

Choosing the Simplest Effective Architecture for an Outsourced Service

An Executive Decision on Necessary Autonomy

Phil Hatch | Akholi

October 23, 2026

Executives should buy the least complex delivery architecture that meets the business need at an acceptable cost and risk. An agentic design earns its place by outperforming feasible simpler alternatives on the same accepted outcome.

Make Autonomy a Choice Rather Than a Requirement

The executive sponsor should define the purchased outcome before approving an agentic architecture. The sourcing question is whether the work needs an agent to select and execute a changing sequence of actions. A service can create substantial value through process simplification, deterministic automation, or bounded human assistance without requiring that additional freedom.

Stanford's successful-deployment case study finds value across simpler and agentic approaches; its selected sample does not estimate population success rates.¹ The implication proposed here is an architecture comparison, not a presumption that agents should always be selected or rejected.

¹ Elisa Pereira et al., *The Enterprise AI Playbook: Lessons from 51 Successful Deployments*, Stanford Digital Economy Lab, April 2026, methodology and chapters 1–3 and 8; purposive successful-case sample, https://digitaleconomy.stanford.edu/app/uploads/2026/03/EnterpriseAIPlaybook_PereiraGraylinBrynjolfsson.pdf.

A hypothetical supplier-invoice service illustrates the choice. Standard invoices with stable rules may need reliable extraction and matching. Exceptions may need a person or an agent that searches several systems. Applying the same architecture to both populations can add cost and authority where neither improves the accepted result.

Compare the Same Business Function

Management should specify the outcome, workload mix, quality threshold, permitted actions, completion deadline, and recovery requirement. Compare feasible alternatives against that common definition. A faster draft is not equivalent to a completed service if the buyer must perform additional review, correction, or system entry.

The study's accounts of manual repair and apparent completion without the corresponding result show why the comparison should follow the full delivery method.² These observations do not prove that a simpler architecture would have performed better. They identify work that a credible comparison must count.

Include buyer and provider dependencies in each option. An apparently inexpensive agent can require new access, clean data, specialist oversight, or a complex recovery route. A simpler option can also fail if its rules are costly to maintain or its exceptions overwhelm people. The comparison should expose these conditions rather than attach a maturity label to the technology.

Use an Architecture Decision Record

Figure 1 proposes a record connecting business need, uncertainty, delegated action, and total operating burden. State why each degree of autonomy is necessary. If the provider cannot explain which uncertainty requires autonomous choice, management should request a simpler design or a test that establishes the added value.

Run representative cases through the alternatives where feasible. Record accepted outcomes, human effort, elapsed time, operating cost, failures, and the work left unresolved. Keep the workload and acceptance rule comparable. Differences in case difficulty should not be mistaken for architectural advantage.

² Phil Hatch, *Agentic AI Outsourcing at Tier 1 Financial Institutions: Full Study Report*, version 1.8.2 (Akholi, forthcoming manuscript dated November 1, 2026), Part 5, “The Five Manual Alteration Methods” and “Disclosure of Manual Alteration,” table 5.8, <https://doi.org/10.2139/ssrn.7164559>; Hatch, *Agentic AI Outsourcing*, Part 3, AI Technology, sec. 4.1, “Accuracy of the Tools’ Own Success and Failure Reports,” indicator C172.

McKinsey’s operating-economics analysis identifies costs from context, refinement, and orchestration beyond model prices.³ The executive comparison should therefore include the complete service cost and variability. A cheaper model call does not establish a cheaper accepted outcome if the design creates repeated tool use and repair.

Figure 1. Choose the necessary level of autonomy

Decision area	Comparison question	Required evidence
Business need	What outcome must improve	Common scope and acceptance rule
Autonomy	Which uncertainty requires choice	Simpler alternative comparison
Operating burden	What supports accepted work	Full cost human effort and exceptions
Expansion	When should the design change	Limits assumptions and review triggers

Author’s proposed operating framework. Adapt to the service and applicable jurisdiction; not an empirically validated standard.

Source: Phil Hatch, proposed framework developed for this paper.

