

PREPRINT: Can the OpenSSF Scorecard be used to measure the security posture of npm and PyPI?

Nusrat Zahan, Parth Kanakiya, Brian Hambleton, Shohanuzzaman Shohan, Laurie Williams
North Carolina State University, USA
[nzahan, pkanaki, bthamble, sshohan, lawilli3]@ncsu.edu

ABSTRACT

The OpenSSF Scorecard project is an automated tool to monitor the security health of open-source software. We used the tool to understand the security practices and gaps in npm and PyPI ecosystems and to confirm the applicability of the Scorecard tool.

1 INTRODUCTION

Recent high-profile supply chain attacks, such as Solarwinds and Codecov, made headlines and directed attention towards the importance of software supply chain security. Therefore, practitioners are increasingly concerned with whether their projects' open-source components are secure.

Though standards, such as the NIST Secure Software Development Framework (SSDF) [1] and OWASP Software Component Verification Standard (SCVS) [2], provide exhaustive lists of security practices, a lack of consensus is observed regarding the implementation, validation, and verification of these practices towards a unified and consistent baseline measurement. Research is being conducted towards the development of different security metrics. However, establishing a baseline measurement is not straightforward since it involves exploring various sources of information, including source code repositories, vulnerability tracking systems, continuous integration/continuous deployment (CI/CD) pipelines, license(s) validity, package release history, and other metrics to develop standards for adoption. Additional challenges arise during the security assessment of packages in a software supply chain, particularly when the packages come from different sources and have different functionality. Then the practitioners have to decide whether or not packages meet the baseline standards based on evidence. Also, practitioners must continue to monitor the open-source-software (OSS) "health" to identify and manage any future risks of the software supply chain breaking down. Therefore, practitioners are now more interested than ever in identifying healthy open-source components and determining the security practices compared to other components within the ecosystem. Towards this end, *the goal of this study is to aid practitioners in producing more secure software products and make informed decisions on the security practices of candidate dependencies by depicting the current security practices and gaps across ecosystems via an empirical study of the OpenSSF Scorecard project.*

The OpenSSF Scorecard project [3] is an automated tool to monitor the security health of the OSS supply chain. The primary goal of this project is to auto-generate a "security score" for OSS projects, using a list of security metrics that can be used to assess the security health of potential dependencies. While projects like Scorecard exist to perform heuristic-based checks of a package's security practices to aid dependency selection, little research has been done to understand the viability of using Scorecard security metrics to identify

existing security gaps and practices in an entire ecosystem rather than just a single package. Observing the pattern of these security measures across one or more ecosystems can assist practitioners in determining how they are performing within the ecosystem, what they can do to improve security standards. Practitioners can also assist in the evaluation of whether a specific security check is effective within that ecosystem. In this work, we studied the Scorecard tool to analyze what security practice patterns are observed in both ecosystems.

2 OPENSSE SCORECARD

The Open Source Security Foundation (OpenSSF) is a cross-industry collaboration with a mission to improve OSS's security. Since the inception on August 3, 2020, OpenSSF has grown to include members from over 75 organizations, including GitHub, Google, IBM, Intel, Microsoft, Red Hat, the OWASP Foundation, and others.

OpenSSF launched the Scorecard project [3] in November 2020 to provide an automated security tool that gives a "security score" for OSS and reduces the manual effort required to analyze a package's security. These results are made available via a BigQuery public dataset, and OSI site.

At the time of the study, Scorecard evaluates 18 security practices ("checks") and assigns a score between 0 to 10 to each. Each check has one of four risk levels: "Critical" risk-weight 10; "High" risk-weight 7.5; "Medium" risk-weight 5; and "Low" risk-weight 2.5. An aggregate confidence score is also returned, which is a weighted average of the individual check scores weighted by risk. The table 1 discusses the 18 Scorecard checks.

3 OTHER COMMUNITY EFFORTS

Guidelines and Standards: The OWASP Software Component Verification Standard (SCVS) [2] is a framework to develop a common set of activities, controls, and security practices that can help in identifying and reducing risk in a software supply chain. There are 6 control families that contain 87 controls for different aspects of security verification or processes. The SCVS has three verification levels, where higher levels include additional controls.

In response to Section 4 of the President's Executive Order (EO) on "Improving the Nation's Cybersecurity (14028)" [5], the U.S. National Institute of Standards and Technology (NIST) updated the **Secure Software Development Framework (SSDF)** [1] to integrate the framework into each SDLC implementation. The framework comprises four groups containing high-level security practices and tasks based on established secure software development models. Each group has a number of practices, which are further split into different tasks. These four groups are-

- **Prepare the Organization (PO):** Practices-5, Tasks-13
- **Protect the Software (PS):** Practices-3, Tasks-4

