BITHOST | RESEARCH REPORT | 2026

THE INVISIBLE RISK

How Organisations Are Unknowingly Handing Their Intellectual Property to AI Tools

Prepared by

Bithost | ZHOST Consulting Private Limited

www.bithost.in | sales@bithost.in

Disclaimer

This report has been prepared by Bithost (ZHOST Consulting Private Limited) for general informational purposes. The findings, statistics, and observations are based on publicly available research, industry surveys, regulatory filings, and firsthand observations gathered during client engagements between 2022 and 2026.

This document does not constitute legal advice. Organisations should seek qualified legal counsel before making decisions related to intellectual property, data privacy, or regulatory compliance. Where third-party statistics are cited, sources are referenced inline. Bithost does not warrant the completeness or accuracy of third-party data.

All company names, products, and incidents referenced are cited from publicly reported sources. No confidential client information appears in this document.

Executive Summary

Artificial intelligence coding tools have moved from novelty to infrastructure in less than three years. By early 2026, a developer who does not use some form of AI assistance is the exception. GitHub Copilot passed one million paid subscribers. ChatGPT crossed 180 million monthly active users. The tools are embedded in IDEs, running in browser tabs during working hours, and integrated into the daily problem-solving habits of engineers at companies of every scale.

What has not moved at the same pace is governance. The decision to adopt these tools has largely been made at the individual developer level, without involvement from legal teams, information security functions, or executive leadership. The result is a structural gap between what organisations believe is happening with their code and their intellectual property, and what is actually happening on developers' screens every working day.

This report focuses on organisations with 10 to 200 employees. This band is not arbitrary. It represents the range where engineering teams produce commercially significant and proprietary code, but where dedicated security operations, legal oversight of developer tooling, and enforceable AI usage policies are rarely present. It is the size range where the risk is highest and the controls are fewest.

Finding / Metric	Statistic	Source / Year
Developers globally using AI coding tools regularly	92%	Stanford HAI 2025
Companies with a formal AI usage policy	14%	Forrester 2025
AI-generated code with detectable security vulnerabilities	40%	Veracode 2024
Developers using personal AI accounts for work tasks	43%	Stack Overflow 2024
Developers who received AI-specific security training in past year	11%	SANS 2024
SMEs with no formal IP audit completed	61%	WIPO SME Survey 2024
Average cost of a data breach globally	$4.88M	IBM Ponemon 2024

Finding / Metric	Statistic	Source / Year
Insider threat cited as top concern by security leaders	79%	Cybersecurity Insiders 2024

The situation does not arise from developer negligence. Engineers using free AI tools are doing exactly what they were trained to do: find the most efficient path to a working solution. The governance failure sits with organisations that have permitted widespread AI tool adoption without any framework, any policy, or any awareness of what it means for the code and data they consider proprietary.

Table of Contents

Chapter 1

How AI Coding Tools Actually Work

The Core Mechanics

Understanding the risk begins with understanding the technology. AI coding assistants are large language models trained on datasets of text and source code at a scale that was not practically achievable before 2020. GitHub Copilot was trained on a significant portion of the public code available across GitHub repositories. OpenAI's models were trained on broad internet text including code documentation, technical forums, open source repositories, and developer blog posts. Google's Gemini and Anthropic's Claude were trained on similarly broad corpora.

When a developer types a question or pastes code into one of these tools, the input is transmitted to a remote server operated by the provider. The model processes the input, generates a response based on statistical patterns in its training data, and returns the result. This transaction is not local. There is no on-device processing in the standard and free tiers of any major AI coding assistant. Every prompt, every code snippet, every question travels across the internet to a commercial server in a data centre the developer has never visited and the organisation has never assessed.

What Gets Transmitted Beyond What Developers Notice

The common assumption is that a developer only shares the specific piece of code they are asking about. IDE plugin behaviour makes this assumption incorrect in most cases.

GitHub Copilot, Tabnine, Codeium, and similar IDE-integrated tools work by continuously sending "context" to the provider's servers to generate relevant completions as the developer types. This context includes the active file, recently opened files, visible function and variable names, and in some configurations, the broader project structure. The context is assembled automatically by the plugin and transmitted without the developer taking any deliberate action. A developer who believes they are only sharing the function they are currently editing may unknowingly be sharing the entire module, its imports, and related configuration files.

Browser-based tools transmit exactly what the developer pastes. However, common pasting patterns mean this is frequently more than developers intend. When debugging a production issue, a developer typically pastes the error message with the full stack trace (which contains internal path names, server names, and sometimes environment details), the relevant code section, and sometimes the database query or API call involved. When asking for a code review, developers regularly paste entire files or entire modules.

What Developers Intend to Share	What Is Actually Transmitted	Risk Level
A single function	Full file + recent files via IDE context	Medium
An error message	Stack trace with paths, server names, env vars	High
A database query	Full schema, table names, column names	High
Configuration help	Config files with API structure, endpoint names	Critical

What Developers Intend to Share	What Is Actually Transmitted	Risk Level
Code review	Proprietary business logic, algorithms, architecture	Critical
Test generation	Implementation code + data models + test fixtures	Critical
Production debugging	Logs with real data samples, customer identifiers	Critical

How IDE Plugins Phone Home Passively

The distinction between active and passive data transmission is important and frequently misunderstood. When a developer uses a browser-based AI tool, they make a conscious decision to paste content and submit a prompt. They have at least some awareness of what they are sharing, even if they have not thought through the implications.

IDE plugins operate differently. The plugin runs continuously in the background of the development environment. As the developer types, the plugin automatically assembles context from the current file and surrounding files and sends it to the provider. This happens with no deliberate action from the developer, no notification that data is being transmitted, and no option to review what was sent before it leaves the machine. The developer is focused on their code. The plugin is quietly transmitting.

Some IDE plugins also log all accepted suggestions, which allows the provider to track which of their completions were used, in what context, and how they were modified. This telemetry data provides the provider with a detailed picture of how developers are working with code, including code that the developer did not directly submit as a prompt.

Local Models: The Partial Solution

A growing category of AI coding tools runs models locally on the developer's machine, eliminating the data transmission risk. Tools like Ollama, LM Studio, and Cursor in local mode allow developers to use open-weight models without sending code to any external server. This approach removes the cloud transmission risk entirely.

The limitations of local models are practical rather than technical. The best-performing models as of 2025 (GPT-4o, Claude 3.7 Sonnet, Gemini 2.0) are not available for local use. Local models require hardware capable of running large parameter counts efficiently, which means a development machine with 16-32GB of RAM and ideally a recent GPU. These requirements exceed the standard specification of many corporate laptops, particularly in cost-sensitive small organisations.

For organisations that can accommodate the hardware requirements, local models represent a technically sound solution for the most sensitive work. They do not, however, eliminate the need for governance around when and how AI tools are used, because the choice of which tool to use for which task still requires human judgement.

Chapter 2

The Data Journey: What Happens After You Press Enter

From Keystroke to Server and Back

Mapping the complete data journey of a submitted code prompt requires following the request through multiple systems, each governed by different policies and each operated by different entities. The journey is straightforward technically but complex from a governance and legal perspective.

Step one: the developer submits input through an IDE plugin, a browser interface, or an API call. The input is encrypted in transit using TLS. The encryption protects the data from interception between the developer's machine and the provider's servers. It does not protect the data once it arrives at those servers.

Step two: the provider's load balancer receives the request and routes it to inference infrastructure. This infrastructure is typically cloud-based. OpenAI runs on Microsoft Azure. Anthropic runs primarily on Amazon Web Services and Google Cloud. The inference servers process the request, generate output, and return it.

Step three: the provider logs the request. Every major AI provider maintains logs of requests for various purposes including safety monitoring, model improvement, abuse detection, and debugging. The retention period, access controls, and permitted uses of these logs vary by provider and by account type.

