Munich, March 25th 2026

IEEE 4th Conference on Secure and Trustworthy Machine Learning

ConCap: Practical Network Traffic Generation for (ML- and) Flow-based Intrusion Detection Systems

Miel Verkerken, Laurens d'Hooge, Bruno Volckaert, Filip de Turck, Giovanni Apruzzese

HÁSKÓLINN Í REYKJAVÍK | REYKJAVIK UNIVERSITY

Munich, March 25th 2026

IEEE 4th Conference on Secure and Trustworthy Machine Learning

ConCap: Practical Network Traffic Generation for (ML- and) Flow-based Intrusion Detection Systems

This is not about Generative AI

Miel Verkerken, Laurens d'Hooge, Bruno Volckaert, Filip de Turck, Giovanni Apruzzese

HÁSKÓLINN Í REYKJAVÍK | REYKJAVIK UNIVERSITY

Network Intrusion Detection Systems

ML / AI / DL

Giovanni Apruzzese
giovannia@ru.is

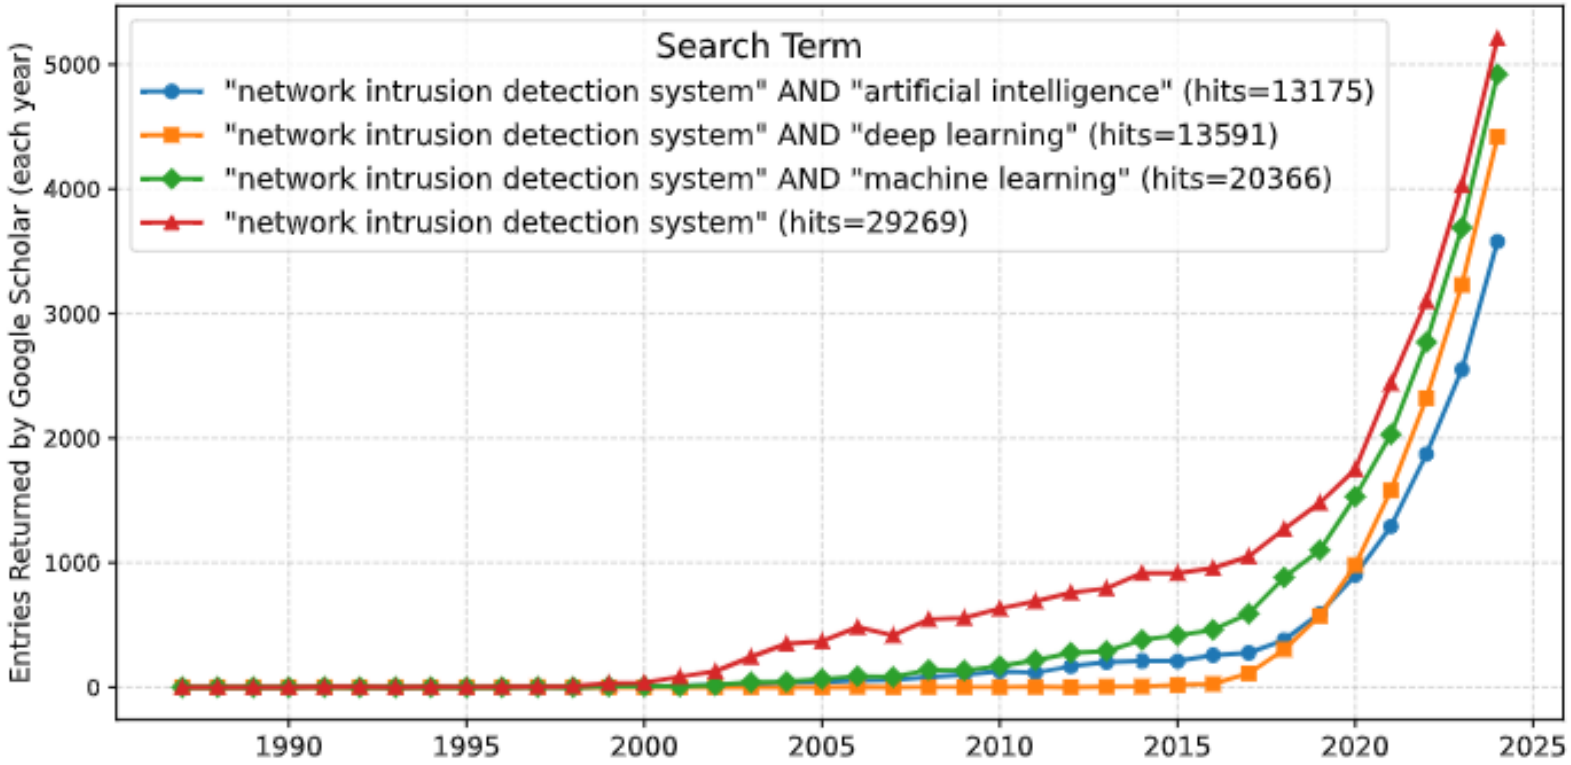

Research interest in NIDS. Queries issued on September 2025.

Giovanni Apruzzese
giovannia@ru.is

The following paper was originally published in the
Proceedings of the 7th USENIX Security Symposium
San Antonio, Texas, January 26-29, 1998

