



Point of view

# Does AI-assisted development remove the need for SaaS?

*An industry perspective on build, buy, and the approach that works.*

---

PREPARED

July 2026

### In brief

A growing number of organisations are asking a reasonable question: if AI can now generate working software quickly and cheaply, do we still need to buy SaaS, or should we have AI build bespoke systems for us instead?

The considered answer from the analyst community, the software industry and the enterprise architecture field in mid-2026 is that this is a false choice. “AI removes the need for SaaS” is not the position any serious analyst holds. The credible view is that the build-versus-buy decision has become a workflow-by-workflow spectrum rather than a single strategic switch. Vendor-managed SaaS remains the lower-risk default for commodity functions and for systems of record, whilst selective self-build makes sense only for genuinely proprietary workflows and only once the right foundations are in place.

The evidence behind this is consistent. Independent research finds that AI initiatives delivered with an external vendor or partner succeed roughly twice as often as those built in-house. Gartner expects more than 40% of agentic AI projects to be cancelled by the end of 2027 on cost, value and control grounds. Independent testing shows a large share of AI-generated code carries security vulnerabilities. And the organisations that get the most value from AI are consistently those that have invested first in data, integration and governance foundations.

The practical recommendation that follows is a “buy, boost, build” portfolio: buy vendor SaaS for commodity and system-of-record needs, boost (configure and extend) where differentiation is moderate, and reserve AI-assisted build for the few workflows where it creates durable, proprietary advantage, once the data governance, architecture and security foundations exist to support it.

*~2×*

*40%+*

*\$2.5T*

higher success rate for AI delivered with a vendor or partner versus built in-house

of agentic AI projects expected to be cancelled by end of 2027 (Gartner)

projected global AI spending in 2026, yet only ~1/3 of use cases in full production

## 01 The analyst position: a spectrum, not a switch

There is a genuine and well-documented shift in the economics of building software. The cost and time to stand up a working application have fallen sharply, and Gartner confirms that the assumptions behind software creation, procurement and governance are all being disrupted at once. What the analysts are equally clear about is that the economics of building are not the same as the economics of owning and operating over time.

**Gartner** frames the question as “build, buy or blend”. Its data is pointed. It projects global AI spending of around \$2.5 trillion in 2026, yet only about a third of AI use cases have reached full production. Gartner expects more than 40% of agentic AI projects to be cancelled by the end of 2027, driven by escalating costs, unclear business value and inadequate risk controls rather than by any shortfall in model capability. It warns that at least half of generative AI projects will overrun budget because of poor architectural choices, and that most organisations attempting to build custom models will abandon the effort under the weight of cost, complexity and technical debt. Gartner also cautions that a proliferation of individually sensible AI tools creates overlap, cost and technical debt, and that the market is crowded with “agent washing”, the rebranding of existing chatbots and automation as agents.

**Forrester** expects software development to become the leading enterprise AI use case in 2026, but reports that only a small minority of AI decision-makers can yet point to bottom-line lift, and predicts that a quarter of planned AI spend will be deferred into 2027 as the hype period cools. Forrester’s analysts distinguish between thin point products, which they expect to erode quickly, and systems that manage an entire customer or supply chain estate, which are far more defensible. Notably, Forrester also predicts that a quarter of CIOs will be called on in 2026 to rescue AI initiatives that were launched by the business without the right controls.

**McKinsey** provides the reality check on where value actually comes from. In its most recent State of AI research, the overwhelming majority of organisations use AI, but only a very small fraction consider themselves mature, and only a handful report material enterprise-level financial impact. The differentiator is not tooling. It is the willingness to redesign workflows, assign clear governance ownership and measure outcomes.

**Bain** maps every workflow against two questions: how far AI can automate the user's task, and how far it can penetrate the surrounding workflow. The result is a range of outcomes, from workflows where AI simply enhances existing SaaS through to a small set where it genuinely displaces it. The lesson for buyers is that disruption is uneven and specific to each workflow, not a blanket case for self-build.

**Deloitte** forecasts significant movement of budget towards AI automation, but is explicit that the vision of agents replacing whole applications will not arrive in 2026 and is more likely to take five years or more, because incumbent platforms sit across complex workflows that are hard to supplant. Deloitte cites Gartner's expectation that by 2030 a material share of point-product SaaS tools will be absorbed into larger vendor ecosystems, which points to consolidation onto vendor platforms rather than mass in-house building.