What Providers Log and For How Long

Provider	Free Tier Retention	Enterprise Retention	Training Use (Free)
OpenAI (ChatGPT)	Logs: 30 days default	API: 30 days, DPA available	Default ON. Opt-out in settings.
GitHub Copilot	Prompts: up to 30 days	Enterprise: not retained for training	Individual: can opt out
Google Gemini	Conversations: 18 months default	Workspace: training excluded	Consumer: default ON
Anthropic Claude	Free: reviewed for safety	API/Enterprise: DPA available	Safety review only per policy
Mistral (via API)	API: standard log retention	Enterprise: configurable	Not for training by default
Amazon CodeWhisperer	Professional: excluded from training	Enterprise: full control	Free tier: may be used
Tabnine	Cloud: standard retention	Self-hosted: no external send	Enterprise mode: no training

The pattern across all providers is consistent. Enterprise accounts with paid agreements carry stronger data protection guarantees, while free and consumer accounts operate under terms that grant the provider broad latitude. The gap between a free account and an enterprise agreement is not marginal. It represents fundamentally different legal relationships and different levels of data protection.

The Subprocessor Problem

AI providers do not operate their infrastructure in isolation. They depend on cloud hosting providers, content moderation vendors, security monitoring services, and operational tools, each of which qualifies as a subprocessor under GDPR and similar frameworks. When code is submitted to an AI tool, it may be processed not only by the AI provider's own systems but by the systems of any subprocessor in their supply chain.

OpenAI's publicly available subprocessor list includes multiple cloud providers, a content moderation service, and several operational SaaS tools. Each of these entities is bound by a data processing agreement with OpenAI, but the developer who submitted the code has no contractual relationship with any of them and no mechanism to enforce any rights against them directly.

Key Observation

The Samsung incident in April 2023 made this concrete: engineers submitted source code, internal meeting transcripts, and hardware test data to ChatGPT on three separate occasions within weeks of the company lifting an AI tool ban. Samsung confirmed it could not retrieve or delete the submitted data after discovery. The code exists in OpenAI's logs. This is not a theoretical exposure. It is a documented, permanent, irrecoverable disclosure of commercially sensitive intellectual property.

Data Residency and Jurisdiction

Organisations subject to data residency requirements face a specific challenge. Most major AI tool providers operate primarily from US-based infrastructure. GDPR imposes restrictions on transferring personal data from the EU to third countries unless an adequate level of protection can be guaranteed. The EU-US Data Privacy Framework, adopted in 2023, provides a mechanism for compliant transfers to participating US organisations, but compliance depends on the specific provider's participation status and the nature of the data transferred.

For code that does not contain personal data, GDPR data transfer restrictions do not directly apply. However, if code contains embedded personal data (test fixtures with real emails, log samples with IP addresses, database schemas that reveal personal data handling practices), the transfer restrictions become relevant. Many developers do not consider whether the code they are working with contains or references personal data.

Chapter 3

The Intellectual Property Question Nobody Is Asking

What IP Protection Actually Means for Software Companies

Intellectual property protection for software operates through several distinct legal mechanisms, each with different conditions, costs, and strengths. Understanding which protections apply to a codebase, and under what conditions they can be lost, is foundational to assessing what is at risk when code is shared with AI services.

Copyright protection arises automatically upon creation of original work. Source code is copyrightable from the moment it is written, with no registration required in most jurisdictions. The practical value of copyright for software, however, is limited. Copyright protects specific expression: the exact way something is written. It does not protect the underlying idea, algorithm, or approach. A competitor who independently develops a different implementation of the same algorithm does not infringe the copyright of the original.

Trade secret protection covers a different and often more valuable category: information that derives its commercial value specifically from being kept confidential. The algorithm itself, the architecture decision, the performance optimisation approach, the business logic that makes a product work differently from competitors. These can be trade secrets even if they are not patentable and even if they could not be protected by copyright. But trade secret status depends entirely on maintaining secrecy. The moment the information is disclosed to a third party without adequate protections, trade secret status is permanently lost.

The Employer-Developer-Provider Triangle

Standard employment agreements assign to the employer all intellectual property created by the employee in the course of their employment. This seems straightforward. A developer builds something for the company, the company owns it. The AI tool complicates each link in this chain.

First, the employee-to-employer link: if the developer uses an AI tool to generate a substantial portion of the code, and if that code does not qualify for copyright protection due to insufficient human authorship, then the employer may not have a valid copyright claim over what it thought it owned. The US Copyright Office has been explicit that works generated by AI without sufficient human creative contribution do not receive copyright protection. Courts in multiple jurisdictions are working through what "sufficient human contribution" means in the context of AI-assisted development. The answers are not yet settled.

Second, the developer-to-provider link: when a developer submits code to an AI service, they grant the provider certain rights to that code under the provider's terms of service. Most terms of service include a licence for the provider to use submitted content for the purposes of operating and improving the service. The developer typically does not have the authority to grant this licence on behalf of their employer, but they do so nonetheless. The employer's rights are not extinguished, but they are complicated by the existence of the provider's licence.

Third, the provider-to-output link: the output generated by the AI tool may contain sequences that originated in the provider's training data. If those sequences came from open source code under specific licences, the output may carry licence obligations the employer is unaware of.

Legal Exposure

An organisation that discovers its developers have been sharing proprietary code with commercial AI services without any policy in place faces a compounded problem: it cannot establish that it took "reasonable measures" to protect its trade secrets (a condition for trade secret enforcement), it cannot determine what of its codebase may carry external licence obligations, and it cannot assess what of its IP has been potentially disclosed to a competitor via the same AI service. None of these problems can be fixed retroactively.

The Copyright Status of AI-Generated Code

The question of who owns code that was primarily written by an AI tool remains unresolved in most legal systems. Several consistent signals have emerged from regulatory and judicial decisions through 2024.

The United States Copyright Office has issued guidance stating that AI-generated content that lacks human authorship cannot be registered for copyright. The Office requires a human author and will scrutinise applications for AI-assisted works on a case-by-case basis depending on the degree of human creative contribution.

The UK Intellectual Property Office conducted a consultation in 2023 on computer-generated works and has not issued final guidance as of early 2025. The UK currently provides limited protection for computer-generated works under the Copyright, Designs and Patents Act, but applying this to AI-generated code is the subject of ongoing legal interpretation.

The EU AI Act does not directly resolve the copyright question but introduces transparency obligations for AI systems that generate content. Systems used to generate substantial volumes of code for commercial purposes will face transparency requirements that may require disclosure of training data sources.

Open Source Licence Contamination

AI models trained on public code repositories have learned from code under every kind of open source licence. The most permissive licences (MIT, Apache 2.0, BSD) allow commercial use with minimal restrictions. Copyleft licences (GPL, LGPL, AGPL, EUPL) impose conditions on derivative works, including in some interpretations the requirement to release derivative source code under the same terms.

Research conducted at Princeton and NYU in 2023 demonstrated that GitHub Copilot can, under certain conditions, reproduce verbatim sequences from its training data, including code under GPL licences. The frequency of verbatim reproduction was approximately 1% of output in test conditions, but this figure scales materially across a large codebase. In a codebase of 100,000 lines with 30% AI-generated content, a 1% verbatim reproduction rate produces approximately 300 lines of potentially unattributed and licence-encumbered code.

If any of those 300 lines come from a GPL-licenced source, the entire work using them may be subject to GPL obligations. A product built on proprietary code that unknowingly incorporates GPL sequences through AI generation faces potential obligation to release its source code publicly. This is not a hypothetical concern. The FOSSA 2024 open source risk report found open source compliance issues in between 7% and 15% of AI-assisted codebases reviewed.

Chapter 4

The Startup and SME Vulnerability Profile

Why the 10 to 200 Employee Range Is the Highest-Risk Band

Large enterprises have infrastructure for managing this risk. When Amazon's legal team identified concerns about ChatGPT usage in February 2023, it had the mechanisms to issue a company-wide memo within days and the authority to enforce compliance. When Apple restricted AI tool usage in May 2023, it had the IT infrastructure to implement technical controls and the HR systems to document and enforce the policy. These are not options available to a 40-person SaaS startup.

