

AI-ASSISTED DEVELOPMENT IN FINLAND 2026

What Finnish software professionals actually do with AI – and what stops them

Witted

SUMMARY

AI-assisted software development in Finland is not a pilot program. It is business as usual. 82% of the software professionals we surveyed use AI tools daily. 71% use AI agents. 68% use AI for architecture and design decisions, not only for writing code.

82%

use AI tools daily

71%

use AI agents

24%

deliver more scope, same price

95%

expect further change in 12 months


Three findings carry the rest of this report.

1. **Agents are the line that separates change from no change.** Of the 27 respondents who use AI agents, exactly one says AI has not meaningfully changed their work. Of the 11 who do not, more than half say nothing has changed.
2. **The barriers are not technical.** The two most-cited obstacles are customer security or contractual restrictions and the cost of tokens and licenses, both at 29%. Tool maturity sits near the bottom at 8%.
3. **Where AI has moved commercial terms, it has mostly moved them in the customer's favor.** 24% now deliver more scope at the same price. 13% face the harder version – customers expecting the same work for less.

None of this is a forecast. It is what practitioners reported doing in August 2026.

AI IS BUSINESS AS USUAL ALREADY


The striking thing about the usage data is not the headline daily-use figure. It is the breadth. Half of respondents use AI across all seven activities we asked about, and 61% across five or more.



Code review and architecture decisions at 68% is the number worth dwelling on. AI is participating in judgment calls – what to build, how to structure it, whether the change is sound – and not only in producing the code.

AGENTS ARE THE DIVIDING LINE

Splitting respondents by whether they use AI agents produces the sharpest contrast in the survey.




There are two honest readings and this survey cannot separate them. Agents may change the unit of work — you stop producing lines and start supervising output. Or the people already leaning in are simply the ones who reached agents first. It may well run both ways: agents take a developer to the next level and change how they work, and the developers whose way of working has already changed are the ones who pick agents up most readily. With a single wave, agents mark the line; we cannot show they caused it.

Either way, the practical implication holds. The question in 2026 is no longer whether a development team uses AI. 82% do, daily. It is whether they have moved from prompting to delegating.

NOT A TECHNOLOGY PROBLEM

39% of respondents selected "no significant barriers" — though 8 of those 15 also ticked a specific obstacle, so read it as nothing decisive rather than nothing at all. For everyone else, the pattern is consistent and it has very little to do with technology.



Tool maturity ranks second from last. The two barriers at the top are both customer-side decisions: what the contract permits, and who pays for the tokens. Neither is solved by a developer trying harder. Both are settled in a procurement conversation.

Cost deserves particular attention, because 29% is higher than most organizations would expect. It means teams have moved past free tiers into real consumption and nobody has decided whose budget line that is. The blocker is named more often by payroll employees (39%) than by freelancers (20%).

THE PERMISSION GAP

The most quotable finding in the survey is not a number. Three respondents, independently and unprompted, described the same situation: they have the capability, and their client environment will not allow it.

"Personally I prefer Claude Code (far superior to the ones I listed and it's not even close), but not allowed at client unfortunately."

— Survey respondent

"I use Claude only on my own projects, not customer projects."

— Survey respondent

"I can't use agents in work but do use them as a practice on my own time. In work I use AI as a "better Google""

— Survey respondent

This is the one barrier entirely within a client organization's control, and it is invisible from the client side unless somebody asks. Nobody in this survey had volunteered it to their client, because nobody had been asked.

It is worth stating what the gap costs. A consultant who runs agents on personal projects and uses AI as "a better Google" at work is delivering at a fraction of their demonstrated capability — and the difference is not skill, budget or willingness. It is a policy that may never have been revisited since it was written.

GOVERNANCE IS NOT THE BRAKE


A common assumption is that formal AI rules slow teams down. This data does not support it.

68% of respondents work under AI workflows, guardrails and tool choices defined at organization or team level. Only 18% say everyone sets their own approach. Where rules are defined, "lack of shared practices" is named as a barrier by 4%. Where everyone improvises, it is the top complaint at 57%.

And defining rules costs nothing in adoption: both groups use AI at the same intensity. Governance here is friction removal, not a handbrake. The organizations that wrote something down are not behind – they simply have fewer people complaining about the absence of shared practice.

PRICING AND SCOPING

37% of respondents have already seen their commercial terms move because of AI. Not next year – now.



The direction matters more than the total. Expanded scope beats rate pressure by nearly two to one. The value is arriving as more delivered output at a stable rate, not as rate compression.


In time-and-materials work – the dominant commercial model in Finnish software consulting – this is close to mechanical. Nobody decides to hand the customer more. An AI-accelerated hour simply produces more per hour billed, so the gain flows to the customer by default unless the rate moves. "More scope at the same price" happens automatically; "same work for less" requires somebody to open a negotiation.

One correlation is worth noting because it is absolute: all nine respondents who deliver more scope at the same price use agents regularly. Nine out of nine. Agent adoption and commercial upside appear to be the same phenomenon seen from two ends.

