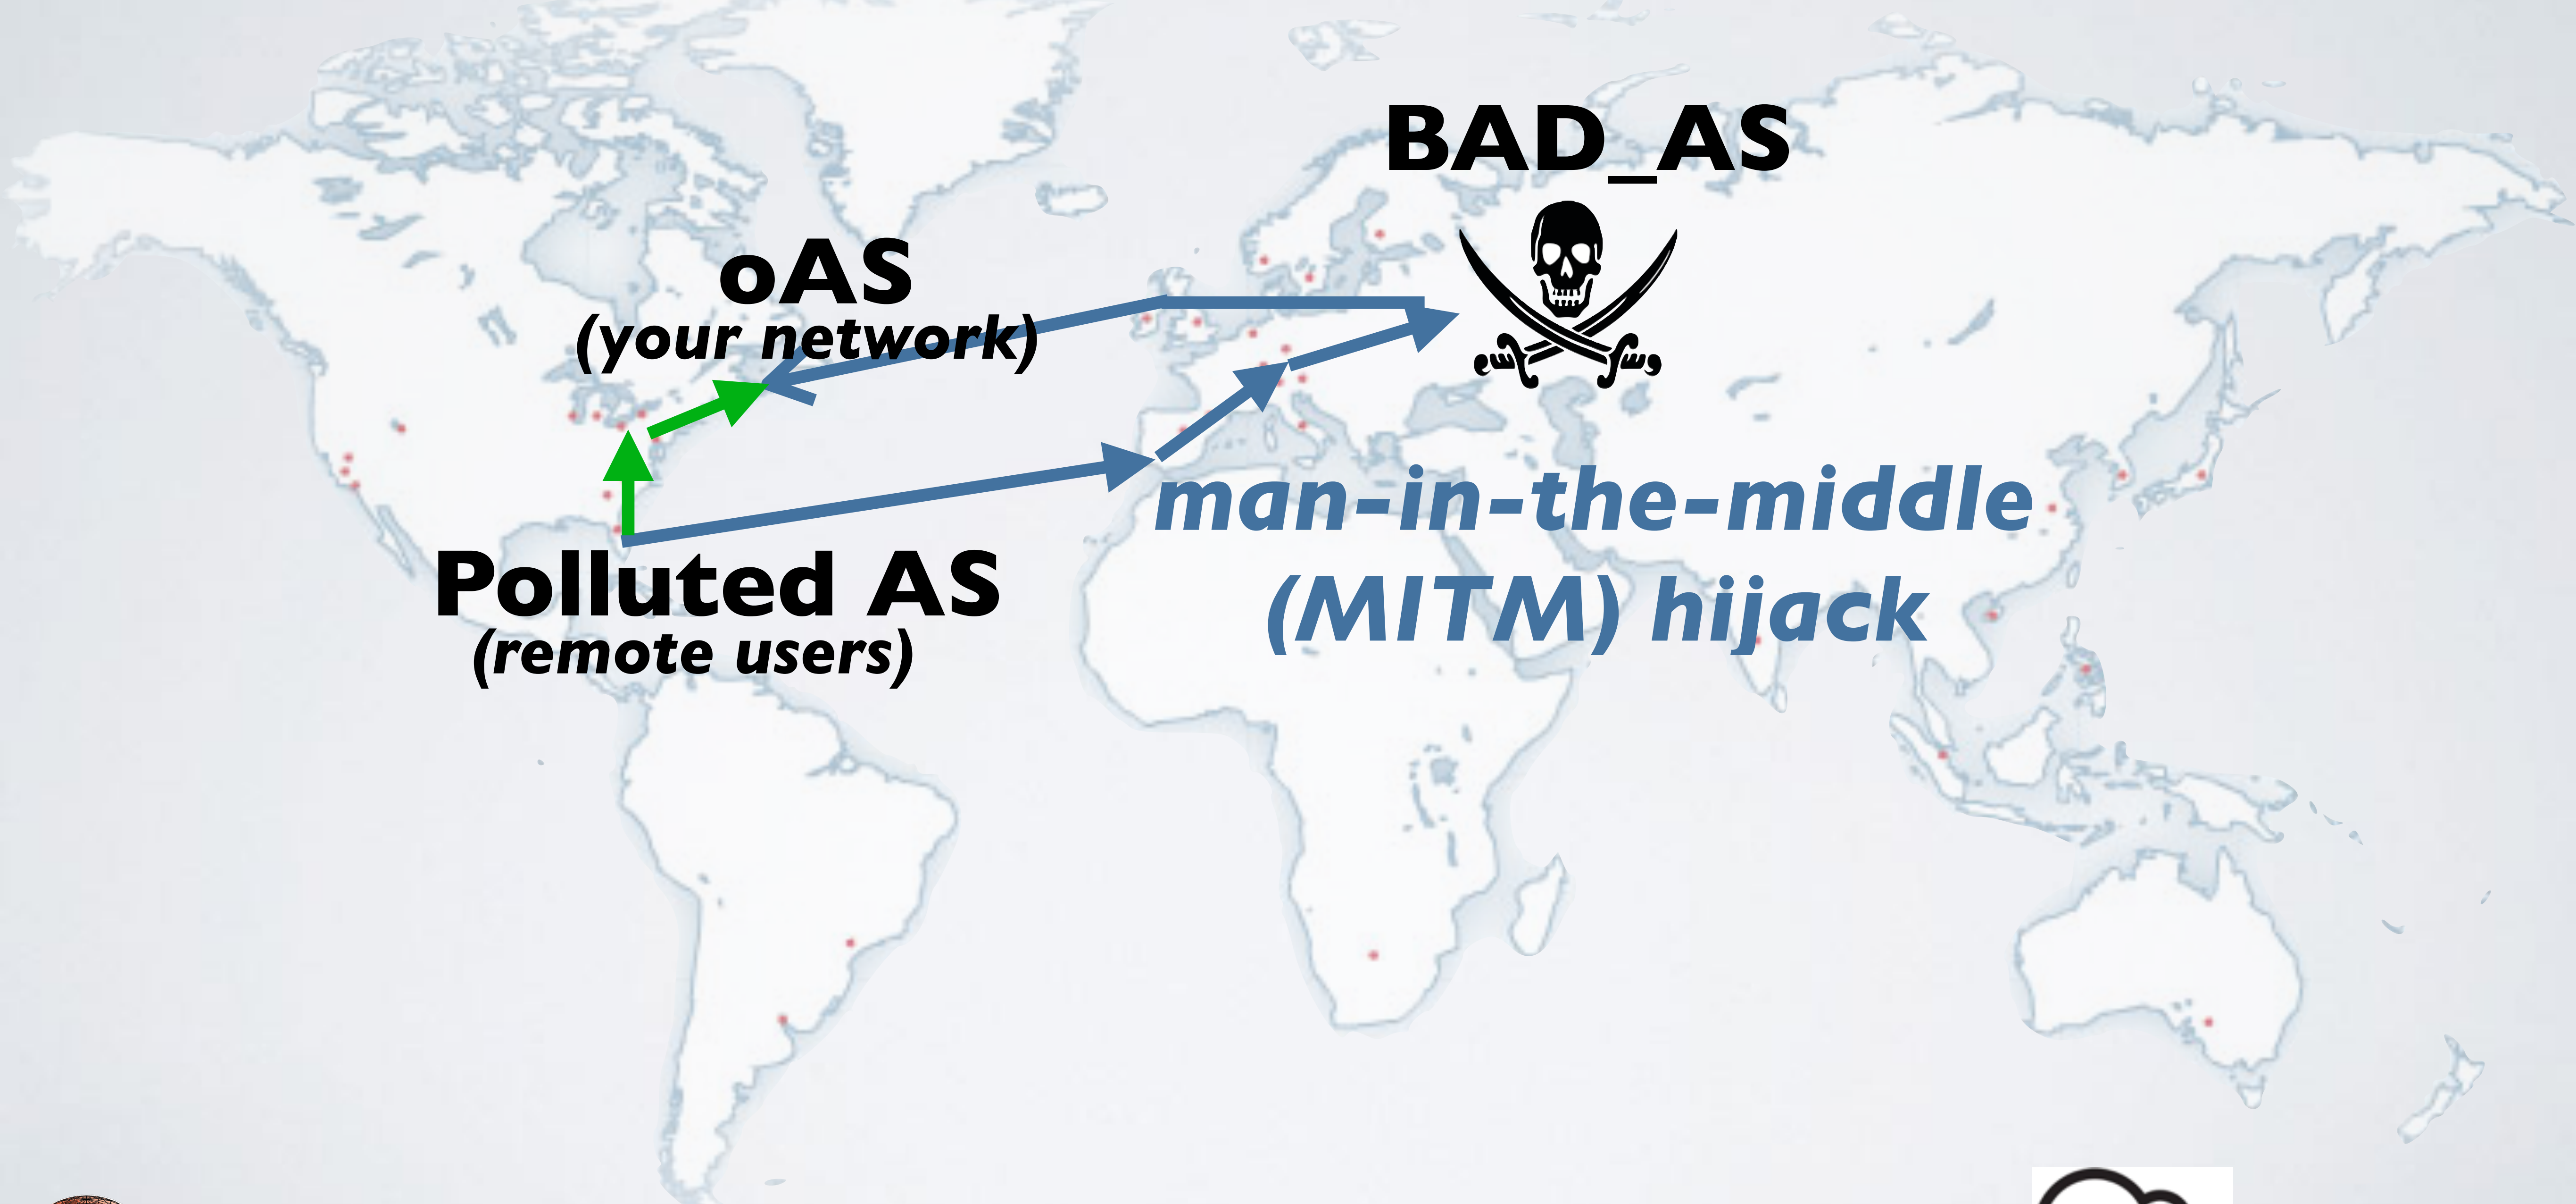

INTERNET ROUTE HIJACKING

many **MITM** *events documented*

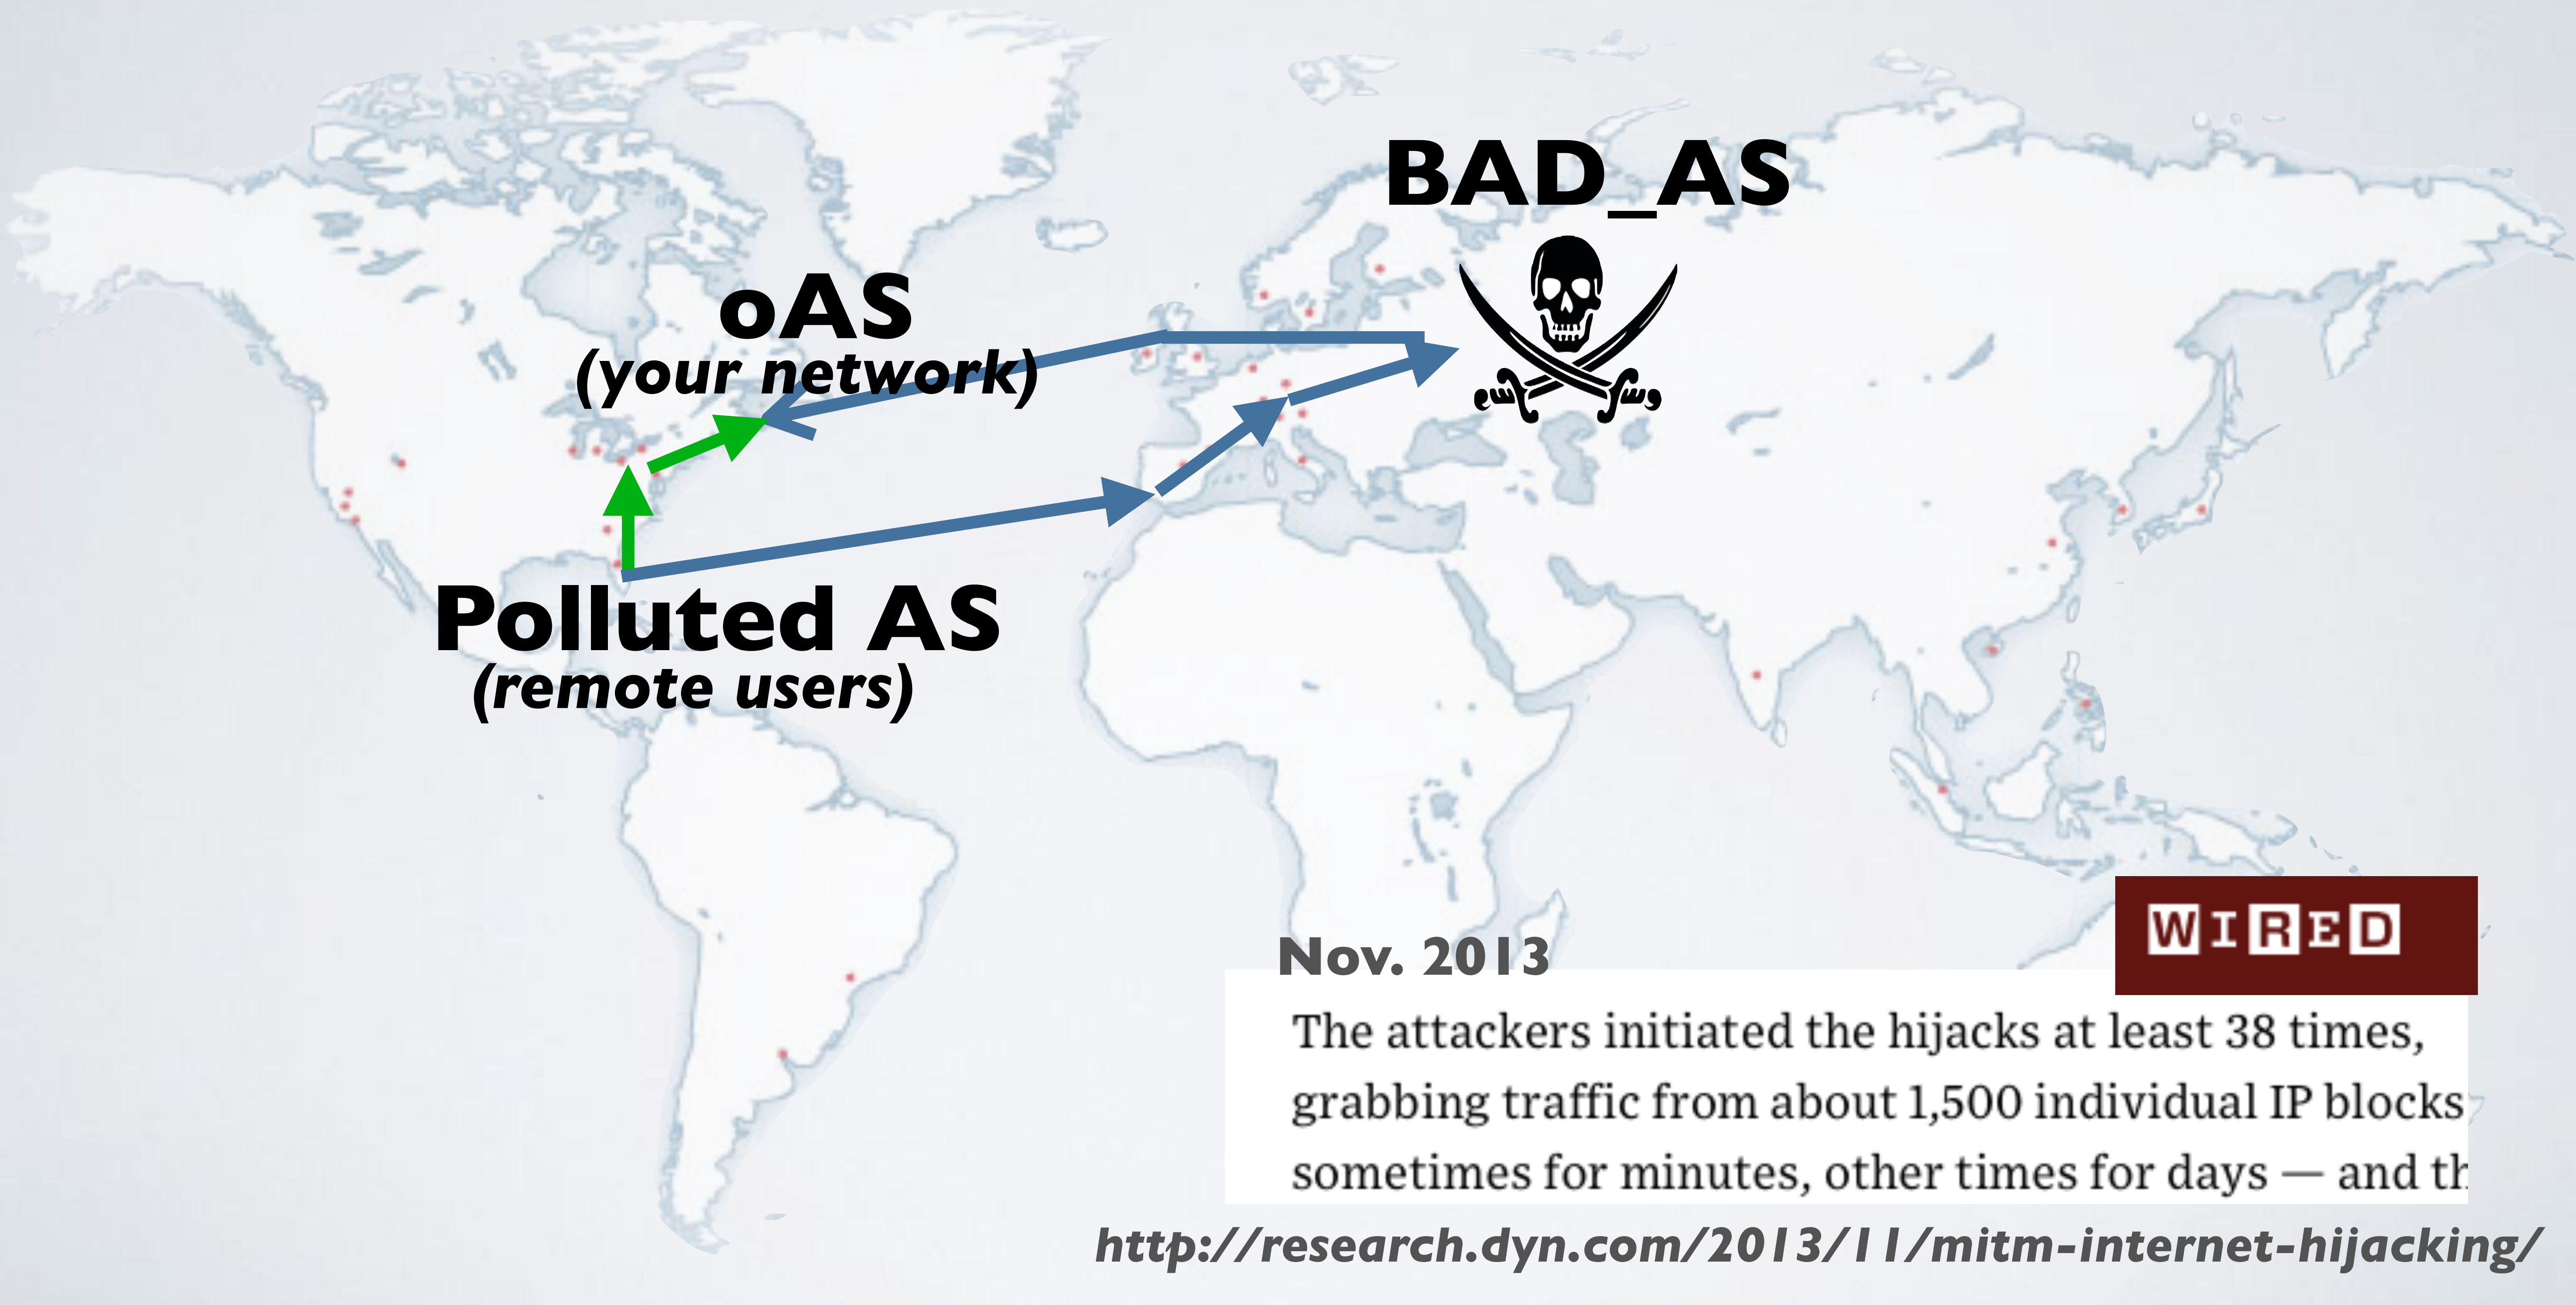

Nov. 2013

The attackers initiated the hijacks at least 38 times, grabbing traffic from about 1,500 individual IP blocks sometimes for minutes, other times for days — and th

