Contracting with the Agentic Workforce

Who keeps the savings when your provider's AI agents do the work

Cliff Justice, Founder, Catalyq · Version 1.1 · 30 October 2026

Key takeaways

1. **Contracts written for people-based delivery leave both the risk and the savings with the provider.** Providers are moving work to AI agents under agreements that say little about who answers for an agent's error or who benefits when fewer people do the work.
2. **Act at renewal or sooner.** Price reviews, gain sharing and productivity commitments decide whether lower delivery costs reach the client, and buyers are already reopening contracts mid-term to secure them.
3. **Accountability follows control, and the provider remains responsible for the tools it selects.** The client answers for the guardrails it sets and the decisions it keeps. The provider answers for its delivery, including AI models it chose but does not own, much as it answers for its subcontractors.
4. **Keep governing after signing.** Agents change as models, data and instructions change, so the client needs records, testing and change rights for the whole term.
5. **The client's own team needs new skills.** Someone has to oversee the provider's agents, own the handoffs between provider and client, and measure the whole process end to end.

Delivery changed faster than the contracts

Most services contracts in force were priced and governed for work done by people. As providers move that work to AI agents, those contracts leave two things with the provider: the risk when an agent makes a mistake, and the savings when fewer people do the work. This paper sets out Catalyq's view of how a client can bring both back under its control.

Genpact's chief financial officer told investors in August that the company is tracking to deliver over a billion dollars in agentic contract value in 2026, five times more than in 2025. ISG's 2026 study of business process outsourcing found that more than half of buyers have moved beyond purely labor-led delivery, though only 14% are getting AI-first outcomes.

HFS reports that most application and infrastructure agreements are still priced on effort, headcount or tickets, including deals signed in 2024 and 2025. In an HFS survey of more than 500 executives, produced in partnership with the provider Firstsource, 42% of BPO commercial models were based on effort or headcount. An agreement priced on people says little about who answers when an agent makes a mistake, what records the client can see, or what happens to the savings when fewer people do the work.

Buyers are dissatisfied with the contracts they have. In an ISG survey of 150 insurance decision-makers, commissioned by an AI software vendor, only 8% of outsourcing buyers would renew their contracts as they stand, 64% would renegotiate to capture AI efficiency, and 23% would rework terms entirely. The sample covers one industry, but the direction matches what HFS reports across sectors.

Law firms including Mayer Brown and Allens have published guidance on contracting for agentic delivery, and Allens warns that legacy outsourcing agreements are unlikely to deal appropriately with AI-powered outsourcing. Few buyers can act on that guidance yet: ISG found that only 21% of enterprises have the skills to govern AI-enabled BPO, while two-thirds expect their providers to lead AI adoption.

What is at stake

The contracts in question cover large sums. In the Americas alone, new and renewed managed services contracts worth $5 million or more a year added up to $23.4 billion of annual contract value in 2025. That figure excludes smaller contracts and other regions, and because these contracts run for several years, the value in force at any time is a multiple of it.

Morgan Lewis notes that suppliers often predict 40% to 70% productivity gains from their AI solutions, and that achieving them is often harder, requiring operating model redesign and retraining. Those figures are suppliers' predictions. Even a small share of the gains that are realized would be material on a spend base of this size.

How those gains are split is contested, and the client's share depends on its contract. Genpact's chief executive told investors that clients moving to agentic delivery showed more than 300 basis points of gross margin expansion for Genpact, which shows that providers are capturing part of the gain, though not whether clients shared in it. An IDC consultant writes that when providers' delivery costs fall,

contract rates rarely follow. Buyer expectations are divided: in an IDC survey cited in the same piece, 55% of IT buyers expected AI-powered managed services to cost more than traditional ones, and 37% expected them to cost less. HFS reports buyers reopening contracts within 24 months of signing to capture AI-driven efficiency sooner.

In the same HFS survey, executives expected outcome-based and shared-revenue pricing to rise from 20% to 39% of BPO commercial models within three years. New contracts will increasingly price for automation from the start, so the largest gap sits in agreements already signed on people-based terms. Every agreement, old or new, will still need governance during the term, because agents keep changing after the contract is signed.

Who is accountable

Responsibility for an agent's actions rests with people and organizations. Most agent errors are operational and get fixed in the normal course of work. Some reach further: an agent that applies the wrong rule to thousands of refunds, journal entries that lead to a misstatement, or a data exposure that brings in regulators. When that happens, the executive team will ask who answers for it, and the contract should already say.

