Xebia

Whitepaper

State of Generative AI Code Assistants in Software Development

v1.3 | February 2024

Table of Contents

Executive Summary

01. Generative AI offers substantial potential to boost productivity in software development, but introduces new concerns in data security, intellectual property rights, and code quality.

02. While legal frameworks for AI are still evolving, the EU and US offer preliminary guidance already.

03. Research indicates software development productivity gains ranging from 10% to 50%, depending on the nature of the task.

04. Despite the widespread use of Gen AI assistants in Europe, only 18% of organisations implement risk management.

05. Selecting the right vendor is crucial for maintaining data safety and confidentiality. Research shows that all free-tier licences are unsuitable for handling proprietary and confidential data.

06. Tabnine Pro emerges as the most secure tool among the four assessed, as its model relies solely on permissive open-source data, posing no greater risk than manual coding. However, it falls short in output quality compared to Amazon's CodeWhisperer Professional and GitHub's Copilot Business, with developers showing a strong preference for the latter. OpenAI's ChatGPT Team, on the one hand, lacks the built-in duplication detection filters offered by its competitors, but on the other, it offers distinctive functions enabling collaboration within the workspace. Interestingly, almost all of the examined tools offer some form of legal protection against copyright infringement under their business licences, provided certain conditions are met.

07. Regardless of the tool chosen, organisations must ensure that the output of Gen AI tools meets quality standards.

08. The recommended strategy for Gen AI adoption is cautious progression, balancing efficiency gains with potential risks. Providing explicit guidelines on permissible tools and their usage, along with comprehensive communication and training, is essential.

Introduction

In the current era of continual digital transformation, organisations strive to improve efficiency while maintaining technical excellence. To tackle this challenge, some experts suggest adopting Generative Artificial Intelligence (Gen AI) and its associated tools, but Xebia takes a more methodical approach and scrutinises these tools with an emphasis on security considerations and associated risks.

To evaluate aspects such as data security, intellectual property rights, and commercial viability, Xebia conducted an all-encompassing inquiry into generative tools tailored for software developers. This endeavour involved not only an assessment of the existing legal framework but also an examination of the contractual stipulations and privacy regulations from the foremost industry vendors. The study additionally examined the effect of Gen AI tools on enhancing developers' daily activities, measuring job satisfaction, the sense of meaningful work and ease of implementation.

This report details the primary conclusions derived from our study.

Key Objectives

The objective of this paper is to provide an overview of the leading Gen AI tools used in software development, with a focus on the following areas:

01. Ensuring the safety and confidentiality of inputs and outputs.

02. Addressing intellectual property rights concerning generated outputs and managing the associated risk of third-party infringement claims.

03. Examining the impact of these tools on developers' productivity and overall job satisfaction.

As a result of the study, Xebia has devised and enacted an internal policy to ensure the secure adoption of Gen AI tools. The following sections provide an overview of our findings and recommendations.

Research Indicates Significant Efficiency Improvements

Generative AI has become a potent asset in the realm of software development. Recent studies highlight how developers using these tools achieve substantial efficiency gains. McKinsey & Company's research shows in Figure 1 a considerable reduction in the time required to complete specific tasks with the assistance of Generative AI.[1]

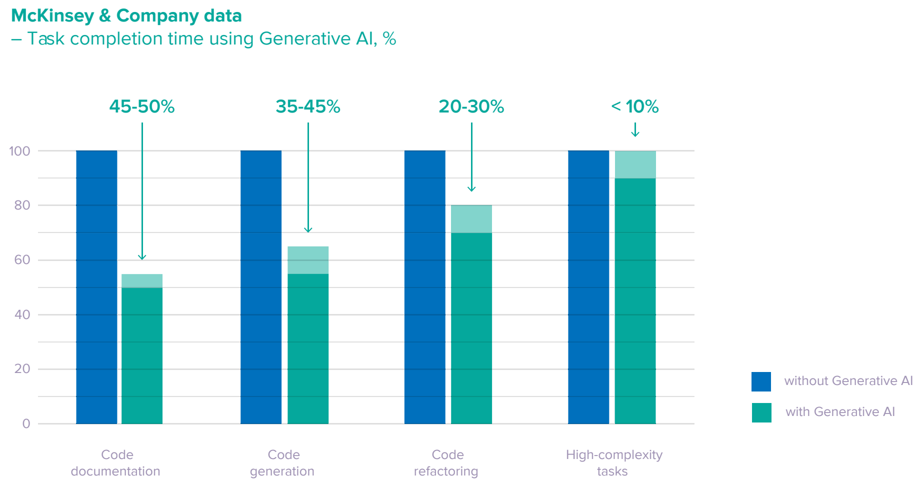

Figure 1
Task completion time savings using Generative AI.

The X-axis depicts the four types of tasks examined, while the Y-axis represents the time needed to accomplish each task. The bars on the left correspond to the initial time to accomplish these tasks (normalised to 100%) and the bars on the right correspond to the time necessary to complete the tasks with a Gen AI code assistant. The figures above each column represent the quantified time savings. Source: McKinsey & Company.

[1] McKinsey & Company, B. K. Deniz, C. Gnanasambandam, M. Harrysson, A. Hussin, S. Srivastava, Unleashinxg developer productivity with Generative AI

To validate the findings of McKinsey & Company, we conducted our own assessment to measure the productivity improvements among a group of developers who use Generative AI tools. Figure 2 shows four charts, each corresponding to one of the four areas investigated by McKinsey & Company.

Figure 2
Task completion time savings using Generative AI – broken down into four areas.