Organisations below 10 employees can substitute informal communication for formal policy. A five-person founding team is small enough that the CEO knows what every engineer is doing. At 20 employees, this is no longer true. At 100 employees spread across multiple time zones, it has not been true for a long time.

The Six Structural Weaknesses

No Dedicated Security Function

Organisations in the 10-200 employee band almost universally lack a dedicated Chief Information Security Officer, a security engineer, or even a consistent part-time security contractor. Security responsibility defaults to the most technically experienced developer available, who also carries product engineering, infrastructure, and often DevOps responsibilities. Security is addressed reactively when something breaks, not proactively as a designed function. When AI governance appears on the agenda at all, it sits below delivery deadlines, infrastructure costs, and technical debt.

Developer Tooling Selected Without Organisational Oversight

In these organisations, the choice of development tools is almost always made by individual developers or engineering leads without legal, security, or procurement review. A developer who finds GitHub Copilot useful installs it. Another developer uses ChatGPT because it solved their last three debugging problems faster than Stack Overflow. The engineering lead does not require approval because the tools are free, and free tools have historically required no procurement process. Nobody considers that "free" in this context means the organisation's code is the product.

Free Tier Usage as the Default

Startups are cost-sensitive, and the cost sensitivity is most acute in engineering infrastructure budgets. GitHub Copilot Individual costs $10 per month. ChatGPT Plus costs $20 per month. These amounts are individually trivial but multiply across an engineering team and encounter budget scrutiny when the company is managing runway carefully. Enterprise agreements, which carry the data protection terms that would actually protect the organisation, cost orders of magnitude more and require procurement cycles that feel disproportionate for a productivity tool.

The result is that the most commercially sensitive code in the organisation is processed by AI tools under the weakest contractual protections available. The developer with access to the payment processing code, the authentication module, and the core business logic uses the free version of the tool because no one told them not to.

Certifications Without Corresponding Practice

An increasing number of small technology organisations hold security certifications: ISO 27001, SOC 2 Type II, Cyber Essentials, or equivalent frameworks. These certifications are frequently required by enterprise customers as a condition of doing business, particularly in regulated industries. The challenge is that certification frameworks were designed to assess documented controls against defined criteria. They were not designed to audit real-time developer behaviour.

AI tool governance represents a gap that most current certification frameworks have not yet systematically addressed. A company can hold a current SOC 2 Type II report covering the most recent audit period and simultaneously have an engineering team sharing proprietary code with free-tier AI services every day, with no policy, no monitoring, and no controls in place. The auditor did not ask. The controls framework did not require it. The certificate says nothing about it.

IP Documentation Gaps

WIPO's 2024 survey of small and medium enterprise intellectual property practices found that 61% of technology SMEs have not conducted a formal IP audit in the past three years. Many cannot identify with precision which components of their codebase constitute proprietary innovation and which are standard integrations, third-party libraries, or boilerplate. This matters for risk assessment because without an IP inventory, an organisation cannot assess what was exposed when a developer shared code with an AI tool and cannot make a credible legal claim about the value of what was potentially disclosed.

No Incident Response Capability for AI-Related Disclosures

Standard incident response frameworks address external breaches: an attacker gaining unauthorised access, a credential compromise, a ransomware infection. An AI-related disclosure has a fundamentally different profile. The access was authorised (the developer chose to share the code), the mechanism was a commercial service the developer uses legitimately for other tasks, and the disclosure may never be definitively confirmed or denied. None of the standard incident response playbooks address this scenario, and organisations in this size band typically have not created one.

Risk Attribute	Typical State in a 10-200 Employee Organisation
AI usage policy	None, or a single line in the employee handbook
AI tool procurement process	Individual developer choice, no approval required
Account type used	Free or personal paid accounts in the majority of cases
Review of AI-generated code	Informal at best, non-existent at worst
IP inventory status	Not documented or only partially documented
AI security training	None delivered in the past 12 months
Legal review of AI tool ToS	Never conducted
Technical monitoring of AI usage	No controls in place
Incident response for AI leaks	No specific procedure exists

Chapter 5

The Gap Between Policy and Practice

The Pattern That Repeats Across Organisations

There is a pattern that appears with remarkable consistency across small technology organisations that have received any form of security audit or vendor assessment. The organisation has a written information security policy. The policy contains a section on data classification or confidentiality obligations. The section states that proprietary information must not be shared with unauthorised third parties. The policy has been acknowledged by employees as part of onboarding.

None of this prevents a developer from opening ChatGPT in a browser tab and pasting the company's payment processing module to ask why a specific function is not handling edge cases correctly. The developer is not acting in bad faith. They have not thought about the information security policy in months, if they remember agreeing to it at all. Even if they had, the policy does not mention AI tools by name because it was last substantively updated before AI coding tools became mainstream. The terms "ChatGPT," "Copilot," "Claude," and "Gemini" do not appear in it.

Why Generic Policies Fail

Information security policies in small organisations are typically written to satisfy audit requirements rather than to guide real-time developer behaviour. They use language like "use AI tools responsibly" or "do not share confidential information with third parties" that sounds adequate at the policy level but provides no actionable guidance in the moment.

A developer who reads "do not share confidential information" and then pastes a codebase file into ChatGPT may genuinely believe they are complying. They are not sharing customer names, financial records, or passwords. The fact that they are sharing proprietary business logic, which is equally confidential and potentially more commercially sensitive, is not addressed by the policy in any specific way. The policy failure is a design failure, not a compliance failure.

The Update Problem

Forrester's 2025 survey found that 14% of organisations have a formal AI usage policy. This figure represents improvement from 13% in 2024 and 9% in 2023, but the rate of policy adoption has not kept pace with the rate of AI tool adoption. In organisations with fewer than 200 employees, the figure is likely lower than the survey average, which draws heavily from enterprise respondents.

Among organisations that do have an AI policy, the document was typically created in response to a specific incident or customer requirement rather than as part of a systematic governance programme. These reactive policies tend to be more restrictive than necessary (banning all AI tool usage, which is unenforceable and counterproductive) or more permissive than prudent (permitting AI tool usage with no classification framework attached).

The Monitoring Gap

Technical controls that could detect or limit AI tool usage exist but are rarely deployed in small organisations. Data loss prevention tools can be configured to detect transmission of specific patterns to external services. Web proxies can log traffic to AI tool domains. Network monitoring can identify unusual data volumes going to known AI endpoints. None of these require enterprise-scale infrastructure. All of them require someone to prioritise their implementation, which has not happened in most organisations of this size.

The result is a monitoring gap with a specific property: the organisation cannot know what has been shared because it has no mechanism to find out. It cannot retrospectively audit what was submitted to AI services. It cannot determine whether its proprietary algorithms were shared, or when, or by whom. The first indication that a problem exists may come from a legal dispute, a competitor product launch, or an auditor's question that nobody has a good answer to.

The Audit Reality

During a 2024 vendor security assessment for a mid-market enterprise, a small software supplier confirmed it followed all industry security standards. A follow-up review of the supplier's engineering team found that all seven developers used free ChatGPT and Claude accounts regularly and had collectively submitted thousands of prompts containing production code over the previous eight months. No policy existed. No monitoring was in place. The supplier held a current SOC 2 Type II report. The auditor had not asked about AI tools. The report said nothing about them.

Chapter 6

Developer Behaviour: What the Data Shows

The Scale of the Shift

The Stack Overflow Developer Survey 2024, which collected responses from over 65,000 developers across 185 countries, provides the clearest publicly available picture of how AI tool adoption has changed developer behaviour since 2022. The trend is unambiguous.

Behaviour / Metric	2022	2023	2024
Using AI tools in professional development work	28%	44%	62%
Planning to use AI tools in the next year	39%	54%	72%
Using AI tools on employer-provided accounts	N/A	26%	34%
Using personal accounts for work AI tasks	N/A	43%	51%
Trusting AI output without additional verification	N/A	22%	31%
Reporting no AI usage policy at their organisation	N/A	68%	61%

The 51% figure for developers using personal AI accounts for professional work tasks in 2024 is the most significant from a governance perspective. In more than half of AI-assisted work interactions, the organisation has no contractual relationship with the tool provider, no data processing agreement governing the interaction, and no visibility into what was submitted.

