

# The Sound of Silence: Mining Security Vulnerabilities from Secret Integration Channels in Open-Source Projects

Ralf Ramsauer

University of Applied Sciences Regensburg  
ralf.ramsauer@oth-regensburg.de

Daniel Lohmann

University of Hanover  
lohmann@sra.uni-hannover.de

Lukas Bulwahn

BMW AG  
lukas.bulwahn@bmw.de

Wolfgang Mauerer

University of Applied Sciences Regensburg,  
Siemens Corporate Research  
wolfgang.mauerer@oth-regensburg.de

## ABSTRACT

Public development processes are a key characteristic of open source projects. However, fixes for vulnerabilities are usually discussed privately among a small group of trusted maintainers, and integrated without prior public involvement. This is supposed to prevent early disclosure, and cope with embargo and non-disclosure agreement (NDA) rules. While regular development activities leave publicly available traces, fixes for vulnerabilities that bypass the standard process do not.

We present a data-mining based approach to detect code fragments that arise from such infringements of the standard process. By systematically mapping public development artefacts to source code repositories, we can exclude regular process activities, and infer irregularities that stem from non-public integration channels. For the Linux kernel, the most crucial component of many systems, we apply our method to a period of seven months before the release of Linux 5.4. We find 29 commits that address 12 vulnerabilities. For these vulnerabilities, our approach provides a temporal advantage of 2 to 179 days to design exploits before public disclosure takes place, and fixes are rolled out.

Established responsible disclosure approaches in open development processes are supposed to limit premature visibility of security vulnerabilities. However, our approach shows that, instead, they open *additional* possibilities to uncover such changes that thwart the very premise. We conclude by discussing implications and partial countermeasures.

## CCS CONCEPTS

• **Security and privacy** → *Human and societal aspects of security and privacy*; **Software security engineering**; • **Software and its engineering** → *Open source model*; **Search-based software engineering**.

## KEYWORDS

vulnerability mining; software repository mining; process analysis

## ACM Reference Format:

Ralf Ramsauer, Lukas Bulwahn, Daniel Lohmann, and Wolfgang Mauerer. 2020. The Sound of Silence: Mining Security Vulnerabilities from Secret Integration Channels in Open-Source Projects. In *2020 Cloud Computing*

*Security Workshop (CCSW '20), November 9, 2020, Virtual Event, USA.* ACM, New York, NY, USA, 11 pages.

## 1 INTRODUCTION

On 14 August 2018, a series of patches was integrated in Linux to provide mitigations for the Level 1 Terminal Fault (L1TF) [42, 43] vulnerability<sup>1</sup> – a speculative execution attack with severe consequences that enable large scale data leakage across virtual machines on Intel-based cloud appliances. While associated CVE entries were already filed in December 2017 [1], the vulnerability was embargoed until 14 August 2018 [17] – the same day of the disclosure and integration of the critical patches for Linux. Unlike ordinary patches, these patches were—for obvious reasons—not discussed and developed on one of Linux’s public communication channels (i.e., mailing lists) beforehand.

However, the fact that a patch was *not* publicly discussed betrays it: we will show that it is possible to detect such patches as soon as they enter a public repository. This gives attackers valuable information advantage to design exploits. For the aforementioned attack, it took another five days until the patches were integrated and rolled out by Debian 9,<sup>2</sup> a popular and wide-spread Linux distribution.

In this paper, we present a methodology to reverse engineer development processes: We collect all publicly available development artefacts and map them against the software repository. Using techniques to connect developer communication with repository entries [37], we are able to uncover commits from non-public secret integration channels with high probability using semi-automatic methods. In this paper, and without loss of generality, we exercise the approach for mail-based development workflows (as used by the Linux kernel, QEMU, GCC, and many other projects), and show how to systematically obtain *off-list patches*: Code changes that were developed outside the official public lists. Besides fixes for security vulnerabilities, we also find that there exist systematic channels to inject code into the Linux kernel while bypassing public discussion. Our method provides two advantages for malicious attackers: (a) it significantly reduces search efforts for fixes of security vulnerabilities, compared to fully manual investigation, and (b) it provides temporal advantage for the design of attacks. We claim following contributions:

to appear in *CCSW '20, November 9, 2020, Virtual Event, USA*

<sup>1</sup> See Linux commit 958f338e96 (hyperlink available in PDF).

<sup>2</sup> See the announcement of Debian kernel 4.9.110-3+deb9u3.

- We present a method to systematically detect development process infringements in open-source projects that works as soon as commits arrive in repositories
- We detect and categorise different types of secret integration channels of the Linux kernel, such as *bypass of development processes* or non-publicly discussed fixes for *security vulnerabilities*
- We discuss methods that mitigate potential threats to open-source software ecosystems

*Outline.* The rest of this paper is structured as follows: We first give an overview of the problem statement in Section 2. In Section 3, we provide a quick introduction to common open-source development practises as they are, for example, implemented by the Linux kernel community. We then present our methodology of mapping development artefacts in Section 4. In Sections 5, we run our analysis and evaluate a certain time window of the Linux kernel. The discussion in Section 6 focuses on potential threats to the ecosystem. Section 7 presents related work. Finally, Section 8 concludes the paper and gives an outlook on future work.

## 2 SECRET INTEGRATION CHANNELS


The openness of the development processes is a key aspect of any open-source software (OSS) project: Almost all development activities happen in public. Since development artefacts (i.e., discussions or patch data on public mailing lists) are observable, this allows us to analyse the process in detail.

However, especially the development of fixes for critical security vulnerabilities intentionally happens behind closed scenes [41]. After their disclosure, fixes silently appear as commits in the repository. Nonetheless, those commits can not be assigned to any prior artefact that relates to its public pre-integration history. Unless the vulnerability is explicitly announced or attracts medial attention, we disprove the common belief that patches typically drown in the noise of other commits in the repository.

Nevertheless, a full coverage of all public available development resources allows us to systematically exclude *regular development noise* in order to separate it from irregularities: We mine for commits in repositories that come from secret integration channels—and detect them just-in-time to design exploits for vulnerabilities. Figure 1 illustrates the chase for missing links: we deduce that commits that can not be assigned to publicly observable artefacts must arise from secret integration channels.

Collaborative development tools, such as version control systems, bug trackers, continuous integration software or mailing lists [28], are central hubs in modern software projects. Those tools gather collateral development artefacts and traces that directly relate to the the final result of the development process: an actual change of code in terms of a commit in a version control system, such as, for example, git [7].

In their everyday work, OSS developers present their patches to the public as part of the integration process. On communication platforms, changes are discussed and reviewed before they are integrated by maintainers, trusted individuals that are authorised to commit changes to official resources [34]. While several web-based collaboration systems aim to ease workflows of development processes [28], especially system software (e.g., operating systems,



**Figure 1: Disclosing secret integration channels. On the left: artefacts on public channels (e.g., patches on mailing lists) are assigned to commits in the repository. On the right: Commits that lack assignable public artefacts arise from secret integration channels.**

system level libraries or compilers) back mailing lists as their predominant tool of choice [6, 19]. Patches are wrapped in mails, sent to public lists and distributed to all subscribers. Everyone is welcome to join the discussion in the mail thread and comment on the patch. Later, the patch is picked up by a maintainer, who integrates it to their repository. It is not untypical that maintainers fine-tune a patch before they apply it [6, 41].

