

STRATEGY REPORT · DIGITAL SERVICES & ARTIFICIAL INTELLIGENCE

After the Billable Day

From time sold to value produced: why AI forces digital services providers to rebuild their economics

Axel Tombereau
Managing Partner, Odyssey

October 2026 · ODY-NR-2026-003

OPEN DISTRIBUTION

AI makes the billable day deflationary

Margin now depends on layered revenue and a cost base managed per unit produced

Recovering demand and healthy order books hide the essential point: as long as billing rests on time spent, every productivity gain from AI reduces revenue. Neither volume nor daily rates will protect margin — only a rebuilt pricing unit and cost base will. Two decision-makers are involved: the chief executive, who owns the model, and the shareholder, who funds the transition.

For chief executives

01 AI-driven demand growth masks structural deflation: productivity gains flow to the client

AI is now judged in the executive committee by its effect on the P&L, no longer by IT alone on technical or functional grounds. The client buys a measurable outcome and uses the same tools as its provider to verify it. While billing rests on time, every gain cuts the value of the next contract: “a contract worth \$100 million may now be worth \$80 million,” observes HCLTech. Absent a gain-sharing clause, that gain goes to the client.

02 Layered revenue combines a fixed-fee base, an outcome-linked fee and asset licensing, and rests on instrumented operations

Roughly two-thirds of revenue sits in service lines that AI reshapes or replaces, based on the published portfolios of Capgemini and Sopra Steria. The build demand that offsets this erosion is temporary, so durable revenue will come from running those AI systems. The base covers fixed costs, the outcome-linked fee captures productivity, licensing monetizes assets independently of headcount; costs are managed per unit produced, inference included. The model enters at each contract renewal, through a new pricing unit, a gain-sharing clause and a measurement protocol. In our simulation, a provider at 9% EBITDA margin falls below 1% within three years on time and materials, and reaches 11% with layered revenue.

03 Four groups of entrants are moving in: incumbents keep their position only by leading the alliances they operate within

Model providers, consulting firms, AI-native entrants, and software vendors are converging on services, as Accenture’s acquisition of Faculty for about £740 million (≈ \$1.0bn) illustrates. That convergence exposes how fragile the incumbents’ advantages have become: mobilizing engineers and operating at scale lose their value once agents do the work. Two advantages remain — access to client decision-makers and knowledge of systems in production — which make providers the natural operators at the center of the value chain. Holding that position requires contractual alliances spanning what the client now expects, from AI strategy through to financial measurement.

For shareholders

04 Leaving time and materials takes a multi-year commitment: temporarily weaker ratios are the price of a durably higher margin

For two to three years, reported revenue declines and working capital rises, since outcome-linked revenue is collected after costs are incurred and some development is capitalized. That trough is an investment, not a loss. Once the transition is complete, EBITDA margin exceeds its pre-AI level, reaching 11% against 9% in our simulation, while it falls below 1% if time and materials is maintained. The difficulty is timing: the trough coincides with exit preparation and outlasts the horizon on which most executives are incentivized. Shareholders have three options: (1) extend the holding period; (2) recapitalize with a patient partner, as Blackstone did at Sia Partners for up to €250 million ($\approx$ \$270m); (3) sell ahead of the trough, leaving the value of the transition to the buyer.

05 Capital now prices the execution corpus, an asset most providers have neither contracted for nor booked

Thrive Holdings, backed by OpenAI, raised about \$2 billion in August 2026 to acquire digital services firms among other targets. Beyond teams and order books, these investors are buying: (1) processes, how the work is actually done, exceptions included; (2) business data the provider handles but rarely owns; and (3) execution traces — what was produced, what was reworked, and why — the material an agent needs to work in a real environment. Ownership must sit with the provider, not the client: an asset whose ownership is neither defined in the contract nor carried on the balance sheet has no value in a transaction.

06 Leadership in 2030-2032 will rest on six measures that no dashboard tracks today

All six follow from management decisions and can be tracked by shareholders: (1) the share of revenue independent of time spent; (2) the share of revenue covered by a signed measurement protocol; (3) fully loaded cost per engagement, inference included; (4) assets owned, capitalized, and licensed; (5) a human rework rate on agent output below 30%; (6) neutrality toward model providers. None requires an acquisition or industrial investment, so all can be started immediately, provided finance and financial control (FP&A) are equipped to measure them. They also give shareholders a scorecard for their holdings, the same one a buyer will apply.

WINDOW TO ACT

By 2030-2032, a provider's standing will depend less on its size than on how quickly it shifts to this value-creating model, that is, on how fast it changes its pricing unit and cost base. The opportunity arises at each contract renewal, while client AI build work still funds the transition.

