BLACKDUCK®

WHITE PAPER

2024 Open Source Risk in M&A by the Numbers

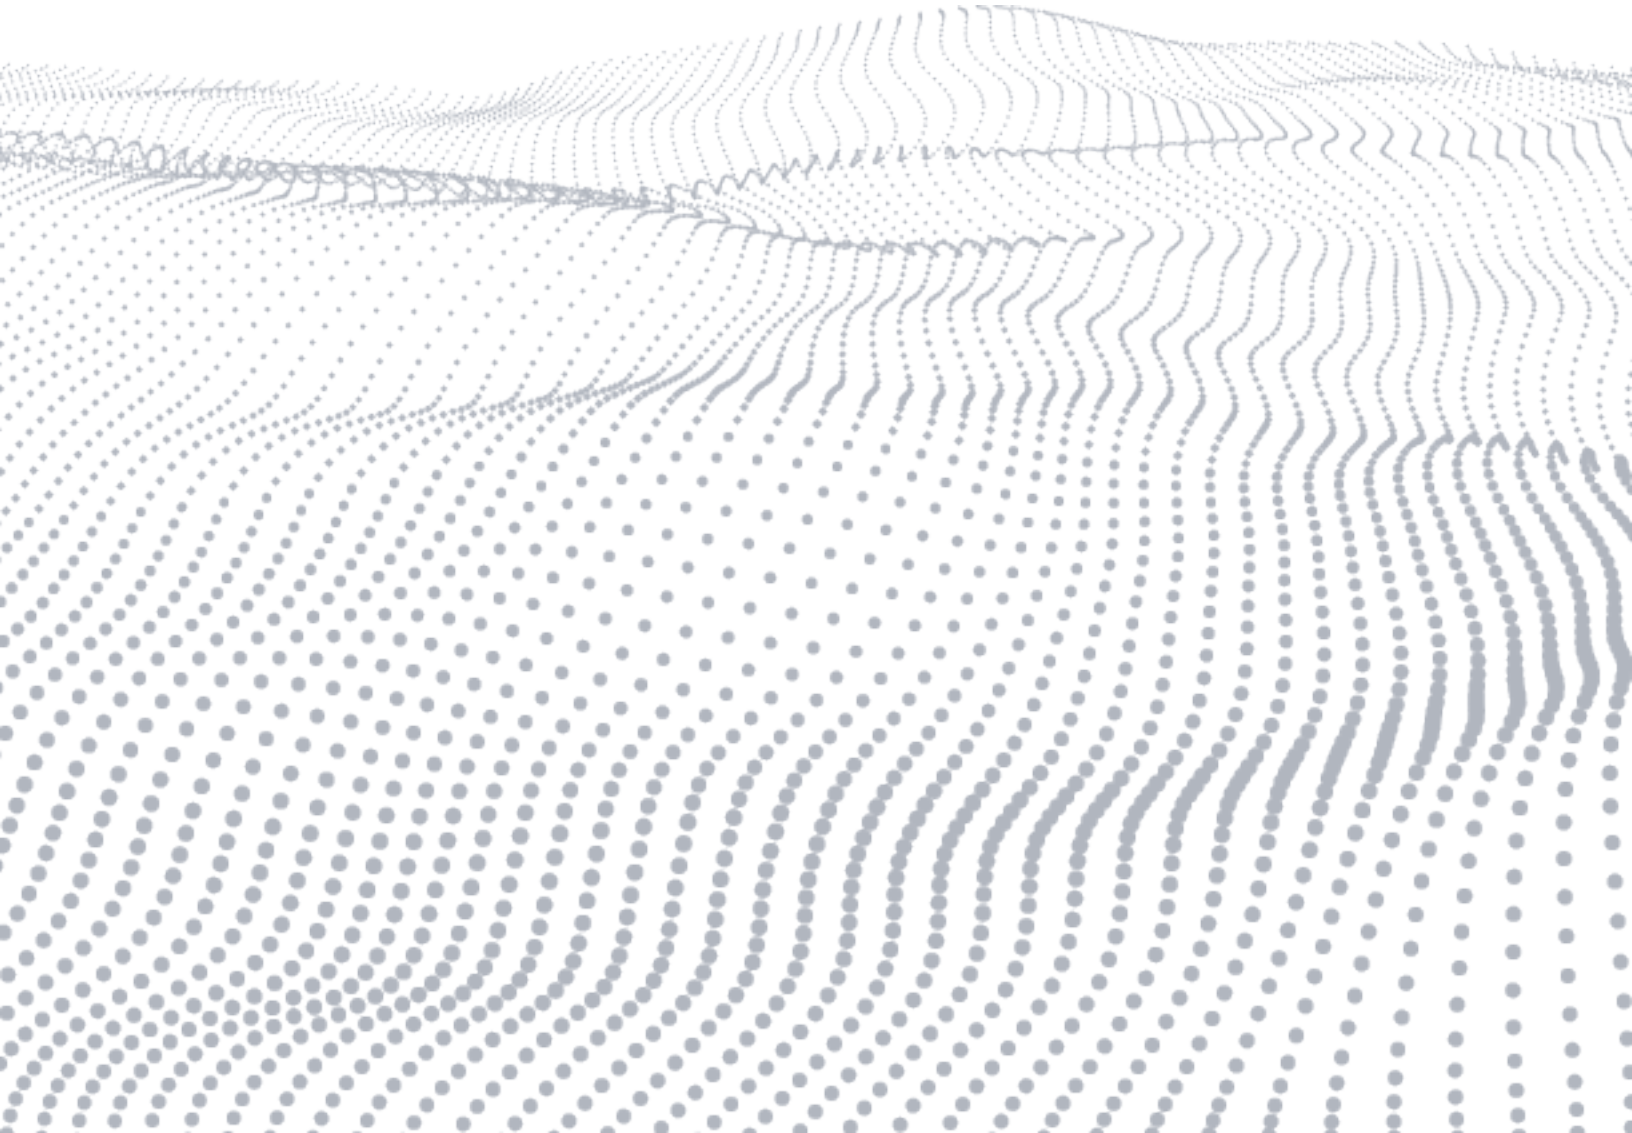

Table of contents

The numbers at a glance

Best practices for firms involved in merger and acquisition (M&A) transactions include audits of the target's code whenever software is a significant part of the value of a deal. Expert third parties that perform multiple types of analyses on the code can help purchasers better understand the technology and capabilities they're buying and identify potential legal, security, and quality issues.

In 2023, our audits found open source in 99% of customer engagements, with an average of 1,635 components discovered per engagement.

Today, open source components and libraries form the backbone of nearly every application in every industry. The reasons are straightforward: Using open source speeds development, drives innovation, and lowers costs, all critical in today's agile software world. The Black Duck® Audit Services team conducts open source audits—as well as analysis of software security and quality—on thousands of codebases (the code and associated libraries that make up an application or service) for its customers each year. Our audits are primarily done in support of M&A transactions, to provide customers with comprehensive, up-to-date Software Bills of Materials of the open source, third-party code, web services, and APIs used in their applications and to enable a view into risks associated with the components.

The 2024 "Open Source Security and Risk Analysis" (OSSRA) report presents analysis from an examination of the anonymized data from over 1,000 commercial codebases audited in 2023. Industries represented in the report include automotive, big data, cybersecurity, enterprise software, financial services, healthcare, the Internet of Things, manufacturing, and mobile apps. The average Black Duck audit engagement in 2023 comprised audits of an average of five codebases.

