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

BGP HIJACKING

stealing/manipulating your routes

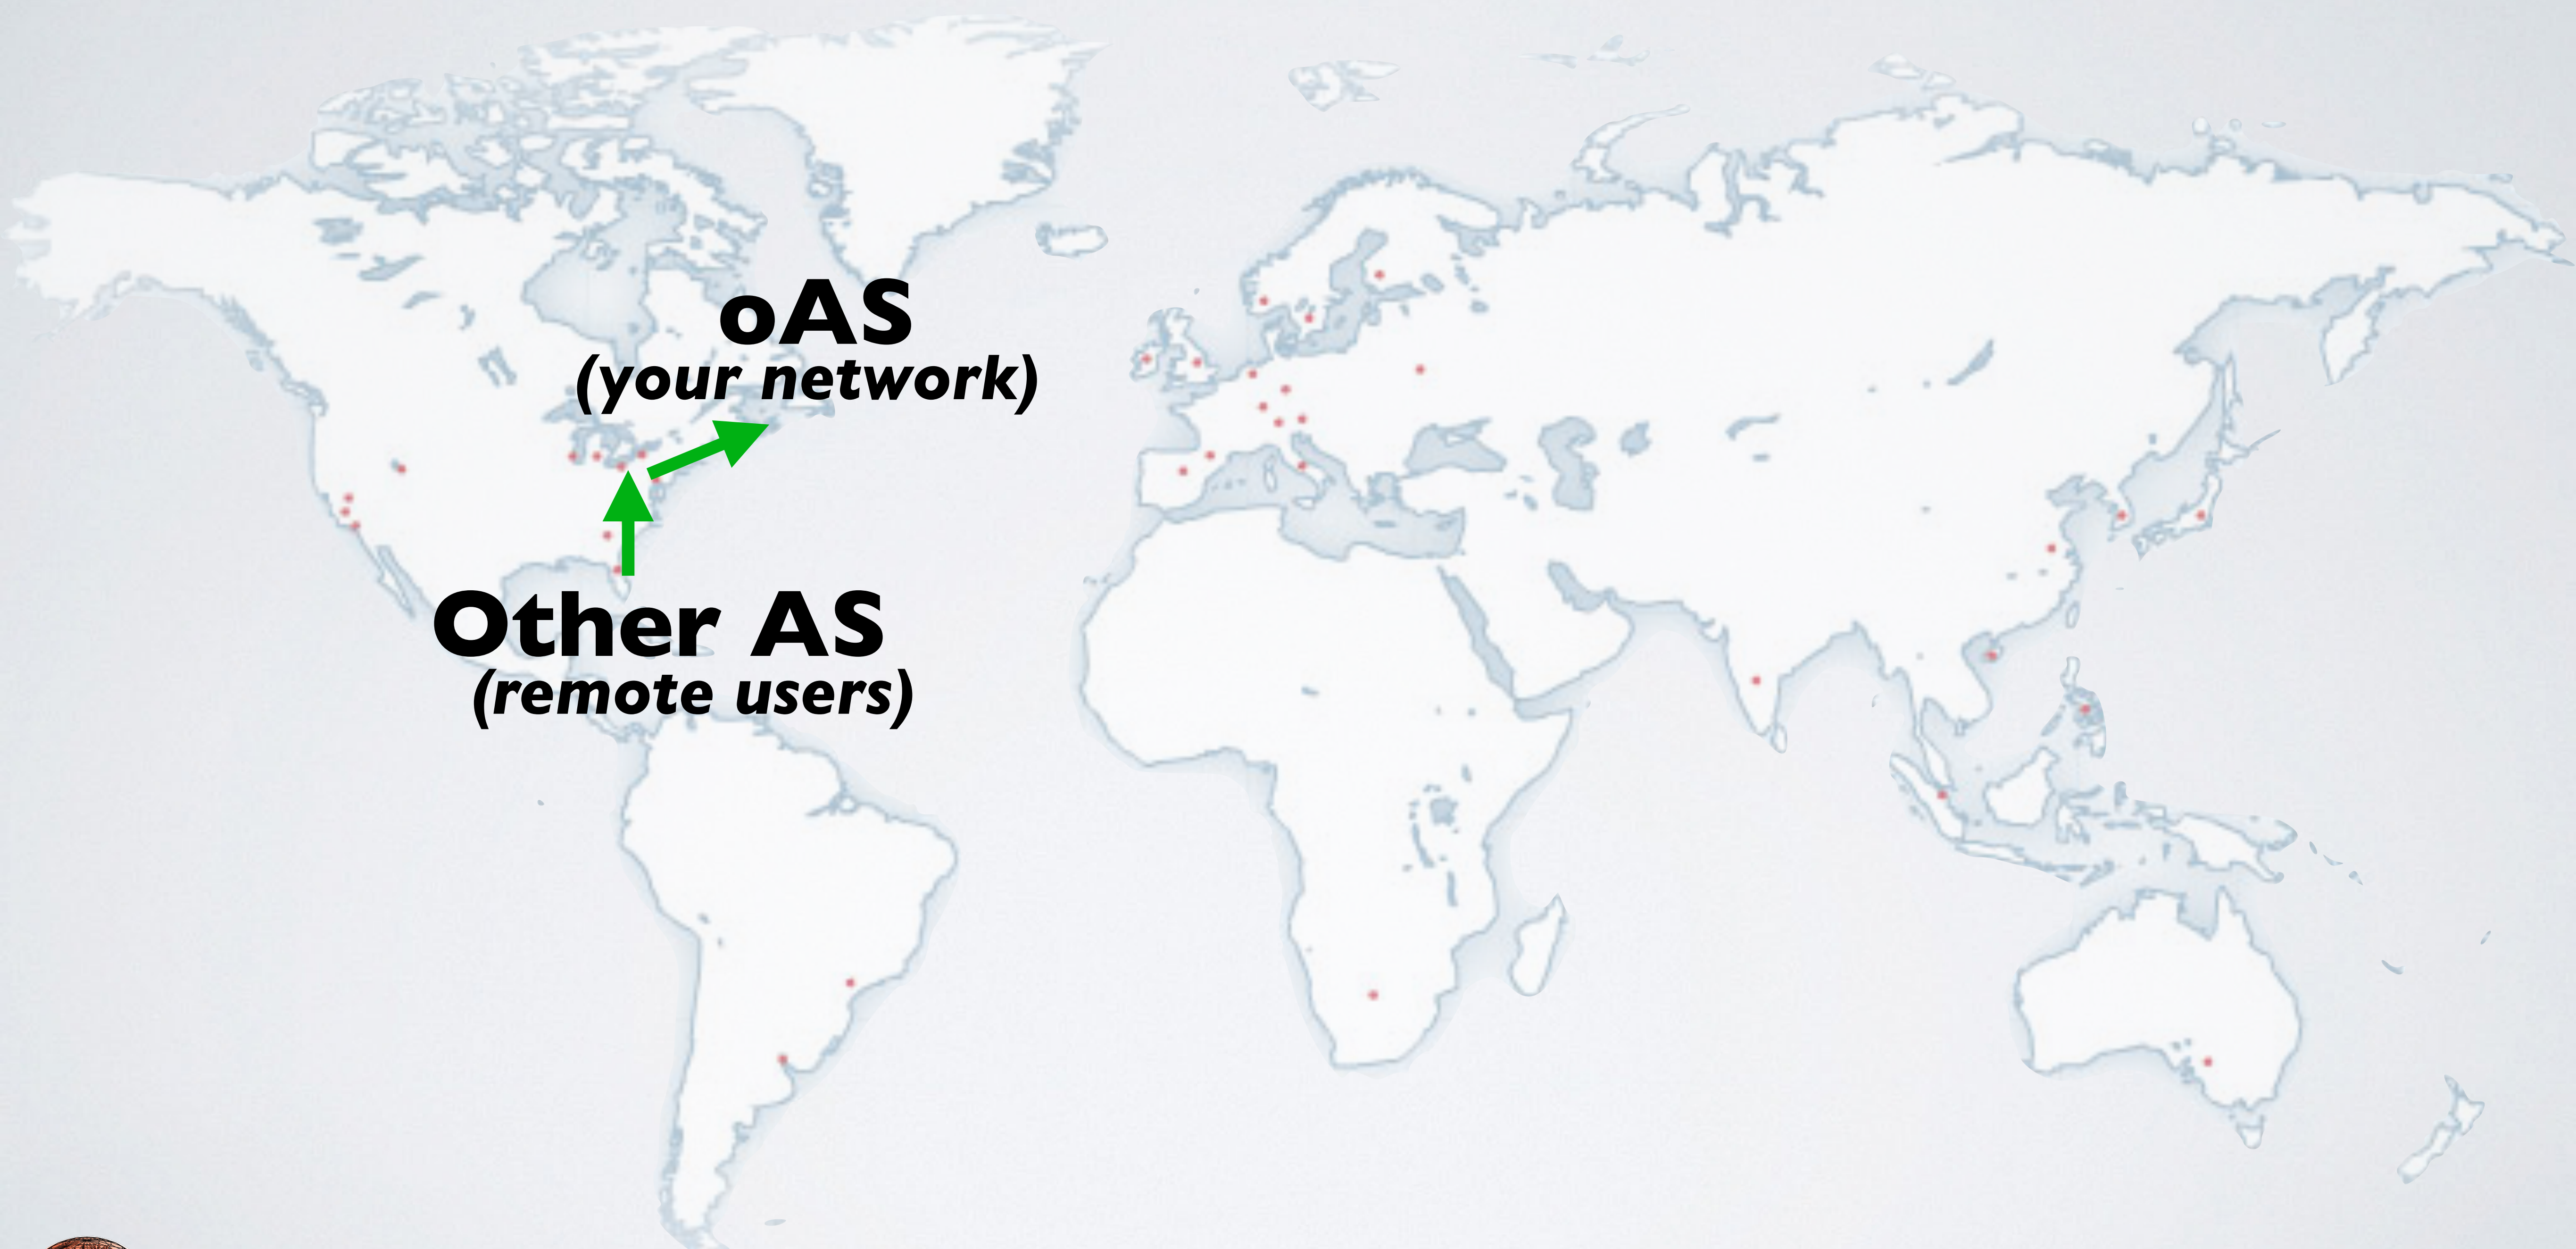

BGP HIJACKING

stealing/manipulating your routes

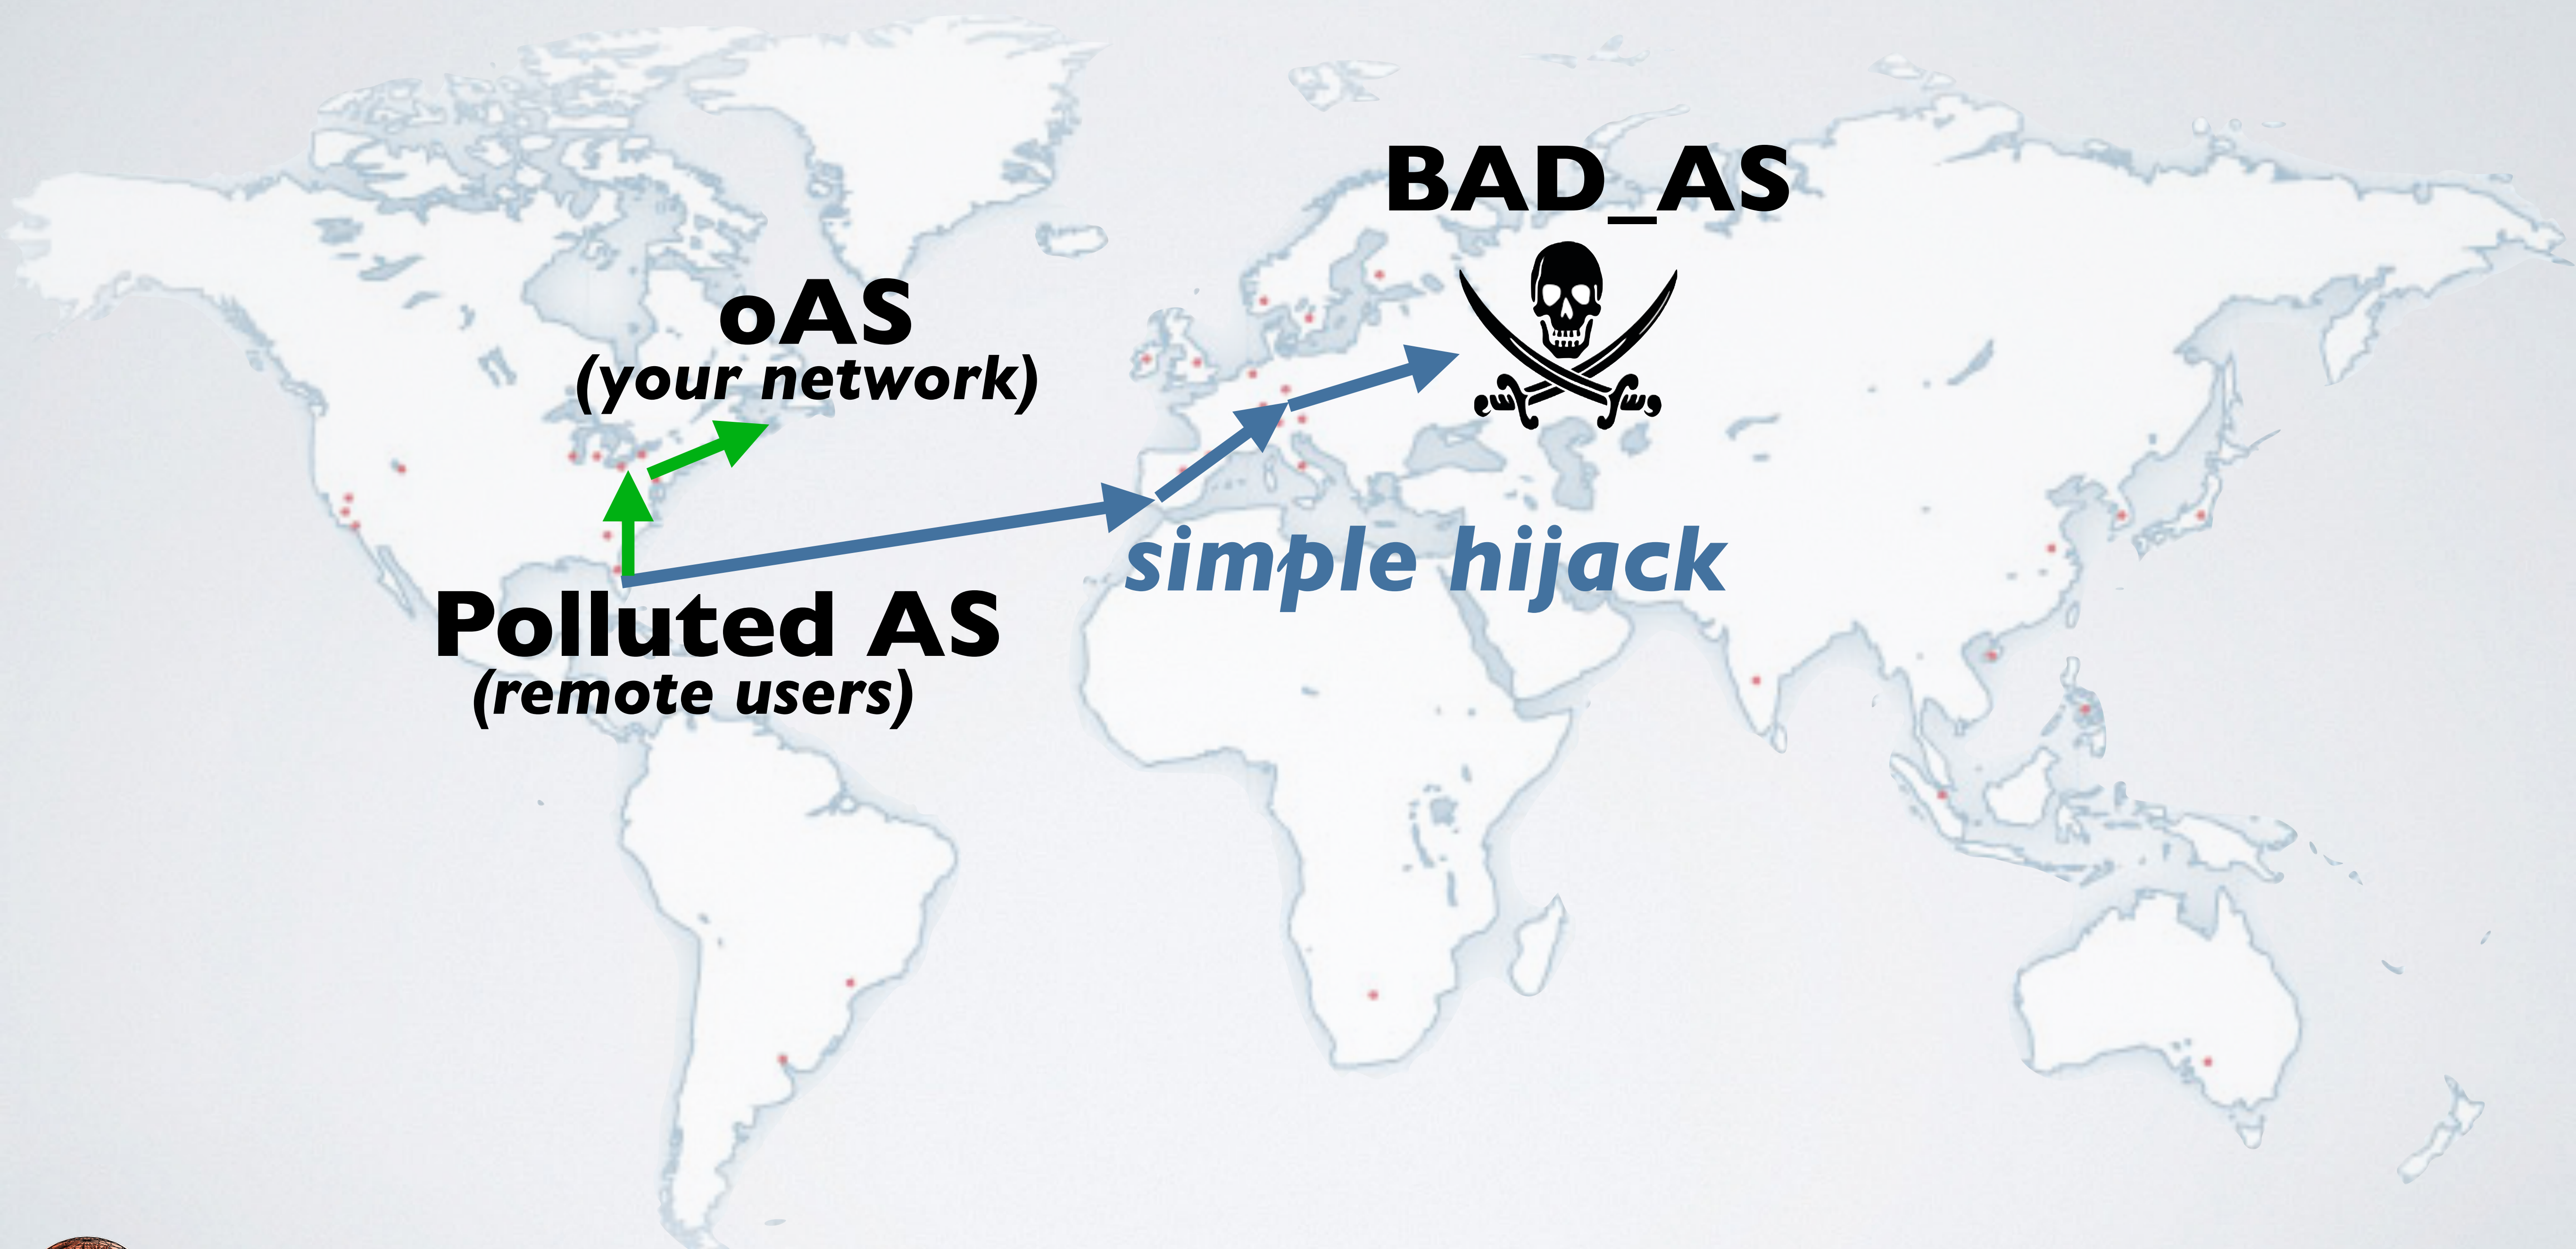

