

Lags in the Release, Adoption, and Propagation of npm Vulnerability Fixes

Bodin Chinthanet · Raula Gaikovina
Kula · Shane McIntosh · Takashi Ishio ·
Akinori Ihara · Kenichi Matsumoto

Received: date / Accepted: date

Abstract Security vulnerability in third-party dependencies is a growing concern not only for developers of the affected software, but for the risks it poses to an entire software ecosystem, e.g., Heartbleed vulnerability. Recent studies show that developers are slow to respond to the threat of vulnerability, sometimes taking four to eleven months to act. To ensure quick adoption and propagation of a release that contains the fix (*fixing release*), we conduct an empirical investigation to identify lags that may occur between the vulnerable release and its fixing release (*package-side fixing release*). Through a preliminary study of 231 package-side fixing release of npm projects on GitHub, we observe that a fixing release is rarely released on its own, with up to 85.72% of the bundled commits being unrelated to a fix. We then compare the package-side fixing release with changes on a client-side (*client-side fixing release*). Through an empirical study of the adoption and propagation tendencies of 1,290 package-side fixing releases that impact throughout a network of 1,553,325 releases of npm packages, we find that stale clients require additional migration effort, even if the package-side fixing release was quick (i.e., package patch landing). Furthermore, we show the influence of factors such as the branch that the package-side fixing release lands on and the severity of vulnerability on its propagation. In addition to these lags we identify and characterize, this paper lays the groundwork for future research on how to mitigate propagation lags in an ecosystem.

Bodin Chinthanet ✉, Raula Gaikovina Kula, Takashi Ishio, Kenichi Matsumoto
Nara Institute of Science and Technology, Japan
E-mail: {bodin.chinthanet.ay1,raula-k,ishio,matumoto}@is.naist.jp

Shane McIntosh
University of Waterloo, Canada
E-mail: shane.mcintosh@uwaterloo.ca

Akinori Ihara
Wakayama University, Japan
E-mail: ihara@wakayama-u.ac.jp

1 Introduction

Vulnerability in third-party dependencies is a growing concern for the software developer. In a 2018 report, over *four million vulnerabilities* were raised to the attention of developers of over 500 thousand GitHub repositories (GitHub, 2018a). The risk of vulnerabilities is not restricted to the direct users of these software artifacts, but it also extends to the broader software ecosystems to which they belong. Examples include the ShellShock (Bennett, 2014) and Heartbleed (Synopsis, 2014) vulnerabilities, which caused widespread damages to broad and diverse software ecosystems made up of direct and indirect adopters. Indeed, the case of Heartbleed emphasized its critical role in the modern web. Durumeric et al (2014) shows that OpenSSL, i.e., the project where the Heartbleed vulnerability originated, is presented on web servers that host (at least) 66% of sites and (at least) 24% of the secure sites on the internet were affected by Heartbleed.

The speed at which ecosystems react to vulnerabilities and the availability of fixes to vulnerabilities is of paramount importance. Three lines of prior works support this intuition:

1. Studies by Hejderup (2015); Howard and Leblanc (2001); Munaiah et al (2017); Nguyen et al (2016); Pashchenko et al (2018); Ponta et al (2018); Williams et al (2018) encourage developers to use security best practices, e.g., project validation, security monitoring, to prevent and detect vulnerabilities in deployed projects.
2. Studies by Cox et al (2015); Decan et al (2017); Kikas et al (2017) show that vulnerabilities can cascade transitively through the package dependency network. Moreover, they observe that security issues are more likely to occur in the field due to stale (outdated) dependencies than directly within product codes.
3. Studies by Bavota et al (2015); Bogart et al (2016); Ihara et al (2017); Kula et al (2018) show that developers are slow to update their vulnerable packages, which is occasionally due to management and process factors.

While these prior studies have made important advances, they have tended to focus on (i) a coarse granularity, i.e., releases, that focused only on the vulnerable dependency and only the direct client. For example, Decan et al (2018) analyzed how and when package releases with vulnerabilities are discovered and fixed with a single direct client, while Decan et al (2018) analyzed releases to explore the evolution of technical lag and its impact. Commit-level analysis similar to Li and Paxson (2017); Piantadosi et al (2019) is important because it reveals how much development activity (i.e., migration effort) is directed towards fixing vulnerabilities compare to the other tasks. Furthermore, there is also a research gap that relates to (ii) the analysis of the package vulnerability fixes with respect to the downstream clients, and not a single direct client.

To bridge these two research gaps, we set out to identify and characterize the release, adoption, and propagation tendencies of vulnerability fixes. We

identify and track a release that contains the fix, which is defined as a *fixing release*. We then characterize the fixing release for each npm JavaScript package in terms of commits for fixing the vulnerability, which is defined as the *package-side fixing release*. Based on semantic versioning (Preston-Werner, 2009), a package-side fixing release is a **package major landing**, **package minor landing**, or **package patch landing**. From a client perspective, we identify *client-side fixing release* lags to classify how a client migrates from a vulnerable version to the fixing version as a **client major landing**, **client minor landing**, **client patch landing**, or **dependency removal**. By comparing the package-side fixing release with the client-side fixing release, we identify lags in the adoption process as clients keep stale dependencies.

Our empirical study is comprised of two parts. First, we perform a preliminary study of 231 package-side fixing releases of npm packages on GitHub. We find that the package-side fixing release is rarely released on its own, with up to 85.72% of the bundled commits being unrelated to the fix. Second, we conduct an empirical study of 1,290 package-side fixing releases to analyze their adoption and propagation tendencies throughout a network of 1,553,325 releases of npm packages. We find that quickly releasing fixes does not ensure that clients will adopt them quickly. Indeed, we find that only 21.28% of clients reacted to this by performing a client patch landing of their own. Furthermore, we find that factors such as the branch upon which a fix lands and the severity of the vulnerability have a small effect on its propagation trajectory throughout the ecosystem, i.e., the latest lineage and medium severity suffer the most lags. To mitigate propagation lags in an ecosystem, we recommend developers and researchers to (i) develop strategies for making the most efficient update via the release cycle, (ii) develop better awareness mechanisms for quicker planning of an update, and (iii) allocate additional time before updating dependencies.

Our contributions are three-fold. The first contribution is a set of definitions and measures to characterize the vulnerability discovery and fixing process from both the vulnerable package, i.e., package-side fixing release, and its client, i.e., client-side fixing release. The second contribution is an empirical study that identified potential lags in the release, adoption, and propagation of a package-side fixing release. The third contribution is a detailed replication package, which is available at <https://github.com/NAIST-SE/Vulnerability-Fix-Lags-Release-Adoption-Propagation>.

1.1 Paper Organization

The remainder of the paper is organized as follows: Section 2 describes key concepts and definitions. Section 3 presents the motivations, approaches, and results of our preliminary study. Section 4 introduces the concepts of fix adoption and propagation lags modeling and tracking with the empirical study to identify them. We then discuss implications of our results and threats to validity of the study in Section 5. Section 6 surveys the related works. Finally, Section 7 concludes the study.

2 Concepts and Definitions




Fig. 1 The relationship between package-side and client-side regarding vulnerability discovery, fixing, and release process of package \mathbb{P} and client \mathbb{X} over time. Red and green releases indicate whether releases are vulnerable or not.

2.1 Package-side Vulnerability Fixing Process

Figure 1 illustrates the timeline of the package-side vulnerability fixing process of package \mathbb{P} in the lower part of the figure. We break down this process into two steps:

Step one: Vulnerability Discovery. Figure 1 shows the vulnerability of package \mathbb{P} being detected after the release of $\mathbb{P}_{V1.1.0}$. As reported by Kula et al. (2018), CVE defines four phases of a vulnerability: (i) threat detection, (ii) CVE assessment, (iii) security advisory, and (iv) patch release. We define the vulnerability discovery as the period between the threat detection and before the patch release. It is most likely that the fixing process starts after the CVE assessment, i.e., a vulnerability has been assigned a CVE number. In addition to a CVE number, a vulnerability report details the affected packages, which can include releases up to an upper-bound of reported versions. In this step, the developers of an affected package are notified via communication channels such as a GitHub issue for GitHub projects.