Buying digital services shifts from capacity to measured outcomes

The buying decision is changing: it moves up to the executive committee, is judged by the effect of AI on the P&L, and passes that performance requirement on to providers.

Providers are establishing themselves as the executors of AI programs for their clients, in a market where sold capacity no longer sets prices. They retain control of execution — deployment, integration, and process and workflow redesign — in support of their clients' AI programs. Listed providers guide to 3 to 6% growth in local currency for 2026 — Accenture 3 to 4%, Cognizant 4 to 6.5% — while digital engineering contracts: Endava guides to a decline of 5 to 6%. Demand is narrowing to a few skills and pricing pressure is setting in.

Value comes from redesigning processes and workflows rather than deploying tools. Nearly three-quarters of companies reporting an effect of AI on operating profitability first reorganized the way they work.

The buying decision moves from the IT department to the executive committee, because the subject now engages corporate strategy, the operating model, and financial performance.

— **Budget trade-offs are now settled in the executive committee,** against operating profit rather than billable days.

— **Clients now have alternatives to providers:** building in-house, since nearly one in three companies have dropped a software purchase they could develop with coding agents, or handing deployment to the services arm of a model provider.

— **The negotiation is no longer about price but about sharing the productivity gain.** The client knows that AI has cut its provider's cost of production; it negotiates how that gain is split, not the daily rate.

Clients demand an outcome they cannot yet measure. One company in five limits its use of AI on cost grounds, and few can measure the value obtained. That measurement is becoming the core of price negotiations.

37 %

of companies report an effect of AI on their EBIT, unchanged over the past year

32 %

have dropped a software purchase they could build with coding agents

60 %

expect to increase their AI investment within the year


Source: McKinsey Global Survey, *The State of AI in 2026*, August 2026 (1,719 respondents, 97 countries); Odyssey analysis.

01 · FROM CAPACITY TO OUTCOMES

Value moves from execution capacity to economic performance, as recorded in the P&L. That shift splits the AI value chain into three segments — frame, execute, measure — and moves value toward the two ends. Execution, where providers have scale, becomes the most contested segment; framing and measurement, where their standing has yet to be established, are where budget and price are decided. What will be left to execute once clients automate part of the production themselves?

EXHIBIT 1

The AI value chain: value moves to the two ends, strategy and measurement, while providers keep their foothold in execution



Source: Odyssey analysis, 2026.

The productivity gain delivered to the client is deducted from the invoice

Earlier waves of digital transformation produced gains that the provider controlled the estimate of, and therefore the split. AI removes that asymmetry: the client observes the gain, quantifies it and claims its share.

The three previous waves lowered the cost of production without eroding margin: *offshore* billed at *onshore* prices, packaged software sold expertise unavailable elsewhere, and cloud billed shared infrastructure by time spent. The client could not always measure.

17 to 25%

operating margin at the large Indian providers, against 17.0% at Accenture and 16.6% at CGI. Within CGI: 24.3% Canada, 15.3% US, 11.2% Europe

THREE RISKS FOR THE PROVIDER

This time AI reaches the provider's own economics and operating model.

- **Substitution of the work.** AI agents absorb part of what used to be billed: development, testing, documentation, maintenance, and also analysis and process redesign.
- **Obsolescence of the role.** Where execution becomes automated, the middle segment of the chain is no longer bought separately: the client runs its own agents or turns to the model vendor. The decision shifts toward framing and measurement, where the provider is not established.
- **Reversed information asymmetry.** The client now measures the gain as accurately as the provider, using the same tools. As long as the billed unit is time, every gain shrinks the base. Absent a gain-sharing clause, the balance of power decides who keeps it.

FOUR CONTRACT VARIABLES

AI moves all four: once time ceases to be the unit of account, the negotiation turns on the outcome produced and on the asset that makes it possible.

- **Volume.** Standardized tasks, which justified hiring junior staff, are automated first: the pyramid narrows at the base.
- **Price.** Clients write the gain into their tenders: a contract worth \$100 million may now be worth \$80 million, according to HCLTech.
- **Quality.** It used to be proven through organization (team, methods, reviews). It is now verified deliverable by deliverable: control becomes a cost line, and the human rework rate replaces utilization as the profitability measure.
- **Time to delivery.** Now that clients know it can be compressed, faster delivery becomes a contractual expectation and ceases to be a pricing argument.

Beyond headcount, the cost base becomes a strategic variable

Changing the pricing unit also shifts the cost base. A provider's economics used to sit in a single line — personnel costs, about 70% of revenue and more than 80% of the cost base. AI replaces it with a composite structure: human work, compute, model licenses, tooling.

