asteyo

WHITE PAPER · AI & NEARSHORE DELIVERY

AI-Augmented, Human-Accountable

What GenAI really changes about nearshore software delivery — and what it does not.

Alexandru Bereczki, CEO · Thomas Wolenski, CSO

June 2026

Asteyo SRL, Cluj-Napoca · office@asteyo.com · asteyo.com

EXECUTIVE SUMMARY

The honest picture of AI-augmented delivery

Two forces are reshaping how software gets built. AI coding tools have moved from curiosity to mainstream in under three years. EU nearshore IT delivery is simultaneously gaining ground as organisations seek engineering depth, time-zone alignment, and accountable governance frameworks.

This paper makes a straightforward argument: GenAI raises the value of senior engineering judgment, not replaces it. The right nearshore partner has engineers who are senior enough to use AI well, governance that makes that usage safe, and process discipline that makes the gains stick.

We are Asteyo — an EU nearshore IT partner founded in 2025, headquartered in Cluj-Napoca, Romania. We are a small, young company in a build-up phase with a limited reference base. What we have is a delivery model built on senior-first staffing, engineering discipline, EU jurisdiction, and founders with 20+ years of enterprise IT experience.

What this paper establishes:

- AI coding tools deliver real gains in specific conditions — and measurably negative outcomes in others.
- The variance is wide, and the direction depends heavily on who is using the tools and on what.
- Experienced engineers in large, mature codebases have been measured as 19% slower with AI access.[7]
- Junior developers using AI without oversight are generating security vulnerabilities at scale: 45% of AI-generated code fails OWASP Top 10 tests.[9]

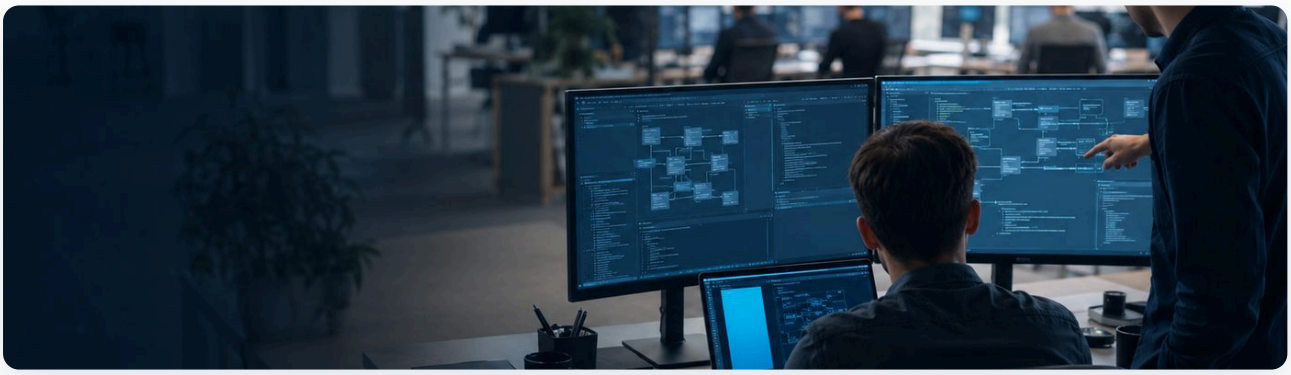

01 · CONTEXT

Two forces, one false choice

Enterprise software delivery economics are being repriced from two directions simultaneously, and most organisations are still treating them as separate decisions.