As one of the world’s largest software undertakings [44], the Linux kernel is the core of a popular and wide-spread operating system and one of largest projects that follows this mail-based development model [41]. More than 10,000 patches are integrated into each major release in eight week cadence.

Nevertheless, a rapidly changing code base, size and complexity inherently results in software defects that can lead to severe software vulnerabilities. In 2019, 170 CVE entries were filed for all different versions and flavours of the Linux kernel, and many more potential security vulnerabilities have been fixed without CVE analysis and assignment [12]. Unavoidably, the *kernel community* has processes on managing critical vulnerabilities.

In contrast to regular development activities, vulnerabilities shall be reported to and discussed on private mailing lists [41]. The rationale behind private discussions is the *responsible disclosure* vulnerability disclosure model: Software producers get the chance to provide fixes for vulnerabilities before they are publicly disclosed [9]. Therefore, security mailing lists are closed-recipients lists to avoid early public attention. Only carefully selected and trusted individuals have permission to join those lists. Security lists are used for coordination, and to setup private communication between reporters and affected subsystems. They can also be used to develop the actual fixes for the issues [27].

Eventually, when the fix is in its final state, it is released for all affected version of the kernel that are supported by the community: This leaves the first publicly visible footprint of the vulnerability: the patch(es) in the repository. Yet, it misses a link to a publicly observable artefact.

### 3 LINUX KERNEL DEVELOPMENT PROCESS

Given enough eyeballs, all bugs are shallow

---

*Linus' Law – by Eric S. Raymond [38]*

This section gives a brief overview of the Linux kernel development process (LKDP). One peculiarity of the LKDP is the large number of contributors (thousands per year) and participants, which lead to the well-known hypothesised connection given above between the decreasing difficulty of detecting bugs with an increasing number of reviewers. Since we abuse the principle to detect patches that have seemingly *not* receive sufficient public attention, it is pertinent to recapture key characteristics of the development process that are relevant for our approach.

#### 3.1 Core Characteristics

Development of the Linux kernel proceeds in two-phase cycles: New code and features are merged during a two-week long *merge window*, which is followed by a two-month long *stabilisation window* [31]. This leads to development cycles of approximately 2.5 months between two major releases. More than 10,000 patches are integrated in each cycle into Linus Torvalds' (the project owner's) git tree, which is commonly called *Linux mainline*. Before code changes (*patches*) are integrated into mainline, they must have been discussed on a public mailing list. This is demanded by the submission guidelines of Linux, and is intended to ascertain good code quality [35]. Because of scalability, availability, robustness, and simplicity, many low-level system software components prefer mail-based communication over using web-based technologies [25]. Other communication channels need not be considered for our purposes.

Similar to a commit in a repository, an email encapsulates a patch that contains a commit message, an informal description of the changes, and a *diff* that specifies insertions and deletions of code—relative to a specific code base. Typically, larger logical changes are split into multiple small patches. This gives a *patch series* whose elements are tied together by a *cover letter*. Cover letters give an informal, higher-level overview of the series. Together with the proper patches, it is sent as a mail thread to maintainers and the corresponding list(s) of the affected subsystem(s) of the project.

Everyone can join the discussion of patches as lists are usually unmoderated. Maintainers who are responsible for the list to which the series is posted, or for a subsystem that the patch addresses, eventually (a) refuse the patch, (b) ask for further refinement of the patch, (c) pick up the patch and commit it to their maintainer tree. Maintainer trees are staging points before code changes are finally integrated mainline. It is not unusual that (b) is repeated over several iterations until the patch series is deemed acceptable for merging.

Because of the massive number of emails and patches, the Linux kernel currently utilises over 200 different mailing lists that are logically partitioned by topic or subsystem. On average, an email is received by one of those lists every 20 seconds.

Maintainers are organised in a semi-formal hierarchy [33]. During a *merge window*, maintainers ask hierarchically higher-level maintainers to *pull their changes*, which is possible in two ways:

Either by picking up and integrating patch data from mailing lists, or by *pulling* code from repositories. Once the top-level maintainer Linus Torvalds pulls and publishes changes, they become part of Linux mainline.

#### 3.2 Lifecycle Management

The latest release of Linux is called the *stable tree*, and is actively supported with bug-fixes until the next mainline release is cut, and becomes the new stable tree. Additionally, the Linux kernel community supports several versions of the kernel in parallel [24] that are referred to as long-term support (LTS) versions. They are based on selected stable trees, and receive official support for up to six years. Figure 2 illustrates the parallel development of mainline Linux and the maintenance of LTS versions.

Linux distributions and vendors usually choose LTS versions as the basis of their kernel (which may additionally contain a substantial amount of added drivers, domain-specific features, and many other additional elements), since they provide a stable and reliable base that will not be subjected to invasive changes (e.g., API changes) during their lifetime. New features are only accepted mainline. Stable and LTS trees may only receive stabilisation patches, bug fixes, or fixes for vulnerabilities.

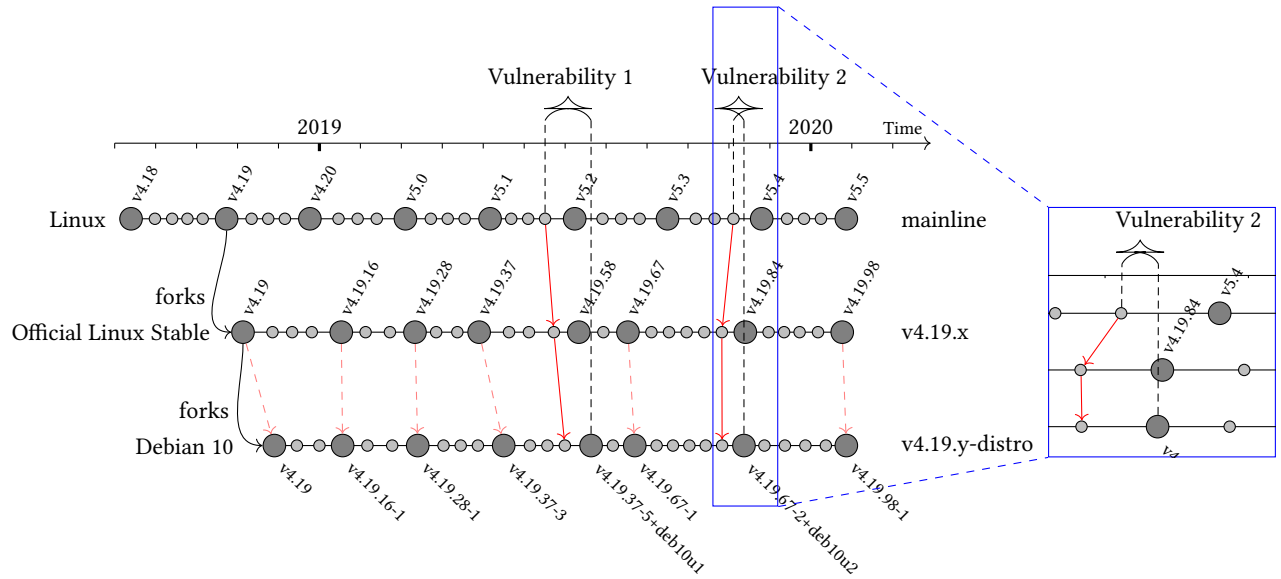
In case patches to LTS versions are also relevant for mainline, they must be, by the *upstream first* convention, integrated in mainline before they are ported back to stable releases. After their release, distributions pick patches from stable versions and apply them to their own kernel repository. From a temporal perspective, the typical pathway of a bug fix is mainline→stable→distribution.