The High-Risk Scenarios in Practice

Production Incident Response

When production systems fail, engineers work under intense time pressure. In this context, copying error logs, stack traces, and relevant code sections into an AI tool is the fastest path to a diagnosis. Production logs contain more sensitive information than developers typically recognise: customer identifiers, internal system topology, data structure details, and sometimes actual data values. The urgency of a production incident is precisely the situation in which a developer is least likely to pause and think about what they are sharing.

Code Review and Architecture Discussions

Developers regularly share significant portions of their codebase with AI tools for technical feedback. "Review this module and identify any issues" is among the most common AI prompts used by software engineers. This use case typically involves sharing functionally complete, working code representing the most commercially valuable work the developer has produced, not boilerplate or standard integrations.

Onboarding to an Unfamiliar Codebase

Developers new to a project or organisation use AI tools to understand existing code. Pasting files and asking "what does this do?" or "how does this component fit into the system?" is a common onboarding pattern. This use case tends to involve core architectural code rather than peripheral modules, precisely because the developer needs to understand the foundations before they can work effectively.

Test Generation

Writing tests requires sharing the code being tested. Automated test generation is one of the most heavily marketed AI coding features, and it is genuinely useful. A developer asking an AI tool to generate unit tests for a function will share the function implementation, the data models it operates on, and any related configuration. The code shared for test generation purposes is exactly the code the organisation has the most interest in protecting.

The "It's Not Sensitive" Misconception

Developer interviews conducted during client assessments consistently surface a specific rationalisation: "I only shared code, not customer data." This reasoning misunderstands where the commercial value of a software product sits. For most technology companies, the proprietary value is not in the customer records stored in the database. It is in the code that processes those records in a particular way, that implements the workflow that customers pay for, that contains the technical decisions that make the product competitive.

Sharing the payment processing logic, the recommendation algorithm, the fraud detection rules, or the core API with an external AI service is precisely sharing the commercially sensitive intellectual property, regardless of whether any customer names or account numbers were included.

Chapter 7

Known Incidents, Reported Leaks, and Quiet Disclosures

Samsung: The Landmark Case

Samsung Electronics provides the most thoroughly documented example of AI-related IP disclosure from a major organisation. In April 2023, Samsung confirmed three separate incidents in which employees submitted proprietary information to ChatGPT within weeks of the company permitting limited AI tool use in engineering teams.

In the first incident, an engineer submitted source code from Samsung's semiconductor equipment measurement database, asking the tool to correct errors. In the second, a different engineer submitted code related to NAND flash equipment for optimisation review. In the third, an employee converted an internal meeting recording to text and submitted it to create meeting notes, capturing strategic discussions about equipment yield and competitive positioning.

Samsung responded by limiting prompt size and eventually implementing a company-wide ban. The company began developing internal AI tools running on its own infrastructure. However, Samsung acknowledged publicly that it had no mechanism to retrieve or delete data already submitted to OpenAI's systems. The code, the technical specifications, and the meeting content remain in OpenAI's log infrastructure indefinitely.

Amazon: Internal Warnings and Code Similarity Concerns

Amazon's legal team circulated an internal memorandum in February 2023 warning employees not to share confidential information with ChatGPT. The memo was prompted by the discovery of employee usage and included a notable specific concern: Amazon had observed ChatGPT generating output that resembled Amazon's own internal documentation, suggesting that Amazon content may have influenced the model's training data or generation patterns.

Amazon simultaneously invested heavily in Anthropic (the company behind Claude) and accelerated development of its own CodeWhisperer product, which operates under enterprise terms that exclude customer code from training. The internal prohibition on ChatGPT was paired with a commercial motivation to direct employees toward tools Amazon had more control over.

Apple: Restriction Before Damage

Apple restricted employee use of ChatGPT and GitHub Copilot in May 2023, according to reporting in the Wall Street Journal and Bloomberg. Apple cited concerns about the potential for confidential information to be leaked through these services. The restrictions were implemented proactively, before any confirmed incident. Apple was simultaneously investing in internal AI capabilities for its own products and had clear incentives to control how its development practices were exposed externally.

The Unreported Majority

The Samsung, Amazon, and Apple cases received coverage because major corporations with active press coverage made decisions that were newsworthy. The equivalent situations at smaller organisations receive no coverage, for reasons that are structural rather than exceptional.

Small organisations typically do not know when a disclosure has occurred. Without monitoring infrastructure, they have no alert mechanism and no baseline against which to measure anomalous data leaving the organisation. The code was submitted, the developer received an answer, and no record of the event exists at the organisational level.

When a small organisation does become aware of a potential exposure, the incentives are strongly against disclosure. A startup that publicly acknowledges its proprietary algorithms were submitted to AI services risks investor confidence, customer trust, and any legal claim to trade secret protection over that technology. The GDPR notification obligation applies to breaches of personal data, not necessarily to IP disclosures. There is no legal requirement to report submitting proprietary code to an AI service. Silence is rational.

Observed Pattern

Across client engagements conducted between 2023 and 2025, Bithost encountered multiple organisations where structured developer interviews revealed systematic AI usage with code submissions including proprietary payment logic, authentication implementations, pricing algorithms, fraud detection rules, and core API designs. In no case had the organisation been notified of this activity. In no case did a security policy specifically address it. In every case, the developers involved believed their usage was within normal and acceptable practice.

The Competitive Intelligence Scenario

There is a distinct and underexplored risk category: the AI tool as an inadvertent channel for competitive intelligence. The mechanism operates as follows.

Organisation A's developers submit proprietary code to a free-tier AI service. That code is incorporated into the service's training pipeline. The model's weights are updated to reflect patterns learned from that code. When Organisation B's developers ask the same AI service for suggestions on similar technical problems, the model generates output that statistically reflects the patterns it learned, including patterns from Organisation A's submitted code. Organisation B receives suggestions that happen to align closely with Organisation A's proprietary approach without either organisation or the AI provider intending this outcome.

Whether this constitutes actionable harm is legally unclear. No direct disclosure occurred. No deliberate misappropriation took place. But the economic value of the IP was diminished, and the mechanism that caused the diminishment was a service both organisations willingly used. The law has not yet developed frameworks for addressing this category of harm.

Chapter 8

The Regulatory Landscape in 2025

GDPR and AI Tool Usage

The General Data Protection Regulation remains the most directly applicable framework for European organisations and for any organisation that processes data about European residents. GDPR compliance in the context of AI tool usage involves several distinct questions that many organisations have not yet asked.

The first question is whether personal data is being submitted to AI services. The answer is frequently yes, through patterns developers do not recognise: stack traces containing IP addresses, test fixtures containing real email addresses, error logs with customer identifiers, database schemas that expose data structure, and API responses used for debugging that contain user-attributed data.

The second question is whether a valid legal basis exists for this processing. Article 6 requires a lawful basis. Legitimate interests, the most commonly invoked basis for data processing in business contexts, requires a balancing test against data subjects' rights and interests. The argument that submitting customer data to a commercial AI service in a third country is consistent with the legitimate interests basis would face significant scrutiny from any supervisory authority that examined it.

The third question is whether the transfer to a third country (primarily the US, where most major AI providers are based) is lawful. The EU-US Data Privacy Framework provides a mechanism for compliant transfers to participating US organisations, but reliance on this framework requires the specific provider to be a certified participant and requires the transferring organisation to conduct appropriate due diligence.