Data Mining Approaches for Intrusion Detection

Wenke Lee and Salvatore J. Stolfo
Columbia University

Giovanni Apruzzese
giovannia@ru.is

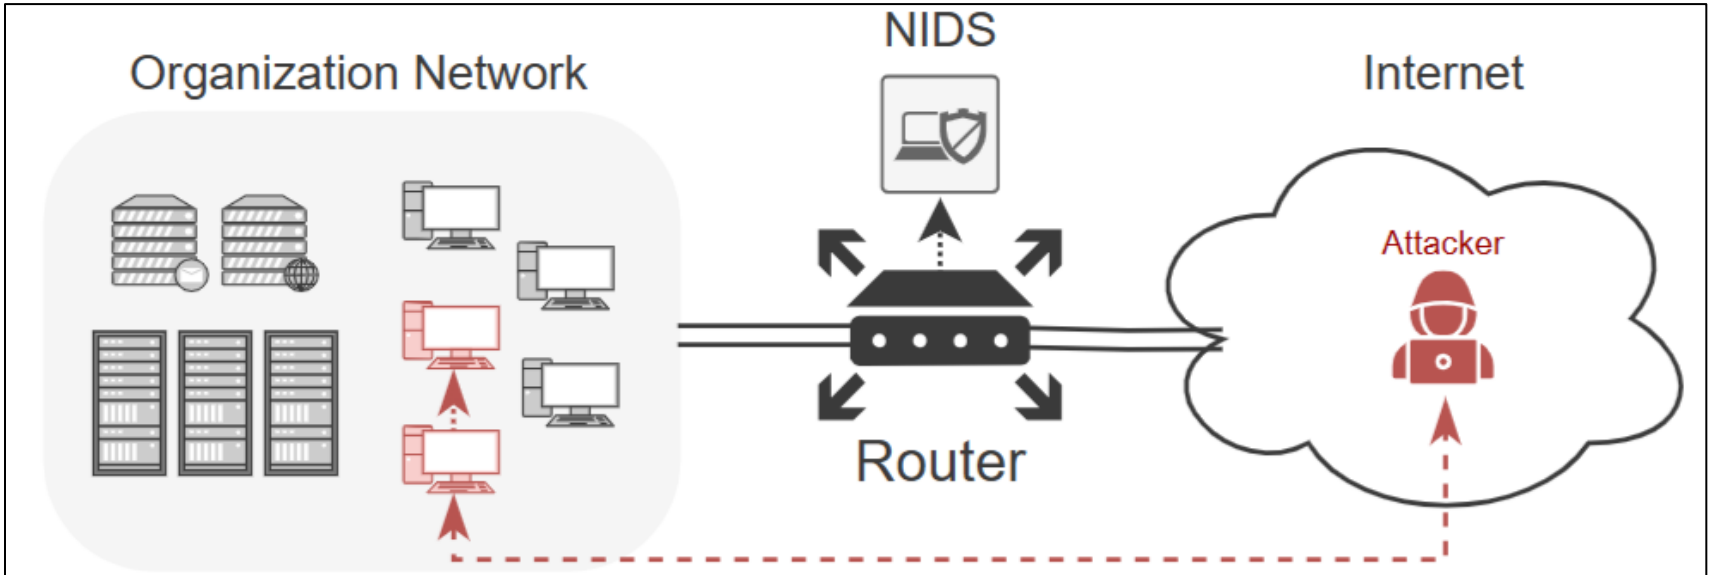

Fig. 1: **Exemplary deployment scenario of an NIDS.**

Giovanni Apruzzese
giovannia@ru.is

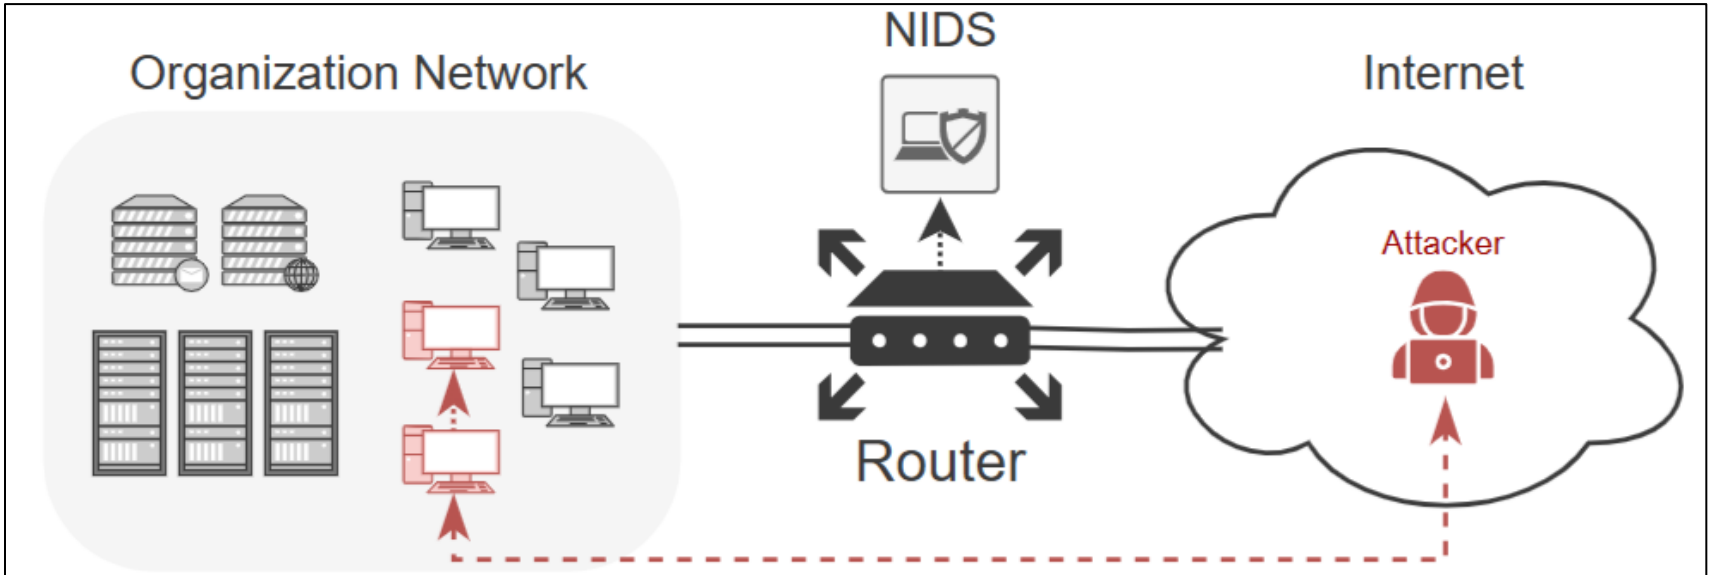

Fig. 1: **Exemplary deployment scenario of an NIDS.**

How can a researcher obtain network-related data for (ML-)NIDS experiments??

Giovanni Apruzzese
giovannia@ru.is

Obtaining network-related data

a. Real-world capture

b. Simulation

c. ‘Benchmark’ Datasets

Giovanni Apruzzese
giovannia@ru.is

Obtaining network-related data

a. Real-world capture

b. Simulation

c. ‘Benchmark’ Datasets ☺

Giovanni Apruzzese
giovannia@ru.is

Obtaining network-related data

a. Real-world capture → availability?

b. Simulation → realism?

c. ‘Benchmark’ Datasets ☺

Giovanni Apruzzese
giovannia@ru.is

2021 IEEE Symposium on Security and Privacy Workshops

Troubleshooting an Intrusion Detection Dataset: the CICIDS2017 Case Study

Gints Engelen, Vera Rimmer, and Wouter Joosen
imec-DistriNet, KU Leuven
Leuven, Belgium
Email:{*firstname.lastname*}@kuleuven.be