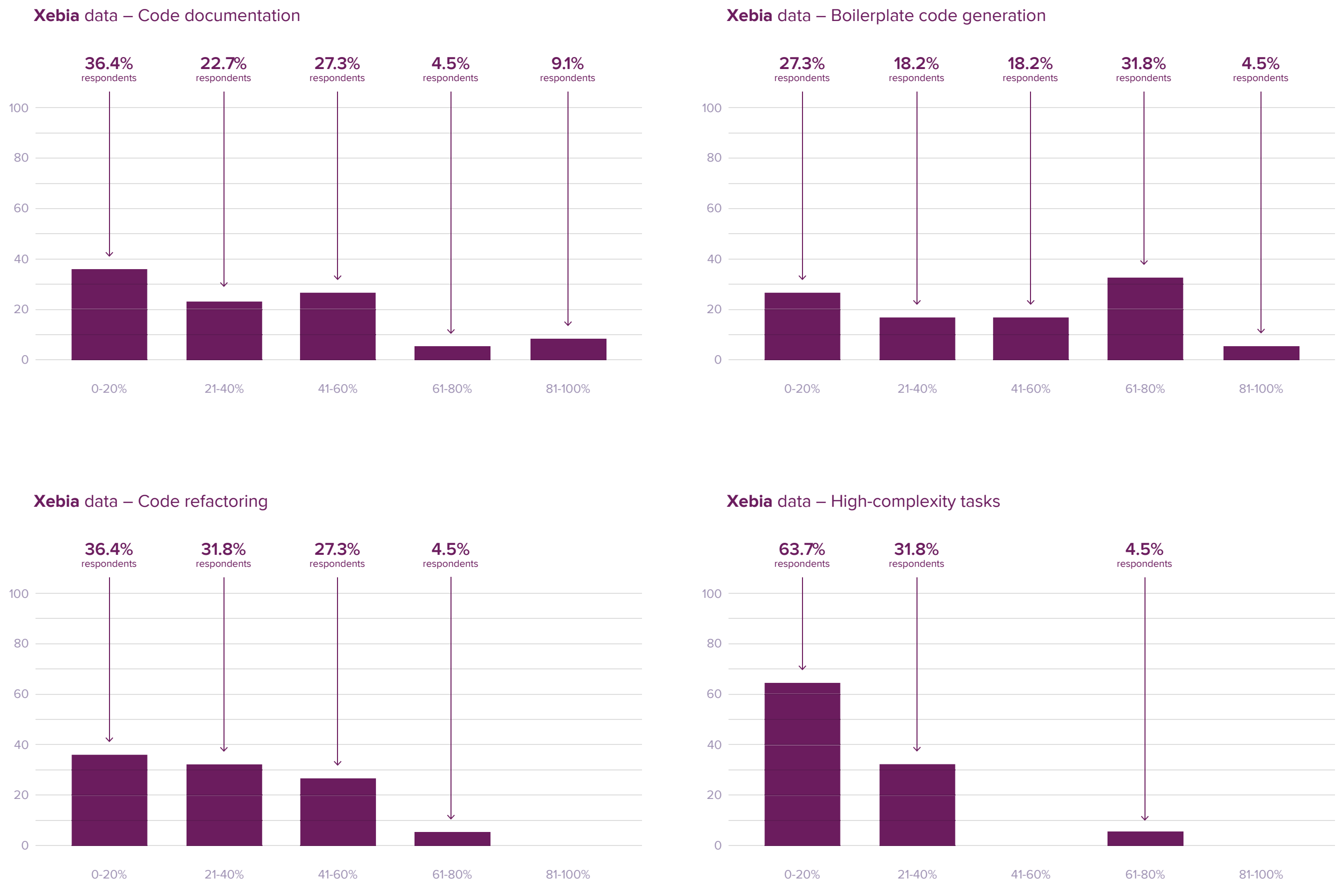
Xebia data – Code documentation

Xebia data – Boilerplate code generation

Xebia data – Code refactoring

Xebia data – High-complexity tasks

The X-axes display five distinct ranges of time savings, while the Y-axes show the percentage of developers who reported that said time saving.

Our observations echo a consistent message: Gen AI yields substantial time savings in software development. McKinsey's research indicates that for less intricate tasks, the time savings achieved through these tools are more significant, sometimes reaching up to a 50% reduction. In contrast, for tasks of higher complexity, the reduction is limited to around 10%. Similarly, Xebia's investigation revealed comparable results, with many participants reporting a 40% or more reduction in time spent on less complex tasks when using Gen AI, whereas for more intricate issues, most developers reported smaller time reductions. Interestingly, the proficiency of less experienced developers may even decline when attempting to solve complex problems with new tools. According to McKinsey's study, these less experienced developers, compared to the control group, might require up to 10% more time to complete such high-complexity tasks while using Gen AI assistance.

The benefits of utilising Generative AI extend beyond bolstering developer productivity. According to both studies, the use of Gen AI significantly contributes to developers' job satisfaction and overall well-being. As figure 3 shows, among developers who incorporated Generative AI into their workflow, an overwhelming 88% (McKinsey) and 81% (Xebia) reported feeling content during work, with the remainder expressing either no clear opinion or disagreeing only slightly. In contrast, McKinsey's study revealed that only 45% of those who did not use these tools shared the same sentiment, while 30% were unhappy. Consequently, Generative AI arguably serves as a catalyst for motivating employees and retaining top talent within organisations.