Table 1: Scorecard Security Metrics and the mapping to the SSDF Framework

Metrics Name (Risk Label)	Security Checks Description	Mapping to SSDF Practices
Dangerous-Workflow (Critical)	Determines if there are dangerous patterns in the package's GitHub workflows due to misconfigured GitHub Actions. Different events with GitHub contexts can trigger workflows. A list of event context data, such as GitHub issues or pull requests, can be controlled by users and, if exploited, may lead to malicious injection.	
Vulnerabilities (High)	Tracks the presence of unfixed vulnerabilities of a package in the Open Source Vulnerabilities (OSV) [4] database.	PW.4, RV.1
Binary-Artifacts (High)	Verifies the presence of executable (binary) artifacts in the repository. Since binary artifacts cannot be reviewed, it is possible to maliciously subvert the executable.	
Token-Permissions (High)	Determines whether the package's automated workflow tokens are set to read-only. This is important because attackers might inject malicious code into the project using a compromised token with write access. If the permission's definitions in each workflow's <code>yaml</code> file are set as read-only at the top level, and the required write permissions are declared at the run-level, the project gets the highest score.	PO.5, PS.1
Code-Review (High)	Evaluates if the package conducts code reviews prior to merging a PR. Code-Reviews may be able to detect malicious code insertion attempts by an attacker. The first step of the check is to see if Branch-Protection is activated with at least one required reviewer. If this fails, the check looks to see if the last 30 commits are Prow, Gerrit, Github-approved reviews or if the merger differs from the committer.	PW.7, RV.1
Maintained (High)	Evaluates if the package is actively maintained and obtains the score based on activities on commits and issues from collaborators, members, or project owners. For example, if a project has at least one commit per week for the preceding 90 days for the latest 30 commits and issues, it will receive the highest score. Inactive projects run the risk of having unpatched code and insecure dependencies.	PW.4
Branch-Protection (High)	Monitors whether GitHub's branch protection settings have been applied to a package's branches. This check enables maintainers to set guidelines to enforce specific workflows, such as requiring reviews or passing particular status checks before acceptance into the main branch. The check is scored on a five-tiered scale. Each tier has multiple checks and must be fully satisfied to gain points at the next tier.	PS.1
Dependency-Update-Tool (High)	Determines whether the repository has enabled dependabot or renovatebot dependency update tool to automate the process of updating outdated dependencies by opening a pull request. Out-of-date dependencies are prone to attacks.	PO.3, PW.4
Signed-Releases (High)	Determines whether the project signed the release artifacts in GitHub by looking for the following filenames in the project's last five releases: <code>*.minisig</code> , <code>*.asc (pgp)</code> , <code>*.sig</code> , <code>*.sign</code> . Signed-Releases attest to the provenance of the artifact.	PS.1, PS.2, PS.3
Pinned-Dependencies (Medium)	Looks for unpinned dependencies in Dockerfiles, shell scripts, and GitHub workflows to verify the project's locked dependencies. Unpinned-Dependency allows auto-updating a dependency to a new version without reviewing the differences between the two versions, which may include an insecure component.	
Security-Policy (Medium)	Looks for a file entitled <code>SECURITY.md</code> (case-insensitive) in directories like the top-level or the <code>.github</code> of a repository to see if the package has published a security policy. Users can learn what constitutes a vulnerability and how to report it securely via a security policy.	RV.1
Packaging (Medium)	Detects language-specific GitHub Actions that upload the package to a related hub and determines if the package is published by GitHub packaging workflows. Packaging makes it easy for users to receive security patches as updates.	
Fuzzing (Medium)	Determines if the project uses fuzzing by checking the repository name in the OSS-Fuzz project list. Fuzzing is important to detect exploitable vulnerabilities.	PW.8
Static Application Security Testing (SAST) (Medium)	Determines if the project uses SAST. These tools can prevent bugs from being inadvertently introduced in the codebase. The checks look for known Github apps such as CodeQL, LGTM, and SonarCloud in the recent merged PRs, or the use of "GitHub/codeql-action" in a GitHub workflow.	PW.7, PW.8
License (Low)	Verifies if the project has published a license by looking for any combination of the following names and extensions in the top-level directory: <code>LICENSE</code> , <code>LICENCE</code> , <code>COPYING</code> , <code>COPYRIGHT</code> and <code>.html</code> , <code>.txt</code> , <code>.md</code> . Scorecard can also detect these files in the <code>LICENSES</code> directory. The lack of a license will hinder any security review and create a legal risk for potential users.	
CII-Best-Practices (Low)	Verifies whether the package has a CII Best Practices Badge, which certifies that it follows a set of security-oriented best practices such as vulnerability reporting policy, automatic process to rebuild the software, SAST, and so on.	PS.1, PS.2, RV.1, PW.5, PW.8
CI-Tests (Low)	Determines if the project runs tests before PRs are merged by looking for a set of CI-system names in GitHub CheckRuns and Statuses in recent 30 commits. CI-Tests enable developers to identify problems early in the pipeline.	RV.1
Contributors (Low)	determines if the project has contributors from multiple organizations by looking at the company field on the GitHub user profile to identify trusted code reviewers. The project must have had contributors from at least three organizations in the last 30 commits to receive the highest score.	

- **Produce Well-Secured Software (PW):** Practices-9, Tasks-16
- **Respond to Vulnerabilities (RV):** Practices-3, Tasks-9).