#### 3.3 Exceptional Vulnerability Handling

The aforementioned public review and integration process allows for an exception when fixes for security vulnerabilities must be handled. The Linux kernel is a key software component of a large class of machines from embedded industrial control appliances to cloud computing servers. Consequently, the Linux kernel community has established standard procedures for responsible disclosure [11, 14].

Linux submission guidelines encourage developers to report exploitable security bugs to the non-public security team mailing list `security@kernel.org`: “For severe bugs, a short embargo may be considered to allow distributors to get the patch out to users; in such cases, obviously, the patch should not be sent to any public lists.” [41]

Similar to the regular public development process, patches for vulnerabilities are iteratively discussed, reviewed and refined – but all related conversations take either place in private email conversations or on closed lists. Once participants agree on a fix [27], or after embargoes are expired, the majority of fixes follow the same procedures as bugs: Patches for mainline and affected stable versions are published at the same time, before they are integrated into distribution repositories. Figure 2 (Vulnerability 1) illustrates the temporal process of a typical vulnerability. There is a second type of coordinated disclosure for severe vulnerabilities that we discuss in Section 6.



**Figure 2: Linux development timeline: Mainline, stable trees and distribution trees are supported in parallel. Typically, fixes for vulnerabilities are first fixed in mainline (cf. Vulnerability 1) and on a stable tree, before they are ported back by distributions. In rare cases (cf. Vulnerability 2), patches appear in distributions before they are published in mainline.**

## 4 METHODOLOGY

The basic idea of our approach is to mine public Linux repositories for patches that have not been discussed on any public mailing list. Previous work [18, 19] associates ongoing development in email threads to commits in repositories such that pre-integration histories of changes can be uncovered. For our purposes, we basically need to reverse the question, and find commits in repositories that do *not* enjoy any traceable pre-integration history.

Temporally, a patch should first appear on a public mailing list before it can be found in the repository. Hence, any new commits in the repository that can not be assigned to emails were integrated through non-public integration channels.

The nature of OSS allows for collecting all required development artefacts. However, assigning patches on mailing lists to commits in repositories is a non-trivial task caused by the lack of reliable, machine-encoded provenance information.

### 4.1 Information Processing

Commits in repositories are identified by their unique commit hash, and patches on mailing lists are identified by their unique *Message-ID*. Yet, the mapping of Message-IDs to commit hashes is lost during the manual integration process of maintainers [7, 19].

The kernel community is aware of the gap of code traceability, which is a frequent subject of meta-discussions [3, 22] on the improvement of the development process. The issue is unsolved by the community at the time of writing.

Recent publications from the software-engineering community attempt to use mining techniques to reconstruct the pre-integration history of software projects in ex-post analyses (see, for instance,

Refs. [19, 37]). Simple textual comparison fails to recover the history with high accuracy [19], as patches on mailing lists may significantly differ from their counterparts in the repository: Maintainers rewrite commit messages, add additional changes to the code, move code, or apply patches against a different state if the code base [7, 19], which can (slightly) change the content of the committed itself. The initial version of a patch can significantly vary from further revisions or from its final state in the repository.


While several approaches to rate similarity rating of patches have been devised [18, 19, 37], each provides a similarity score  $\text{sim}$  for a pair of patches. For the comparison of patches, we do, in the first place, not differentiate between patches as commits in repositories or patches as mails on mailing lists. Let  $\mathcal{M}$  be the set of patches on mailing lists and  $\mathcal{C}$  be the set of commits in the repository. Let further  $\mathcal{U}$  be *universe of patches*  $\mathcal{U} = \mathcal{M} \cup \mathcal{C}$ . We can define  $\text{sim}$  as:

$$\text{sim} : \mathcal{U} \times \mathcal{U} \rightarrow [0, 1] \quad (1)$$

where 0 denotes no similarity, and 1 denotes textual equivalence. The operator  $\text{sim}$  considers various tuneable aspects for the comparison of two patches. We choose a method due to Ramsauer et al. [37] that provides high accuracy.

The overall rating for the similarity of two patches in this approach consists of a similarity score of the commit message and a similarity score for the diff. Both are weighted by a heuristic factor. The score for the comparison of commit messages and diffs is mainly based on token-based Levenshtein string distances [29] (for other details that are not relevant for our purpose, we refer to Ref. [37]).

The approach turns assigning patches in mails to commits in repositories to a problem in graph theory: The universe  $\mathcal{U}$  forms



**Figure 3: Creating clusters of similar patches. Patches that exceed the threshold  $t = 0.8$  form subgraphs of similar patches. Cluster  $\alpha$ : contains patches on mailing lists as well as commits, cluster  $\beta$  only contains patches on mailing lists. Cluster  $\gamma$  is not mapped to any mail artefact and a potential off-list patch.**

the vertices of an undirected and weighted graph  $G = (\mathcal{U}, E)$ . The weight  $w(e)$  of an edge  $e = \{e_1, e_2\} \in E$  is determined by a function  $\text{sim}$  that rates the similarity of two patches  $e_1$  and  $e_2$ :

$$w(e) := \text{sim}(e_1, e_2) \quad (2)$$

Determining all edges in  $E$  requires  $\binom{|\mathcal{U}|}{2}$  computational expensive calls of  $\text{sim}$ . Therefore, prefiltering strategies mitigate the combinatorial explosion for a high number of patches  $|\mathcal{U}|$ : Patches are only compared if they, for example, modify at least one common file. Further prefiltering methods are explained in [37].

The graph  $G$  is used to derive an undirected and unweighted subgraph  $G' = (\mathcal{U}, E')$  that only contains edges exceeding a certain threshold  $t$  for the edge weight:

$$E' = \{\{e_1, e_2\} \in \mathcal{U} \mid \text{sim}(e_1, e_2) > t\} \quad (3)$$

$G'$  consists of connected components that divide  $\mathcal{U}$  into partitions of similar patches, that is, equivalence classes. We identify those equivalence classes as  $\sim_S$ :

$$[x]_S = \{y \in \mathcal{U} \mid x \sim_{G'} y\}, \quad (4)$$

where  $\sim_{G'}$  denotes reachability. Note that  $|[x]_S| > 0$ . Figure 3 illustrates the creation of clusters of similar patches. We distinguish between three different types of clusters:

#### 1. Unintegrated Patches: $[x]_S \subset \mathcal{M}$ .

All members of this equivalence class can only be found on mailing lists. Members are, for example, different revisions of the same patch.

No commit in repositories can be found in this category. This means that the patch has either not been integrated yet (as the discussion is, for example, still ongoing), or, that the patch has been rejected and is not object to integration.

#### 2. Integrated Patches: $\exists x_1, x_2 \in [x]_S : x_1 \in \mathcal{M} \wedge x_2 \in \mathcal{C}$ .

This category describes a finished integration process, as members of the equivalence class can be found on both: the mailing list

and the repository. Again,  $[x]_S \cap \mathcal{M}$  describes several revisions of the patch,  $[x]_S \cap \mathcal{C}$  denotes the assigned commits in the repository.

Note that several commits in the repository may be assigned: While backports also match to mainline commit, there may even been multiple mainline commits, as a patch can be picked up by multiple maintainers and appear as multiple commits.

