

A Third Declined to Buy. Renewals Never Moved.

Testing McKinsey's 32 percent against twelve vendors' disclosed retention, backlog and renewal rates, and pricing the three-year in-house build that replaces them

David Soden

Independent analyst and practitioner, enterprise automation and applied AI
davidsoden.com/market-assessment

About this report

Every substantive claim in this report carries an evidence classification and, where applicable, a numbered source. Facts are separated from vendor claims, third-party estimates, the author's assessments, and untested hypotheses. Forward-looking sections use scenarios with observable tripwires rather than point forecasts. Stated assumptions are listed before the analysis begins.

PUBLISHED FOR PUBLIC READING | SHARE FREELY, UNMODIFIED, WITH ATTRIBUTION

© 2026 David Soden, davidsoden.com/market-assessment. All rights reserved. You may share this file complete and unmodified, and quote short passages with attribution. Republishing under another name, excerpting into a commercial deliverable, resale, and use as machine learning training data require written consent. Full terms appear under Rights and Permitted Use on the final page.

Market Assessment reports are published and commissioned at davidsoden.com/market-assessment. This document is analysis, not investment, legal, or procurement advice.

Scope and evidentiary standard

This report tests one claim: that agentic coding tools are displacing packaged software purchases at enterprise scale. The claim entered wide circulation in late August 2026 on the back of a single number, McKinsey's finding that 32 percent of surveyed organizations declined to buy software because they could build it internally. The number has since been repeated across trade press, analyst notes and vendor competitive decks, usually without its question wording attached.

Coverage runs across three vendor groups: enterprise application vendors (Salesforce, SAP, ServiceNow, Workday, HubSpot), analytics and data platform vendors (Snowflake, Datadog), and developer tooling (Atlassian, GitLab), with C3.ai, Asana and DocuSign added as smaller reads on the same question. The fourth subject is the cost of the internal build that would replace any of them, priced over three years.

Published numbers in this category need discounting for three specific reasons. Survey findings about procurement intent are collected from panels of people who are not the procurement decision maker and are weighted for country, not for software spend. Vendor retention metrics are defined differently by every vendor that reports one, so net revenue retention at Snowflake and dollar net retention at DocuSign are not the same measurement even though both are printed as a percentage. And the cost work on custom software maintenance that everyone cites is decades old, predates cloud infrastructure entirely, and has never been re-run against agent-generated code.

Where the evidence here is a claim rather than a number, it carries a label saying so. Where a number exists but measures something narrower than the sentence it is being used to support, that gap is stated rather than smoothed over.

Label	Definition	How to treat it
FACT	Verifiable in the cited primary source: a filing, press release, official documentation, or standards body announcement.	Reliable as stated. Check the source if the exact wording matters.
VENDOR CLAIM	Reported by a vendor about its own business or product. Self-measured and not audited by any third party.	The vendor said it. Directionally useful, not a measurement. Definitions of the metric are usually undisclosed.
THIRD-PARTY ESTIMATE	A research firm forecast or survey result. Methodology is proprietary and generally not reproducible.	Use for direction and order of magnitude only. Do not build a model on it.
ASSESSMENT	The author's reasoned judgment from the cited evidence. Falsifiable, and the reasoning is shown so it can be argued with.	This is analysis, not reporting. Disagree with it where your evidence differs.
HYPOTHESIS	A proposition offered for testing. Not currently supported by sufficient evidence to assert.	Treat as a research question, not a conclusion. Each one names what would confirm or refute it.
SCENARIO	A structured future state with a stated subjective probability. The probability is the author's judgment, not a calculated figure.	Use the leading indicators attached to each, not the probability. The indicators are observable; the probability is not.

Assumptions this analysis rests on

Five premises are assumed rather than shown. Each names what fails if it is wrong.

First, that a vendor losing renewals to customer-built replacements would eventually have to say so. Public software companies disclose churn drivers when a cohort becomes material, because failing to is a securities problem and not just a communications one. If displacement is real and material, it surfaces in a risk factor, a downsell cohort, or an analyst answer. If this is wrong, and vendors are absorbing build-driven churn silently inside blended retention figures, then the vendor disclosure section of this report proves nothing and the survey answer stands unchallenged.

Second, that two quarters is long enough for a purchase decision made between May and June 2026 to show up somewhere in a forward metric. Enterprise contracts renew annually, so a declined renewal in mid-2026 would compress current remaining performance obligation before it compresses revenue. If enterprise renewal cycles are slower than assumed, or if the declines were concentrated in new purchases rather than renewals, the absence of a signal is early rather than meaningful.

Third, that the three-year cost model uses United States onshore engineering rates. An organization staffing the same build from a lower-cost location changes the arithmetic by roughly

half on the largest line item, and the break-even against a licence moves with it. Every conclusion about build economics in this report is a conclusion about onshore builds.

Fourth, that internal applications handling regulated data inherit the same audit obligations as purchased ones. A custom application in scope for SOC 2, HIPAA or PCI does not get a lighter treatment because it was written internally, and in practice gets a heavier one because there is no vendor attestation to inherit. If an organization's auditors accept a materially lighter control set for internal tools, the compliance line in the model is overstated.

Fifth, that the productivity research on agent-assisted development generalizes from open-source and general enterprise settings to greenfield internal application work. This is the weakest of the five. The METR trial studied experienced maintainers working in repositories they already knew, which is close to the opposite of a greenfield internal build. If agents are substantially better at net-new application code than the published research measures, the engineering line in the model is too high and the break-even falls.

The call

ASSESSMENT Agentic coding is changing what enterprises build, not yet what they renew. A third of surveyed organizations say they declined a purchase, but no vendor has disclosed a single renewal lost to an internal build, and forward commitments across the twelve I checked are growing rather than shrinking. The displacement that is real sits at the edges: the departmental tool, the reporting layer, the second seat tier nobody uses. Systems of record are protected by data gravity and certification, not by switching cost alone. I would change position on the first earnings call where a vendor attributes a named churn or downsell cohort to customer-built replacements. Watch gross retention, not net.

My judgment on the evidence assembled here, nothing more. There is data I can't see and data I may have missed. New evidence can move me, including to the opposite position, and I reserve the right to move.

Executive summary

The gap between what buyers say about build versus buy and what sellers report about renewals is the whole story. Both sets of numbers are real. They are answers to different questions, and the loudest reading of the survey number is not supported by anything a vendor has filed.

FINDING 1 The 32 percent measures a decision, not a cancellation.

FACT McKinsey's question asked whether an organization had decided against buying one or more software products or features because the functionality could be built internally with agentic coding tools. Nobody was asked whether the replacement shipped, whether it passed an audit, or whether a contract was terminated. A deferred evaluation and a cancelled renewal score identically. ^{[1] [3]}

FINDING 2 The sample is a standing executive panel weighted by country GDP.