Giovanni Apruzzese
giovannia@ru.is

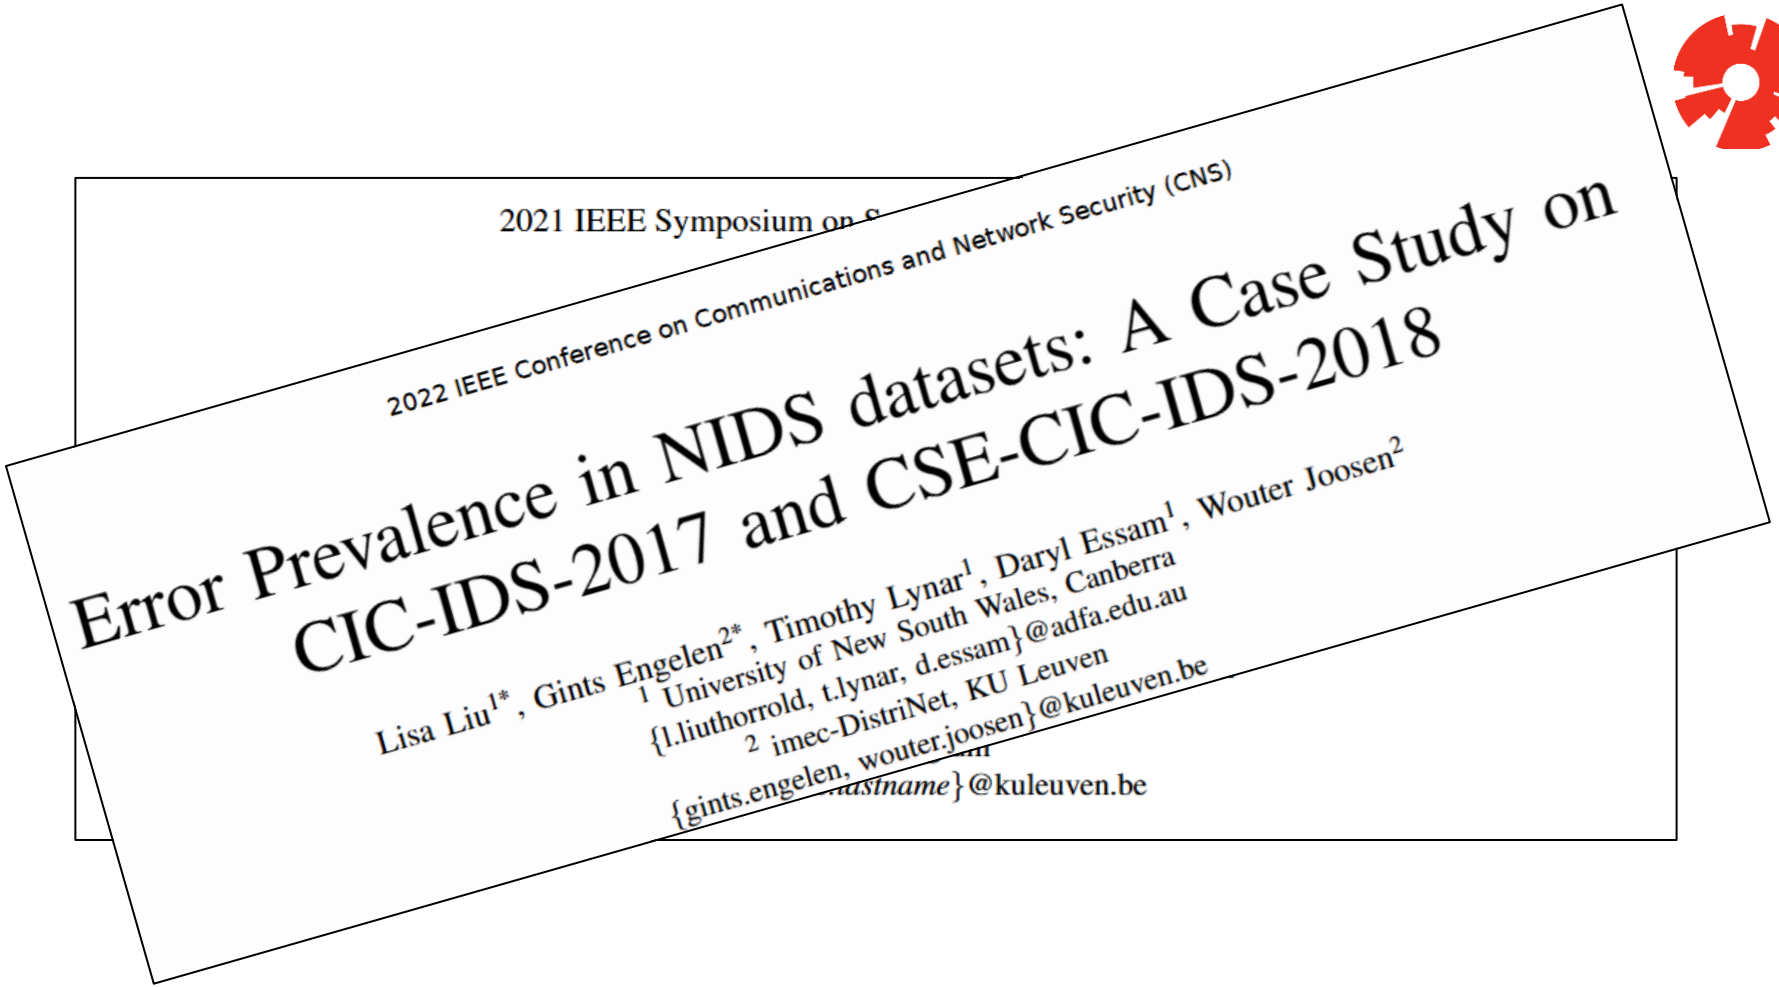

Giovanni Apruzzese
giovannia@ru.is

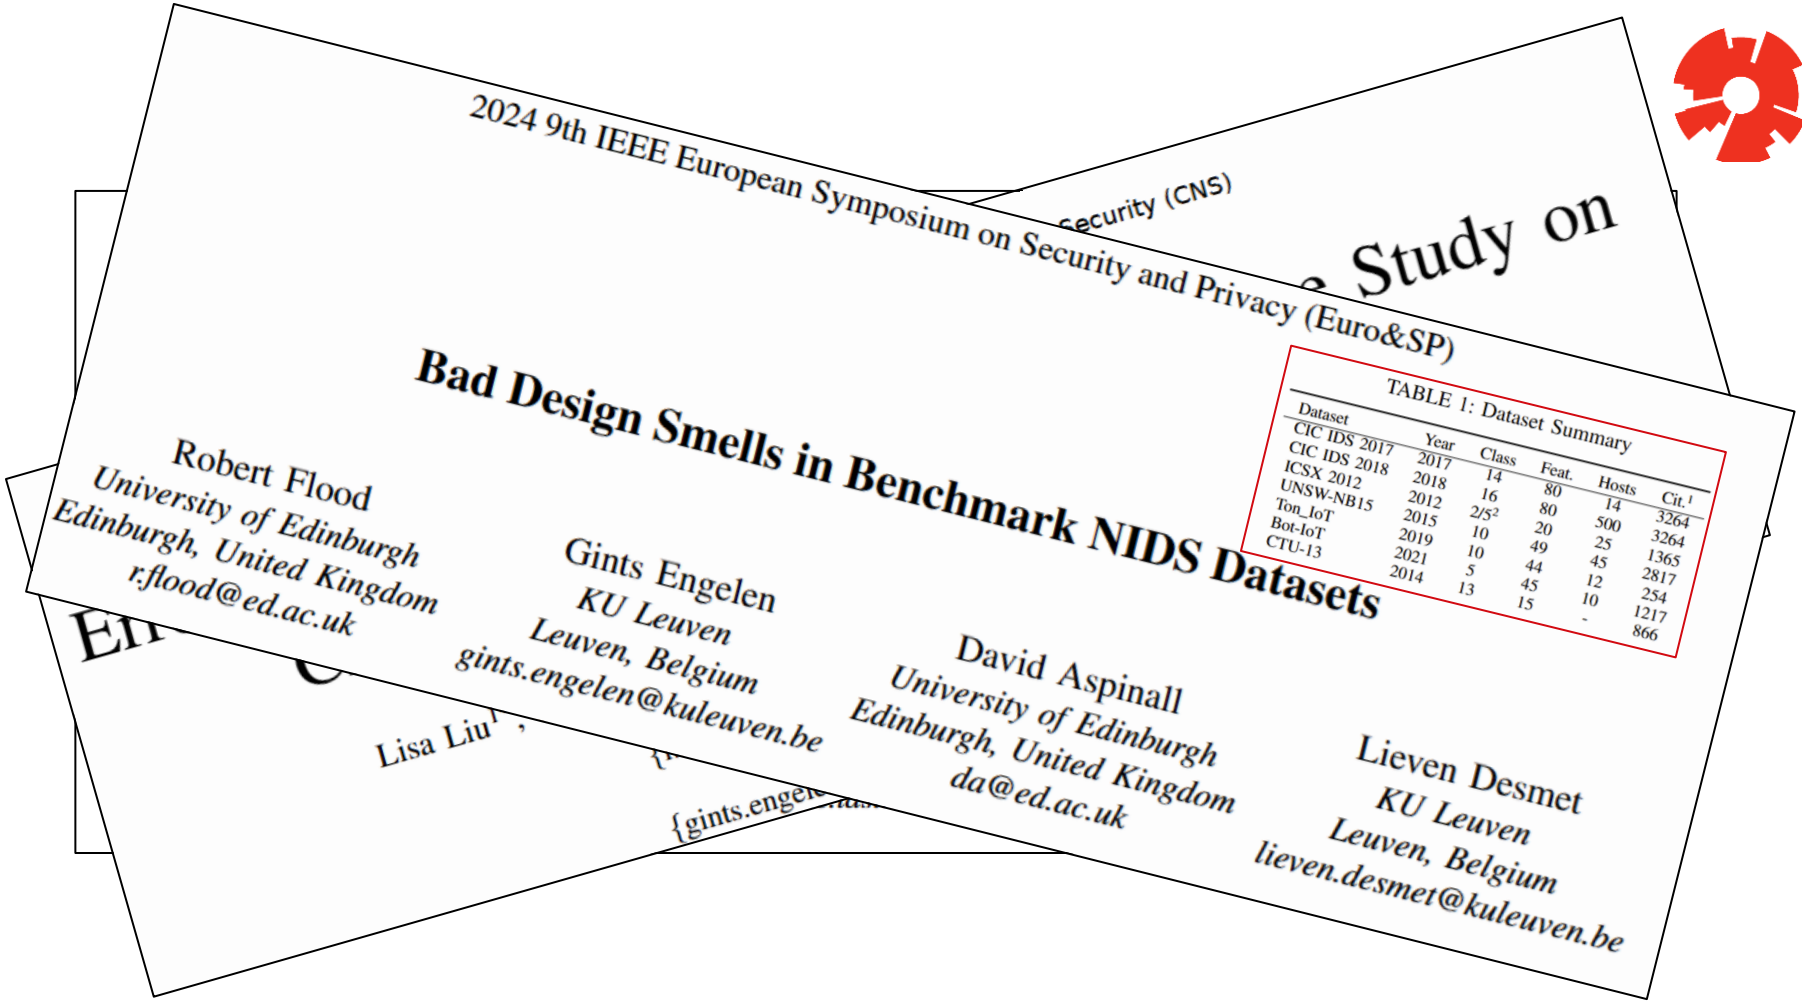