Generative AI in the development workflow

By mid-2025, 84% of developers were using or planning to use AI tools, up from 76% a year earlier, with 51% of professional developers using them daily.[1] The productivity promise is genuine in the right context: controlled trials have measured task completion 40–55% faster on isolated, greenfield tasks.[2,3] Yet the same survey found more developers actively distrust AI tool accuracy (46%) than trust it (33%), with only 3% reporting high trust.[1]

EU nearshore IT delivery

Romania has an estimated 200,000 to 250,000 software engineers and graduates around 10,000 ICT professionals a year. DACH and UK enterprises that partner with Romanian nearshore providers typically realise total-cost savings of 40–50% against equivalent hires in Frankfurt, London, or Zurich — without the ten-hour time-zone gap or geopolitical risk that offshore models carry.[5]

The framing being presented to engineering leaders — adopt AI aggressively and your headcount problem disappears, or choose cheap nearshore and keep humans cheap — does not hold up in either direction. AI coding tools do not replace engineering judgment. They amplify it: upward when senior engineers wield them deliberately, and downward when junior developers use them unsupervised against complex production systems.

The 2024 DORA State of DevOps Report, drawing on over 39,000 professionals, found that a 25% increase in AI adoption is associated with a 1.5% decrease in delivery throughput and a 7.2% decrease in delivery stability — a direct consequence of teams skipping the process fundamentals that AI cannot substitute.[6]

84%
developers using or planning AI tools
Stack Overflow 2025 [1]

51%
using AI tools daily
Stack Overflow 2025 [1]

46%
distrust AI tool accuracy
Stack Overflow 2025 [1]

7.2%
delivery stability decrease with AI adoption
DORA 2024 [6]

Where AI helps, where it does not

The honest picture of AI coding productivity spans a wide range of findings. The variance is the signal, not a methodological inconvenience to be averaged away.

Where the gains are real

The most-cited controlled study found participants using Copilot completed a specific HTTP server implementation 55.8% faster than the control group.[2] A six-week field study at ANZ Bank found Copilot-assisted developers completed tasks 42.4% faster overall.[3] These gains are real in structured settings — for boilerplate generation, test scaffolding, documentation, and refactoring well-understood patterns.

Where the gains are absent — or negative

A 2025 RCT by METR recruited 16 experienced open-source developers working on large repositories averaging 22,000 GitHub stars and over one million lines of code, with 246 real tasks assigned to AI-allowed or AI-disallowed conditions. Tooling: Cursor Pro with Claude 3.5/3.7 Sonnet. Result: developers given AI access took **19% longer** on average. They had predicted 24% faster before the study, and believed it had made them 20% faster after. The measured direction ran opposite to their intuition.[7]

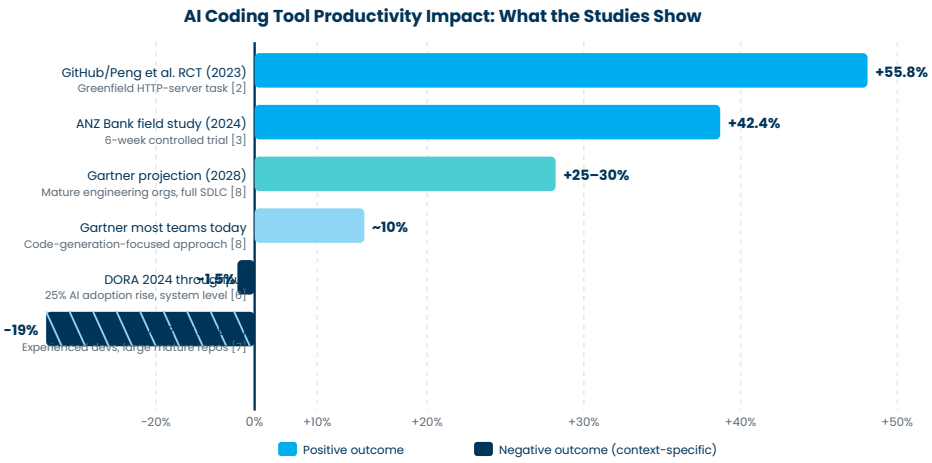

Figure 1: Measured productivity outcomes across major AI coding tool studies, 2022–2025. Variance reflects task type, codebase complexity, and developer seniority. No single figure generalises across all contexts.