http://research.dyn.com/2013/11/mitm-internet-hijacking/

BGP UNIVERSE

before ARTEMIS

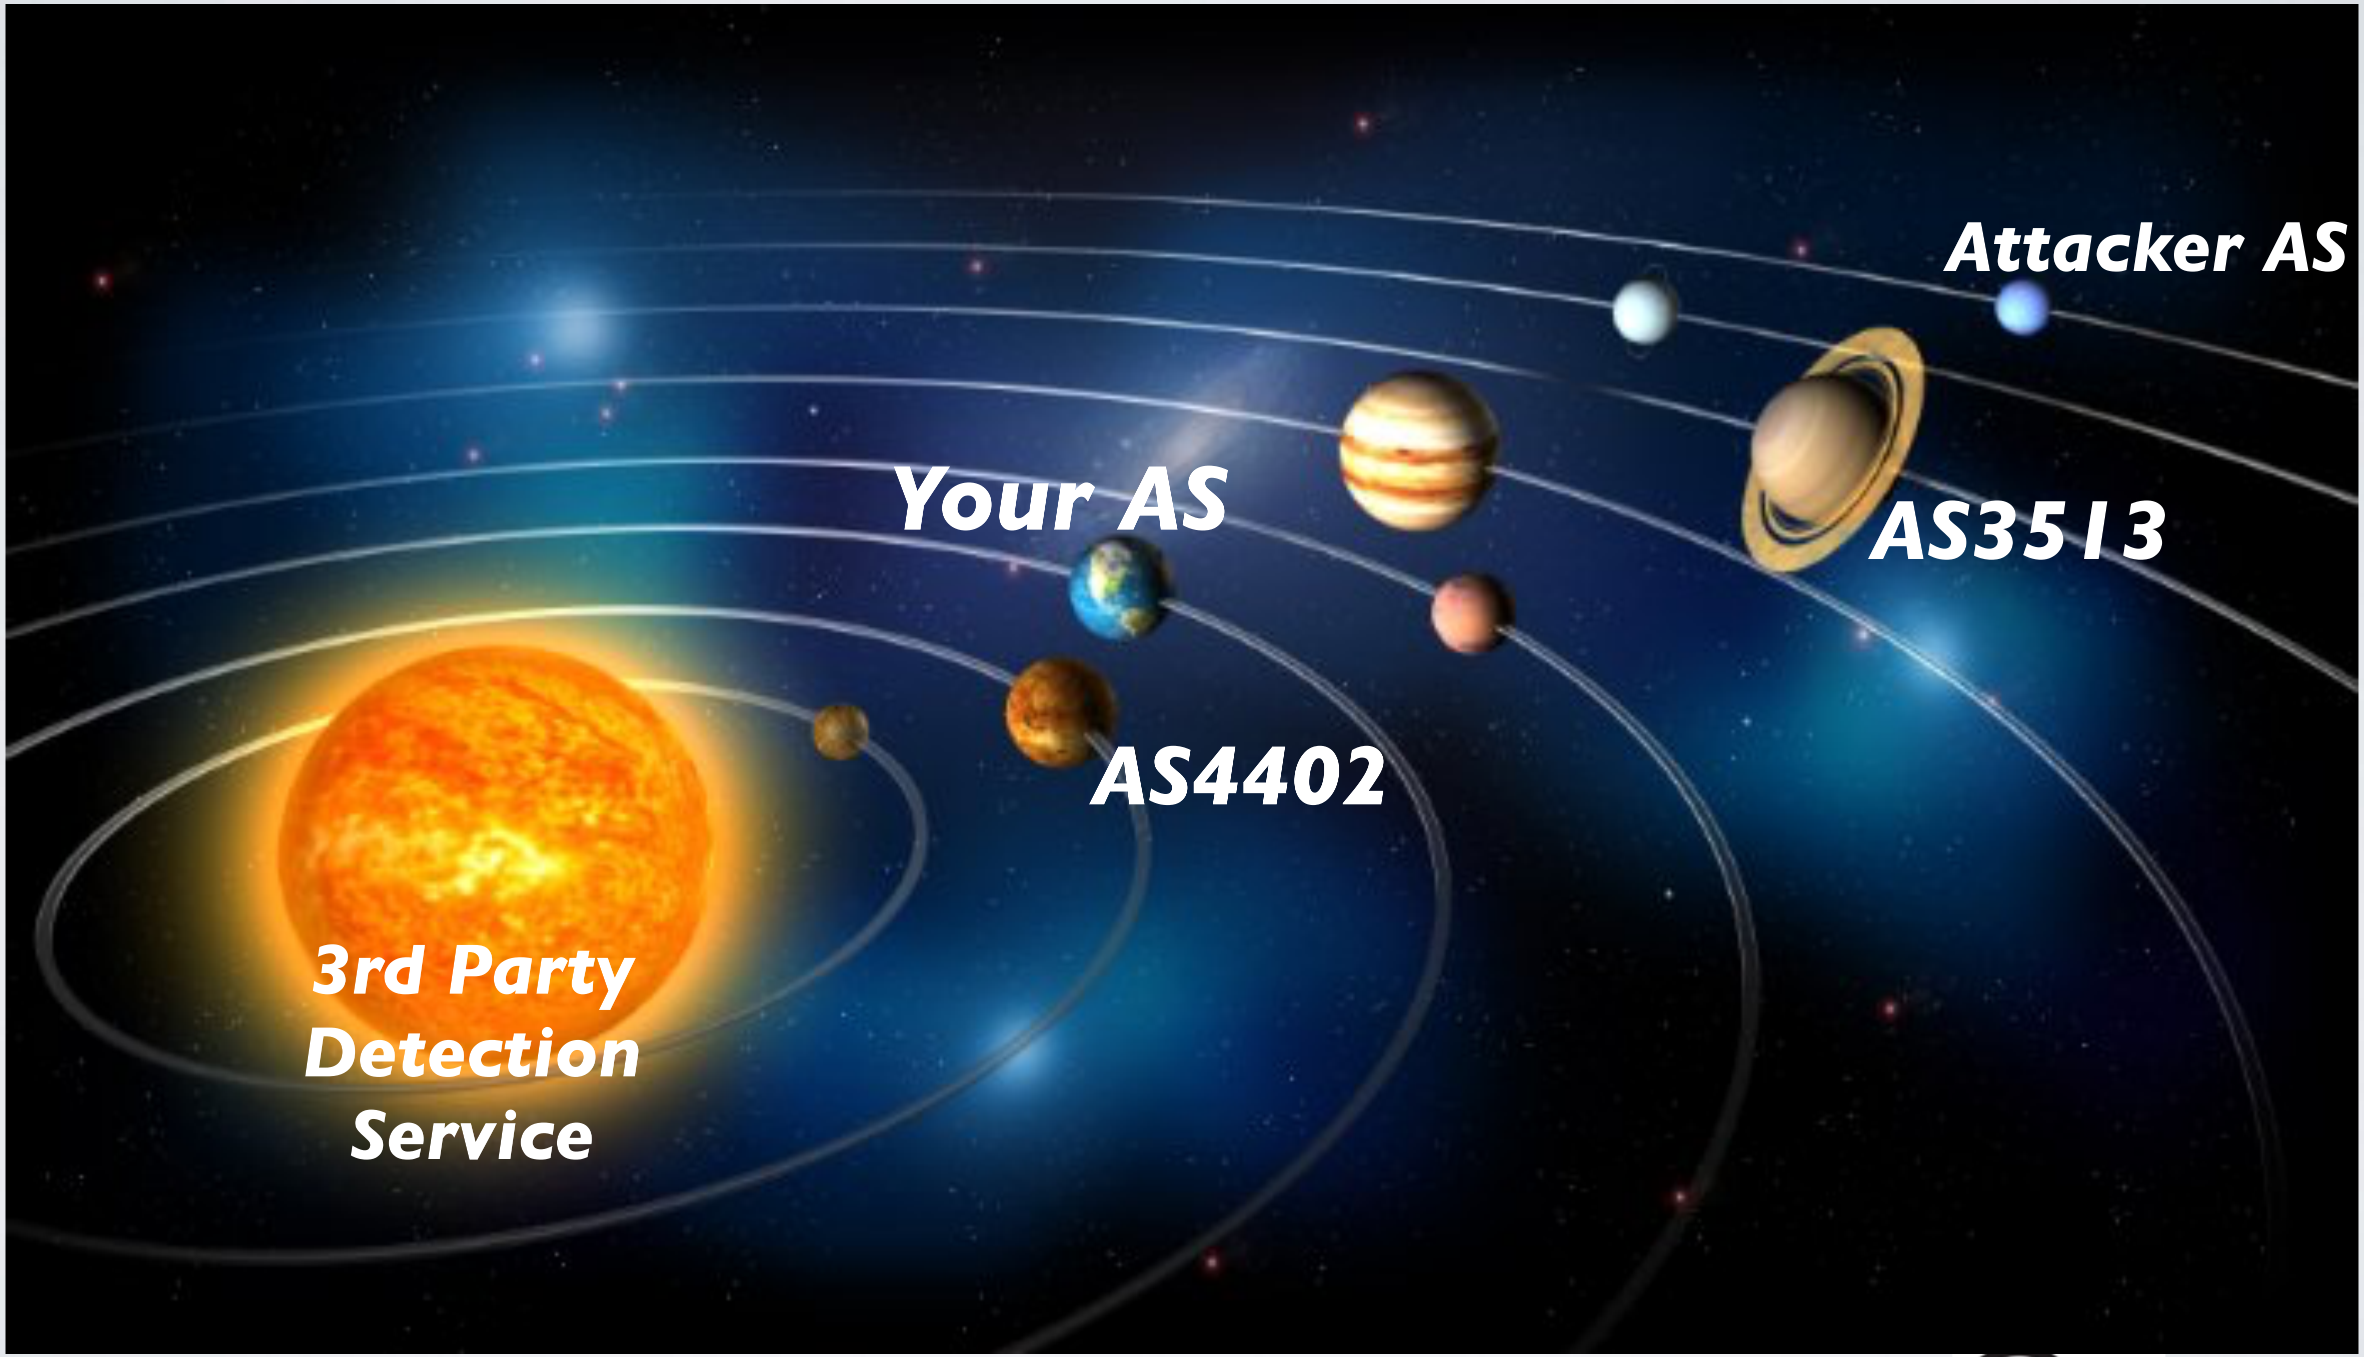

THIRD PARTY SERVICES

both theoretical and practical issues

- **Evasion**
  - Only simple attack configurations are considered
- **Accuracy**
  - Potential for lots of false positives
  - or alternatively lots of false negatives
- **Speed**
  - Manual verification then manual mitigation
- **Privacy**
  - Need to share private information