Giovanni Apruzzese
giovannia@ru.is

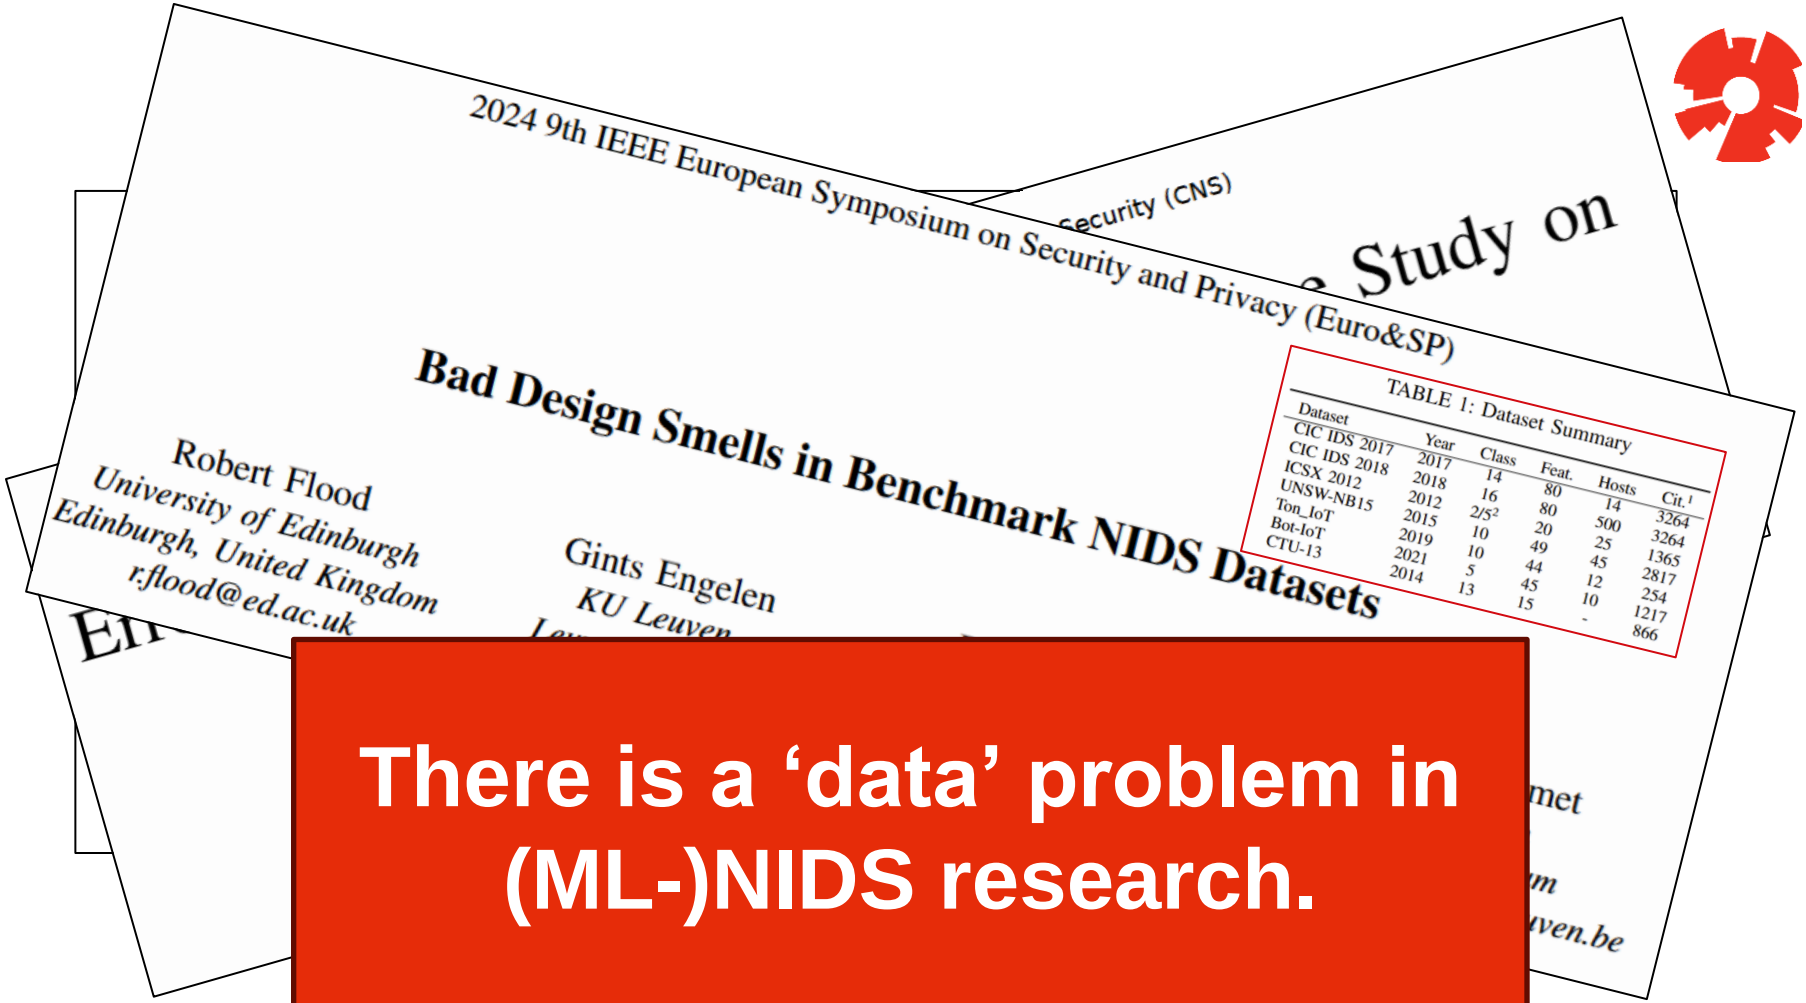

There is a ‘data’ problem in (ML-)NIDS research.