BGP HIJACKING

stealing/manipulating your routes

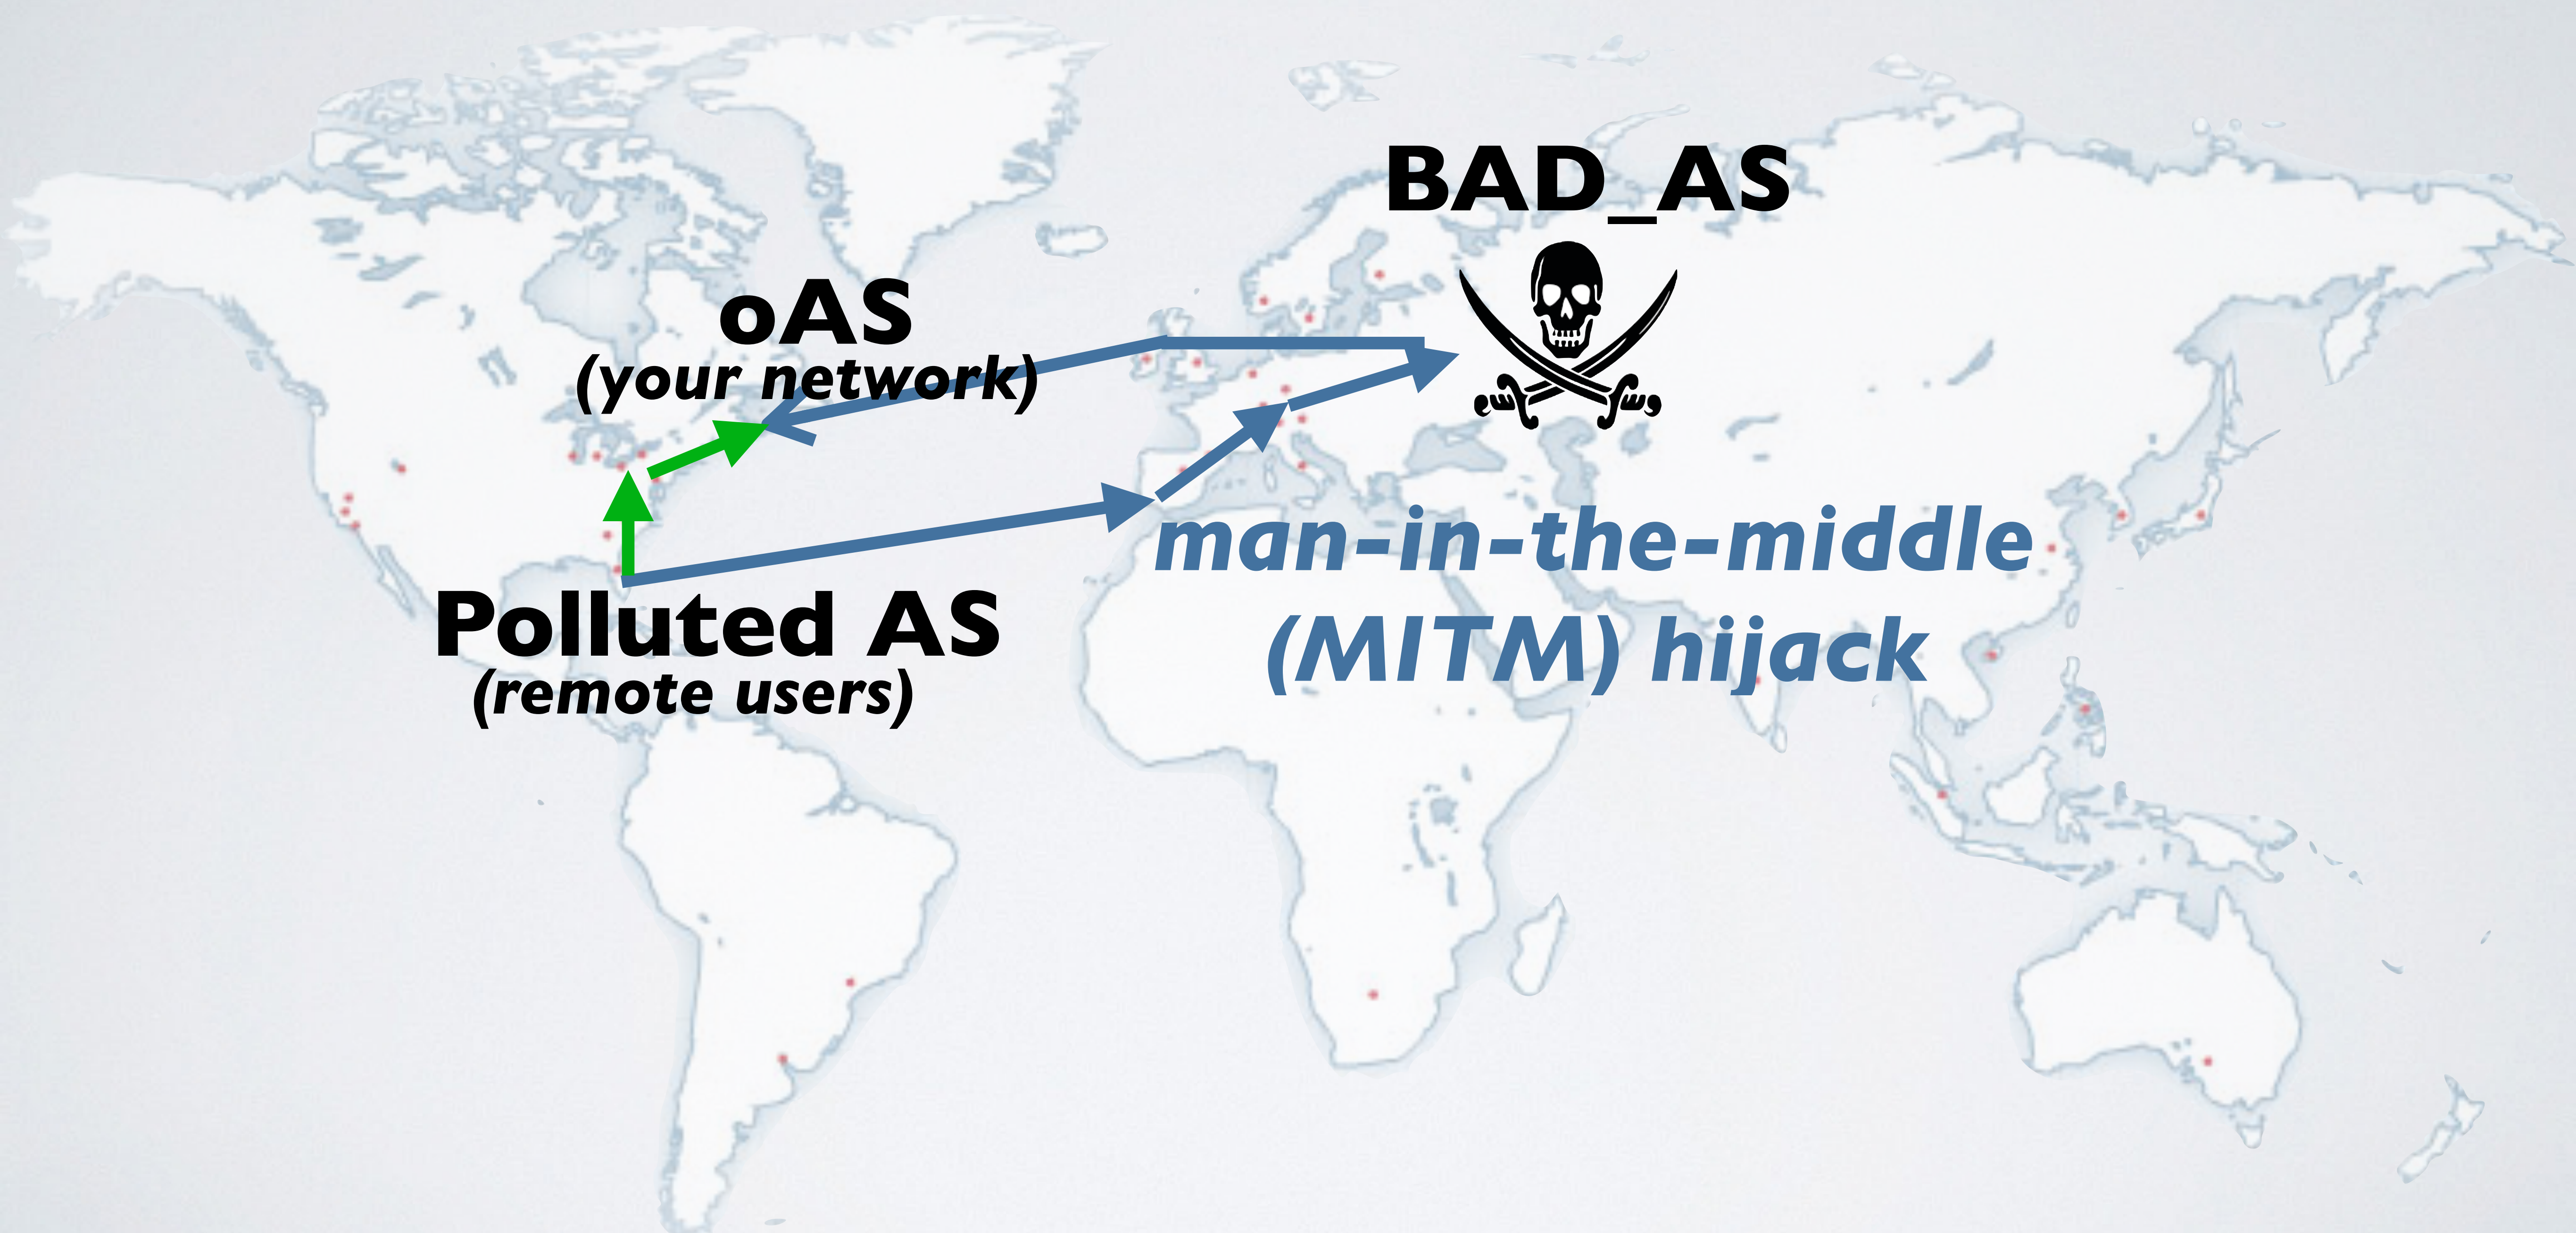

MANDATORY SLIDE WITH NEWS HEADLINES, DATES, BIG NAMES, ...

Place here your favorite recent headline

Place here your favorite recent headline