FACT 1,719 responses across 97 nations, fielded 4 May to 8 June 2026, drawn from McKinsey's Global Survey Panel of more than 33,000 participants and weighted by each respondent's nation's contribution to global GDP. The weighting corrects for geography. It does not correct for software spend, seniority, or whether the respondent has any visibility into a renewal calendar. ^{[1] [4]}

FINDING 3 Twelve vendors, no corroboration.

FACT Across Salesforce, SAP, ServiceNow, Workday, Snowflake, Datadog, Atlassian, GitLab, HubSpot, Asana, DocuSign and C3.ai, not one filing or transcript from the June to September 2026 reporting run attributes churn, downsell or slowed renewal to customers building replacements with agentic coding tools. ServiceNow's renewal rate held at 98 percent. Atlassian's cloud revenue grew 31 percent on seat expansion, which is the opposite signal. ^{[6] [10]}

FINDING 4 Forward commitments are growing across every vendor checked.

FACT Atlassian's remaining performance obligations rose 44 percent, Snowflake's 30 percent, SAP's current cloud backlog 26 percent, ServiceNow's current remaining performance obligation 21 percent, and Salesforce's 14 percent. Workday is the slowest at 8 percent on total subscription backlog, and management attributed that to bookings mix rather than to lost renewals. ^{[5] [6] [7] [8] [10] [11]}

FINDING 5 Aggregate software spending is accelerating, not decelerating.

FACT Gartner's July 2026 forecast has worldwide IT spending up 14.2 percent in 2026 to \$6.37 trillion, with software growing 15.1 percent. If a third of enterprises were withdrawing purchases at any material dollar value, the software segment would be the first place it appeared. It has not appeared. ^[19]

FINDING 6 The real softness is in seat-priced collaboration, and it predates agentic coding.

ASSESSMENT HubSpot's net revenue retention sits at 102 percent, down a point year over year, and Asana's at 97 percent. Both are soft. Both companies attribute the softness to budget scrutiny and seat optimization, and both metrics have been drifting down since well before agentic coding tools were purchasable. Reading them as build displacement requires ignoring the trend line they sit on. ^{[12] [14]}

FINDING 7 A three-year in-house replacement runs roughly \$1.26 million against \$540,000 of licence.

THIRD-PARTY ESTIMATE Modelling a 300-seat application at \$180,000 a year, and pricing engineering, agentic tooling, security review, compliance evidence, hosting and one staff departure at published unit costs, the build lands near 2.3 times the buy over three years. The build pays only where the annual licence exceeds roughly \$420,000. ^{[27] [28] [29] [35]}

FINDING 8 Agents make the cheapest line item cheaper.

ASSESSMENT First-version development is the smallest recurring cost in any three-year model. The maintenance literature has put post-release work at 60 to 80 percent of lifecycle cost for four decades, and GitClear's 2026 analysis of 623 million code changes finds

refactoring line moves down 70 percent and duplicated code blocks up 81 percent as AI authorship rises. Cheaper first versions and more expensive second years is a worse trade than it looks. ^[21]

FINDING 9 The near-term cost of building is a security cost, not a licence saving.

FACT IBM's 2026 breach study of 602 breached organizations found shadow AI involved in 43 percent of incidents, up from about 20 percent a year earlier, with those incidents averaging \$5.39 million against a \$4.99 million global average. Retool's own customer survey reports builders shipping outside IT oversight because development now outpaces procurement. ^{[20] [25] [26]}

FINDING 10 McKinsey's industry split is the best available evidence for the protection thesis.

FACT Technology reported 41 percent and healthcare payers and providers 39 percent, against insurance at 19 percent and the public and social sector at 17 percent. The spread tracks regulatory dependency and certification burden almost exactly, which is a stronger and more useful finding than the headline average it was averaged into. ^[3]

What the 32 percent actually measures

FACT McKinsey published The State of AI: Global Survey 2026 on 25 August 2026. The online instrument was in the field from 4 May to 8 June 2026 and returned 1,719 responses from 97 nations. Thirty-six percent of respondents said they work for organizations with more than \$1 billion in annual revenue. The data are weighted by each respondent's nation's contribution to global GDP. Trade coverage of the finding began appearing within days of publication. ^{[1] [2]}

FACT The finding at issue reads that nearly a third of respondents, 32 percent, report their organizations have decided against buying one or more software products or features because they could be built internally with agentic coding tools. ^[1]

Four things that sentence does not say

ASSESSMENT It does not say a contract was terminated. Deciding against buying covers a renewal a CIO killed, an evaluation a team shelved, a module a customer declined to add to an existing subscription, and a proof of concept nobody ever funded. Those four have wildly different revenue consequences for a vendor and identical scores on this question. ^{[1] [3]}

ASSESSMENT It does not say the replacement exists. The question is about a procurement decision, not a shipped system. An organization that declined a purchase in May 2026 on the belief that agents could build the equivalent, and has nothing running in September, answers yes. ^[3]

ASSESSMENT It does not say how many, or how much. One or more software products or features is satisfied by a single declined seat add-on. A respondent who cancelled a seven-figure ERP renewal and a respondent who declined a \$400 a month scheduling tool are indistinguishable inside the 32 percent. ^[1]


ASSESSMENT It does not establish that the respondent knows. McKinsey describes the panel as representing the full range of regions, industries, company sizes, functional specialties and tenures, which is a claim about breadth and not about procurement authority. No published breakdown establishes what share of the 1,719 sit anywhere near a renewal calendar. Respondents are answering for their organization about a decision most of them did not make. ^{[1] [4]}

The sample base, and what the weighting corrects for

FACT McKinsey's Global Survey Panel is a standing community of more than 33,000 executives and managers who have opted in to be polled. 1,719 responses against that frame is a response rate near five percent, and respondents self-select into any given wave. ^[4]

ASSESSMENT This is a non-probability panel, which is normal and disclosed, and it is fine for direction. The problem is the specific correction applied. Weighting by a nation's share of global GDP fixes the geographic skew of who answered. It does nothing about the two skews that matter for this question: the over-representation of people enthusiastic enough about AI to answer an AI survey, and the absence of any weighting by software spend, so a 200-person firm's declined scheduling tool carries the same weight as a bank's declined core platform. ^{[1] [4]}

ASSESSMENT The high-performer cut sharpens the point. McKinsey defines AI high performers as respondents attributing at least 5 percent of EBIT to AI and describing the impact as significant, about 6 percent of the sample. Nearly half of that 6 percent reported skipping a purchase, against 31 percent of everyone else. The people most invested in the AI story are the people most likely to report having acted on it, which is what selection bias looks like when it is working. ^[3]

FIGURE 1 The industry split is a better finding than the average

THIRD-PARTY ESTIMATE Share of respondents reporting their organization declined one or more software purchases because the functionality could be built internally with agentic coding tools. Self-reported procurement intent from a self-selected executive panel, weighted by country GDP. Sector cell sizes are not published, so the tails are the least reliable bars here. ^[1] ^[3]

How the number degrades in transmission

FACT One widely syndicated summary of the same report renders the finding as nearly one third of surveyed companies choosing to build AI coding tools in-house rather than buy software, which describes a different market entirely. Another correctly reports the decision but drops the qualifier that it measures intent rather than outcome. ^[34]


ASSESSMENT This matters more than pedantry. By the time the statistic reaches a board deck or a competitive battlecard it has usually lost the words one or more, products or features, and decided against buying, leaving something that reads as a third of enterprises cancelling

software. Nothing in the survey supports that reading, and the survey's own authors did not offer it. [\[1\]](#) [\[34\]](#)

What twelve vendors disclosed

Every figure in this section comes from a company press release, an 8-K, or an earnings call in the June to September 2026 reporting run. Where a metric is a vendor's own non-standard definition, that is said rather than assumed away.