Automating these practices is essential to implementing security practices at scale. The Scorecard tool allows us to automate the measurement of security practice metrics at scale. We investigated whether the 18 security practices defined by Scorecard [3] complement the Executive Order (EO) and SSDF framework as part of secure SDLC practices for organizations. To that end, two authors individually mapped each metric to SSDF practices and compared the findings. We found that out of the 18 Scorecard security metrics, 13 can be mapped to the SSDF framework’s practices. Table 1 showed the mapping between each Scorecard metric and SSDF practices. Note that each SSDF practice consists of a number of tasks, hence, a practice can be linked to more than one Scorecard metrics.

Tools: Open Source Insights (OSI) [6] is a Google-developed and hosted tool to aid practitioners in grabbing information about the source code location, package metadata, licenses, releases, and vulnerabilities of open-source products. OSI scans millions of open-source packages from different ecosystems, constructs dependency graphs, and annotates the metadata in a dashboard. Apart from the package’s metadata, the OSI dashboard also shows statistics about a package’s direct or transitive dependencies. On the OSI website, users can view the vulnerability mapping of a package as well as the vulnerability mapping with associated dependencies. In addition to the package metadata, OSI has also incorporated Scorecard security metrics to help understand package security practices.

4 METHODS

This section discusses the data sourcing and generation process of this study. We compiled a package list and relevant metadata from the npm and PyPI ecosystems to collect the security score for those packages from Scorecard tool.

4.1 Ecosystem Package Metadata

Package Name: To begin, we collected a list of all package names available in both ecosystems. We sourced the list of npm packages names (1,494,105) from study [7] and the list of PyPI package names (3,65,450) was collected using PyPI API [8] in April 2022.

Dependents data: The number of dependents reflects the importance of a project by quantifying how many other projects use it. We collected dependent information from the OSI API [6]. In this work, we collected dependent information to prioritize the packages list for manual review.

4.2 OpenSSF Scorecard Score

The Scorecard tool only runs on source code hosted by GitHub. Hence, to obtain the Scorecard scores for a given package, the first step was to map the package to its respective source code location. To retrieve the source code location for both ecosystems, we use the OSI API [6]. We collected unique GitHub repositories of 767,389 npm and 191,158 PyPI packages. The package-to-repository mapping is not always a 1:1 match. Multiple packages can be found in a single repository. In total, we collected 947,936 npm packages

with 767,389 unique GitHub repositories and 211,088 PyPI packages with 191,158 unique GitHub repositories.

Then, Scorecard runs a weekly scan of open-source packages to generate the security score of those packages. However, we could not directly utilize this data for both ecosystems because, at the time of this study, Scorecard scores were only generated on 760k of 947K npm and 10K of 211K PyPI packages. Therefore, we submitted a pull request to the Scorecard repository, adding the GitHub repositories of missing packages to collect the scores from both ecosystems. The weekly Scorecard scan was able to run on those GitHub repositories after the Scorecard team successfully merged the PR.

Out of the 947,936 npm packages and 211,088 PyPI packages, we collected the generated score of 832,422 npm packages and 191,483 PyPI packages. We reviewed 50 randomly-chosen packages where the Scorecard failed to generate scores and found that we did not have access to those GitHub repositories. We collected the Scorecard score on May 09, 2022.

For each package, we could obtain 15 out of 18 Scorecard security metrics and their aggregate score, with the missing 3 metrics being the CI-Test, SAST, and CONTRIBUTOR checks. The Scorecard team took out these three checks to scale the weekly job since computing these metrics is API intensive, and GitHub rate limiting can be a bottleneck for the weekly run. As a result, we could not collect data for these three metrics.

4.3 Ecosystem Security Practices

We observe each ecosystem’s security practices and patterns by analyzing the Scorecard security checks and their assigned values. We measure the frequency of packages with scores for the 15 security practices, depicted in Figure 1 in npm and PyPI, respectively. We consider each of the practices as a security check metric.

The notation -1 in Figure 1 indicates that Scorecard could not get conclusive evidence of implementing practices, or perhaps an internal error occurred due to a runtime error in Scorecard. The inconclusive outcome is graded as -1 instead of 0. Since a value of 0 will affect the package’s aggregate score, Scorecard assigned a value of -1 to avoid the penalty of failing a check. Seven of the 15 security checks had packages with a score of -1 . The notation “1-10” denotes the percentage of packages achieving scores ranging from 1 to 10. A higher % (green cell in Figure 1) shows that most packages implement the practice. A lower % (red cell) indicates a higher % of packages failed the practice and received a score of 0. For example, in PyPI, the Fuzzing check had 0.1% packages scored between 1-10 and 18% package scored -1 , hence, 81.9% packages scored 0 and have not implemented the practice yet. The mean and standard deviation are measured to understand an ecosystem’s central tendency and spread of score distribution.

