
WHITE PAPER · NO. 01 · SEPTEMBER 2026

The Shifts in Software

A LiteByte White Paper

On the economics of building business software,
and why the threshold for what is worth building
has moved.



01 THE OLD ARITHMETIC

Fixed and marginal. One cost is paid once whatever happens next. The other is paid again for every customer who arrives. The wider the distance, the stronger the pull towards one product sold to everybody.


Software was expensive, so it had to be generic

Nobody set out to sell a company software that only half fits it. The generic product was not a failure of imagination. It was the answer to an arithmetic problem that held for most of four decades.

Building a serious business application costs a great deal before a single customer touches it. Copying it costs close to nothing. The OECD puts the position directly: software and other data dependent technologies combine high fixed costs with very low costs per additional unit, which is why they show such large economies of scale.¹

Every commercial decision in the industry follows from that shape. If the build is expensive and the copy is cheap, you build once and sell often. You standardise, because variation puts the expensive part back in. You chase volume, because volume is what repays the build. Multitenancy, the architecture behind most cloud software, made the logic concrete: one running copy of a product serves thousands of customers at once, so the cost of building and operating it is spread across all of them. AWS and Microsoft still describe their own architecture in those terms, as a set of trade offs between cost, isolation and how far a product can bend for any one customer.²

FIGURE 1 · THE LOOP THAT BUILT THE INDUSTRY



THE AMORTISATION LOOP. Each step exists to protect the step before it. Standardisation is not a design preference; it is what keeps a large one off investment recoverable.

One clarification, because the shorthand is usually wrong. The marginal cost of serving another customer was never actually zero. Cloud capacity, storage, support, customer success and now model inference all rise with use. The real contrast was between a large investment made once and a much smaller one made repeatedly.³ That gap was wide enough to organise an industry around.

02 THE BARGAIN

What large means here. A thousand employees or more in the ONS data, ten thousand or more in the Census data. For forty years the ability to get software that fits has tracked headcount and balance sheet. This paper is about why that is changing.

What the customer gave up

Every business that bought a standard product accepted the same trade, usually without naming it: capability you could never have funded alone, inside a product designed for a market rather than for you.


Who actually accepted it depended on size, and the two ends of the market barely resemble each other.

At the top, the largest organisations never really accepted it, because they have always built their own. ONS puts UK investment in computer software and databases at £57.2 billion in 2023, of which £42.9 billion was own-account: software an organisation produced for itself rather than bought.⁴ Three quarters of the spend goes on building, not buying.

Where they do buy, they buy the product and then pay to bend it. Configuration is included; anything real is a change request. Supplier rate cards on the UK government's Digital Marketplace put senior developers between £450 and £600 a day, so a change needing a few months of specification, build, testing and release lands near £30,000 before anyone has argued about scope.⁵ For a large firm it is a line item on a page nobody reads.

At the other end, smaller firms get neither route. There is no in-house team to build with and no budget to bend what they buy, so they take the product as shipped, work around it in spreadsheets, or never buy. The SME Digital Adoption Taskforce found they see products as built for larger organisations, returning to cost, relevance, switching and implementation.⁶ The OECD's 2026 UK review lands in the same place: cost, appropriateness, skills and trust.⁷

FIGURE 2 · SOFTWARE ADOPTION BY FIRM SIZE, EU ENTERPRISES, 2025



SOFTWARE ADOPTION TRACKS HEADCOUNT. The same survey, the same year, three categories of business software. Source: Eurostat, Larger enterprises used more e-business apps in 2025, dataset isoc_eb_iip.

Eurostat's 2025 enterprise survey puts the gap on the software itself.⁸ Among EU enterprises with 250 or more employed, 89% run ERP, 65% CRM and 69% business intelligence. Among those with ten to forty nine, the figures are 41%, 25% and 11%.⁹

03 WHERE THE COST FELL

A tool, not an outcome.

The same class of assistant produced a 55% saving in one trial and a 19% loss in another. What separated them was not the model. It was the people using it, and how practised they were with it.

Where the cost actually fell


Building software got dramatically cheaper, and the evidence for that is stronger than the evidence for almost anything else in this paper.

GitHub's controlled experiment on a bounded task measured a 55% cut in completion time.¹⁰ Microsoft led three randomised field experiments covering 4,867 developers at Microsoft, Accenture and a Fortune 100 electronics manufacturer, and found a 26% increase in completed tasks among those given AI assistance.¹¹ These are not vendor claims. They are randomised trials run inside real engineering organisations, and they point the same way.

What they do not say is that the saving turns up on its own. METR randomised experienced open source maintainers and found them 19% slower with AI, and slower while believing they had been faster.¹² Most had used the tool for only a few dozen hours, and METR itself cannot rule out that the benefit takes longer than that to arrive. The result gets quoted as a rebuttal. It is better read as the obvious thing about any powerful tool: it is only as good as the hands it is in.

FIGURE 3 · THREE RANDOMISED TRIALS, ONE KIND OF TOOL

MEASURED EFFECT OF THE AI ASSISTANT



THE TOOL IS ONLY AS GOOD AS THE HANDS IT IS IN. Three randomised studies of AI coding assistants, ordered by result.

DORA's framing, drawn from close to five thousand professionals, is the right one: AI amplifies the engineering system around it, strong or weak.¹³ In practised hands inside sound teams it delivers what the trials recorded. In unpractised hands it can cost time.

The cost of producing software built for one business has fallen a long way. It falls furthest in the hands of engineers who know how to use the tools, which is why the judgement did not disappear.

That matters because the work that got cheaper includes exactly the work that used to be too expensive to justify. Implementation was never the whole bill, and lifecycle cost still covers architecture, data, integration, assurance and security.¹⁴ The frontier moves anyway, because it does not need every line of that list to fall.

04 THE FRONTIER

What sits inside cost.