FIGURE 2 Net retention, most recent disclosed quarter



FACT These are not the same measurement. Snowflake and Datadog report consumption-weighted net revenue retention, GitLab a dollar-based net retention rate, DocuSign a rounded dollar net retention for direct customers, HubSpot and Asana subscription net revenue retention. Datadog disclosed a range in the low 120s rather than a point figure; 121 is the midpoint and is the author's rendering. Compare the direction, not the level. [\[8\]](#) [\[9\]](#) [\[12\]](#) [\[13\]](#) [\[14\]](#) [\[15\]](#)

ASSESSMENT The ordering is the finding. The top of that chart is consumption pricing and the bottom is seat pricing, and the two groups have been separating for three years. Nothing about agentic coding is required to explain it. Consumption vendors bill more when workloads grow, and AI workloads are growing. Seat vendors bill more when headcount grows, and enterprise headcount is not. [\[8\]](#) [\[9\]](#) [\[12\]](#) [\[14\]](#)

The seat-priced vendors, where displacement should show first

FACT HubSpot reported second-quarter 2026 revenue of \$911.7 million, up 20 percent, with customers up to 306,446 from 267,982. Net revenue retention was 102 percent, down one point year over year. Management attributed the softness to budget optimization headwinds and guided full-year net revenue retention roughly flat, with customer dollar retention stable in the high 80s. ^[12]

FACT Asana reported second-quarter fiscal 2027 revenue of \$216.4 million, up 10 percent, with dollar-based net retention of 97 percent, improved from 96 percent and up for a fifth consecutive quarter. Customers spending at least \$100,000 a year retained at 98 percent. ^[14]

ASSESSMENT Asana improving for five straight quarters is close to fatal for the displacement reading of that number. If agentic coding were eating collaboration software, the vendor with the softest retention in the group and the most obviously replaceable product would be getting worse, not better. It is getting better slowly from a low base, which is what a company looks like after a headcount correction has finished working through the renewal book. ^[14]

FACT Atlassian closed fiscal 2026 with fourth-quarter revenue of \$1,766 million, up 28 percent, and cloud revenue of \$1,213 million, up 31 percent, which the company attributed to migrations and seat expansion in Jira and Confluence. Subscription annual recurring revenue reached \$6.6 billion, up 23 percent, and remaining performance obligations \$4.8 billion, up 44 percent. Customers at \$3 million of annual recurring revenue grew 50 percent to 164. ^[10]

ASSESSMENT Atlassian is the cleanest counter-example available. It sells seat-priced developer tooling to exactly the population that has agentic coding tools in hand, and its seats are expanding at 31 percent while its customers deploy agents. Whatever developers are building with agents, they are not building their own issue tracker. ^[10]

The systems of record

FACT ServiceNow reported second-quarter 2026 subscription revenue of \$3,877 million, up 24.5 percent, with current remaining performance obligation of \$13.20 billion, up 21 percent, and a renewal rate of 98 percent. ^[6]

FACT Salesforce reported second-quarter fiscal 2027 revenue of \$11.3 billion, up 11 percent, with current remaining performance obligation of \$33.5 billion, up 14 percent in constant currency, and total remaining performance obligation of \$66.3 billion, up 11 percent. Full-year guidance was raised to a range of \$46.1 billion to \$46.4 billion. ^[5]

FACT Workday reported second-quarter fiscal 2027 subscription revenue of \$2.471 billion, up 13.9 percent, with 12-month subscription revenue backlog up 14.2 percent and total subscription revenue backlog up 8.0 percent. The company attributed the total backlog figure to a shift toward customer-based bookings and to industry mix in net new business, and said AI drives more than 25 percent of new annual contract value. ^[11]

ASSESSMENT Workday's 8 percent total backlog growth is the softest forward number in the set and the one most worth watching, because a headcount-priced human resources system is the

textbook target for this thesis. Management's explanation is a bookings-mix explanation, not a demand explanation, and no analyst on the call pushed them toward a build-displacement answer. Taking the softness as evidence of internal replacement would mean rejecting the company's stated reason with nothing to put in its place. ^[11]

FACT SAP reported current cloud backlog of EUR 22.9 billion for the second quarter of 2026, up 26 percent at constant currencies, and guided to a slight deceleration from the 25 percent recorded in 2025. Cloud gross margin was 74.6 percent against 75.2 percent a year earlier. ^[7]

The one place an analyst named the risk

FACT UBS downgraded SAP from Buy to Neutral on 26 August 2026 while raising its price target to EUR 201, citing slow progress on agentic AI delivery, with 17 of 200 targeted AI agents deployed by year-end, decelerating cloud backlog growth, and the first fall in cloud gross margin before stock-based compensation since 2021. The note flagged the risk of customers pursuing do-it-yourself AI adoption instead. ^{[17] [18]}


ASSESSMENT This is the strongest sell-side articulation of the thesis in the period, and it is worth reading precisely. UBS is not saying SAP is losing renewals to internally built systems. It is saying that if SAP ships agents slowly enough, customers will build their own agents on top of SAP data rather than buy SAP's. That is a threat to attach-rate expansion, not to the core subscription, and it is a very different financial event from a cancelled renewal. ^{[17] [18]}

FACT In the same week, JPMorgan cut Intuit to Neutral with a price target of \$331, down from \$605, and Bank of America cut it to Neutral at \$360, both citing share loss to lower-cost AI-based alternatives spreading from TurboTax into QuickBooks. ^[17]

ASSESSMENT Intuit is the closest thing to a real displacement case in the period and it is not a build case. Customers are not writing their own tax software with agents. They are switching to cheaper AI-native competitors. That mechanism is competitive substitution, which software has always faced, and it has nothing to do with the McKinsey question. Conflating the two is the most common error in the current commentary. ^[17]

