

The decision that matters now is **who you build with.**

A five-minute read for operators deciding how to put AI to work.

01 -- THE QUESTION CHANGED

For twenty years the safe advice was to buy software, never build it. That advice made sense when custom development took years and seven figures. Both conditions have changed. Off-the-shelf tools now carry documented waste that compounds at every renewal. The cost of building systems fitted to your operation has fallen hard. And the failure pattern of companies that build alone is now well mapped, which means it can be avoided. The decision that matters in 2026 is who you build with.

02 -- WHAT YOU ACTUALLY PAY FOR WHEN YOU BUY

Generic software is built for the average business in its category. Your business is specific. You pay for that gap three ways.

You pay for what you never use. Companies with fewer than 500 employees run an average of 152 software subscriptions, and industry benchmark data shows only about half of paid licenses get used at all. Usage telemetry across 615 software products found that 80 percent of features are rarely or never touched.

~54%

of paid SaaS licenses are ever used.

ZYLO SAAS MANAGEMENT INDEX, 2026

80%

of software features are rarely or never used,
measured across 615 products.

PENDO FEATURE ADOPTION REPORT

You pay more every year for the same gap. SaaS list prices have been rising near 13 percent annually, with increases concentrated at renewal, when switching is hardest. The product does not fit your operation any better at the higher price.

You pay in labor to glue it together. Each tool holds a slice of your data, and your people reconcile the slices by hand. Researchers writing in Harvard Business Review measured knowledge workers switching between applications roughly 1,200 times per day, losing just under four hours a week to reorientation. That is about 9 percent of the working year spent finding your place again.

Buying still wins for commodity functions. Payroll, accounting, and email are solved problems, and you should buy them. The waste concentrates in the systems that run your core operation: quoting, scheduling, dispatch, job tracking, client workflows. That is precisely where average fits you worst, and where you are already paying people to work around the software.

Buy the commodity. Build the core. The systems that run your actual operation are the ones generic software fits worst.

03 -- WHY BUILDING ALONE GOES WRONG

You have likely seen an internal project fail, or know someone who has. The data explains what you saw. Peer-reviewed research covering thousands of IT projects shows cost overruns follow a power law. The typical project lands near plan. Then one in six becomes an outlier, with cost overruns averaging 200 percent. The distribution has a fat tail, and the tail is where projects kill budgets.

Those blowups have known causes, and the causes form before any code is written. When RAND interviewed 65 veteran AI engineers about why projects fail, 84 percent pointed to leadership-driven root causes: solving the wrong problem, shifting priorities midstream, automating a process that was already broken.

Operations-heavy businesses hit two additional walls. The first is key-person risk. Internal systems tend to live in one or two heads, and the business is exposed the day those heads leave. The second is maintenance. A system you build alone is a system you staff forever, in a talent market where you compete against software companies.

Custom systems succeed when those causes are designed out from the start: modular scope, short delivery cycles, and a process fixed before it gets automated. Experienced builders carry that playbook from dozens of prior projects. First-time internal teams write it as they go, at your expense.

04 -- WHAT CHANGED THE MATH

Goldman Sachs CIO Marco Argenti said this year that AI has made it cheap enough to build smaller applications in-house, that internal tools which took months now ship in days, and that the bank has already terminated some vendor contracts as a result. That is a cost statement, on the record, from one of the most conservative technology buyers in the world.

The pattern is broad. In a late-2025 survey of 817 technology teams, 35 percent had already replaced at least one SaaS product with something they built, and 78 percent expected to build more in 2026. The categories under the most pressure are workflow automation, internal admin tools, and CRM. Those are operational systems, the exact category where generic software fits worst.

One caveat: AI speeds development unevenly. Controlled studies show large gains on well-scoped work and smaller gains on complex systems. The leverage is real, and it accrues to teams that have shipped these builds many times before.

05 -- WHAT THE SUCCESSFUL MINORITY DOES DIFFERENTLY

A 2025 MIT study of enterprise AI adoption found that roughly 95 percent of corporate GenAI pilots produced no measurable P&L impact. The roughly 5 percent that worked shared a pattern. Organizations that built with specialized external partners reached successful deployment about 67 percent of the time. Purely internal builds succeeded about one third as often. The winners also behaved differently as buyers: they demanded deep customization to their workflows and judged vendors on operational results.

Partnered builds succeed at roughly twice the rate of internal ones.

MIT NANDA, STATE OF AI IN BUSINESS 2025

The second half of the pattern predates AI by decades. Michael Hammer warned in 1990 that automating a broken process locks the breakage in. MIT economists later showed that returns on new technology depend on redesigning the work around it, and that bolting new tools onto old processes can produce negative returns.