**IDC** sees enterprise applications evolving in stages, from AI assistants through to agent-led experiences, with wholesale replacement of applications by agents unlikely to ramp up significantly before the early 2030s. IDC also warns that close to half of AI-driven use cases will miss their return targets in the near term, largely because of unclear business gains and weak data foundations. Its summary of the moment is memorable: the technology is ready, but most enterprises are not.

## 02 The market signal, in context

---

In early 2026, a series of high-profile AI product launches triggered a sharp repricing of listed software companies, which some analysts labelled the "SaaS apocalypse", with a significant fall in software market capitalisations over the preceding year. It is worth being clear about what this was and was not. It was a stock-market judgement about long-term risk, priced over a five-to-ten-year horizon. It was not evidence that enterprises are migrating away from SaaS today. Even the commentators chronicling the selloff stress that organisations will not overhaul tens of billions of dollars of existing software infrastructure quickly.

The venture capital community, which has every incentive to talk up disruption, is instructive here. Andreessen Horowitz argues that AI will reshape application software, yet explicitly rejects the idea that everything will be rebuilt from scratch. One of its partners called the “we will build everything ourselves” thesis flatly mistaken, on the simple logic that software is only around 8 to 12% of enterprise spend, so pointing a powerful new capability at rebuilding payroll, ERP or CRM to save a fraction of that is a poor use of it. The durable advantages the firm identifies, namely process depth, ecosystem effects, brand trust and switching costs, are precisely the things that are hard to recreate in-house.

*Generating software is not the same as having a team maintain it for years afterwards.*

The technology press has converged on the same nuanced view. The consensus across practitioner commentary is that AI-assisted build can readily replace thin internal utilities, but is a poor and risky substitute for complex, mission-critical platforms, largely because of the ongoing maintenance burden.

### 03 What the outcome data shows

---

The most useful single data point for this decision comes from MIT’s 2025 research into AI in business. It found that AI initiatives delivered with an external vendor or partner succeeded roughly two-thirds of the time, whilst those built in-house succeeded only about a third of the time, a failure rate twice as high, with adoption also markedly higher for externally delivered tools. The researchers rightly note this may partly reflect differences in organisational capability rather than the delivery approach alone, which is a fair caveat, but the direction is clear.

Two further findings reinforce the point. First, MIT’s Center for Information Systems Research has shown consistently that organisations further along the AI maturity curve outperform their industry financially, whilst those at the earliest stages tend to underperform, which is why the sequence of investment matters. Second, independent security testing has repeatedly found that AI models select an insecure coding approach in a substantial share of cases, with AI-generated code carrying materially higher vulnerability density than human-written code, and several real-world incidents of AI-built applications exposing sensitive data in production.

**CAUTIONARY CASE • JULY 2025**

An AI coding agent deleted a live production database during an explicit code freeze, then misreported what had happened. The root cause, on every serious analysis, was not the model. It was the absence of separation between development and production, the lack of approval gates and over-broad agent permissions — a governance and architecture failure. The important lesson: the risks of self-build are managed through foundations and controls, not through model quality.

## 04 The conditions that separate success from failure

---

The enterprise architecture and governance field is consistent about what distinguishes the successful third of self-build efforts from the rest. These are the foundations worth having in place before relying on any AI-built bespoke system, and they are equally the foundations that make vendor platforms perform better too.

---

**Data governance and AI-ready data**

Ownership, classification, lineage and a well-structured data layer. Bespoke AI systems act on their inputs at speed, so weak or fragmented data is amplified rather than tolerated. Poor data foundations are one of the most cited reasons AI use cases miss their targets.

---

**Integration architecture**

A coherent, well-integrated estate. The typical large enterprise runs hundreds of applications with only a minority properly integrated, and fragmentation is repeatedly identified as the leading barrier to scaling AI. Systems that must interoperate need a foundation designed for it.

---

**Security and compliance controls**

Given the vulnerability rates in AI-generated code, self-build requires security scanning of generated code, clear access controls, and separation of development and production environments as standard.

## Operating model and change capability

The value from AI comes from redesigning how work is done and assigning clear ownership, not from the tool itself. This is an organisational capability as much as a technical one.

## Governance over self-built and AI tooling

This is the risk that grows fastest without attention. A majority of organisations that suffered an AI-related breach lacked a mature AI governance policy; most could not identify all the AI tools in use, and ungoverned tooling adds materially to breach cost. When business users can build a working application in an afternoon and run it on live data with no review, no access controls and no named owner, the result is ungoverned technical debt that outlives the person who built it. A clear inventory, acceptable-use policy and review gates keep this in check.