Discovery, architecture, implementation, data, integration, assurance, security, deployment, maintenance, compute and risk. AI compresses several. It leaves others where they were.

The line that moved

Every piece of software a company could build sits on one test. Build it when the expected benefit, discounted, is greater than the lifecycle cost, discounted. Below that line it is a bad idea however good it sounds in a meeting.

The benefit side did not change. The cost side did, and a band of projects that were always technically possible and never economically sensible crossed over. In reality benefit has probably risen too, because AI makes software more useful, agentic systems most of all, but that is a separate argument, taken up in The long run shift.

FIGURE 4 · THE SOFTWARE VIABILITY FRONTIER



LOWER COST DOES NOT CREATE VALUE. Opportunities are ranked by the value they would create. Cheaper building makes none of them more valuable; it lets a band of them clear the threshold for the first time.

A worked example makes it concrete. Suppose a workflow specific to your business would generate £20,000 of value across its life: hours recovered, errors avoided, a little margin protected. At a £60,000 build cost, refusing to build it is the correct decision. At £6,000 it becomes correct to build it. Nothing about the workflow changed. The frontier did.

The interesting category is not software that became cheap. It is software that was always possible and never worth building.

This applies at more than one scale, which is the part most often missed. An entire bespoke system can cross the line. So can a single feature inside a product you already own.

Which is why “how much cheaper is it?” is the weaker question. An eighty per cent saving on something nobody was going to build is worth nothing. A thirty per cent saving that carries a project over the threshold changes what a business can do.

05 RENT, BEND OR OWN

Rent it, bend it, or own it

Customisation is where the cost change lands hardest, and it is the part of the argument most often put badly.

Not a discount. The point is not that the same thing got cheaper. It is that the same budget now reaches a different object, and you keep what it produces.

For twenty years, fit had a price list. You rented the standard product, and where it did not match how you worked you paid the vendor or a partner to change it, one feature at a time, at rates set by what that work genuinely cost. There was no third column, because building the whole thing cost far more than bending somebody else's.

FIGURE 5 · WHAT THE MONEY BUYS

THREE WAYS TO GET SOFTWARE THAT FITS

Rent it as it comes	Rent it and bend it	Build it and own it
WHAT YOU PAY A subscription per seat, paid for as long as you use it	WHAT YOU PAY The subscription, plus a change request for anything real	WHAT YOU PAY A project cost, then hosting and the cost of changing it
WHAT IT FITS Fits the average of the market it was built for	WHAT IT FITS Fits closer, never exactly, and drifts on every upgrade	WHAT IT FITS Fits the process you actually run
WHAT YOU HOLD At the end you hold a renewal date	WHAT YOU HOLD At the end you hold the same renewal date	WHAT YOU HOLD At the end you hold the code, the data and the rights

The middle column was the only route to fit for most of the last two decades, and it was priced for firms that could carry it.

THE DISTANCE BETWEEN THE COLUMNS IS WHAT MOVED. When producing working software gets cheaper, the budget that used to buy a handful of change requests starts to cover a system built around the process itself.

That last part deserves its own sentence. A subscription buys continuing access to somebody else's product for as long as you keep paying. A build can leave you holding the code, the data and the rights to both. For the money a mid-sized firm already spends renting and bending, buying and owning is now a real option rather than a rhetorical one.

This is the most profound shift in software we are seeing. Building your own platforms and automations was unheard of outside the largest firms, because the price was simply too high. What was left was a standard product off the shelf, or the cost of bending somebody else's. For close to the same money, a business can now have software designed for it.

That is what we call Truly Tailored at LiteByte. We work with you to design and produce software at a low cost, shaped around how your business actually runs, and you hold the IP rights and the ownership at the end of it.

06 TWO MARKETS

Not the same as “SaaS is dead”. Standardised products stay economically superior wherever the requirement is genuinely common. What moves is the boundary, not the model.

Consumption consolidates, production fragments


Two things that sound contradictory happen at once, because they are the same cost change seen from opposite ends of the market.

Most companies run several products bought from several companies, because each capability came from whoever standardised it best and none were built to speak to one another. HMRC opened a call for evidence on exactly that in March 2026, describing businesses running accounting software alongside CRM, point of sale, e-commerce, booking and payments.¹⁵

The change coming is not five products becoming one with a shorter vendor list. It is that the connective layer stops being rented at all. As more of it is built for the business and handed over, five subscriptions from five companies become one system the company holds. Capability comes back in house, liaising with five account managers stops being anybody’s job, and the cost falls with it. LiteByte positions itself as the partner for that move from renting to owning, keeping the expertise a business still needs in place.

Unification is where the operational gain sits. Systems bought one at a time cannot talk to each other, and the vendors who sold them had little reason to let them. Once the layer is the business’s own, the records join up and a question can cross all of them at once.

FIGURE 6 · RENTED STACK, OWNED LAYER



TWO DIFFERENT KINDS OF SOFTWARE. The layer that encodes how a business runs can come back in house. The layer whose specification is written by a regulator has no reason to.

There is a limit here and it is a sharp one. Where the specification is written by somebody other than the customer, every firm genuinely needs the same thing. Banking, insurance, payroll, tax filing, clinical records: the requirement comes from a regulator, so that software is standardised because the obligation is, not because anybody compromised. Cheaper production should push those markets towards deeper specialisation inside each vertical, not towards every firm writing its own. Tailoring belongs where the requirement is the business’s own: how it sells, quotes, schedules and reports.

07 THE LONG VIEW

Where this ends up. Not a new screen. A system you speak to rather than learn your way around, with the value sitting in the model underneath rather than the pages on top.

The long run shift

Click, search, ask, delegate. Four ways to get something out of a computer, in the order the industry moved through them. The last two change what software has to be. Our view is that more of it will be built for an AI layer to work, not only a person.