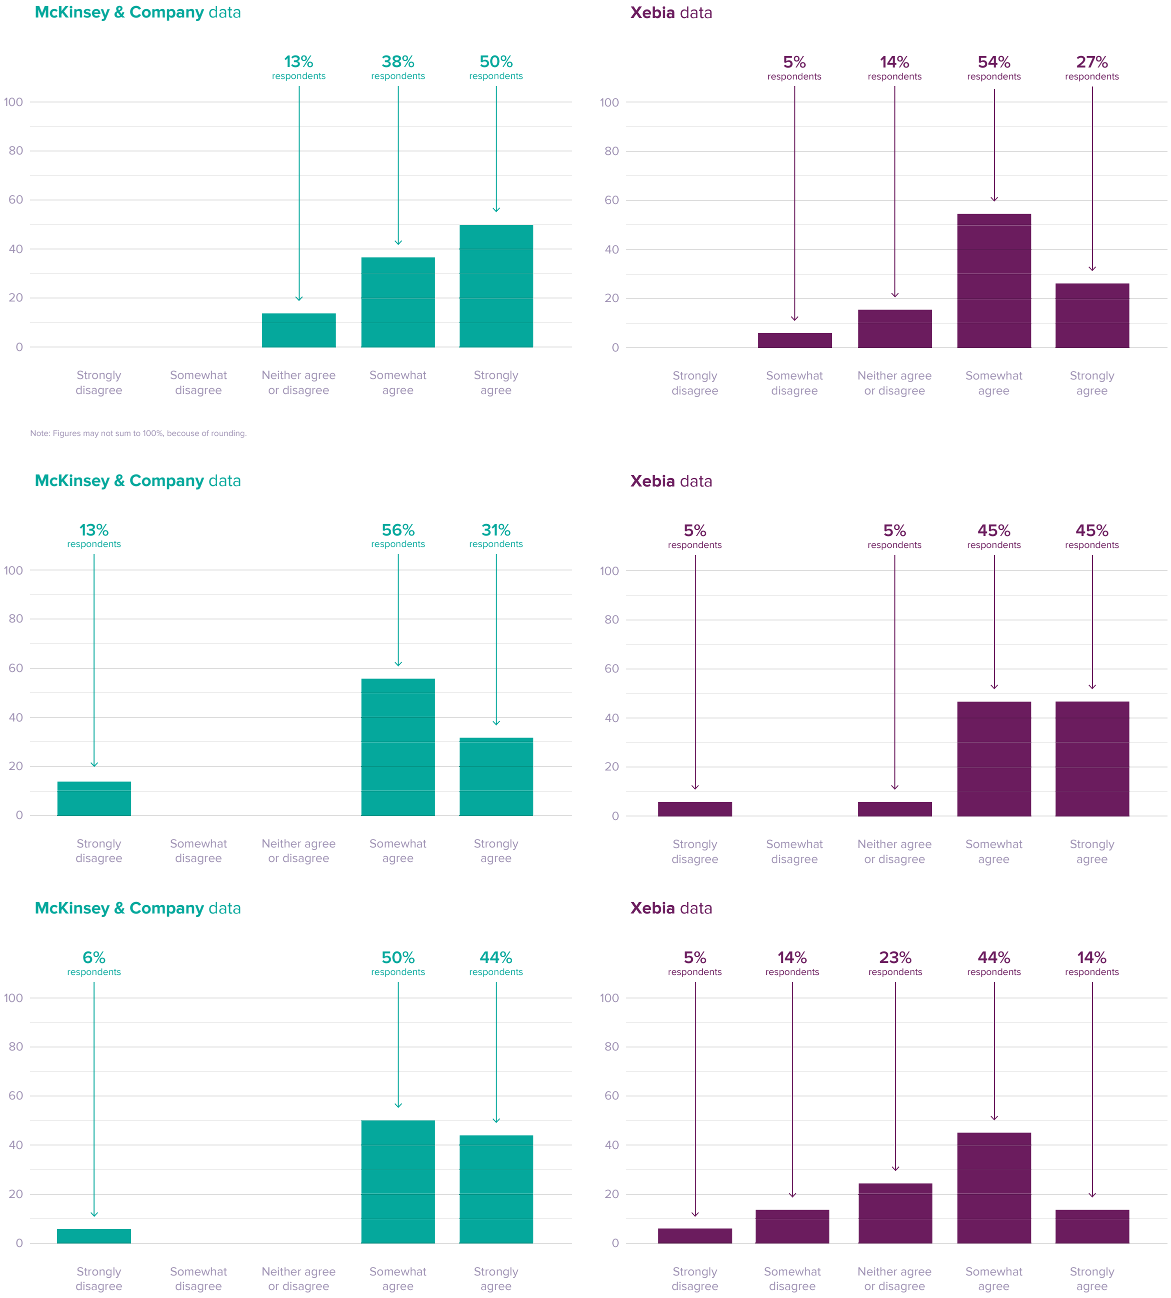

Note: Figures may not sum to 100%, because of rounding.

Figure 3
Satisfaction assessments when using Gen AI tools.

The three charts for both studies show how many developers share positive feelings about using these tools at work.

"Using Gen AI-assisted tools makes me feel happy about my work."

"Using Gen AI-assisted tools allows me to focus more on satisfying and meaningful work."

"Using Gen AI-assisted tools allows me to enter a 'flow' state quicker."

Risks Associated with Generative AI Usage

The adoption of Generative AI technologies in software development has significant security implications and raises data privacy concerns. Organisations should address vulnerabilities and privacy issues that may arise when adopting new tools. However, the research conducted by Opinium for the Ricoh group found that so far only 18% of organisations have implemented risk management measures around Gen AI deployment.[2]

The challenges associated with Gen AI tools fall into four core areas:

Output quality concerns

The quality of the produced output guides the subsequent decisions and actions of software developers, each of which can either augment or compromise the overall project's trajectory. Some of the concerns around quality are:

- **Hallucinations – Gen AI may be misleading**
  Gen AI tools often exhibit unwavering confidence in their responses, even when producing fictitious data. In the most severe instances, the output is entirely fabricated, constituting what is known as "Gen AI's hallucination". This phenomenon is not an outlier but rather a prevalent occurrence, rendering the task of distinguishing the truth from falsehood inherently challenging. Even when not entirely fabricated, the quality of outputs, such as generated code, may be subpar or harbour security vulnerabilities. The risk becomes more pronounced when engineers use Generative AI in conjunction with technologies outside their area of expertise. This requires an enhanced examination of AI-generated artifacts, including peer reviews. Implementing standard precautions, such as integrating code quality and security scanning into delivery pipelines, is highly recommended.

- **Biases in output – Gen AI may favour one solution over another**
  This concern stems from the inherent bias that may be encoded within the model used to generate output, which is predominantly determined by the data sources used for training.[3] While the veracity of Gen AI output may not be disputed, its alignment with the overarching context can become controversial. In numerous instances, the generated solutions may exhibit differences from prevailing methodologies or non-conformity with internal corporate protocols. Code-generating models may also tend to endorse particular technologies, thereby issuing non-optimal recommendations, even when an objectively superior alternative exists.

- **Dated models – Gen AI's outputs may be based on out-of-date information**
  Gen AI models generate outputs based on information that may become obsolete. The significant time investment required to train the model exacerbates this vulnerability. Gen AI model used by a given tool may not be equipped with the latest insights into industry trends and developments, including new frameworks, tools, programming libraries or updates to programming language versions. As a result, the output could inadvertently generate obsolete or invalidated solutions, resulting in developers not adhering to the latest programming paradigms, best practices, security criteria or speed enhancements. The reliability and security of the software produced could therefore be compromised.

[2] Ricoh Europe, Enhancing employee experience with automation and AI
[3] Harvard Business School, ChatGPT: Did Big Tech Set Up the World for an AI Bias Disaster?

Data security and confidentiality

While data in the context of Gen AI tools is normally perceived as the text entered into a chat interface, it is pivotal to recognise that the code base itself can act as an input as well. Depending on the terms and conditions, any input data may be used for further training of the underlying model. In the direst of circumstances, data introduced as an input could inadvertently re-emerge as part of an output in response to another user's query.[4]

Data safety concerns may also arise due to the transmission of telemetry data for the provider's analytical purposes. Therefore, it is necessary to carefully configure the tool and conduct a thorough review of its default settings and data persistency policy. For entities prioritising data security, the adoption of locally hosted solutions, exemplified by on-premises models, is the optimal course of action.

Probability of Intellectual Property Rights infringement

Generative AI models derive their efficacy from extensive training data, primarily sourced from vast public code repositories. Nonetheless, concealed within this strength lies a potential legal conundrum. The introduction of code from non-permissive open-source licences, such as the General Public License (GPL), to a company's repository could lead to significant legal repercussions.

To ensure legal compliance, it is imperative to scrutinise the origin of the source code used to train these models. Setting a tolerable risk threshold is advisable for organisations. Neglecting this due diligence may culminate in the integration of generated output that breaches the intellectual property rights, and, in the grimmest scenario, a company could face legal liabilities worth millions of dollars.

Only a minority of Gen AI-powered assistant providers assures that their tools are exclusively trained on permissive repositories, reducing the risk of infringing on intellectual property rights when using generated content. Conversely, other vendors offer alternative mechanisms, such as filtering processes, to curtail the risk. Some consider taking legal action to protect their customers in the event of third-party claims, subject to certain conditions being met.

