

# The Unexplored Terrain of Compiler Warnings

Gunnar Kudrjavets  
University of Groningen  
Groningen, Netherlands  
g.kudrjavets@rug.nl

Nachiappan Nagappan  
Microsoft Research\*  
Redmond, WA, USA  
nnagappan@acm.org

Aditya Kumar  
Snap, Inc.  
Santa Monica, CA, USA  
adityak@snap.com

Ayushi Rastogi  
University of Groningen  
Groningen, Netherlands  
a.rastogi@rug.nl

## ABSTRACT

The authors' industry experiences suggest that compiler warnings, a lightweight version of program analysis, are valuable early bug detection tools. Significant costs are associated with patches and security bulletins for issues that could have been avoided if compiler warnings were addressed. Yet, the industry's attitude towards compiler warnings is mixed. Practices range from silencing all compiler warnings to having a zero-tolerance policy as to any warnings. Current published data indicates that addressing compiler warnings early is beneficial. However, support for this value theory stems from grey literature or is anecdotal. Additional focused research is needed to truly assess the cost-benefit of addressing warnings.

## CCS CONCEPTS

• **Software and its engineering** → *Software defect analysis; Software design trade-offs; Empirical software validation.*

## KEYWORDS

Defect prevention, compiler warning, Clang, GCC, MSVC

### ACM Reference Format:

Gunnar Kudrjavets, Aditya Kumar, Nachiappan Nagappan, and Ayushi Rastogi. 2022. The Unexplored Terrain of Compiler Warnings. In *44th International Conference on Software Engineering: Software Engineering in Practice (ICSE-SEIP '22)*, May 21–29, 2022, Pittsburgh, PA, USA. ACM, New York, NY, USA, 2 pages. <https://doi.org/10.1145/3510457.3513057>

## 1 OPPORTUNITY AND MOTIVATION

One of the earliest stages of software development during which bugs can be detected is when new code changes are compiled. A compiler can flag *potential* issues found in the code. The cost of fixing a software defect increases significantly during the later phases of the development cycle [14]. It is optimal to correct problems as early as possible because even minor bugs can result in catastrophic

consequences [17]. Classical memory safety related bugs, common to programming languages such as C and C++ (e.g., a double-free), are the reasons behind approximately 70% of security updates Microsoft issues each year [9]. For Microsoft, the cost of fixing a bug resulting in a security bulletin is approximately \$100,000 [5, p. 11]. Modern compilers (e.g., Clang, GCC, and MSVC) can detect typical programming mistakes such as integer overflows or underflows, out-of-bounds errors, memory management problems, etc., either during the compilation phase or via using runtime sanitizers. Acting on compiler warnings enables engineers to fix defects early and prevent the cascading set of failures the bugs would have otherwise caused. Despite the potential benefits of heeding compiler warnings, industry attitudes are mixed. For Linux kernel development, it took thirty years to start treating warnings as errors [15]. Google takes a somewhat contrarian approach by aiming to never issue compiler warnings because they find that developers ignore them [16, p. 427]. The warnings are either enabled as errors or never shown in the compiler output.

In our experience, demonstrating the value of fixing compiler warnings or changing organizational culture to treat compiler warnings as a first-class defect prevention tool is challenging. Often, changes in attitude and engineering processes only take place after damaging events (e.g., critical services becoming inaccessible, irrecoverable data loss, zero-day exploits) have already manifested.

## 2 EXISTING EVIDENCE AND GUIDANCE

Existing research and empirical data about the benefits of fixing compiler warnings is minimal. Most of the data comes from grey literature related to writing secure code or anecdotal knowledge passed down from experienced practitioners.

Microsoft practices recommend using the highest level of warnings to inspect code for potential security vulnerabilities and compile “cleanly” without any errors or warnings [4, 7]. The downside of not fixing the warnings is articulated in a case study on a large code base where integer-related warnings had been disabled. Analysis reveals that about 20% of the hidden warnings contain potentially exploitable conditions [6]. A post-mortem analysis from Facebook finds that enabling all compiler warnings as errors reveals issues such as memory leaks, infinite recursion, and catastrophic bugs where a compiler would “optimize” away critical functions [8]. Reducing the attack surface and finding opportunities to clean up code during the maintenance phase is another suggested application for utilizing compiler warnings [11].