Clicking means learning somebody's information architecture: menus, tabs, a screen for each thing and a training session to explain where everything lives. Searching means the destination still has to exist; you just reach it faster. Both assume somebody decided in advance what you would want and built a place to keep it.

Asking does not. Power BI Copilot, Looker Conversational Analytics and Tableau Agent now take a question rather than a route to the answer.¹⁶ Delegating goes further again, to an agent that holds permissions and takes actions inside the system instead of returning a chart. Salesforce is already selling voice agents wired into CRM records and workflow.¹⁷

The economics mirror the viability frontier exactly. Business intelligence asks a question before you do: will this be asked often enough to justify a report? Question, query, metric, filter, chart, dashboard. Every step is a small fixed cost paid in advance, which is why the unusual question never gets built for. A conversational layer assembles part of that chain at the moment of asking, so the long tail of questions stops being uneconomic.

FROM OUR OWN WORK

We built an operating layer of this kind for an estate and hospitality client: customer records, bookings, quoting and delivery in one system rather than four, with a conversational agent sitting on the data model. Staff ask it questions in ordinary language and get answers from live records, including questions nobody had built a report or a feature for.

Three things follow from that. The first is reach: a screen only shows what somebody decided in advance to build a screen for, so a question had to earn a button, a tab or a chart before anyone could ask it, whereas an AI layer on the data model answers the question you actually have.

The second is speed, because a layer like this does not only answer questions, it carries out the work behind them. If someone asks it to raise an employee's allocation on a project by ten per cent, it makes that change wherever the change belongs, while a person is still deciding which system to open.

The third is that safety has to be designed in from the beginning. Permissions, scope and access rules belong in the architecture rather than bolted on afterwards, and getting them right takes somebody who understands the system underneath. This is where the judgement matters more, not less.

08 WHAT FOLLOWS

What stays scarce.

Architecture, data design, security, regulatory reading, and working out which problem is worth solving at all. The cheaper the building gets, the more of the job that is.

What becomes scarce

If code becomes abundant, value moves to whatever stays scarce. Architecture. Data design. Threat modelling. Identity and permissions. Regulatory interpretation. Integration. Verification. Domain understanding. And, before any of it, working out which workflow is worth encoding at all.

Security research gives the concrete version. Pearce and colleagues found substantial vulnerability rates in generated code under security relevant prompts, and Perry and colleagues found assisted participants writing less secure code while feeling more confident about it.¹⁸

There is a commercial consequence we take seriously. When implementation compresses, learning speed becomes the binding constraint. The question stops being how fast a team can build the requested feature and becomes how fast it can work out which feature was worth requesting.

Where this could be wrong

- 1 The saving depends on the hands it is in.** Three randomised trials are not a law, and a team that has not learned to work with these tools will not see the gains the trials recorded.
- 2 Cost can move rather than disappear.** Faster generation pushes work downstream into review, debugging, security remediation and technical debt. Plausible looking code is not production grade code.
- 3 Agentic systems have to be configured carefully.** Architecture, permission scope and access rules are what keep an agent inside the work it was given, and the consequences of getting that wrong can be severe. The ICO flags over broad access, unnecessary database connectivity and the complexity of automated decisions; the EU AI Act, the Cyber Resilience Act and UK data law raise the governance requirement rather than lower it.¹⁹

The old economics rewarded software designed once and rented to as many organisations as possible. Cheaper building does not abolish that model. It narrows the range of situations in which it is the only sensible one. As the threshold falls, more businesses can justify software shaped around their own work and keep what it produces (once a strategy reserved for the largest firms), more suppliers can serve narrower markets, and an AI layer over the top becomes a faster, more capable way of working than the screens it replaces.

These are the shifts we think are worth watching, and worth acting on. They are underway now across software and the wider technology sector, and the decisions they change are ones most businesses are already facing.

NOTES

Nineteen citations carry the argument. The wider source bank that follows exists so that every statement here can be traced, and so that weaker claims could be removed before publication rather than after.