To learn why a check passed or failed, we manually reviewed 25 sample GitHub repositories from each ecosystem for each practice. We ranked each metric by the highest number of dependents and selected the top 25 packages. One author reviewed 50 GitHub repositories (25 from each ecosystem) totaling 750 repositories for the 15 checks. A second reviewer then verified the findings by selecting 100 repositories at random. We used the Cohen Kappa statistic to test the inter-rater reliability and achieved a score of 0.961. We resolved our disagreement after discussing our findings, and the

Security Check	Package Score Frequency				Score Stat			
	npm	PyPI	npm	PyPI	npm	PyPI	npm	PyPI
	-1	-1	1--10	1--10	Mean	Mean	STD	STD
✓ Dangerous-Workflow	0.0%	0.0%	99.8%	99.7%	10.0	10.0	0.5	0.5
Pinned-Dependencies	0.1%	0.3%	99.8%	99.2%	9.5	8.7	1.3	2.1
✓ Vulnerabilities	0.3%	0.2%	99.7%	99.8%	10.0	10.0	0.0	0.0
✓ Binary-Artifacts	0.0%	0.0%	99.5%	98.6%	9.9	9.7	0.8	1.4
✓ Token-Permissions	0.0%	0.0%	84.4%	71.7%	8.4	7.2	3.6	4.5
✓ License	0.0%	0.0%	68.6%	88.0%	6.9	8.8	4.6	3.3
✓ Code-Review	0.3%	0.2%	30.8%	34.8%	1.4	1.5	2.7	2.7
✓ Maintained	0.0%	0.0%	13.9%	24.1%	1.1	1.8	2.9	3.5
✓ Branch-Protection	1.7%	1.5%	10.2%	10.3%	0.5	0.5	1.7	1.6
✗ Dependency-Update-Tool	0.0%	0.0%	5.5%	2.9%	0.6	0.3	2.3	1.7
✓ Security-Policy	0.0%	0.0%	3.2%	2.6%	0.3	0.3	1.8	1.6
✗ Packaging	99.0%	94.1%	1.0%	5.9%	10.0	10.0	0.0	0.0
✗ CII-Best-Practices	0.0%	0.0%	0.2%	0.2%	0.0	0.0	0.2	0.1
✗ Signed-Releases	97.5%	93.0%	0.1%	0.5%	0.3	0.7	1.6	2.5
✗ Fuzzing	17.9%	18.0%	0.0%	0.1%	0.0	0.0	0.1	0.3
% of packages passing a security check with a score between 1-10. Towards red cell indicate that a higher % of packages failed the check with a score of 0.								
% of packages had inconclusive results, hence, achieved a score of -1, instead of 0-10								

✓ Packages can use Scorecard to verify these practices immediately

✗ Packages may not use Scorecard to verify these practices immediately due to check's inherited reliance on other systems
To display accurate stats of ecosystem practices, the Mean and STD(standard deviation) columns did not consider the [-1] value

Figure 1: npm and PyPI Ecosystems Security Practices measured by Scorecard Tool

first reviewer cross-reviewed other repositories to make changes if required. Then, we needed to examine more packages to understand a given score for Vulnerabilities and Code-Review metrics. We again chose further packages by highest dependent order.

In the following subsection, we discuss each security check for both ecosystems and the frequency statistics.

4.3.1 Dangerous-Workflow: This check detects the following two patterns in workflows: untrusted code checkout; and script injection with untrusted context variables. More than 99% packages passed the check. However, we observed 1,938 npm packages and 508 PyPI packages where Scorecard found vulnerable code patterns. Out of 50 repositories used for manual analysis, we had 8 packages with -1, all of which were the outcome of internal errors, and 11 packages with vulnerable code patterns in workflows, hence, scored 0. Among them, 3 npm packages had untrusted code checkout patterns, and 5 PyPI and 2 npm packages had warnings about script injection. At the end of this section, we provide a case study explaining how an attacker can exploit such patterns in workflows.

4.3.2 Pinned-Dependencies: In both ecosystems, more than 99% of packages had a practice of using at least one pinned dependency. Among these, 81% npm packages and 66% PyPI packages got a score of 10, indicating that they do not have any unpinned dependencies in listed directories. The score, however, may not reflect an accurate statistics. We observed that the tool does not check package.json and package-lock.json files in a npm package repository. A package.json is a JSON file that exists at the root of npm packages, containing the metadata relevant to the project to manage the project's dependencies, scripts, versions. We found packages could achieve a score of 10 even if the package had multiple unpinned dependencies in the JavaScript package's package.json file, indicating Scorecard findings do not depict the accurate status of pinned dependencies in an ecosystem. We also observed that Scorecard does not verify the presence of Dockerfiles, shell scripts, and GitHub workflows files in a repository. If a repository did not have any files of those types, a package would receive a score of 10 for not having an unpinned dependency on those missing files.

4.3.3 Vulnerabilities: More than 99% packages did not have any open vulnerabilities in the OSV database. Hence, they scored 10. Scorecard found 7 npm packages and 5 PyPI packages with unfixed vulnerabilities. In addition, 2,703 npm packages and 322 PyPI packages got a score of -1 for inconclusive results. Our manual repository review selected repositories where the package had inconclusive scores or open vulnerabilities, ranked by number of dependents. Note that, we did not review packages with scores of 10 since, these packages did not have any open vulnerabilities reported in the OSV database. The reason behind the negative score (-1) was that those repositories were empty. In total, we found 39/50 empty repositories. One package had 10 open vulnerabilities with a score of 0, and 9 packages had 1 vulnerability open with a score of 9.