Step two: Vulnerability Fix and Release. Figure 1 shows the vulnerability of package \mathbb{P} being fixed and released as $\mathbb{P}_{V1.1.1}$. We define the vulnerability fix and release as the period where developers spend their efforts to identify and mitigate the vulnerable code. Once the fix is ready, developers merge that fix to a package repository. In most GitHub projects, developers will review changes via a GitHub pull request. Semantic versioning convention

is also used to manage the release version number of a package (Preston-Werner, 2009). We define the *fixing release* as the first release that contains a vulnerability fix ($\mathbb{P}_{V1.1.1}$). We also define *package-side fixing release* to describe the fixing release of the vulnerable package and to show how a package bumped the release version number from a vulnerable release to a fixing release. There are three kinds of package-side fixing release based on the semantic versioning: (i) **package major landing** (major number is bumped up), (ii) **package minor landing** (minor number is bumped up), and (iii) **package patch landing** (patch number is bumped up). As shown in Figure 1, package-side fixing release of package \mathbb{P} is package patch landing, i.e., $\mathbb{P}_{V1.1.0}$ to $\mathbb{P}_{V1.1.1}$.

2.2 Client-side Fixing Release

Prior work suggests that lags in adoption could be the result of migration effort (Kula et al., 2018). Thus to quantify this effort, we compare a vulnerable release that a client is using, i.e., listed in the report, against the package-side fixing release to categorize as a *client-side fixing release*. Developers of the vulnerable package fulfilled their responsibility, thus the adoption responsibility is left to the client.

Figure 1 shows the timelines of a fixing release and its clients. As illustrated in the figure, client \mathbb{X} suffers a lag in the adoption of a package-side fixing release by switching dependency branches, i.e., client minor landing. Also, client \mathbb{X} directly depends on package \mathbb{P} . Client \mathbb{X} is vulnerable ($V2.0.0$) due to its dependency ($\mathbb{P}_{V1.0.1}$). To mitigate the vulnerability, package \mathbb{P} creates a new branch, i.e., minor branch, which includes the package-side fixing release ($\mathbb{P}_{V1.1.1}$). Client \mathbb{X} finally adopts the new release of package \mathbb{P} which is not vulnerable ($V1.1.2$). We consider that client \mathbb{X} has a lag, which was not efficient because it actually skipped the package-side fixing release ($\mathbb{P}_{V1.1.1}$). Instead, client \mathbb{X} adopted the next release ($\mathbb{P}_{V1.1.2}$). A possible cause of lags is the potential migration effort needed to switch branches, i.e., from $\mathbb{P}_{V1.0.1}$ to $\mathbb{P}_{V1.1.2}$. The migration effort for major or minor changes may include breaking changes or issues in the release cycle.

2.3 Motivating Example

Figure 2 shows a practical case of the vulnerability fixing process which affects a library for network communication, i.e., `socket.io`. Figure 2(a) and Figure 2(b) show step one, where the vulnerability is reported as a GitHub issue.¹ and summarized in snyk.io.² The vulnerability report contains detailed information regarding the identified problem, its severity, and a proof-of-concept to confirm the threat. In this example, `socket.io` was vulnerable to a medium severity

¹ <https://github.com/socketio/socket.io/issues/856>

² <https://snyk.io/vuln/npm:socket.io:20120323>

Vulnerability DB > npm > socket.io

Insecure Randomness

Affecting `socket.io` package, versions <0.9.7

Overview

`socket.io` is a node.js realtime framework server. Affected versions of the package are vulnerable to Insecure Randomness due to the cryptographically insecure `Math.random` function which can produce predictable values and should not be used in security-sensitive context.

Remediation

Upgrade `socket.io` to version 0.9.7 or higher.

References

- [GitHub Changelog](#)
- [GitHub PR](#)
- [GitHub Issue](#)
- [GitHub Commit](#)
- [Google Group Forum](#)

CVSS SCORE

5.3

Severity

MEDIUM SEVERITY

CREDIT

Martin Thomson

CVE

[CVE-2017-16031](#)

CWE

[CWE-330](#)

SNYK ID

`npm:socket.io:20120323`

DISCLOSED

22 Mar, 2012

PUBLISHED

13 Feb, 2017

Twitter Facebook LinkedIn Google+ Email

(a) `socket.io` vulnerability report with medium severity.

ID generation vulnerability #856

martinthomson opened this issue on Apr 27, 2012 · 0 comments

The generation of new IDs for new sockets is vulnerable to prediction attacks. Potentially, though highly unlikely, collisions are also possible.

Mentioned in the comments of #487.

The `Manager.generateId()` method generates identifiers using the default `Math.random()` pseudorandom number generator, which produces predictable output. Knowledge of the `socket.io` identifier is sufficient information for a client to receive information from a `socket.io` server. With the current implementation, an attacker that is able to predict subsequent identifiers can make a request and gain potentially private information.

(b) GitHub Issue reporting the vulnerability.

Fix for ID generation vulnerability #856 #857

Merged `rauchg` merged 4 commits into `socketio/socket.io` from `martinthomson:fix856` on Apr 27, 2012

Conversation · Commits · Checks · Files changed

martinthomson commented on Apr 27, 2012

Using crypto.randomBytes to generate a session ID. 96 bits of randomness, 24 bits of sequence number. Using the URL safe base64 encoding RFC 4648 to represent the value.

Backward compatibility for node 0.4 uses an insecure generation method with `Math.random()`. I could fix this too, but I don't believe that this version of node is still popular enough to justify the effort.

(c) Pull request containing the fixing commit.

Making ID generation securely random

master (#857) · 2.2.0 · 0.9.7

martinthomson committed on Apr 27, 2012

1 parent fe6dd87 commit 67b4eb9abdf111dfa9be4176d1709374a2b4ded8

(d) Fixing commit to mitigate the vulnerability.

0.9.7

bb900d4

0.9.7

rauchg tagged this on Jul 25, 2012 · 872 commits to master since this tag

Release 0.9.7

(e) New package-side fixing release has been released.

Fig. 2 Developer artifacts that mitigate a vulnerability (`socket.io`) on GitHub.

vulnerability. We also found that the reporter is the same person who also created the fix. Figure 2(c) and 2(d) show step two. Figure 2(c) shows a fix that will be merged into the code base. The fix is submitted in the form of a pull request.³ Figure 2(d) shows that there are four commits in a pull request

³ <https://github.com/socketio/socket.io/pull/857>

with one commit actually fixes the vulnerability⁴. The other three commits were found to be unrelated, e.g., "removing fixes for other bug". Figure 2(e) shows that the package-side fixing release was available on July 25, 2012.⁵

Interestingly, there is a lag in the vulnerability fix and release step. This example shows that there is an 89 days period between when the fix was created and released for any client to use. We found that `socket.io` merged its fixes on April 27, 2012, however, it was actually released on July 25, 2012 and classified as a package patch landing, i.e., V0.9.7.

3 Package-side Fix Commits and Landing: Preliminaries

From the motivating example in Section 2.3, we reveal how much development activity is directed towards fixing vulnerabilities compare to the other tasks. Thus, we conduct a preliminary study to characterize package-side fixing release at the commit-level which including (i) changes of package-side fixing release, (ii) contents of package-side fixing release. We first highlight the motivation, approach, and analysis to answer our preliminary questions. We then show our data collection and finally provide the results. The following two preliminary questions guide the study:

(PQ₁) What is the prevalence of package patch landing?

- *Motivation* Our motivation for PQ_1 is to analyze the package-side fixing release. Different from Decan et al (2018), we manually investigate the version changes in the package-side fixing release itself. Our assumption is that *every fix is applied as a package patch landing*.
- *Approach* The approach to answer PQ_1 involves a manual investigation to identify the package-side fixing release, i.e., package major landing, package minor landing, package patch landing. This is done in three steps.
 1. The first step is to extract the fix-related information on GitHub. The extracted information is captured into three types as (i) an issue, (ii) a commit, and (iii) a pull request.
 2. The second step is to identify the release that contains a fix. This step involves an investigation of the package history. From the link in the first step, the first author manually tracked the git commit history to identify when the fix was applied.
 3. The final step is to identify a difference between a vulnerable release and a package-side fixing release. We compare a vulnerable release, i.e., listed in the report, against a package-side fixing release to categorize the changes: (i) package major landing, (ii) package minor landing, and (iii) package patch landing.
- *Analysis* The analysis to answer PQ_1 is the investigation of package-side fixing releases. We use a summary statistic to show the package-side fixing

⁴ <https://github.com/socketio/socket.io/commit/67b4eb9abdf111dfa9be4176d1709374a2b4ded8>

⁵ <https://github.com/socketio/socket.io/releases/tag/0.9.7>

Table 1 A summary of package-side dataset information for preliminary study.

npm Vulnerability Report Information	
Disclosures period	Apr 9, 2009 – Aug 7, 2020
# vulnerability reports	2,373
# samples of vulnerability reports (with fix references)	231
# vulnerable packages	172

release distribution. Furthermore, an interesting example case from the result is used to confirm and explain our findings.