Giovanni Apruzzese
giovannia@ru.is

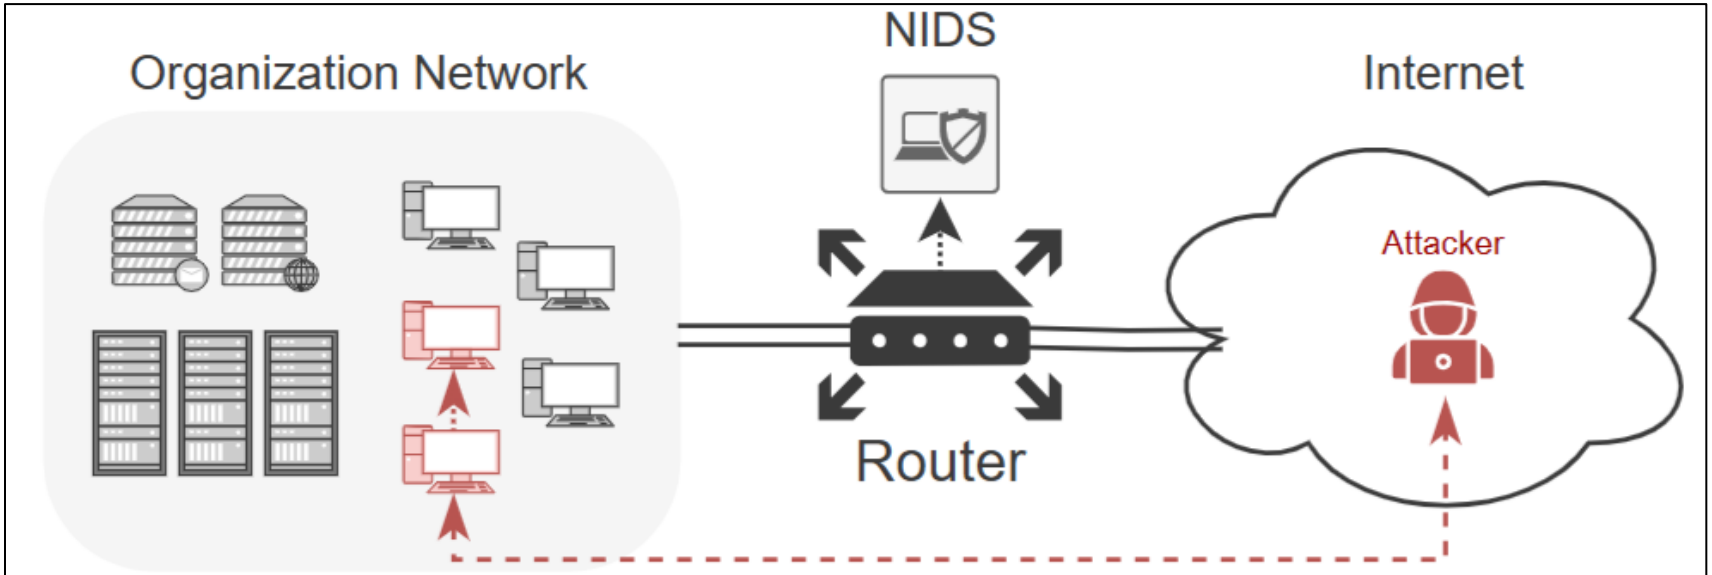

Fig. 1: **Exemplary deployment scenario of an NIDS.**

Giovanni Apruzzese
giovannia@ru.is

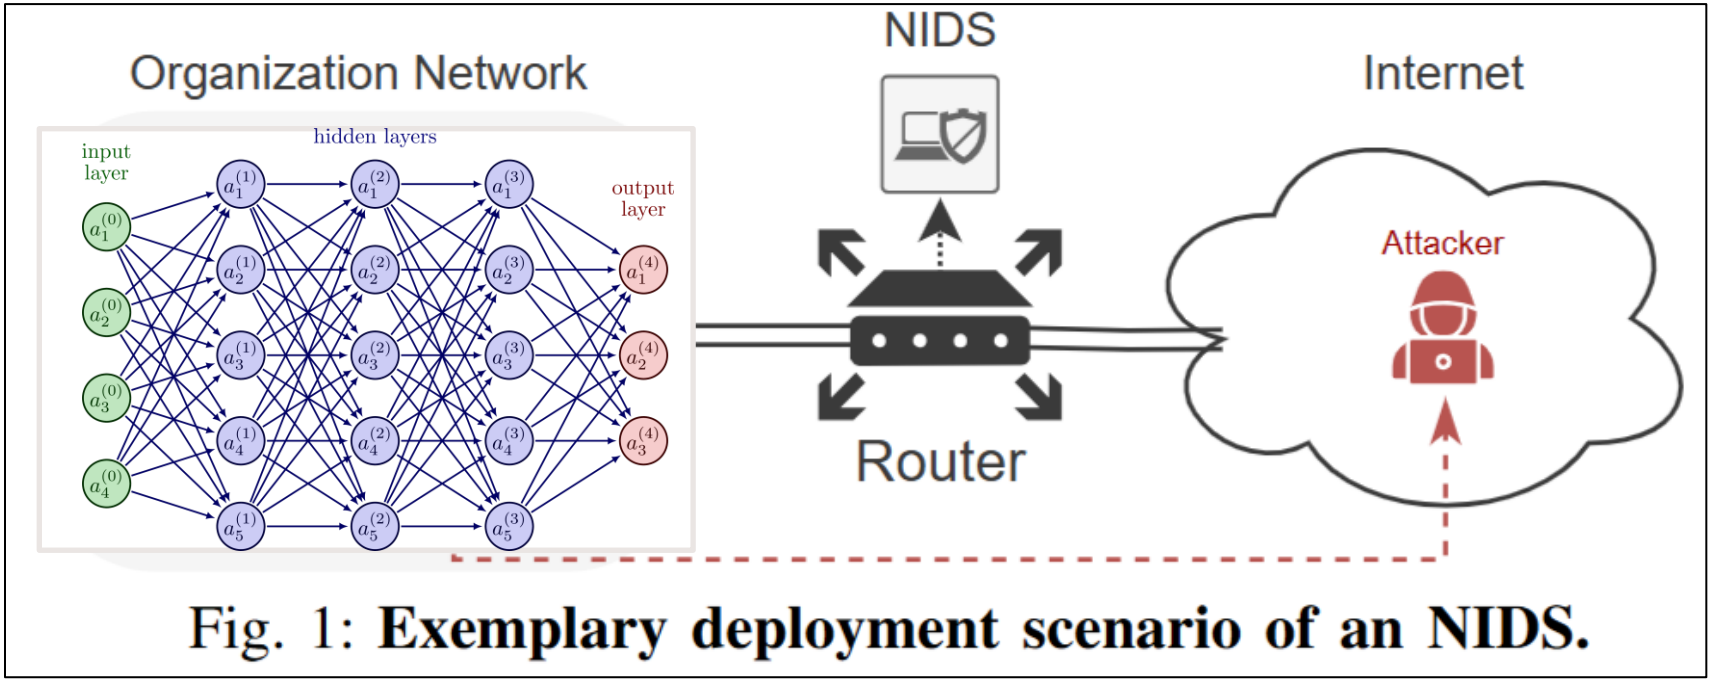

Fig. 1: **Exemplary deployment scenario of an NIDS.**

Giovanni Apruzzese
giovannia@ru.is

Giovanni Apruzzese
giovannia@ru.is

Obtaining network-related data

a. Real-world capture

b. Simulation

c. Benchmark

Giovanni Apruzzese
giovannia@ru.is