4.3.4 Binary-Artifacts: More than 99% packages had a score greater than 0. The manual review of 50 repositories found eight packages with a score of 0 and noticed that these packages had more than nine binary artifacts with a mean and standard deviation of 78.25 and 87.17, respectively. These packages were umbrella projects encompassing a variety of tools and libraries. Clients are forced to use these binary artifacts directly with the inability to inspect those files for dangerous behaviors. Another 32 packages in manual review were given a score from 1 to 10 based on the number of binary artifacts ranging from 0 to 9. A score of 10 means no binaries, a score of 9 means the presence of one binary, and the scores continue to decrease toward 1 as the number of binary artifacts increased towards 9. We also found a false positive in one npm package repository, where Scorecard identified 108 binaries, two of which were .txt files.

4.3.5 Token-Permissions: In this check, npm yielded a more promising result: nearly 84% of packages have read and write permissions declared in workflows, compared to 71% of PyPI packages. Our manual review found similar patterns as we observed in Pinned-Dependencies. Fourteen (14) packages did not have any GitHub Actions specified in the repository, but Scorecard assigned 10 to those packages for Token-Permissions. Here, the score was 10 because the tool does not check the presence of workflows in the repository.

4.3.6 License: We observed that 68% of npm packages and 88% of PyPI packages had published licenses in GitHub repository, indicating, npm has a higher tendency to avoid licensing in the repository. Our manual review revealed that 4 npm packages and 8 PyPI packages had a license mentioned in the repository, specifically in `Readme.md` and `setup.py` files. However, Scorecard could not identify them. We found the following unique file names and extensions from our review: GPL-2.0, LICENSE, LICENCE, LICENSE.txt, LICENSE.rst, LICENSE.PSF, LICENSE.APACHE, LICENSE.BSD, LICENSE.md, LICENSE-MIT.

4.3.7 Code-Review: 30% of npm packages and 34% of PyPI packages had code review practices in their repository. One reason behind failing this check would be that the check is not applicable if the package has one maintainer. The trait is expected in the case of npm because most packages in npm are small and maintained by one maintainer. Study [7] reveals that, in 2021, 1.5 million npm packages had an average of 1.7 maintainers. However, our manual

review found otherwise, 9 packages scored 0 and had no code review practices even though they had more than one contributor in GitHub repositories. We also had -1 in 5 sample repositories where the repos were empty. To verify this pattern, we reviewed an additional 10 repositories with -1 . These repositories were empty on GitHub. Hence, indicating why Scorecard assigned -1 as an inconclusive result. In total, we found 2,695 npm and 321 PyPI empty repositories with -1 where the other four checks Vulnerabilities, Branch-Protection, Packaging, and Signed-Releases including Code-Review had -1 .

4.3.8 Maintained: Our findings show that more than 85% packages in npm and 75% PyPI packages were unmaintained in GitHub. What is more crucial is that for npm, unmaintained packages may have a more extended period than 90 days, as study [7] revealed that in 2021, more than 58% of packages in the npm registry were unmaintained over two years. Our manual inspections were consistent with Scorecard data where 9/50 packages were inactive in a range of 1 year to 7 years.

4.3.9 Branch-Protection: Only 10% of packages passed this check in each ecosystem, indicating these repos had at least one tier of branch protection applied. Hence, 90% npm and PyPI packages had branch protection disabled in the repository. The numbers are considerably high, indicating that a large number of packages in both ecosystems did not create a branch protection rule in repositories. Out of five tiers of scoring- "Enabling branch protection", "inhibits force to push, and branch deletion" are tier 1 check. Then, the presence of at least one reviewer(tier 2), enabling status checks(tier 3), the presence of a second reviewer(tier 4), and admin dismisses the stale review (tier 5) are the other tiers. When Scorecard is run without an administrative access token, the requirements that require admin privileges are ignored to avoid penalizing a package score.

Our manual review found that Scorecard checks the default branch and any branch that was used for creating a release and uses GraphQL API to verify the protection. However, we verified the branch-protection by looking into the GitHub branches api [9]. We found 13/50 packages had a score of -1 due to internal error because Scorecard: a) looked for the incorrect branch name that did not exist in the repository; b) could not locate the branch even though it existed; c) the main branch had a different name than the "main" or "master"; and d) branch protections were disabled in main and release branch.

4.3.10 Dependency-Update-Tool: 94% of npm packages and 97% of PyPI packages failed this check because dependabot and renovatebot were not used as dependency update tools. A project that uses other tools or manually updates dependencies, will obtain a score of 0 on this check, just like other packages with outdated dependencies. This check can only confirm if the dependency update tool is enabled; it can not confirm if the dependency-update-tool is running or if the tool's pull requests are merged.

4.3.11 Security-Policy: Only 3.2% npm and 2.5% PyPI packages have a `security.md` file. After looking into 50 sample packages, we observed that: a) 25 packages do not adhere to standard security policies; and b) 11 packages have a different reporting procedure for vulnerabilities. Users can, for example, submit bugs in other places such as GitHub issues, specific email addresses, and different

bug databases outside of GitHub, or use a different security policy reporting file security.rst.