(PQ₂) What portion of the release content is a vulnerability fix?

- *Motivation* Extending PQ_1 , our motivation for PQ_2 is to reveal what kinds of contents are bundled within the package-side fixing release. We complement recent studies, but at the commit-level (Decan et al. 2018; Decan et al. 2018; Hejderup, 2015). We would like to evaluate the assumption that *commits bundled in a package-side fixing release are mostly related to the fixing commits*.
- *Approach* The approach to answer PQ_2 involves a manual investigation of contents inside the package-side fixing release. This is done in three steps.
 1. The first step is to gather information of a fixing commit. In this case, we tracked the git commit history similar to the second step of PQ_1 .
 2. The second step is to list commits in a package-side fixing release. We use GitHub comparing changes tool to perform this task⁶.
 3. The final step is to identify the type of commits bundled in a package-side fixing release. Similar to PQ_1 , the first author manually tracked and labeled the commit as either (i) fixing commit or (ii) other commit. To label commits, the first author uses source codes, commit messages and GitHub pull request information. For validation, other co-authors confirmed the results, i.e., one author found the evidence and the other validated.
- *Analysis* The analysis to answer PQ_2 is to examine the portion of fixing commits in a package-side fixing release. We show the cumulative frequency distribution to describe the distribution of fixing commits for 231 package-side fixing releases. We use a box plot to show fixing commit size in terms of lines of code (LoC). Similar to PQ_1 , we show interesting example cases from the result.

3.1 Data Collection

Our dataset contains the vulnerability reports with fix-related information on GitHub. For the vulnerability reports, we crawled the data from snyk.io (Snyk,

⁶ <https://help.github.com/en/github/committing-changes-to-your-project/comparing-commits>

Table 2 A summary statistic of package-side fixing release distribution in PQ_1 .

Package-side fixing release	# of fixing releases
Major package landing	17 (7.36%)
Minor package landing	47 (28.14%)
Patch package landing	149 (64.50%)
	231

[2015] that were originally disclosed in CVE and CWE database. For the fix-related information, we focus on packages from the npm JavaScript ecosystem that is one of the largest package collections on GitHub ([NPM, 2010]) and also has been the focus of recent studies ([Abdalkareem et al., 2017; Decan et al., 2017; Decan et al., 2018; Decan et al., 2018; Hejderup, 2015; Kikas et al., 2017]). At the time of this study, we extracted fix-related information links directly from snyk.io (e.g., GitHub issue, commit, pull request).

As shown in Table 1, we crawled and collected all reports from April 9, 2009 to August 7, 2020, i.e., in total 2,373 reports. To identify the reports with fix-related information, we removed reports that (i) do not have the fixing release or (ii) do not provide any fix-related information. After that, we randomly select around 237 reports (10% of 2,373) and manually filter reports that the vulnerable package does not follow semantic versioning. In the end, the dataset for PQ_1 and PQ_2 analysis consists of 231 reports that affect 172 packages.

3.2 Results to the Preliminary Study

(PQ_1) *What is the prevalence of package patch landing?*

Table 2 shows the evidence that not every fix is applied as a patch. This evidence contradicts our assumption. We find that 64.50% of fixes are a package patch landing. On the other hand, we find that 7.36% and 28.14% of fixes are package major landing and package minor landing respectively. From our result, we suspect that some releases, especially package major landing and package minor landing might contain unrelated contents to the fixing commits.

The example case is a package major landing of an HTTP server framework, i.e., `connect` (V2.0.0). This package was vulnerable to Denial of Service (DoS) attack ([Snyk, 2017a]). Under closer investigation, we manually validated that the other fixes were bundled in a package-side fixing release, including API breaking changes (i.e., removed function).⁷ This fix also takes 53 days before it gets released. We suspect that this may cause a lag in the package-side fixing release, especially if the project has a release cycle.

⁷ <https://github.com/senchalabs/connect/blob/fe28026347c2653a9602240236fc43a8f0ff8e87/History.md#200--2011-10-05>




Fig. 3 We find that 91.77% out of 231 fixing releases have fixing commits up to 14.28% of commits in a package-side fixing release.




Fig. 4 LoC of the fixing commits for 231 vulnerabilities. We find that there are only few fixing fix commits in the package-side fixing release, i.e., median of 10.

Summary: We find that the vulnerability fix is not always released as its own patch update. We find that only 64.50% of fixing releases are a package patch landing. The rest of fixing releases are either package major landing, i.e., 7.36%, or package minor landing, i.e., 28.14%.

(PQ₂) What portion of the release content is a vulnerability fix?

Figure 3 is evidence that fixes are usually bundled with other kinds of changes. We find that 91.77% out of 231 fixing releases have up to 14.28% commits that related to the fix, which means that 85.72% of commits were unrelated. Figure 4 shows that the fix itself tends to contain only a few lines of code, i.e., median of 10 LoC. Similar to the commit-level analysis, the package-side fixing release tends to contain a lot of changes, i.e., median of 219 LoC. These

results complement the finding of PQ_1 about package-side fixing release might contain unrelated changes to the fixing commit.

We show two examples to investigate the content of a fix and its size. The first example is a package patch landing of a simple publish-subscribe messaging for a web, i.e., `faye` (Snyk, 2017b). Under closer manual inspection, we find that there is one commit that updates the default value of variables.⁸ However, a package patch landing includes a total of 45 commits that is not related to the fix.⁹ In the second example, we show that the actual fix is only a few lines of code. The `npm` package, which is the command line interface of a JavaScript package manager (Snyk, 2017c), took seven lines of code to fix the vulnerability.¹⁰

Summary: We find that a small portion of the release contents related to the vulnerability fix, with 91.77% of 231 fixing releases have up to 85.72% unrelated commits. Furthermore, we found that the fix itself tends to have only a few lines of code (i.e., median of 10 lines of code).

4 Client-side Lags after the Package-side Fixes

The results of our preliminary study characterize the package-side fixing release, where we find that (i) up to 64.50% of vulnerability fixes are classified as a package patch landing and (ii) up to 85.72% of commits in a release are unrelated to the actual fix. Based on these results, we suspect that potential lags might occur while the package-side fixing release get adopted by the clients and transitively propagate throughout the dependency network. Hence, we perform an empirical evaluation to explore potential lags in the adoption and propagation of the fix.

4.1 Model and Track Lags


To explore potential lags in both adoption and propagation, we model and track the package-side fixing release and client-side fixing release as illustrated in Figure 5.

Released and Adopted by Version - We identify lags in the adoption by analyzing the prevalence of patterns between a package-side fixing release and client-side fixing release, which is similar to technical lag (Zerouali et al (2018)) and based on semantic versioning. The definition of *package-side fixing release* was explained in Section 2.1 which describes how the package bumped the release version number. Note that pre-releases or special releases are not considered in


⁸ <https://github.com/faye/faye/commit/e407e08c68dd885896552b59ce65503be85030ad>

⁹ <https://github.com/faye/faye/compare/0.8.8...0.8.9>

¹⁰ <https://github.com/npm/npm/commit/f4d31693>



(a) An example of a lag in the propagation, caused by cascade delays from upstream clients.



(b) An example of a Latest Lineage (LL) and a Supported Lineage (SL) classify to track the freshness of a package-side fixing release.

Fig. 5 These figures show the terms that are used to model and track the lags.

this study. We then define a new term called a *client-side fixing release*. Client-side fixing release describes how clients bumped the version of an adopted package up from vulnerable version to fixing release. There are four kinds of client-side fixing release: (i) **client major landing** (major number of an adopted package is bumped up), (ii) **client minor landing** (minor number of an adopted package is bumped up), (iii) **client patch landing** (patch number of an adopted package is bumped up), and (iv) **dependency removal** (adopted package is removed from a client dependency list).

