Alpha

Build vs. Buy in 2026: How AI Shifted the Risk Landscape for Insurers

August 2026

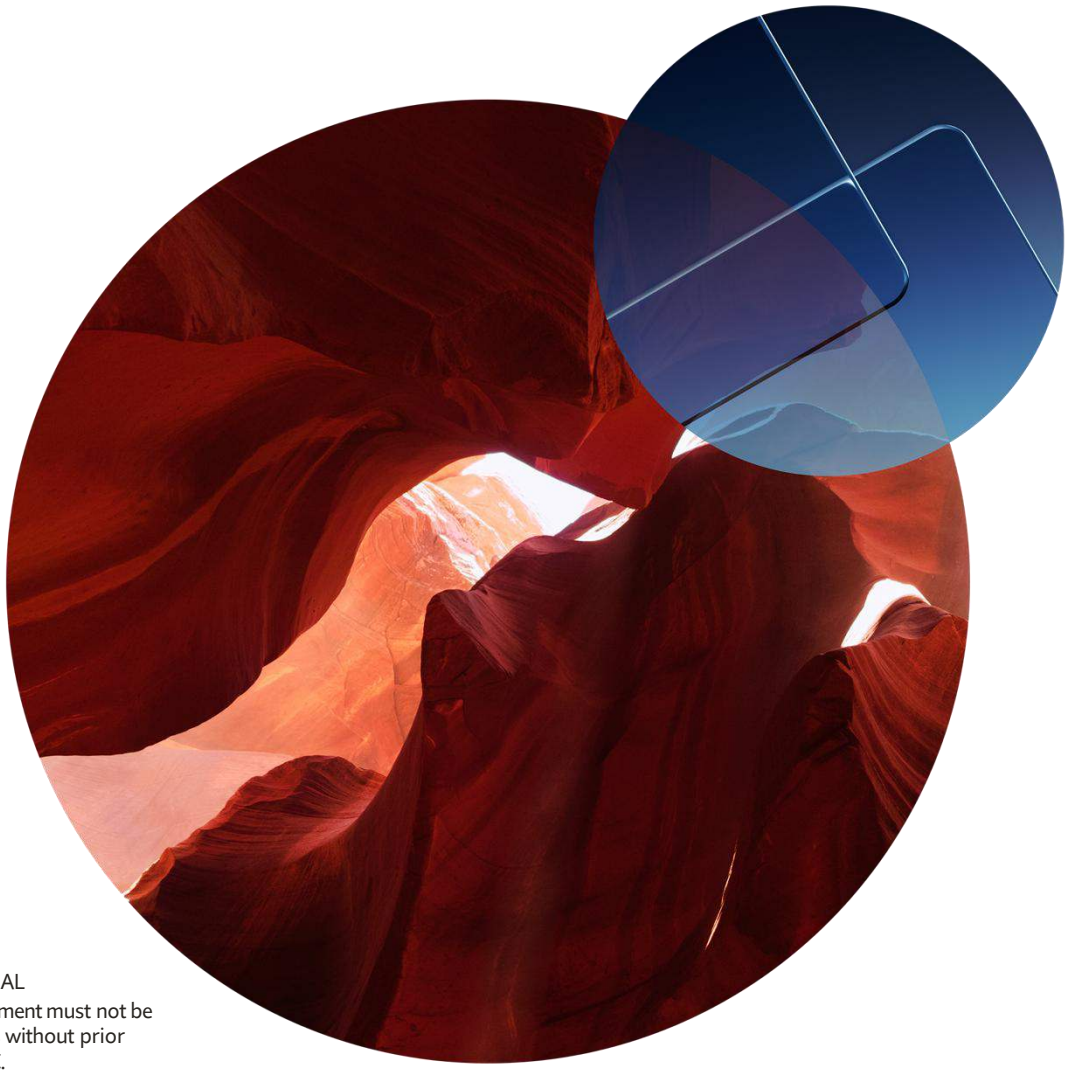

Build vs. Buy in 2026

The Situation

"Build vs. Buy" has long been a decision weighing internal software development capabilities and budget against the cost of contracting with a vendor. With AI in the mix, the decision is now a question of where the capability sits in relation to regulated decision-making and the runtime cost of leveraging AI. For insurers, rating, pricing, and underwriting workflows now offer integrations with AI tooling, which can complicate the regulatory load associated with each capability.

