Cryptojacking Detection: A Novel Two-Step Approach Integrating Machine Learning and DNS Monitoring

Anishka Vissamsetty[1]

Abstract— **The rapid emergence of cryptocurrencies and blockchain technology has introduced not only innovative opportunities, but also cybersecurity challenges. One such threat is cryptojacking, in which cybercriminals exploit victims' computing power for unauthorized cryptocurrency mining. This study proposes a novel two-step approach for efficient cryptojacking detection: the first step employs machine learning to identify fluctuations in CPU and memory utilization, with a focus on accurate detection by considering network traffic in the second step. DNS logs are analyzed to identify new domains during CPU spikes, cross-referenced against a curated dataset. The proposed model achieved an accuracy of 89% using gradient-boosted decision trees on real-world data. This fusion of machine learning and network monitoring provides a potent defense against the growing menace of cryptojacking, with practical applications spanning corporate environments, cloud infrastructure, and personal devices.**

I. Introduction

In recent years, cryptocurrency has become popular as a new form of digital currency. Cryptocurrency is a virtual currency that uses a computer network to process transactions. Unlike traditional physical currencies, cryptocurrency does not rely on a central authority, such as a bank or government, to maintain it. Cryptocurrency relies on a ledger called blockchain to record a list of transactions. Blockchain utilizes cryptographic techniques, a practice of securing communication and protecting data.

Cryptocurrency mining is the process by which new cryptocurrency is created and transactions, or the transfer of cryptocurrency between two parties, are initiated. Transactions are first sent to a network of computers, or nodes, where they are checked for legitimacy. Only valid transactions are then grouped into blocks, forming the basis of the blockchain. Miners compete to solve complex mathematical puzzles in order to create the next block. This process of solving the puzzle is known as Proof of Work (POW) and requires significant computational power and resources. The first miner to successfully solve the puzzle adds their block to the blockchain, earning cryptocurrency rewards and transaction fees. The idea behind providing rewards is to incentivize miners to use their computational power and resources to maintain the cryptocurrency network [1].

The rapid proliferation of cryptocurrencies and blockchain technology has introduced innovative possibilities in the digital landscape. For example, it has created faster global transactions, automation, technological advancements, and investment opportunities.

With the advancements in blockchain, cybercriminals have devised a threat to digital security: cryptojacking. Cryptojacking is an attack where criminals infect a user's device with malware to run cryptocurrency software and generate money. For criminals, cryptojacking is a lucrative crime because they are able to generate a large amount of cryptocurrency. Cryptojacking can occur for anywhere between 10 minutes to months, depending on how long the attacks go unnoticed. [2] However, cryptojacking may lead to performance degradation, device damage, financial losses, and data breach risks for victims. In 2023, the total number of cryptojacking attacks in the United States was 215 million, a 340% increase from a year ago [3]. The number of cryptojacking incidents continues to increase with new advancements in technology such as increased GPU usage and availability of mining tools. Moreover, cryptojacking attacks can often go unnoticed for large periods of time since they operate covertly.

Case Study

A recent cryptojacking attack by Gui-vil, an Indonesian threat group, demonstrates the process of cryptojacking. In 2023, Gui-vil exploited AWS EC2 instances to facilitate illegal crypto-mining operations. The group was able to gain initial access to AWS servers by using publicly compromised credentials from vulnerable GitLab instances. They gained access keys which allowed them to enter the AWS environment. Using compromised credentials is one way to gain initial access, but other methods include phishing, supply chain attacks, or infecting websites.

The group performed internal reconnaissance to explore the AWS services and maintain presence. Gui-vil ran large instances in the victims' AWS organization that they used for crypto-mining, which cost tens of thousands of dollars daily. The group then attempted to evade detection by modifying logs and disabling monitoring or alerts [4].

Gui-vil's case study shows just one instance of illegal crypto-mining and its damaging effects.

Proposal

This study aims to propose a two-step approach that can efficiently and accurately detect cryptojacking and its associated evasion techniques to avoid detection. The first step is employing a machine learning model to identify fluctuations in victims' memory and CPU utilization rates. Statistical data analysis has shown that the CPU utilization rate of mining scripts ranged from 5% to 90%, and half of the attacks had rates less than 70% [5]. Therefore, CPU monitoring alone is inaccurate and may result in false positives because CPU usage can increase due to ordinary workload spikes. Thus, the second step involves network monitoring to collect Domain Name System (DNS) logs. DNS logging is used to verify if a cryptojacking attack is occurring. By recording all DNS connections, the logs can

[1]A.V. is with Notre Dame High School, 596 S Second St, San Jose, CA 95112 (corresponding author to email: anishka18v@gmail.com).

be checked to verify if there are any malicious domains within a network.