Klarna is the clearest recent case of both truths at once. AI plus aggressive process simplification let a regulated, bank-licensed company drop Salesforce and Workday for a lighter stack of custom and alternative tools. Its AI assistant handled work equivalent to 700 customer service agents in its first month. Then quality slipped where automation had outrun process design, and the company publicly rehired human agents in 2025. Custom systems work. Sequencing and judgment determine whether they work for you.

Avolis was built for this decision. We work as a forward-deployed partner inside operations-heavy businesses, and our method maps directly onto the failure modes above. Diagnose comes first: we map how your operation actually runs and fix the process before anyone writes code, because automating a broken process is the oldest failure in the book. Build comes second: modular systems, shipped in short cycles, shaped to your workflows rather than the industry average. Embed comes last: we stay until your team runs the system without us, because a system nobody adopts returns nothing.

THREE PATHS FOR YOUR CORE SYSTEMS

	BUY OFF-THE-SHELF	BUILD ALONE	BUILD WITH A PARTNER
Fit to your operation	Built for the industry average; you adapt to it	Exact fit, if the project survives	Exact fit, designed from your real workflows
Cost over time	Rises ~13% per year; roughly half of licenses go unused	Unpredictable; one in six projects overruns by 200% on average	Fixed-scope phases with a defined asset at the end
Biggest risk	Paying forever for software that never quite fits	Key-person loss, scope creep, permanent maintenance burden	Choosing the wrong partner
Who owns the asset	The vendor	You	You
Deployment success evidence	~95% of GenAI pilots show no P&L impact	Succeeds one third as often as partnered builds	~67% success rate (MIT NANDA 2025)

The first step is small by design. The AI Readiness Evaluation is a short, fixed-scope assessment that maps where your operation loses money to misfit software and manual workarounds, what a custom build would return, and in what order to sequence it. If the numbers say your current stack is the right answer, the evaluation will say that too. You will know either way before committing to anything.

NEXT STEP

Start with the **AI Readiness Evaluation.**

Fixed scope, concrete findings, and a build-or-keep answer you can hold us to.

Contact us

(623) 323-1232
team@avolis.ai

Zylo, [2025 SaaS Management Index press release](#) — 152 apps at companies under 500 employees. [zylo.com](#)

Zylo, ["How Much Is Wasted on SaaS Spend?"](#) — 54% of licenses used, 2026 data. [zylo.com](#)

Pendo, [2019 Feature Adoption Report](#) — 80% of features rarely or never used. [pendo.io](#)

Vertice, [SaaS Inflation Rate tracker](#) — ~13% annual SaaS price inflation. [vertice.one](#)

Murty, Dadlani, Das, ["How Much Time and Energy Do We Waste Toggling Between Applications?"](#), Harvard Business Review, 2022. [hbr.org](#)

Flyvbjerg & Budzier, ["Why Your IT Project May Be Riskier Than You Think"](#), Harvard Business Review, 2011 — one in six projects, 200% average overrun ([open version](#)). [hbr.org](#)

Flyvbjerg et al., ["The Empirical Reality of IT Project Cost Overruns"](#), Journal of Management Information Systems, 2022 — power-law distribution, 5,392 projects. [arxiv.org](#)

RAND Corporation, ["The Root Causes of Failure for Artificial Intelligence Projects"](#), 2024 — 84% cite leadership-driven root causes. [rand.org](#)

MIT NANDA, ["The GenAI Divide: State of AI in Business 2025"](#) — 95% of pilots, 67% vs one-third deployment success ([coverage](#)). [mlq.ai](#)

Marco Argenti (Goldman Sachs CIO) on Bloomberg Odd Lots, April 2026, [via AI Street](#). [ai-street.co](#)

Retool, ["AI Build vs. Buy Report 2026"](#) — 35% replaced a SaaS tool; 78% building more. [retool.com](#)

Michael Hammer, ["Reengineering Work: Don't Automate, Obliterate"](#), Harvard Business Review, 1990. [hbr.org](#)

Brynjolfsson & Hitt, ["Beyond Computation"](#), Journal of Economic Perspectives, 2000 — complementary organizational change. [wharton.upenn.edu](#)

Klarna, [AI assistant results press release](#), 2024. [klarna.com](#)

CX Today, [Klarna stack change, accurately reported](#), 2024. [cxtoday.com](#)

CX Dive, [Klarna rehiring human agents](#), 2025. [customerexperiencedive.com](#)



The Avolis Research Group

Avolis's in-house research practice, focused on how operations-heavy small and mid-sized businesses actually adopt AI. It synthesizes primary economic research, government survey data, and results from real implementations into practical, vendor-neutral guidance.