What Changed?

Writing the Code

Front-end software development has become much more accessible thanks to AI and vibe coding, a development workflow where applications are created using natural language prompts. Instead of spending money on application development, small teams can generate software with minimal long-term investment. With vibe coding, technology teams can focus more on backend infrastructure management, cloud integration, and system maintenance after the application is built.

But vibe coding is not a flawless end-all solution for software development needs. AI models learn their patterns and behaviors from large, sometimes insecure, codebases and data sources and might lack the defensive intuition that experienced developers learn over time. There's an important caveat with vibe coding: technology teams can generate code faster and cheaper, but the code should still undergo human review before go-live.

The Risk Equation

Thanks to AI, Automated Decision-Making Technology (ADMT) is a new variable added to the risk equation for software development. ADMT encompasses machine learning-driven tools that process personal customer information to substantially replace human decision-making, often resulting in changes to employment, lending, housing, or financial outcomes. ADMT could allow AI agents to make core business decisions for insurers, including rating, underwriting, and pricing. However, the risk of making a non-compliant decision is high because of the proximity to sensitive personal information.

To date, there are limited regulations around the use of ADMT in financial decisions, which leaves the risk landscape wide open. California passed legislation in January 2026 instituting strict transparency and disclosure laws on firms electing to use ADMT for consumer-related decisions, but few states have followed suit.

There are many cases where the lack of regulation could cause a legal nightmare for firms leveraging ADMT. As an example, an AI agent gathering information to underwrite a life insurance policy cross-references publicly available data, but that data could include extraneous information not directly related to the underwriting task. The AI agent might engage in proxy discrimination, whereby a decision is made using "neutral" customer information (i.e. credit score, medical records) that is indirectly correlated with characteristics in a protected class (i.e. race, ethnicity, sexuality).

The ownership of the regulatory implications of ADMT depends on where the system or application sits in a particular workflow. Rating, pricing, and underwriting capabilities usually live in dedicated external engines. If an application orchestrates workflow and exchanges data with those engines and does not make the decision itself, the heavy compliance burden remains with the engine. But if you build an application yourself, you own the responsibility of managing the regulatory load.

Trade-Offs

Costs & Contracting

If you build the capability, you own everything you build for life. This includes the deployment pipeline, defect resolution, cloud infrastructure, system integrations, security posture, and audit trail. Maintenance becomes more expensive after 18 months, and you are responsible for defining and staffing the operating model. You also define the governance model and manage regulatory compliance.

If you buy the capability, you "rent" an application and inherit the vendor's SaaS infrastructure, patching, security attestations, and compliance scaffolding, leaving the regulatory burden with the vendor. However, there is little room for differentiation, and this path requires long-term contracting. If you want to exit, sunsetting a vendor application is typically a 12–18-month project.

Operating Burden: Running Agentic AI

Owning the build now means owning a cost category that did not exist a few years ago: the runtime cost of the AI itself. As insurers move from single AI calls to agentic workflows – where one AI agent hands a task to another – cost accumulates at each step of the workflow. A single customer interaction can turn into multiple AI calls across different functions, which adds up quickly.

Traditional IT budgets are not designed for this compounding cost. Each agent in a chain tends to repeat work (re-reading context, re-querying APIs the previous agent already called), and a single misconfigured retry threshold can exhaust a monthly budget in hours. Industry research suggests agentic systems can consume five to nine times more tokens per workflow than standard generative AI. In one cited example, an ecommerce firm that implemented an agentic supply chain organizer saw their monthly cost increase from roughly $5,000 to $50,000, because the model got stuck in recursive loops during high-volume periods.