II. Methods

Data Collection and Processing

The data used in this study was obtained from NVIDIA's Anomalous Behavior Profiling dataset, which contains 80,241 logs of benign and malware workflows [6]. This dataset contains NVIDIA GPU and CPU metrics sampled at regular time intervals. There are 31 features, including timestamps, memory usage, power consumption, temperature, and clock frequencies. The labels were encoded into numerical values, with 0 representing an attack and 1 representing no attack or a benign state. The dataset was then standardized using the Standard Scaler, which uses the Z-score transformation on every input variable. The Z-formula calculates the number of standard deviations a data point is away from the mean value, providing a measure of its position within a distribution.

Machine Learning

The dataset was first split into training and test sets, with 80% used for training and 20% used for testing. Several supervised machine-learning models were then trained on the data. These models included logistic regression, a statistical model that calculates the probability of binary classification tasks; random forest classifier, an ensemble learning method that constructs multiple decision trees during training; and gradient-boosted decision trees, a model that uses boosting on trees to iteratively learn and minimize errors and variance [7]. Gradient-boosting is an algorithm in which the predictions of weak learners are combined to improve overall performance. These models were chosen because they are well-suited for classification tasks and have a range of complexities, from linear to ensemble methods.

DNS Logging Collection

When a crypto-mining operation is orchestrated, the miner initiates an outbound connection to the mining pool using specific protocols and network communication. A mining pool is a joint group of cryptocurrency miners who share their resources to maximize rewards. Mining pools have associated server addresses or domain names that miners use to connect their software.

The second step of this approach proposes to collect DNS logs originating from a machine, as shown in Figure 1. DNS logging reveals detailed data on DNS traffic, including domain name queries, timestamps, and other records. We continuously monitor and collect DNS logs, creating a time-series model for each machine.

,"query_timestamp":"2023-08-05T23:44:54Z","query_name":"google.com.","query_type":"A"}
,"query_timestamp":"2023-08-05T23:44:54Z","query_name":"google.com.","query_type":"AAAA"}
,"query_timestamp":"2023-08-06T00:56:07Z","query_name":"pool-phx.supportxmr.com.","query_type":"AAAA","query_class"}
,"query_timestamp":"2023-08-06T00:56:07Z","query_name":"pool.supportxmr.com.","query_type":"A"}
,"query_timestamp":"2023-08-06T00:56:07Z","query_name":"pool-phx.supportxmr.com.","query_type":"A"}
,"query_timestamp":"2023-08-06T00:56:07Z","query_name":"pool.supportxmr.com.","query_type":"AAAA"}

Figure 1. DNS Logging Example

Detection Phase

The cryptojacking detection phase combines the machine learning model and DNS logging collection. Figure 2 shows a flowchart describing the entire detection process. First, if the model detects a spike in CPU usage, then the DNS logs are checked to see if new domains are generated based on the timestamps. "New domains" are defined as domains visited during the CPU spike but not before. The newly generated domains are then compared to a subset of 100,000 domains in the Alexa Top Domains list [8]. This dataset is composed of the top URLs visited in terms of unique users and page views. Security researchers commonly use the Alexa Top Domains list to obtain samples of "benign" domains. A subset of 100,000 was chosen because only the top 100,000 domains were previously found to be statistically meaningful, whereas the domains ranking below 100,000 do not contain enough data to be commonly used. This means that the domain names for mining pools are not found in the top 100,000 [9]. If the new domains are part of the subset, it is not classified as a cryptojacking attack; otherwise, if they are not part of the subset, it is classified as a cryptojacking attack.

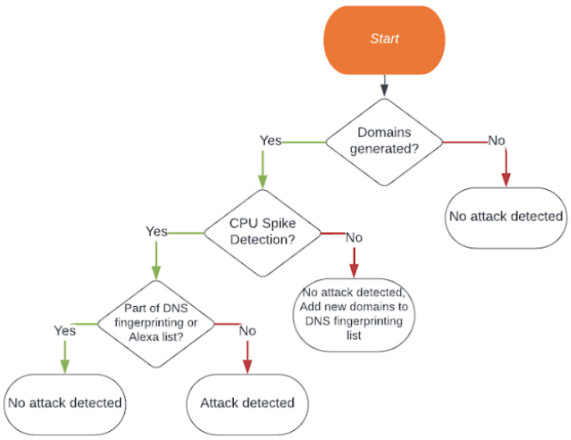

Figure 2. Flowchart describing multi-step cryptojacking detection method

III. Results

The three machine learning models were tested on the test set, consisting of a random sample of 870 rows. The gradient-boosted trees performed the best with an accuracy of 88.5%. The random forest classifier performed similarly with an accuracy of 88.0%, and the logistic regression achieved 86%. The gradient-boosted trees algorithm likely

performed the best because it is able to minimize errors and variance by combining the predictions of weak learners.