If an error creates material liability or public damage to the company's reputation, the question reaches the board, and the board cannot hand it to the provider. Under Delaware's Caremark doctrine, directors must make a good-faith effort to oversee risks central to the business. Lawyers at WilmerHale note that no Delaware court has yet decided a Caremark claim based on a failure to oversee AI, and legal commentators agree that delegating AI work to a vendor does not delegate the board's oversight duty. In an EY review of leading companies' disclosures, about 40% now assign AI oversight to a board committee, up from 11% in 2024.

Accountability follows control. This has been the working rule of outsourcing for decades: the party with the right and the practical ability to make a decision or prevent a failure answers for the result. It still works when agents do the work, provided the contract spells out who controls what. The rule decides who bears the cost between the parties. To customers and regulators, the client remains answerable for what is done in its name, whatever the contract says.

Decision or outcome	Client answers for	Provider answers for
Risk appetite and guardrails	Setting them and writing them into the contract	Operating within them

Decision or outcome	Client answers for	Provider answers for
Choice of AI model and tools	Any technology it mandates	Everything it selects for its own delivery, including models it does not own
An agent's authority	Approving any expansion	Staying within approved limits
Sign-off on output	Decisions it reserves to itself	Meeting service levels and correcting errors
Incidents	Its right to suspend an agent	Remediation and reversal
Value delivered	Tracking it through the contract's savings terms	Reporting cost and productivity honestly

Where the rule is contested. Providers argue that they cannot answer for AI models they do not control, and in a narrow sense they are right: no provider can guarantee how another company's model will behave. Buyers heard the same argument when providers moved work onto cloud platforms, and the answer then still applies. The provider chose the model, built its service around it and charges for the result. It should commit to what it controls (testing, guardrails, records, remedies and notice of model changes) and price or share any remaining risk openly. Mayer Brown takes the same position, treating the model provider like a subcontractor the provider selected. It also notes that providers may price the exposure or seek narrow exceptions, and that a provider should be excused for technology the client required it to use if it raised its concerns beforehand.

Agent failures often have causes on both sides, and a client that approves guardrails shares responsibility for what it approved. Contracts can handle shared causes with joint root-cause reviews, fees placed at risk and earned back, and gain sharing that rewards both parties for fewer errors. Some losses involve no third party at all, such as an erroneous payment or a corrupted record, so the agreement should say how those losses are recovered. Three fixes are common: treat recovery of erroneous payments and the cost of correction as direct damages that the waiver of consequential losses does not exclude; set separate or higher caps for breaches of data, confidentiality, intellectual property and regulation, as Morgan Lewis recommends; and close the gap where a direct operational loss falls outside the third-party indemnities.

Unexpected success also needs an owner. A sharp fall in the cost of the work should reach the client's price through whichever savings mechanism the contract sets. A surprisingly good result also deserves review. An agent that approves more claims raises throughput and may also raise loss exposure.

Choosing a provider that delivers through agents

When a provider will run agents on a client's work, the controls around those agents belong in the selection criteria, because they decide what happens to the client's data and how fast the provider's cost of delivery can fall. Analyst evaluations such as ISG's study of about 140 agentic AI service providers help build a shortlist. Five requests then test each finalist's claims:

1. Walk through one real agent action from start to finish: who approved it, what access it used, where it was logged and how it would be reversed.
2. Commit to a price per unit of work on comparable work, and to how that price will fall over the term as automation grows.
3. Show the most recent change of AI model and the test results that came with it.
4. Show the data terms of its AI suppliers, including restrictions on training with client data.
5. Describe the human oversight for work like yours, with references from clients at a similar stage.

These controls sit alongside domain expertise, transition capability and price, and they cannot make up for weak knowledge of the process being automated. They matter most in the commercial negotiation: the provider with the most mature controls has the most room to cut its own costs, so that is where to push hardest on the price path over the term.

What the contract needs to say

Catalyq's approach rests on two principles.

- **Define the service by the business function and its outcomes.** Agents, models and instructions will change over the term, so the contract should anchor on what gets done, and the price should follow the work delivered.
- **Assume the agent will change.** The same input to an AI model can produce different outputs, and behavior shifts as models, data and instructions change. A test passed at go-live says little about behavior a year later, so the client needs rights that last the whole term: records, retesting and approval of changes.