4.3.12 Packaging: Only 1% of npm packages and 5.8% of PyPI items passed the packaging check. Since the software can be packaged in multiple ways, the challenges of coordinating several package release protocols may prohibit developers from releasing packages on GitHub Actions, which can be one reason for the limited number of packaging in the GitHub packaging workflows. At the time of this study, Scorecard could not query the package registries directly. Hence, packages that do not use GitHub actions get -1 instead of 0 . Note that a package's aggregate score will be penalized if it has a score of 0 and inconclusive or -1 , have no effect on the aggregate score. Our manual inspection identified only 2 npm packages and 6 PyPI packages used GitHub packaging workflow, while 47/50 packages had releases on GitHub. Additionally, the Scorecard failed to detect four packages (2 from each ecosystem) that had a publishing GitHub workflow. The names of these files are [publish, ci, release].yaml.

4.3.13 CII-Best-Practices: Scorecard found the CII Best Practices Badge in just 1,665 (0.2%) npm and 341(0.1%) PyPI packages. The CII Best Practices program is a way for Free/Libre and Open Source Software (FLOSS) projects to demonstrate that they follow best practices. Projects can voluntarily self-certify to report how they follow each best practice. According to the CII Best Practice Program website only 4,766 FLOSS projects have reported their security policies and received different degrees of badges, indicating why both ecosystems fell short on this check.

4.3.14 Signed-Releases: Only 578 (0.1%) npm and 936 (0.5%) PyPI packages had signed releases. Moreover, almost 100% packages failed this check. The low number of Signed-Releases in GitHub repositories are expected behavior for both ecosystems, as package developers release versions to the package registry (npmjs.org or pypi.org) rather than code hosting platforms like (github.com). Additionally, we observed that Github, PyPI, and npm each has different regulations to control package release to a registry. To publish in both registries, one must take additional steps to confirm the release, which can be incompatible with their workflow [10]. For instance, the GitHub registry accepts only scoped packages. Therefore, if a JavaScript package is currently named X, it must be renamed @username/X to publish in GitHub.

Scorecard assigns -1 instead of 0 if the tool can not detect the signed release. In addition, our manual review revealed that Scorecard often verifies older signed versions rather than checking for signatures on the new five releases. For example, one package received an 8/10 score, meaning 4/5 of recent releases of that package had signed artifacts. However, we found the signed artifacts that were from older versions, which contradicts the defined rules of Scorecard. Then we also observed repositories tagged commits as a release rather than creating a release on GitHub. However, none of the commits were GitHub verified, and Scorecard does not identify tagged releases.

4.3.15 Fuzzing: Both ecosystems fell short on this check. Scorecard validates fuzzing exclusively through the tracking of packages in OSS-Fuzz project. OSS-Fuzz has been tested only in 650 open-source packages as of July 2022 and a package that uses fuzzing with

other tools would fail the check similar to Dependency-Update-Tool check, indicating why the npm and PyPI ecosystems failed this check. Out of 650 open-source packages that use OSS-Fuzz, we found 50 npm packages and 104 PyPI packages. Despite the fact that this check was passed by only a few packages, PyPI has more fuzzing practice (50 percent more) than npm. One reason why npm packages do not use fuzzing could be that fuzzing JavaScript(JS) engines is tricky and requires expertise. Instead of processing user-supplied seeds, JS engines scan and interpret user seeds into an abstract syntax tree (AST) [11] which impacts the performance of fuzzers. Our manual analysis yielded no different results from what we expected. Only two PyPI packages used OSS-Fuzz, and 48 other packages had no fuzzer. Then, Fuzzing had most of the -1 after Signed-Releases and Packaging, but in our manual analysis re-running 12 packages with -1 , scored 0 in new run, indicating run time error occurred during the first run.

4.4 Deep Dive on Dangerous Workflow Check:

The only metric with a "Critical" risk level was Dangerous Workflow which checks for dangerous coding patterns in the code files of GitHub Actions. Potentially dangerous misuse of the GitHub workflows may lead to malicious authors (i.e. attackers) being able to perform data theft and data integrity breaches. The usage of an injection attack in the issue title and issue comment is a simple demonstration of such an attack vector, allowing the attacker to break the action environment and launch a process on the runner environment.

We explored if Scorecard can accurately detect these patterns in a GitHub repository. Instead of attacking existing repositories that failed this check, we executed the Scorecard tool on a dummy GitHub repository where we build a workflow with an intentionally-vulnerable issue action, inspired by [12]. Our vulnerable workflow (Figure 2) is executed on a GitHub runner whenever a new issue is created by anyone. For our attack, we crafted a malicious issue title in the dummy repository as malicious author to open a reverse shell connection in our local machine.

For reverse shell, we created the following issue in the dummy repository from a different GitHub account user-New malicious issue title" && bash -i >& /dev/tcp/4.tcp.ngrok.io/ngrok endpoint 0>&1 && echo". Here, we used ngrok service on local machine to expose local server ports to the Internet. Whenever the issue is created on the dummy repository, GitHub workflows will print the issue details to the log, and label it as "New Issue" using a PAT (Personal Access Token). Here, the line 14 in step 1 of Figure 2 "ISSUE TITLE: {{github.event.issue.title}}" is vulnerable to command injection because the hosted runners replaces the macros {{ ... }} blindly and echo "{{github.event.issue.title}}" becomes echo "{{New malicious issue title}}", thus, giving an attacker to run a reverse shell inside the hosted runner as part of the arbitrary code execution capabilities. An attacker can read sensitive files like .credential from the runner folder.