Intellectual Property Rights to Gen AI outputs

The emergence of artificial entities that produce creative output poses an unprecedented challenge to the existing legal framework. The current legislative landscape offers limited clarity regarding the regulation of AI-generated content. This uncertainty is further underscored by the assertion from the U.S. Copyright Office, which holds that Gen AI lacks the necessary authorship attributes to be deemed a creator of creative output.[5]

Within the domain of code generated by Gen AI tools, service providers either simply declare they does not claim output ownership rights, or they transfer the ownership rights to the requesting user, although without conferring full copyright ownership. The concept of authorless copyrights, as it stands, remains an elusive legal construct, with no established jurisprudential consensus regarding the copyrightability of such code.[6] Consequently, pending any statutory clarification, it is judicious to presume that AI-generated artifacts remain unshielded by copyright protection, although users may possess ownership rights. Notably, even in instances where ownership rights are transferred, the potential for the generated output to infringe upon the intellectual property rights of third parties persists.

[4] Bloomberg, Samsung Bans Staff's AI Use After Spotting ChatGPT Data Leak
[5] Bloomberg Law, IP Issues With AI Code Generators
[6] TechTarget, Is AI-generated content copyrighted?

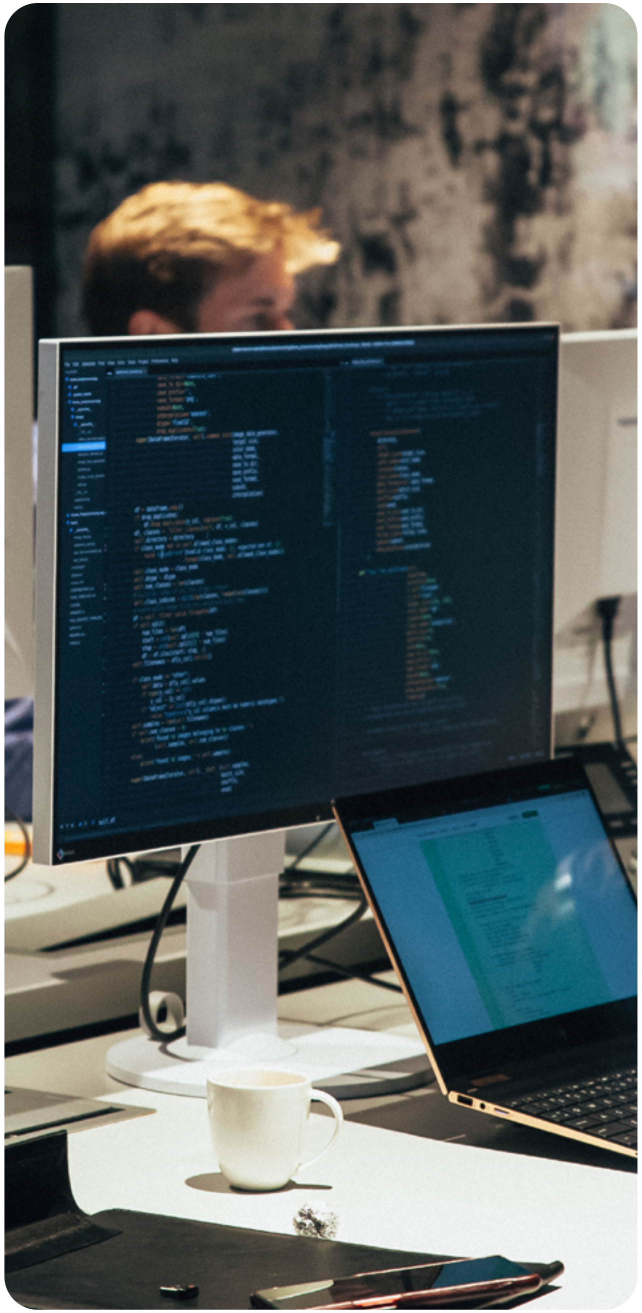

Current Regulatory Frameworks

The area of Generative AI currently lacks regulation due to its rapid development. Both the European Union and the US have embarked on a comprehensive examination of the responsible use of artificial intelligence, but there are still no official legal requirements in force.

At the forefront of this endeavour in the EU stands the AI Act.[7] Incepted in 2021, the AI Act is anticipated to navigate the legislative corridors of the EU. The central tenet of the AI Act is the self-classification of AI systems based on their inherent risk. Its enforcement is set to commence during a transitional period scheduled for late 2023, therefore the AI Act could come into effect by late 2025 or early 2026.

In the US, the White House Office released the Blueprint for an AI Bill of Rights in October 2022.[8] Its goals are to prove that everyone is protected from the unsafe use of automated systems, to prevent algorithmic discrimination, to ensure data privacy, and to provide an explanation of any automated system in plain language. Currently, it serves as a set of recommendations rather than a legally binding statute, akin to the AI Act status in the EU.

These new regulations, especially EU's AI Act, bear a striking resemblance to GDPR. Just as the latter became a global benchmark, the new act has a similar potential in the field of AI. Non-compliance with the regulations can prove as costly for businesses as with GDPR, thus it is vital that all take note of the new law. In case of the AI Act, the severity of the penalties for non-compliance depends on the type of offense, ranging from €10-30 million or, if the offender is a company, up to 1-6% of its total worldwide annual turnover for the preceding financial year, whichever sum is greater.[9]

Notably, the new regulations could have a greater impact on Gen AI tool providers than on their users. However, the providers may need to modify their terms of use in response to the new legislation, and that could have a direct impact on organisations using them. Although the new law is still not active, it is recommended to keep track of the new AI regulations and implement them to prevent costly alterations in the future.

[7] European Parliament, Artificial Intelligence Act: Briefing
[8] The White House Office, Blueprint for an AI Bill of Rights
[9] Publication Office of the European Union, Proposal for the Artificial Intelligence Act, Article 71

Research on Code Generating Tools

Verification criteria and methodological framework

The primary goal of this study was to develop a robust framework to assess challenges and research questions about Gen AI code assistants. A comprehensive literature review was carried out: this encompassed studies, whitepapers, articles, and documents from entities engaged in Gen AI research and development. Moreover, an assessment of available Gen AI code assistants was conducted, scrutinising user manuals, privacy policies, and terms of service. Valuable insights were gleaned from interviews with key players such as software developers, AI experts, legal authorities, and business professionals. Feedback from software developers was also collected through surveys, giving a holistic view of their experiences with these tools.

Our evaluation focused on how Gen AI tools manage sensitive data, specifically their capacity to maintain the confidentiality of proprietary information covering storage, transmission, and processing. We also examined the intellectual property rights related to code output from these tools, seeking clarity on whether the tools grant ownership rights to developers or organisations, or retain certain rights. This clarity is crucial for both legal compliance and alignment with our organisational goals. Lastly, we assessed the readiness of Gen AI tool vendors to protect their customers from third-party claims, examining the depth of legal support and documentation they provide for handling intellectual property disputes.