RESEARCH GOAL. We seek to develop an *open-source* tool that enables the automatic *generation* (and *labeling*) of network traffic data that resembles that of a real network.

Giovanni Apruzzese
giovannia@ru.is

ConCap (Container Capture)

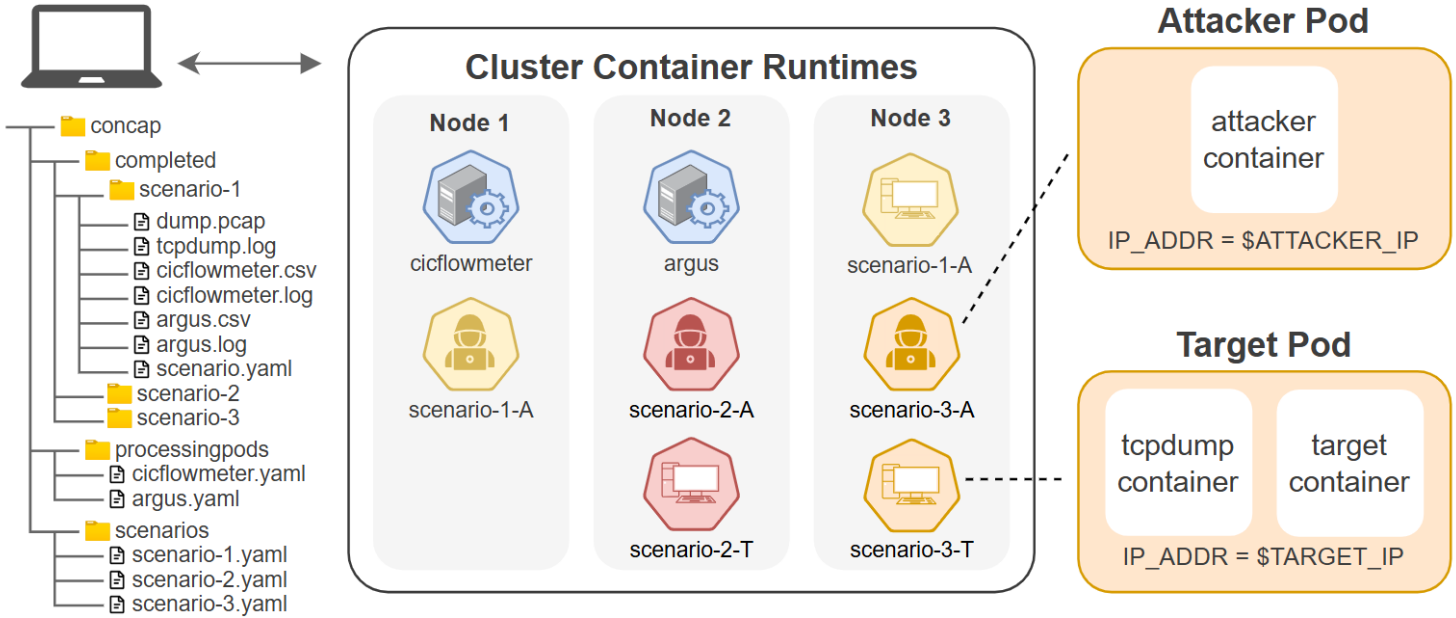

Fig. 2: **Overview of ConCap.** [left] ConCap configured with two NetFlow extractors and three scenarios, [mid] executing all scenarios simultaneously on the cluster. [right] A view of a running scenario's attacker and target pods.