Notes

1. OECD, *Digital Economy Outlook 2024, Volume 1*. oecd.org/en/publications/oecd-digital-economy-outlook-2024-volume-1_a1689dc5-en
2. AWS Well-Architected SaaS Lens, *Silo, Pool and Bridge Models*; Microsoft Azure Architecture Center, *Tenancy models for a multitenant solution*. docs.aws.amazon.com/wellarchitected/latest/saas-lens/silo-pool-and-bridge-models.html
3. OECD, *Digital Economy Outlook 2024*, on fixed versus incremental cost; see also OECD, *Vectors of Digital Transformation*.
4. ONS, *Investment in intangible assets in the UK: 2023*. Computer software and databases £57.2bn, of which £42.9bn own-account and £14.3bn purchased. ons.gov.uk/economy/economicoutputandproductivity/productivitymeasures/bulletins/investmentinintangibleassetssintheuk/2023
5. Supplier SFIA rate card published under G-Cloud 14 on the UK government's Digital Marketplace, listing senior developer day rates from £450 to £600. Cited as a published rate, not a government-set price; the £30,000 figure is arithmetic from it, a few months of one senior developer, not a survey. assets.applytosupply.digitalmarketplace.service.gov.uk/g-cloud-14
6. UK Department for Business and Trade, *SME Digital Adoption Taskforce: final report*, 2025. gov.uk/government/publications/sme-digital-adoption-taskforce-final-report
7. OECD, *SME Technology Adoption in the United Kingdom*, 2026. oecd.org/en/publications/sme-technology-adoption-in-the-united-kingdom_5f25ce2a-en
8. Eurostat, *Larger enterprises used more e-business apps in 2025*, 20 May 2026. Among enterprises with 250 or more employed, 89% used ERP, 65% CRM and 69% business intelligence; among those with 10 to 49 employed the figures were 41%, 25% and 11%. ec.europa.eu/eurostat/web/products-eurostat-news/w/ddn-20260520-1
9. Eurostat, *E-business integration*, Statistics Explained, dataset isoc_eb_iip. Size classes are 10 to 49 persons employed (small) and 250 or more (large); the medium class of 50 to 249 is omitted from the figure. ec.europa.eu/eurostat/statistics-explained/index.php?title=E-business_integration
10. GitHub, *Quantifying GitHub Copilot's impact on developer productivity*; Peng et al., arXiv.
11. Cui, Demirer, Jaffe, Musolff, Peng and Salz, *The Effects of Generative AI on High-Skilled Work: Evidence from Three Field Experiments with Software Developers*, Management Science. 4,867 developers at Microsoft, Accenture and an anonymous Fortune 100 electronics manufacturer; 26% increase in completed tasks.
12. METR, *Measuring the Impact of Early-2025 AI on Experienced Open-Source Developer Productivity*, arXiv 2507.09089. Participants typically used the AI tool for a few dozen hours before and during the study. METR update, February 2026. metr.org/blog/2025-07-10-early-2025-ai-experienced-os-dev-study
13. DORA, *State of AI-assisted Software Development 2025*. dora.dev/dora-report-2025
14. NIST SP 800-218, *Secure Software Development Framework*; NIST SP 800-218A, generative AI extension.
15. HMRC and DBT, *Call for Evidence: Business Systems Integration*, 2026. gov.uk/government/calls-for-evidence/business-systems-integration
16. Microsoft, *Copilot for Power BI overview*; Google Cloud, *Conversational Analytics in Looker*; Tableau, *Tableau Agent FAQ*.
17. Salesforce, *Agentforce Voice*, on voice agents tied to CRM data, actions and workflows. salesforce.com/uk/agentforce/voice
18. Pearce et al., *Asleep at the Keyboard?*, arXiv 2108.09293; Perry et al., *Do Users Write More Insecure Code with AI Assistants?*, arXiv 2211.03622.
19. ICO, *Tech futures: Agentic AI*; EU Regulation 2024/1689 (AI Act); EU Regulation 2024/2847 (Cyber Resilience Act); Data (Use and Access) Act 2025. ico.org.uk/about-the-ico/research-reports-impact-and-evaluation

Full source bank

Two hundred and forty items, retained in full. The first seventy or so are the sources most likely to carry a direct citation; the rest are here for verification, background, counterargument and future work. Where an item is first-party vendor material it should be cited for product functionality or the vendor's own data, never as independent proof of economic benefit. Publisher labels are live links where a URL was recorded.

GOVERNMENT, OFFICIAL STATISTICS AND ECONOMIC INSTITUTIONS

- 1 OECD, Digital Economy Outlook 2024 Vol. 1: unusually direct discussion of software/data technologies as high-fixed-cost, low-additional-cost technologies. [OECD](#)
- 2 OECD, SME Technology Adoption in the UK, 2026: core UK evidence on cost, appropriateness, trust, skills and SME adoption. [OECD](#)
- 3 UK DBT, SME Digital Adoption Taskforce final report, 2025: primary UK policy evidence on adoption barriers. [GOV.UK](#)
- 4 Eurostat, Larger enterprises used more e-business apps in 2025: definitive 41/89 ERP, 25/65 CRM and 11/69 BI comparison. [Eurostat](#)
- 5 Eurostat, Digital economy and society: enterprises, wider EU enterprise ICT series. [Eurostat](#)
- 6 Eurostat digital-economy database: underlying enterprise ICT datasets for further cuts by size/country. [Eurostat](#)
- 7 Eurostat ICT usage methodology: methodology for interpreting enterprise adoption figures. [Eurostat](#)
- 8 OECD, Vectors of Digital Transformation: foundational "scale without mass"/near-zero marginal replication treatment. [OECD](#)
- 9 OECD, Going Digital: Shaping Policies, Improving Lives: broader digital economics and innovation context. [OECD](#)
- 10 OECD, The Digital Transformation of SMEs: SME digitalisation, scale and adoption constraints. [OECD](#)
- 11 OECD, Data Shaping Firms and Markets: SME data gap, intangibles and concentration. [OECD](#)
- 12 OECD Economic Outlook 2025: useful treatment of investment in software/intangibles and corporate hurdle rates. [OECD](#)
- 13 OECD, Innovation in the Software Sector: historical sector economics and volume effects. [OECD](#)
- 14 OECD Going Digital Toolkit: comparative adoption/productivity indicators. [OECD](#)
- 15 OECD ICT Access and Usage by Businesses: cross-country business digitalisation data. [OECD](#)
- 16 HMRC/DBT, Business Systems Integration call for evidence, 2026: excellent primary evidence for fragmented business systems and integration burdens. [GOV.UK](#)
- 17 HMRC, Business Systems Integration landing page: consultation metadata and status. [GOV.UK](#)
- 18 ONS, E-commerce and ICT activity: historic UK business technology-adoption dataset. [ONS](#)
- 19 ONS, AI affecting UK businesses research, 2025: current official UK AI-business research programme. [ONS search](#)
- 20 ONS, Business Insights and Conditions Survey: high-frequency business technology/adoption questions. [ONS](#)
- 21 ONS technology/digital economy data hub: broader UK business digitalisation evidence. [ONS](#)
- 22 DBT, Business population estimates: denominator and SME population context. [GOV.UK](#)
- 23 European Commission SME definition: official EU size thresholds and SME importance. [European Commission](#)
- 24 EU Digital Decade policy programme: official digitalisation targets and monitoring. [European Commission](#)
- 25 Digital Decade country reports: country-level business digitalisation evidence. [European Commission](#)
- 26 European Commission, Digitalisation in Europe: official statistical overview. [European Commission](#)
- 27 US Census, Annual Business Survey: firm-level innovation/technology adoption source. [US Census](#)
- 28 US Census, Business Trends and Outlook Survey: high-frequency business AI-use evidence. [US Census](#)
- 29 US Census, Business Formation Statistics: useful context for testing entrepreneurship effects. [US Census](#)
- 30 US BEA Digital Economy: official US digital-economy measurement. [BEA](#)
- 31 BEA Intellectual Property Products: software/R&D as investment in national accounts. [BEA](#)
- 32 US BLS, Software Developers Occupational Outlook: official employment/outlook reference. [BLS](#)
- 33 US BLS, Computer Programmers Occupational Outlook: useful comparison between programming and broader software-development roles. [BLS](#)
- 34 US SBA Office of Advocacy, Small Business FAQs: US SME context. [SBA](#)
- 35 World Bank, Digital Progress and Trends: global digitalisation context. [World Bank](#)
- 36 IMF, Digitalisation and productivity research portal: macroeconomic background. [IMF](#)
- 37 UK Cyber Security Breaches Survey: authoritative evidence on cyber burdens facing organisations. [GOV.UK](#)
- 38 UK AI sector studies: official measurement of the UK AI supplier ecosystem. [GOV.UK](#)
- 39 UK AI Opportunities Action Plan: policy context for AI diffusion. [GOV.UK](#)
- 40 CMA AI Foundation Models review: UK competition and market-structure perspective. [CMA](#)
- 41 CMA Cloud Services Market Investigation: cloud concentration, switching and interoperability context. [CMA](#)
- 42 European Commission, cloud-computing enterprise statistics: complementary cloud-adoption evidence. [Eurostat](#)
- 43 European Commission, AI use by enterprises: size differences in enterprise AI adoption. [Eurostat](#)
- 44 European Investment Bank, Investment Survey: digital investment and finance constraints. [EIB](#)
- 45 OECD Financing SMEs and Entrepreneurs: financing constraints relevant to fixed software investment. [OECD](#)
- 46 OECD Productivity database: background for technology/productivity comparisons. [OECD](#)
- 47 OECD Digital Security of SMEs work: complements adoption economics with security burden. [OECD](#)
- 48 ENISA, Cybersecurity for SMEs: European SME-specific cyber challenges. [ENISA](#)