Commercial vs free licensing

This report emphasises the use of Gen AI tools in a professional setting, excluding the free or "personal" versions of these tools since they do not meet basic confidentiality standards. Notably, most free licences utilise input data to train their models by default, while only some of them provide an opt-out option under a cost-free tier. Not having full control over how data is processed makes free-tier licences unsuitable for handling proprietary and confidential information.

Amazon CodeWhisperer Professional

CodeWhisperer is one of the majors Gen-AI-driven software development assistants. Its primary function is generating code suggestions based on user input. Its affiliation with Amazon is crucial in projects leveraging Amazon Web Services.[10] However, the utility of CodeWhisperer extends beyond the confines of Amazon-centric projects and can function as a general-purpose development assistant in diverse contexts. For organisations that already use Amazon services, two compelling factors encourage its adoption:

- Firstly, the pre-established trust in Amazon as a vendor reduces the need for extensive security verifications, as certain data sharing has already occurred;
- Secondly, Amazon's pricing policy, characterised by progressive discounts for higher expenditure, makes the selection of this vendor a cost-effective choice.

Category	Assessment	Notes
Data storage	**Positive**	The Professional version processes user code snippets, comments, and file content solely to provide and maintain the service. The processed content is not used or retained for further service improvement purposes.
Input safety and confidentiality	**Positive**	The Professional version does not store or use the processed content to train the model or reproduce suggestions for other users.
Output copyright and licensing issues	**Warning**	On rare occasions, the tool may produce code snippets similar to those in the training data, which consists of non-permissive open-source repositories.
Output ownership	**Positive / Warning**	The vendor declares that a user owns the output produced by the tool, but that output is not protected by intellectual property rights, at least not in the US and EU. Additionally, there is no assurance that the output will not violate third parties' IPRs.
Defence against third-party claims	**Positive**	CodeWhisperer's reference tracker identifies code recommendations that may be similar to the training data. If enabled, and if a couple of other judicious conditions are met, Amazon will defend all of its Professional plan users against third-party claims of intellectual property infringement. This includes payment of any adverse final judgement or settlement, with no cap on liability.

The first aspect to consider is data storage. CodeWhisperer demonstrates commendable adherence to privacy standards, as it avoids the persistent storage or subsequent use of user input, including code snippets, comments and file content, within the Amazon ecosystem. Its sole function is to provide services, with no ulterior motive to enhance the tools.

When considering copyrights and licences, our investigation reveals that the use of the Professional Tier of CodeWhisperer can occasionally produce output similar to the code snippets used in the training dataset. This issue is heightened by the use of non-permissive open-source licences, notably the General Public License (GPL), in the training dataset of the Amazon model. Consequently, the risk of inadvertently generating code subject to GPL-like licences is not insignificant.

While the likelihood of such an event is quite low, its materialisation could have significant legal and financial consequences. One convenient aspect of Amazon's tool in this regard is that it assigns ownership of the output code to the user.

... continues on the next page

[10] Medium, GitHub Copilot vs. AWS CodeWhisperer: A Comparative Analysis

In the realm of mitigating risks associated with third-party assertions of intellectual property rights infringements, Amazon classifies CodeWhisperer Professional as an "Indemnified Generative AI Service."

This designation signifies Amazon's commitment to legally protect its users from allegations that CodeWhisperer's outputs infringe upon the intellectual property rights of third parties. Furthermore, Amazon pledges to cover the expenses associated with any adverse final judgement or settlement, explicitly stating that these financial obligations are exempt from any limitations on damages specified in the agreement with AWS.

However, eligibility for such indemnification is contingent upon adherence to specific conditions.[11] These require that the claim does not arise from the user's input violating another entity's intellectual property rights, that the user acts in good faith, and that the user activates all available filtering mechanisms.

The filtering mechanisms operate as an automated reference tracking system designed to detect similarities between code suggestions and non-permissive training data, thereby minimising the risk of infringement.

GitHub Copilot Business

This tool leverages the Gen AI model jointly developed by OpenAI, Microsoft, and GitHub. Copilot Business is trained on a diverse dataset of publicly accessible source code, including an extensive repository of code publicly available on GitHub. This comprehensive training dataset encompasses, however, instances governed by non-permissive open-source licences, such as the General Public License (GPL).

When examining data storage practices, GitHub's tool showcases admirable data privacy measures for the data provided by users, such as prompts, and suggestions returned by the tool. User prompts are used solely to generate suggestions, and are discarded once a suggestion is returned. The business licence also guarantees that prompts and suggestions received remain protected, preventing their use for making recommendations to other users or for any other processing or sharing activities.

GitHub also collects user engagement data, i.e. information on events generated when interacting with the tool. These are retained for 24 months[12], and are used by GitHub and Microsoft to provide and further improve the service, as well as to detect potential abuses or violations of the terms of use. In addition, GitHub allows its business users to enter into a Data Protection Agreement[13], which increases the transparency of data processing.

The posture regarding intellectual property rights aligns closely with Amazon's precedent, but it differs within the sphere of the ownership rights. Copilot Business does not claim ownership of any suggestion, but it also does not seek to determine whether a suggestion can be owned by anyone at all. The company notes that the output ownership rights may depend on many factors, including but not limited to the laws of the relevant country, or the length of the suggestion. Regardless, there is no claim of output ownership.

In terms of the risk of intellectual property rights infringement, GitHub reports that a scenario where the code output is similar to the training data can occur, although it happens in less than 1% of cases. This issue results from the non-permissive licences incorporated during the model's training phase.

... continues on the next page

[11] Amazon, AWS Service Terms

[12] GitHub, GitHub Copilot Business Privacy Statement

[13] GitHub, GitHub Data Protection Agreement

In order to mitigate the risk, Copilot includes a duplicate detection filter to suppress certain suggestions that match publicly available code for snippets of at least 150 characters. With the filter enabled, suggestions that match the criteria will not be returned to the user.

In the realm of legal safeguards, GitHub was the first of the major vendors to offer active legal protection against third-party claims, but it only does so under certain conditions. Firstly, such protection is contingent upon the comprehensive utilisation of Copilot Business's duplicate detection filter.[14] Secondly, GitHub's acceptable use policies prohibit the employment of content that infringes upon the proprietary rights of any entity as a prompt, thereby necessitating user verification of input materials. Thirdly, GitHub's general terms delineate liability solely for code that remains unmodified, as provided by GitHub.[15] Despite the absence of empirical instances where such protection has been enacted, these stipulations bear significance, considering that developersoften modify AI-generated code, potentially voiding GitHub's legal safeguards.