#### 3. Off-list Patches: $[x]_S \subset \mathcal{C}$ .

No public development artefact can be assigned to the commit(s) in  $[x]_S$ . Besides false positives results of the heuristic, this category contains commits that arise from non-public integration channels.

With patches in group 3, we are able to identify commits that come from secret integration channels. Those commits will be subject of our analysis in Section 5.

## 4.2 Data Acquisition

The Linux Kernel community officially provides mailing list archives.<sup>3</sup> As archiving method, they use the public-inbox storage format.<sup>4</sup> The public-inbox approach stores mails in git repositories and provides a convenient data exchange format as standard tooling can be used to search for or to extract mails from the repository. Different mailing lists are stored in separate repositories.

However, a low amount of false positives because of misses requires full coverage of all mailing list data for the time frame of interest. While official resources reach back to early days of Linux, archives do not cover all mailing lists. Only a subset of  $\approx 100$  lists of over 200 referenced lists of the project are provided by the Linux Foundation.

Therefore, we subscribed to all 200 publicly available lists and collect mailing list data since May 2019. Our archives receive regular updates and are publicly available<sup>5</sup>. We use the open-source tool PaStA<sup>6</sup> for the construction of the commit hash  $\leftrightarrow$  message-id map.

## 5 ANALYSIS

In contrast to a *just-in-time* online analysis that constantly monitors new incoming mails and commits on a regular basis, we perform the detection of off-list patches as an ex-post analysis of a predefined time window. From a retrospective view, we can examine if a commit would have been detected as an off-list patch if a just-in-time online analysis would have been performed.

### 5.1 Overview

We are naturally limited by the availability of artefacts for the choice of time window for the analysis. For the time window of emails, we consider the date since creation of our data collection (2019-May-01) until we performed the analysis (2019-Dec-01).

Patches typically take weeks to months until they are integrated to the repository [18]. To select commits in the repository that are relevant for the analysis, we need to be aware that git, the version

<sup>3</sup> Available at <http://lore.kernel.org>

<sup>4</sup> see <https://public-inbox.org/README.html>

<sup>5</sup> Available at <https://github.com/orgs/linux-mailinglist-archives/>

<sup>6</sup> Available at <https://github.com/lfd/PaStA>

control system that is used by the Linux kernel, distinguishes between two temporal events: the author date and the commit date. The commit date is the date when the commit has been applied to the developer's (local) repository. Rewriting a repository's history can affect commit dates. The author date is the date when the commit was originally made (e.g., the date when the code was committed by the original author to their repository) or, in case of an email-based workflow, the timestamp when the email was sent (i.e., the Date: header of a mail). Hence, we integrate all commits with an author date within the same time window as chosen for emails. We respect all commits that meet the abovementioned criterion up to Linux version 5.4 (released 2019-Nov-24).

In that time window, we found 516,197 different messages,  $\approx$  40% of them contain actual patches. However, not all mails that contain patches are relevant for the analysis. Messages contain mails from bots, pull requests, backports and other noise. The tool PaStA filters those messages by applying appropriate heuristics.

The remaining messages are first compared against each other to find clusters of similar patches (i.e., several revisions of the same logical patch) and are then compared and mapped against 30,396 commits in the repository. Thresholds significantly influence the precision of the results. We chose thresholds in alignment with Ref. [37]. In the time window of our analysis, we mapped  $\approx$  96% of all commits against patches from mailing list and therefore regular development noise, while 1,240 commits were not assigned to any Message-Id.

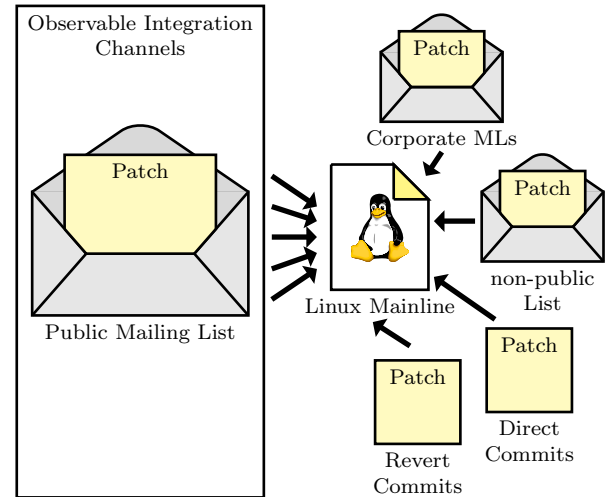
A commit with a missing mapping to a message can fall into one of the following categories: (1) The heuristic failed to detect the patch (false negative) (2) The original patch was sent to the list before we started recording mailing list data (miss of discussions) (3) *off-list patches* – patches that were integrated through a non-public channel.

## 5.2 Off-list Patches

With a manual investigation of the remaining 1,240 commits, we were able to find different categories for off-list patches in the Linux kernel repository. Figure 4 illustrates different types of off-list integration channels.

*Revert commits.* A revert commit is a commit that reverts a previous commit in the repository's history. They are used, for instance, to eliminate new features or enhancements if they cause undesired side effects or if they are in a defective or an incomplete state. It is often the preferred choice to revert the commit, as it is more efficient and less error prone to simply revert corresponding changes rather than to provide expensive or complex fixes, especially at the end of a development cycle. A refined version of the commit can later be integrated during the next development cycle.<sup>7</sup>

Many maintainers do not send reverting patches to mailing lists. They either integrate the reverting patch directly, or they send a response to the original thread of introducing patch that it will be reverted while omitting the actual reverting patch. Hence, we lack the reverting patches on mailing lists.



**Figure 4: Public observable (left) and non-public integration channels (right).**

Such reverting patches can automatically be detected, as the subject line of the commit message contains the keyword `Revert` by convention. For the time window of the analysis, we detected 64 off-list revert commits in the repository.

*Commits by Repository Owners.* Repository owners have a special role in projects: They have the permission to push code to official resources.

In case of Linux, Linus Torvalds is the owner of the official repository and the last approving authority. It is his final decision to judge if a patch or pull request is integrated mainline.

This, in turn, allows him for integrating or reverting patches ad libitum. It is not unusual that he reverts patches without discussion or short before the release of a new version.<sup>8</sup>

Torvalds sees himself as the manager of the Linux kernel – and no longer as active developer. Nevertheless, he sometimes integrates code or fixes without any prior public discussion. We can automatically detect those commits, as project owners are known. For the time window of the analysis, we detected 48 off-list patches from Linus Torvalds. None of them contained security-related fixes.

However, the phenomenon of bypassing public review processes can also be observed at other maintainers.

*Bypass of public review processes.* During our analysis, we found several regular patches that have never been sent to any public mailing list. To exclude false negatives of the heuristic, we contacted 18 different authors and collected affirmative answers from 14 authors – four did not answer.

All of them confirmed our finding that their patch(es) have never been posted on a public mailing list. For example, we found that one maintainer committed 40 patches to their repository in the time window of our analysis. The author confirmed our assumption and commented that they did not expect it to be that many.

<sup>7</sup> Example: Linux Commit 69bf4b6b54fb: [...] and it's [the bug] not immediately obvious why it happens. It's too late in the rc cycle to do anything but revert for now.