Therefore, firms require a new discipline to govern agentic spend. One approach is cost attribution, or FinOps for AI. Cost attribution binds every agent action to an accountable owner, budget, workflow, and business outcome and treats each action as an auditable unit of work. When you buy, the vendor largely absorbs this complexity and prices it into the subscription. When you build, the price becomes one more part of the operating model you must design, staff, and fund for the life of the application.

Owning the build now means owning a cost category that did not exist a few years ago: the runtime cost of the AI itself.

Decision Framework

Guiding Questions

In today's Build vs. Buy discussion, two central questions should guide the decision:

1. Is the capability a standard function, or does it create real differentiation?
2. Does the solution make a regulated decision, or connect with the systems responsible for one?

Those answers point to four practical courses of action.

Option 1: Configure (Buy) Commodity That Makes a Regulated Decision

When the capability does not provide competitive advantage, the lower-risk path is typically to tailor an established vendor platform rather than develop one internally. Examples include a third-party rating engine, suitability engine, or underwriting rules engine, where product rules and pricing parameters are configured within the vendor's pre-built environment.

Consider a suitability engine: it is a configurable tool that evaluates application data and product details against the suitability rules you establish. Your role would be to manage how the engine fits into the broader system architecture and determine how its workflow should operate within the existing front-office model, so the team can concentrate on improving suitability assessment instead of gathering and reconciling information manually. Vendor platforms may also include configurable reporting dashboards, enabling you to set metrics for work volume, quote-to-bind timing, suitability review duration, and customer satisfaction.

Under this model, the vendor owns maintenance of the core decisioning engine, product updates, and the supporting controls and functionality. You still remain engaged on regulatory expectations, but the vendor carries the primary compliance burden.

Option 2: Configure/Light Build Commodity that Interacts with Regulated Systems

For capabilities such as workflow orchestration, customer experience management, or data services that sit around the regulated decision rather than making it, insurers can buy a standard product or service and make limited customizations where needed.

For example, a new business workflow tool that tracks the tasks associated with an application is a great use case for this option. Commodity new business workflow tools are built with capabilities that manage hand-off, assignment, and prioritization across the submission, quote, referral, underwriting review, and bind phases of a new business application.

With this use case, you own the integration with your existing systems, the process design based on your business workflow, user adoption, and any light extensions required to fit your operating model. The vendor remains responsible for building the underlying product, standard features, and ongoing technical maintenance. This option is great for improving or streamlining an

established workflow or data ingestion process while remaining outside of the regulated path. Because the capability does not make the regulated decision itself, the risk is lower, and insurers can preserve internal engineering capacity for more valuable or differentiating use cases.

Option 3: Build Differentiator that Makes Regulated Decisions

If your rating, pricing, or underwriting capability is a true differentiator and makes a regulated decision, the highest-risk and most self-sufficient option is to build the capability in-house. This option may seem like the most attractive because of AI's market appeal, but relying on automated decisioning tools increases the compliance burden significantly.

This option reduces the long-term cost of contracting with a vendor, but you are responsible for the platform's design, rules development, testing, governance, and maintenance forever. Building in-house only makes sense if (1) you have mature governance and controls for regulated systems in place, and (2) the value of differentiation outweighs the cost, complexity, and scrutiny associated with the build.

If you're a mature, nationally credited insurer with bespoke product offerings across multiple jurisdictions, it might make sense to build an AI-enabled underwriting system yourself. However, be weary of overinvesting in AI-enabled decision workflows, because costs compound exponentially, as explained above.

Overall, we would advise against this option, because relying on AI to make core business decisions through a self-built platform is a slippery slope in terms of audits and exposure to scrutiny.