The OSSRA report presents data on codebases—the code and associated libraries that make up an application. The average transaction that involved a Black Duck audit in 2023 included an average of five codebases. As this paper focuses on M&A transactions, data is presented by transaction, meaning the software being acquired via the transaction. For example, the average number of open source components found per transaction was 1,635.

The Black Duck Audit Services team audits thousands of codebases for our customers each year, with the primary aim of identifying software risks during M&A transactions. Despite 2023's slowdown in tech M&As, tech acquirers continue to rely on Black Duck for software due diligence insights and advice.

Acquirers in M&A deals want to understand what risks may be associated with the software they're acquiring—specifically around licensing, security, and the quality of the open source used in that software. The audits found open source in the software associated with 99% of transactions, and an average of 1,635 components were discovered per transaction. On average, 77% of this "proprietary software" was in fact open source and third-party code. Given the time-to-market, cost savings, and development advantages of leveraging open source components, it's no surprise that companies continue to rely heavily on open source as part of their software development process. But the large number of discrete components speaks to the challenge of tracking it all.

85% of transactions included components with license conflicts, most frequently Creative Commons Attribution ShareAlike 4.0 licenses.

The vulnerabilities and license compliance issues discovered in the codebases were almost as pervasive as open source itself. In 2023, 85% of transactions included components with license conflicts. The most common conflict was related to the Creative Commons Attribution ShareAlike 4.0 license, closely followed by ShareAlike 3.0.