Junior without oversight: the compounding risk

The most significant risk is AI deployed as a substitute for engineering depth rather than a multiplier on it.

45%
AI-generated code with security vulnerabilities
Veracode 2025 [9]

86%
XSS failure rate in AI-generated code
Veracode 2025 [9]

66%
Developers cite "almost right, not quite" as top frustration
Stack Overflow 2025 [1]

Veracode's 2025 GenAI Code Security Report tested over 100 large language models across 80 coding tasks in four programming languages and found that 45% of AI-generated code contains security vulnerabilities, including OWASP Top 10 failures.[9] AI models generated insecure code in 86% of scenarios involving cross-site scripting, and Java — the dominant language in DACH enterprise backends — showed the highest per-language failure rate at 72%.[9]

Developers who rely on AI tools without critical review consistently rate insecure AI-generated solutions as secure at higher rates than equivalent human-written code — a false sense of security. The 2025 Stack Overflow survey found that 66% of developers cite "solutions that are almost right, but not quite" as the dominant frustration: a friction cost that falls entirely on whoever is responsible for review.[1]

- **AI tools require review infrastructure.** Security gates and engineers capable of evaluating AI output are not optional overhead — they are load-bearing elements of any AI-augmented workflow.
- **CI/CD becomes more important, not less.** If AI generates code at higher volume and speed, the processes that catch errors before production must scale proportionally.
- **Junior developers using AI unsupervised** produce output that looks more complete and confident than it is. The review pressure is harder to skip precisely because the code appears plausible.

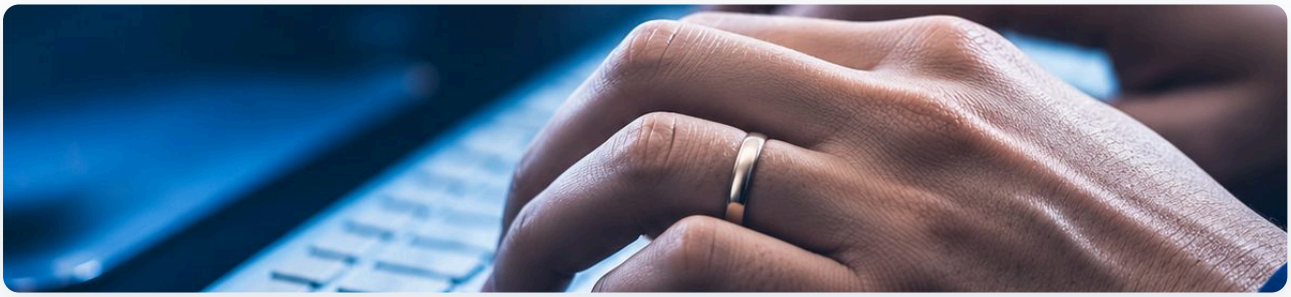

IP, data, and EU AI Act accountability

AI coding tools that process proprietary code raise concrete governance questions that are now board-level for regulated organisations.

IP and data exposure

Every AI coding assistant that processes proprietary code transmits that code to a third-party inference endpoint. For individual and team-tier accounts, most current tools use code inputs for model improvement by default, with opt-out available only at enterprise tier. For organisations whose competitive moat includes proprietary trading logic, underwriting models, or core product algorithms, the question is concrete: what code is being processed, under which terms, stored where?

Accountability in practice

In client engagements, our engineers use the client's own tools and environment — the tooling decision belongs to the client. At the founder level, Alexandru Bereczki and Thomas Wolenski personally use ChatGPT and Claude as part of their daily work. We are building our AI-usage framework around three principles: scope proprietary client code under explicit controls; review tooling selection against each client's regulatory environment; maintain engineer accountability for every output regardless of how it was generated.