Furthermore, eligibility for GitHub's indemnification requires immediate written notification of the claim and granting autonomous authority over its defence and settlement to the company, practices that align with industry standards. Subject to the satisfaction of the aforementioned conditions, GitHub pledges to pay the amount of any resulting adverse final judgement or approved settlement.

Noteworthy, in the context of liability constraints, GitHub expressly excludes the obligation to defend against third-party claims from any limitations applicable to standard subscription products.

Category	Assessment	Notes
Data storage	Positive	The Business version sends code snippets to GitHub solely for suggestions, and then deletes them once the suggestion is provided. However, it retains user engagement data for 24 months, and uses this data to improve the service and detect abuse.
Input safety and confidentiality	Positive	GitHub's privacy statement guarantees that any user suggestions and prompts are not shared or used as code suggested for other Copilot users.
Output copyright and licensing issues	Warning	According to internal research by GitHub, Copilot may generate code with perfect matches to the training data, including non-permissive open-source repositories, although the likelihood is low and less than 1%.
Output ownership	Positive / Warning	The vendor does not claim the output ownership rights. The output is not protected by intellectual property rights, at least not in the US and EU. Additionally, there is no assurance that the output will not violate third parties' IPRs.
Defence against third-party claims	Positive / Warning	GitHub provides legal protection against third-party claims of intellectual property infringement, subject to certain conditions being met. The tool must have all duplication detection filtering features enabled and be used in accordance with acceptable use policies. The generated code must remain unmodified, as provided by GitHub. This requirement, combined with the tendency of developers to customise the code received, raises the question of whether GitHub's indemnification is practically applicable. There is no cap on damages liability.

[14] GitHub, GitHub Copilot Product Specific Terms

[15] GitHub, GitHub General Terms

Tabnine Pro

Tabnine, a relatively lesser-known vendor, offers a tool of the same name. We examined the Pro version of it. Notably, a distinguishing feature of Tabnine is the exclusive use of permissive open-source repositories for its model training, setting it apart within the landscape of generative software tools.

When assessing data storage methods, Tabnine Pro does not store code snippets or user inputs. Its approach ensures short-term data retention, and the lack of extended storage prevents the use of this data for training their public model.[16] In line with common practices among generative tool vendors, Tabnine grants users ownership of the output, though with a unique twist.

Specifically, Tabnine provides users with a perpetual non-exclusive license for the generated code. The difference becomes evident when examining Tabnine's stance on copyrights and licences used in their model's training.

Category	Assessment	Notes
Data storage	**Positive**	The Pro version sends user input and code snippets to the tool for precise and pertinent suggestions, but Tabnine does not store this information. Telemetric data is sent to Tabnine by default, but it is possible to opt out.
Input safety and confidentiality	**Positive**	The Pro version does not use any user data or code for the public model training or to generate suggestions for other users.
Output copyright and licensing issues	**Positive**	Tabnine uses only permissive open-source repositories to train the model.
Output ownership	**Positive**	Tabnine grants the user a non-exclusive, perpetual and royalty-free license to use generated code.
Defence against third-party claims	**N/A**	Tabnine uses only permissive open-source repositories to train the model.

As previously mentioned, Tabnine's model is trained exclusively on permissive open-source repositories, characterised by their lenient rules on software use, modification, and distribution. As a result, using Tabnine does not pose a greater risk of intellectual property rights violations than manual code writing by a developer.

Given these factors, Tabnine Pro users can be confident that using the tool is unlikely to lead to legal issues related to the use of the code produced by the software.

[16] Tabnine, Privacy Policy for Tabnine Pro

ChatGPT Team

While not exclusively a code generation tool, OpenAI's ChatGPT stands out with one of the largest and most potent models among Generative AI tools, offering four plans including a Team plan tailored for professionals working in teams of at least two users. This plan provides a dedicated workspace for team-only access. It also allows for custom models (custom GPTs) to be created and then associated and shared within team's workspace, and the latter is not possible on the lower Team plan. Custom GPTs allow for precise adaptations to software development, incorporating specific preferences, business contexts, technology stacks, or internal documents like project guidelines or coding standards. This ensures consistency in team input processing by maintaining these customisations within the fine-tuned GPT model.

When it comes to data storage and protection, the Team plan is subject to the Enterprise privacy policy.[17] It is a distinguishing feature that ensures customer content is not used as training material for OpenAI models, a practice divergent from the Plus and Free plans for individuals, which by default utilise user data for training purposes.

OpenAI does retain non-API user inputs, aiding in tracking ongoing conversation threads, and stores said conversation history on their servers with encryption at rest, but it provides mechanisms for data deletion. Any deleted conversation is to be effectively removed from OpenAI's infrastructure within 30 days.

The deletion policy extends to data submitted to fine-tuned custom GPT model. Additionally, OpenAI allows the execution of a Data Processing Addendum for ChatGPT Team users, enhancing data handling transparency.

Category	Assessment	Notes
Data storage	Positive / Warning	The Team plan stores conversations and user inputs as a conversation thread, but it permits users to delete the data.
Input safety and confidentiality	Positive	Due to the Team plan being subject to OpenAI's Enterprise privacy policy, the customer content does not serve as training material for OpenAI models.
Output copyright and licensing issues	Negative	OpenAI used a mixture of permissive and non-permissive code repositories for the training of its models, and offers no built-in filters to distinguish code that is proposed under restrictive and permissive licences. Compared to its competitors, this raises the likelihood of violating licensing agreements if the tool produces output similar to the training data.
Output ownership	Positive / Warning	The vendor declares that a user owns the output produced by the tool, but that output is not protected by intellectual property rights, at least not in the US and EU. Additionally, there is no assurance that the output will not violate third parties' IPRs.
Defence against third-party claims	Positive	OpenAI offers a „Copyright Shield" to all of its business customers, which effectively indemnifies professional users against any damages and settlement amounts payable to a third party resulting from a claim of intellectual property rights infringement. Judicious terms apply, and there is no cap on liability.

[17] OpenAI, Enterprise Privacy Policy

The Team plan is governed by OpenAI's Business terms of use[18], identical to those applied to ChatGPT's highest Enterprise plan, what consequently is its most notable differentiator. It introduces specific constraints alongside obligations, but those are deemed to be judicious.

Responsibility for the integrity and legality of input content, as well as the scrutiny of output for its accuracy and suitability, including necessitating human review, still rests with the customer. Furthermore, customers must recognise the potential for non-uniqueness in outputs, acknowledging the possibility of similar content being generated to other users as well. If these conditions are met, among a few others, a customer can profit from the indemnification provided by OpenAI against damages and settlements pertaining to third-party intellectual property rights infringement claims. Importantly, OpenAI's indemnification obligations are excluded from the damages liability cap. However, exceptions apply, e.g. if the infringement claim arises from customer's input or customer's non-compliance with the terms, applicable laws or industry standards.
It is also worth to note that, in certain scenarios, OpenAI reserves the right to settle claims independently of customer consent. Notwithstanding, this indemnity policy represents a major improvement compared to the terms that govern ChatGPT's usage under non-professional plans, where OpenAI's liability is close to none. This improvement matters greatly, given that the tool's model is trained on a diverse dataset, including content from non-permissively licensed open-source repositories like the General Public License (GPL). As a result, there is a possibility of generating code that resembles that from permissive and non-permissive sources used for training. Furthermore, ChatGPT does not provide any built-in filters, such as duplication detection or code references, which are available on tools offered by GitHub and Amazon. These measures do not eliminate the risk, but they do provide an initial line of defence.