FACT C3.ai reported quarterly revenue of \$52.4 million against \$70.3 million a year earlier, a net loss of \$92.8 million, and professional services revenue down to \$3.2 million from \$10.0 million, alongside bookings up 73 percent sequentially and federal bookings up 138 percent year over year. Chief executive Tom Siebel positioned C3 Code, an agentic application builder that turns prompts into enterprise applications, as the growth engine. ^[16]

ASSESSMENT C3.ai is the inverse trade and it is instructive. A vendor whose packaged application revenue fell 25 percent is now selling the tool customers would use to build the replacement. If build displacement is real, the money moves to whoever sells the shovels, and C3's bookings recovery is a weak early instance of that. It is weak because 22 agreements and a sequential bookings number are not a trend, and because a company down 25 percent on revenue has an obvious incentive to reposition around whatever is currently fundable. ^[16]

FIGURE 3 Forward commitments, year-over-year growth in the latest reported quarter

FACT Four different disclosure definitions plotted together: remaining performance obligations, current remaining performance obligations, current cloud backlog at constant currency, and subscription revenue backlog. Levels are not comparable across vendors; growth direction is. Salesforce cRPO is stated at constant currency and includes the Informatica acquisition. ^{[5] [6] [7] [8] [10] [11]}

ASSESSMENT Contracted future revenue is where a wave of declined renewals would appear first, months ahead of reported revenue, because a customer who declines to renew stops adding to the obligation immediately. Every vendor in that chart is adding to it. The slowest is growing at 8 percent and has given a mix-based reason. On the disclosure available in September 2026 there is no vendor-side corroboration of the survey finding at all. ^{[5] [6] [7] [8] [10] [11]}

ASSESSMENT The absence is not proof of absence. Displacement concentrated in vendors too small to be covered, in purchases never made rather than contracts cancelled, or in the long tail of departmental subscriptions below any disclosure threshold, would be invisible in exactly this way. That is the honest limit of this section, and it is why the recommendation below is to watch gross retention rather than net. ^{[5] [12]}

Pricing the build that replaces the licence

The published cost work everyone cites for this comparison is old. The finding that maintenance consumes 60 to 80 percent of software lifecycle cost dates to the maintenance literature of the

1970s and 1980s and has been restated by the IEEE Computer Society since. It predates cloud infrastructure, continuous delivery, SOC 2 as a commercial requirement, and every line of agent-generated code ever written. It is the best benchmark available and it is not a good one.

THIRD-PARTY ESTIMATE The scenario priced here is a 300-seat departmental application replacing a purchased product at \$180,000 a year, which is \$600 per seat per year and typical of mid-market customer relationship management, IT service management or analytics. Three years of licence at flat pricing is \$540,000, or about \$573,000 with a 3 percent annual uplift. [\[27\]](#) [\[28\]](#) [\[29\]](#) [\[35\]](#)

Cost line	Year 1	Year 2	Year 3	Three-year
Engineering, fully loaded	\$375,000	\$250,000	\$250,000	\$875,000
Agentic coding tools and tokens	\$12,600	\$12,600	\$12,600	\$37,800
Security review and penetration test	\$20,000	\$15,000	\$15,000	\$50,000
Compliance evidence and audit	\$45,000	\$35,000	\$35,000	\$115,000
Hosting and infrastructure	\$30,000	\$30,000	\$30,000	\$90,000
Staff turnover and re-ramp	\$30,000	\$30,000	\$30,000	\$90,000
Build total	\$512,600	\$372,600	\$372,600	\$1,257,800
Purchased licence, for comparison	\$180,000	\$180,000	\$180,000	\$540,000

THIRD-PARTY ESTIMATE Engineering assumes 1.5 full-time equivalent years in year one and one thereafter, at a \$250,000 fully loaded onshore cost. That figure sits inside the published range: the Bureau of Labor Statistics puts median software developer pay near \$133,000 with senior roles at \$150,000 to \$185,000, and current benchmarks put the fully loaded cost of a senior United States engineer at roughly \$250,000 to \$265,000 once benefits, payroll tax, equipment, recruiting and overhead are counted. [\[29\]](#) [\[35\]](#)


THIRD-PARTY ESTIMATE Agentic tooling assumes three engineers at \$350 a month all in. Published 2026 benchmarks put GitHub Copilot Enterprise at \$39 per user per month and Cursor Business at \$32, but seat price is not the cost. Once agentic use begins, total monthly spend per developer at thousand-seat scale lands between \$200 and \$500, and Claude Code deployments average \$150 to \$250 per developer per month. This is the line most likely to be understated in a business case and the smallest line in the model either way. [\[28\]](#)

THIRD-PARTY ESTIMATE Compliance and security assume the application falls in scope for a SOC 2 Type II. A first-year programme including readiness, penetration testing, tooling and internal time runs \$30,000 to \$150,000, with the audit fee itself \$30,000 to \$75,000 for mid-market scope and annual renewals at 75 to 90 percent of the initial fee. A SOC 2 penetration test for a standard software-as-a-service scope runs \$8,000 to \$25,000. The \$45,000 used here sits at

the low end and assumes an existing compliance programme the application joins rather than one it creates. ^[27]

THIRD-PARTY ESTIMATE Turnover is priced at one departure across three years. Tech sector tenure runs roughly two to three years against a 4.1 year national average, LinkedIn's 2025 data puts the direct cost of replacing a software engineer near \$43,700 in recruiter fees, signing bonuses and onboarding, and Gallup's replacement range is 50 to 200 percent of salary for senior technical roles. The \$90,000 total is deliberately conservative and excludes the cost that actually bites, which is the departure of the only person who understood why the agent wrote it that way. ^[36]

FIGURE 4 Three years, one 300-seat application



THIRD-PARTY ESTIMATE Thousands of US dollars. The author's model, built from published unit costs, not from observed projects. Onshore engineering rates; an offshore team roughly halves the largest bar and moves the break-even with it. Excludes the opportunity cost of the engineering capacity consumed, which is the largest omission here. ^{[27] [28] [29] [35] [36]}


ASSESSMENT The build lands at roughly 2.3 times the buy, and the annualised run cost of about \$419,000 is the number worth carrying out of this report. A purchased product has to cost more than roughly \$420,000 a year before an onshore in-house replacement is the cheaper option on cash alone, before any adjustment for risk, and that threshold sits above the entire departmental software portfolio of most mid-market organizations. ^{[27] [28] [29] [35] [36]}

ASSESSMENT Note where agents actually help. They compress the engineering line in year one, which is \$375,000 of a \$1.26 million total, or 30 percent. Halving first-version development time saves about 15 percent of the three-year cost and changes no decision. Everything agents currently do well sits in the smallest part of the model, and everything they do badly sits in the largest. ^{[21] [22] [24]}

Why the run cost is the hard part

FACT GitClear's 2026 analysis of 623 million code changes from 2023 to 2026 found refactoring line moves down 70 percent against 2022 levels, cross-file function calls down 35 percent, and long-term legacy maintenance work down 74 percent, while within-commit copy and paste rose 41 percent, duplicated code blocks rose 81 percent, error-masking constructs rose 47 percent, and two-week code churn rose 15 percent. ^[21]

