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§¶
*University of Ghent (IDLab, imec), §University of Liechtenstein, ¶Reykjavik University

(miel.Verkerken, laurens.dhooge, bruno.volckaert, filip.deturck)@ugent.be, giovanni.apruzzese@uni.li

Abstract

Network Intrusion Detection Systems (NIDS) have been studied in research for almost four decades. Yet, despite thousands of papers claiming scientific advances, a non-negligible number of recent works suggest that the findings of prior literature may be questionable. At the root of such a disagreement is the well-known challenge of obtaining data representative of areal-world network—and, hence, usable for security assessments.

We tackle such a challenge in this paper. We propose ConCap, a practical tool meant to facilitate experimental research on NIDS. Through ConCap, a researcher can set up an isolated and lightweight network environment and configure it to produce network-related data, such as packets or NetFlows, that are automatically labeled—hence ready for fine-grained experiments. ConCap is rooted on open-source software and is designed to foster experimental reproducibility across the scientific community by sharing just one configuration file. Through comprehensive experiments on 10 different network activities, further expanded via in-depth analyses of 21 variants of two specific activities and of 100 repetitions of four other ones, we empirically verify that ConCap produces network data resembling that of a real-world network. We also carry out experiments on well-known benchmark datasets as well as on a real "smart-home" network, showing that, from a cyber-detection viewpoint, ConCap's automatically-labeled NetFlows are functionally equivalent to those collected in other environments. Finally, we show that ConCap enables to safely reproduce sophisticated attack chains (e.g., to test/enhance existing NIDS). Altogether, ConCap is a solution to the "data problem" that is plaguing NIDS research.

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.

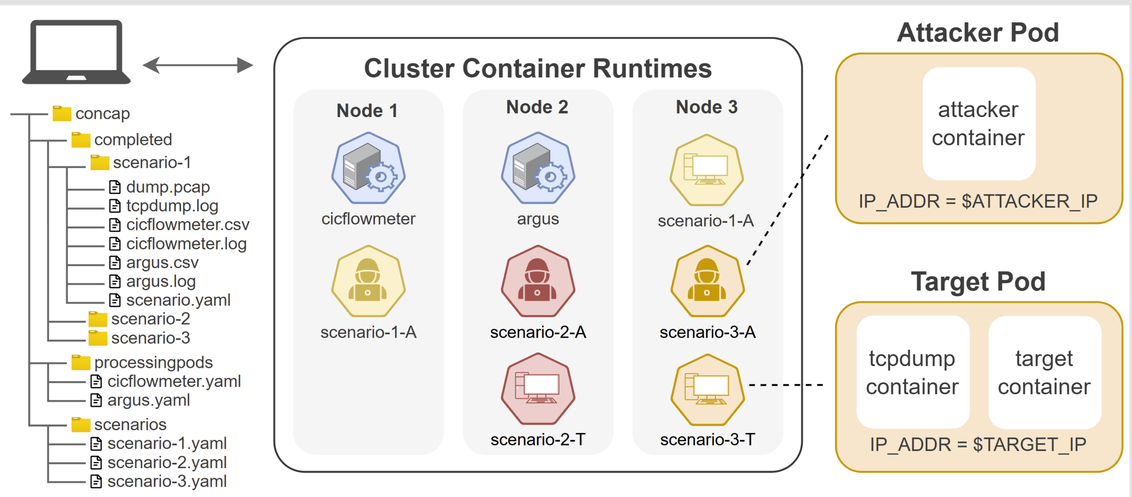

Fig. 2: **Overview of ConCap.** [left] ConCap configured with two NetFlow extractors and three scenarios, [mid] executing all scenarios simultaneously on the cluster. [right] A view of a running scenario's attacker and target pods.

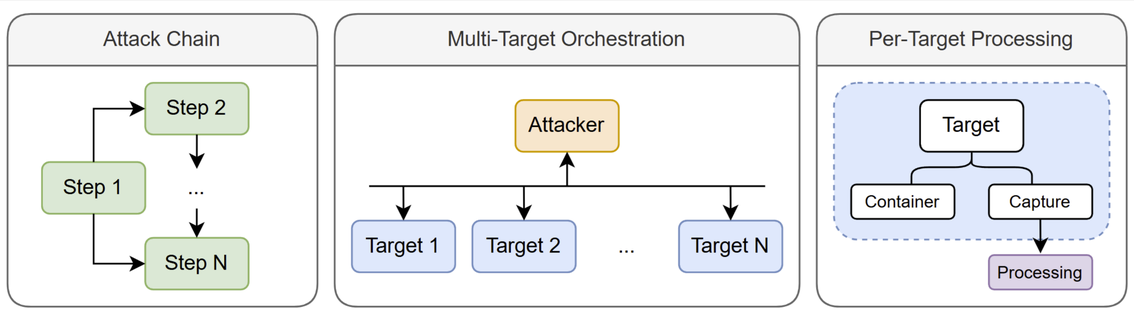

Fig. 3: **Multi-target scenario with ConCap.** [left] ConCap supports advanced multi-step attack chains, [mid] performed over multiple targets. [right] The traffic is captured and processed per target, enabling per-target labeling.

TAKEAWAY. ConCap is the first open-source tool for generating realistic, labeled network traffic for NIDS research. It offers attack flexibility, automated NetFlow generation and labeling, and parallel scenario execution with fine-grained control. A short demo is available in our repository [23].

Repository: https://github.com/idlab-discover/ConCap

RQ1: ***Does ConCap generate traffic that resembles that of a physical network?***

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

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.

RQ2: ***To what extent is the traffic generated by ConCap deterministic?***

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

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].

Empirical Demonstration and Applications of ConCap

We carry out an extensive experimental campaign, showing that:
- ConCap can be used to replicate prior research findings
- ConCap enables creation of ML-based NIDS from scratch.
- ConCap can be used for assessments of NIDS against unseen attacks.

We also measure the **operational requirements** for running ConCap---both on a commodity laptop, and on a Web cluster. In both cases, the startup time is less than 3 seconds. Containers require less than 2MiB of RAM for their execution, and the overall memory footprint of ConCap is below 40MB. The CPU utilization is always below 10%.