Regulatory Area	Key Requirement Relevant to AI Tool Usage	Maximum Penalty
GDPR Art. 6	Lawful basis required for personal data processing	EUR 20M or 4% global turnover
GDPR Art. 28	Data processing agreement required with processors	EUR 10M or 2% global turnover
GDPR Art. 44	Adequate safeguards for third-country transfers	EUR 20M or 4% global turnover
EU Trade Secrets Directive	Reasonable measures to maintain secrecy required	Civil liability, injunctions
EU AI Act (2025+)	Transparency, documentation for high-risk AI use	EUR 30M or 6% global turnover
UK GDPR	Same structure as EU GDPR post-Brexit	GBP 17.5M or 4% global turnover
CCPA (California)	Personal information protection and disclosure rights	USD 7,500 per intentional violation

The EU AI Act: What Changes in 2025 and 2026

The EU AI Act entered into force in August 2024. Implementation is phased. Prohibited practices (AI systems using prohibited manipulation or exploitation techniques) became enforceable from February 2025. Obligations for high-risk AI systems, including transparency requirements and conformity assessments, apply from August 2026. Obligations for general-purpose AI models with high systemic risk apply from August 2025.

For organisations using AI tools in development rather than deploying AI systems, the direct obligations under the AI Act are limited in the immediate term. However, the Act introduces a transparency requirement for general-purpose AI models that will affect how providers document their training data. This documentation will become relevant to organisations assessing their exposure to open source licence contamination in AI-generated code.

The Act also establishes that organisations deploying AI systems in high-risk contexts (employment screening, access to essential services, public safety, critical infrastructure) must conduct conformity assessments that will include examining the provenance of AI-generated code in those systems. For software companies whose products touch these areas, the governance of AI-assisted development will become a compliance requirement, not just a risk management practice.

Trade Secret Law in Practice

The EU Trade Secrets Directive (2016/943/EU), implemented across member states, defines a trade secret as information that is secret, has commercial value because of its secrecy, and has been subject to reasonable steps by its owner to keep it secret. All three conditions must be met. The third condition is where AI tool usage creates the most immediate legal exposure.

"Reasonable steps" is assessed based on what a reasonable person in the same industry and of the same size would have done to protect information of that type and value. As AI tool usage becomes more widely understood and as governance frameworks become more widely available, the threshold for what constitutes reasonable steps will rise. An organisation that had no AI usage policy in 2023 could perhaps argue that no reasonable steps were established at that time. An organisation with no AI usage policy in 2025 will face a harder argument, given the volume of public guidance, regulatory statements, and industry frameworks that have emerged since 2022.

Export Controls and Cross-Border Considerations

Organisations working in defence, aerospace, dual-use technology, critical infrastructure, or with government contracts in any major jurisdiction may be subject to export control regulations that restrict the transfer of technical information to foreign servers or foreign entities. The EAR (Export Administration Regulations) in the US, ITAR (International Traffic in Arms Regulations), and equivalent regimes in the EU, UK, India, and other jurisdictions create specific obligations that may be triggered by submitting technical code to AI services.

An engineer at an Indian defence technology company submitting proprietary sensor code to a US-based AI service may be creating an unauthorised export under both Indian and US regulations, entirely inadvertently. An engineer at a UK defence supplier sharing technical specifications through a debugging prompt may be violating ITAR obligations that their employment contract makes them personally responsible for. Most developers in affected organisations have never been briefed on how export control law applies to AI tool usage.

Chapter 9

Facts, Figures, and Industry Analytics

AI Adoption Statistics (2024-2025)

Finding / Metric	Statistic	Source / Year
Developers using AI coding tools professionally	62%	Stack Overflow 2024
Developer teams with AI tools in daily workflow	47%	GitHub State of Octoverse 2024
ChatGPT monthly active users	300M+	OpenAI early 2025
GitHub Copilot paid subscribers	1.8M+	Microsoft earnings call Q1 2025
Percentage of code in some repositories written by Copilot	46%	GitHub Octoverse 2024
Developers planning to expand AI tool usage in next 12 months	72%	Stack Overflow 2024
Organisations with at least one AI tool in dev stack	81%	JetBrains Developer Ecosystem 2024
Global AI developer tools market size	$4.6B	MarketsandMarkets 2024
Projected market size by 2028	$12.6B	MarketsandMarkets 2024
Enterprise spend per developer on AI tools annually	$1,200-2,400	Gartner estimate 2025

Security and Vulnerability Statistics

Finding / Metric	Statistic	Source / Year
AI-generated code with detectable security vulnerabilities	40%	Veracode 2024
Security issues introduced through Copilot suggestions in research	36%	Stanford Security Lab 2024
Developers who bypass security controls under deadline pressure	68%	SANS 2024
Average time to identify a data breach	194 days	IBM Ponemon 2024
Average total cost of a data breach globally	$4.88M	IBM Ponemon 2024
Average cost of a breach for organisations under 500 employees	$3.31M	IBM Ponemon 2024
Insider threat cited as top concern by security professionals	79%	Cybersecurity Insiders 2024
Organisations that suffered an insider-related incident in past year	71%	Cybersecurity Insiders 2024

Finding / Metric	Statistic	Source / Year
Proportion of breaches involving human error or insider action	68%	Verizon DBIR 2024
GDPR fines issued across EU (cumulative total to end 2024)	EUR 4.5B+	GDPR Enforcement Tracker 2025

Governance and Policy Statistics

Finding / Metric	Statistic	Source / Year
Organisations with a formal AI usage policy	14%	Forrester 2025
Enterprises with board-level AI governance structure	22%	Gartner 2025
SMEs with any form of AI governance documentation	9%	Forrester 2025
Organisations that reviewed AI tool ToS before adoption	8%	ISACA 2024
Legal teams involved in selecting developer AI tools	18%	Thomson Reuters 2024
Developers who received AI-specific security training in past year	11%	SANS 2024
Organisations with technical monitoring of AI tool usage	9%	Forrester 2025
DevSecOps programmes that include AI tool governance	7%	Gartner 2024
Companies with AI tool usage in employee code of conduct	23%	Mercer Workforce Survey 2024
Organisations tracking which AI tools developers use	16%	IDC 2024

Intellectual Property and Compliance Statistics

Finding / Metric	Statistic	Source / Year
Copyright offices rejecting AI-only authorship claims	Multiple	USPTO/UKIPO/EUIPO 2023-2025
Copilot verbatim reproduction of licenced code in test conditions	Approx 1%	Princeton/NYU 2023
Open source compliance issues in AI-assisted codebases	7-15%	FOSSA 2024
IP-related disputes involving AI-generated content	300+ filed	Global count through 2024
Technology SMEs with no formal IP audit in past 3 years	61%	WIPO 2024
Trade secret value as portion of tech SME IP portfolio	35-60%	WIPO estimate 2024

Finding / Metric	Statistic	Source / Year
Software development organisations with code classification policy	28%	IDC 2024
Organisations aware of AI training data source for tools they use	11%	ISACA 2024

AI Tool Risk by Provider Category (2025)

Tool Category	Data Sovereignty	Training Risk (Free)	IP Protection Level	Recommended For
Free cloud LLM (browser)	Low	High	Minimal	Public code only
Paid personal tier	Low-Medium	Medium	Limited	Non-sensitive tasks
Enterprise cloud tier	Medium-High	Low (excluded)	DPA available	Most tasks with policy
Self-hosted open model	Full control	None (no transmission)	Complete	All sensitive work
IDE plugin (free)	Low	Depends on config	Minimal	Non-proprietary only
IDE plugin (enterprise)	Medium-High	Low (excluded)	DPA included	Standard development

Chapter 10

Sector-Specific Risk Analysis

Why Sector Context Matters

AI tool governance risk is not uniform across sectors. The combination of regulatory exposure, IP sensitivity, data sensitivity, and contractual obligations varies significantly depending on what the organisation does and who it serves. This chapter profiles five sectors where the risk profile is particularly acute.

Financial Technology and Payments

Financial technology companies occupy a sector where proprietary algorithms represent the core commercial value and where regulatory obligations are layered on top of standard IP concerns. A payment processing algorithm, a fraud detection model, or a credit scoring approach is simultaneously the company's primary competitive asset, a potential trade secret, and a system that may be subject to explainability requirements under consumer protection regulation.