"Snippets"—lines of code that have been copied and pasted into source code—are quite frequently found by the Black Duck Audit team. These snippets are often taken from the popular blog site Stack Overflow, which automatically licenses all publicly accessible user contributions under Creative Commons Attribution ShareAlike. Unfortunately, the blanket license also covers code snippets posted on the site. We say "unfortunately" because Creative Common licenses are not intended for software, with Creative Commons explicit about this in its FAQ: "We recommend against using Creative Commons licenses for software." The CC-SA license can be read in some situations as having a similar "viral" effect (that is, any work derived from a copyleft-licensed work must also be licensed under the same copyleft terms) as the GNU Public License (GPL) and can become a concern from a legal standpoint.

97% of transactions included at least one unpatched open source vulnerability.

Unpatched software vulnerabilities are one of the biggest cyberthreats organizations face, and unpatched open source components in software add to security risk. Ninety-seven percent of transactions included at least one unpatched open source vulnerability, with a mean of 439 vulnerabilities per transaction. Ninety-four percent of the transactions contained at least one high-risk vulnerability. High-risk vulnerabilities are those that have been actively exploited and already have either a documented proof-of-concept exploit or classification as a remote code execution vulnerability.

There's no single answer for the increase in high-risk vulnerabilities between this year and last, when the percentage of high-risk vulnerabilities was 74%. One possible explanation is the economic downturn limiting the number of resources available to locate and patch vulnerabilities. However, it's indicative that nearly all—99%—the transactions had codebases that were found to contain components 10 versions (or more) behind the most current available version of the component. The simple conclusion is that the majority of open source consumers simply aren't updating the components they use.

Our audits showed that jQuery was the #1 component containing vulnerabilities. Seventy-four percent of the audited codebases contained the jQuery component. It should be noted that jQuery is not inherently insecure. In fact, it is a well-maintained open source library with a large community of users, developers, and maintainers. But according to the audits, jQuery was the component most likely to have vulnerabilities, even though all the jQuery vulnerabilities listed in the 2024 OSSRA report have available patches. To reiterate, this is not meant as an indictment of the quality of jQuery, but rather of the processes that companies use to monitor and patch vulnerabilities. It is important for users of jQuery—and indeed users of all open source—to be aware of the potential security risks associated with the use of older versions of software, and to take steps to mitigate those risks.

In 58% of transactions, the software included components that had no new development activity in the last two years.

Of the customer engagements conducted by the Black Duck Audit Services team in 2023 that included risk assessments, 58% of the software audited included components that had no new development activity in the last two years. Ninety-eight percent had open source more than four years out-of-date, that is, using open source with newer versions available—often with many newer versions available.

Besides contributing to security risk, getting too far behind in versioning increases the danger of functional risk. That is because the more versions that need to be leapfrogged when upgrading a component, the harder the integration can become. Thus, falling behind can be costly and risky when an upgrade becomes necessary to eliminate a vulnerability.

Best practices in the use of open source software require developers to understand which components and associated licenses are in their code and what obligations result from their use of that open source. However, tracking open source manually can be an impossible task for any organization.

Customer engagements in 2023 conducted by the Black Duck Audit Services team found a huge number of open source components (an average of 1,635 components per transaction), each with versions, vulnerabilities, and licenses that need to be tracked.

This is a practical example of the importance—if not absolute necessity—of automated open source management. At that scale, a company cannot rely on manual processes, and an acquirer cannot assume that a target is comprehensively tracking the open source it uses.

The reality is that many companies, particularly smaller ones, don't have the necessary processes and tools to manage their developers' use of open source; it's simply not a high priority to the typical startup. Although we've seen improvement in the last five years, it is still the case that most targets are unable to produce an open source Software Bill of Materials easily and routinely. When they do, it is rarely accurate and never complete. And these companies don't tend to track security vulnerabilities (as evidenced in the forthcoming statistics).

Software composition analysis for M&A due diligence

Knowing what open source code is in a company's codebase is crucial for properly managing its use and reuse, ensuring compliance with software licenses, and staying on top of patching vulnerabilities—all essential steps in reducing business risk. From an M&A perspective, a code audit enables a buyer to understand risks in the software that could affect the value of the intellectual property and the remediation required to address those risks, as well as to plan out a roadmap going forward. Savvy sellers may employ an audit proactively to avoid surprises in due diligence, particularly given the amount of unknown open source in a typical company's code.

An open source audit can be invaluable for companies wanting a better understanding of the code's composition. Using a technique known as software composition analysis (SCA) and a range of sophisticated tools, expert auditors comprehensively identify the open source components in a codebase and flag legal compliance issues related to those components, prioritizing issues based on their severity. The audit identifies known security vulnerabilities that affect the open source components, as well as information such as versions, duplications, and the state of a component's development activity. It also provides clues as to the sophistication of a target's software development processes. Open source is so ubiquitous today that if a company isn't managing that part of software development well, it raises questions as to how well it is managing other aspects.