BGP UNIVERSE

before ARTEMIS

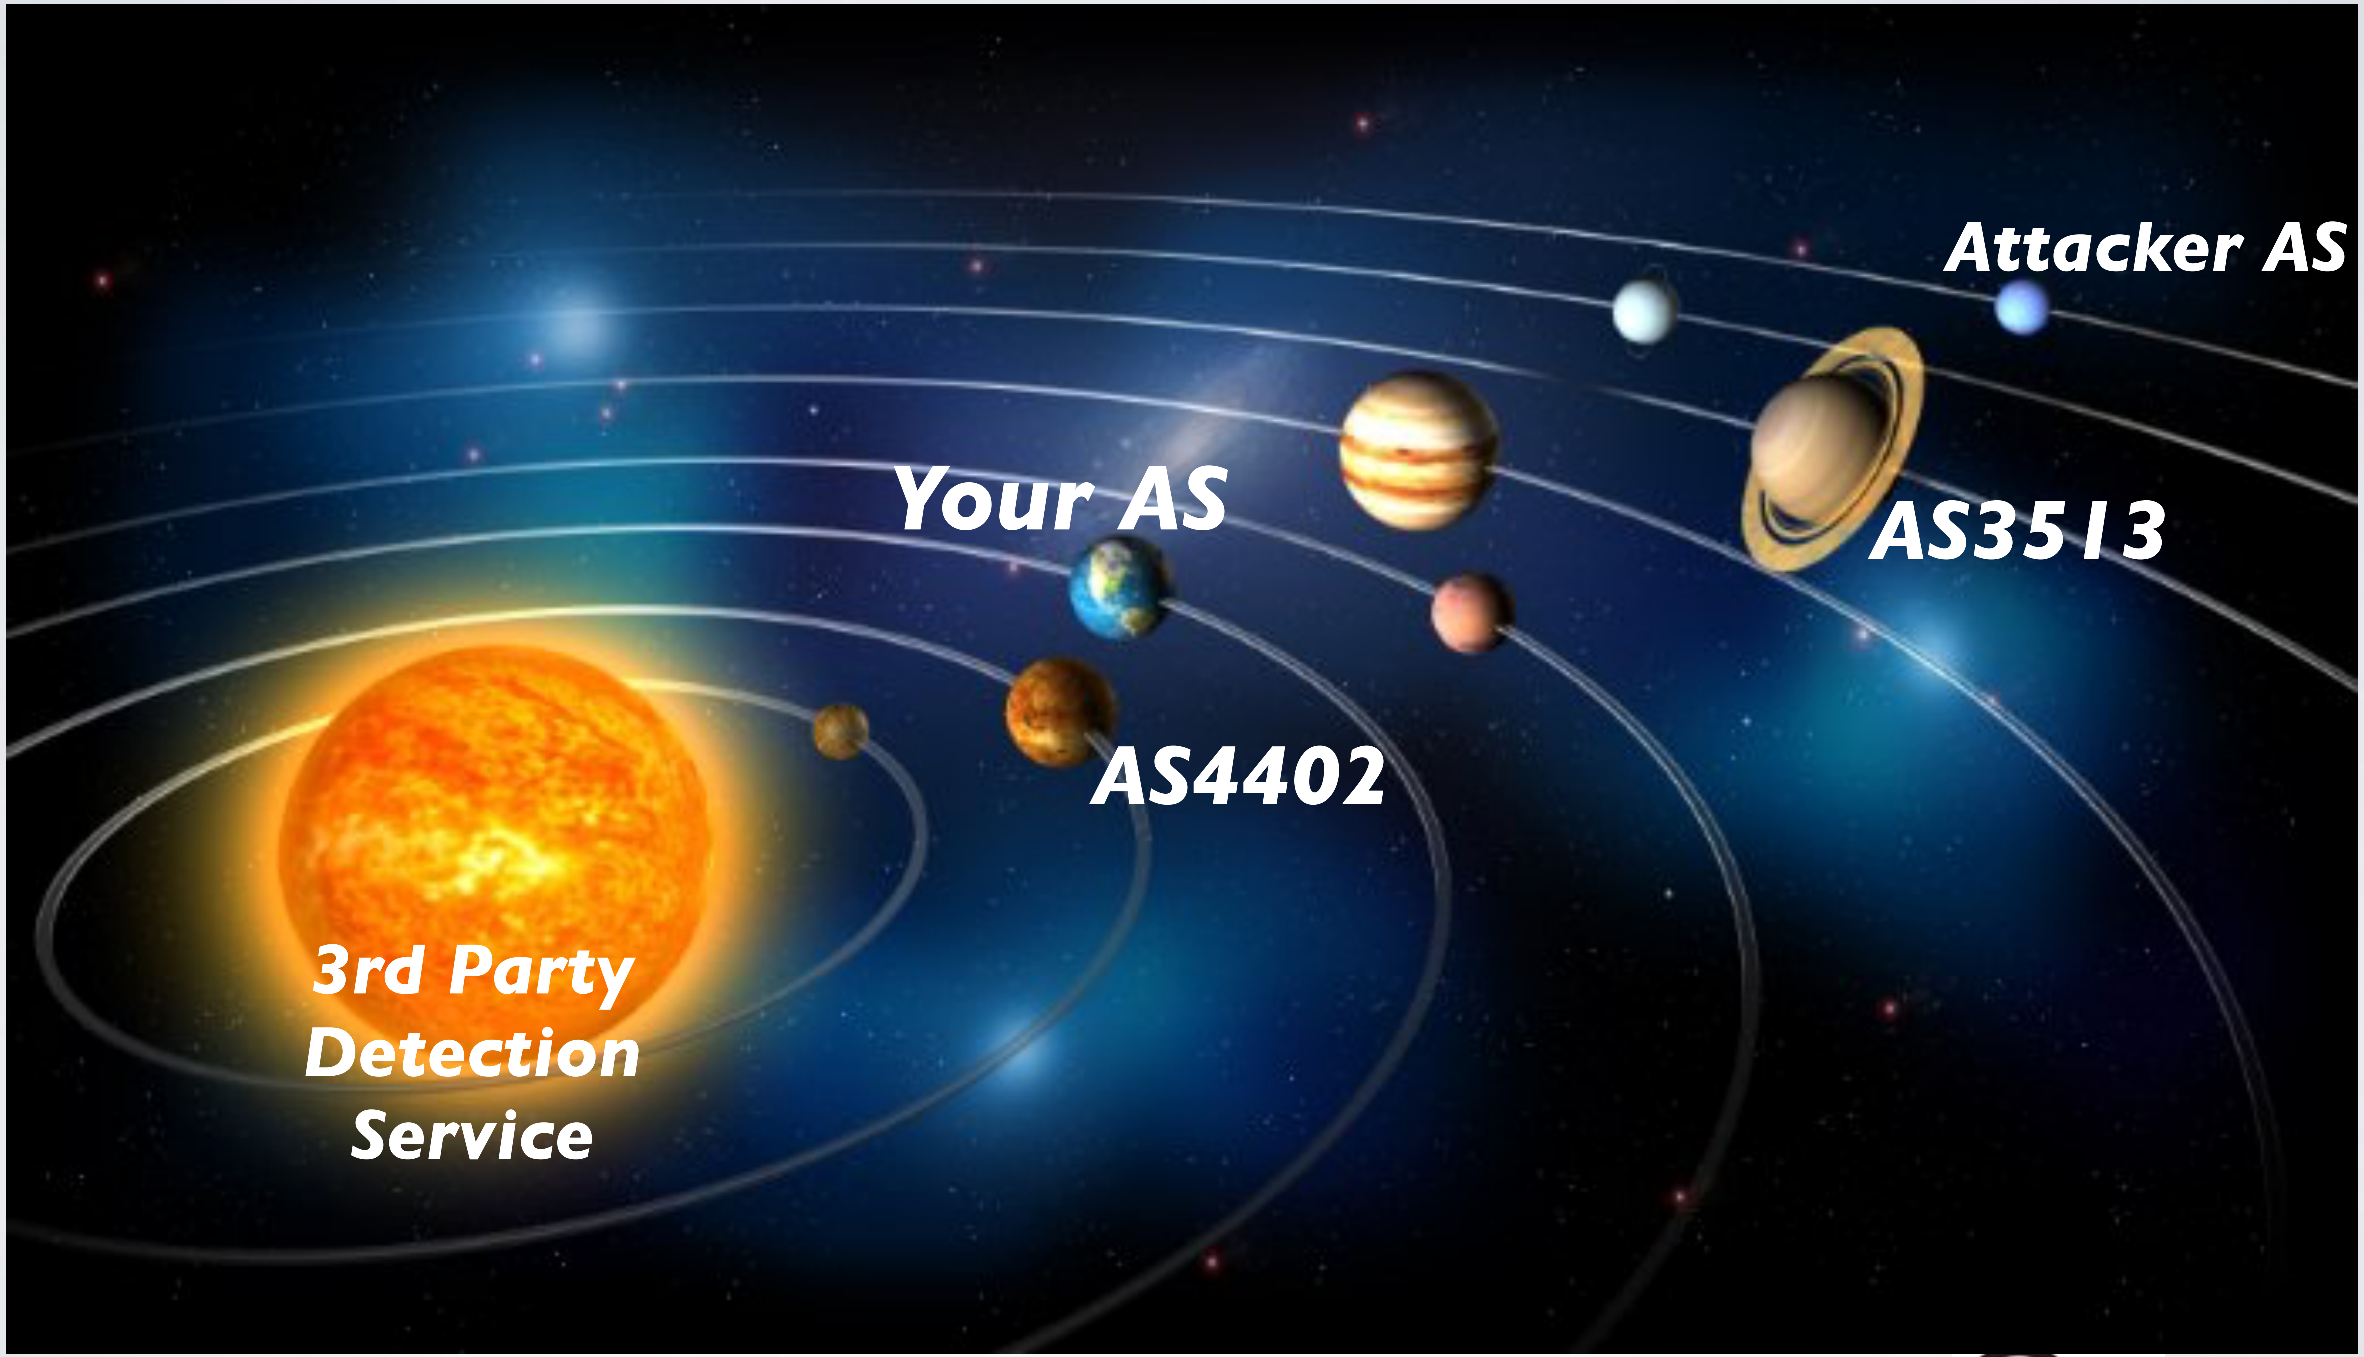

ARTEMIS IN A NUTSHELL

a ptolemaic revolution :-)