FIGURE 5 Seven code quality signals as AI authorship rises



FACT Percentage change against baseline across 623 million code changes, 2023 to 2026. GitClear sells developer analytics and has a commercial interest in the maintainability argument; the measurements are drawn from public repositories and the method is published, which is more than most competing claims offer. This is correlation with rising AI authorship, not a controlled comparison. ^[21]

ASSESSMENT Read that chart as a maintenance cost forecast rather than a code quality complaint. Duplication and error-masking are precisely the properties that make a change in year two expensive, because a fix has to be found in four places and a swallowed exception has to be diagnosed before it can be repaired. An organization that builds in 2026 at agent speed is buying a year-two and year-three cost profile that nobody has yet priced, because no codebase built this way is three years old. ^[21]

FACT METR's randomized controlled trial of 16 experienced open-source developers across 246 real tasks found them 19 percent slower with early-2025 AI tools, while forecasting a 24 percent speedup beforehand and estimating a 20 percent speedup afterwards. A February 2026 follow-up on later tools showed some evidence of speedup, but METR reported selection effects that made the central estimate unreliable. ^[22]

ASSESSMENT The 39-point gap between measured and perceived productivity is the most important number in this report for anyone approving a build. Every business case for an internal replacement is written by the people who felt 20 percent faster. If the perception gap holds at anything like that size in enterprise settings, build estimates are not merely optimistic, they are optimistic in a way that internal review cannot catch, because the reviewers have the same perception. ^[22]

FACT The 2026 DORA work on return from AI-assisted development models a 39 percent first-year return with an eight-month payback for a 500-person engineering organization, while assuming change failure rate rises from 5 percent to 6 percent after adoption, which alone contributes \$344,000 of negative downtime impact in the model. DORA's 2025 survey of nearly 5,000 technology professionals found 90 percent using AI at work and 30 percent reporting little or no trust in the code it generates. Practitioner commentary in the same period lands on the same point from the quality side, arguing that agent output needs a human review checkpoint and a named owner of the release decision rather than autonomous execution. ^{[23] [24] [33]}

ASSESSMENT DORA's framing, that AI amplifies whatever engineering discipline already exists, is the right lens for a build decision. An organization with mature testing, version control and continuous integration converts agent output into shipped software. An organization without them converts it into rework. Most organizations considering replacing a purchased product with an internal build are doing so because they do not have a large disciplined engineering function, which is the condition under which the amplifier works against them. ^{[23] [24]}

Which categories are exposed and which are not

Exposure here means the probability that a competent internal team with agentic tooling can produce something a business unit will accept in place of the purchased product, within a budget cycle. It is not a judgment about product quality.

Category	Exposure	What decides it
Internal admin tools, forms, approval workflows	High	No external data dependency, no certification, no audit inheritance. Already the target of low-code, and the population most likely to have been bought on a corporate card rather than through procurement.
Reporting and dashboard layers over owned data	High	The data is already inside the organization. The purchased product supplies presentation, and presentation is the thing agents produce most convincingly on a first pass.

Category	Exposure	What decides it
Point integrations and internal API glue	High	Short-lived, low blast radius, and the work agents are demonstrably best at. This is where displacement is real and where the dollars are smallest.
Departmental project and task tracking	Medium-high	Low data gravity and a low bar for acceptance, but Atlassian's 31 percent cloud growth on seat expansion says the incumbents are still winning this in practice.
Customer relationship and marketing automation	Medium	Buildable in outline, but the value sits in deliverability reputation, the connector estate and data accumulated over years, none of which an agent writes.
Observability and analytics platforms	Low-medium	Ingest economics and retention scale defeat the build. The cost that matters is storage and query, not the application, which is why these vendors price on consumption.
Financial close, consolidation and ERP	Low	Audit dependency. External auditors test the control, and a system with no vendor attestation to inherit puts the entire evidence burden on the organization.
Payroll and employment tax	Very low	Jurisdictional rate maintenance is a permanent subscription to somebody else's regulatory research. The software is not the product; the tax tables are.
Clinical, laboratory and regulated life sciences	Very low	Validation under GxP, and a change control regime that makes agent-speed iteration a liability rather than an advantage.
Core banking, payments and insurance policy admin	Very low	Certification and regulatory filing dependency, plus a failure mode that is a regulatory event rather than an outage.

ASSESSMENT Three forces do the protecting and they are not equally durable. Data gravity is the weakest, because it is a cost of moving rather than a prohibition, and it erodes as extraction gets cheaper. Certification is stronger, because it converts a build decision into a multi-year programme with an external gatekeeper. Regulatory dependency is strongest, because it is a permanent obligation to track somebody else's rule changes, and no volume of generated code discharges it. ^{[1] [3]}

ASSESSMENT Switching cost is doing less work here than vendors would like. An organization with agentic tooling can regenerate a great deal of custom integration code cheaply, which is exactly the moat that integration-heavy vendors have relied on. What it cannot regenerate is an attestation, a certification, or a jurisdictional rate table maintained by someone contractually liable for it. Vendors whose defence is complexity should expect that defence to weaken. Vendors whose defence is liability should not. ^{[10] [21]}

FACT McKinsey's own sector split corroborates this ordering without being designed to. Technology at 41 percent and professional services at 38 percent sit at the top, both categories with in-house engineering capacity and light external certification. Insurance at 19 percent and the public and social sector at 17 percent sit at the bottom, both with heavy filing and procurement regimes. ^[3]

HYPOTHESIS Proposition: the categories reporting the highest decline rates are the categories with the lowest average contract value, so the 32 percent headline overstates dollar displacement by a wide margin. This is confirmed if any vendor or analyst publishes declined-purchase rates weighted by contract value and the weighted figure comes in materially below 32 percent. It is refuted if a weighted figure comes in at or above 32 percent, which would mean displacement is concentrated in large contracts and this report's read is wrong. ^{[1] [3]}

The one case study everyone cites

FACT Klarna announced in September 2024 that it was ending its Salesforce and Workday relationships as part of an AI overhaul. Subsequent reporting established that Klarna moved human resources to Deel and assembled customer relationship management from other software-as-a-service products alongside internal work, rather than replacing them with AI-built systems. In March 2025 chief executive Sebastian Siemiatkowski said he did not think other companies would follow, adding that it might be the opposite of the end of Salesforce. ^{[31] [32]}

ASSESSMENT The canonical build-displacement case turned out to be a vendor consolidation. That is worth holding in mind whenever a new one is announced, because the announcement and the correction get very different amounts of attention, and the correction is where the operational detail lives. ^{[31] [32]}

Scenarios to 2028

The probabilities below are the least reliable content in this report and are the author's subjective weights. The earliest visible signs are the useful part, because each is observable without private information.