EU AI Act: timeline update

The original August 2026 compliance deadline for Annex III high-risk AI systems has been deferred. Under the Digital Omnibus, the Council and Parliament reached a provisional agreement on 7 May 2026 to move the Annex III obligations from 2 August 2026 to 2 December 2027.[10] The change takes legal effect once formally adopted and published in the Official Journal, expected before 2 August 2026. The direction of travel is unchanged: accountability for AI-generated artefacts must trace to a human decision-maker. Development teams in regulated domains need to know whether their AI toolchain creates documentation or audit-trail obligations. The deferral buys time to build those controls; it does not remove the requirement.

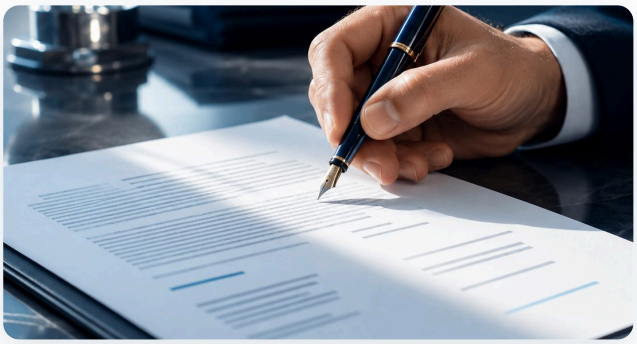

The structural argument for experienced delivery

The pattern across the evidence is consistent. AI coding tools are a multiplier. Multipliers amplify the quality of the input they act on.

- **Senior engineers become more productive.** AI handles mechanical scaffolding and boilerplate, freeing experienced engineers for architecture, requirements decomposition, security design, and code review. Gartner projects 25–30% productivity gains by 2028 — but only in organisations with mature engineering practices.[8]
- **Junior engineers without oversight become a liability.** AI does not substitute for years of pattern recognition, contextual judgment, and debugging experience. It makes junior output look more complete and confidently formatted than it is.
- **Architecture and accountability become the scarce inputs.** AI can write a function. It cannot own the consequences of that function in a production system. Someone must specify the right architecture, review the AI-generated implementation, understand its failure modes, and stand behind it when it breaks.
- **Code review changes character.** The question shifts from "did a human write this?" to "does a human understand this, stand behind it, and take responsibility for it?" The DORA evidence, METR study, and Veracode security data all converge on this point.

The capacity requirement does not diminish under AI adoption — it concentrates. Specifically, the need for senior capacity that uses AI well, because the judgment gap between a senior engineer and an unsupervised junior widens when AI makes junior output look superficially complete.

06 · ASTEYO'S STANCE

AI-augmented, human-accountable: honest about what we are

Asteyo was founded in 2025. We have a limited reference base, we are still building our client portfolio, and we are in a deliberate investment phase. We think you deserve to know that before deciding whether to engage with us.

What we have is a delivery model built from the start to the standard that enterprise engineering requires — because our founders have spent 20+ years on the enterprise side and know what that looks like.

- **Senior engineers from day one.** Our delivery model starts with experienced engineers, not a bench of juniors supplemented by AI output. Senior engineers are the engineers for whom AI assistance produces the most reliable productivity gains, and who are best positioned to catch the outputs that need catching.
- **Structured AI usage around clear principles.** In client engagements, our engineers work within the client's own tools and environment. We are building our AI-usage framework around explicit controls for proprietary client code, regulatory-environment review of tooling selection, and engineer accountability for every output.
- **Code review and CI/CD as non-negotiables.** Every AI-assisted pull request goes through the same review gates as human-authored code. Our CI/CD pipelines include automated security scanning (SAST), dependency vulnerability checks, and test coverage thresholds.
- **ISO/IEC 27001:2022 and ISO 9001:2015 certification.** Issued by SYSTEMA (IAS/IAF-accredited), valid to December 2028. Certified scope: custom, client-oriented software development. For clients completing a vendor security questionnaire, it is a meaningful starting point.
- **EU jurisdiction, GDPR-native.** All delivery infrastructure, data handling, and contractual relationships operate under EU law — the default operating environment, not an add-on.