Figure 5(a) shows an example of the two terms defined above. First, we find that the package-side fixing release for package \mathbb{P} is classified as a package patch landing. This is because of the difference between a fixing release ($\mathbb{P}_{V1.1.1}$) and its vulnerable release ($\mathbb{P}_{V1.1.0}$). Furthermore, we find that the client-side fixing release for client \mathbb{X} is a client minor landing. This is because of the difference between the adopted fixing release ($\mathbb{P}_{V1.1.2}$) and its previous vulnerable release ($\mathbb{P}_{V1.0.1}$).

Propagation Influencing Factors - We define *Hop* as the transitive dependency distance between a package-side fixing release and any downstream clients that have adopted this fix, i.e., one, two, three, and more than or equal to four hops. As shown in Figure 5(a), client \mathbb{X} is one hop away from package \mathbb{P} . We consider two different factors to model and track lags in the propagation:

1. *Lineage Freshness*: refers to the freshness of the package-side fixing release as inspired by Cox et al (2015) and Kula et al (2018). Figure 5(b) shows two types of lineage freshness based on the release branches including: *Latest Lineage (LL)*: the client has adopted any package-side fixing release on the latest branch, and *Supported Lineage (SL)*: the client has adopted any package-side fixing release not on the latest branch. Our assumption is that a package-side fixing release in the latest lineage is adopted faster than a package-side fixing release in a supported lineage, i.e., suffer less lags. Figure 5(b) shows that three versions of package \mathbb{P} (V1.0.2, V1.0.3, V1.1.3) are classified as SL.
2. *Vulnerability Severity*: refers to the severity of vulnerability, i.e., H = high, M = medium, L = low, as indicated in the vulnerability report (as shown in Figure 2(a) from Section 2). Our assumption is that a package-side fixing release with higher severity is adopted quicker, i.e., less lags.

4.2 Empirical Evaluation

The goal of our empirical study is to investigate lags in the adoption and propagation. We use these two research questions to guide our study:

(RQ₁) Is the package-side fixing release consistent with the client-side fixing release? Our motivation for RQ₁ is to understand whether developers are keeping up to date with the package-side fixing releases. We define that package-side and client-side fixing releases are consistent if client-side fixing release follow package-side fixing release. For example, client minor landing and package minor landing combination is consistent, but client major landing and package patch landing combination is not consistent. Our key assumption is that the inconsistent combination requires more migration effort than the consistent one, which in turn is likely to create lags.

(RQ₂) Do lineage freshness and severity influence lags in the fix propagation? Our motivation for RQ₂ is to identify the existence of lags during a propagation. Concretely, we use our defined measures, i.e., propagation influencing factors, to characterize a propagation lag. Our assumption is that a package-side fixing release on the latest lineage with high severity should propagate quickly.

Data Collection - Our data collection consists of (i) vulnerability reports and (ii) the set of cloned npm package and client git repositories. We use the same 2,373 vulnerability reports as shown in our preliminary study which crawled from snyk.io (Snyk [2015]). As inspired by Wittern et al. (2016), we cloned and extracted information of npm package and client from public GitHub repositories. In this study, we consider only normal dependencies listed in the `package.json` file to make sure that the packages are used in the production environment. Hence, other types of dependencies including: (1) `devDependencies`, (2) `peerDependencies`, (3) `bundledDependencies`, and (4) `optionalDependencies` are ignored in this study since they will not be installed in the downstream clients in the production or cannot be retrieved directly from the npm registry. To perform the lags analysis, we first filter reports that do not have the fixing release. We then used the package name and its GitHub link from the reports to automatically match cloned repositories.

As shown in Table 3, our data collection included 2,373 vulnerability reports that disclosed from April 9, 2009 to August 7, 2020. There are 1,290 reports that already published the fixing releases which affect 786 different packages. The statistics of vulnerable packages and reports are presented in the table. For package and client repositories, we collected a repository snapshot from GitHub on August 9, 2020 with 152,074 repositories, 611,468 dependencies, and 1,553,325 releases.

Table 3 A summary of the data collection which used to populate the dataset to answer RQ_1 and RQ_2 .

npm JavaScript Ecosystem Information	
Repository snapshot	Aug 9, 2020
# package & client GitHub repositories	152,074
# total dependencies (with downstream)	611,468
# total packages releases	1,553,325
npm Vulnerability Report Information	
Disclosures period	Apr 9, 2009 – Aug 7, 2020
# reports	2,373
# reports (with fixing release)	1,290
# vulnerable packages	786
# vulnerabilities per package	
- mean	1.64
- median	1.00
- SD	2.58
# high severity vulnerabilities	566
# medium severity vulnerabilities	647
# low severity vulnerabilities	77

Approach to Answer RQ_1 - The data processing to answer RQ_1 involves the package-side fixing release and client-side fixing release extraction. Similar to PQ_1 , we first identify the package-side fixing release by comparing a vulnerable release and a fixing release. To track the client-side fixing release, we then extract the direct clients' version history of the vulnerable packages. A client

Table 4 A summary of dataset information for the empirical study to answer RQ_1 and RQ_2 .

RQ_1 Dataset	
# vulnerability reports (follow semver)	410
# vulnerable packages (follow semver)	230
# direct clients (follow semver)	5,417
# filtered clients (break semver)	4,000
RQ_2 Dataset	
# vulnerability reports (with fix released)	617
# vulnerable packages (with fix released)	344
# downstream clients	416,582

Table 5 A summary number of filtered clients grouped by their update pattern in RQ_1 . There are 4,000 packages and clients that excluded in the RQ_1 .

semver Update Patterns	# clients
Major only	94
Minor only	293
Patch only	2,812
No change	801
Filtered packages & clients	4,000

is deemed vulnerable if its lower-bound dependency falls within the reported upper-bound as listed in a vulnerability report.

To ensure quality, we additionally filter out packages and clients that did not follow semantic versioning as shown Table 5. Our key assumption is to keep packages and clients that follow a semantic version release cycle, i.e., packages and clients should have all the update patterns of major landing, minor landing, and patch landing. As a result, 4,000 packages and clients were filtered out from the dataset. As shown in Table 4, our final dataset for RQ_1 consists of 410 vulnerability reports that affect 230 vulnerable packages and 5,417 direct clients.

The analysis to answer RQ_1 is the identification of lags in the adoption. We show the frequency distribution of client-side fixing release in each package-side fixing release. In order to statistically validate our results, we apply Pearson’s chi-squared test (χ^2) (Pearson, 1900) with the null hypothesis ‘the package-side fixing release and the client-side fixing release are independent’. To show the power of differences between each package-side fixing release and client-side fixing release combination, we investigate the effect size using Cramér’s V (ϕ'), which is a measure of association between two nominal categories (Cramér, 1946). According to Cohen (1988), since the contingency Table 6 has 2 degrees of freedom (df*), effect size is analyzed as follows: (1) $\phi' < 0.07$ as Negligible, (2) $0.07 \leq \phi' < 0.20$ as Small, (3) $0.20 \leq \phi' < 0.35$ as Medium, or (4) $0.35 \geq \phi'$ as Large. To analyze Cramér’s V, we use the `researchpy` package.¹¹

¹¹ <https://pypi.org/project/researchpy/>

Table 6 A contingency table shows the frequency distribution of client-side fixing release for each package-side fixing release. We find that (i) there is a dependency between package-side fixing release and client-side fixing release and (ii) only the case of package minor landing is consistent.

	Package major landing (66)	Package minor landing (99)	Package patch landing (245)	All
Client major landing	448 (46.18%)	453 (30.00%)	657 (22.37%)	1,558
Client minor landing	2 (0.21%)	761 (50.40%)	1,082 (36.84%)	1,845
Client patch landing	0 (0.00%)	0 (0.00%)	625 (21.28%)	625
Dependency removal	520 (53.61%)	296 (19.60%)	573 (19.51%)	1,389
All	970	1,510	2,937	5,417

Approach to Answer RQ2 - The data processing to answer RQ_2 involves propagation influencing factors extraction. There are three steps to track downstream clients and classify lineage freshness and severity. First, we build and traverse in a dependency tree for each package-side fixing release using a breadth-first search (BFS) approach. The meta-data is collected from each downstream client which includes: (i) version, (ii) release date, and (iii) dependency list, i.e., exact version and ranged version. We then classify whether or not a client is vulnerable using an approach similar to RQ_1 . Our method involves removing duplicated clients in the dependency tree, which is caused by the npm tree structure. Second, we classify the lineage freshness of a fixing release by confirming that it is on the latest branch. Finally, we extract the vulnerability severity from the report. As shown in Table 4, our final dataset for RQ_2 consists of 617 vulnerability reports, 344 vulnerable packages with fixing releases, and 416,582 downstream clients.