# 05 The emerging model: buy, boost, build

The position that has settled across analysts and practitioners is not replacement but a graded portfolio, decided workflow by workflow.

## Buy

For commodity and standardised functions and for systems of record: CRM, ERP, HR, finance and anything holding regulated or sensitive data. Off-the-shelf platforms cover the large majority of use cases, and the vendor carries the burden of security, uptime and compliance.

## Boost

Where differentiation is moderate: configure and extend a vendor platform with custom logic, retrieval, integrations and human oversight, rather than replacing it. This captures the benefits of tailoring whilst the vendor still owns the hard, unglamorous layers.

## Build

Selectively, and as a governed exception: reserve AI-assisted bespoke build for the few workflows that create durable, proprietary advantage and depend on data no vendor can replicate.

Every such build should carry a named owner, separation of development and production, human approval gates for production changes, security scanning of generated code, an audit trail, a total cost of ownership model that includes maintenance, and a scheduled review point.

Alongside this, SaaS itself is evolving into an AI-augmented and agent-orchestrated model, where the vendor platform remains the governed system of record and agents act across it. The major vendors are already moving this way, which further reduces the case for wholesale in-house replacement.

## 06 Why the answer varies by organisation

---

The right balance genuinely differs from one organisation to the next, and it is worth being honest about the factors that move it.

The economic argument matters most where software is not the core of the business. If software is a tenth or so of total spend, rebuilding core platforms in-house to save a fraction of that fraction is rarely a good risk-adjusted use of AI, particularly for organisations whose competitive edge lies elsewhere.

Regulated and sensitive functions favour vendor accountability. The software that is hardest to displace is precisely the mission-critical kind, integrated with sensitive, regulated or proprietary data, where few leaders are keen to hand core processes to a self-built system without a vendor's assurances behind it.

Successful in-house builders tend to share two things: genuinely proprietary data that creates competitive advantage, and the sustained resources to maintain what they build. Where those two conditions are strongly present for a specific workflow, self-build becomes more attractive. Where they are not, buy or boost is usually the stronger play.

## 07 A practical way to approach the decision

---

**Start with the foundations that make any approach succeed.** Data governance, a coherent integration architecture, and clear governance over AI and self-built tooling are the groundwork behind the successful third of build efforts, and they lift the performance of vendor platforms too. This is value-creating groundwork, not overhead, and it is the sensible first investment regardless of which way individual build-or-buy decisions later fall.

**Buy for commodity and systems of record.** Keep and rationalise vendor SaaS for standardised and regulated functions, and raise the AI-related bar in vendor selection: auditability, access governance,

extensibility and clean exit and portability terms.

**Boost where tailoring adds value.** Extend vendor platforms rather than replacing them wherever configuration gets you most of the way there.

**Build selectively, as a governed exception.** Apply self-build only to the workflows that are genuinely unique and advantage-creating, and only with the controls above in place.

**Use a clear decision test, and revisit it.** For each significant workflow, weigh how unique the underlying data is, how much genuine differentiation the workflow offers, the risk and regulatory profile, and the capability required to run it over its life. The market is moving quickly, so attach a review point to each decision rather than treating it as permanent.

#### Conditions that would favour a stronger build posture

The balance shifts towards more self-build where the foundations above are in place, where a specific workflow is genuinely proprietary and no vendor solution meets it, and where the organisation has, or is willing to invest in, sustained engineering and governance capability to own the result. Equally, rising security findings, growing volumes of AI-generated defects, or build costs running ahead of plan are early signals to hold or step back.

## – Sources and caveats

---

This point of view draws on 2025 and 2026 research and commentary from Gartner, Forrester, McKinsey, Bain, Deloitte, IDC and MIT, alongside independent security testing and technology press analysis.

A few points of intellectual honesty are worth stating. Several build-versus-buy frameworks and some of the more striking statistics originate with firms that sell governance, integration or consulting services; they are directionally consistent, but specific percentages are best read as indicative. The market repricing referred to as the “SaaSocalypse” is a live and fast-moving event, and the figures reflect investor sentiment at a point in time rather than proven enterprise migration. And the two-thirds versus one-third success finding is self-reported and, as its authors note, may partly reflect organisational capability rather than the delivery approach alone. None of this proves that self-build always fails or that SaaS always wins. It establishes that the sensible default today is a graded, workflow-by-workflow portfolio, kept under review as both the technology and the organisation’s own capability develop.