- 49 UK Productivity Institute digital-adoption research hub: supporting UK productivity literature. [Productivity Institute](#)
- 50 European Data Portal/Open Data maturity work: data infrastructure context. [data.europa.eu](#)

LARGE EMPIRICAL, EXPERIMENTAL AND PEER-REVIEWED RESEARCH

- 51 Cui et al., Effects of Generative AI on High-Skilled Work: three randomised developer field experiments, 4,867 developers. [Microsoft Research](#)
- 52 Same study, SSRN version: paper/revision trail. [SSRN](#)
- 53 METR, Early-2025 AI and experienced open-source developers: essential counterevidence. [METR](#)
- 54 METR paper, arXiv 2507.09089: full methodology/results. [arXiv](#)
- 55 METR 2026 update: newer tools appear more favourable but uncertainty remains. [METR](#)
- 56 Peng et al., Impact of AI on Developer Productivity: GitHub Copilot controlled experiment. [arXiv](#)
- 57 GitHub productivity experiment explainer: first-party summary of the 55% result. [GitHub](#)
- 58 DORA 2025 State of AI-assisted Software Development: organisational amplifier framing. [DORA](#)
- 59 Google Cloud DORA 2025 announcement: nearly 5,000 respondents and qualitative research context. [Google Cloud](#)
- 60 DORA 2024 report: baseline immediately before the 2025 agentic shift. [DORA](#)
- 61 DORA 2023 report: useful evidence on user-centricity and organisational performance. [DORA](#)
- 62 DORA research programme: longitudinal software-delivery research archive. [DORA](#)
- 63 Anthropic, Economic Index: Software Development: large corpus of coding interactions and automation/augmentation patterns. [Anthropic](#)
- 64 Anthropic Economic Index: wider task-level AI labour evidence. [Anthropic](#)
- 65 Noy & Zhang, Experimental Evidence on Productivity Effects of Generative AI: influential controlled productivity study outside coding. [Science](#)
- 66 Brynjolfsson, Li & Raymond, Generative AI at Work: large field evidence on customer-support productivity. [NBER](#)
- 67 Dell'Acqua et al., Navigating the Jagged Technological Frontier: evidence that AI gains depend heavily on task boundaries. [SSRN](#)
- 68 Pearce et al., Asleep at the Keyboard?: security analysis of Copilot-generated code. [arXiv](#)
- 69 Pearce et al., IEEE Security & Privacy publication: peer-reviewed venue record. [IEEE](#)
- 70 Perry et al., Do Users Write More Insecure Code with AI Assistants?: controlled human study. [ACM](#)
- 71 Perry et al. preprint: accessible full research record. [arXiv](#)
- 72 Fu et al., Security Weaknesses of Copilot-Generated Code in GitHub Projects: empirical in-the-wild security study. [ACM](#)