Giovanni Apruzzese
giovannia@ru.is

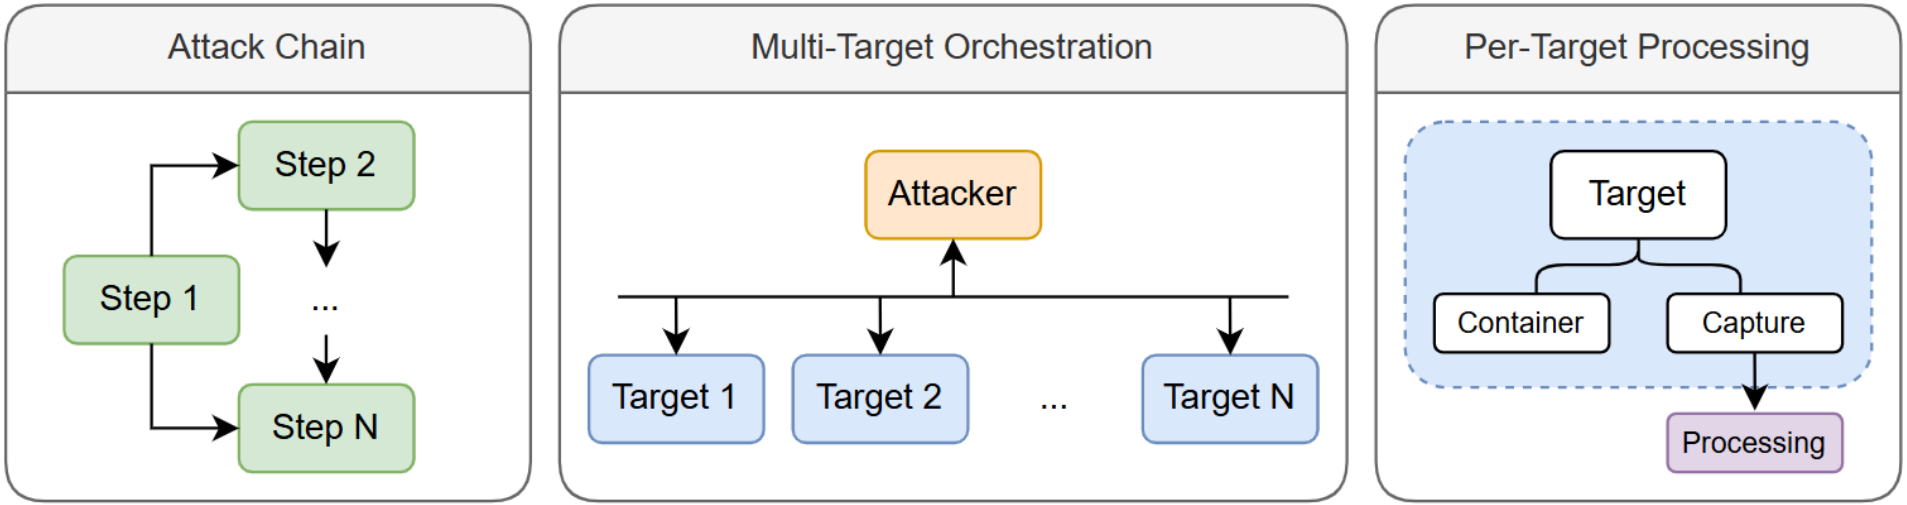

Fig. 3: **Multi-target scenario with ConCap.** [left] ConCap supports advanced multi-step attack chains, [mid] performed over multiple targets. [right] The traffic is captured and processed per target, enabling per-target labeling.

Giovanni Apruzzese
giovannia@ru.is

RESEARCH QUESTION #1: Does `ConCap` generate network traffic that resembles that of a physical network?

Giovanni Apruzzese
giovannia@ru.is

TABLE I: **Broad analysis of various network activities.** We quantitatively and qualitatively compare the number of packets and NetFlows generated by our bare-metal and ConCap environments across 10 network activities.

Tool	Number of Packets		CICFlowMeter Flows		Argus Flows	
	Bare-metal	ConCap	Bare-metal	ConCap	Bare-metal	ConCap
ping						
dig						
mysql						
curl/ftp						
nmap						
nmap (-sV)						
patator (SSH)						
patator (FTP)						
slowloris						
wfuzz						

Giovanni Apruzzese
giovannia@ru.is

TABLE I: **Broad analysis of various network activities.** We quantitatively and qualitatively compare the number of packets and NetFlows generated by our bare-metal and ConCap environments across 10 network activities.

Tool	Number of Packets		CICFlowMeter Flows		Argus Flows	
	Bare-metal	ConCap	Bare-metal	ConCap	Bare-metal	ConCap
ping	20	20	1	1	10	10
dig	10	10	5	5	5	5
mysql	37	27	1	1	1	1
curl/ftp	4910	1675	2	2	2	2
nmap	5	5	2	2	2	2
nmap (-sV)	128	103	10	10	11	11
patator (SSH)	30 960	31 485	680	680	680	680
patator (FTP)	2093	1860	70	70	70	70
slowloris	21 300	21 305	1500	1500	1500	1500
wfuzz	2310	2086	15	15	15	15

Giovanni Apruzzese
giovannia@ru.is

RESEARCH QUESTION #1: Does ConCap generate network traffic that resembles that of a physical network?

ANSWER TO RQ1. Our packet- and NetFlow-level analyses revealed that the network traffic generated by ConCap resembles that of a physical network. Deviations are due to expected differences in the network channel, which are impossible to control—but do not affect the payload content.

Giovanni Apruzzese
giovannia@ru.is

RESEARCH QUESTION #2: To what extent is the traffic generated by `ConCap` deterministic?

Giovanni Apruzzese
giovannia@ru.is

TABLE III: **RQ evaluation.** We repeat the scenarios with `ConCap` 100 times.

Environment	Attack	Packets		Number of Flows	
		Count	Sum of Bytes	CICFlowMeter	Argus
`ConCap`	Ping scan	20 ± 0	1960 ± 0	1 ± 0	10 ± 0
	Basic Port Scan	13 ± 0	694 ± 0	5 ± 0	6 ± 0
	Full Port Scan	2504 ± 6	239 401±4631	1098 ± 1	1092 ± 1
	SSH Bruteforce	26 960±354	4 759 703±23 353	680 ± 0	679 ± 0

Giovanni Apruzzese
giovannia@ru.is

TABLE III: **RQ evaluation.** We repeat the scenarios with `ConCap` 100 times.

Environment	Attack	Packets		Number of Flows	
		Count	Sum of Bytes	CICFlowMeter	Argus
Bare-Metal	Ping scan	20 ± 0	1960 ± 0	1 ± 0	10 ± 0
	Basic Port Scan	13 ± 0	736 ± 0	5 ± 0	6 ± 0
	Full Port Scan	2751 ± 88	271 551±6235	1091 ± 43	1093 ± 43
	SSH Bruteforce	30 935 ± 37	5 060 797±2417	680 ± 0	679 ± 0
`ConCap`	Ping scan	20 ± 0	1960 ± 0	1 ± 0	10 ± 0
	Basic Port Scan	13 ± 0	694 ± 0	5 ± 0	6 ± 0
	Full Port Scan	2504 ± 6	239 401±4631	1098 ± 1	1092 ± 1
	SSH Bruteforce	26 960±354	4 759 703±23 353	680 ± 0	679 ± 0

Giovanni Apruzzese
giovannia@ru.is

RESEARCH QUESTION #2: To what extent is the traffic generated by `ConCap` deterministic?

ANSWER TO RQ2. Traffic generated by `ConCap` is deterministic content-wise, but the nondeterministic nature of networking results in small variations in packets and bytes. Variations are due to how data is acknowledged (e.g., using a separate TCP packet) and should not impact the outcome of a scientific experiment. `ConCap` does not just simulate network traffic, it generates it such that it resembles physical networks—whose real-world behavior is unpredictable [21].

Giovanni Apruzzese
giovannia@ru.is

(some) Applications of ConCap

Giovanni Apruzzese
giovannia@ru.is

TABLE V: **Real-world equivalence.** Network traffic generated with ConCap is compatible with a real-world "smart home" network.

Train Set	$\mathcal{B}$+$\mathcal{P}$			$\mathcal{B}$+$\overline{\mathcal{P}}$		
Test Set	$\mathcal{B}$	$\mathcal{P}$	$\overline{\mathcal{P}}$	$\mathcal{B}$	$\mathcal{P}$	$\overline{\mathcal{P}}$
DT	<0.001	1.000	0.899	<0.001	>0.999	>0.999
RF	0.000	1.000	0.899	0.000	1.000	1.000
XGB	0.000	>0.999	0.889	<0.001	1.000	1.000
SVM	0.000	1.000	1.000	<0.001	1.000	1.000
DNN	<0.001	1.000	1.000	<0.001	1.000	1.000

Giovanni Apruzzese
giovannia@ru.is

TABLE IV: **Reproducibility of prior work.** We show ConCap-generated data is functionally equivalent to that of existing benchmark datasets.