Nevertheless, for organisations weighting the advantages and disadvantages of any given tool it is reassuring to know OpenAI now extends its so-called „Copyright Shield" over all their business customers, offering a layer of protection against copyright infringement risks.

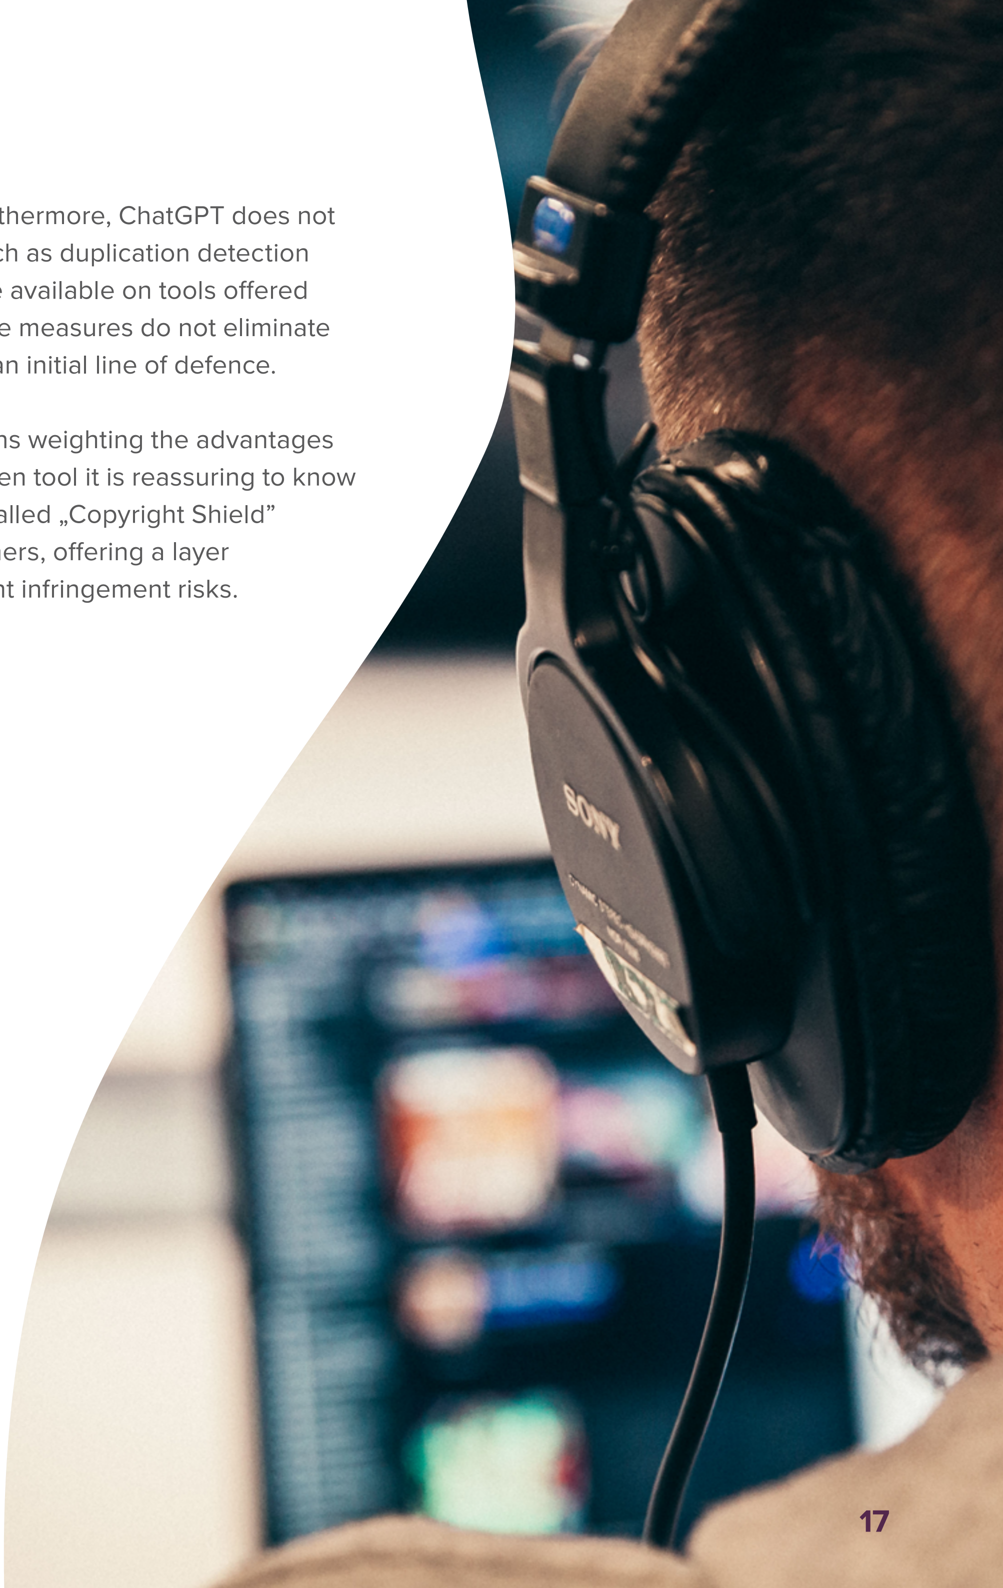

[18] OpenAI, Business Terms

Summary of Gen AI Tools Analysed

Category	Amazon CodeWhisperer Professional	GitHub Copilot Business	Tabnine Pro	ChatGPT Team
Data storage	Positive	Positive	Positive	Positive / Warning
Input safety and confidentiality	Positive	Positive	Positive	Positive
Output copyright and licensing issues	Warning	Warning	Positive	Negative
Output ownership	Positive / Warning	Positive / Warning	Positive	Positive / Warning
Defence against third-party claims	Positive	Positive / Warning	N/A	Positive

Xebia's Recommendations for Code-Generating Tools

After assessing various tools, it is clear that each offers its own benefits and challenges. The right tool depends on specific needs, methods, and risk tolerances, as each tool's potential drawbacks vary in severity.

