USENIX Security '24 Artifact Appendix: You Cannot Escape Me: Detecting Evasions of SIEM Rules in Enterprise Networks

Rafael Uetz[*] Marco Herzog[*] Louis Hackländer[*] Simon Schwarz[†] Martin Henze[‡,*]

[*]*Fraunhofer FKIE* [†]*University of Göttingen* [‡]*RWTH Aachen University*

A Artifact Appendix

A.1 Abstract

To facilitate reproduction of our experiments as well as extensions and improvements of AMIDES, our artifact comprises five components: (1) the AMIDES source code along with automation scripts, (2) the set of SIEM rules from the public Sigma repository that we analyzed for possible evasions, (3) the set of *matches* that we created, i.e., SIEM events triggering the aforementioned detection rules, (4) the set of *evasions* that we created, i.e., *matches* adapted such that the executed commands achieve the exact same goal without triggering the respective rule, and (5) a set of *benign* SIEM events that we created using the open-source testbed SOCBED (since we are not allowed to share events from the real enterprise network). Together, these components enable reproduction of our experiments and thus confirmation of our claims.

A.2 Description & Requirements

A.2.1 Security, privacy, and ethical concerns

The artifact does not pose a security risk when downloaded and executed since none of the potentially malicious command lines within the SIEM events and rules are run during the experiments. However, security software such as endpoint protection products might still raise alerts due to these command lines being contained in the event and/or rule files.

Due to ethical concerns, we will not make our full set of evasions publicly available (as stated in the paper). However, our artifact contains a small number of evasions for testing purposes, namely, those given as examples in the paper in Tables 1 and 3. Please contact us if you require the full set of evasions for your research.

A.2.2 How to access

The artifact is available on GitHub[1]. Please start by cloning or downloading the repository on a commodity computer running Linux or macOS.

[1] https://github.com/fkie-cad/amides/releases/tag/v1.0.0

A.2.3 Hardware dependencies

Assuming small training and validation datasets such as those provided with the artifact, AMIDES runs on a commodity computer with a minimum of 8 GB of RAM and requires around 2 GB of disk space. For larger training and validation datasets, more RAM and disk space are required. A fast CPU has a positive impact on the duration of training and validation.

A.2.4 Software dependencies

AMIDES is written in Python. The repository contains a list of Python package requirements that need to be installed in order to use AMIDES. All of the requirements can be installed from PyPI using pip. For convenience, the repository contains a Dockerfile and automation scripts to build containers running AMIDES and reproduce our experiments. Building and operating the containers requires a Docker installation. The containers should run on any operating system, however, our automation scripts are currently written for Linux and macOS. Please refer to the README file for further information.

A.2.5 Benchmarks

The majority of experiments in our paper (all except "Applicability to Other Rule and Event Types") are based on real benign SIEM events from a large enterprise network that are strictly prohibited to be taken off the premises. However, as stated in §6 "Datasets and Ethical Considerations", we additionally created synthetic benign events using the open-source testbed *SOCBED* (which were also used for the aforementioned experiment) to facilitate reproduction of our experiments and confirmation of our claims. These synthetic benign events, along with all other data required for reproducing our experiments (cf. §A.1) are contained in the repository except for the full set of evasions (cf. §A.2.1).

A.3 Set-up

A.3.1 Installation

After cloning or downloading the repository, please execute the steps described in the "Building the Quickstart Environ-

ment" section of the README file to build the image and containers that will run AMIDES and reproduce our experiments. Alternatively, AMIDES can be installed locally, which is described in the "Installation" section of the README file.

A.3.2 Basic Test

Running the Installation (§A.3.1) and Experiments (§A.4.2) instructions without errors ensures that the quickstart environment and AMIDES are functioning properly. If any of the automated steps fail, corresponding output messages will be generated. In case AMIDES was installed locally, the successful execution of its unit tests indicates that all components are functioning. Please refer to the "Testing" section of the README file for instructions on how to execute unit tests.