	Train Set	$\mathcal{B}$+$\mathcal{P}$			$\mathcal{B}$+$\overline{\mathcal{P}}$		
	Test Set	$\mathcal{B}$	$\mathcal{P}$	$\overline{\mathcal{P}}$	$\mathcal{B}$	$\mathcal{P}$	$\overline{\mathcal{P}}$
CICIDS17	DT	<0.001	0.997	0.917	<0.001	0.980	1.000
	RF	0.000	0.998	1.000	0.000	0.986	1.000
	XGB	<0.001	0.998	0.869	<0.001	1.000	1.000
	SVM	<0.001	0.993	0.907	<0.001	0.993	1.000
	DNN	0.000	0.998	0.920	<0.001	0.996	1.000
CICIDS18	DT	0.000	1.000	0.859	<0.001	1.000	1.000
	RF	0.000	1.000	0.859	<0.001	1.000	1.000
	XGB	0.000	1.000	0.859	<0.001	0.984	1.000
	SVM	0.000	1.000	0.907	<0.001	1.000	1.000
	DNN	0.000	1.000	0.907	<0.001	1.000	1.000

Giovanni Apruzzese
giovannia@ru.is

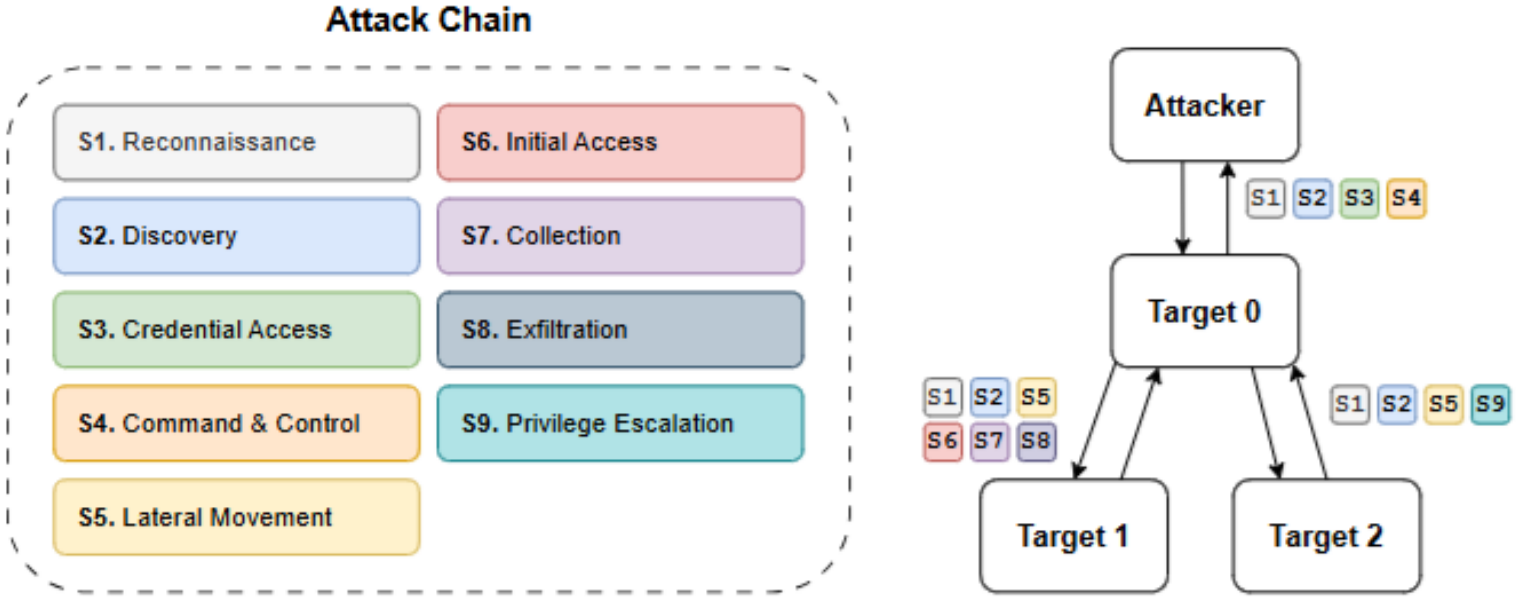

Fig. 5: **Using ConCap to reproduce complex attack chains envisioned in MITRE ATT&CK.** The attacker first compromises an exposed target before performing lateral movement to two internal hosts.

Giovanni Apruzzese
giovannia@ru.is

TABLE VI: **Testing existing ML-NIDS against unseen CVEs with ConCap.** We develop benchmark ML-NIDS using CICIDS17 (baseline tpr=0.999 with fpr=0.001) and CICIDS18 (baseline tpr=0.999 with fpr=0.001). Then, we use ConCap to generate and label NetFlows of three CVEs (we report the # Packets and # NetFlows of each capture) involving completely different attacks than those used to "train" the ML-NIDS, and we test them against all of our models (RF, DT, SVM, DNN, HGB). The results show the average tpr (and std. dev.) across all models.

Attack	Traffic Statistics		CICIDS17	CICIDS18
	# **Packets**	# **NetFlows**	tpr (std. dev)	tpr (std. dev)
CVE-2024-47177	44658	4800	0.205 (0.389)	0.550 (0.396)
CVE-2024-36401	10372	640	0.199 (0.399)	0.214 (0.378)
CVE-2024-2961	391350	1280	0.112 (0.224)	0.011 (0.021)

Giovanni Apruzzese
giovannia@ru.is

Operational requirements

- Memory footprint: <40MB
- CPU-utilization: <10%
- Initialization time: <3s

Giovanni Apruzzese
giovannia@ru.is

Conclusions

- ConCap is an open-source tool for network-traffic generation (https://github.com/idlab-discover/ConCap)
- Traffic generated by ConCap resembles that of a real-world network and can be used for ML-related experiments (due to automated and 100%-accurate NetFlow labeling)
- ConCap ‘unlocks’ the possibility to test new attacks in a safe and research-friendly environment

Giovanni Apruzzese
giovannia@ru.is

Munich, March 25th 2026

IEEE 4th Conference on Secure and Trustworthy Machine Learning

ConCap: Practical Network Traffic Generation for (ML- and) Flow-based Intrusion Detection Systems

Miel Verkerken, Laurens d'Hooge, Bruno Volckaert, Filip de Turck, Giovanni Apruzzese

HÁSKÓLINN Í REYKJAVÍK | REYKJAVIK UNIVERSITY

Common Setup and Workflow. The experiments discussed in this section entail two distinct environments.

- *Physical network.* Two "bare-metal" physical hosts, each having six Intel Core i5-9400, 32GB RAM running Ubuntu 20.04.6 and connected by a 1 Gbit switch.
- `ConCap`. A Kubernetes cluster (v1.29.0) with 1 control plane and 3 worker nodes. Each node has 16GB RAM, four Intel Xeon E5-2640v4, running Ubuntu 22.04.3. The machines are interconnected by a 10 Gbit switch.

Giovanni Apruzzese
giovannia@ru.is