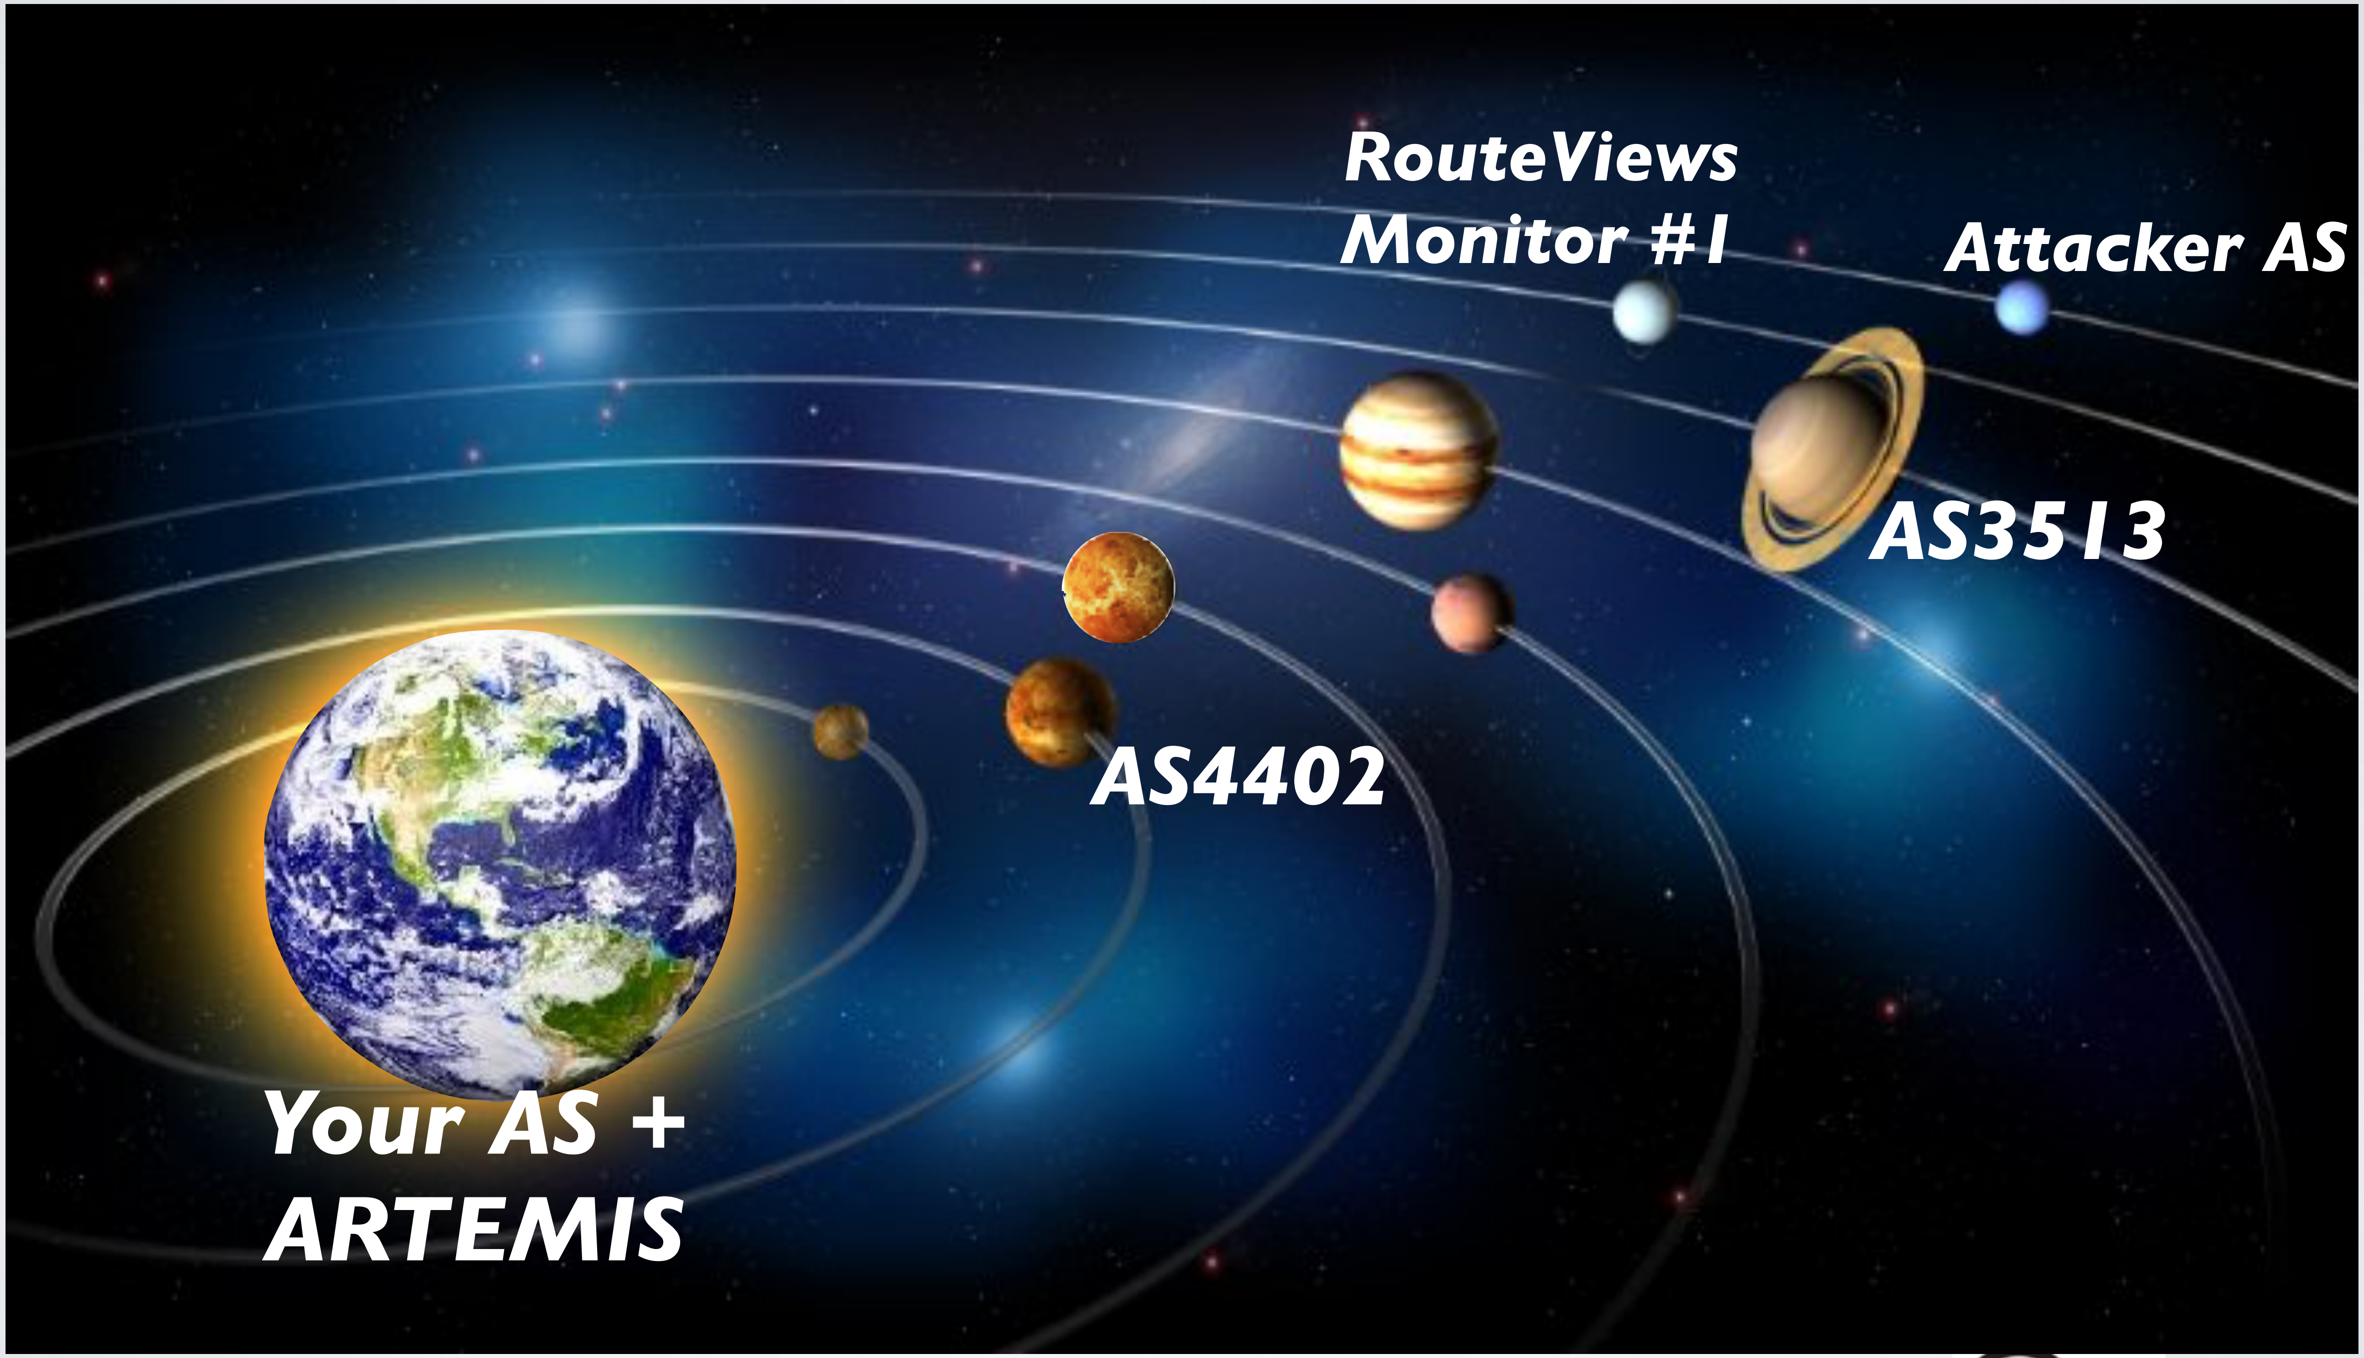

ARTEMIS IN A NUTSHELL

..then suddenly everything makes sense

- **Evasion**
  - Covers *all* attack configurations
- **Accuracy**
  - *0% FP, 0% FN:* for most attack configurations
  - 0% FN for the remaining ones (alternatively manage FP-FN trade-off)
- **Speed**
  - Automated mitigation: neutralize attacks in a *minute*
- **Privacy & Flexibility**
  - *full privacy*
  - per-prefix + per-event type, configurable mitigation