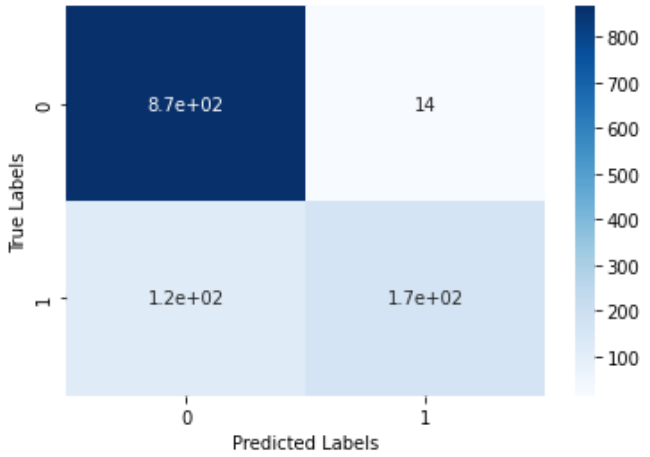

Figure 3. Confusion Matrix of Gradient-Boosted Trees Algorithm

Figure 3 shows a confusion matrix of the gradient-boosted trees, where 0 represents an attack and 1 represents no attack. We observe a high false positive rate, likely because of CPU utilization rates increasing due to workload spikes and non-malicious activity. However, with the DNS monitoring, the solution filters out all the false positives.

$$Precision = \frac{TP}{TP+FP} \quad (1)$$

$$Recall = \frac{TP}{TP+FN} \quad (2)$$

The precision for this model was 0.98 and the recall was 0.88, as calculated by formulas (1) and (2), respectively. This suggests that the model correctly identified 88% of all cryptojacking attacks, and when it predicted an attack, it was correct 98% of the time. The precision and recall indicate an effective and accurate model.

Feature Importance

Feature importance is a technique used in machine learning to assess the relevance of input features in making predictions. A score is calculated for each feature of the gradient boosting model, and this score determines the extent to which a feature contributes to the model's predictive accuracy. Different metrics were used to calculate each score, including the number of nodes, the number of root nodes, the sum of the split scores, and the average minimum depth of the first occurrence of a feature [10]. The top ten most important features are shown in Figure 4.

1	nvidia_smi_log.gpu.fb_memory_usage.free
2	nvidia_smi_log.gpu.bar1_memory_usage.free
3	nvidia_smi_log.gpu.clocks.video_clock
4	nvidia_smi_log.gpu.fb_memory_usage.used
5	nvidia_smi_log.gpu.bar1_memory_usage.used
6	nvidia_smi_log.gpu.clocks.sm_clock
7	nvidia_smi_log.gpu.pci.tx_util
8	nvidia_smi_log.gpu.clocks.graphics_clock
9	nvidia_smi_log.gpu.utilization.memory_util
10	nvidia_smi_log.gpu.temperature.memory_temp

Figure 4. Table of top 10 most important features of the gradient-boosted decision trees

These features are related to the memory usage on the GPU, the rate of transactions, the percent of time the GPU is being used, and the temperature and clock frequencies. High values for these features generally correspond to a cryptojacking attack because the computer is being used extensively. In addition, using GPU resources can generate heat, which is why the temperature should be monitored, and unusual clock frequency can indicate malicious activity [11].

When the gradient-boosted model was retrained with only the above subset of features, the resulting accuracy was 89%. Hence, the model is able to achieve similar accuracy using only this subset of features.

Application

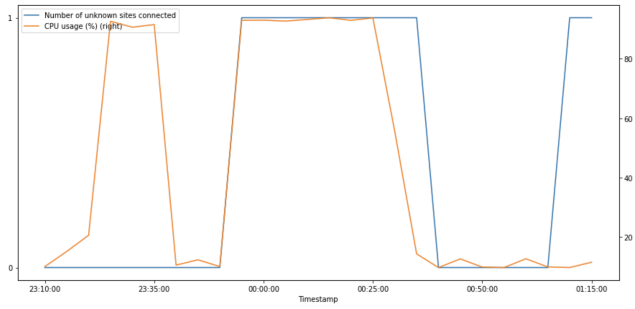

Figure 5. Graph showing a real application of CPU usage and DNS logs

Figure 5 shows an application of CPU usage and DNS logs collected on a machine running cryptojacking software. The orange line shows the CPU load in percent and the blue line shows the number of unknown sites a machine is connected to. When the CPU usage is high and the machine is connected to an unknown site (between roughly 23:45 and 00:35), it indicates a cryptojacking attack. From 23:10 to 23:40, although the CPU load is high, the machine is not connected to unknown sites, and therefore, it does not indicate an attack.