<sup>8</sup> Example: Release of Linux 5.3: Linux Commit 72dbcf7215

Most patches were only minor stylistic fixes, but we also found invasive patches. The author agreed that those patches would have required a public review process. We questioned maintainers why they skipped the official review process. Their typical answer was that they *accidentally forgot to send the patch*.

While many of those commits contain uncontroversial changes like documentation, style or typographical fixes, other commits contain in-depth fixes for subsystems. One maintainer explained that they picked up a fix from another subsystem that is also valuable for their area of responsibility. However, all responding maintainers agreed that those patches should have been publicly discussed.

*Established non-public integration channels.* Besides maintainers that directly commit patches without discussion, we also found subsystems that tend to bypass public review processes.

We have evidence from our observations that some subsystems deliberately bypass public review processes. For example, there are whole architectures and subsystems that are in the responsibility of certain companies.<sup>9</sup> A corporate representative has the role as an official mainline subsystem maintainer, which gives them the possibility to send pull requests to Linux—by trust. Within those subsystems, we can find off-list patches from authors other than the maintainer. Still, those patches can not be found on any public mailing lists. At the same, the author’s and maintainer’s email address show that both work for the same company.

From such artefacts in commits, we conclude the existence of non-public company internal review and integration processes. However, those patches do intentionally bypass public review process.

One maintainer confirmed our assumption and underlined that they *forgot to add the public list*, and that *normally all patches are discussed on the public mailing list before they land*.

*Security Vulnerabilities.* The remaining commits contain fixes for security vulnerabilities. According to Linux’s security process (explained in Section 3), patches for security vulnerabilities should be discussed on private non-public communication channels.

Typically, the majority of those patches drown in the noise of thousands of other commits. To prevent simple keyword-based search heuristics, commit messages are worded neutrally, links to CVE entries are only sometimes mentioned in the commit message [26].

To confirm our assumption that we hit security vulnerabilities through non-public integration channels, we contacted 12 authors. A list of the related and confirmed vulnerabilities can be found in Table 1. All of them confirmed that those patches are security related and that they have either been discussed on the non-public security mailing list, or been sent directly to the maintainer.

We calculated, in days, how long it takes for Debian 10 (Buster) and Ubuntu 18.04 (Bionic Beaver Hardware Enablement Kernel) to apply the patch to the distribution’s fork of the Linux kernel. Positive numbers denote a potential temporal advantage for an attacker, negative numbers mean that the distribution applied the patch before it was disclosed to public. The categories of vulnerabilities contain denial of service attacks, buffer overflows, privilege escalation, and buffer over-reads.

<sup>9</sup>We do not want publicly point to those subsystems.

In our analysis, we found, among others, fixes for the spectre-like attacks CVE-2019-11135 [39] and CVE-2019-1125 [2]. Ubuntu integrated both fixes before they were publicly disclosed, while Debian only integrated fixes for CVE-2019-11135 before they were publicly disclosed.

We also found patches for an *easy to exploit*<sup>10</sup> denial-of-service attack for ARM64-based Cavium systems. The vulnerability has no assigned CVE entry. It took almost two months for Ubuntu Bionic to integrate the patch. At the time of writing, Debian Buster, as well as the affected 4.19 Linux LTS tree, still lack appropriate fixes.

For the majority of vulnerabilities, our approach gives an attacker a temporal advantage from 2 to 179 days. While most patches for vulnerabilities are included on the stable Linux LTS trees, some distributions still lack patches for the corresponding vulnerabilities.

## 6 DISCUSSION

In this section, we first examine validity and potential weaknesses of our approach, and then discuss how our results affect OSS development processes. We conclude with suggestions how they can be adapted to accept (and deal with) risks that are anyway unavoidable, and concentrate on handling highly critical issues as good as possible.

### 6.1 Validity

*Analysis method.* Our work conducts an ex-post analysis. We consider a time window of seven months worth of mailing list data, and compare it against the corresponding time window in the repository. This allows us to judge from a *future* perspective if a patch *would have* been detected as an off-list patch at the time it was integrated into the repository.

Nevertheless, the retrospective position is only required to determine the practicability of the approach: It is straightforward to extend our methods to apply just-in-time, which is obviously necessary to abuse any undistributed security fixes. Periodic, frequent updates of the repository and mailing list data ascertain valid and current data, and are a mere technical detail. New incoming commits must be compared against the available mailing list data. If a patch is not an off-list patch, then the corresponding mailing list entry must be available at the moment of the analysis. As soon as a commit is pushed to a public available repository, our method allows to determine if the commit comes from a private channel.

In a private discussion, Greg Kroah-Hartman, maintainer, among others, of the stable and LTS trees of the Linux kernel, states the undocumented procedures how patches are distributed behind the scenes [27]. The exchange strategies vary depending on the involved maintainer(s) and the issue at hand: One possibility is to exchange patches via private email. Another method is to distribute patches as git bundles, a technique that allows for exchanging elements of a git repository without relying on public remote servers, while it still guarantees stable commit hashes to maintain unique patch identifiers. In a third method Linus Torvalds pulls patches from a maintainer tree. Since such trees are publicly available, this method opens a further temporal advantage for attackers,

<sup>10</sup>According to an assessment by the author of the fixes [45]

**Table 1: A list of vulnerabilities that were detected by our approach and confirmed by corresponding authors. A negative period means that patches were integrated by distributions before the vulnerabilities were disclosed. SecML means if the patch was routed through the security mailing list, or privately discussed with maintainers.**

CVE-2019	Description	Patches	SecML	Ubuntu 18.04	Debian 10
NA	DoS vulnerability for Cavium systems	4	no	59d	n/a
NA	smack: use after free	1	no	14d	n/a
13233	x86/insn-eval: use after free	1	yes	54d	61d
13272	pot. privilege escalation	1	yes	25d	12d
12817	ppc: inter-process memory leak	2	yes	-5d	44d
1125	x86/speculation: spectre v1 swaggs	4	yes	-5d	2d
14283	floppy: out-of-bounds read	2	yes	71d	18d
14283	floppy: DoS / div by zero	2	yes	71d	18d
11833	ext4: leak of sensitive data	1	yes	71d	29d
NA	s390: pot. leak of sensitive data	1	yes	45d	n/a
NA	apparmor: out of bounds by user-controlled data	1	no	179d	60d
11135	x86/tsx/speculation: TSX async abort side channel	9	yes	-1d	-1d

as a just-in-time analysis can also monitor patches from maintainers' repositories.

*Generalisability.* The primary concern of this paper is an in-depth analysis of patch flow into the Linux kernel repository from non-public resources by using peculiarities of its mail-based development process. However, our approach is neither limited to Linux as analysis target, nor to mailing lists as means of discussion. Variants of this process are used by many other system-level OSS projects, for instance GCC, QEMU, U-Boot, LLVM, busybox, and many others. Except for handling some technical details and taking minor process differences (e.g., the use of multiple parallel communication channels) into account, our approach can be directly applied to such systems, albeit we do not consider an according evaluation in this paper.

The exact reasons for the existence of non-public integration channels depend on the project. Especially in projects with smaller communities, maintainers often tend to directly commit code changes without public announcement or discussion (e.g., busybox), as up-front public discussion is often considered time-consuming and dispensable. However, this limitation is mitigated by the fact that projects with smaller communities only receive a moderate amount of patches. Especially critical system software typically demands adherence to public review processes, regardless of community size.