If you're on the buy side of a tech M&A transaction, an open source audit should absolutely be part of the software due diligence process. Acquirers need to identify problematic open source in the target's code before the transaction terms are set, and a trusted third-party audit is the best way to get a deep, comprehensive view. Prospective sellers should prepare for questions about the composition of their code and how well they have managed open source security and license risk. Proactive sellers can prepare for an acquisition by having their software audited in advance.

Open source license compliance risk

Like all software, open source components are governed by licenses that vary in terms of rights, obligations, and restrictions. Failure to comply with open source licenses can put businesses at risk of litigation and could compromise their intellectual property. Generally, as part of a definitive agreement, sellers need to represent that they have the rights to any software they are using (and there may be even more explicit representations and warranties regarding open source). Yet few sellers are totally "clean" when it comes to open source license compliance.

Organizations can manage and comply with license requirements only if they can identify the open source components and confirm that the use of those components is consistent with the terms of the applicable license. Just as with security vulnerabilities, it's impossible to manage license compliance risk without identifying all the components in the software.

Open source license compliance remains critical

Based on Black Duck audit data, investors in all software verticals should be concerned about open source licensing and the potential risk of litigation or threat to their intellectual property rights due to failure to comply with an open source license. Black Duck analyses indicate that the 20 most popular licenses cover approximately 98% of the open source in use. But whether the software in question uses one of those popular licenses or some variant, the license matters.

In 68% of transactions, auditors found open source with customized licenses, or components freely available on the internet but with no discernible license at all.

License risk arises when software includes open source with licenses that conflict with the overall license of the codebase. For example, the GPL is an extremely common license that often governs components in commercial software and requires distributors to make source code available. But commercial software vendors typically do not offer to provide source code or complete source code, which creates a conflict with that license.

Sometimes an open source component has a so-called "custom license" in which the developer created their own licensing language or added language to a variant of a standard license. Customized open source licenses might place undesirable requirements or limitations on the licensee and will often require legal evaluation for possible IP issues or other implications. For example, the JSON license is based on the permissive MIT license, but the JSON license adds the distinction that "the software shall be used for good, not evil." The ambiguity of this statement leaves its meaning up to interpretation, posing a particular concern in M&A scenarios where acquirers are hesitant to inherit this type of indistinct legal risk. And there are many examples of developers being similarly "creative" with licenses.

Black Duck audit engagements conducted in 2023 found that 68% of transactions included open source with customized licenses or no license at all. If third-party code is used without a license, this raises legal concerns. In the U.S. and many other jurisdictions, creative work—including software—is placed under exclusive copyright by default. Unless there's a license that specifies otherwise (or the copyright holders grant permission), no other party can use, copy, distribute, or modify the software without the risk of litigation.

Broken down by industry, the sector with the highest percentage of codebases that contained open source license conflicts (92%) was the computer hardware and semiconductors sector. The marketing tech sector had the lowest percentage of codebases with open source license conflicts at 19%.

Permissive and reciprocal licenses

Open source licenses fall on a spectrum from permissive to reciprocal. Permissive licenses place minimal obligations on companies that redistribute the associated software. By contrast, reciprocal (also known as "copyleft") licenses require the licensee to make any improvements or enhancements available to the public under similar terms. In some cases, the entirety of the work that incorporates the licensed software, even a small portion of it, may fall under the reciprocal obligation.

Permissive licenses

Permissive open source licenses generally require only that the licensee attribute the original portions of the licensed code to the original developers both in the code and in documentation. Of course, to provide such an attribution, the licensee must be aware that they're using the licensed code, so it is no surprise that most targets are not fully compliant with permissive licenses. While this is generally considered a lesser risk than, say, distributing GPL-licensed code, it is still an issue that most acquirers will want to address in their plans to meet their own corporate standards, typically stricter than those of most targets.

Reciprocal licenses

On the other end of the spectrum, codebases containing reciprocal licenses are quite problematic for an acquirer. Many of these licenses require associated code to be made available to the public under the same license. A licensee that violates a reciprocal license could be at risk of litigation and may be required to disclose all the source code of the application. Such issues often come to light when a licensee is acquired in an M&A transaction, and acquiring companies will want to remediate either before or after close.

IP compliance risk introduced by AI coding tools

Arising with the use of AI-powered coding suggestion tools are questions around ownership, copyright, and licensing of the generated code. For example, a class-action lawsuit filed against GitHub, Microsoft, and OpenAI claims that GitHub Copilot—a cloud-based AI tool that offers developers autocomplete-style suggestions as they code—violates both copyright law and software licensing requirements. The lawsuit further claims that the code suggested by Copilot uses licensed materials without attribution, copyright notice, or adherence to the original licensing terms.