Developers at fintech companies working on transaction processing code, risk models, or regulatory reporting systems frequently need to debug complex logic and optimise performance. These are precisely the use cases for which AI tools are most useful and most likely to be used without explicit thought about what is being shared. A developer asking an AI tool to help optimise a fraud detection rule is sharing the detection logic that the company considers its core IP.

Additional exposure arises from compliance obligations. PCI DSS (Payment Card Industry Data Security Standard) requirements for protecting cardholder data extend to any system that processes, stores, or transmits cardholder data. Submitting code that handles payment card data to an external AI service may create a PCI DSS compliance question that the company's QSA (Qualified Security Assessor) has not been asked about.

Healthcare Technology

Healthcare software companies face IP risk compounded by regulatory obligations around medical device software and health data protection. HIPAA in the US, the EU Medical Devices Regulation, and equivalent frameworks in other jurisdictions create specific requirements for the documentation, validation, and control of software used in clinical contexts.

A developer at a healthcare software company submitting clinical decision support logic to an AI service is simultaneously potentially sharing proprietary algorithmic IP, potentially submitting code that handles protected health information, and potentially creating a documentation gap in the software validation record that the product's regulatory approval depends on. The validation record for medical device software typically requires that all development tools be qualified and their use documented. An unqualified AI tool in the development workflow is a regulatory gap.

E-commerce and Retail Technology

E-commerce technology companies build competitive advantage through recommendation algorithms, dynamic pricing engines, personalisation systems, and supply chain optimisation. These systems represent the commercial heart of the business and are typically the subject of significant engineering investment. They are also precisely the systems that developers most frequently want AI assistance with, because the logic is complex and the debugging surface is large.

A developer submitting a recommendation algorithm or a dynamic pricing model to an AI service is sharing the technical logic that produces the competitive differentiation the business is built on. The fact that the code appears abstract (numerical models, scoring functions, weighting parameters) does not reduce its sensitivity. In many cases, the abstract representation is more revealing than the user-facing product, because it exposes the specific decisions the organisation made about how to optimise for commercial outcomes.

Government Technology and Public Sector Contractors

Organisations building software for government clients face a distinct combination of security classification requirements, procurement restrictions, and export control obligations. Many government technology contracts include specific provisions prohibiting the use of foreign-hosted AI services for development work related to the contract. These provisions are often not on the radar of the individual developer who receives a task and reaches for the fastest tool available.

Beyond contractual restrictions, government technology often involves infrastructure or process logic that falls within export control classifications. Engineers working on border control systems, law enforcement tools, defence procurement software, or critical national infrastructure face specific obligations that most AI tool usage would violate. The consequences of violation are not limited to contract termination. Personal liability under export control law, including criminal liability in the most serious cases, is a real exposure.

Early-Stage Startups Seeking Investment

Startups approaching venture capital fundraising or strategic acquisition face a specific IP risk that intersects with AI tool usage. Sophisticated investors and acquirers conduct technical due diligence that increasingly includes assessment of AI tool governance. The question is not just whether the codebase has IP problems. The question is whether the founders can credibly claim that the IP is clean, unencumbered, and defensible.

A startup that has been using free-tier AI tools without governance cannot make this claim confidently. It cannot certify that no proprietary code was submitted to training-enabled services. It cannot certify that no AI-generated code in the product carries open source licence obligations. It cannot demonstrate that its trade secret protections were maintained through reasonable measures. Due diligence findings in any of these areas can reduce valuation, introduce conditions into term sheets, or in the most serious cases cause deals to collapse.

Chapter 11

The True Cost of Inaction

Beyond the Headline Risk

The framing of AI tool risk as "IP theft" or "data breach" captures the dramatic end of the risk spectrum but misses the more common and more insidious costs. Most organisations will not experience a confirmed incident where a competitor demonstrably benefited from proprietary code that was shared with an AI service. This makes the risk easy to dismiss as theoretical. The costs that are almost certain to occur are subtler but financially material.

The Valuation Impact

For organisations in growth mode, the clearest financial consequence of poor AI governance is reduced valuation at fundraising or exit. Investor due diligence has evolved rapidly since 2022. Technical due diligence teams now ask specific questions about AI tool usage in development workflows, the account types used, and whether governance policies exist and are enforced.

A startup that cannot answer these questions confidently faces adjustments to its valuation through IP risk haircuts or indemnification requirements that transfer liability to the founders. In a competitive fundraising environment where investors are comparing opportunities, a clean IP story is a differentiator. A messy IP story involving uncontrolled AI usage is a reason to pass.

The Compliance and Audit Cost

Organisations operating in regulated sectors or serving enterprise customers are increasingly required to demonstrate AI governance as part of vendor assessments and certification renewals. ISO 27001 bodies are updating guidance to include AI tool usage in scope. SOC 2 examinations are beginning to incorporate AI governance questions. Enterprise customers in financial services, healthcare, and government are inserting AI governance requirements into vendor agreements.

The cost of retroactively building governance infrastructure under time pressure from an audit finding or a customer requirement is substantially higher than the cost of building it proactively. An organisation asked to produce evidence of AI governance controls for a due diligence process in six weeks faces a choice between a rushed and inadequate response and a delayed deal.

The Talent and Culture Cost

Developers who work at organisations with clear, well-reasoned AI governance policies experience less friction and less ambiguity than those who work without policies. A developer who knows exactly which tools they can use for which tasks, under what account types, and for which categories of code does not need to make daily judgement calls that create risk. A developer without guidance either makes those calls independently (creating risk) or avoids AI tools entirely to stay safe (creating productivity cost).

The talent dimension extends to recruitment. Developers with security awareness and professional standards actively prefer organisations that have thought through AI governance. The signal value of a well-designed AI policy is that the organisation takes engineering professionalism seriously. The absence of a policy is a signal of a different kind.

Cost Category	Probability	Financial Impact Estimate
Due diligence IP risk finding (M&A)	High for active deals	Valuation reduction 5-20%
Investor due diligence complication	Medium-High	Condition, delay, or deal failure
GDPR fine (personal data in submissions)	Low-Medium	EUR 10K to EUR 500K+ for SMEs
Enterprise customer audit failure	Medium	Contract loss or remediation cost
Trade secret enforcement failure	Low (only if dispute arises)	Loss of IP enforcement right
Open source compliance remediation	Medium	Engineering cost 20-100 days
Security certification scope expansion	High as frameworks update	Audit cost increase 15-30%
Reputational damage (if incident disclosed)	Low (events are rare)	Customer loss, hiring difficulty

Chapter 12

The Risk Scorecard

Assessing Your Organisation's Exposure

The following scorecard is designed as a structured first-pass diagnostic, not a comprehensive audit framework. Complete it honestly based on the current state of your organisation, not the state you intend to reach. Each dimension is scored from 0 (adequate) to 4 (significant exposure).

Score 0-20: Low to moderate risk. Policy development and awareness work is the priority. No immediate action required but a governance programme should be initiated within six months.

Score 21-30: Moderate to high risk. Immediate policy development and tool procurement review is required. Legal counsel should be engaged for a preliminary IP and compliance assessment.

Score above 30: High to critical risk. Legal review of IP exposure, urgent policy development, and technical controls implementation should be initiated within 30 days.

Risk Dimension	Score 0 (Low)	Score 2 (Medium)	Score 4 (High)
AI usage policy exists and is current	Formal, current, enforced	Informal or outdated	Does not exist
Account types in use	Enterprise only	Mixed personal/enterprise	Free/personal only
Data classification framework for code	Documented and applied	Informal understanding	No classification
Technical monitoring of AI tool usage	Controls in place	Visibility without controls	No monitoring
Developer training on AI risks	Specific training < 6 months	General security training only	None in past year
IP inventory and documentation	Full documented inventory	Partial documentation	Not documented
Legal review of AI tool terms	Reviewed and assessed	Not recently reviewed	Never reviewed
Code review process for AI output	Formal review required	Informal review practice	No review process
GDPR/data protection assessment for AI	AI included in DPIA process	Partial assessment	Not assessed
Incident response for AI disclosures	Specific procedure exists	General IR plan only	No procedure

Industry Baseline

Based on observations across multiple client engagements in the 10-200 employee technology sector during 2023-2025, the typical organisation in this bracket scores between 28 and 36 before any governance work is undertaken. A score above 30 indicates an IP and compliance risk profile that is not consistent with reasonable data protection practice, regardless of what security policies say on paper.