Our idea of development process reverse engineering is also applicable to processes that do not build upon mailing lists: If *any* publicly available development artefacts (e.g., pull requests, entries in issue trackers, ...) are available that include relevant data before their integration, then reverse process engineering uncovers any irregularities, in particular, deliberate violations of the development process.

*Scalability.* In a time window of roughly seven months, we found 30,396 relevant commits in the repository (authored after 2019-May-01 and integrated before Linux v5.4, released 2019-Nov-24). Within those commits, we found 1,240 potential off-list patches. By applying heuristics to exclude revert patches and commits by

project owners, we were able to exclude further 112 commits. With our approach, we filtered  $\approx 96\%$  of regular development noise.

Nevertheless, 1,128 commits required manual analysis, which may seem to imply a considerable impediment to a fully automatic system at first glance. However, commits span a time window of 207 days. On a daily basis, this accounts to manual investigation of (rounded up) six commits per day. Assuming, in accordance to our personal experience gathered, that an experienced developer can decide within a minute or two if a patch addresses a vulnerability, then the daily time investment would only require a reasonable amount [36] of around ten minutes.

Not enjoying the benefits of our system would require a fully manual inspection of all incoming commits, which is unrealistic: The official repository of the Linux kernel (merge commits are already excluded) received 70,632 commits between release v5.0 and v5.4. The development between those releases took 329 days. On average, 215 patches were integrated per day. Assuming the same amount of time required for manual investigation, an experienced developer would need more than three hours of concentrated reviewing per day. Hence, we argue that our approach is suitable for real-world scenarios, as it significantly reduces the amount of time that is required for manual review.

However, the time to find *some* security-related fixes could be reduced even to zero by employing simple heuristics, such as filtering for well-known author or institution names: For instance, out of the 12 fixes we identified, 3 originated from Jann Horn (GPZ). While this might have been pure coincidence, we argue that learning about the social structure behind Linux could be exploited in this respect.

*Internal Validity.* Our method systematically uncovers non-public integration channels and identifies commits that are potential fixes for security vulnerabilities. However, the method fails for vulnerabilities that are discussed in public before integration.

Statistical data on how many patches are sent to private security mailing lists, or how many critical vulnerabilities are discussed in public are not available. Hence, it is hard to calculate the accuracy of the approach since the recall is not available. Yet, we found 12



vulnerabilities in our analysis, which underlines the practical utility of our approach.

However, it is worth mentioning that counting or searching for CVE entries for a certain time window is neither an appropriate method of accounting the number of vulnerabilities in a system nor an alternative method to automatically find security vulnerabilities: Only a fraction of kernel security fixes get CVEs [13, 26] assigned. CVEs are also known to be abused as *integration shortcuts* [27],<sup>11</sup> and do on occasion not even address real vulnerabilities [12].

*Construct Validity.* We discussed our method with experts of the closed Linux security mailing list. They confirmed validity of our approach to gain information on non-public integration channels.

## 6.2 Consequences

*Fixes for Vulnerabilities.* The primary success criterion for our approach is simple: Can attackers gain temporal advance to design exploits? We argue that this is the case if the patch can be found in public resources before software distributors roll out patches: Reverse engineering of the development process allows for aimed targeting of commits that would otherwise *hide* between thousands of other commits.

As mentioned in Section 3, the majority of patches for vulnerabilities first appear in the Linux mainline and stable trees before distributions pick up the relevant patches. From a temporal perspective, patches first appear on mainline and stable trees, and are then integrated by distributions (cf. 2, Vulnerability 1). We call this the *mainline first* disclosure model.

However, there is an exception for highly critical vulnerabilities: Before their public disclosure, patches are secretly disclosed to the kernel maintainers of the distributions, which buys them time to prepare their kernel tree to roll out updates (cf. Figure 2, Vulnerability 2) as soon as an embargo ends. In this way, a patch can be integrated to the distribution’s tree before it is published mainline.

This method ensures that affected systems can receive fixes as soon as the vulnerability is officially disclosed. Yet, this process requires time-consuming and extensive coordination between maintainers of distributions and the kernel community, since a strict temporal publishing coordination is required to make the approach effective. Coordination efforts are even more complex when hardware bugs (such as bugs in speculative execution [23, 32]) are involved, as multiple operating systems can be affected. This additionally requires cross-community coordination—between different operating systems (variants of BSD, Windows, macOS), commercial and non-commercial vendors, and, under exceptional circumstances [23], even with compiler manufacturers. This process is therefore only considered in rare cases.

We call this process the *distro first* disclosure model, as patches are integrated by distributions before they are officially published mainline.

According to Kroah-Hartman [27], there is no clear definition of the disclosure process, and no definitive criteria for circumstances when the *distro first* model should be used. As an ad-hoc process,

<sup>11</sup>For instances, processes of commercial companies that must be passed before contributions can be placed in open source projects can contain shortcuts for critical vulnerabilities, and “critical” is equated with “has CVE assigned”.

subsystem maintainers decide how to handle a fix: patches can, for example, be routed through maintainer trees to Linus Torvalds, or Linus merges the patch directly, depending on the area of the kernel that was involved.

To give an example, fixes for flaws in the speculative execution model (cf. CVE-2019-11135 [39] and CVE-2019-1125 [2]) of modern CPUs were entirely developed and rolled out to distributions in private. Our approach can still detect that the patches stem from off-list channels as soon as they are available in a repository – but at that point in time, patched binaries are already available for the public. Nevertheless, our method can still provide some valuable temporal advance as the availability of patches does not imply immediate deployment in the field.

However, the *mainline first* disclosure model is used for the majority of fixes for vulnerabilities. As distributions maintain forks of the Linux kernel, and manually select patches that are integrated from mainline, it can take up to months for patches to be integrated (cf. Table 1). In particular, selecting patches for local forks on a case-by-case basis misses relevant fixes that are available on LTS.

For these cases, the integration process of distribution kernels can be considered as *security by obscurity*, since (a) the patches do not follow a coordinated disclosure process to distributions to protect affected systems before their official publication, and (b) the existence of the actual fixes is obfuscated by private discussion and regular development noise.

We hence argue that release strategies of distributions should be reconsidered, as we have demonstrated that distributions are vulnerable for attacks over long periods of time.

Furthermore, we argue that fixes for vulnerabilities should be publicly discussed *after* their disclosure. While preliminary versions for severe vulnerabilities that require *distro first* integration should be developed under the *distro first* model, we recommend using a full disclosure model in all other cases. Early versions of fixes for vulnerabilities can still be discussed on secret lists, but they should be publicly reviewed after their embargo.

A public review process can enhance the software quality of the fix per se—after all, this is the main concern of public discussion—, but can also avoid the the inadvertent introduction of *additional* vulnerabilities by fixing one vulnerability, which is unfortunately a real pattern [12]. Public discussion before integration would also defeat our mechanisms, which is eventually desirable.

*Code Infiltration.* In addition to detecting fixes for vulnerabilities, we also encountered hidden integration channels *besides* security mailing lists, such as maintainers or companies that—systematically or inadvertently—bypass official submission procedures, for instance by direct maintainer commits without external review, or company-internal review. The existence of such channels, shows that trusted individuals can easily infiltrate the project, and secretly introduce malicious artefacts (while this possibility is given, our method allows for finding concrete instances, which is otherwise not possible). The existence of such commits contradicts one of the key promises of an open development model.