We tested our dummy repository with the Scorecard tool to verify whether the Scorecard identifies the vulnerable pattern in our workflow, and step 2 in Figure 2 shows that the Scorecard was able to identify the vulnerable pattern, referring to the exact line number. This case study is an example of how the dangerous

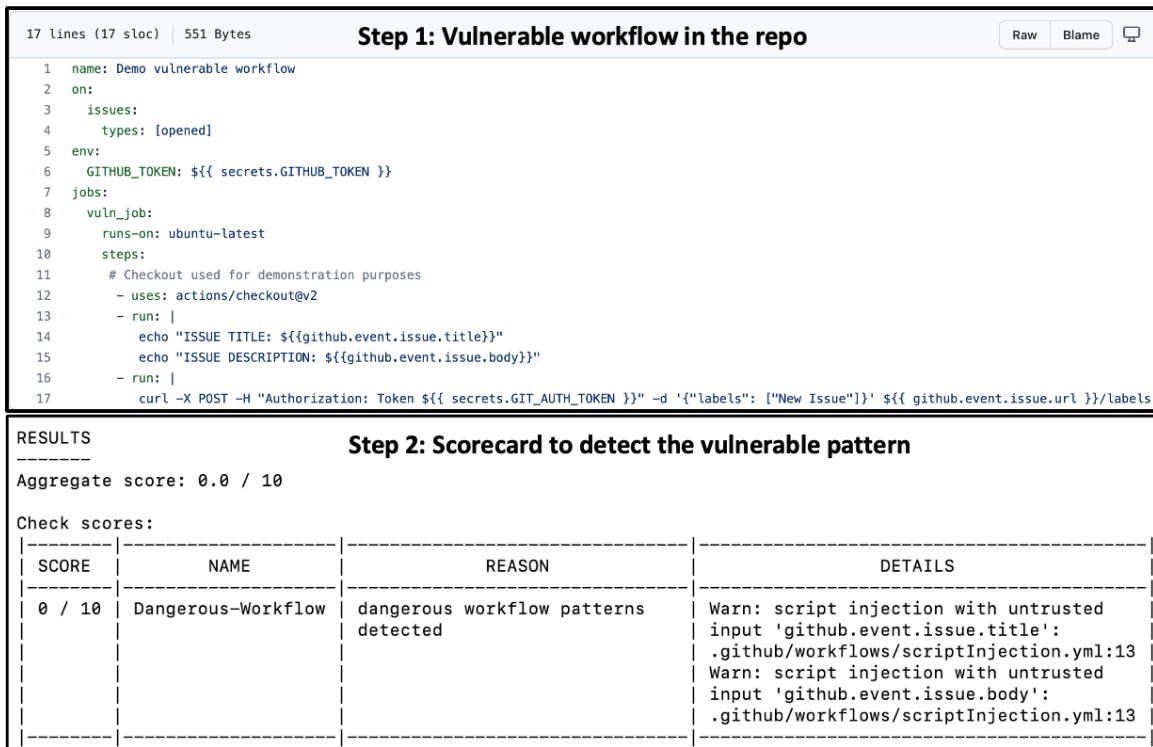


Figure 2: Case study on Dangerous Workflow to detect vulnerable code pattern

workflow check can help to identify vulnerable code patterns. We found 2,446 npm (1,938) and PyPI (508) packages that might be vulnerable to similar malicious attacks.

5 DISCUSSION:

Our work is an effort to utilize the OpenSSF Scorecard framework and tool to understand and measure cross-ecosystem package security practices and gaps. The Scorecard project seeks to provide a head start for everyone, including expert and non-expert decision-makers. The tool can be used to understand how a package performs against other OSS packages in an ecosystem and how one can identify the gap in security practices and improve the security posture. For example, if a package has a set of dependencies, then the package owner can assess all these candidate dependencies by using Scorecard tool to understand the standard security practices. Suppose one of the dependencies has open vulnerabilities with no maintained or code review status. In that case, the project owner can decide whether to accept these risks or work with the relevant stakeholders to improve the package security.

We have also observed and been told that the Scorecard team welcomes new security metrics and discussions that indicate the Scorecard is evolving with time. Scorecard adoption at the community level will help with increasing package security, identifying Scorecard’s existing issues, and determining how Scorecard may evolve and improve in response to diverse ecosystem particularities.

5.1 Security Check Metrics Evaluation

The metrics in the Scorecard are of different levels of usefulness: some are abstract, whereas others are concrete. A deliberate degree of freedom is necessary to calculate the aggregate score. For example, CII-Best-Practices open up debate about whether they are beneficial in some cases since it is a voluntary effort of maintainers. The maintainers need to self-certify a report on how they follow different security practices. Therefore, failing this check may not necessarily always advocate package do not implement these security practices. On the other hand, practitioners can use the Scorecard tool to measure the security practices of Dangerous-Workflow, Vulnerabilities, Binary-Artifacts, Token-Permissions, License, Code-Review, Maintained, Branch-Protection and Security-Policy metrics.