Chapter 13

What Good Governance Looks Like

The Goal Is a Sustainable Framework, Not a Ban

Organisations that respond to AI tool risk by prohibiting all AI tool usage create a different and in some respects worse problem. Prohibition without alternatives is not enforceable in an engineering culture where AI tools have become genuinely integral to productivity. Developers find workarounds. Personal mobile hotspots bypass corporate web filters. Private accounts bypass monitoring. The risk exposure continues; the monitoring capability is lost.

Effective governance creates a framework in which AI tool usage happens in a controlled, documented, and auditable way. Developers retain access to tools that make them more productive. The organisation has visibility and contractual protections. The most sensitive code is handled under the most protective conditions. This is achievable without material reduction in developer productivity. In many cases, moving from uncontrolled free-tier usage to structured enterprise usage actually improves outcomes because better tools become available under better terms.

Component 1: Baseline Assessment

The first action is an honest inventory of current AI tool usage. Developer interviews, review of browser history or web proxy logs if available, and direct questions to engineering leads about what tools are in use and for what purposes. This is not an audit designed to identify and punish non-compliance. It is a baseline to understand actual exposure.

Most organisations find the baseline assessment surfaces patterns they were not aware of: specific tools being used for specific high-sensitivity tasks, account types that provide fewer protections than assumed, and historical submissions that include categories of code the organisation would not have deliberately authorised for external sharing.

Component 2: Data Classification Framework

A practical three-level classification system covers most development environments. Level one: code that does not contain proprietary business logic, is not commercially sensitive, and does not reference personal data. Standard library integrations, boilerplate configurations, and publicly documented API implementations typically fall here. Level one code can be used with any AI tool.

Level two: code that implements proprietary features, contains business logic, or handles personal data, but where enterprise account protections are sufficient. Most development work falls here. Level two code requires an enterprise account with a data processing agreement before submission to cloud AI services.

Level three: code that represents the organisation's core competitive differentiation, that implements security-critical functionality (authentication, encryption, access control), or that handles particularly sensitive personal data categories. Level three code should only be processed by locally hosted AI tools or should be excluded from AI assistance entirely.

Component 3: Enterprise Tool Procurement

The commercial cost of enterprise AI agreements is substantially lower than the cost of a single regulatory investigation, a due diligence finding, or an open source compliance remediation exercise. GitHub Copilot Business costs $19 per user per month. GitHub Copilot Enterprise costs $39 per user per month. An engineering team of 10 developers running

Copilot Enterprise pays $4,680 per year for explicit training exclusions and a data processing agreement.

The procurement decision should be preceded by legal review of the specific data processing addendum offered by each provider. Enterprise agreements vary between providers and have been updated multiple times. The DPA, not the product page or the sales conversation, governs the actual obligations.

Component 4: Developer Training

Training on AI tool governance cannot be a subsection of generic annual security awareness training. It needs to be specific, scenario-based, and relevant to the tools developers actually use. Training that refers to "AI services" in the abstract without naming the tools developers use every day does not change behaviour.

Effective training content addresses the specific risks in the use cases developers encounter: debugging production issues with AI assistance, code review with AI assistance, generating tests, using AI for architecture questions. For each use case, the training clarifies which category of code is involved, which tool configurations are appropriate, and what to do when uncertain.

Component 5: Technical Controls

Policy without technical controls relies entirely on voluntary compliance. Technical controls that can reinforce the policy include web proxy logging to identify AI tool domain traffic, DLP configuration to detect code submission patterns, Copilot policy controls in GitHub organisations that restrict which features are enabled, and local model deployment for work on high-sensitivity code.

The controls do not need to be punitive or surveillance-oriented. Logging that the organisation uses to understand AI tool usage patterns (how much traffic, which tools, what volume) is qualitatively different from logging designed to catch and discipline employees. The former builds the governance picture the organisation needs. The latter damages engineering culture.

Component 6: Incident Response Procedure

The AI disclosure incident response procedure addresses a specific scenario that standard IR frameworks do not. Key elements: how to identify what was shared and through which service; how to assess whether the shared content constitutes a trade secret, personal data, or regulated technical information; what GDPR notification obligations apply; what legal remedies are available; and how to document the incident for potential regulatory enquiry.

The procedure should be exercised at least annually in a tabletop exercise that involves the engineering lead, a legal representative, and the most senior security-accountable person in the organisation. Organisations that have never worked through the scenario in a low-stakes environment will be unprepared when a real incident occurs.

Chapter 14

Building an AI-Ethical Engineering Culture

Why Culture Matters More Than Policy

Policies without culture are documents. Culture without policies is inconsistent. The organisations that manage AI tool risk most effectively combine written governance frameworks with engineering cultures in which security consciousness and professional responsibility are genuine values rather than compliance obligations.

An engineering culture that treats AI governance seriously is not one in which developers feel surveilled or restricted. It is one in which developers understand why certain boundaries exist, have access to the best available tools within those boundaries, and feel confident making independent judgements about edge cases because they understand the principles rather than just the rules.

Leadership Signal

The most powerful driver of engineering culture is the behaviour of technical leadership. A CTO or VP of Engineering who uses enterprise AI accounts, talks openly about why the organisation made specific tool choices, and treats AI governance as a professional standard rather than a compliance burden signals that this matters. A leader who uses their personal ChatGPT account for work and waves off governance questions with "we need to move fast" signals the opposite.

The behaviour does not need to be dramatic. Visible choices in normal workflow communicate cultural expectations more effectively than policy documents. Which tools appear in onboarding materials? What account types are shown in team demos? Are AI tool governance questions given time in engineering all-hands meetings? These signals accumulate.

Onboarding as the Governance Gateway

Developer onboarding is the most reliable mechanism for establishing AI governance expectations. A new engineer who goes through structured onboarding that includes a specific session on AI tool usage, the organisation's classification framework, which tools are approved on which account types, and why this matters arrives with the right mental model from the first day. This is substantially easier than trying to change existing habits in a team that has been working without governance for years.

Onboarding content should include practical exercises rather than policy recitation. A developer who has worked through a real example of classifying code and choosing the appropriate AI tool under that classification remembers the framework better than one who read it in a document.

Feedback Loops and Continuous Improvement

AI tool capabilities are changing rapidly. A governance framework built for the tool landscape of early 2024 may not adequately address the capabilities and risk profile of tools available in late 2025. Governance needs a review cycle that is faster than the annual security policy update cycle.

Quarterly reviews of the approved tool list, the classification framework, and the enterprise agreements in place are a reasonable minimum. These reviews should involve both a technical perspective (what new capabilities have emerged, what new tools are developers asking about) and a legal/security perspective (what has changed in provider terms, what regulatory developments are relevant).

AI Ethics as Competitive Differentiation

In B2B markets, AI governance is becoming a positive signal rather than merely the absence of a negative one. Enterprise buyers assessing software vendors increasingly view documented AI governance as evidence of organisational maturity and professional standards. A vendor that can credibly demonstrate how it governs AI tool usage in development, what IP protections are in place, and how it ensures the code delivered to customers is clean and auditable has an advantage over one that cannot answer these questions.

The organisations that build AI governance infrastructure in 2025 will not only avoid the costs associated with uncontrolled usage. They will be positioned to use their governance posture as a differentiator in procurement processes, in fundraising, and in the talent market.

Chapter 15

How Bithost Can Help

Making AI Use Standard, Ethical, and Auditable

Bithost works with technology organisations to build the governance layer that makes AI-assisted development safe, auditable, and defensible. The work is grounded in what is actually happening in the organisation, not in a template framework applied regardless of context.