The Copilot case highlights the legal complexities surrounding the use of AI-generated code. For software developers, refraining from using AI-assisted coding tools until the issue is resolved by legal or government decision is obviously the safest way to avoid an action for license or copyright violations, but the reality is that many developers continue to use them. An open source audit as part of software due diligence will highlight instances where an AI has "copied" code from open source projects.

Open source security risk

Large organizations may manage hundreds to thousands of software assets, ranging from mobile apps to cloud-based systems to legacy systems running on premises. That software is typically a mix of commercial off-the-shelf packages and custom-built codebases, both of which increasingly contain open source components.

As noted earlier, the Black Duck Audit Services team found open source in 99% of transactions in 2023. Here's the reality: If an organization builds, sells, or simply uses software, it's safe to assume that the software contains open source.

In 97% of transactions in which we were involved, we found unpatched open source vulnerabilities, with a mean of 439 vulnerabilities uncovered per engagement. Ninety-four percent of the transactions had at least one high-risk vulnerability. "High-risk" indicates that a vulnerability has been actively exploited, has documented proof-of-concept exploits, or has been classified as a remote code execution vulnerability.

Open source projects usually issue small updates at a much higher frequency than the average commercial software vendor. When these updates contain security updates, companies need to have a strategy to adopt them rapidly. But because open source updates need to be "pulled" by users, many companies consuming open source components don't apply the patches they need, exposing their business to the risk of attack and their applications to potential exploits. This is understandable—these companies typically aren't even aware of the outdated components in their code.

The data indicates that development teams may be struggling with the dynamic nature of open source security risk, especially with the increase in open source use. An open source component with no known vulnerabilities doesn't necessarily stay that way a year, month—sometimes not even a week—later.

The audits also demonstrate that many organizations are startlingly behind in using the latest version of any given open source component. Ninety-eight percent of our 2023 M&A transactions included open source components more than four years out-of-date. Ninety-nine percent of the transactions were not using the most current version of the component—exposing their codebases to security risks and other issues. Firms acquiring companies, undergoing mergers, entering joint ventures, or even managing underlying technology supply chains must also consider the risks they are potentially inheriting.

Summary

The thousands of audits conducted by the Black Duck Audit Services team have consistently revealed that almost every codebase contains open source code. As documented in this paper, 99% of transactions audited in 2023 included open source components.

The Black Duck Audit Services team generally audits codebases from software-heavy companies, as opposed to enterprises that use software to support their business. The primary value of software companies is their proprietary code. The ratio of open source to proprietary code in their codebases, while still quite high at 77%, is eclipsed by how much open source is used in large enterprises. The figures cited by analysts such as Forrester, which generally look at enterprise IT groups for their reports, consistently find that over 90% of IT organizations use open source software in mission-critical workloads, and that open source often comprises up to 90% of new codebases.

With the growth of open source usage comes risk, due primarily to organizations lacking the needed tools and processes to recognize what—and how much—open source is in their internal and public-facing applications.

Failure to comply with open source licenses can put businesses at significant risk of litigation and jeopardize their ownership rights to their software. Perhaps more importantly, especially in the context of M&A transactions, failure to comply with open source licenses more likely to result in demands for additional acquirer protections, pre-close remediation, escrows, purchase price reductions, or deal delay. Eighty-five percent of transactions in 2023 involved components with license conflicts. Sixty-eight percent had open source with customized licenses or no discernible license at all. In other words, license issues are pervasive in the M&A deals we see, and they are potentially costly.

Significant monetary and brand risk can be buried in the open source components of an acquired application. Evaluating that risk as part of an acquirer's due diligence must be a factor in the decision-making process.

By identifying open source code and third-party components and licenses, an open source audit can alert your firm to potential legal and security issues in an M&A transaction. With an open source audit, you can also

- Avoid surprises
- Mitigate legal exposure
- Understand risks that may affect software asset values
- Resolve potential issues before they affect the transaction
- Build appropriate protections into the deal terms
- Plan integration and remediation

For more information, download the 2024 "Open Source Security and Risk Analysis" report. Learn more about Black Duck audits.

About Black Duck

Black Duck® offers the most comprehensive, powerful, and trusted portfolio of application security solutions in the industry. We have an unmatched track record of helping organizations around the world secure their software quickly, integrate security efficiently in their development environments, and safely innovate with new technologies. As the recognized leaders, experts, and innovators in software security, Black Duck has everything you need to build trust in your software. Learn more at www.blackduck.com.