Bound the Architecture That Wins

An agentic option may be justified where the sequence is variable, evidence is distributed, and the service can recognize success and recover from failure. Management should still decide what the agent may not do. A useful planning capability does not automatically justify financial execution, customer commitments, or unrestricted access.

For regulated operations, compare options under the applicable entity and activity requirements. EU, UK, Swiss, and US services may have different approval, privacy, record, and professional-duty conditions. The common business outcome should remain visible, while the operating design reflects the obligations of the location in which it is delivered.

The provider should identify the assumptions under which its preferred architecture remains suitable. Changes in volume, exception mix, tool access, or legal scope may alter the decision. Record the conditions that trigger a fresh comparison rather than assume the initial architecture is permanent.

³ Lari Härmäläinen et al., “Is That AI Agent Worth It? Agentic Economics and the Modern Operating Model,” McKinsey & Company, July 13, 2026, “Why the operating expenditure bill keeps rising” and “The implications for competitive dynamics,” <https://www.mckinsey.com/capabilities/quantumblack/our-insights/is-that-ai-agent-worth-it-agentic-economics-and-the-modern-operating-model>.

Approve Evidence Rather Than Architectural Fashion

A reasonable objection is that comparative testing delays deployment and may overlook future capability. Use a bounded pilot and stage the commitment. The organization can preserve an option to expand while requiring the current service to justify its present cost and authority. Future capability should remain a forecast with explicit assumptions.

The executive approval should name the selected architecture, rejected alternatives, comparative evidence, operating restrictions, and review triggers. Where evidence is incomplete, state what must be learned and cap the exposure during that learning period. A vendor demonstration alone should not close the investment decision.

The simplest effective architecture is the one that meets the accepted business need with the least unnecessary operating burden. It may include agents, people, conventional software, or a deliberate combination. The board should require management to explain the choice through service evidence rather than the ambition of the technology label.

Evidence and Source Note

The study is a nonrepresentative cross-section of 850 vendor delivery-team respondents on 150 engagements serving Tier 1 financial institutions, collected during January–June 2026. Neither banks nor vendors corroborated the accounts. It establishes reported experience, not causation, legal breach, or market prevalence. Its application here is a mechanism to investigate, not an estimate of conditions in other industries. The author-supplied version 1.8.2 is dated November 1, 2026 and is cited as a forthcoming manuscript consulted October 23, 2026.⁴

This is an analytical decision paper based on the supplied study and selected primary public sources checked through October 23, 2026. The original contribution is an outsourcing architecture decision record comparing necessary autonomy, accepted outcomes, retained effort, and conditions for expansion. The proposed method has not been independently validated. Legal requirements depend on the entity, activity, jurisdiction, and effective provisions; supervisory observations, research, and professional commentary have distinct authority.

Bibliography

Hatch, Phil. *Agentic AI Outsourcing at Tier 1 Financial Institutions: Full Study Report*. Version 1.8.2. Akholi. Forthcoming manuscript dated November 1, 2026. Author-supplied manuscript consulted October 23, 2026. <https://doi.org/10.2139/ssrn.7164559>.

Hämäläinen, Lari, Mark Patel, Sven Blumberg, Tanguy Catlin, and Wasim Lala. “Is That AI Agent Worth It? Agentic Economics and the Modern Operating Model.” McKinsey & Company. July 13, 2026. <https://www.mckinsey.com/capabilities/quantumblack/our-insights/is-that-ai-agent-worth-it-agentic-economics-and-the-modern-operating-model>.

Pereira, Elisa, Alvin Wang Graylin, and Erik Brynjolfsson. *The Enterprise AI Playbook: Lessons from 51 Successful Deployments*. Stanford Digital Economy Lab. April 2026. https://digitaleconomy.stanford.edu/app/uploads/2026/03/EnterpriseAIPlaybook_PereiraGraylinBrynjolfsson.pdf.

⁴ Hatch, *Agentic AI Outsourcing*, Introduction, “Executive Summary” and “Evidentiary Basis and Its Limits.”