SCENARIO A Edge erosion, core intact 50 percent

Agentic builds absorb internal tools, glue code and reporting layers, while systems of record renew normally. Aggregate software spend keeps growing because AI-attached expansion at incumbents exceeds the departmental tail lost to internal builds.

Consequence. Seat-priced collaboration and departmental vendors see gross retention decline by one to three points while systems-of-record vendors hold. Total software spending growth stays above 10 percent.

Earliest visible sign. A second-tier collaboration vendor guides down on gross retention specifically, without a corresponding macro or headcount explanation.

SCENARIO B The run cost lands 25 percent

Organizations that built in 2025 and 2026 hit year-two maintenance on agent-generated codebases carrying the duplication and error-masking profile GitClear measures, and quietly re-buy. Build displacement peaks and partially reverses.

Consequence. A visible cohort of return-to-vendor deals, and a shift in agentic tool marketing from generation toward maintenance and comprehension of existing code.

Earliest visible sign. A vendor names win-back from customer-built replacements as a growth source on an earnings call, which would be the first time either direction has been quantified publicly.

SCENARIO C Displacement becomes visible in disclosure 15 percent

The survey answer converts into contract action at scale during the 2027 renewal cycles, and at least two covered vendors attribute a named churn cohort to customer-built replacements.

Consequence. Gross retention deteriorates across seat-priced application vendors simultaneously, and a build-displacement risk factor appears in multiple annual filings rather than in one.

Earliest visible sign. The first specific attribution on an earnings call. This is the tripwire that would overturn the call in this report.

SCENARIO D Substitution, not construction 10 percent

The competitive damage arrives from AI-native vendors underpricing incumbents rather than from customers building, which is the mechanism already visible in the Intuit downgrades. Packaged software is disrupted, but by other packaged software.

Consequence. Incumbent pricing pressure and margin compression without any deterioration in build-versus-buy survey answers, and a widening gap between incumbent and AI-native vendor growth rates.

Earliest visible sign. A second large incumbent downgraded specifically on share loss to an AI-native competitor rather than on AI delivery execution.

Leading indicators worth monitoring

Each of these is observable from public disclosure. The first column is the event, and the last says where to look for it.

If this happens	It means	Moves toward	Where to watch
A vendor separates gross retention from net retention in disclosure	Expansion is masking churn and management expects to be asked	Scenario C	Quarterly supplements
A build-displacement risk factor appears in two or more annual filings	Counsel now considers it material and disclosable	Scenario C	Form 10-K risk factors
ServiceNow renewal rate drops below 97 percent	The most durable renewal metric in enterprise software has moved	Scenario C	Quarterly press release

If this happens	It means	Moves toward	Where to watch
Atlassian cloud seat expansion decelerates below 15 percent	Developers are displacing the tooling they use to build	Scenario C	Quarterly shareholder letter
Gartner cuts the software segment below 10 percent growth	Aggregate spend is finally reflecting withdrawn purchases	Scenario C	Quarterly IT spending forecast
Agentic coding vendors reposition marketing toward legacy comprehension	The industry has met the year-two maintenance wall	Scenario B	Vendor product launches
A vendor names win-back from internal builds as a growth source	The reversal has become large enough to sell against	Scenario B	Earnings call transcripts
McKinsey's 2027 wave reports this figure falling	The 2026 reading was enthusiasm rather than plan	Scenario B	State of AI, mid-2027
A second incumbent is downgraded on AI-native share loss	The damage is substitution, not construction	Scenario D	Sell-side rating changes
A breach disclosure names an internally built AI application	The security cost is arriving ahead of the licence saving	Scenario A	Form 8-K, state filings

Implications

For CIOs approving application renewal cycles

ASSESSMENT Carry the \$420,000 threshold into the renewal conversation. Below roughly that annual licence value, an onshore internal replacement does not pay over three years on the model here, and the burden of proof sits with whoever is proposing to build. Above it, the build deserves a real evaluation with the run cost fully loaded and the compliance scope settled in advance rather than discovered in the second year. ^{[27] [28] [29] [35]}

ASSESSMENT The more urgent exposure is not the renewal you might cancel. It is the applications your teams have already built without telling you. IBM found shadow AI present in 43 percent of breached organizations against 20 percent a year earlier, and 68 percent of breached organizations had no policy governing AI use at all. Run the inventory before you run the build-versus-buy analysis, because the second is a planning exercise and the first is a live risk. ^{[25] [26]}

ASSESSMENT Where you do build, price the exit. The failure mode is not the application that does not work, it is the application that works, becomes load-bearing, and then loses the two people who understood it. Require a named maintenance owner and a documented decommissioning path as a condition of approval, and expect both to be harder to secure than the build budget. ^{[21] [36]}

For software vendor CFOs modelling seat attrition

ASSESSMENT Nothing in the disclosure record justifies modelling build-driven attrition as a distinct line yet. What it justifies is separating gross from net retention in your own internal reporting immediately, because expansion is currently masking whatever is happening underneath, and by the time it stops masking it you will be explaining a trend rather than getting ahead of one. [\[12\]](#) [\[14\]](#)

ASSESSMENT The exposed revenue is the seat nobody uses. Displacement pressure is arriving first as seat rationalization dressed as budget optimization, which is what HubSpot and Asana are describing. Audit your own low-utilization seat base and price it as at risk, because that is the population a customer can eliminate without replacing anything at all, let alone building it. [\[12\]](#) [\[14\]](#)

ASSESSMENT If your defence is integration complexity, it is weakening. Agentic tooling regenerates custom connector code cheaply, which is the specific moat that erodes fastest. If your defence is an attestation, a certification, or a maintained regulatory dataset, it is holding, and it is worth making that dependency more visible in how you package and price rather than assuming buyers already understand it. [\[10\]](#) [\[21\]](#)

For equity analysts covering enterprise software

ASSESSMENT Treat the 32 percent as a sentiment indicator and not a demand indicator, and say so when it appears in a client question. The measured object is a stated procurement decision from a self-selected panel weighted by country GDP, with no contract value attached and no verification that anything shipped. It has real information content about intent and none about revenue. [\[1\]](#) [\[4\]](#)

ASSESSMENT The metric that would actually settle this is gross retention, and most vendors do not disclose it. Ask for it. A vendor with 120 percent net retention and deteriorating gross retention is losing customers and hiding it inside expansion, and that is the only pattern consistent with quiet build displacement. ServiceNow's 98 percent renewal rate is the cleanest disclosure in the group precisely because it is a gross measure. [\[6\]](#)

ASSESSMENT Keep the two theses apart in the models. Build displacement and AI-native substitution have opposite implications for total category spend: the first shrinks it, the second reallocates it. The Intuit downgrades are substitution and the UBS note on SAP is closer to attach-rate risk. Neither is evidence for the build thesis, and both have been cited as though they were. [\[17\]](#) [\[18\]](#)