The analysis to answer RQ_2 is the identification of lags in the propagation. We show a summary statistic of lags in terms of days, i.e., the mean, the median, the standard deviation, and the frequency distribution, with two influencing factors. In order to statistically validate the differences in the results, we apply Kruskal-Wallis non-parametric statistical test (Kruskal and Wallis, 1952). This is a one-tailed test.¹² We test the null hypothesis that ‘lags in the latest and supported lineages are the same’. We investigate the effect size using Cliff’s δ , which is a non-parametric effect size measure (Romano et al., 2006). Effect size are analyzed as follows: (1) $|\delta| < 0.147$ as Negligible, (2) $0.147 \leq |\delta| < 0.33$ as Small, (3) $0.33 \leq |\delta| < 0.474$ as Medium, or (4) $0.474 \leq |\delta|$ as Large. To analyze Cliff’s δ , we use the `cliffsDelta` package.¹³

4.3 Results to the Empirical Study

(RQ₁) Is the package-side fixing release consistent with the client-side fixing release? Our results are summarized into two findings. First, Table 6 shows the evidence that most of package-side fixing releases are package

¹² <https://github.com/scipy/scipy/issues/12231#issuecomment-660404413>

¹³ <https://github.com/neilernst/cliffsDelta>

Table 7 A result of statistical test for RQ_1 . We find that differences between each distribution are significant and have a large level of effect size.

Statistic	Value
Pearson’s chi-squared test (χ^2)	1,484.48
p-value	< 0.001
Cramér’s V (ϕ')	0.37

patch landings. As shown in the first row of a table, we find that 245 out of 410 fixing releases have package patch landings (highlighted in red). We also find that there are 66 package major landings and 99 package minor landings. This finding complements the result of PQ_1 .

Second, Table 6 shows the evidence that there is a dependency between package-side fixing release and client-side fixing release variables. However, there is no consistency across package-side fixing release and client-side fixing release. As highlighted in *Client patch landing* row of Table 6, we find that there are only 21.28% of clients adopt a package patch landing as client patch landings. Instead, clients are more likely have client minor landings, i.e., 36.84% of clients (highlighted in red). For the case of package major landing, there are 53.61% of clients remove their dependencies to avoid vulnerability (highlighted in yellow). The majority of clients that still adopt the package major landing are around 43.18% as client major landing. The only case that we find consistent is package minor landing which 50.40% of clients adopt the fix as client minor landing (highlighted in green).

For the statistical evaluation, we find that there is an association between the package-side fixing release and the client-side fixing release. Table 7 shows that our null hypothesis on ‘the package-side fixing release and the client-side fixing release are independent’ is rejected (i.e., $\chi^2 = 1,484.48$, $p\text{-value} < 0.001$). From the Cramér’s V effect size (ϕ'), we got a value of 0.37 which shows the large level of association.

Answer to RQ_1 : No for the case of package patch landing, but yes for the others. We find that only 21.28% of package patch landing are adopted as a client patch landing in clients. Instead, 36.84% of package patch landing are adopted as a client minor landing. The evidence suggests that since clients keep stale dependencies, more migration effort is required to fix that client due to the potential risk from backward incompatible changes (i.e., client major landing or client minor landing).

(RQ_2) Do lineage freshness and severity influence lags in the fix propagation? Our results are summarized into two findings. First, Table 8 shows the evidence that the lineage freshness influences lags in a propagation. As highlighted in red, we find that LL has more lags than SL in terms of days for every hops, e.g., median of lags for the first hop: 164 days > 89 days.

Table 8 A summary statistic of lags in the propagation (# days) categorized by lineage freshness to show the difference between lags in LL and SL. Lags in the table is not accumulative.

	# hop	# clients	Mean	Median	SD
LL	1	18,444	311.69	164.00	370.92
	2	55,044	299.70	157.00	359.80
	3	74,257	255.21	130.00	313.86
	≥ 4	239,128	212.11	112.00	270.62
	386,873				
SL	1	2,880	217.90	89.00	298.30
	2	5,675	269.55	139.00	318.41
	3	5,948	216.43	101.00	276.27
	≥ 4	15,206	181.55	96.00	240.11
	29,709				

Table 9 A summary statistic of lags in the propagation (# days) categorized by vulnerability severity to show the difference of lags between high, medium, and low severity vulnerability fixes. Lags in the table is not accumulative.

	# hop	# clients	Mean	Median	SD
H	1	5,569	187.09	91.00	239.72
	2	16,160	212.04	103.00	277.36
	3	18,444	190.99	88.00	255.29
	≥ 4	34,189	184.91	94.00	242.13
	74,362				
M	1	14,320	350.54	194.00	399.91
	2	39,758	341.29	193.00	386.66
	3	55,625	280.18	151.00	331.87
	≥ 4	210,571	215.75	116.00	274.23
	320,274				
L	1	1,435	219.39	123.00	248.54
	2	4,801	214.78	127.00	246.28
	3	6,136	184.29	98.00	223.92
	≥ 4	9,574	180.46	94.00	234.38
	21,946				

Second, Table 9 shows the evidence that the vulnerability severity influences lags in a propagation. As highlighted in green, we find that the high severity fixing release has the least lags than others in every hop, e.g., the first hop: 91 days. We also find that the medium severity fix has the most lags than others as highlighted in red, e.g., the first hop: 194 days.

For the statistical evaluation, we find that lags in the latest and supported lineage showed to have a significant (p-value < 0.001), but negligible to small association. Table 10 shows that our null hypothesis on whether ‘lags in the latest and supported lineages are the same’ is rejected, i.e., the first hop to the more than the fourth hop for medium severity, the second hop and more than the fourth hop for low severity; and the second hop for high severity.

Table 10 A comparison of lags in the propagation between clients that adopt the latest lineage and supported lineage fixing release, i.e., by the median. We find that difference between each distribution is significant mostly in the case of medium severity. The effect sizes of those differences are negligible and small level. (LL: median of the latest lineage, SL: median of the supported lineage, *: p-value < 0.001).

# hop	H	M	L
1	LL > SL negligible	LL > SL * small	LL > SL negligible
2	LL > SL * negligible	LL > SL * negligible	LL > SL * small
3	LL > SL negligible	LL > SL * negligible	LL > SL negligible
≥ 4	LL > SL negligible	LL > SL * negligible	LL < SL * negligible

Answer to RQ_2 : Yes, lineage freshness and severity influence lags in the propagation. We find that fixing releases that occur on the latest lineage and medium severity suffer the most lags.

5 Discussion

5.1 Lessons Learned

This section discusses three main implications based on our results in PQ_1 , PQ_2 , RQ_1 , and RQ_2 . These are presented as lessons learned and could have implications for both practitioner and researcher.