Option 4: AI-First Build

Leveraging AI to build a differentiating platform that interacts with regulated systems instead of having AI *make* the regulated decision itself is often the most attractive option, especially now that AI makes front-end and workflow development faster and cheaper. This option is best when you are large insurer offering multiple products across jurisdictions, because an AI-first build allows you to efficiently and accurately configure bespoke product, suitability, or underwriting rules while remaining outside of the regulated path.

Building with AI is faster, more efficient, and generally less expensive than full-time software developers thanks to vibe coding (unless you rely on agentic workflows whose costs compound significantly). A common pitfall of the AI-first build is that insurers neglect change controls and governance, assuming AI will ensure workflows are regulatorily compliant and that the code will undergo the same rigorous testing that the traditional SDLC enforces.

Since you own the build in this case, you need strong controls around all product testing, integration, operational support, and maintenance. While you remain outside of regulators' scrutiny, you are responsible for ensuring your product is functional, marketable, and produces accurate results. Our experts at Alpha FMC can help you narrow down the vendor landscape to select a vibe code platform that meets your business and operational requirements.

Alpha's Advantage

We have deep industry expertise in Property & Casualty and Life & Annuities insurance that informs our strategic advice. We offer guidance in defining a business decision framework, developing a target operating model, and implementing scalable solutions. Alpha helps firms assess where to configure proven platforms, where to pursue targeted build, and where AI-led innovation can create real competitive advantage.

If you're interested in learning more about how Alpha can shepherd your firm into the future of insurance, please contact us via LinkedIn or our website.

Contributors

Taylor Fay | Partner

Taylor is an expert in leveraging technology to re-engineer new business and underwriting processes for insurers, resulting in increased speed-to-quote, enhanced agent/broker engagement, and boosted sales volume. He is skilled at translating complex business needs into product innovation, building high-performing agile teams, and managing cross-departmental delivery to achieve superior results.

E: taylor.fay@alphafmc.com

P: (612) 840-9711

Will Newton | Senior Client Director

Will is a seasoned insurance and technology leader in Alpha's Insurance practice, with deep expertise across MGA and program operations, carrier workflow, technology delivery, and strategic business consulting. He has led engagements with tier-one carriers and brokers, delivering across operational, strategic, and data-driven initiatives.

E: william.newton@alphafmc.com

P: (617) 455-1616

Abby Cohen | Consultant

Abby is a Consultant with Alpha's Insurance practice. She has experience developing and executing large-scale operations and technology transformation strategies for national insurers. Abby specializes in technical project management for software development engagements. Within Alpha, Abby is actively involved in marketing efforts for the Insurance practice.

E: abigail.cohen@alphafmc.com

P: (718) 362-4068

References

- Badenhausen, Tess. “NCOIL Special Committee on Race in Insurance Underwriting Holds Virtual Interim Meeting: Adopted Definition of ‘Proxy Discrimination.’” NCOIL. March 8, 2021. https://ncoil.org/news/ncoil-special-committee-on-race-in-insurance-underwriting-holds-virtual-interim-meeting-adopted-definition-of-proxy-discrimination/.

- “CCPA Updates, Cybersecurity Audits, Risk Assessments, Automated Decisionmaking Technology (ADMT), and Insurance Regulations.” https://cppa.ca.gov/regulations/ccpa_updates.html

- Cloud Security Alliance AI Safety Initiative. “CSA Research Note: AI-Generated Code Vulnerability Surge 2026.” April 4, 2026. https://labs.cloudsecurityalliance.org/research/csa-research-note-ai-generated-code-vulnerability-surge-2026/.

- National Association of Insurance Commissioners. “Use of Artificial Intelligence Systems by Insurers.” April 3 2026. https://content.naic.org/insurance-topics/artificial-intelligence.

- Nussbaum, Itay. “Attack Surface.” Apiiro. February 3, 2026. https://apiiro.com/blog/attack-surface/.

- Revenium. “AI Attribution.” June 10, 2026. https://www.revenium.ai/post/ai-attribution/.

Confidence in the complex