A.4 Evaluation workflow

A.4.1 Major Claims

(C1): AMIDES detects a majority of our crafted evasions without any false alerts (cf. §6.1 "Classification Performance" and Figure 3, plot "AMIDES").
(C2): AMIDES' classification performance keeps up with a (much more costly) benchmark approach (cf. §6.1 "Comparison with Benchmark Approach" and Figure 3).
(C3): AMIDES helps security analysts to attribute its evasion alerts to potentially evaded SIEM rules (cf. §6.2 and Figure 4).
(C4): AMIDES degrades gracefully in detection performance when the training set is tainted with attacks (cf. §6.3 "Influence of Tainted Training Data" and Figure 5).
(C5): AMIDES is applicable to multiple SIEM rule and event types (cf. §6.3 "Applicability to Other Rule and Event Types" and Figure 6).

A.4.2 Experiments

Please follow the instructions given in the "Running Experiments using the Quickstart Environment" section of the README file. The corresponding container will then automatically execute all four experiments (E1-E4, see below), reproduce the above-mentioned plots as PDF files, and place them in the specified folder. Note that the file names of the generated plots include the major claims they are addressing (i.e., C1-C5). In case AMIDES was installed locally, please refer to the "Running Experiments" section of the README file, which also contains more details on the experiments.

Since the benign SIEM events in the repository are not from a real enterprise network but generated by a testbed (cf. §A.2.5), the plots corresponding to Figures 3, 4, and 5 will look different compared to the paper. More precisely, the results on the testbed-generated data are significantly better because the number of benign events is much lower compared to the real enterprise events, leading to an easier classification task for AMIDES (cf. §6.1). Still, these results confirm our claims and facilitate future research.

The evaluation comprises the four subsequently mentioned experiments and requires approximately 45 human-minutes, 20 compute-minutes, and around 2 GB of disk space. As a congruence check, we provide the correct output in the document `Paper Supplement.pdf` in the repository.

(E1): *Classification Performance* [2 compute-minutes]: AMIDES' evasion detection performance is compared to a benchmark approach that was trained using matches instead of SIEM rules. Result: `figure_3_c1_c2_misuse_classification.pdf` in the `amides/plots/process_creation` folder.
(E2): *Rule Attribution* [2 compute-minutes]: AMIDES' rule attribution performance is evaluated by assessing if detected evasions are correctly assigned to the corresponding Sigma detection rules. Result: `figure_4_c3_rule_attribution.pdf` in the `amides/plots/process_creation` folder.
(E3): *Tainted Training Data* [10 compute-minutes]: AMIDES' evasion detection performance is evaluated after training data has been tainted using evasions. The tainting is repeated for different fractions of events and different events for each fraction. Result: `figure_5_c4_tainted_training.pdf` in the `amides/plots/process_creation` folder.
(E4): *Other Rule and Event Types* [2 compute-minutes]: AMIDES' evasion detection performance is evaluated for three additional rule and event types (Windows PowerShell, Windows Registry, and Web). Result: `figure_6_c5_classification_other_types.pdf` in the `amides/plots` folder.

A.5 Notes on Reusability

AMIDES is fit for application in enterprise networks. To this end, users can perform the required training with their own data (i.e., a SIEM ruleset and a set of up-to-date benign SIEM events). The resulting model can then be loaded and applied to SIEM events by the open-source log data processor *Logprep*[2], for which we implemented and published an AMIDES processor[3]. For more information on how to create models for AMIDES from scratch, refer to the "Running Custom Experiments" section of the README file.

A.6 Version

Based on the LaTeX template for Artifact Evaluation V20231005. Submission, reviewing and badging methodology followed for the evaluation of this artifact can be found at https://secartifacts.github.io/usenixsec2024/.

[2] https://github.com/fkie-cad/logprep
[3] https://logprep.readthedocs.io/en/latest/user_manual/configuration/processor.html#amides