1. **Release cycle matters.** According to the results of PQ_2 , fixing commits are small with less than 14.28% of fixing commits in the package-side fixing release. We suspect that vulnerability fix repackaging is a cause for lags. Hence, developers of the vulnerable packages are recommended to release fixes as soon as they have applied the fix, if not, they should highlight these fixes when bundling the fix. Additionally, from 10 randomly selected vulnerabilities, we found that discussions between developers did not include an explicit mention of the vulnerability, i.e., GitHub issue, commit, and pull request. Since developers bundled the fix with other updates, developers may have been unaware. *In summary, researchers should provide strategies for making the most efficient update via the release cycle. For example, (i) releasing an emergency patch that does not introduce any new features for security fixes to maintain backward compatibility and (ii) providing a security support for an active version, i.e., long-term-release version. Furthermore, practitioners can upgrade security fixes as first class citizens, so that the vulnerability fix can travel quicker throughout the ecosystem.*

2. Awareness is important. According to PQ_1 and RQ_1 , 64.50% of fixes are a package patch landing. However, clients are more likely to have client major landing and client minor landing, i.e., 22.37% and 36.84%, than the patch client landing, i.e., 21.28%. Security fixes need to be highlighted in the update note, as a possible reason is for failure to update because client developers are more interested in major features that are highlighted in an update. Recently, some open source communities start to make tools to highlight the vulnerability problems in a software ecosystem. GitHub (GitHub, 2017, 2020) made a new function for notifying a new vulnerability from the dependency list of clients by using a bot. However, GitHub stated that the tool will not be able to catch everything and send the alert notification within a guaranteed time frame. Also, npm (NPM, 2018a) made a new command for listing the vulnerability information in downloaded dependencies of clients and try to automatically fix them called `npm audit`. However, there are some cases that `npm audit` does not work. For example, the immediate dependency does not adopt the package-side fixing release from the vulnerable package (NPM, 2018b). In these cases, a manual review is required. Thus, client developers have to wait for the propagation of the fixing release, i.e., a lag exists. According to RQ_1 , 1,389 of 5,417 clients, i.e., 25.64%, decided to remove vulnerable dependencies to mitigate the risk of vulnerabilities. Instead of waiting for the vulnerability fix, clients might remove the vulnerable dependency if they are able to find a similar package as a replacement or do not want to use the vulnerable dependency anymore.

From a result of PQ_2 , explicit package-side fixing release with a highlight of the vulnerability is needed to speed up the adoption. *In summary, researchers and practitioners need to provide developers more awareness mechanisms to allow quicker planning of the update. The good news is current initiatives like GitHub security are trending towards this.*

3. Migration cost effort. From our first finding of RQ_2 , the package-side fixing release in the latest lineage suffers more lags than the supported lineage in terms of days. A possible reason is that developers of clients consider a package in the supported lineage is more worthwhile to adopt the new release than the latest lineage regardless of the fix. In terms of security, according to the official documentation of npm (NPM, 2016), when a security threat is identified, the following severity policy is put into action: (a) P0: Drop everything and fix!, (b) P1: High severity, schedule work within 7 days, (c) P2: Medium severity, schedule work within 30 days, (d) P3: Low severity, fix within 180 days. Surprisingly, the second finding of RQ_2 shows evidence that low severity fixes are adopted quicker than medium severity fixes. A possible reason for the quicker adopting of low severity could be because the fix is easier to integrate into the application. *In summary, researchers and practitioners that are package developers in npm seem to require additional time before updating their dependencies.*

5.2 Threats to Validity

Internal Validity - We discuss three internal threats. The first threat is the correctness of the tools and techniques used in this study. We use the listed dependencies and version number as defined in the `package.json` meta-file. The threat is that sometimes some dependencies are not listed or the semantic version is invalid and vice-versa, so we applied a filter to remove clients that do not follow the semantic versioning, thus making this threat minimal. The second threat is the tools used to implement our defined terminology (i.e., `numpy`, `scipy`, `gitpython`, and `semantic-version`). To mitigate this, we carefully confirmed our results by manually validating the results for RQ_1 and RQ_2 , then also manually validating results with statistics on the npm website. For the existing tool for suggesting the package-side fixing release like `npm audit`, we found that this tool was inappropriate to use in this work due to its limitation of the `package-lock.json` file is required for analyzing repositories, which only 2.27% of repositories of the dataset and the tool assumes the latest information. Unlike our work, we analyzed the data available in the historical snapshot. As the correctness of dependency relations depends on getting all dependencies, the final internal threat is the validity of our collected data. In this study, our ecosystem is made up of packages and clients that were either affected directly or transitively by at least a single vulnerability. We also make sure that the packages and their repositories are actually listed on a npm registry. We are confident that the results of RQ_2 are not affected by invalid data.

External Validity - The main external threat is the generality of other results to other ecosystems. In this study, we focused solely on the npm JavaScript ecosystem. However, our analysis is applicable to other ecosystems that have similar package management systems, e.g., PyPI for Python, Maven for Java. Immediate future plans include studying the lags in other ecosystems. Another threat is the sample size of the analyzed data. In this study, we analyzed only 1,290 vulnerability reports with package-side fixing releases from 2,373 extracted reports from snyk.io. This small size of sample data might not be able to represent the population. However, we are confident of the data quality and reduced bias as we followed strict methods to validate by two authors for PQ_1 , PQ_2 , and RQ_1 data.

Construct Validity - The key threat is that there may be other factors apart from the two factors, i.e., lineage freshness and vulnerability severity. These factors are based on prior studies, i.e., measuring of dependency freshness from Cox et al (2015), exploring the impact of vulnerability to transitive dependencies from Kikas et al (2017), responding to a vulnerability NPM (2016). For future work, we would like to investigate other factors.

6 Related Work

Complementary related works are introduced throughout the paper, in this section, we discuss some key related works.

On Updating Dependencies - These studies relate to the migration of libraries to the latest versions of libraries. With new libraries and newer versions of existing libraries continuously being released, managing a system library dependencies is a concern on its own. As outlined in [Bogart et al \(2016\)](#); [Rae-maekers et al \(2012\)](#); [Teyton et al \(2012\)](#), dependency management includes making cost-benefit decisions related to keeping or updating dependencies on outdated libraries. Additionally, [Bavota et al \(2015\)](#); [Hora et al \(2015\)](#); [Ihara et al \(2017\)](#); [Robbes et al \(2012\)](#); [Sawant et al \(2016\)](#) showed that updating libraries and their APIs are slow and lagging. [Decan et al \(2017\)](#) showed the comparison of dependency evolution and issue from three different ecosystems. Their results showed that these ecosystems faced the dependency update issue which causes a problem to downstream dependencies, however, there is no perfect solution for this issue. [Kula et al \(2018\)](#) found that such update decisions are not only influenced by whether or not security vulnerabilities have been fixed and important features have been improved, but also by the amount of work required to accommodate changes in the API of a newer library version. [Decan et al \(2018\)](#) performed an empirical study of technical lag in the npm dependency network and found that packages are suffered from the lags if the latest release of dependencies are not covered by their version ranges in *package.json* file. They also found that semantic versioning could be used in order to reduce the technical lag. [Mirhosseini and Parnin \(2017\)](#) studied about the pull request notification to update the dependencies. Their results showed that pull request and badge notification can reduce lags, however, developers are often overwhelmed by lots of notifications. Another study by [Abdalkareem et al \(2017\)](#) focused on the trivial packages and found that they are common and popular in the npm ecosystem. They also suggested that developers should be careful about the selection of packages and how to keep them updated. **Our work** focuses on the lag of the fixing release adoption and propagation with the influencing factors (i.e., lineage freshness and severity). We also expand our study from prior works by taking transitive clients (i.e., downstream propagation) into consideration. Our work complements the findings of prior work, with the similar goal of encouraging developers to update.

On Malware and Vulnerabilities - These studies relate to the security vulnerability within a software ecosystem from various aspects. [Decan et al \(2018\)](#) explored the impact of vulnerability within the npm ecosystem by analyzing the reaction time of developers from both vulnerable packages and their direct dependent packages to fix the vulnerabilities. They also considered the reactions of developers from different levels of vulnerability severity. They found that the vulnerabilities were prevalent and took several months to be fixed. Several studies also explored the impact of vulnerability within various ecosystems by analyzing the dependency usage ([Hejderup, 2015](#); [Kikas et al, 2017](#); [Lauinger et al, 2017](#); [Linares-Vásquez et al, 2017](#)). Some studies tried

to characterize the vulnerability and its fix in various ecosystems other than the npm ecosystem (Li and Paxson, 2017; Piantadosi et al, 2019). There is a study about the relationship between bugs and vulnerabilities, to conclude that the relationship is weak (Munaiah et al, 2017). In order to increase the developers' awareness of the security vulnerability, some studies tried to create a tool to detect and alert vulnerability when it disclosed (Cadariu et al, 2015; GitHub, 2018b). There is a study about addressing the over-estimation problem for reporting the vulnerable dependencies in open source communities and the impact of vulnerable libraries usage in the industry (Pashchenko et al, 2018). Additionally, some studies tried to predict the vulnerability of software systems by analyzing the source code (Alhazmi et al, 2007; Chowdhury and Zulkernine, 2011; Shin and Williams, 2008). *Our work* takes a look at the package-side fixing release and its fix at the commit-level of npm package vulnerabilities, instead of the release-level in prior studies. We also propose a set of definitions to characterize the package-side fixing release and client-side lags, which covers both direct and transitive clients. Prior works, instead, only analyze the direct clients.