ISO 27001 SYSTEMA · ISO 9001 SYSTEMA · **EU Jurisdiction** · **GDPR-Native**

ISO/IEC 27001:2022 & ISO 9001:2015 · SYSTEMA (IAS/IAF-accredited) · Valid Dec 2025–Dec 2028

07 · HOW ASTEYO DELIVERS

Three stages, de-risked from the start

Asteyo structures every engagement around a three-stage path that de-risks the start, builds capability, and sustains performance — whether you are a DACH Mittelstand company building a new platform or a Scale-Up managing a growing engineering footprint.

STAGE 1

Align and Pilot (Weeks 1–6)

A single well-scoped workstream: a priority integration, a service extraction, a test automation layer, or a specific feature domain. We demonstrate our engineering standards in your actual context. AI tooling is introduced and configured to your data governance requirements during this stage.

STAGE 2

Build and Integrate (Months 2–6)

The validated model scales to a dedicated pod or multiple workstreams. Senior engineers who led the pilot own knowledge transfer and architecture continuity. We integrate into your existing CI/CD, review, and planning processes rather than running a parallel track.

STAGE 3

Operate and Evolve

For clients requiring ongoing managed services, AMS, or dedicated team continuity, we operate as a stable engineering function. Senior retention on long-running engagements is a structural priority: the engineers who know the codebase stay on it.

Service coverage spans Java, Python, .NET, React, Angular, Vue, Node.js, and legacy COBOL modernisation; QA and test automation; DevOps and platform engineering; Azure, AWS, and GCP cloud; managed services and AMS; and SAP.

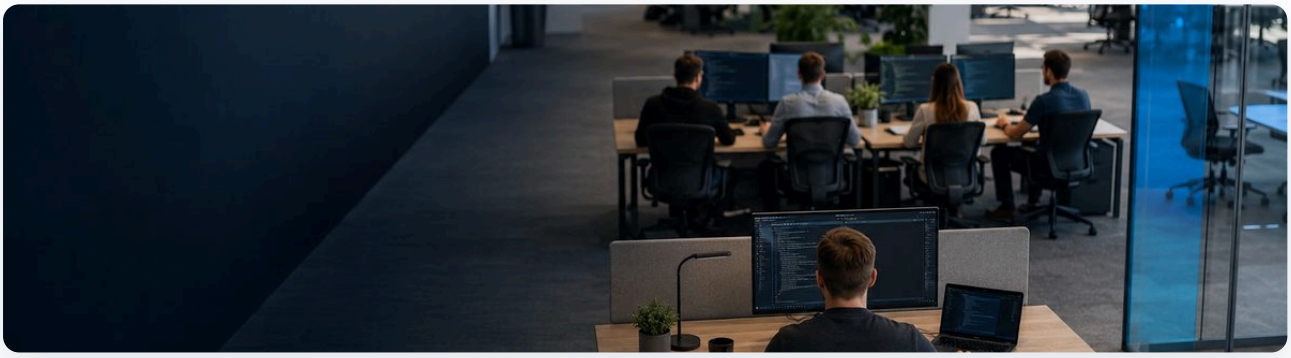

asteyo

"

What AI does is shift the scarce input. When code generation is cheap, the scarce thing becomes the experienced engineer who can specify what to generate, review whether the output is correct, and stand behind it in production.

08 · NEXT STEP

Start with a conversation

Our engagements start with a scoped pilot that demonstrates the model in your context before you commit to scale. We work with DACH Mittelstand engineering leaders and Scale-Up CTOs who want a partner that gives them a straight answer.

Questions worth asking any nearshore partner:

- What is your senior-engineer ratio, and how do you define “senior”?
- What security gates exist for AI-generated code before it reaches your CI/CD pipeline?
- Who specifically is accountable for AI-assisted output in a production incident?
- What data governance controls apply when engineers use AI tools on client codebases?
- Are your contracts structured to meet DORA Article 30 audit and exit requirements?