- 73 Vaithilingam et al., Expectation vs Experience with GitHub Copilot: human-computer interaction evidence. [ACM](#)
- 74 Barke et al., Grounded Copilot: qualitative/interaction model for programmer use of code assistants. [ACM DL](#)
- 75 Forsgren et al., SPACE of Developer Productivity: warns against reducing productivity to code/output measures. [ACM Queue](#)
- 76 Forsgren et al., DevEx: What Actually Drives Productivity: developer experience and friction framework. [ACM Queue](#)
- 77 Accelerate / State of DevOps research archive: foundational delivery-performance evidence. [DORA](#)
- 78 Stack Overflow Developer Survey 2024: AI adoption versus trust gap. [Stack Overflow](#)
- 79 Stack Overflow Developer Survey 2025: subsequent developer/AI adoption evidence. [Stack Overflow](#)
- 80 Stack Overflow 2024 AI analysis: 76% using/planning AI but weaker trust in accuracy. [Stack Overflow](#)
- 81 Goldfarb & Tucker, Digital Economics: foundational JEL synthesis of falling digital economic costs. [AEA](#)
- 82 Varian, Buying, Sharing and Renting Information Goods: information-goods ownership/rental economics. [Wiley](#)
- 83 Bakos & Brynjolfsson, Bundling Information Goods: classic economics of near-zero marginal-cost goods. [INFORMS](#)
- 84 Shapiro & Varian, Information Rules: foundational information-goods strategy reference. [Harvard Business Review Press](#)
- 85 Haskel & Westlake, Capitalism Without Capital: intangible investment and scaling economics. [Princeton](#)
- 86 Corrado, Hulten & Sichel, Measuring Capital and Technology: foundational intangible-capital measurement. [Wiley](#)
- 87 Andrews, Criscuolo & Gal, The Best versus the Rest: frontier/laggard productivity diffusion. [OECD](#)
- 88 Bloom, Sadun & Van Reenen management research: complementary organisational capital context. [NBER](#)
- 89 Brynjolfsson, Hitt & Yang, Intangible Assets and Computers: complementary investment around IT. [Brookings](#)
- 90 Bresnahan, Brynjolfsson & Hitt, IT, Workplace Organisation and Skilled Labour: technology complements organisational change. [QJE](#)
- 91 Bessen, Learning by Doing / technology adoption research: useful adoption and complementary-skill background. [Boston University](#)
- 92 Generative AI Fuels Solo Entrepreneurship, but Teams Still Lead at the Top: 160,000+ Product Hunt launches; early producer-fragmentation evidence. [arXiv](#)
- 93 AI as “Co-founder”: GenAI for Entrepreneurship: firm-registration evidence; treat as emerging working paper. [arXiv](#)
- 94 Vibecoding and Digital Entrepreneurship: emerging research on lowered digital-product creation barriers. [arXiv search](#)
- 95 NBER AI research archive: additional economics studies as the field matures. [NBER](#)
- 96 Stanford AI Index 2026: broad empirical AI adoption/cost/capability context. [Stanford HAI](#)
- 97 Stanford AI Index 2025: prior-year benchmark for trend comparisons. [Stanford HAI](#)
- 98 Stanford Digital Economy Lab: frontier research on AI and productivity. [Stanford](#)
- 99 MIT Initiative on the Digital Economy research: complementary economics/productivity literature. [MIT IDE](#)
- 100 AEA digital economics literature: economics-journal source base for additional theory. [AEA](#)

STANDARDS, SECURITY, REGULATION, INTELLECTUAL PROPERTY AND ACCOUNTING

- 101 NIST SP 800-218, Secure Software Development Framework: foundational lifecycle security practices. [NIST](#)
- 102 NIST SP 800-218A: GenAI/foundation-model extension to secure software development. [NIST](#)
- 103 NIST AI RMF 1.0: AI governance/risk-management foundation. [NIST](#)
- 104 NIST Generative AI Profile: GenAI-specific risk-management profile. [NIST](#)
- 105 NIST Cybersecurity Framework 2.0: general cybersecurity governance. [NIST](#)
- 106 NIST Zero Trust Architecture, SP 800-207: identity/access architecture reference. [NIST](#)
- 107 CISA Secure by Design: security as design responsibility rather than downstream testing. [CISA](#)
- 108 CISA Secure by Design Pledge: practical product-security commitments. [CISA](#)
- 109 NCSC, Guidelines for Secure AI System Development: UK-led international AI-security guidance. [NCSC](#)
- 110 NCSC cloud-security principles: cloud architecture/security background. [NCSC](#)
- 111 NCSC logging guidance: observability/security-monitoring complement. [NCSC](#)
- 112 ENISA AI cybersecurity work: European AI-security research and guidance. [ENISA](#)
- 113 OWASP Top 10 for LLM/GenAI Applications: application-level GenAI threat catalogue. [OWASP](#)
- 114 OWASP Top 10 Web Application Security Risks: broader application-security baseline. [OWASP](#)
- 115 OWASP SAMM: software-assurance maturity model. [OWASP](#)
- 116 OpenSSF Scorecard: open-source dependency security checks. [OpenSSF](#)
- 117 OpenSSF Best Practices: secure open-source dependency practice. [OpenSSF](#)
- 118 SLSA framework: software-supply-chain integrity. [SLSA](#)
- 119 SPDX specification: software bill-of-materials standard. [SPDX](#)
- 120 CycloneDX specification: SBOM and dependency transparency. [CycloneDX](#)
- 121 ISO/IEC 27001: information-security management system. [ISO](#)
- 122 ISO/IEC 27002: information-security controls. [ISO](#)
- 123 ISO/IEC 42001: AI management-system standard. [ISO](#)
- 124 ISO/IEC 23894: AI risk-management guidance. [ISO](#)
- 125 ISO/IEC 25010: software/product quality model. [ISO](#)
- 126 ISO/IEC/IEEE 12207: software lifecycle processes. [ISO](#)
- 127 CIS Critical Security Controls: operational security baseline. [CIS](#)
- 128 UK Copyright, Designs and Patents Act 1988: statutory copyright basis. [legislation.gov.uk](#)
- 129 UK IPO, Ownership of copyright works: critical commissioned-work ownership guidance. [GOV.UK](#)
- 130 UK Government, Knowledge Asset classes and types: expressly covers software/source code as copyright material. [GOV.UK](#)
- 131 UK IPO, Using your IP rights: ownership/transfer/value guidance; warns that IP without economic benefit may have limited value. [GOV.UK](#)