PUBLIC MONITORING INFRASTRUCTURE

enables visibility of all significant events

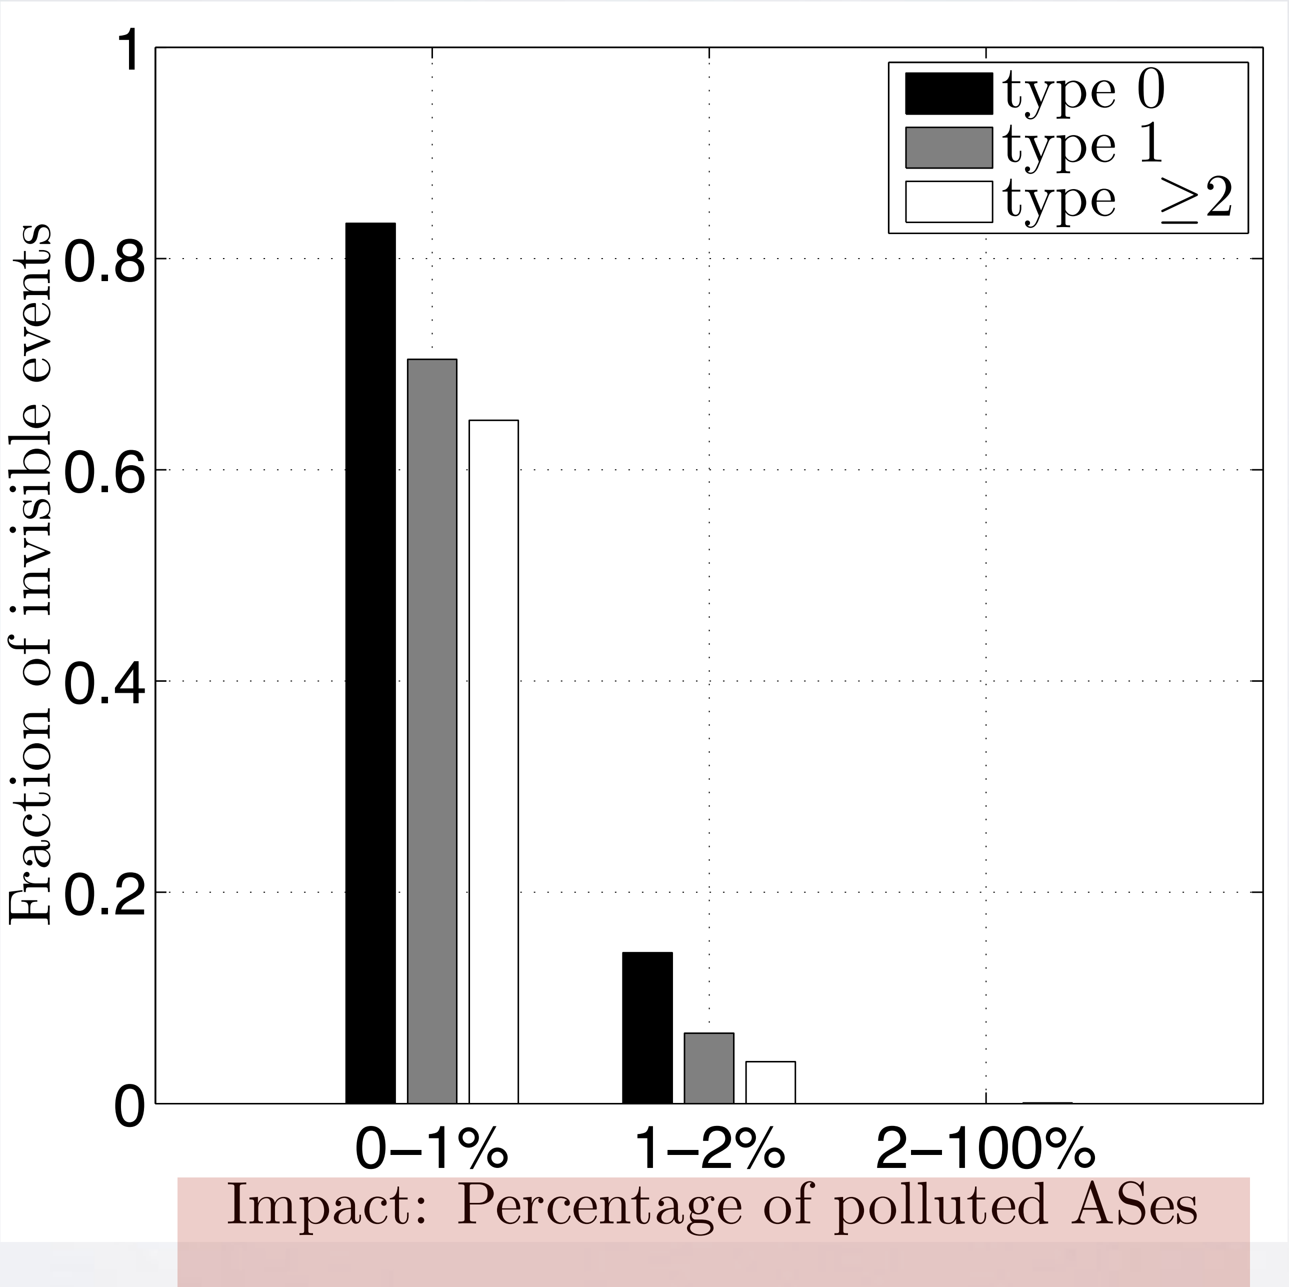

- In the paper:
  - by type of service
  - Impact
  - Speed

BGP HIJACKING TAXONOMY

3 dimensions

- **1)** Based on how the “attacking” AS Path looks like
  - **Type 0** hijack: <*prefix:* ***BAD_AS****, …>* *(a.k.a. “prefix origin hijack”)*
  - **Type 1** hijack: <*prefix: oAS,* ***BAD_AS****, …>*
  - **Type 2** hijack: <*prefix: oAS, AS1,* ***BAD_AS****, …>*
  - …
  - **Type N** hijack: <*prefix: oAS, AS1, …,* ***BAD_AS****, …>*
  - **Type U** hijack: <*prefix: unaltered_path>*

- **2)** Based on the prefix: announced ***prefix*** *or* ***sub-prefix****, or* ***squatting***

- **3)** Based on what happens on the data-plane: *Black Holing (****BH****), Imposture (****IM****), Man in the Middle (****MM****)*

ATTACK COVERAGE

ARTEMIS vs previous literature

TABLE 1: Comparison of BGP prefix hijacking detection systems/services w.r.t. ability to detect different classes of attacks.

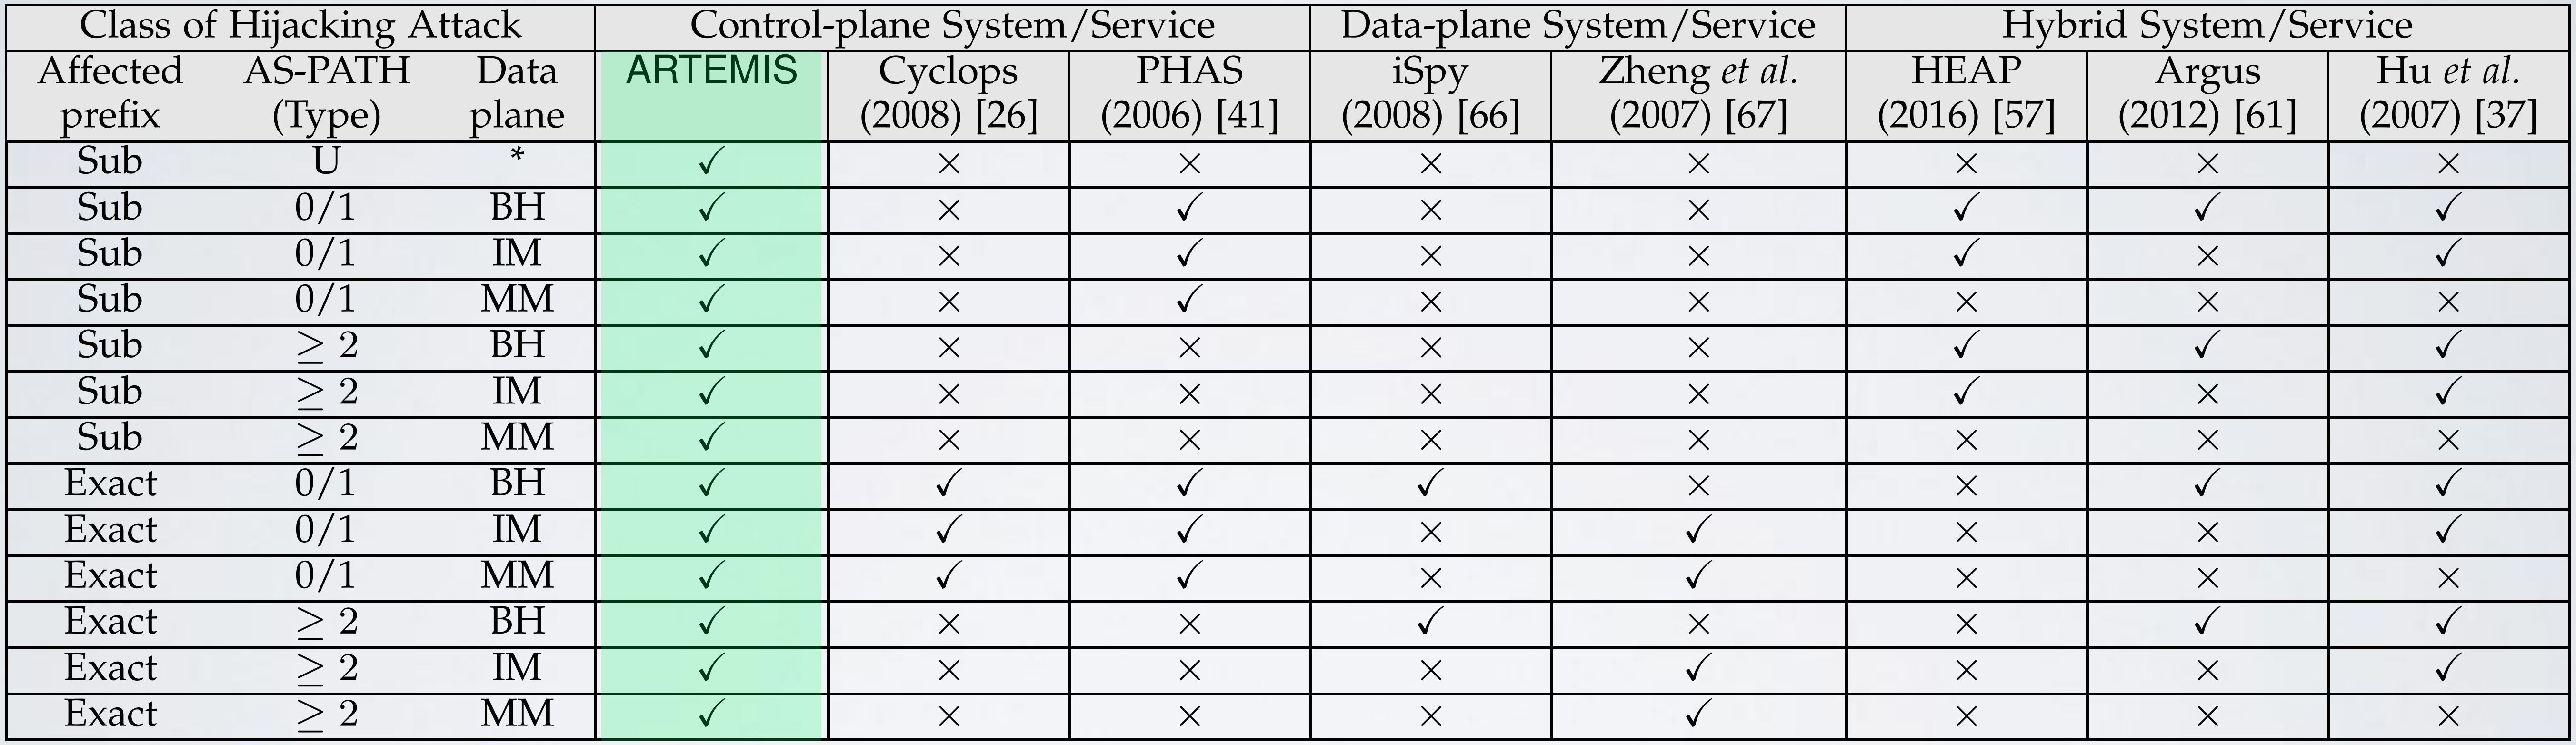

Class of Hijacking Attack			Control-plane System/Service			Data-plane System/Service		Hybrid System/Service		
Affected prefix	AS-PATH (Type)	Data plane	ARTEMIS	Cyclops (2008) [26]	PHAS (2006) [41]	iSpy (2008) [66]	Zheng *et al.* (2007) [67]	HEAP (2016) [57]	Argus (2012) [61]	Hu *et al.* (2007) [37]
Sub	U	*	✓	×	×	×	×	×	×	×
Sub	0/1	BH	✓	×	✓	×	×	✓	✓	✓
Sub	0/1	IM	✓	×	✓	×	×	✓	×	✓
Sub	0/1	MM	✓	×	✓	×	×	×	×	×
Sub	≥ 2	BH	✓	×	×	×	×	✓	✓	✓
Sub	≥ 2	IM	✓	×	×	×	×	✓	×	✓
Sub	≥ 2	MM	✓	×	×	×	×	×	×	×
Exact	0/1	BH	✓	✓	✓	✓	×	×	✓	✓
Exact	0/1	IM	✓	✓	✓	×	✓	×	×	✓
Exact	0/1	MM	✓	✓	✓	×	✓	×	×	×
Exact	≥ 2	BH	✓	×	×	✓	×	×	✓	✓
Exact	≥ 2	IM	✓	×	×	×	✓	×	×	✓
Exact	≥ 2	MM	✓	×	×	×	✓	×	×	×