We contacted maintainers for subsystems for which we found such patches, and they confirmed our assumption that they integrated code without prior public review. While maintainers are aware of that they sometimes intentionally bypass the process,

they were surprised of the magnitude of unreviewed patches—the confirmed “record” is more than 40 per half-year per author, the estimated number for unconfirmed cases is higher.

## 7 RELATED WORK

Software vulnerability life cycle analysis is related to this area work, and a well-researched topic [5, 16, 40].

Huang et al. [16] find a considerable delay between disclosure of vulnerabilities and the availability of fixes. Based on a case study of six different projects, they found an average time of 52 days from vulnerability disclosure to releasing an actual fix. However, they also find that almost half of the vulnerabilities are fixed within one week.

In 2010, Arora et al. [5] argue that instant disclosure of a vulnerability forces vendors to speed up the release of a fix by 35 days.

Shahzad et al. [40] analyse the life cycle of vulnerabilities that are filed in software vulnerability data sets. In their large-scale analysis that includes a big variety of different projects, they find that the amount of time required to fix vulnerabilities decreased from 1998 to 2011: Since 2008, 80% of all vulnerabilities are fixed by vendors before their disclosure. Yet, their study does not consider that providing a patches is only a first step, but necessitates integration in software distributions, and actual deployment by users.

In a large-scale empirical study, Li and Paxson [30] investigate bug-fixes for security vulnerabilities in open-source projects. For their comprehensive analysis, they assign 3,094 CVE entries in the National Vulnerability Database (NVD) to 4,080 commits in 682 unique git repositories. Mining for links to commits in the CVE description establishes the approximate connection between CVE entry and commit hash. Later, they extract characteristics of security-related commits. They find that security fixes are less complex and more localised than non-security fixes. Furthermore, they find that 70% of security-related patches were committed before public disclosure and conclude that development and deployment processes provide a window of opportunity for exploitation. However, for a responsible disclosure process, it is necessary that patches must be developed (and committed) before disclosure. Yet, the date of a commit is not necessarily the date of its public visibility. In this work, we showed that developers intentionally distribute and release patches on secret channels before they finally publicly publish the repositories. Attackers do not have the opportunity for prior exploitation in those cases. In this work, we respect this fact and use the time difference of the public availability of a binary software release and the date of the public disclosure as the basis for our analysis.

Kroah-Hartman argues that only a small fraction of Linux kernel security fixes are assigned to CVE entries [26]. From 2006-2018, 1005 CVEs were assigned to the kernel. He argues that, on average, bugs with CVE entries are 100 days fixed in mainline before they get a CVE assigned. Furthermore, he argues that the amount of vulnerabilities of vendor distributions can significantly be reduced by choosing LTS versions of Linux.

Insider attacks, such as infiltration, or compromises of organisational structures, are well-known in literature [8, 21]. We showed a practical *outsider attack* that exploits the openness of the development model itself by using its development artefacts to conclude to

systematic integration of patches that lack public discussion. In [4], Anderson argues that the security of a development model should *not* depend whether it is open or closed.

The software engineering community uses artefact mining techniques to draw quantitative conclusions on development processes [20] or to determine various software performance indicators [10, 15].

## 8 CONCLUSION

We showed that reverse engineering of public development processes allows to detect code that arises from non-public integration channels. Our approach removes 96% of regular development noise and points to hot spots that contain fixes for critical security vulnerabilities. With our method, we were able to detect 12 vulnerabilities in Linux in a time window of seven months. We collected responses from all authors that confirm our presumptions. Attackers can use this information to gain temporal advantage, as they can design exploits before affected systems receive patches.

Furthermore, we found evidence that some subsystems and maintainers intentionally bypass the regular development process. Therefore we argue that it is possible to systematically infiltrate malicious code to the kernel by bypassing the (mandatory) public review processes. We shared our findings with the Linux kernel community and discussed possibilities of potential mitigations.

Our future work will focus on just-in-time online analyses and automated process monitoring: Automatic notifications to maintainers or authors can help to raise the awareness of the importance of public code review processes.

## 9 ACKNOWLEDGEMENTS

We thank Greg Kroah-Hartman for giving us the opportunity to discuss the topic with Linux kernel security officers. We also thank authors of off-list patches for their detailed answers, discussions on their patches and fruitful conversation. We do not mention them by name.

We also thank Christian Weber and Florian Loher for their assistance with creating figures and illustrations.

This work was supported by the iDev40 project and the German Research Council (DFG) under grant no. LO 1719/3-1. The iDev40 project has received funding from the ECSEL Joint Undertaking (JU) under grant no. 783163. The JU receives support from the European Union’s Horizon 2020 research and innovation programme. It is co-funded by the consortium members, grants from Austria, Germany, Belgium, Italy, Spain and Romania.

## REFERENCES

- [1] 2017. CVE-2018-3615, CVE-2018-3620, and CVE-2018-3646. <https://cve.mitre.org>.
- [2] 2019. CVE-2019-1125. <https://cve.mitre.org>.
- [3] Doug Anderson. 2019. Allowing something Change-Id (or something like it) in kernel commits. <https://lists.linuxfoundation.org/pipermail/ksummit-discuss/2019-August/006739.html>.
- [4] Ross Anderson. 2002. Security in Open versus Closed Systems—The dance of Boltzmann, Coase and Moore. In *Open Source Software: Economics, Law and Policy*.
- [5] Ashish Arora, Ramayya Krishnan, Rahul Telang, and Yubao Yang. 2010. An Empirical Analysis of Software Vendors’ Patch Release Behavior: Impact of Vulnerability Disclosure. *Information Systems Research* 21 (03 2010).