\*This work was initiated while Nachiappan Nagappan was with Microsoft Research. He is currently with Facebook, Inc.



This work is licensed under a Creative Commons Attribution International 4.0 License.

ICSE-SEIP '22, May 21–29, 2022, Pittsburgh, PA, USA

© 2022 Copyright held by the owner/author(s).

ACM ISBN 978-1-4503-9226-6/22/05.

<https://doi.org/10.1145/3510457.3513057>

We can find only one paper investigating the correlation between compiler warnings and defects [10]. The study finds experimental evidence that “[a] large number of compiler warnings of a source file is an indicator that the file contains also an above-average number of defects.” The conclusion is based on a limited amount of data and uses a version of GCC from 2006. Given the advances in compiler technology during the last 16 years and the size of industrial code bases (e.g., in 2017 the Windows code base contained 3.5 million files) we need additional studies utilizing the latest versions of compiler toolsets and larger projects [2]. The remaining discoverable research related to compiler warnings is focused on their correctness, readability, and validity [1, 13].

### 3 OBSERVATIONS FROM INDUSTRY

The trend we observe is that *an engineer’s experience and seniority are directly related to his or her attitude towards fixing warnings*. The more experience with the cost and consequences of basic programming errors the engineers have, the more appreciative they are of ensuring the correctness of the code as early as possible.

From a technical point of view, *we rarely observe projects treating warnings as errors and triggering build breaks as a result*. Turning on *all possible warnings* is mainly done by engineers developing compilers themselves. Very few projects in industry have a zero-tolerance policy towards the presence of compiler warnings. A rare example is safety-critical code, e.g., software developed by NASA [3]. We have not been able to find any public data regarding standards related to compiler warnings in other companies producing safety-critical software, e.g., Airbus, Boeing, Tesla, etc.

A variety of reasons contribute to compiler warnings either not being fixed or deprioritized. The main reason is the *lack of empirical evidence* to show either correlation or causal relationship between fixing compiler warnings and decrease in defect density. Another key reason is the *cost of adapting stricter compiler warning levels to legacy code*. Techniques such as treating warnings as errors are time-consuming to implement unless projects established this policy from the very beginning. We cannot discount the *impact on an engineer’s career* as well. The lack of external motivation to fix the compiler warnings is often caused by the fact that preemptively fixing compiler warnings does not get rewarded as well as the post hoc activity associated with debugging and bug fixing. The *repetitive nature of fixing the compiler warnings* is another factor making long-term code quality improvement initiatives unpopular. The number of warnings to be analyzed may reach into hundreds, thousands or even tens of thousands depending on the size of the code base. A key reason related to engineers not willing to fix compiler warnings is *distrust in the validity of the warnings* due to past experiences with false positives. This belief can be summarized as “if warnings would indicate *real problems*, then they would be errors instead”.

### 4 FUTURE RESEARCH DIRECTIONS

We recommend that researchers partner with practitioners working on open- and closed-source software to focus on following topics:

- (1) *Explore the current state*. What are the default warning levels, attitudes and sets of beliefs toward fixing the warnings? Are they influenced by software’s technical abstraction level?

- (2) *Investigate the relationship (or lack thereof) between compiler warnings and defects, team productivity, and product risk*.
- (3) *Establish baseline metrics related to compiler warnings*. For example, warnings per file, per KLOC, change in the ratio of warnings with the application of stricter levels of compilation, number of suppressed warnings per KLOC?
- (4) *Rank warning categories according to their precision and recall*. Propose a recommended set of warnings per compiler.
- (5) *Conduct case studies about projects having zero-tolerance policy towards warnings*. Is the approach cost-effective outside the scope of safety-critical software?
- (6) *Evaluate the economics (e.g., negative impact) of fixing warnings*. SQLite development team finds that “[m]ore bugs have been introduced into SQLite while trying to get it to compile without warnings than have been found by static analysis” [12].
- (7) *Variation between programming languages*. How similar are or should be warnings for low-level (e.g., C), functional (e.g., OCaml), or scripting (e.g., Ruby) languages?