ACCURATE DETECTION

becomes trivial in most of the cases

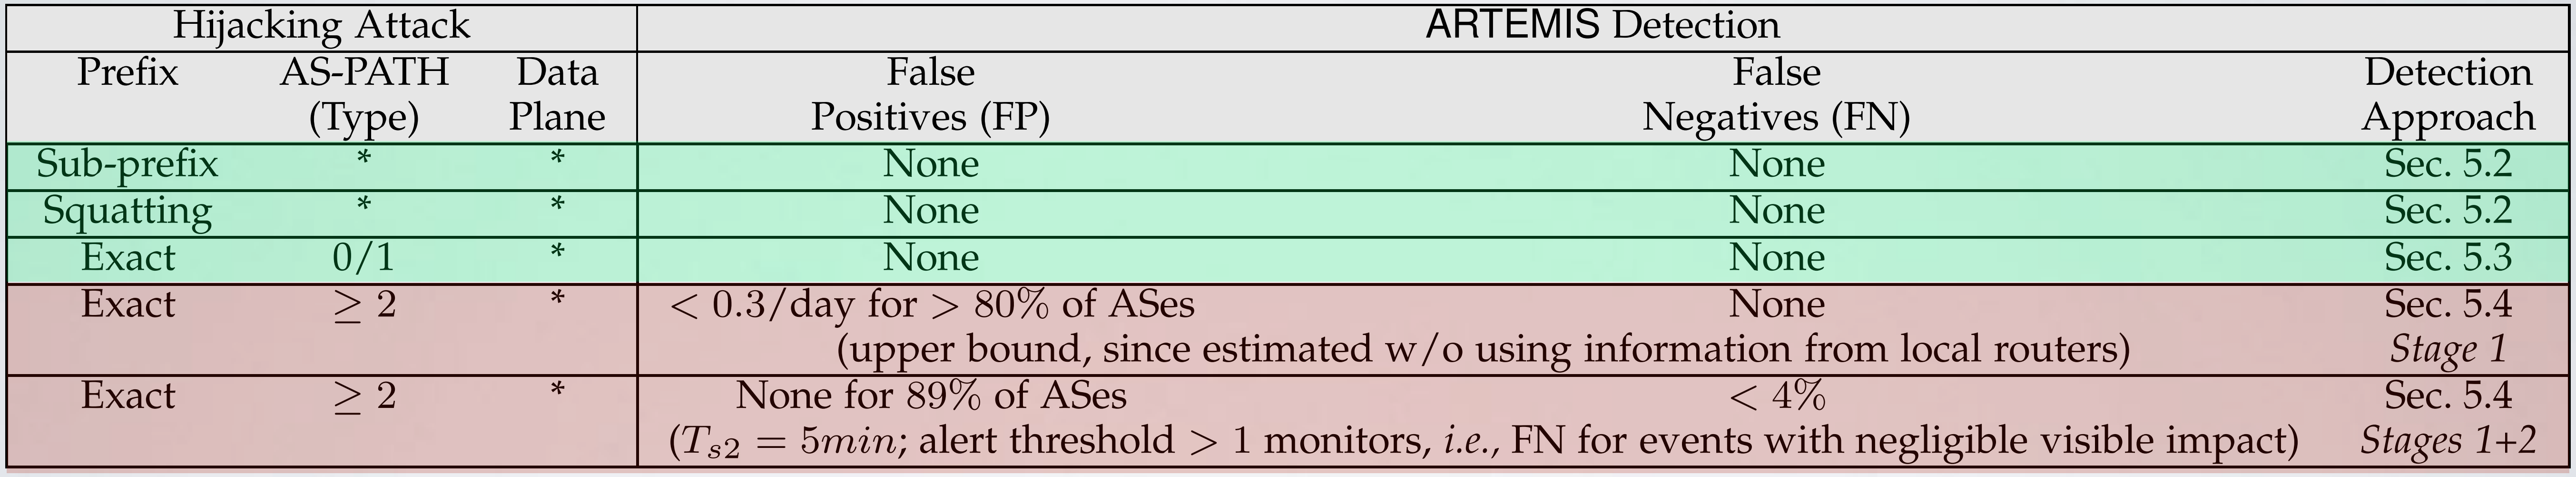

Hijacking Attack			ARTEMIS Detection		
Prefix	AS-PATH (Type)	Data Plane	False Positives (FP)	False Negatives (FN)	Detection Approach
Sub-prefix	*	*	None	None	Sec. 5.2
Squatting	*	*	None	None	Sec. 5.2
Exact	0/1	*	None	None	Sec. 5.3
Exact	≥ 2	*	< 0.3/day for $> 80\%$ of ASes (upper bound, since estimated w/o using information from local routers)	None	Sec. 5.4 *Stage 1*
Exact	≥ 2	*	None for 89% of ASes ($T_{s2} = 5min$; alert threshold > 1 monitors, *i.e.*, FN for events with negligible visible impact)	$< 4\%$	Sec. 5.4 *Stages 1+2*

TYPE ≥ 2 HIJACKS

Stage I

- Triggered when: a BGP update (for a monitored prefix) whose AS-PATH contains a N-hop AS-link (N ≥ 2) that is not included in the previously verified AS-links list

- Legitimate if this link has been observed in the *opposite direction* in the AS-links list from monitors and local BGP routers (10 months history).

<prefix: oAS, neighborAS, ***BAD_AS****, …> attack announcement*

<any prefix: …, ***BAD_AS,*** *neighborAS, …,* ***BAD_AS****, …> pre-attack fails*

<any prefix: …, ***BAD_AS,*** *neighborAS, …,* ***2ndBAD_AS****, …> pre-attack ok*

TYPE ≥ 2 HIJACKS

Stage 1 and Stage 2

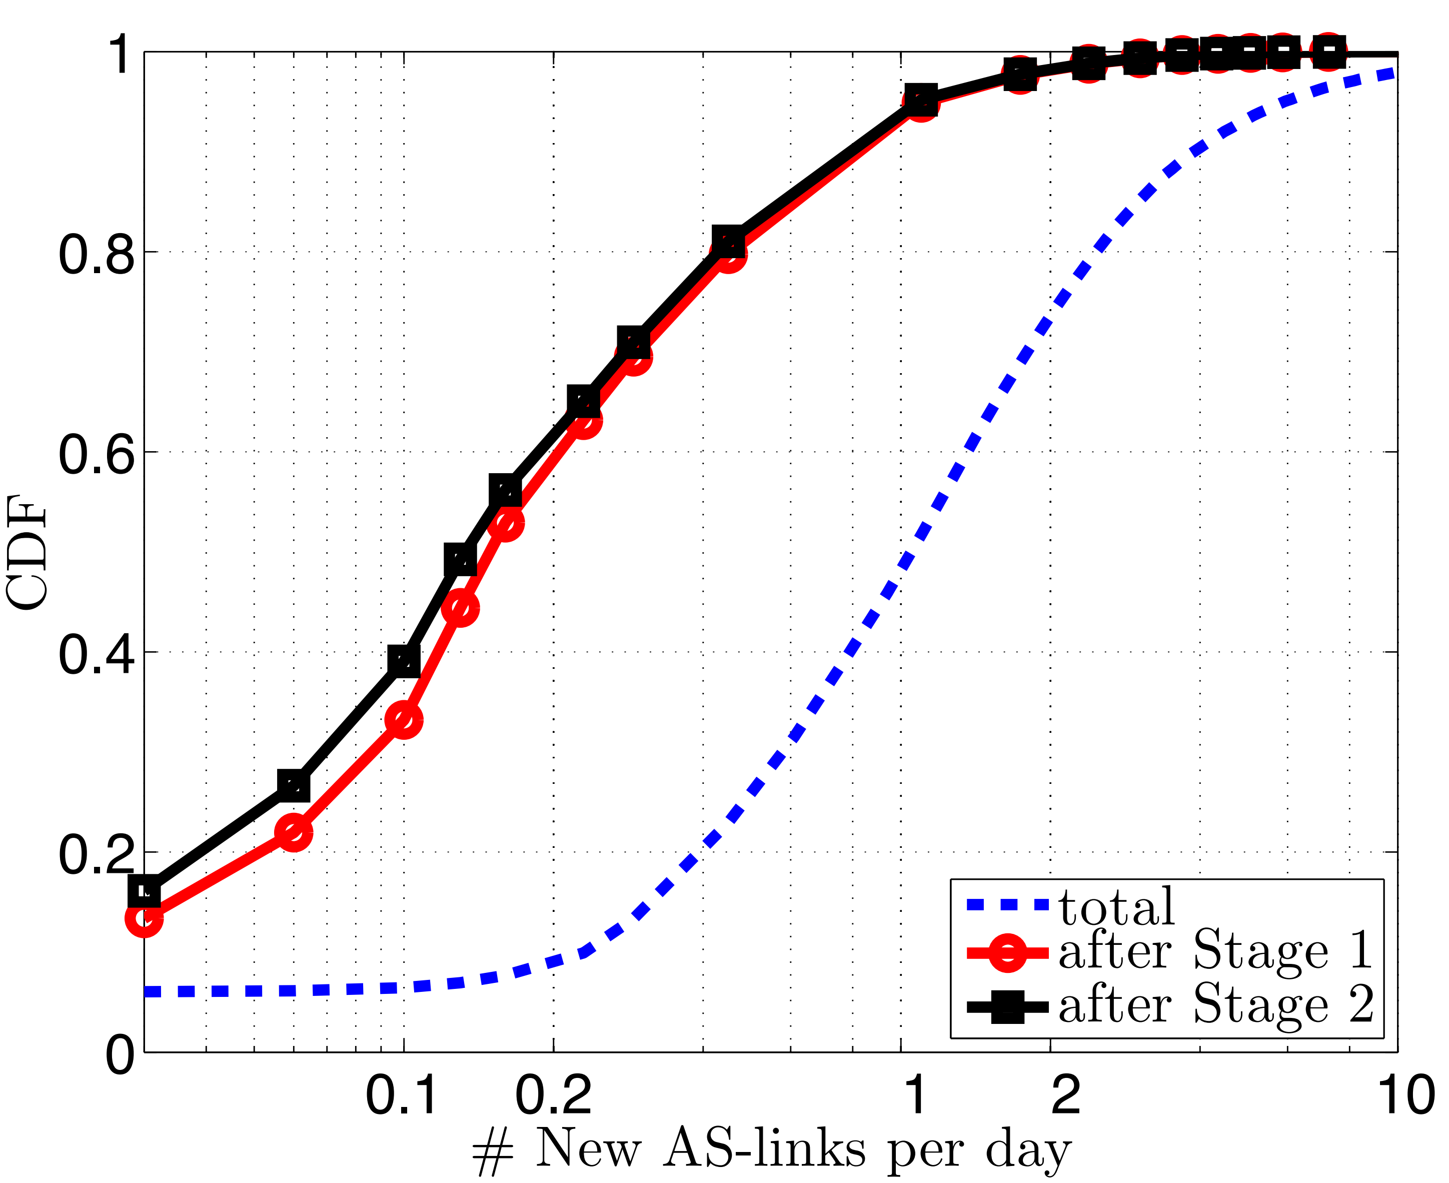

- **Stage 2**
  - Wait 5 minutes
  - Recheck tables
  - …

TABLE 5: Simulation results of the reduction (%) of false positives by *Stage 2*, due to the information from monitors and local routers.

position of new link:	2^{nd} hop	3^{rd} hop
only monitors	0.2%	4.6%
monitors+local router(s)	24.2%	31.8%