Place here your favorite recent headline

SOLUTIONS IN USE (1/2)

Proactive: RPKI

- Only 8% of prefixes covered by ROAs [1]
- Why? → limited adoption & costs/complexity [2]
- Does not protect the network against all attack types

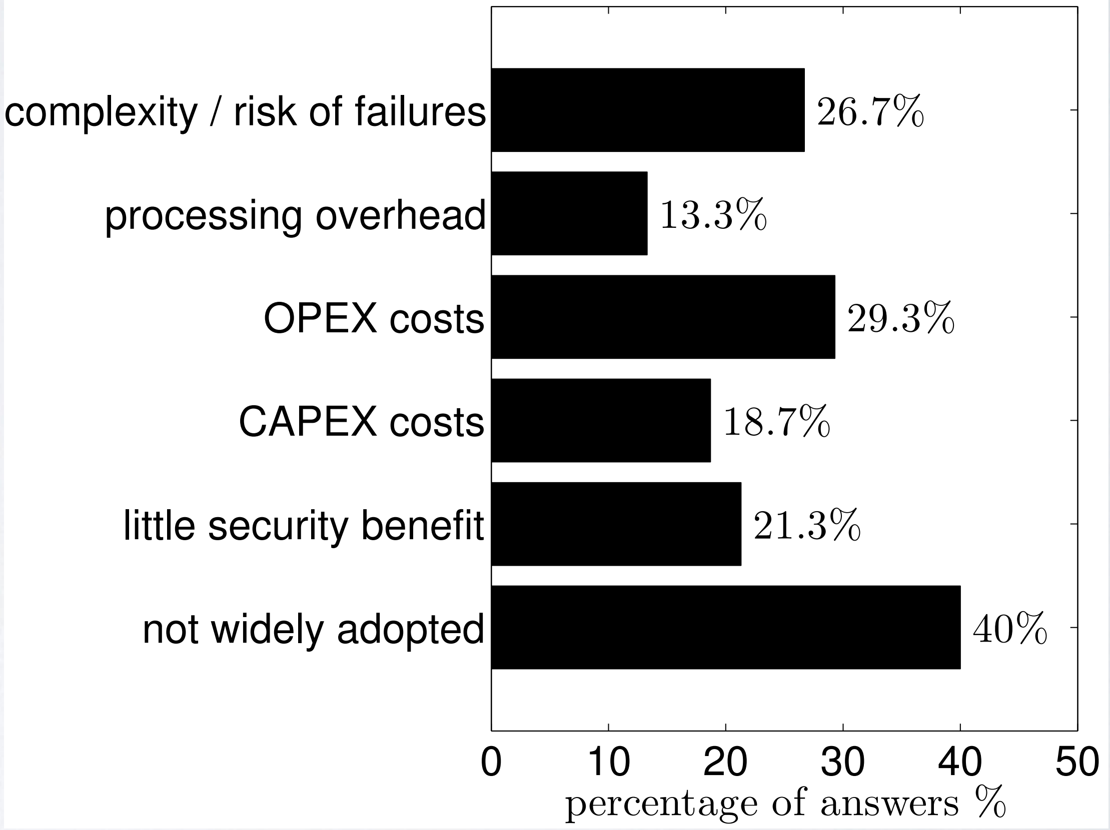

Reasons for not using RPKI [2]

[1] NIST. RPKI Monitor https://rpki-monitor.antd.nist.gov/. May 2018
[2] P. Sermpezis, et. al., "A survey among Network Operators on BGP Prefix Hijacking", in ACM SIGCOMM CCR, Jan 2018.