The 13% facing "same work for less" is small, but it is the leading edge of a real risk. If it grows faster than the "more scope" side, the market re-prices. One wave gives a level, not a trend – which is the main argument for running this survey again.

TOOLS AND MODELS

Development environments



Below that: Opencode (2 respondents); Cursor, Kiro, M365 Copilot and client in-house tools (1 each).

Models



Below that: locally-run models such as Qwen (2 respondents); Grok, Llama and Mistral / Devstral (1 each).

Two observations. An agentic command-line tool has overtaken the incumbent IDE assistant in this population, which says something about how the work is now done rather than about which vendor won. And multi-model is normal: 55% use two or more models and 50% two or more environments, with write-ins including locally-run models. Practitioners are choosing per task, not standardizing on a vendor.

WHAT PRACTITIONERS PUSH BACK ON

A survey where every comment is enthusiastic is a marketing document. This one is not, and the dissent is worth reading carefully, because it comes from people using these tools daily rather than from sceptics at a distance.

“AI is not faster for everything. Creates a lot of tech debt. Can’t handle complexity very well. Doesn’t understand like a human. It’s useful as a tool but these current models are probably creating more problems long term than they solve. Short term they are helpful but software is long term.”

— Survey respondent

That is the argument every experienced developer in the market is already making internally, and it deserves to be taken at face value: speed at the point of writing is not the same as total cost over a system's life.

“AI still hallucinates and makes mistakes/design choices that prevent us from completely handing over development to AI agents. It is already possible to drop the middle-man (developers) and just have AI agents completely do the work, but the end result is just not reliable enough, in my experience. But who knows how long it takes until this is possible. Then there’s the question power overconsumption and how it scales in the long run.”

— Survey respondent

The sentiment data carries the same tension. 95% expect their process to change further within 12 months. At the same time, half of respondents are at least somewhat concerned that AI will reduce demand for their own skills within two years. 45% are somewhat concerned and 5% very concerned. Another 24% expect demand for their skills to rise.

These are the same people using AI fluently every day. They are not resisting the technology. They are using it and thinking clearly about what it means for them, and any leader reading this will recognize both attitudes in their own team.

HOW TO DRIVE AI-ASSISTED DEVELOPMENT IN YOUR ORGANIZATION

Four things follow from the data, each of them a decision somebody already has the authority to make.

1. **The constraint is procurement, not people.** Customer security restrictions and token cost are the top two barriers at 29% each; tool maturity is near the bottom at 8%. The developers are ready. The contract and the budget line are not.
2. **Decide who pays for tokens.** Cost is a first-order blocker, and it is named more often by employees than by freelancers – which suggests independents may be absorbing it personally to stay fast. An unowned token budget is a capability cap.
3. **Write the rules down. Any rules.** 4% complain about missing shared practices where rules exist, against 57% where they do not, and defining them costs nothing in adoption.
4. **If your team is not running agents, that is the gap.** 96% of agent users report changed work, against 45% of non-users, and every respondent delivering more scope at the same price uses agents regularly.

Underpinning all four: ask your delivery team what they cannot use. Three respondents independently described having agent skills their client will not allow at work. Nobody volunteered it, because nobody was asked. That question is free, and it tells you which of the four to start with.

THE ONE THAT GOT AWAY FROM THE AVERAGES

Most of this report is percentages, which is the right way to read a survey. But one comment describes a change in kind rather than degree, and it is the clearest illustration of what the pricing data is measuring.

"A year ago I did 100% of coding by hand. Now AI makes 99% of the coding, I just do some random little pieces. The speed increase and lowered cost of code production has enabled my customer to improve ways of working and build tools that were not sustainable to build before AI. I checked and have built a tool consisting of 25k lines of code in about 40 hours that would not have been prioritised before. And if it had, it would have taken months to build."

— Survey respondent

Read that as a scope claim rather than a productivity claim. The tool was not built faster — it was built at all, having previously not been worth prioritizing. That is a different economic event from the same work getting cheaper, and it is the version of the AI story that most rewards a software organization's attention.

METHODOLOGY

Survey of 38 Finnish software professionals working through Witted – Witted Mavericks payroll employees and Witted Partners freelancers working through their own companies. Fielded August 2026. 38 respondents against a population of 173 – roughly a 22% response rate.

Respondents are predominantly fullstack and backend developers (68% fullstack, 76% fullstack or backend), plus architects, AI and data specialists, a designer and project managers. 71% work with client organizations of 1,000 or more employees. Client sector was recorded but is not analyzed here: the public-sector subgroup was too small to support a finding. Seniority and years of experience were not asked.

The survey was anonymous. No names or contact details were collected, and free-text comments are reproduced verbatim as anonymous respondent comments. Employment model: 20 freelancers, 17 payroll employees, one respondent describing themselves as a consultant – 47% of respondents employed-side. Freelancing is the more common of the two models in Witted's Finnish frontoffice, so payroll employees are likely over-represented here, which probably pulls the sentiment findings slightly more cautious than the wider group would.

Percentages are of all respondents unless stated otherwise, and multi-select questions sum above 100%. This is a practitioner pulse, not a statistically representative study of the Finnish market.

Witted