- [6] Nicolas Bettenburg, Ahmed E Hassan, Bram Adams, and Daniel M German. 2015. Management of community contributions. *Empirical Software Engineering* 20, 1 (2015).
- [7] Christian Bird, Peter C Rigby, Earl T Barr, David J Hamilton, Daniel M German, and Prem Devanbu. 2009. The Promises and Perils of Mining Git. In *6th IEEE International Working Conference on Mining Software Repositories (MSR)*.
- [8] Matt Bishop, Heather M Conboy, Huong Phan, Borislava I Simidchieva, George S Avrunin, Lori A Clarke, Leon J Osterweil, and Sean Peisert. 2014. Insider threat identification by process analysis. In *IEEE Security and Privacy Workshops*.
- [9] Hasan Cavusoglu, Huseyin Cavusoglu, and Srinivasan Raghunathan. 2004. Emerging Issues in Responsible Vulnerability Disclosure. In *Workshop on Information Technology and Systems (WITS)*.
- [10] Alexander Chávez, Isabella Ferreira, Eduardo Fernandes, Diego Cedrim, and Alessandro Garcia. 2017. How Does Refactoring Affect Internal Quality Attributes? A Multi-Project Study. In *Proceedings of the 31st Brazilian Symposium on Software Engineering (SBES&AZ17)*.
- [11] Steve Christey and Chris Wysopal. 2002. Responsible vulnerability disclosure process. <https://tools.ietf.org/html/draft-christey-wysopal-vuln-disclosure-00>. *IETF draft* (2002).
- [12] Jonathan Corbet. 2019. What to do about CVE numbers. *Linux Weekly News (LWN)* (2019). <https://lwn.net/Articles/801157/>
- [13] Jake Edge. 2019. CVE-less vulnerabilities. *Linux Weekly News (LWN)* (2019). <https://lwn.net/Articles/791855/>
- [14] Stefan Frei, Dominik Schatzmann, Bernhard Plattner, and Brian Trammell. 2010. Modeling the security ecosystem—the dynamics of (in) security. In *Economics of Information Security and Privacy*.
- [15] Hadi Hemmati, Sarah Nadi, Olga Baysal, Oleksii Kononenko, Wei Wang, Reid Holmes, and Michael W Godfrey. 2013. The MSR Cookbook: Mining a decade of research. In *10th Working Conference on Mining Software Repositories (MSR)*.
- [16] Zhen Huang, Mariana DAngelo, Dhaval Miyani, and David Lie. 2016. Talos: Neutralizing Vulnerabilities with Security Workarounds for Rapid Response. In *IEEE Symposium on Security and Privacy (SP)*.
- [17] Intel Corporation. 2018. Resources and Response to Side Channel L1 Terminal Fault. <https://www.intel.com/content/www/us/en/architecture-and-technology/l1tf.html>.
- [18] Yujuan Jiang, Bram Adams, and Daniel M German. 2013. Will my patch make it? And how fast? Case study on the Linux kernel. In *10th Working Conference on Mining Software Repositories (MSR)*.
- [19] Yujuan Jiang, Bram Adams, Foutse Khomh, and Daniel M. German. 2014. Tracing Back the History of Commits in Low-Tech Reviewing Environments: A Case Study of the Linux Kernel. In *Proceedings of the 8th ACM/IEEE International Symposium on Empirical Software Engineering and Measurement (ESEM)*.
- [20] Mitchell Joblin, Sven Apel, Claus Hunsen, and Wolfgang Mauerer. 2017. Classifying developers into core and peripheral: An empirical study on count and network metrics. In *IEEE/ACM 39th International Conference on Software Engineering (ICSE)*.
- [21] Florian Kammüller and Christian W Probst. 2013. Invalidating policies using structural information. In *IEEE Security and Privacy Workshops*.
- [22] Shuah Khan. 2019. [MAINTAINERS SUMMIT] Patch version changes in commit logs? <https://lists.linuxfoundation.org/pipermail/ksummit-discuss/2019-June/006543.html>.
- [23] Paul Kocher, Jann Horn, Anders Fogh, Daniel Genkin, Daniel Gruss, Werner Haas, Mike Hamburg, Moritz Lipp, Stefan Mangard, Thomas Prescher, Michael Schwarz, and Yuval Yarom. 2019. Spectre Attacks: Exploiting Speculative Execution. In *40th IEEE Symposium on Security and Privacy (S&P)*.
- [24] Greg Kroah-Hartman. 2007. Linux kernel development. In *Proceedings of the Linux Symposium*.
- [25] Greg Kroah-Hartman. 2016. Why kernel development still uses email. *Linux Weekly News (LWN)* (2016). <https://lwn.net/Articles/702177/>
- [26] Greg Kroah-Hartman. 2019. MDS, Fallout, Zombieland & Linux. In *Embedded Linux Conference Europe (ELCE)*.
- [27] Greg Kroah-Hartman. 2020. personal communication.
- [28] Filippo Lanubile, Christof Ebert, Rafael Prikladnicki, and Aurora Vizcaino. 2010. Collaboration tools for global software engineering. *IEEE Software* 27, 2 (2010).
- [29] Vladimir I Levenshtein. 1966. Binary codes capable of correcting deletions, insertions, and reversals. In *Soviet physics doklady*.
- [30] Frank Li and Vern Paxson. 2017. A large-scale empirical study of security patches. In *Proceedings of the ACM SIGSAC Conference on Computer and Communications Security*.
- [31] Linux Kernel Community. 2020. Linux – How the development process works. <https://www.kernel.org/doc/html/latest/process/2.Process.html>.
- [32] Moritz Lipp, Michael Schwarz, Daniel Gruss, Thomas Prescher, Werner Haas, Anders Fogh, Jann Horn, Stefan Mangard, Paul Kocher, Daniel Genkin, Yuval Yarom, and Mike Hamburg. 2018. Meltdown: Reading Kernel Memory from User Space. In *27th USENIX Security Symposium*.
- [33] Wolfgang Mauerer. 2010. *Professional Linux kernel architecture*. John Wiley & Sons.
- [34] Wolfgang Mauerer and Michael C Jaeger. 2013. Open source engineering processes. *it-Information Technology* 55, 5 (2013).
- [35] Andrew Meneely and Laurie Williams. 2009. Secure open source collaboration: an empirical study of Linus' law. In *Proceedings of the 16th ACM conference on Computer and communications security*.
- [36] Emerson Murphy-Hill, Edward K. Smith, Caitlin Sadowski, Ciera Jaspán, Collin Winter, Matthew Jorde, Andrea Knight, Andrew Trenk, and Steve Gross. 2019. Do Developers Discover New Tools on the Toilet?. In *Proceedings of the 41st International Conference on Software Engineering (ICSE)*.
- [37] Ralf Ramsauer, Daniel Lohmann, and Wolfgang Mauerer. 2019. The List is the Process: Reliable Pre-Integration Tracking of Commits on Mailing Lists. In *IEEE/ACM 41st International Conference on Software Engineering (ICSE)*.
- [38] Eric Raymond. 1999. The cathedral and the bazaar. *Knowledge, Technology & Policy* 12, 3 (1999).
- [39] Michael Schwarz, Moritz Lipp, Daniel Moghimi, Jo Van Bulck, Julian Stecklina, Thomas Prescher, and Daniel Gruss. 2019. ZombieLoad: Cross-Privilege-Boundary Data Sampling. In *Proceedings of the ACM SIGSAC Conference on Computer and Communications Security (CCS)*.
- [40] Muhammad Shahzad, Muhammad Zubair Shafiq, and Alex X Liu. 2012. A large scale exploratory analysis of software vulnerability life cycles. In *34th International Conference on Software Engineering (ICSE)*.
- [41] The Kernel Community. 2020. Submitting patches: the essential guide to getting your code into the kernel. <https://www.kernel.org/doc/html/latest/process/submitting-patches.html>.
- [42] Jo Van Bulck, Marina Minkin, Ofir Weisse, Daniel Genkin, Baris Kasikci, Frank Piessens, Mark Silberstein, Thomas F. Wenisch, Yuval Yarom, and Raoul Strackx. 2018. Foreshadow: Extracting the Keys to the Intel SGX Kingdom with Transient Out-of-Order Execution. In *Proceedings of the 27th USENIX Security Symposium*. USENIX Association. See also technical report Foreshadow-NG [43].
- [43] Ofir Weisse, Jo Van Bulck, Marina Minkin, Daniel Genkin, Baris Kasikci, Frank Piessens, Mark Silberstein, Raoul Strackx, Thomas F. Wenisch, and Yuval Yarom. 2018. Foreshadow-NG: Breaking the Virtual Memory Abstraction with Transient Out-of-Order Execution. *Technical report* (2018). See also USENIX Security paper Foreshadow [42].
- [44] Minghui Zhou, Qingying Chen, Audris Mockus, and Fengguang Wu. 2017. On the Scalability of Linux Kernel Maintainers' Work. In *Proceedings of the 11th Joint Meeting on Foundations of Software Engineering*.
- [45] Marc Zyngier. 2019. personal communication.