TYPE ≥ 2 HIJACKS

Stage 2 w/ FN of small impact

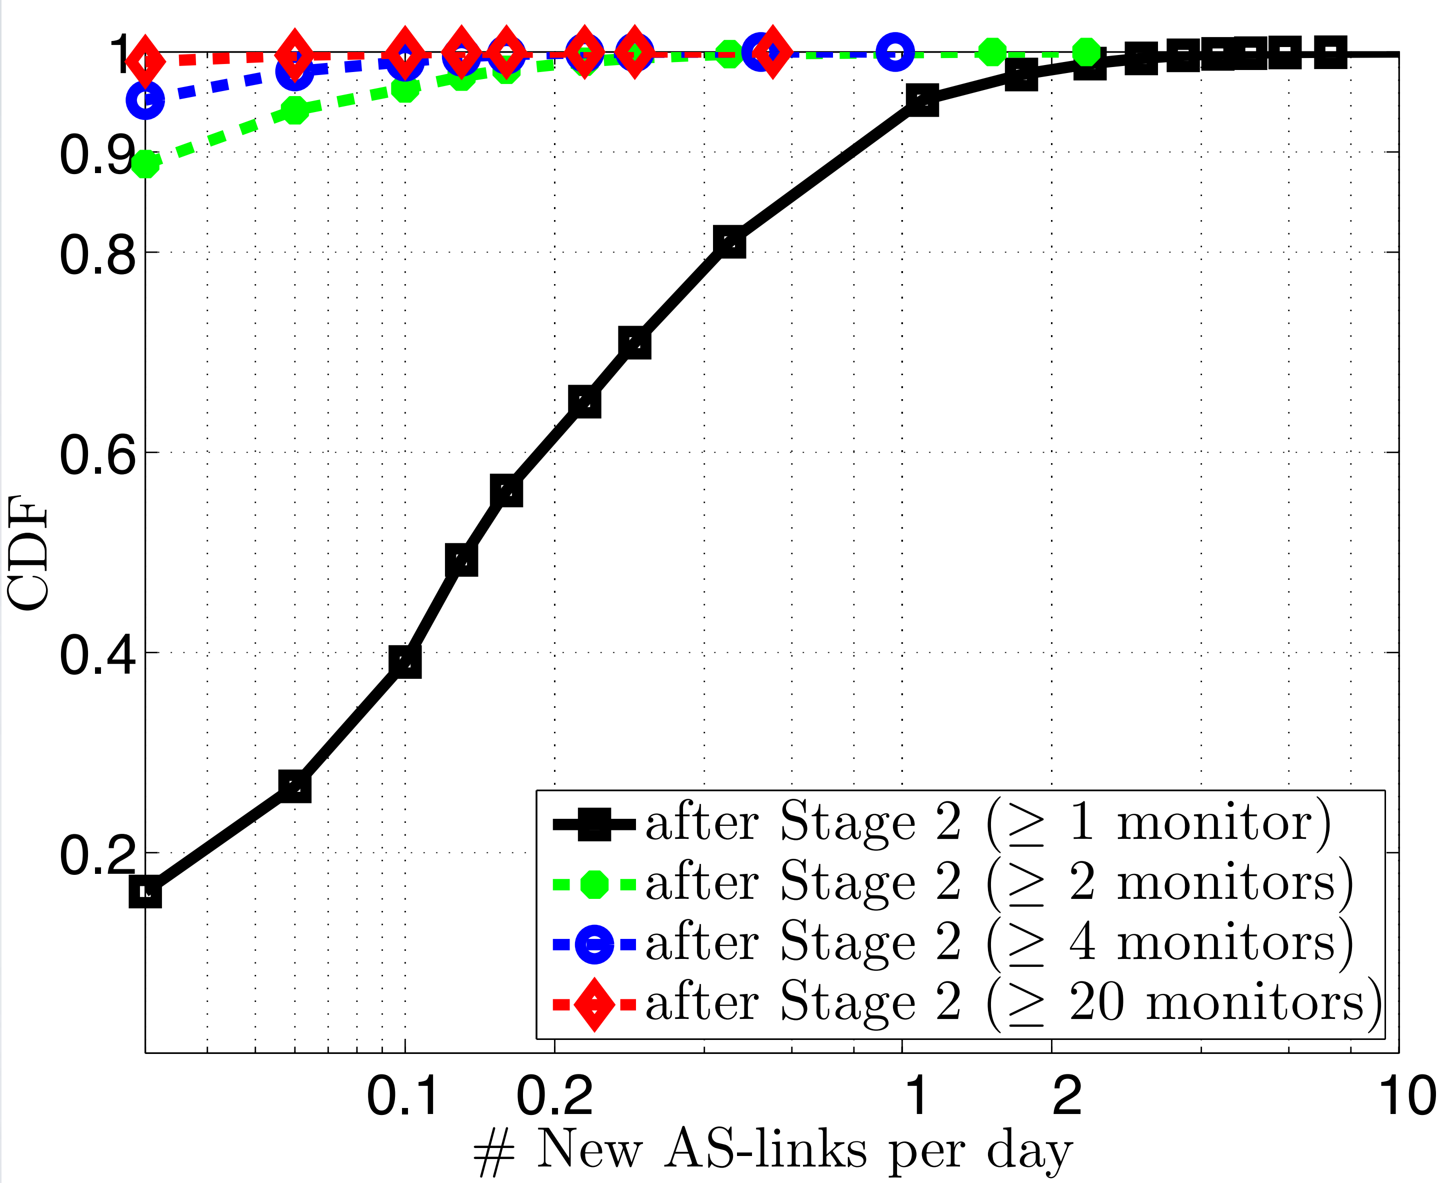

- **Stage 2**
  - wait 5 minutes
  - Recheck tables
  - Optional: decisions based on observable impact

MITIGATION

in the paper: simulation + experiments on the actual Internet

- DIY: de-aggregate while you can!
- When you can't, maybe ask help to the DoS mitigation guys

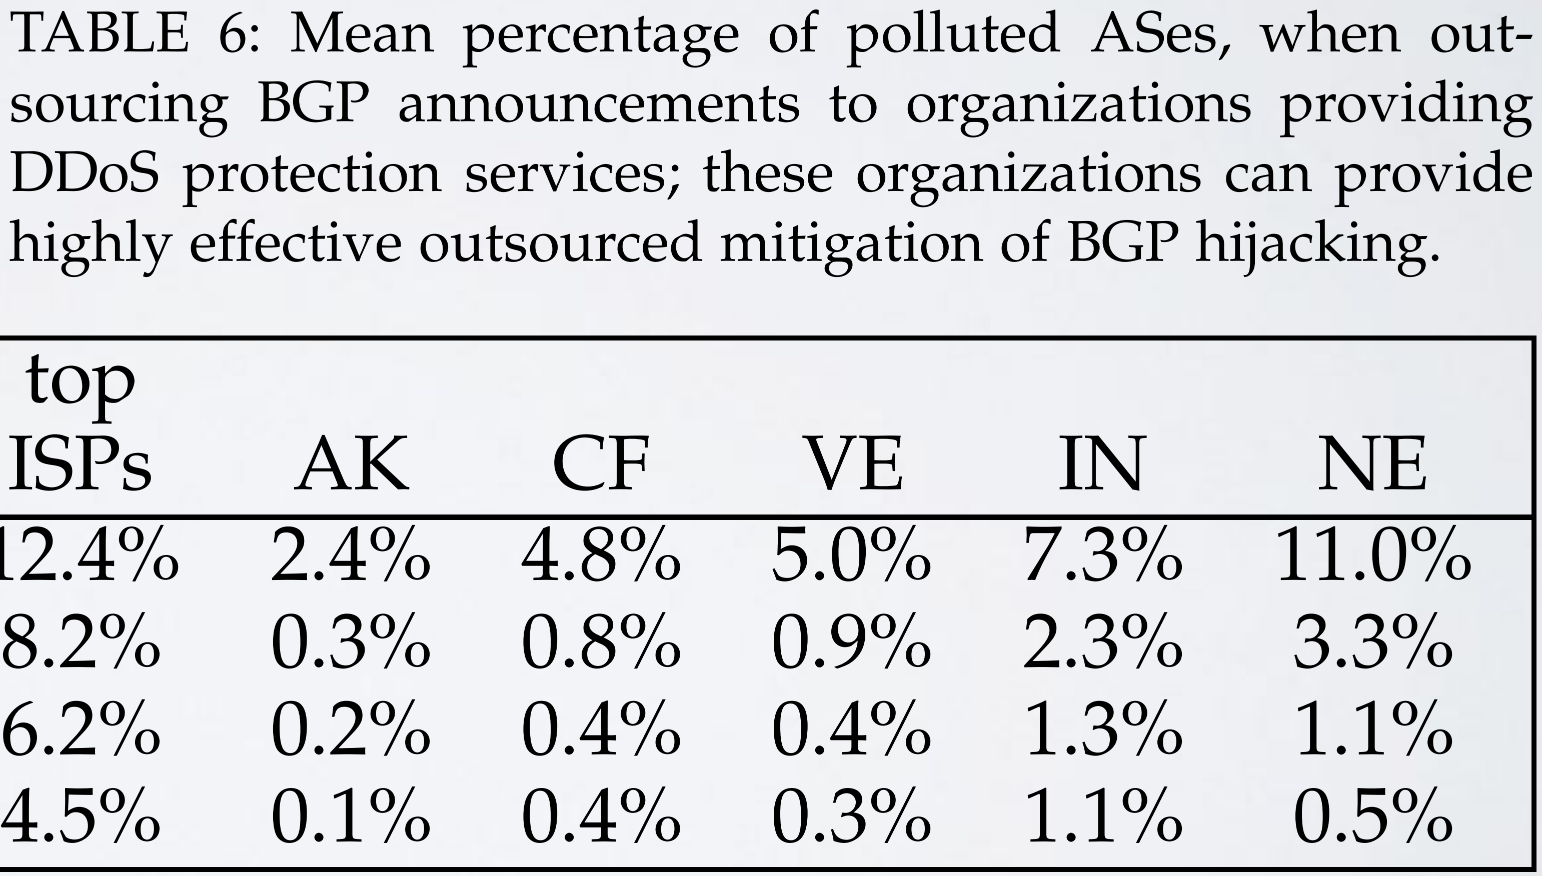

TABLE 6: Mean percentage of polluted ASes, when outsourcing BGP announcements to organizations providing DDoS protection services; these organizations can provide highly effective outsourced mitigation of BGP hijacking.

	without outsourcing	top ISPs	AK	CF	VE	IN	NE
Type0	50.0%	12.4%	2.4%	4.8%	5.0%	7.3%	11.0%
Type1	28.6%	8.2%	0.3%	0.8%	0.9%	2.3%	3.3%
Type2	16.9%	6.2%	0.2%	0.4%	0.4%	1.3%	1.1%
Type3	11.6%	4.5%	0.1%	0.4%	0.3%	1.1%	0.5%

ARTEMIS TOOL

soon available

- Open source
- Based on CAIDA BGPStream
- **EU** side of development sponsored by RIPE NCC
- Implementation challenges
  - automated configuration
  - mitigation

THANKS

alberto@caida.org