SOLUTIONS IN USE (2/2)

Reactive: 3rd Party Services

- **Comprehensiveness**: detect only simple attacks
- **Accuracy**: prone to false positives (FP) & false negatives (FN)
- **Speed**: manual verification & then manual mitigation
- **Privacy**: need to share private info, routing policies, *etc.*

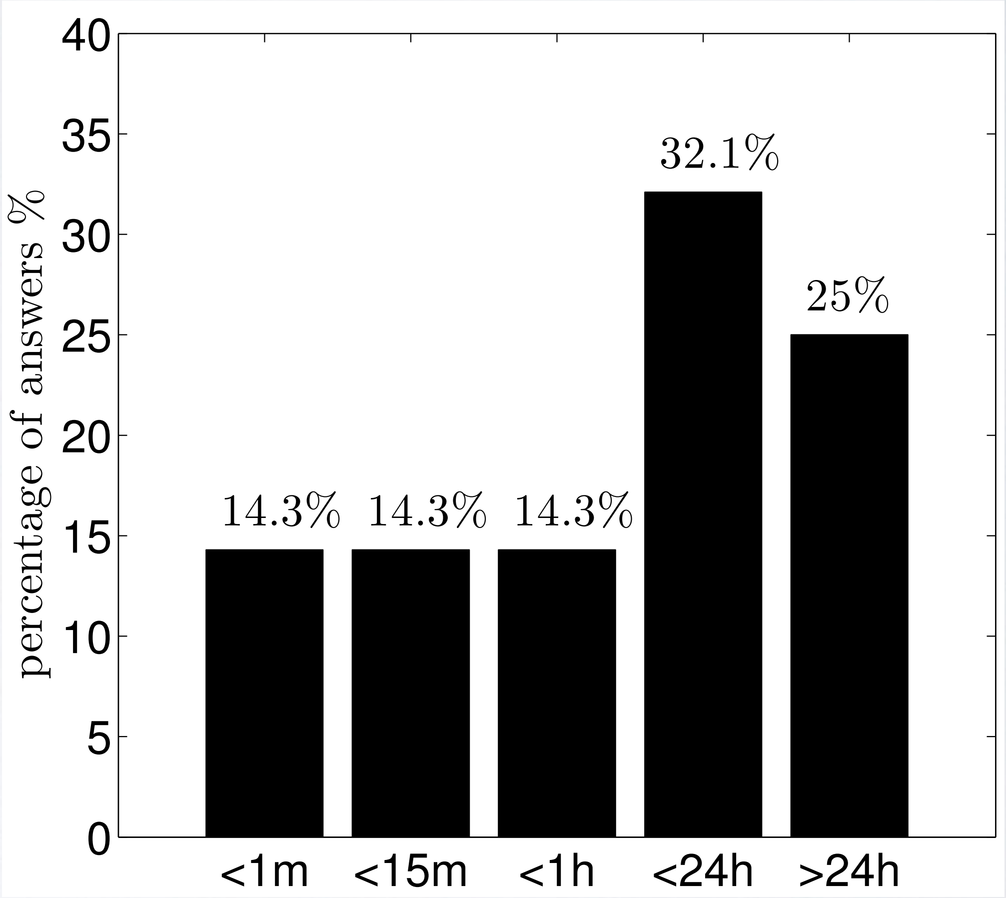

How much time an operational network was affected by a hijack [2]

[2] P. Sermpezis, et. al., "A survey among Network Operators on BGP Prefix Hijacking", in ACM SIGCOMM CCR, Jan 2018.