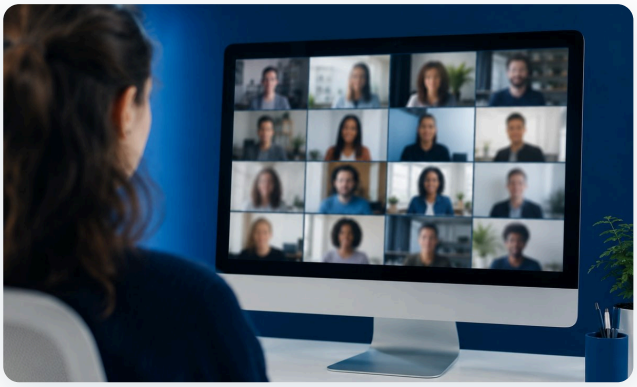

Next Step

Reach us at office@asteyo.com or at asteyo.com. The right next step is a conversation.

office@asteyo.com · asteyo.com

Asteyo — Reliability · People · Partnerships. We get things done.

REFERENCES

Sources & notes

Figures are cited to their primary or best-available public sources. Survey figures and regulatory scope estimates represent published data that may vary by source and methodology.

1 **Stack Overflow**, *2025 Developer Survey — AI section*, Dec 2025. 49,000+ respondents. survey.stackoverflow.co/2025/ai — 84% using/planning AI (up from 76%); 51% daily use; 33% trust accuracy, 46% distrust, 3% high trust.

2 **Peng, S. et al.**, *The Productivity Effects of Generative AI: Evidence from a Field Experiment with GitHub Copilot*, MIT GenAI Initiative, 2023. arxiv.org/pdf/2302.06590 — Greenfield HTTP-server task, RCT: 55.8% faster with Copilot.

3 **Harness / ANZ Bank**, *The Impact of GitHub Copilot on Developer Productivity: A Case Study*, 2024. harness.io/blog/… — Six-week field study: 42.4% faster; beginners +52%, advanced +40%.

4 **Compound.Law / Lurus.AI**, *GitHub Copilot GDPR and IP Analysis*, 2025. compound.law

5 **Romania talent figures** — secondary nearshore-market sources (Alcor BPO; DevsData): est. 200,000–250,000 engineers, ~10,000 ICT graduates/year; 40–50% cost advantage vs. DACH/UK. Density: Eurostat (2024) 2.8% of Romanian employment.

6 **Google/DORA**, *Accelerate State of DevOps Report 2024*. dora.dev/research/2024 — 39,000+ respondents. 25% AI adoption increase associated with 1.5% throughput decrease and 7.2% stability decrease.

7 **METR (Model Evaluation and Threat Research)**, *Measuring the Impact of Early-2025 AI on Experienced Open-Source Developer Productivity*, July 2025. metr.org/blog/2025-07-10-… and arxiv.org/pdf/2507.09089 — RCT: 16 experienced developers, 246 tasks, large mature repos (avg 22,000 stars, 1M+ LOC). AI access = 19% slower on average. Tooling: Cursor Pro + Claude 3.5/3.7 Sonnet.

8 **Gartner**, *Predicts 2025: Navigating the Rise of AI in Software Engineering*, 2024–2025. gartner.com/… — 25–30% gains by 2028 in mature orgs; most current teams see ~10%.

9 **Veracode**, *2025 GenAI Code Security Report*. veracode.com/… — 100+ LLMs, 80 tasks, 4 languages: 45% of AI-generated code fails OWASP Top 10; XSS failure 86%; Java highest per-language rate 72%.

10 **Council of the EU (Consilium)**, press release "Artificial Intelligence: Council and Parliament agree to simplify and streamline rules," 7 May 2026. consilium.europa.eu/… — Digital Omnibus: Annex III high-risk obligations deferred from 2 Aug 2026 to 2 Dec 2027.

Contact: office@asteyo.com · asteyo.com