

The Delegation Boundary in Commercial Underwriting

In submission to bind, high AI fit and retained decision authority never meet.

Raymond Carlson

CARLSON BUSINESS VALUE ENGINEERING™
U.S. COPYRIGHT REG. NO. TXU 2-524-043

AI can do the work. Who holds the authority?

None of the three architecture domains substitutes for the others. Strategic architecture sets the value and risk of each capability. Operational architecture sets workforce participation, accountability and the authority governing the work. Technology architecture identifies where language models, rules, analytical AI or agents can perform or support it.

A high fit for language models does not transfer decision authority. A high business value does not imply high demand for a language model. A low-authority capability may not need a language model at all.

THE QUESTION

The enterprise should not ask only where AI can do the work. It must also ask where AI can act, and under whose authority.

Three firms are answering this question. None of them finds the boundary.

PA Consulting¹

PA Consulting argues that the more valuable and the more risky a piece of work is, the more a person should be involved in it, and it recommends mapping which roles do what in each workflow. We agree that value and risk matter. But value and risk do not tell you who must make the decision. In the submission-to-bind model, the highest-value capability among those AI does best is generating product recommendations. It is rated high on four value dimensions, including risk and compliance, and the carrier can still hand it to AI. Roles and workflows are also the wrong place to record the answer, because they change every time a carrier reorganizes. The capability does not.

Decerto²

Decerto, which builds underwriting rules engines, says AI should handle the data work and underwriters the judgment work, with a rules engine routing each submission to bind, decline or refer. That is the right instinct, but data and judgment are not where the line falls. Recommending a premium audit and recommending products both take judgment, and the carrier can still delegate them. The line falls where the carrier approves and binds.

Workday³

Workday's new Work Intelligence breaks each job into the tasks people do and connects those tasks to more than 75,000 skills, to show where AI can help. That describes what the workforce does and what AI could do. It cannot tell a carrier what it may let AI decide, because authority is not a skill. Authority belongs to the capability, whoever performs it.

¹ Ben Edwards, "Why AI Transformation Is Really Workforce Transformation," in *The New Rules of Execution: Redesigning Work to Unlock AI Value*, Raconteur, in partnership with PA Consulting, October 19, 2026.

² Mariusz Zagajewski, "AI in Underwriting: Will It Replace Underwriters? A 2026 Playbook for P&C Carriers," Decerto, updated May 5, 2026.

³ Workday, "Workday Introduces Work Intelligence to Help Customers Design Work for AI," October 13, 2026.

High AI fit and retained authority never meet.

This edition tests the proposition against a single operating model: commercial P&C submission to bind, from clearing a submission through binding, renewal and premium audit. It is the operating core of a commercial carrier, and the model descends from a four-year underwriting transformation at a global commercial insurer. Of its capabilities, 561 carry both a language model fit rating and an authority rating.

LANGUAGE MODEL FIT	HIGH AUTHORITY	MEDIUM	LOW	TOTAL
High	0	0	28	28
Medium	6	36	11	53
Low	1	52	427	480

Commercial P&C submission-to-bind model. Functional capabilities rated on both dimensions.

Every one of the 28 capabilities where language models fit best carries low authority. None of the seven capabilities where the carrier keeps high authority is a strong fit: six are a medium fit and one is low.

The boundary shows one row down. Among capabilities with a medium fit, 42 of 53 retain medium or high authority. As the work moves toward the underwriting judgment, authority returns to the underwriter.

THE SCOPE

The broader corpus holds models built for different purposes and at different operational depths. This edition deliberately tests the proposition against one operating model rather than presenting a blended statistic.

The verbs draw the line.

Every one of the seven high-authority capabilities approves or binds. The 28 capture, maintain, search, create, document, recommend and generate.

HIGH AUTHORITY · 7	BEST LANGUAGE MODEL FIT · 7 OF 28
Approve an application	Capture underwriting rationale
Bind coverage	Maintain manuscript forms and wording
Approve binding of facultative reinsurance	Search for existing party information
Approve binding of multiple policies	Create a quote proposal
Approve non-renewal of a policy	Document the rationale for credits and debits
Approve a multinational program as lead entity	Recommend a premium audit
Approve a multinational program as servicing entity	Generate product recommendations

BUSINESS CAPABILITY		
Ability to assess underwriting risk		
FUNCTIONAL CAPABILITY	AI FIT	AUTHORITY
Ability to capture risk evaluation assessments	High	Low
Ability to determine terrorism applicability	Medium	Medium
Ability to determine insurance to value	Medium	Medium
Ability to override risk exposure information	Medium	Medium
Ability to assess insured products	Medium	Medium

THE RECORD AND THE JUDGMENT

Within a single business capability, language models fit the record of the risk assessment better than the judgments inside it.

Take the work off the desk. Leave the authority with the underwriter.

The 28 are not jobs AI replaces. They are a delegation layer: the record, preparation and communication around the underwriting decision. AI can take that work off the underwriter's desk without taking underwriting authority away from the underwriter.

Value does not settle the question either. The highest-value capability among the 28, generating product recommendations, is rated high for revenue, risk and compliance, customer experience and strategic agility, and it still carries low authority.

01

Delegate the record

The 28 high-fit, low-authority capabilities are where AI can run the work end to end.

02

Augment the judgment

The 36 medium-fit, medium-authority capabilities are where AI supports a decision the underwriter still owns.

03

Reserve the approval

The seven high-authority approvals and binds stay with the authority the carrier grants.

The remaining 490 capabilities are a weak fit for language models or carry little authority. Where they automate, rules engines and analytical AI do the work within the same authority limits, and much of it stays as it is.

WEDNESDAY

Trust Starts at Hire takes up the next question. If the capability sets the authority required, who grants it, and how is it represented?

AI can take the work off the underwriter's desk. It cannot take the authority the carrier grants the underwriter.