AI Usage Audit

We conduct structured interviews with development teams and a technical review of current tooling to produce an accurate baseline of AI tool usage across the organisation. The audit identifies account types in use, maps use cases to code sensitivity categories, and produces a prioritised risk assessment. Most organisations find that the audit surfaces patterns they were not aware of and provides the first honest picture of actual exposure. The typical audit engagement runs two to three weeks.

Policy and Governance Framework Development

We develop AI usage policies that are specific and actionable rather than generic. The policies are written in language that developers understand and can apply without consulting a lawyer. They include specific guidance for each code classification level, approved and prohibited tool configurations, and clear escalation paths for edge cases. We work with your legal team or connect you with qualified legal counsel to ensure the policy is consistent with applicable law in your jurisdiction and with your contractual obligations to customers.

Enterprise Tool Procurement and Vendor Assessment

We assess the AI tools your team currently uses or wants to use, review the relevant data processing agreements and terms of service, and provide a structured comparison of the protections offered under each tier. Where enterprise agreements are appropriate, we support the procurement and legal review process. We do not accept commissions or referral payments from AI tool providers. Our assessments reflect your interests, not theirs.

Developer Training

We deliver training sessions designed specifically for development teams that address AI tool risks in the context of real development workflows. The training is scenario-based, practically focused, and updated to reflect current tool capabilities and current regulatory guidance. It is not generic security awareness content repurposed for AI. It addresses the specific scenarios where developers in your team are most likely to create exposure.

Technical Controls and Monitoring

For organisations requiring technical enforcement of AI usage policies, we design and implement appropriate controls. This may include DLP configuration, network-level filtering with approved tool exceptions, monitoring of AI tool API usage, integration of governance into CI/CD pipelines, and local model deployment for sensitive development work. Controls are designed to support developers rather than obstruct them.

Ongoing Advisory

The AI tool landscape and the regulatory environment around it are both changing faster than any static document can track. Bithost provides ongoing advisory services to clients who need a partner that keeps pace with developments and translates them into practical guidance for their specific situation. Clients on advisory retainers receive quarterly briefings on material developments, proactive notification when tool provider terms change, and on-call support for governance questions as they arise.

Ready to understand your actual exposure?

Start with an AI usage audit. Most clients complete it in two to three weeks and leave with a clear picture of where they stand and what to do about it.

No template frameworks. No generic output. An assessment built for your team, your tools, and your code.

Bithost

ZHOST Consulting Private Limited

www.bithost.in

sales@bithost.in

Appendix A: Key Terms and Definitions

Trade Secret

Information that derives its commercial value from being kept confidential, where the owner has taken reasonable steps to maintain that confidentiality, and where the information is not generally known to or readily ascertainable by persons who would benefit commercially from its disclosure. Source code, algorithms, business logic, and technical architectures commonly qualify. Unlike patents, trade secrets are not registered and cannot be restored once disclosed.

Data Processing Agreement (DPA)

A contract between a data controller (the organisation using an AI tool) and a data processor (the AI tool provider) governing how personal data is handled on behalf of the controller. Required under GDPR Article 28 wherever a processor handles personal data on the controller's behalf. Enterprise AI agreements typically include a DPA. Free tier accounts do not.

Large Language Model (LLM)

A machine learning model trained on large datasets of text to generate human-like output in response to prompts. ChatGPT, Claude, Gemini, and the models underlying GitHub Copilot are all LLMs. They generate output by predicting statistically likely sequences of text given an input, based on patterns observed in training data.

Training Data

The dataset used to train an AI model. For code-generation models, training data typically includes large volumes of public code from repositories, documentation, and technical writing. Some providers also use interaction data from user submissions to refine models. The use of submitted code as training data varies by provider and account tier.

Copyleft Licence

An open source licence condition requiring that derivative works be released under the same or compatible licence terms. The GNU General Public Licence (GPL) is the most common copyleft licence. If code incorporates sequences from GPL-licensed training data, the codebase using that code may be subject to GPL obligations including source code disclosure requirements.

EU AI Act

European Union Regulation (EU) 2024/1689 establishing a risk-based framework for AI systems. Prohibitions on high-risk practices applied from February 2025. High-risk system obligations apply from August 2026. Introduces transparency requirements for general-purpose AI models that will affect how providers disclose training data sources.

SOC 2

Service Organisation Control 2, a voluntary compliance framework developed by the American Institute of CPAs. Assesses a service organisation's controls over security, availability, processing integrity, confidentiality, and privacy based on the Trust Services Criteria. A SOC

2 report reflects controls as assessed during the audit period and does not provide real-time or continuous assurance.

Appendix B: References and Sources

Reference	Document Title or Description	Year
Stanford HAI	AI Index Annual Report, developer adoption data	2025
Stack Overflow	Annual Developer Survey, AI tools and usage section	2024
GitHub	State of the Octoverse, AI-assisted development statistics	2024
Forrester Research	AI Governance in the Enterprise survey	2025
Veracode	State of Software Security: AI-generated and AI-assisted code	2024
SANS Institute	Security Awareness Report, developer behaviour data	2024
IBM/Ponemon Institute	Cost of a Data Breach Report	2024
WIPO	Intellectual Property and Small and Medium Enterprises Survey	2024
Gartner	AI in Software Engineering forecast and governance survey	2025
IDC	Worldwide AI Developer Tools and Platforms Spending Guide	2024
GDPR Enforcement Tracker	Aggregated fine data across EU supervisory authorities	2025
Princeton/NYU	Codex and Copilot memorisation and licence compliance study	2023
FOSSA	Open Source Risk Report: AI-assisted development and licence compliance	2024
MarketsandMarkets	AI Code Tools and Development Platforms market forecast	2024
Thomson Reuters	Future of Professionals: Legal technology and AI usage survey	2024
Cybersecurity Insiders	Insider Threat Report	2024
ISACA	State of Cybersecurity: AI risks and governance	2024
Verizon Business	Data Breach Investigations Report	2024
JetBrains	The State of Developer Ecosystem, AI tools section	2024
Wall Street Journal	Apple restricts employee AI tool use	May 2023
Reuters	Samsung confirms ChatGPT data leak incidents	April 2023
Bloomberg	Amazon internal memo warning on ChatGPT confidentiality risks	Feb 2023
EU Official Journal	Regulation (EU) 2024/1689 (EU AI Act)	2024

Reference	Document Title or Description	Year
Mercer	Global Workforce Trends: AI in the workplace and governance	2024
Stanford Security Lab	Security properties of AI-generated code analysis	2024

Appendix C: AI Tool Comparison Matrix

This matrix provides a reference comparison of the major AI coding tools available as of early 2025. Tool terms and features change frequently. Verify current terms directly with each provider before making procurement decisions.

Tool	Free Training Use	DPA Available	Enterprise Tier	Local Option	Best For
ChatGPT (OpenAI)	Default ON	API Enterprise	Yes	No (cloud)	General Q&A, debugging
GitHub Copilot	Individual: opt-out	Business/Enterprise	Yes	No	IDE completion, code gen
Claude (Anthropic)	Safety review only	API/Teams/Enterprise	Yes	No	Complex reasoning, review
Gemini (Google)	Consumer: default ON	Workspace tier	Via Workspace	No	GCP integration
Codeium (Windsurf)	Individual: limited	Enterprise tier	Yes	Enterprise option	IDE, free alternative to Copilot
Tabnine	Cloud: limited	Enterprise	Yes	Yes (self-hosted)	Air-gapped, on-premise needs
Amazon Q Developer	Professional: excluded	AWS BAA available	Yes	No	AWS-heavy stacks
Ollama + local models	N/A (no transmission)	N/A	N/A	Yes (fully local)	Sensitive code, full control
Cursor (local mode)	N/A (local)	N/A	Business tier	Yes	Sensitive work, IDE experience
LM Studio	N/A (local)	N/A	N/A	Yes (fully local)	Local model management

www.bithost.in | sales@bithost.in