### REFERENCES

- [1] Titus Barik, Denae Ford, Emerson Murphy-Hill, and Chris Parnin. 2018. How Should Compilers Explain Problems to Developers?. In *Proceedings of the 2018 26th ACM Joint Meeting on European Software Engineering Conference and Symposium on the Foundations of Software Engineering* (Lake Buena Vista, FL, USA) (ESEC/FSE 2018). Association for Computing Machinery, New York, NY, USA, 633–643. <https://doi.org/10.1145/3236024.3236040>
- [2] Brian Harry. 2017. The Largest Git Repo on the Planet. Retrieved January 10, 2022 from <https://devblogs.microsoft.com/bharry/the-largest-git-repo-on-the-planet/>
- [3] Gerard J. Holzmann. 2006. The Power of 10: Rules for Developing Safety-Critical Code. *Computer* 39, 6 (June 2006), 95–97. <https://doi.org/10.1109/MC.2006.212>
- [4] Michael Howard. 2006. A Process for Performing Security Code Reviews. *IEEE Security & Privacy Magazine* 4, 4 (July 2006), 74–79. <https://doi.org/10.1109/msp.2006.84>
- [5] Michael Howard and David Leblanc. 2002. *Writing Secure Code* (2nd ed.). Microsoft Press, USA.
- [6] Michael Howard and David Leblanc. 2007. *Writing Secure Code for Windows Vista(TM) (Pro - Step By Step Developer)*. Microsoft Press, USA.
- [7] Michael Howard and Steve Lipner. 2006. *The Security Development Lifecycle*. Microsoft Press, USA.
- [8] Aditya Kumar. 2019. *-Wall Found Compilation Errors*. Facebook. Retrieved January 10, 2022 from <https://llvm.org/devmtg/2019-10/slides/Kumar-WallFoundCompilationErrors.pdf> Presented at the LLVM Developers’ Meeting.
- [9] Miller Matt. 2019. Trends and Challenges in the Vulnerability Mitigation Landscape. USENIX Association, Santa Clara, CA. 13th USENIX Workshop on Offensive Technologies.
- [10] Raimund Moser, Barbara Russo, and Giancarlo Succi. 2007. Empirical Analysis on the Correlation Between GCC Compiler Warnings and Revision Numbers of Source Files in Five Industrial Software Projects. *Empirical Software Engineering* 12, 3 (May 2007), 295–310. <https://doi.org/10.1007/s10664-006-9029-x>
- [11] Thomas Pearse and Paul Oman. 1995. Maintainability Measurements on Industrial Source Code Maintenance Activities. In *Proceedings of International Conference on Software Maintenance*. IEEE Comput. Soc. Press, Washington, DC, USA, 295–303. <https://doi.org/10.1109/ICSM.1995.526551>
- [12] SQLite. 2020. *How SQLite Is Tested*. SQLite Consortium. Retrieved January 10, 2022 from <https://www.sqlite.org/testing.html#staticanalysis>
- [13] Chengnian Sun, Vu Le, and Zhendong Su. 2016. Finding and Analyzing Compiler Warning Defects. In *Proceedings of the 38th International Conference on Software Engineering* (Austin, Texas) (ICSE ’16). Association for Computing Machinery, New York, NY, USA, 203–213. <https://doi.org/10.1145/2884781.2884879>
- [14] Gregory Tassej. 2002. *The Economic Impacts of Inadequate Infrastructure for Software Testing*. Diane Publishing Company.
- [15] Linus Torvalds. 2021. *Enable ‘-Werror’ by default for all kernel builds*. Linux. Retrieved January 10, 2022 from <https://git.kernel.org/pub/scm/linux/kernel/git/torvalds/linux.git/commit/?id=3fe617ccaf6f5bb33c2391d6f4eeb41c1fd0151>
- [16] Titus Winters, Tom Manshreck, and Hyrum Wright. 2020. *Software Engineering at Google: Lessons Learned from Programming Over Time*. O’Reilly, Boston.
- [17] Michael Zhivich and Robert K. Cunningham. 2009. The Real Cost of Software Errors. *IEEE Security Privacy* 7, 2 (2009), 87–90. <https://doi.org/10.1109/MSP.2009.56>