This model can be used by collecting data from a machine and deploying code available on an open-source

GitHub repository. The AWS Cloud Watch metrics can be used to monitor CPU usage and collect DNS logs.

IV. Discussion

This paper presents a novel solution to detect malicious crypto-mining effectively. This approach addresses the limitations of existing approaches by combining machine learning techniques with network monitoring, enabling a comprehensive and accurate detection process.

Current solutions are ineffective because the majority of them rely heavily on CPU monitoring. Attackers are easily able to evade and bypass detection mechanisms that rely on CPU monitoring by limiting usage, introducing idle states, or sharing resources. Evasions often lead to false negatives because low CPU activity is not reported as an attack. Additionally, CPU monitoring can lead to false positives because CPU usage can fluctuate due to workload spikes and other non-malicious activities. The proposed approach mitigates these effects by integrating machine learning with network monitoring. The machine learning model is beneficial because it enables adaptation — the model can learn from evolving strategies and a variety of factors. Some of the most important and relevant features found are memory usage, utilization, and temperature, while other techniques only monitor CPU utilization rates. In addition, the use of DNS logging can identify whether a system is connected to malicious domains, confirming whether a cryptojacking attack is occurring.

Some limitations of this study are that it has currently not been tested on large amounts of data. In addition, the data used in this study only comprised data involving the Monero cryptocurrency because of a lack of publicly available data. Since there are various types of cryptocurrencies, especially as new techniques and algorithms continue to be built, a broader range of datasets should be used to make the model more versatile and adaptable. As an extension of this research, future studies should seek to collect more data and perform more testing experiments. Companies may hope to build an API to automatically collect CPU metrics and DNS logs from a user's computer, making this solution employable.

This research holds both theoretical and practical applications in the field of cybersecurity. The use of machine learning highlights the future and possibilities of supervised machine learning algorithms – especially gradient-boosted trees. Furthermore, this tool has a wide scope for use in both organizations and personal devices. Companies that use a large amount of computing power can deploy this software in their systems, preventing unauthorized mining. In addition, individuals can utilize this software on their devices to prevent attacks, offering a defense against cryptojacking across various sectors.

V. References

[1] F. Fang et al., "Cryptocurrency trading: a comprehensive survey," Financial Innovation, vol. 8, no. 1, Feb. 2022, doi: https://doi.org/10.1186/s40854-021-00321-6. Available: https://jfin-swufe.springeropen.com/articles/10.1186/s40854-021-00321-6

[2] "Cryptojacking," *www.interpol.int*. https://www.interpol.int/en/Crimes/Cybercrime/Cryptojacking

[3] D. Thomas, "Cryptojacking Cases Increase Almost 400% In a Year," BeInCrypto, Jul. 26, 2023. https://beincrypto.com/cryptojacking-increases-fourfold-europe-us/

[4] Ahl, I. (2023, May 22). Permiso | Blog | Unmasking GUI-Vil: Financially Motivated Cloud Threat Actor. Permiso.io. https://permiso.io/blog/s/unmasking-guivil-new-cloud-threat-actor/

[5] G. Xu et al., "Delay-CJ: A novel cryptojacking covert attack method based on delayed strategy and its detection," Digital Communications and Networks, May 2022, doi: https://doi.org/10.1016/j.dcan.2022.04.030.

[6] "Morpheus/models/datasets/training-data/abp-sample-nvsmi-training-data.json at branch-23.11 · nv-morpheus/Morpheus," *GitHub*. https://github.com/nv-morpheus/Morpheus/blob/branch-23.11/models/datasets/training-data/abp-sample-nvsmi-training-data.json (accessed Aug. 15, 2023).

[7] "What is Boosting? Guide to Boosting in Machine Learning - AWS," *Amazon Web Services, Inc.* https://aws.amazon.com/what-is/boosting/

[8] "Alexa Top 1 Million Sites," *www.kaggle.com*. https://www.kaggle.com/datasets/cheedcheed/top1m

[9] W. Rweyemamu, T. Lauinger, C. Wilson, W. Robertson, and E. Kirda, "Clustering and the Weekend Effect: Recommendations for the Use of Top Domain Lists in Security Research,"

[10] "CLI User Manual — Yggdrasil Decision Forests documentation," *ydf.readthedocs.io*. https://ydf.readthedocs.io/en/latest/cli_user_manual.html (accessed Aug. 15, 2023).

[11] What is Cryptojacking and how does it work? (2023b, May 18). Usa.kaspersky.com. https://usa.kaspersky.com/resource-center/definitions/what-is-cryptojacking