ARTEMIS

self-managed detection & mitigation

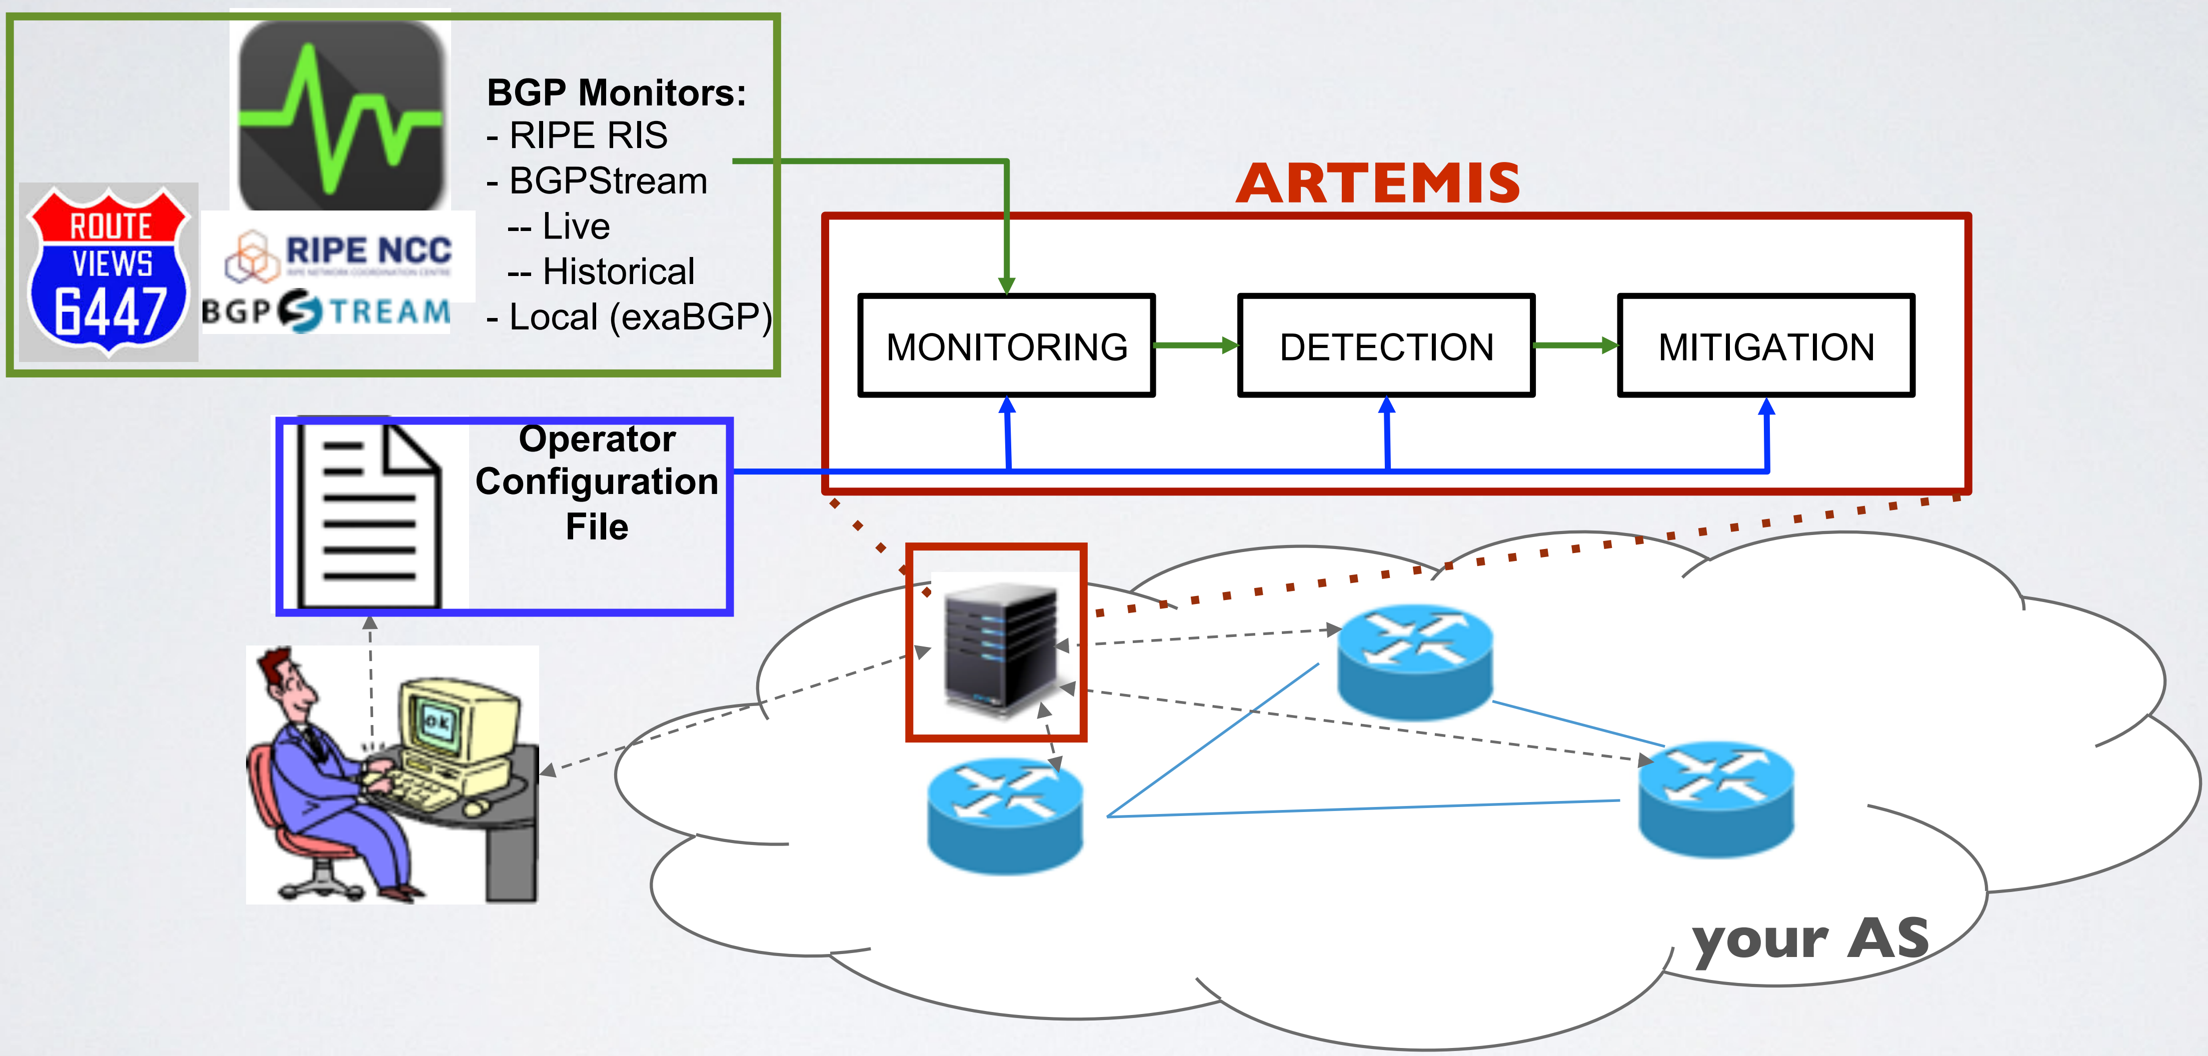

A VIEW SHIFT..

..and suddenly everything makes sense

3rd Party

- **Evasion**
  - Detect only simple attacks
- **Accuracy**
  - Potential for lots of *FPs*
  - or alternatively lots of *FNs*
- **Speed**
  - Manual verification & then manual mitigation
- **Privacy**
  - Need to share private information

ARTEMIS

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

PUBLIC MONITORING INFRASTRUCTURE

enables visibility of all significant events

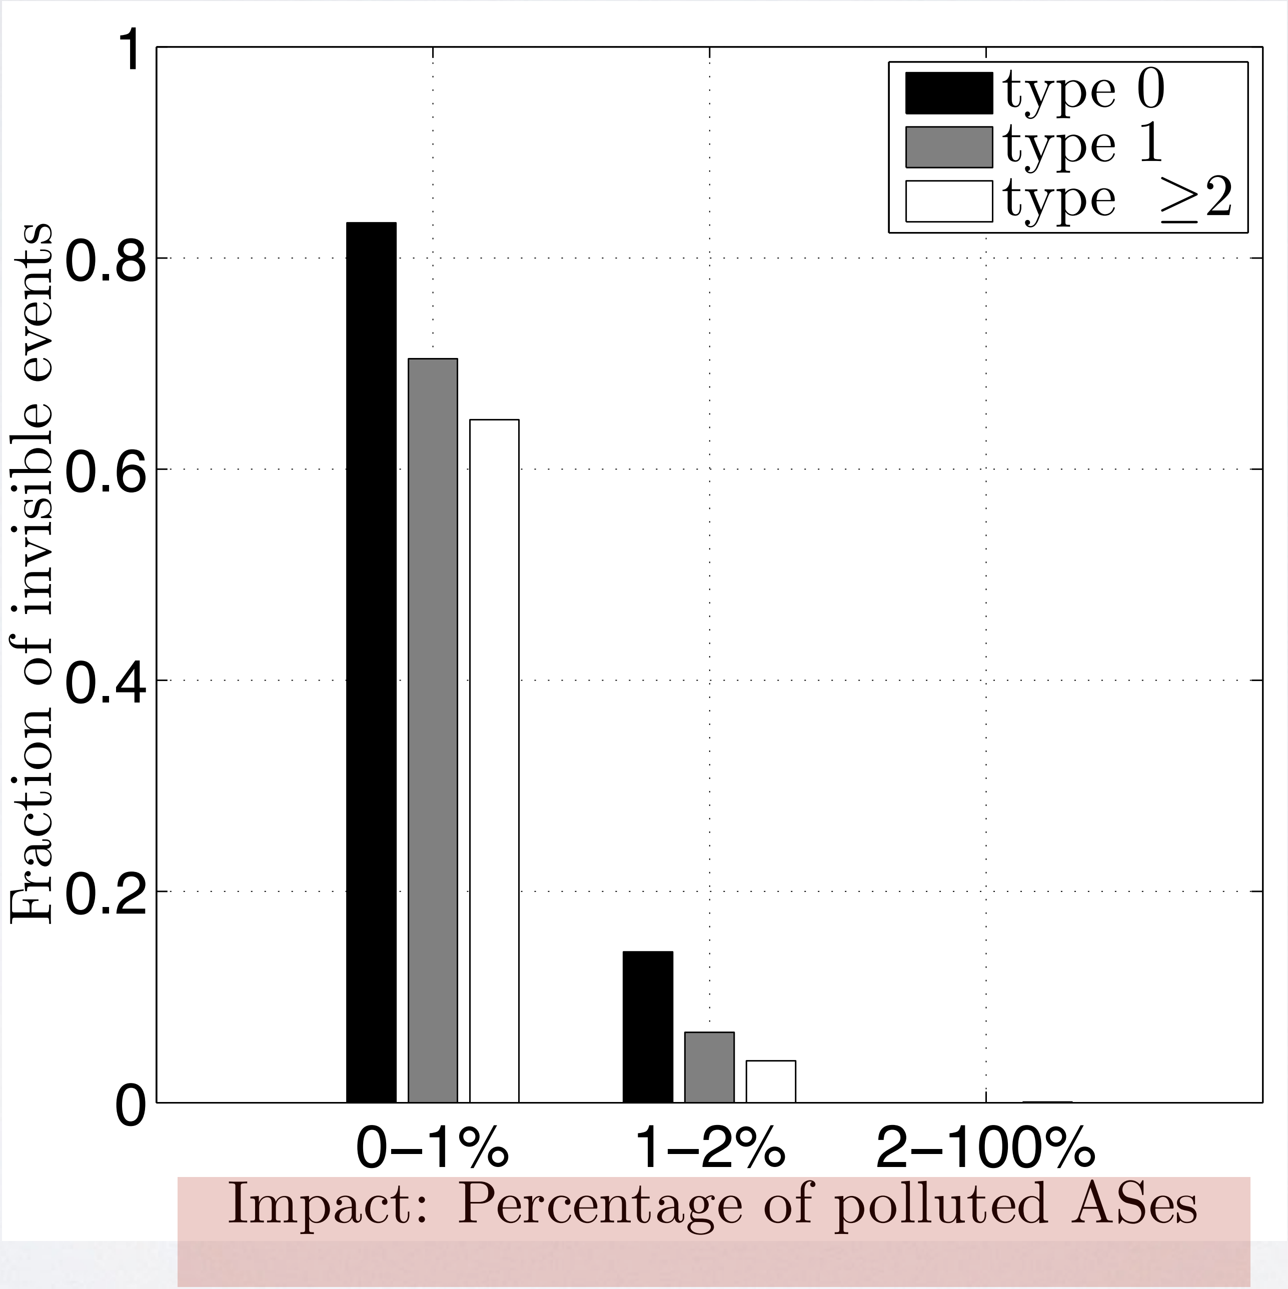

- In the paper:
  - by type of service
  - Impact
  - Speed

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: …,* ***BAD_AS****, oAS*>
  - **Type 2** hijack: <*prefix: …,* ***BAD_AS****, AS1,oAS*>
  - …
  - **Type N** hijack: <*prefix: …,* ***BAD_AS****, … AS1, oAS*>
  - **Type U** hijack: <*prefix: unaltered_path*>