For procurement leads pricing build alternatives

ASSESSMENT When a business unit brings you a build proposal, ask for four numbers the proposal will not contain: the fully loaded run cost in years two and three, the compliance scope and who signs the evidence, the named maintenance owner, and the decommissioning cost. In the model here those four together are 70 percent of the three-year total, and a proposal built on first-version effort has priced the smallest part of the problem. [\[27\]](#) [\[28\]](#) [\[29\]](#)

ASSESSMENT Use the threat credibly rather than loudly. The build alternative is a real negotiating position now in a way it was not in 2024, and vendors know the survey number as well as you do. It is most credible on renewals above roughly \$420,000 a year, on products with light certification exposure, and where you can name the team that would do the work. Deployed outside those conditions it is a bluff a competent account executive will call. ^{[27] [28] [29] [35]}

ASSESSMENT Forrester's January 2026 buying research puts procurement in the decision seat in 53 percent of business buying cycles, with 13 internal stakeholders on a typical purchase. That is leverage, and it also means a build decision made inside one engineering team has bypassed a governance process a purchase would have gone through. Treat internal builds as procurement events and route them accordingly. ^[30]

For engineering managers owning internal builds

ASSESSMENT The METR result is the one to internalize: developers who were 19 percent slower believed they were 20 percent faster. That gap does not get caught by internal review, because the reviewers share the perception. If you are estimating a build, take the estimate your team produces and treat the confidence attached to it as uninformative rather than reassuring. ^[22]

ASSESSMENT DORA's amplifier finding sets the honest entry condition. If you have automated testing, disciplined version control, continuous integration and small batch sizes, agent output converts into shipped software and the build is worth considering. If you do not, the same output converts into rework and change failure, and the correct answer is to buy the product and spend the capacity on the foundations instead. ^{[23] [24]}

ASSESSMENT Instrument for duplication and churn from day one, not from the first incident. The measurable signature of the problem is public: duplicated blocks, error-masking constructs and two-week churn all rising while refactoring falls. Those are trackable in your own repository from the first commit, and they are the earliest warning available that a build is accruing a year-two cost nobody has budgeted. ^[21]

What this document cannot answer

Four questions matter more than anything settled above, and none can be answered from public sources.

What is the dollar value of the declined purchases? The survey counted respondents, not contracts. Without contract values attached, 32 percent of organizations declining one or more products is compatible with almost any total, from a rounding error to a material share of category spend. Answering it needs either a re-run of the instrument with value bands, or procurement data from a panel of enterprises.

How many of the declined purchases produced a working replacement? The question measured intent in May and June 2026 and nobody has gone back. A follow-up wave asking the same respondents what actually shipped, what it cost, and whether it is still running would be the single most valuable piece of primary research in this category, and it would take about six weeks.

What does a three-year-old agent-generated codebase cost to maintain? Nobody knows, because none exists. Every maintenance figure in this report, including the model, is extrapolated from human-authored code or inferred from quality signals. The first honest answer arrives in 2027 or 2028 and will come from organizations that built early and are willing to say what happened.

What is happening below the disclosure threshold? Twelve public vendors were checked. The categories this report identifies as most exposed, internal tools and departmental subscriptions, are largely sold by private companies that disclose nothing. If displacement is concentrated there it is invisible to this method by construction, and the way to see it is a spend-analysis panel across mid-market software budgets rather than more filings.

Half-life of this assessment

Six months. Every vendor figure here is a point-in-time reading from the June to September 2026 reporting run and is superseded by the next quarter. The retention numbers, the backlog growth rates and the sell-side positions on SAP and Intuit should all be re-read against the December 2026 and March 2027 results before anything in this report is quoted. The scenario probabilities are calibrated to what was disclosed as of early September 2026 and will go stale first.

Twelve months. The McKinsey figure itself, which will be re-fielded in the 2027 wave, and the read on whether the 2026 answer converted into contract action. The three-year cost model's unit prices, particularly agentic tooling, where the pricing model has already moved once in 2026 from per-seat toward consumption billing and is not settled. GitClear's quality signals, which will either continue to deteriorate or stabilize as tooling improves, and either outcome changes the maintenance argument materially.

Twenty-four months. The category exposure ordering. Data gravity erodes slowly and certification regimes change on legislative timescales, so the table separating exposed from protected categories should hold in shape even as individual rows move. The first real maintenance cost data on agent-generated codebases also arrives inside this window and is the thing most likely to change the build economics in either direction.

Architectural, and unlikely to decay. That maintenance dominates software lifecycle cost. That an audited system needs someone accountable for the evidence. That regulatory rate maintenance is a subscription to somebody else's ongoing work rather than a piece of code. These have survived every previous cycle where a new tool made writing software cheaper, including fourth-generation languages, offshore development and low-code, and there is no mechanism in agentic coding that repeals them.

Limitations

This report was produced independently. The author holds no position in and receives no compensation from any vendor named, and no vendor reviewed the analysis before publication or was given advance sight of it.

The primary McKinsey report was not directly retrievable during the research window. Its methodology figures, the 32 percent finding and the industry breakdown are drawn from secondary reporting of the survey, cross-checked across four independent summaries that agree on the sample size, the field dates, the headline figure and the sector splits. The exact question wording quoted in this report comes from one of those summaries rather than from a published question set, and no full question set appears to be public. That is the weakest link in the first section and it is load-bearing, because the argument turns on what the question asked.

Vendor figures are drawn from company press releases, 8-K filings and earnings coverage rather than from complete transcript review. A build-displacement remark made in an unscripted answer and not picked up in coverage would have been missed. Twelve vendors is a defensible sample of the public enterprise, analytics and developer tooling categories, and it is not the whole category.

The three-year cost model is the author's construction from published unit costs. It has not been validated against observed projects, because no dataset of completed agentic replacement builds with full three-year cost histories exists. Reasonable people will move the engineering, compliance and turnover assumptions, and the model is presented line by line so those changes can be made without rebuilding it. Offshore staffing moves the largest line by roughly half. The model excludes the opportunity cost of the engineering capacity consumed, which is the most consequential omission in it.

Absence of vendor corroboration is treated here as meaningful evidence, and it carries a specific weakness. Vendors do not disclose churn causes they cannot cleanly attribute, and a customer who declines to renew rarely tells the vendor it was because they built the replacement. A real effect could be present in the numbers and absent from the explanations, which is why the tripwires in this report favour gross retention over anything management says about it.