Tabnine Pro stands out as the safest choice. Users of this tool face a minimal risk of legal issues, especially concerning intellectual property rights or software license breaches. However, some developers find the suggestions provided by Tabnine less useful than those returned by other tools.

GitHub's Copilot Business is a favourite among developers, thanks to its ease of use, flexibility, and wide range of features. Through its partnership with Microsoft, GitHub was the first to offer legal protection against third-party claims. It is a significant risk mitigation factor, however, to benefit from the indemnity, users must follow strict rules when using the output. One of the requirements prohibits alteration of the generated code. This casts doubt on the practical application of the legal protection offered. The fact that there are no reported cases of GitHub defending its customers against such claims yet provides no further comfort.

Amazon's CodeWhisperer Professional might be a good pick for businesses already using Amazon's services. It excels at integrating with AWS services and offers competitive costs due to Amazon's pricing strategies. The tool can identify potential code issues related to non-permissive open-source licences and highlight them. Above all, it is now classified as one of the Gen AI services entitled to execute Amazon's defence of claims and indemnity clause. As long as the built-in reference tracking mechanism is enabled and a few other reasonable requirements are met, Amazon promises to protect its users against third-party claims of intellectual property and cover any related expenses.

The ChatGPT Team cannot be easily compared to competitors' tools. It is not designed to streamline software development processes, and, unlike other tools, it does not provide any form of detection or filtering of generated code. It can therefore suggest code that might infringe on non-permissive open-source licences, as these were used for OpenAI's model training. However, the Team plan significantly improves the data processing and retention practices. The customer input by default does not serve as training material, rendering the Team plan suitable for commercial purposes. Additionally, customers can now benefit from indemnification against damages and settlements related to third-party intellectual property infringements offered by OpenAI, provided that specific and reasonable requirements are met.

Recommendations for Safe Adoption of Gen AI

Incorporating Generative AI tools into an organisation is a complex and layered process. Companies typically adopt one of three main strategies:

01. Delay the adoption of Gen AI until the technology matures and legislative frameworks are established, or until there is a clear consensus within the industry on risk management.

02. Fully commit to Gen AI, using the tools without significant restrictions.

03. Proceed with a cautious adoption of Gen AI, calibrating the trade-offs between efficiency gains and potential risks within their particular environment.

The selection of a strategy for Generative AI integration greatly depends on the context of the organisation, influenced by factors such as its tolerance for risk, the sector it operates within, and its strategies for growth and potential exit. Regardless of the chosen path, it is essential to equip employees with clear guidelines on the usage of Gen AI tools. Yet, strikingly, only about 16% of organisations have crafted such internal guidelines, according to a study by Opinium for Ricoh.

Postponing the implementation of Generative AI tools out of caution can impede innovation and cede advantage to competitors who exploit Gen AI's potential. This hesitation may also cause employee dissatisfaction and pose challenges in enforcing such a delay. Moreover, there's a real risk that employees might sidestep such restrictions, a concern underscored by Opinium's research for Ricoh showing that 48% of Europe's workforce already employ Gen AI tools in their roles, with 18% incorporating them into their daily activities.

Consequently, while this cautious approach may seem to mitigate immediate risks, it must be weighed against the possibility of losing a competitive edge. Embracing Gen AI tools without restrictions is an approach that does not suit most businesses. Allowing employees to independently use these tools can significantly raise the organisation's risk concerning confidentiality and intellectual property rights. Although the likelihood of such risks is generally low, the consequences could be highly detrimental and expensive. For most organisations, the optimal strategy is to calibrate the rate and extent of Gen AI tools adoption to align with their operational context. Accomplishing this requires establishing explicit guidelines that define the permissible tools and their appropriate usage, alongside providing thorough communication and training for staff members.

At Xebia, we have adopted the latter approach by implementing guidelines and developing training programs on Gen AI for our employees. We started by categorising Gen AI applications into two main types, each with its own set of rules:

01. **Inspiration and Learning:** Under this category, the use of proprietary data as input is strictly forbidden, and any outputs produced by Gen AI cannot be used as part of any project deliverables.

02. **Analysis and Artifact Creation:** This category is reserved for approved tools that meet our high-security benchmarks for commercial use. Regardless of the tool used, it is mandatory for the outputs to undergo a rigorous peer-review process to ensure accuracy and compliance.

We allow our employees to utilise any Gen AI tool for inspiration and learning purposes. However, when it comes to processing and creating materials that are proprietary, only a vetted selection of tools is authorised. Additionally, there are stringent guidelines on the proper configuration of these tools.[19, 20, 21]

As organisations navigate the complexities of integrating Gen AI into their operations, they might face hurdles similar to the ones we have encountered. In light of this, we suggest a suite of strategies for consideration:

01. Develop a clear policy outlining the permissible uses of Gen AI within your organisation, ensuring it remains current with evolving legislation and technological advancements. Align this policy with your organisation's specific context and the level of risk deemed acceptable, particularly regarding confidentiality and intellectual property rights.

02. Train your employees to foster an understanding of the advantages, potential risks, and the company's stance on Gen AI tools. Embed this training within the onboarding process to guarantee that new staff are knowledgeable from the start. Maintain regular communication to keep all employees aware of any shifts in the technological landscape or company policy.

03. Exercise discernment when selecting Gen AI tools, steering clear of potential traps related to licensing and configuration. These details can be crucial. Weigh the benefits of Software as a Service (SaaS) against on-premises solutions, according to what is most suitable for your needs.

04. Assure that the output from Gen AI tools aligns with your organisation's established quality benchmarks and protocols. This can be achieved by setting up rigorous review mechanisms and tailoring the delivery pipeline accordingly.

[19] GitHub, Configuring GitHub Copilot settings
[20] Amazon, Configuring Amazon CodeWhisperer Professional settings
[21] Tabnine, Configuring Tabnine Pro telemetry settings

Summary

Generative AI has the potential to significantly boost productivity in software development.

Although it poses certain risks, these can generally be reduced to manageable levels with the right implementation. For organizations in fast-paced and competitive fields, the cost of delaying its adoption could be substantial.

Xebia

About Xebia

Xebia is an IT Consultancy and Software Development Company that has been creating digital leaders across the globe since 2001. With offices on every continent, we help the top 250 companies worldwide embrace innovation, adopt the latest technologies, and implement the most successful business models. To meet every digital demand, Xebia is organized into chapters. These are teams with tremendous knowledge and experience in Agile, DevOps, Data and AI, Cloud, Software Development, Security, Quality Assurance, Low Code, and Microsoft Solutions. In addition to high-quality consulting and state-of-the-art software, Xebia Academy offers the training that modern companies need to work better, smarter, and faster. Today, Xebia continues to expand through a buy and build strategy. We partner with leading IT companies to gain a greater foothold in the digital space.

Find more information on how Xebia is driving innovation at xebia.com.