Mining-related Studies - These studies relate to the mining techniques in the software repository and software ecosystem. The first step of software repository mining is data collection and extraction. Researchers need to have data sources and know about which part of the data can use in their work. In the case of the npm package repository, we can extract the information of packages from `package.json` meta-file (Mirhosseini and Parnin, 2017; Wittern et al, 2016). In the case of the security vulnerability, we can collect data from Common Weakness Enumeration (CWE) (Mitre Corporation, 2018b) and Common Vulnerabilities and Exposures (CVE) (Mitre Corporation, 2018a) database (Alhazmi et al, 2007; Cadariu et al, 2015; Chowdhury and Zulkernine, 2011; Lauinger et al, 2017; Linares-Vásquez et al, 2017; Munaiah et al, 2017). To study the issues within the software ecosystem, we also define the traversal of the downstream clients by using the dependency list of clients. These studies introduce some techniques to model the dependency graph (Bavota et al, 2015; Hejderup, 2015; Kikas et al, 2017). *Our work* uses similar mining techniques to extract the dependencies as well as construct the ecosystem. In our work, we manually extract and investigate the commits to understand the contents of the fix.

7 Conclusion

Security vulnerability in third-party dependencies is a growing concern for software developers as the risk of it could be extended to the entire software ecosystem. To ensure quick adoption and propagation of a fixing release, we conduct an empirical investigation to identify lags that may occur between the vulnerable release and its fixing release from a case study of npm JavaScript ecosystem. We found that the package-side fixing release is rarely released on its own, with up to 85.72% of the bundled commits in a package-side fixing

release being unrelated to the fix. We then found that a quick package-side fixing release (i.e., package patch landing) does not always ensure that a client will adopt it quicker, with only 17.69% of clients matching a package patch landing to a client patch landing. Furthermore, factors such as the lineage freshness and the vulnerability severity have a small effect on its propagation.

In addition to theses lags that we identified and characterized, this paper lays the groundwork for future research on how to mitigate these propagation lags in an ecosystem. We suggest that researchers should provide strategies for making the most efficient update via the release cycle. Practitioners also need more awareness to allow quicker planning of the update. Potential future avenues for researchers include (i) a developer survey to a better understanding of the reason for releasing and adopting fixes, (ii) a performance improvement plan for highlighting the fixing release tool, (iii) a tool for managing and prioritizing vulnerability fixing process.

Acknowledgment

This work was supported by the Japan Society for Promotion of Science (JSPS) KAKENHI Grant Numbers JP18H04094, JP18H03221, 20K19774, 20H05706.

References

- Abdalkareem R, Nourry O, Wehaibi S, Mujahid S, Shihab E (2017) Why Do Developers Use Trivial Packages? An Empirical Case Study on npm. In: Proceedings of the 11th Joint Meeting on Foundations of Software Engineering (ESEC/FSE), pp 385–395
- Alhazmi OH, Malaiya YK, Ray I (2007) Measuring, analyzing and predicting security vulnerabilities in software systems. *Computers and Security* 26(3):219–228
- Bavota G, Canfora G, Di Penta M, Oliveto R, Panichella S (2015) How the Apache Community Upgrades Dependencies: An Evolutionary Study. *Empirical Software Engineering (EMSE)* 20(5):1275–1317
- Bennett JT (2014) Shellshock in the Wild - Shellshock in the Wild. <https://www.fireeye.com/blog/threat-research/2014/09/shellshock-in-the-wild.html>, (Accessed on 02/08/2021)
- Bogart C, Kästner C, Herbsleb J, Thung F (2016) How to Break an API: Cost Negotiation and Community Values in Three Software Ecosystems. In: Proceedings of the 24th International Symposium on the Foundations of Software Engineering (FSE), pp 109–120
- Cadariu M, Bouwers E, Visser J, van Deursen A (2015) Tracking known security vulnerabilities in proprietary software systems. In: Proceedings of the 22nd International Conference on Software Analysis, Evolution, and Reengineering (SANER), pp 516–519

- Chowdhury I, Zulkernine M (2011) Using complexity, coupling, and cohesion metrics as early indicators of vulnerabilities. *Journal of Systems Architecture (JSA)* 57(3):294–313
- Cohen J (1988) *Statistical Power Analysis for the Behavioral Sciences*. Routledge
- Cox J, Bouwers E, Eekelen M, Visser J (2015) Measuring dependency freshness in software systems. In: 2015 IEEE/ACM 37th IEEE International Conference on Software Engineering (ICSE), pp 109–118
- Cramér H (1946) *Mathematical Methods of Statistics*. Princeton University Press
- Decan A, Mens T, Claes M (2017) An empirical comparison of dependency issues in OSS packaging ecosystems. In: *Proceedings of the 24th International Conference on Software Analysis, Evolution, and Reengineering (SANER)*, pp 2–12
- Decan A, Mens T, Constantinou E (2018) On the evolution of technical lag in the npm package dependency network. In: *Proceedings of the 34th International Conference on Software Maintenance and Evolution (ICSME)*, pp 404–414
- Decan A, Mens T, Constantinou E (2018) On the impact of security vulnerabilities in the npm package dependency network. In: *Proceedings of the 15th International Conference on Mining Software Repositories (MSR)*, pp 181–191
- Durumeric Z, Li F, Kasten J, Amann J, Beekman J, Payer M, Weaver N, Adrian D, Paxson V, Bailey M, Halderman JA (2014) The matter of heart-bleed. In: *Proceedings of the 2014 Conference on Internet Measurement Conference (IMC)*, ACM
- GitHub (2017) About security alerts for vulnerable dependencies. <https://help.github.com/articles/about-security-alerts-for-vulnerable-dependencies/>, (Accessed on 02/08/2021)
- GitHub (2018a) How security alerts are keeping your code safer. <https://github.blog/2018-03-21-security-alerting-a-first-look-at-community-responses/>, (Accessed on 02/08/2021)
- GitHub (2018b) Viewing and updating vulnerable dependencies in your repository. <https://help.github.com/articles/viewing-and-updating-vulnerable-dependencies-in-your-repository/>, (Accessed on 02/08/2021)
- GitHub (2020) Keep all your packages up to date with dependabot - the github blog. <https://github.blog/2020-06-01-keep-all-your-packages-up-to-date-with-dependabot/>, (Accessed on 02/08/2021)
- Hejderup J (2015) *In Dependencies We Trust: How vulnerable are dependencies in software modules?* Master's thesis, Delft University of Technology
- Hora A, Robbes R, Anquetil N, Etien A, Ducasse S, Valente MT (2015) How Do Developers React to API Evolution? The Pharo Ecosystem Case. In:

- Proceedings of the 31st International Conference on Software Maintenance and Evolution (ICSME), pp 251–260
- Howard M, Leblanc D (2001) Writing Secure Code. Redmond, WA, USA
- Ihara A, Fujibayashi D, Suwa H, Kula RG, Matsumoto K (2017) Understanding When to Adopt a Library: A Case Study on ASF Projects. In: Proceedings of the 13th International Conference on Open Source Systems (OSS), pp 128–138
- Kikas R, Gousios G, Dumas M, Pfahl D (2017) Structure and Evolution of Package Dependency Networks. In: Proceedings of the 14th International Conference on Mining Software Repositories (MSR), pp 102–112
- Kruskal WH, Wallis WA (1952) Use of Ranks in One-Criterion Variance Analysis. *Journal of the American Statistical Association* 47(260):583–621
- Kula RG, De Roover C, German DM, Ishio T, Inoue K (2018) A generalized model for visualizing library popularity, adoption, and diffusion within a software ecosystem. In: Proceedings of the 25th International Conference on Software Analysis, Evolution and Reengineering (SANER), pp 288–299
- Kula RG, German DM, Ouni A, Ishio T, Inoue K (2018) Do developers update their library dependencies? *Empirical Software Engineering (EMSE)* 23(1):384–417
- Lauinger T, Chaabane A, Arshad S, Robertson W, Wilson C, Kirda E (2017) Thou Shalt Not Depend on Me: Analysing the Use of Outdated JavaScript Libraries on the Web. In: Proceedings of the 24th Network and Distributed System Security Symposium (NDSS)
- Li F, Paxson V (2017) A large-scale empirical study of security patches. In: Proceedings of the 24th ACM SIGSAC Conference on Computer and Communications Security (CCS), p 2201–2215
- Linares-Vásquez M, Bavota G, Escobar-Velásquez C (2017) An Empirical Study on Android-related Vulnerabilities. In: Proceedings of the 14th International Conference on Mining Software Repositories (MSR), pp 2–13
- Mirhosseini S, Parnin C (2017) Can automated pull requests encourage software developers to upgrade out-of-date dependencies? In: Proceedings of the 32nd IEEE/ACM International Conference on Automated Software Engineering (ASE), pp 84–94
- Mitre Corporation (2018a) CVE - Common Vulnerabilities and Exposures (CVE). <https://cve.mitre.org/>, (Accessed on 02/08/2021)
- Mitre Corporation (2018b) CWE - Common Weakness Enumeration. <https://cwe.mitre.org/>, (Accessed on 02/08/2021)
- Munaiah N, Camilo F, Wigham W, Meneely A, Nagappan M (2017) Do bugs foreshadow vulnerabilities? An in-depth study of the chromium project. *Empirical Software Engineering (EMSE)* 22(3):1305–1347
- Nguyen VH, Dashevskiy S, Massacci F (2016) An automatic method for assessing the versions affected by a vulnerability. *Empirical Software Engineering (EMSE)* 21(6):2268–2297
- NPM (2010) npm. <https://www.npmjs.com/>, (Accessed on 02/08/2021)
- NPM (2016) Responding to Security Threats and Critical Updates . <https://www.npmjs.com/policies/security#>

- [responding-to-security-threats-and-critical-updates](#), (Accessed on 02/08/2021)
- NPM (2018a) Auditing package dependencies for security vulnerabilities. <https://docs.npmjs.com/auditing-package-dependencies-for-security-vulnerabilities>, (Accessed on 02/08/2021)
- NPM (2018b) Security vulnerabilities found requiring manual review. <https://docs.npmjs.com/auditing-package-dependencies-for-security-vulnerabilities#security-vulnerabilities-found-requiring-manual-review>, (Accessed on 02/08/2021)
- Pashchenko I, Plate H, Ponta SE, Sabetta A, Massacci F (2018) Vulnerable Open Source Dependencies: Counting Those That Matter. In: Proceedings of the 12th International Symposium on Empirical Software Engineering and Measurement (ESEM), pp 42:1–42:10
- Pearson K (1900) X. on the criterion that a given system of deviations from the probable in the case of a correlated system of variables is such that it can be reasonably supposed to have arisen from random sampling. The London, Edinburgh, and Dublin Philosophical Magazine and Journal of Science 50(302):157–175
- Piantadosi V, Scalabrino S, Oliveto R (2019) Fixing of security vulnerabilities in open source projects: A case study of apache http server and apache tomcat. In: Proceedings of the 12th IEEE Conference on Software Testing, Validation and Verification (ICST), pp 68–78
- Ponta SE, Plate H, Sabetta A (2018) Beyond Metadata: Code-centric and Usage-based Analysis of Known Vulnerabilities in Open-source Software. In: Proceedings of the 34th International Conference on Software Maintenance and Evolution (ICSME), pp 58–68
- Preston-Werner T (2009) Semantic Versioning 2.0.0. <https://semver.org/>, (Accessed on 02/08/2021)
- Raemaekers S, van Deursen A, Visser J (2012) Measuring software library stability through historical version analysis. In: Proceedings of the 28th International Conference on Software Maintenance (ICSM), pp 378–387
- Robbes R, Lungu M, Röthlisberger D (2012) How Do Developers React to API Deprecation?: The Case of a Smalltalk Ecosystem. In: Proceedings of the 20th International Symposium on the Foundations of Software Engineering (FSE), pp 56:1–56:11
- Romano J, Kromrey JD, Coraggio J, Skowronek J, Devine L (2006) Exploring methods for evaluating group differences on the nsse and other surveys: Are the t-test and cohen’sd indices the most appropriate choices. In: Annual meeting of the Southern Association for Institutional Research, Citeseer
- Sawant AA, Robbes R, Bacchelli A (2016) On the reaction to deprecation of 25,357 clients of 4+1 popular Java APIs. In: Proceedings of the 32th International Conference on Software Maintenance and Evolution (ICSME), pp 400–410

- Shin Y, Williams L (2008) An empirical model to predict security vulnerabilities using code complexity metrics. In: Proceedings of the 2nd International Symposium on Empirical Software Engineering and Measurement (ESEM), pp 315–317
- Snyk (2015) Vulnerability DB. <https://snyk.io/vuln>, (Accessed on 02/08/2021)
- Snyk (2017a) Denial of Service (DoS) - npm:connect. <https://snyk.io/vuln/npm:connect:20120107>, (Accessed on 02/08/2021)
- Snyk (2017b) Insecure defaults in faye | snyk. <https://snyk.io/vuln/npm:faye:20121107>, (Accessed on 02/08/2021)
- Snyk (2017c) Symlink attack due to predictable tmp folder names in npm | snyk. <https://snyk.io/vuln/npm:npm:20130708>, (Accessed on 02/08/2021)
- Synopsys (2014) Heartbleed Bug. <http://heartbleed.com/>, (Accessed on 02/08/2021)
- Teyton C, Falleri JR, Blanc X (2012) Mining Library Migration Graphs. In: Proceedings of the 19th Working Conference on Reverse Engineering (WCRE), pp 289–298
- Williams L, McGraw G, Migue S (2018) Engineering Security Vulnerability Prevention, Detection, and Response. *IEEE Software* 35(5):76–80
- Wittern E, Suter P, Rajagopalan S (2016) A Look at the Dynamics of the JavaScript Package Ecosystem. In: Proceedings of the 13th International Conference on Mining Software Repositories (MSR), pp 351–361
- Zerouali A, Constantinou E, Mens T, Robles G, Gonzalez-Barahona J (2018) An empirical analysis of technical lag in npm package dependencies. In: Proceedings of the 17th International Conference on Software Reuse (ICSR), pp 95–110



Bodin Chinthanet is a Ph.D. student at Nara Institute of Science and Technology, Japan. His research interests include empirical software engineer-

ing and mining software repositories. He is focusing on the security vulnerabilities in software ecosystems and how developers react to vulnerabilities in their software projects. Website: <https://bchinthanet.com/>.



Raula Gaikovina Kula is an Assistant Professor at Nara Institute of Science and Technology. He received the Ph.D. degree from Nara Institute of Science and Technology in 2013. His interests include Software Libraries, Software Ecosystems, Code Reviews and Mining Software Repositories.



Shane McIntosh is an Associate Professor at the University of Waterloo. Previously, he was an Assistant Professor at McGill University. He received his Ph.D. from Queen's University. In his research, Shane uses empirical methods

to study build systems, release engineering, and software quality. Website: <http://shanemcintosh.org/>.



Takashi Ishio received the Ph.D. degree in information science and technology from Osaka University in 2006. He was a JSPS Research Fellow from 2006-2007. He was an assistant professor at Osaka University from 2007-2017. He is now an associate professor of Nara Institute of Science and Technology. His research interests include program analysis and software reuse.



Akinori Ihara is a lecturer at Wakayama University in Japan. His research interests include empirical software engineering, open source software engineering, social software engineering and mining software repositories (MSR). His work has been published at premier venues like ICSE, MSR, and ISSRE. He

received the M.E. degree (2009) and Ph.D. degree (2012) from Nara Institute of Science and Technology. More about Akinori and his work is available online at <http://www.wakayama-u.ac.jp/~ihara/>.



Kenichi Matsumoto is a professor in the Graduate School of Science and Technology at Nara Institute Science and Technology, Japan. His research interests include software measurement and software process. Matsumoto has a Ph.D. in information and computer sciences from Osaka University. He is a senior member of IEEE, and a member of the IEICE, and the IPSJ. Contact him at matumoto@is.naist.jp.