- 132 UK Business guidance, Safeguard your IP: current IP valuation/protection guidance. [Business.gov.uk](#)
- 133 IFRS IAS 38 Intangible Assets: primary accounting source for development expenditure/software intangibles. [IFRS](#)
- 134 IFRS SIC-32 Website Costs: application of IAS 38 to software/web development stages. [IFRS](#)
- 135 FRC FRS 102: principal UK accounting standard for many non-IFRS entities; current 2026 framework. [FRC](#)
- 136 IAS 36 Impairment of Assets: relevant where capitalised software loses economic value. [IFRS](#)
- 137 IFRS 3 Business Combinations: acquired intangible-asset treatment, relevant to M&A claims. [IFRS](#)
- 138 IASB Intangible Assets research project: useful because accounting for modern software/intangibles remains under active review. [IFRS](#)
- 139 EU Artificial Intelligence Act, Regulation 2024/1689: authoritative EU AI regulation. [EUR-Lex](#)
- 140 EU Cyber Resilience Act, Regulation 2024/2847: cybersecurity obligations for covered digital products. [EUR-Lex](#)
- 141 EU GDPR, Regulation 2016/679: data-protection baseline for EU deployments. [EUR-Lex](#)
- 142 EU Data Act, Regulation 2023/2854: access, portability and data-use framework. [EUR-Lex](#)
- 143 NIS2 Directive: cybersecurity requirements for relevant sectors/entities. [EUR-Lex](#)
- 144 Digital Operational Resilience Act: relevant to financial-sector LiteByte clients. [EUR-Lex](#)
- 145 EU Product Liability Directive 2024/2853: updated software/digital-product liability context. [EUR-Lex](#)
- 146 UK Data (Use and Access) Act 2025: current UK data-law reform context. [legislation.gov.uk](#)
- 147 ICO, AI and data-protection guidance: authoritative UK deployment requirements. [ICO](#)
- 148 ICO AI resources and risk toolkit: practical governance tools. [ICO](#)
- 149 ICO Tech Futures: Agentic AI: exceptionally relevant current risk analysis. [ICO](#)
- 150 ICO Generative AI consultation response: data-controller, purpose, accuracy and rights issues. [ICO](#)

VENDOR ARCHITECTURE AND PRODUCT DOCUMENTATION

- 151 AWS SaaS Lens: silo, pool and bridge: multitenant architecture/economics. [AWS](#)
- 152 AWS SaaS Lens: overall SaaS architecture guidance. [AWS](#)
- 153 AWS SaaS tenant isolation strategies: isolation-versus-sharing architecture. [AWS](#)
- 154 Microsoft Azure Well-Architected SaaS methodology: explicitly connects multitenancy, COGS and pricing. [Microsoft](#)
- 155 Microsoft multitenant tenancy models: cost/isolation/responsiveness trade-offs. [Microsoft](#)
- 156 Microsoft Azure multitenant architecture guide: broader technical reference. [Microsoft](#)
- 157 Google Cloud, GKE multitenancy: resource sharing and isolation. [Google Cloud](#)
- 158 Google Cloud, Spanner multi-tenancy: database-level tenancy patterns. [Google Cloud](#)
- 159 Google Cloud, SaaS architectures: multi-tenant design patterns. [Google Cloud](#)
- 160 Microsoft Power BI Copilot overview: current conversational BI capability. [Microsoft](#)

- 161 Power BI Copilot and semantic models: accuracy/model-preparation caveat central to the thesis. [Microsoft](#)
- 162 Power BI, prepare semantic models for AI: metadata/verified-answer requirements. [Microsoft Learn](#)
- 163 Power BI natural-language Q&A: predecessor/continuity of intent-based analytics. [Microsoft](#)
- 164 Google Looker Conversational Analytics overview: governed natural-language analytics. [Google Cloud](#)
- 165 Looker Conversational Analytics GA announcement: product direction and BI-bottleneck framing. [Google Cloud](#)
- 166 Looker Data Agents: domain-specific conversational analysis configuration. [Google Cloud](#)
- 167 Looker Embedded Conversational Analytics: agent interface inside third-party products. [Google Cloud](#)
- 168 Tableau Agent FAQ: conversational visualisation/data exploration and governance caveats. [Tableau](#)
- 169 Salesforce Agentforce: action-oriented CRM agents. [Salesforce](#)
- 170 Salesforce Agentforce Voice: voice → CRM data/actions/workflows. [Salesforce](#)
- 171 Microsoft Copilot Studio: enterprise agent/workflow construction. [Microsoft](#)
- 172 Microsoft Dynamics 365 Copilot: CRM/ERP agentic direction. [Microsoft](#)
- 173 Amazon Q in QuickSight / Quick Suite: natural-language enterprise analytics. [AWS](#)
- 174 Amazon Q Business: enterprise agent/retrieval/action layer. [AWS](#)
- 175 Google Gemini Code Assist: first-party AI coding-agent documentation. [Google Cloud](#)
- 176 GitHub Copilot documentation: coding-agent capabilities and limitations. [GitHub](#)
- 177 GitHub responsible use of Copilot: review/testing requirements for generated code. [GitHub](#)
- 178 GitHub Copilot coding agent: evidence of movement beyond autocomplete into delegated tasks. [GitHub](#)
- 179 GitHub Copilot code review: AI-assisted review capability. [GitHub](#)
- 180 AWS Q Developer: coding-agent/productivity tooling. [AWS](#)
- 181 Anthropic Claude Code documentation: agentic software-development workflow. [Anthropic](#)
- 182 Databricks AI/BI Genie: conversational analytics over governed enterprise data. [Databricks](#)
- 183 Snowflake Cortex Analyst: natural-language structured-data querying. [Snowflake](#)
- 184 Oracle Analytics AI capabilities: additional enterprise BI evidence. [Oracle](#)
- 185 SAP Joule: enterprise-agent layer over business systems. [SAP](#)
- 186 ServiceNow Now Assist: workflow-centric enterprise GenAI. [ServiceNow](#)
- 187 HubSpot Breeze: CRM AI-agent direction among mid-market software. [HubSpot](#)
- 188 Zoho Zia: AI across integrated business-software suite. [Zoho](#)
- 189 Microsoft Fabric Copilot: enterprise data/analytics AI layer. [Microsoft](#)
- 190 Google Cloud agentic AI architecture centre: architecture patterns and governance. [Google Cloud](#)