- **2)** Based on the prefix announced: ***exact****,* ***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 Rule	Needed Local Information	Detection Approach
Sub-prefix	*	*	None	None	Config. vs BGP updates	Pfx.	Sec. 5.2
Squatting	*	*	None	None	Config. vs BGP updates	Pfx.	Sec. 5.2
Exact	0/1	*	None	None	Config. vs BGP updates	Pfx. + ASN (+ neighbor ASN)	Sec. 5.3

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 Rule	Needed Local Information	Detection Approach
Sub-prefix	*	*	None	None	Config. vs BGP updates	Pfx.	Sec. 5.2
Squatting	*	*	None	None	Config. vs BGP updates	Pfx.	Sec. 5.2
Exact	0/1	*	None	None	Config. vs BGP updates	Pfx. + ASN (+ neighbor ASN)	Sec. 5.3
Exact	≥ 2	*	< 0.3/day for $> 73\%$ of ASes	None	Past Data vs BGP updates (bidirectional link)	Pfx.+ Past AS links	Sec. 5.4 *Stage 1*
Exact	≥ 2	*	None for 63% of ASes ($T_{s2} = 5min$, $th_{s2} > 1$ monitors)	$< 4\%$	BGP updates (waiting interval, bidirectional link)	Pfx.	Sec. 5.4 *Stage 2*

hard problem in remaining cases (fake link 2 hops or more from origin + exact prefix hijack)

FAKE LINK (TYPE ≥ 2) HIJACKS

Detection: Stage 1

• Triggered when the AS-PATH of a BGP update (for a monitored prefix) 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).

NOW: <*your prefix:* …, ***ASX, ASY***, *oAS*> *announcement with new link attached to 1-hop neighbor ASY*

HISTORY: <*any prefix:* …, ***ASY, ASX***, …> *reverse link exists; it was announced by ASY*

FAKE LINK (TYPE ≥ 2) HIJACKS

Detection: Stage I

- Only way for an attacker to fake a link in the opposite direction is to announce a loop

NOW:

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

HISTORY:

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

- Can be evaded though, if the attacker controls more than one AS

HISTORY:

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

Center for Applied Internet Data Analysis
University of California San Diego
caida

Foundation for Research and Technology-Hellas
University of Crete,
FORTH
INSTITUTE OF COMPUTER SCIENCE

FAKE LINK (TYPE ≥ 2) HIJACKS

Detection: Stage 1 - there is more..

- We also require that there is **no common ASN** appearing in **each and every** observed AS path on the left of *(i)* the new link and on the left of *(ii)* the reverse link in the history

NOW: <*your prefix:* …, ***BAD_AS, ASX, ASY***, *oAS*> *announcement with new link*

HISTORY: <*any prefix:* …, ***ASY, ASX***, …> *e.g., there is at least one path without BAD_AS*

FAKE LINK (TYPE ≥ 2) HIJACKS

Detection: Stage 1

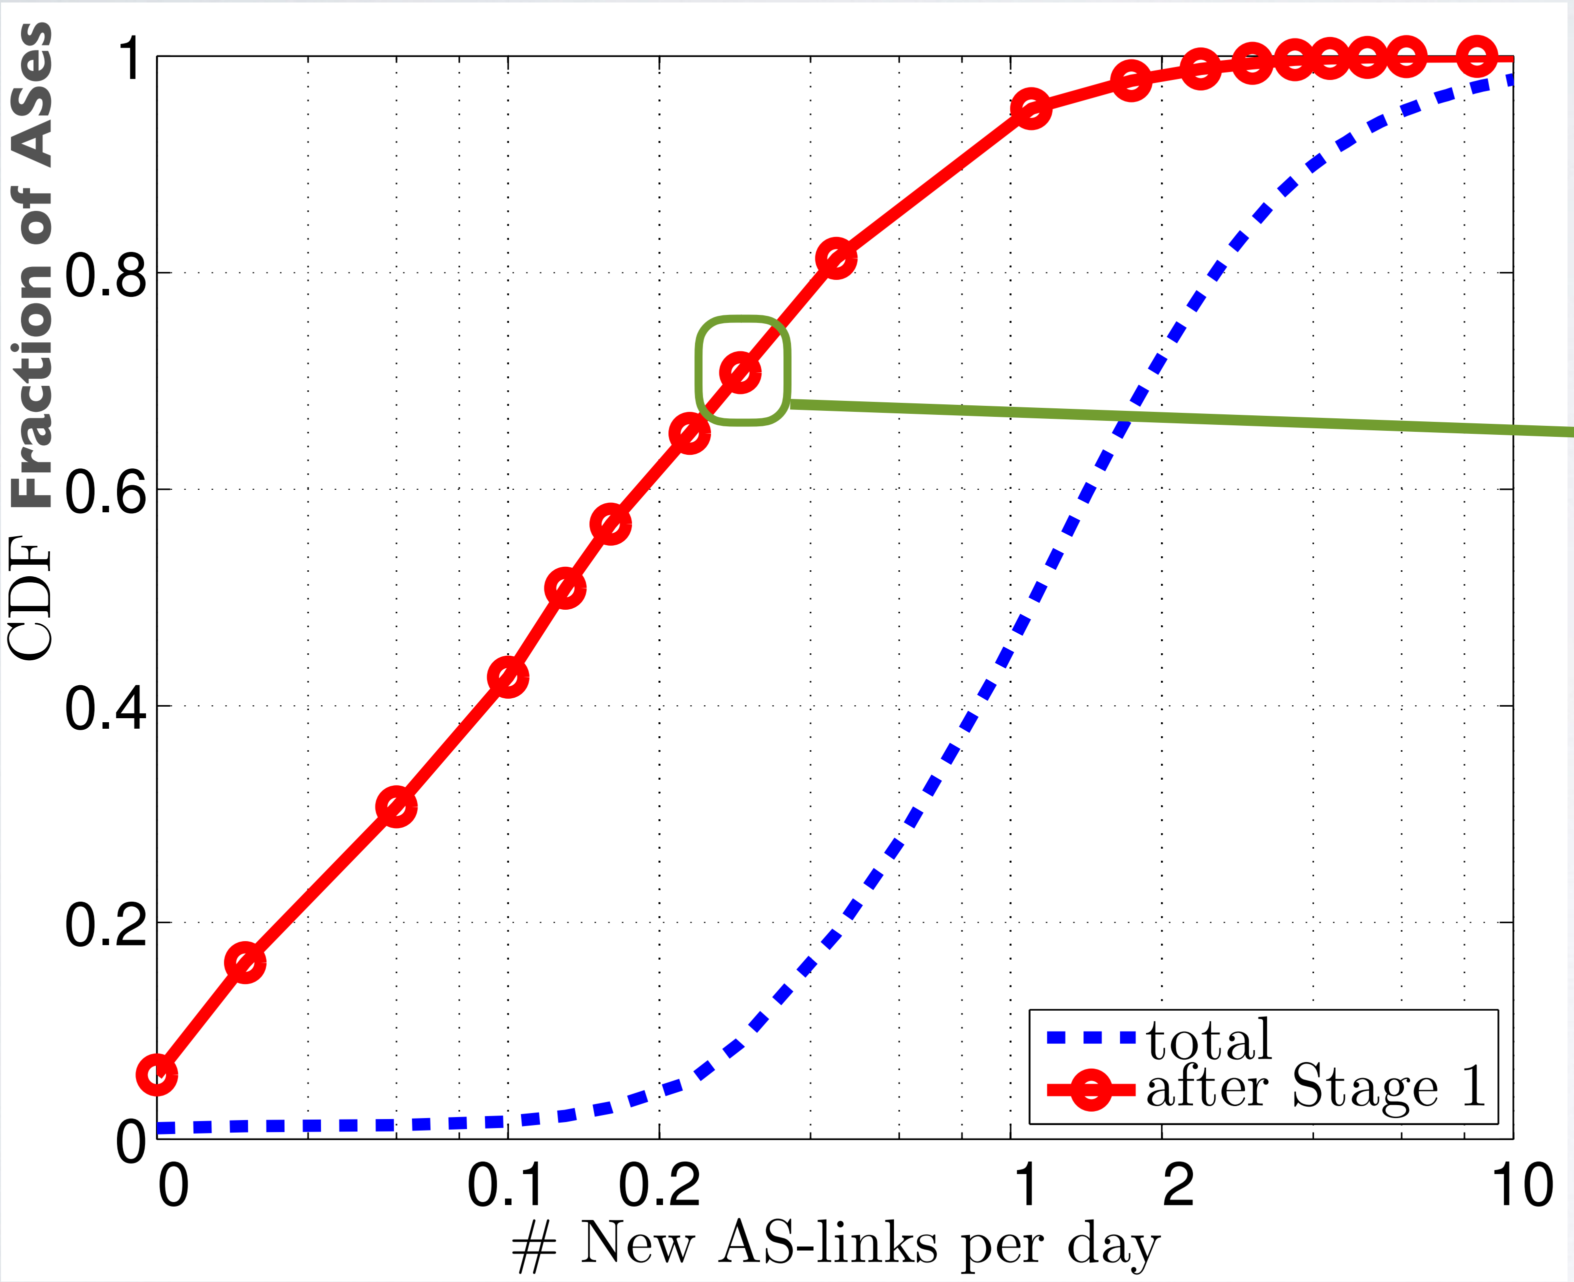

We emulated ARTEMIS Stage1 for 30 days for *each* AS originating prefixes in March 2017 (data from 438 monitors)