Then, the standards specified by the Scorecard tool for Dependency-Update-Tool, Fuzzing, Signed-Releases and Packaging security checks, were weakly adopted in npm and PyPI. At the time of the study, Scorecard assigned -1 if they could not verify the evidence of Signed-Releases and Packaging practices, hence -1 does not influence the aggregate score of a package. Scorecard intends to integrate the package registry to improve these metrics measurement. Dependency-Update-Tool, Fuzzing metrics can be measured differently due to the possibility of using different tools. However, no agreement on the list of the tools or how to verify whether a package uses any of them can be found in OSS communities. For better OSS security, the community needs to agree or standardize

these metrics so that Scorecard can integrate the standard practices. We do acknowledge, however, such agreement could be challenging. In that case, Scorecard may separate these metrics from the aggregate score calculation. If a package implements these practices, the package may get a bonus point instead of directly impacting the aggregate score. Steps like these will aid in depicting a more accurate picture of ecosystem security best practices.

The Pinned-Dependencies check requires revision based on different ecosystems. For example Pinned-Dependencies do not check the package. json and package-lock. json files for the dependency version. Then Signed-Releases, and License matrices can be improved. Signed-Releases, for example, ignore the tagged release commit and do not search the package registry directly for the releases. Then, enhancing the list of keywords can help boost the License stats. For better ecosystem evaluation, we should filter out packages with empty repositories, but Scorecard generated aggregated scores for those repos because metrics like Dangerous-Workflow, Binary-Artifacts, Pinned-Dependencies, and Token-Permissions looked for the existence of specific properties in GitHub, which empty repositories did not contain. Hence, Scorecard assigned a score of 10 instead of 0, and -1 , whereas the other 11 metrics values were between 0, -1 . These metrics open up debate about whether Scorecard should check for the existence of GitHub repositories or GitHub workflows first, before verifying the good or bad practices. For example, if a package has no GitHub workflows, the tool will score 10 in Dangerous-Workflow and Token-Permissions metrics. However, it does not actually verify that the package follows good workflow patterns. After reporting our findings, the Scorecard team acknowledged and agreed to improve Scorecard to enable automated testing more efficiently.

The Scorecard tool gives us a way to measure package security practices. However, practitioners have to adopt these practices to improve package security. Both ecosystems have a gap in practicing Code-Review, Maintained, License Branch-Protection and Security-Policy practices in the GitHub repository. In terms of License and Maintained, PyPI outperformed the npm ecosystem. Only 68% npm packages had a published license in the repository, compared to over 88% of PyPI packages. Only 13% npm and 24% PyPI packages were actively maintained. Additionally, 90% packages in both ecosystems did not have default branch protection enabled in their repositories. Practitioners should improve the security practices measured by these metrics.

6 CONCLUSION:

This research attempts to abstract from the GitHub repository level and bring security check metrics to the ecosystem level for baseline measurement. However, finding complete baseline measurements is often challenging because projects within one ecosystem use different tools and regulations and have wildly varying activity levels. The Scorecard provides a way to measure security practices to establish a baseline for OSS security measurement. Even if some metrics are not functional for all packages, knowing about them will inspire and direct practitioners on what to do if they want to adopt these practices or identify the gap that is preventing them from doing so.

7 ACKNOWLEDGMENTS

This work was supported and funded by Cisco. We thank the OpenSSF Scorecard Team for their valuable feedback and assistance in generating Scorecard data for such a vast number of repositories.

REFERENCES

- [1] M. Souppaya et al., "Secure software development framework (ssdf) version 1.1," *NIST Special Publication*, vol. 800, p. 218, 2022.
- [2] OWASP, "Software component verification standard (scvs)," <https://owasp-scvs.gitbook.io/scvs/>, 2020.
- [3] Scorecard, "Security scorecards for open source projects," <https://github.com/ossf/scorecard>, 2021.
- [4] Google, "Open source vulnerability database," <https://osv.dev/>, 2021.
- [5] D. The White House, "Executive order on improving the nation's cybersecurity," <https://www.federalregister.gov/documents/2021/05/17/2021-10460/improving-the-nations-cybersecurity>, 2021.
- [6] OSI, "Open source insight(osi)," <https://deps.dev/>, 2022.
- [7] N. Zahan et al., "What are weak links in the npm supply chain?" in *2022 IEEE/ACM 44th ICSE-SEIP*. IEEE, 2022, pp. 331–340.
- [8] PyPI, "Pypi api for package name," <https://pypi.org/simple/>, 2022.
- [9] GitHub, "The github branches api," api.github.com/repos/OWNER/REPO/branches/branch-name, 2022.
- [10] D. Wermke et al., "Committed to trust: A qualitative study on security & trust in open source software projects," in *2022 IEEE Symposium on Security and Privacy (SP)*.
- [11] S. T. Dinh et al., "Favocado: Fuzzing the binding code of javascript engines using semantically correct test cases," in *Proceedings 2021 Network and Distributed System Security Symposium, Virtual*, 2021.
- [12] N. Dotam, "Vulnerable github actions workflows," <https://www.legitsecurity.com/blog/github-privilege-escalation-vulnerability>, 2022.