VERTICAL SOFTWARE, SAAS PROLIFERATION AND MARKET ANALYSIS

- 191 Constellation Software corporate site: canonical vertical-market-software example. [Constellation](#)
- 192 Constellation Software investor relations: filings and operating-group economics. [Constellation](#)
- 193 Constellation: vertical market software company description: mission-critical specialised-software model. [Constellation](#)
- 194 Constellation ESG/operating structure: decentralised vertical-business model. [Constellation](#)
- 195 Volaris Group: large portfolio of vertical-market software companies. [Volaris](#)
- 196 Topicus.com: European vertical-software operating model. [Topicus](#)
- 197 Tidemark, 2025 Vertical & SMB SaaS Benchmark: 59% multi-product; useful specialist-software benchmark. [Tidemark](#)
- 198 Tidemark Vertical SaaS Knowledge Project: extensive sector case-study archive. [Tidemark](#)
- 199 Stripe/Tidemark vertical SaaS work: payments/platform expansion as vertical product strategy. [Stripe](#)
- 200 Okta, Businesses at Work 2025: application-estate evidence from Okta's customer data. [Okta](#)
- 201 Okta, Businesses at Work 2026: current application-estate trends. [Okta](#)
- 202 BetterCloud, State of SaaS 2025: vendor survey on SaaS sprawl and consolidation. [BetterCloud](#)
- 203 BetterCloud, State of SaaS 2026: evidence that AI can re-expand portfolios after consolidation. [BetterCloud](#)
- 204 Zyllo, SaaS Management Index 2026: software spend, discovery and shadow-AI evidence. [Zyllo](#)
- 205 Zyllo SaaS/AI research library: further spend/consolidation benchmarks. [Zyllo](#)
- 206 Productiv software-management research: utilisation and portfolio rationalisation. [Productiv](#)
- 207 Torii SaaS-management research: application sprawl and shadow-IT perspective. [Torii](#)
- 208 Vendr software-buying benchmarks: pricing/procurement and SaaS portfolio evidence. [Vendr](#)
- 209 Flexera State of ITAM: software asset-management and spend governance. [Flexera](#)
- 210 Flexera State of the Cloud: cloud-spend and multi-cloud context. [Flexera](#)
- 211 Bessemer State of the Cloud: cloud/SaaS business-model analysis. [Bessemer](#)
- 212 Battery Ventures State of the OpenCloud: software-market operating metrics. [Battery](#)
- 213 Meritech SaaS benchmarks: public-company software economics and margins. [Meritech](#)
- 214 Clouded Judgement: high-quality public SaaS financial benchmarking. [Clouded Judgement](#)
- 215 McKinsey, Economic potential of generative AI: secondary macro context; not a core causal source. [McKinsey](#)
- 216 McKinsey GenAI explainer: broad mechanism/business-adoption background. [McKinsey](#)
- 217 BCG AI/code productivity research: useful consultancy triangulation, secondary tier. [BCG](#)
- 218 Bain software/SaaS research: consolidation and software-business-model background. [Bain](#)
- 219 Deloitte State of Generative AI in the Enterprise: enterprise adoption/risk surveys. [Deloitte](#)
- 220 PwC AI Jobs Barometer: task/skills/productivity labour-market context. [PwC](#)
- 221 KPMG AI Quarterly Pulse Survey: enterprise implementation sentiment. [KPMG](#)
- 222 Menlo Ventures, State of Generative AI in the Enterprise: enterprise AI spending/use-case evidence. [Menlo Ventures](#)
- 223 a16z enterprise AI / vertical SaaS analysis: useful strategic interpretation, not core evidence. [a16z](#)
- 224 Sapphire Ventures AI/software research: software-industry strategy and market formation. [Sapphire Ventures](#)
- 225 Thoughtworks Technology Radar: practitioner evidence on AI-assisted software engineering patterns. [Thoughtworks](#)
- 226 GitLab Global DevSecOps Report: developer/tool-chain survey evidence. [GitLab](#)
- 227 JetBrains Developer Ecosystem Survey: large developer-practice dataset. [JetBrains](#)
- 228 Sonatype State of the Software Supply Chain: dependency/supply-chain risk evidence. [Sonatype](#)
- 229 Snyk AI/security research: generated-code security and developer-practice context. [Snyk](#)
- 230 Cloud Security Alliance AI research: enterprise AI-security/governance perspective. [CSA](#)
- 231 FinOps Foundation, State of FinOps: usage-sensitive cloud/AI cost governance. [FinOps Foundation](#)
- 232 CNCF Annual Survey: cloud-native platform adoption, relevant to lower deployment/infrastructure friction. [CNCF](#)
- 233 CircleCI State of Software Delivery: delivery-performance benchmarking. [CircleCI](#)
- 234 LinearB engineering benchmarks: practitioner-level software delivery benchmarks. [LinearB](#)
- 235 DX developer-productivity research library: DevEx and productivity evidence synthesis. [DX](#)
- 236 Atlassian developer-experience research: engineering friction and platform-team evidence. [Atlassian](#)
- 237 Gartner SaaS management research: portfolio-management market context; often paywalled, use selectively. [Gartner](#)
- 238 IDC SaaS/cloud research: independent market-data complement. [IDC](#)
- 239 Forrester software/platform research: market consolidation/platform strategy context. [Forrester](#)
- 240 CB Insights AI/startup research: funding/startup formation triangulation; secondary source. [CB Insights](#)