Term	What to ask for
Guardrails under change control	The decisions agents may make, value limits, the data and systems they may touch, actions that always need human approval, and actions never permitted

Term	What to ask for
Records and audit	Logs, decision traces, model and version identifiers, test results, and records of exceptions and human intervention, kept for a stated period
Named ownership	A named owner at the provider for each agent workflow, and the client approver for its scope
Suspension and reversal	A client right to suspend an agent or cut its access, a process to revert or return to manual work, and correction of errors at the provider's cost
Model changes	Prompt notice, the likely effect and testing of any substitute model; client approval of material changes, treated as given if the client does not respond within an agreed period; and a client right to reduce the agent's authority or suspend the function while a problem is resolved, with service level relief and agreed costs for any manual fallback
Ongoing testing	Regression tests after every material model, instruction or workflow change, with results reported
Data use across the chain	No training on client data, passed down to model providers, and covering prompts, outputs, logs and test data
Performance measures	Outcome service levels, plus reported handoff, rework and error rates
Human oversight	Minimum oversight staffing and skills, with any cut in human review going through change control
Price per unit of work	A price per invoice, claim or ticket that falls as automation grows, with the pricing assumptions named, benchmarking, and repricing when the assumptions change. Open-book reporting of the provider's own cost only where the client chooses gain sharing on that cost
Exit	Ownership of, or a license to, what was built for the client (its agent configurations, instructions, workflows, test sets and runbooks) and copies of the intervention records; a transition license and documentation for the provider's own platform components; and the agents kept running during transition
Ownership of AI work	Client ownership of, or a license to, the client-specific prompts, configurations and workflows, kept separate from the provider's pre-existing platform. Material generated by AI may carry limited copyright protection, so the contract should state the rights directly
Term and termination	A term no longer than the automation plan needs, termination for convenience with fees that decline over the term, a stated treatment of any unrecovered automation investment, and benchmarking in long deals

Term length matters more than it used to. Some providers are pushing for contracts of five, seven or even ten years to recoup their AI investments, according to Brad Peterson of Mayer Brown. A longer term should come with a price that falls on a stated schedule.

Who keeps the savings. In a people-based contract, savings reach the client through rate cards, headcount changes and annual productivity commitments. Agents weaken that link, because the provider's cost can fall quickly while the price stays tied to a baseline set at signing. Three mechanisms restore it, and most agreements should use at least one.

Mechanism	How it works	Suits
A committed price path	The provider commits at signing to price reductions over the term, sized to its automation plan	A client that wants certainty and can judge the plan
Gain sharing on open books	The provider shares the cost of delivering the work on an open-book basis, with defined cost elements, audit rights and an allowance to recover its investment, and an agreed share of any fall below the baseline returns to the client	Work where the pace of automation is uncertain
An automation reopener	Defined milestones, such as agents taking over a named process step or a set share of volume, trigger a price review	Long contracts signed before the provider's plans are firm

Each mechanism works best on a price per unit of work that the client can track and benchmark. Open books on the provider's own cost belong only where the client chooses gain sharing.

What matters most for a smaller buyer. A mid-market company with one or two outsourcing contracts does not need every term in full. Four carry most of the value: a price review or savings share tied to automation, a right to know where AI does the work and who reviews it, provider responsibility for correcting defective work, and exit terms that let the work move to another provider. The rest can be scaled to the size of the contract.

Example: a finance and accounting agreement signed in 2026

This example comes from a real agreement filed publicly with the SEC, with identifying details withheld. The client is a small public health care company. The provider supplies outsourced finance leadership and accounting services to small public companies. The agreement was signed in 2026, and its services

are delivered mainly by people today, so it shows how a recent contract handles AI before agents do much of the work.

What the agreement gets right. It addresses AI directly. The provider may not put confidential or material non-public information into public AI tools that train on submitted data without the client's written approval. It may use enterprise AI tools that do not train on client data, after telling the client which tools it will use, and it remains primarily responsible for its use of them. A senior professional from the provider serves as the company's chief financial officer, so a named person signs and certifies the filings whatever tools are used.