Sources

- [1] McKinsey & Company. "The State of AI: Global Survey 2026." QuantumBlack, 25 August 2026. *Analyst firm* mckinsey.com
- [2] ETCIO / The Economic Times. "Organizations build software in-house using agentic coding tools instead of buying." 30 August 2026. *Independent press* cio.economictimes.indiatimes.com
- [3] The Daily Brief. "A Third Skipped a SaaS Buy. Now Price the Run Cost." beri.net, August 2026. Source of the quoted question wording and the sector breakdown. *Second-hand (unverified)* beri.net
- [4] McKinsey & Company. "McKinsey Global Surveys." Panel description and survey practice. *Analyst firm* mckinsey.com
- [5] Salesforce, Inc. "Salesforce Delivers Record Second Quarter Fiscal 2027 Results." 26 August 2026. *Vendor primary* salesforce.com
- [6] ServiceNow, Inc. "ServiceNow Reports Second Quarter 2026 Financial Results." Form 8-K, July 2026. *Vendor primary* sec.gov
- [7] SAP SE. "SAP Quarterly Statement Q2 2026." 23 July 2026. *Vendor primary* sap.com
- [8] Snowflake Inc. "Financial Results for the Second Quarter of Fiscal 2027." Form 8-K, 2 September 2026. *Vendor primary* sec.gov
- [9] Datadog, Inc. "Datadog Announces Second Quarter 2026 Financial Results." August 2026. *Vendor primary* datadoghq.com
- [10] Atlassian Corporation. "Fourth Quarter and Fiscal Year 2026 Shareholder Letter." Form 8-K, August 2026. *Vendor primary* sec.gov

- [11] Workday, Inc. "Workday Announces Fiscal 2027 Second Quarter Financial Results." 27 August 2026. *Vendor primary* workday.com
- [12] HubSpot, Inc. "HubSpot Reports Q2 2026 Results." 5 August 2026. *Vendor primary* hubspot.com
- [13] GitLab Inc. "GitLab Reports Second Quarter Fiscal Year 2027 Financial Results." 1 September 2026. *Vendor primary* gitlab.com
- [14] Asana, Inc. "Asana Announces Second Quarter Fiscal 2027 Results." 3 September 2026. *Vendor primary* asana.com
- [15] Investing.com. "DocuSign Inc Q2 2027 Earnings Call Highlights." September 2026. *Independent press* investing.com
- [16] Lauchlan, Stuart. "How's C3.ai's turnaround going? Some signs of life in the old dog yet." *diginomica*, 3 September 2026. *Independent press* diginomica.com
- [17] Yahoo Finance. "5 big analyst AI moves." 29 August 2026. Covers the UBS downgrade of SAP and the JPMorgan and Bank of America downgrades of Intuit. *Independent press* finance.yahoo.com
- [18] Investing.com. "SAP dips as UBS downgrades stock on lack of AI catalysts, weak investor sentiment." 26 August 2026. *Independent press* investing.com
- [19] Gartner. "Gartner Forecasts Worldwide IT Spending to Grow 14.2% in 2026, Totaling \$6.37 Trillion." 27 July 2026. *Analyst firm* gartner.com
- [20] Retool. "The Build vs. Buy Shift: How Vibe Coding and Shadow IT Have Reshaped Enterprise Software." February 2026. 817 respondents, all Retool customers and builders, surveyed in late 2025. *Vendor-published comparison (commercial interest)* businesswire.com
- [21] GitClear. "The Maintainability Gap: 2026 AI Code Quality Research." 623 million code changes, 2023 to 2026. *Practitioner aggregate (self-selected)* gitclear.com
- [22] METR. "Measuring the Impact of Early-2025 AI on Experienced Open-Source Developer Productivity." 10 July 2025, with a February 2026 follow-up. *Independent press* metr.org
- [23] DORA and Google Cloud. "State of AI-assisted Software Development 2025." Nearly 5,000 respondents. *Analyst firm* dora.dev
- [24] InfoQ. "New DORA Report Claims Strong Engineering Foundations Drive AI Return on Investment." 11 May 2026. *Independent press* infoq.com
- [25] IBM. "Cost of a Data Breach Report 2026." 602 breached organizations, March 2025 to February 2026, 17 industries and 16 countries. *Vendor-published comparison (commercial interest)* ibm.com
- [26] Cybersecurity Dive. "As data breaches grow costlier, ungoverned AI creates new risks." 2026. Source of the shadow AI incidence and cost figures. *Independent press* cybersecuritydive.com
- [27] Scrut. "How much does SOC 2 compliance cost in 2026? A practitioner's breakdown." 2026. *Third-party comparison (commercial interest possible)* scrut.io
- [28] DX. "AI coding assistant pricing and ROI guide (2026): costs, benchmarks, and what the data shows." 2026. *Third-party comparison (commercial interest possible)* getdx.com
- [29] US Bureau of Labor Statistics. "Software Developers, Quality Assurance Analysts, and Testers." Occupational Outlook Handbook. *Standards body primary* bls.gov
- [30] Forrester Research. "Forrester's 2026 Buyer Insights: The State of Business Buying." 21 January 2026. *Analyst firm* forrester.com
- [31] CX Today. "Klarna Didn't Replace Salesforce and Workday with AI; It Replaced Them with Alternative SaaS Apps." December 2024. *Independent press* cxtoday.com
- [32] TechCrunch. "Klarna CEO doubts that other companies will replace Salesforce with AI." 4 March 2025. *Independent press* techcrunch.com
- [33] Ferreira, Helder. "What agentic actually means for software quality, beyond the buzzwords." *DevPro Journal*, 1 September 2026. The author is a director of product management at Sembi. *Independent press* devprojournal.com
- [34] Value Add Pulse. "McKinsey enterprise AI ROI 2026: survey data, key stats." 2026. Cited only as an example of the finding degrading in transmission. *Second-hand (unverified)* valueaddvc.com
- [35] NextGen Coding Company. "Engineer Cost Index 2026: Onshore vs Nearshore vs Offshore." 2026. *Third-party comparison (commercial interest possible)* nextgencodingcompany.com
- [36] Rockstar Developer University. "Developer Turnover and Retention Statistics." 2026. Aggregates LinkedIn and Gallup figures on tenure and replacement cost. *Second-hand (unverified)* rockstardeveloperuniversity.com

Rights and permitted use

© 2026 David Soden, davidsoden.com/market-assessment. All rights reserved.

This document, including its structure, analysis, findings, evidence classification, and framing, is the original work of David Soden, published at davidsoden.com/market-assessment.

It is published for public reading. You are free to share the complete, unmodified file, to link to it, and to quote short passages in commentary, reporting, or discussion, provided the author and the site above are credited.

The following require prior written consent: republishing this work in whole or in substantial part under another name, masthead, or brand; excerpting or modifying it in a way that alters the analysis or removes the evidence classifications; incorporating any portion of it into a paid, commissioned, or client-facing deliverable; resale or distribution behind a paywall; and use of this document as training data for machine learning models.

Commissioned Market Assessment reports, commercial licensing, and briefings on this material are available.

This document is analysis, not investment, legal, or procurement advice. It was not commissioned, reviewed, or funded by any company named in it, and the author holds no position in any of them.

Market Assessment reports, commissioned research, and briefings: davidsoden.com/market-assessment