73% of the ASes saw less than 1 suspicious event every 3 days

FAKE LINK (TYPE ≥ 2) HIJACKS

Detection: Stage 2

- Trades latency for additional info
- Wait 5 min (*configurable*) to:
  1. Leverage new information from monitors and local routers
     ~30% improvement (in simulation) w/ data from local routers
  2. **Estimate the impact** of the event based on how many monitors see it
  3. Can be configured to not generate alert (or alert only but not auto-mitigate, etc.) for events with low impact
     Trades removing **FPs** for potential **FNs** w/ small impact

FAKE LINK (TYPE ≥ 2) HIJACKS

Detection: Stage 2

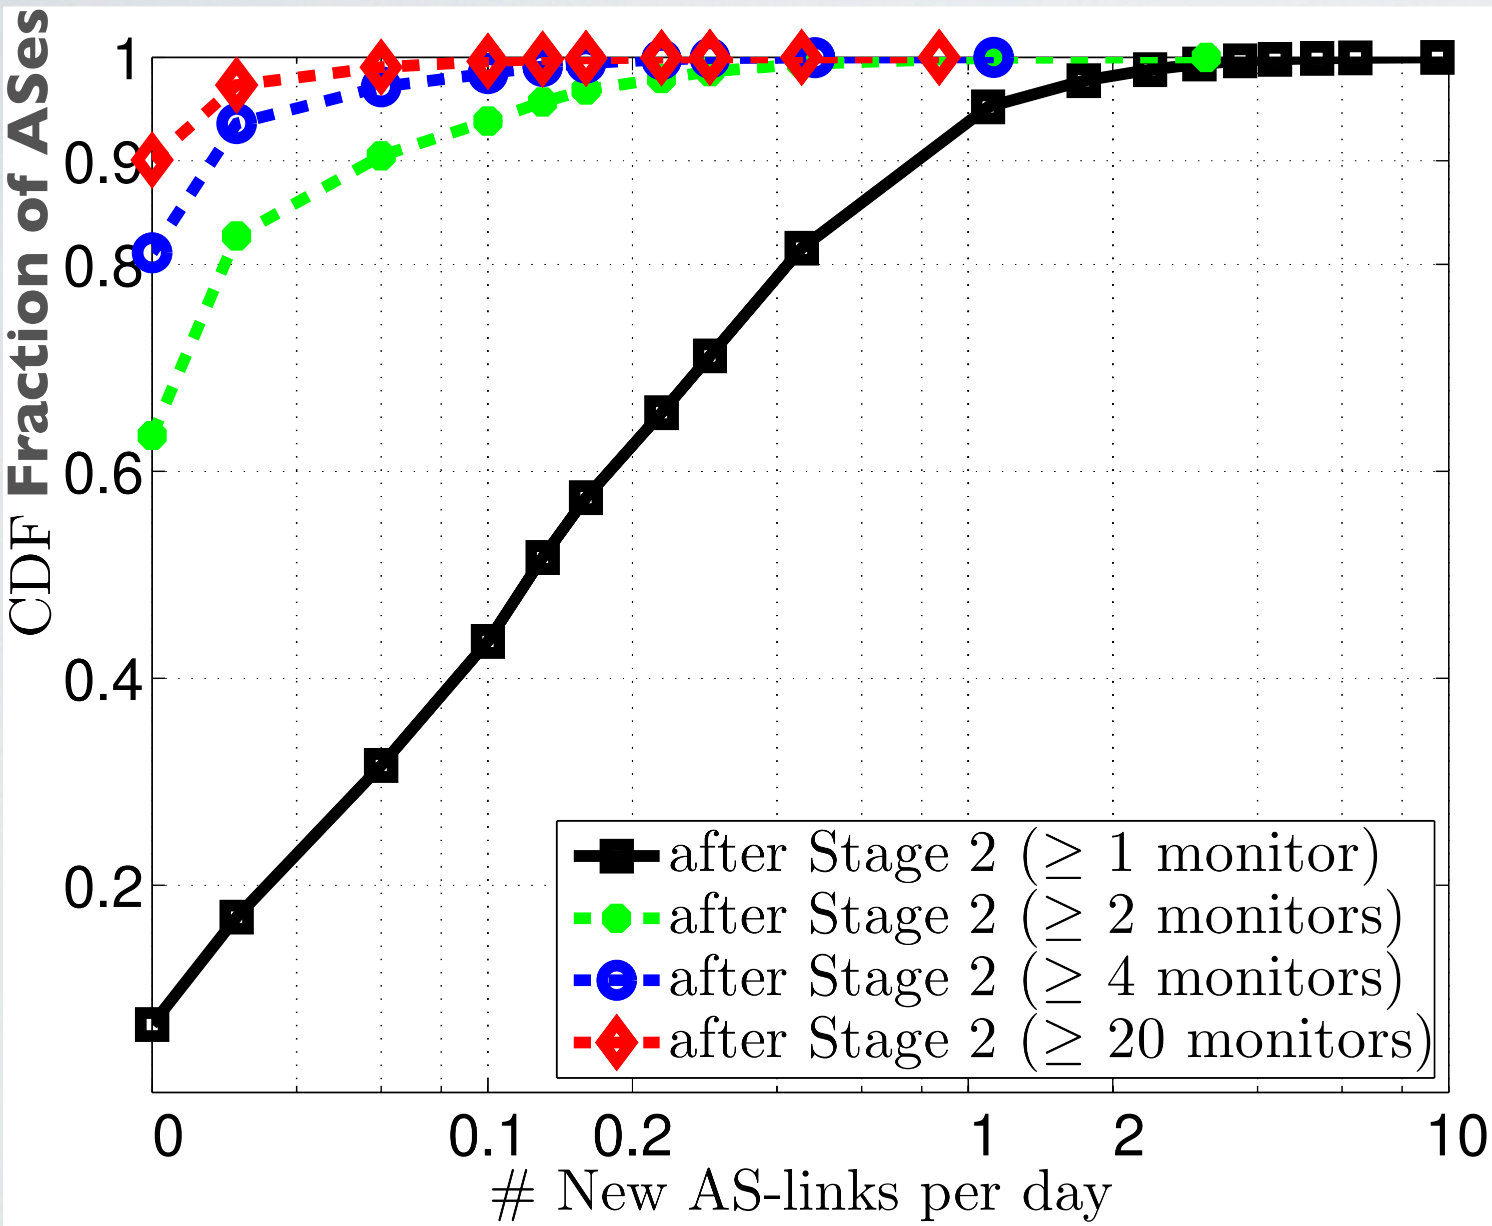

We emulated ARTEMIS Stage1+2 for 30 days for each AS originating prefixes in March 2017 (data from 438 monitors)

The majority of the "unverified new links" that pass Stage 1 are seen by only 1 monitor

If, e.g., the operator decides to ignore *[or treat differently]* events seen by < 4 monitors (blue curve) the vast majority (81%) of ASes would not see a single *[relevant]* alert in the whole month

MITIGATION

in the paper: simulations + experiments on the actual Internet

- DIY: de-aggregate while you can!
  - only possible down to /24 granularity

- When you can't, maybe ask help to the DoS mitigation guys

Percentage of polluted ASes when fighting an exact-prefix hijack without or with outsourcing to large ISPs or DoS mitigators

	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%

OPENSOURCE ARTEMIS TOOL

stay tuned - work in progress

- open source
- based on CAIDA BGPStream
- Devel partially sponsored by "RIPE NCC Community Projects 2017"
- Implementation challenges
  - automated configuration
  - mitigation

RIPE NCC
RIPE NETWORK COORDINATION CENTRE

THANKS

alberto@caida.org

https://arxiv.org/abs/1801.01085

http://www.inspire.edu.gr/artemis/