Where it would fall short as AI does more of the work

Issue	What the agreement says	A proportionate ask at renewal
Pricing	A fixed monthly retainer tied to the services of a named lead	A price review at renewal that takes account of any automation in the work
Records of AI work	Notice of which tools are used, with no record of what they produced	A short log of which deliverables AI helped prepare and who reviewed them
Human review	Not specified for AI-prepared work	Confirmation that a qualified person reviews AI-prepared filings and board materials before sign-off
Changes in tools	Notice when tools are first used	Notice before material changes in the tools used on client work
Ownership of AI work	The provider keeps its prompts, workflows, templates and AI tooling	Copies of the client-specific workflows and checklists a successor would need
Remedies	Service credits, capped at a small share of the monthly fee, as the sole remedy for missed service levels	Reperformance of defective work at the provider's cost, however it was delivered, outside the service credit cap and within the general liability cap
Exit	Return of books, records and deliverables, with a short transition period	The same, plus the client-specific workflows above

What the example shows. Even a carefully written 2026 agreement treats AI mainly as a confidentiality matter. Pricing, review and exit terms are still written for work done by people. For a contract of this size the fixes are modest, and they are easiest to agree before the provider's use of AI grows.

After signing: governing delivery for the term

Agents keep changing after signature, so the client needs a regular governance routine that uses the rights the contract gives it.

- **Quarterly review of what agents did.** Review decision traces and exception records for a sample of agent work, alongside the usual service level report.
- **Change control for models and instructions.** Treat a model change like any material change: notice, a test report, client approval, and a fallback if results slip. Model providers can retune, restrict or retire a model with little notice, so the provider should keep a tested substitute ready.
- **Price per unit of work, tracked over time.** Compare the price the client pays per invoice, claim or ticket with benchmarks and with the provider's automation plan, and use the contract's savings mechanism when the gap widens.
- **Rehearsed incidents.** Agree in advance who can suspend an agent, how its actions are reversed and how fast work returns to people, and test it once a year.
- **Authority reviews.** Any widening of what agents may decide, such as higher value limits or new systems, goes through the same approval as the original guardrails.

The same controls apply to agents the client runs itself, so internal and provider agents can be overseen to one standard.

The client's own team

Agent-based delivery changes the job of the team that manages the provider. That team now oversees a system it can see only as far as the contract requires the provider to report.

Companies rarely outsource a function end to end. They outsource selected processes or parts of them and keep the rest, so agent work crosses the boundary between provider and client. A provider's agents might capture and code invoices while the client's team approves payments, or a provider's agents might triage tickets that in-house staff resolve. That creates five priorities.

1. **Treat the handoffs as the main risk.** Errors gather where a provider's agent passes work to the client's people or systems. Define who owns the work at each handoff, how exceptions move, and how fast escalations are picked up.
2. **Get the client environment ready.** Provider agents need identities, permissions and access to the client's systems, along with clean data and documented processes. Run that readiness as its own project, with a plan and an owner, before agents go live, so a later shortfall is not disputed.

3. **Build oversight skills.** The team moves from tracking service levels to reading decision traces, approving guardrail changes, reviewing test results and signing off model changes.
4. **Keep the process knowledge recoverable.** As agents take over a process, the know-how moves into the provider's configurations and instructions. Strong exit terms and periodic process documentation keep the option to bring the work back or move it.
5. **Measure the whole process end to end.** A provider can meet its service levels while handoffs degrade the overall result. Track cycle time, error rate and cost per transaction across the full process, including the steps the client keeps.

Questions for boards and executives

1. Which executive is accountable for AI across the business, including providers' agents, and which board committee oversees it?
2. Which of our services contracts already involve AI agents, or will at renewal?
3. For each of those contracts, does the agreement say who answers for agent errors, what records we can see, and how savings are shared?
4. How will the board learn of an agent incident at a provider, and what threshold sends it there?
5. Is our own team equipped to oversee agent-based delivery?

A first 90 days for sourcing leaders

1. List every services contract where the provider uses AI or plans to, with renewal dates.
2. Review the top five by value against the contract terms in this paper, and note where each is silent.
3. Ask each provider for its price per unit of work on your contract and how it expects that price to change as automation grows.
4. Agree an interim governance rhythm with those providers: records, model change notice and an incident contact.
5. Choose a savings model and target terms for the next renewal, and start the conversation before the renewal notice period opens.

About Catalyq. Catalyq is a technology platform that helps organizations reshape their operational footprint and services ecosystem in the AI era, starting with the contracts that enable them. Catalyq offers contract reviews against

a standard that covers many of the terms discussed in this paper, and readers should weigh its recommendations with that interest in mind.

Read online: catalyq.com/research/contracting-with-the-agentic-workforce

This paper is for general information and is not legal advice. Figures from vendor-sponsored research are identified as such in the text.