70 to 80%

of a provider's cost base sits in personnel costs alone

The pyramid narrows at the base. The tasks that justified hiring juniors are automated first, while demand moves toward profiles the pyramid does not produce: control of agent output, data architecture, operation of agent fleets.

-6,7%

monthly decline in the unit price of inference at constant capacity

Losing its base, a provider loses its school: which profiles to hire, how to train them without junior assignments, how to redeploy staff toward supervision all remain unanswered.

The *offshore* trade-off loses its rationale once production is automated. The hourly cost gap ceases to be the first margin lever; the trade-off now concerns how offshore capacity, automated capacity and local supervision complement one another.

Inference cost enters the P&L without the provider setting its price. Token prices are set by vendors, leaving providers exposed to their volatility, and the model hierarchy is reshuffled with each generation. The result is dependence on the chosen model, which is costly to replace, and exposure to regulatory, geopolitical and energy risk.

Renting or owning compute capacity becomes a capital decision again : dedicated infrastructure, partnership with a neocloud, pay-per-use rental. Providers ran their own data centers in the 2010s before ceding that position to specialists with deeper capital. AI reopens the question on different terms — unit cost of inference, data sovereignty — and this time the decision shapes the sale price.

These trade-offs redefine the operating model. Running an engagement will mean assembling a structured setup — agents, human supervision, compute capacity — and tracking its unit cost and human rework rate, rather than staffing a team and tracking its utilization.

Two-thirds of a provider's revenue sits in lines that AI reshapes or replaces in volume

The service portfolio is unevenly exposed: the risk concentrates where AI transforms both the work and the unit that bills it.


The portfolio is set to move from AI *build* to instrumented *run*. One third of activities are resilient — consulting, managed services, engineering, cybersecurity — because the provider holds a defensible position or an asset there. The remaining two-thirds see their pricing unit change first: that is where the margin of the next three years is decided.

EXHIBIT 2

Eight service lines, three fates: a provider's portfolio tested against AI

TYPICAL REVENUE MIX OF A EUROPEAN PROVIDER

based on the published portfolios of Capgemini and Sopra Steria, among the few that split revenue by service line



Light blue: resilient. Steel blue: reshaped. Midnight: replaced in volume. Two-thirds of revenue sits in the last two — the scope tested by the simulation that follows.

Source: weights based on Capgemini (2025 annual report) and Sopra Steria (2025 results) — among the few providers that publish a service-line split; Accenture reports only consulting and managed services, CGI by geography. Each group's own taxonomy, not directly comparable; fates and units: Odyssey analysis, 2026.

(1): success fee indexed on a contracted saving (baseline defined at framing)

The billed unit and the use of released capacity make the margin, not productivity

The first driver of margin is the measurement protocol on which the price charged to the client is negotiated.

EXHIBIT 3

What becomes of the two-thirds exposed: three P&Ls at three years, at identical productivity, for a 500-person provider

Common assumptions — a 500-person provider, blended onshore-nearshore delivery, \$60m revenue (\$120k per head), of which \$30m in the AI scope. Personnel costs at 70% of revenue, against cost of revenues of 72.3% at EPAM in Q1 2026. A 30% productivity gain on that scope releases \$9m of billable capacity. Inference rises from 1% to 5% of revenue, i.e. \$3.0m. Discounts apply only to the AI scope: -15% reduces revenue by \$4.5m. Only the pricing unit agreed at renewal separates the three scenarios.

In \$m unless stated · Headcount in FTE · EBITDA margin as % of revenue · Horizon: financial year T+3

	Starting point — T&M, before AI	A. T&M maintained — productivity is handed back to the client	B. Unit of work — 15% discount on the AI scope	C. Contractual base + outcome-linked fee + asset licensing
Revenue	60,0	51,0	55,5	58,5
of which AI scope	30,0	21,0	25,5	28,5
Headcount (FTE)	500	425	425	449
Personnel costs	42,0	35,7	35,7	36,2
External purchases	6,0	8,4	8,4	8,7
of which inference	0,6	3,0	3,0	3,3
Other costs	6,6	6,6	6,6	7,0
EBITDA	5,4 (9,0 %)	0,3 (0,6 %)	4,8 (8,6 %)	6,6 (11,3 %)
Main driver	—	Ability to adjust headcount — a favorable assumption	Level of discount on the AI scope: -10% leaves 6.3, -20% brings it to 3.3	\$2.0m reinvested in assets, of which 1.5 capitalized; without capitalization, 5.1 (8.7%)
FROM A (\$0.3M) TO C (\$6.6M): WHERE THE \$6.3M EBITDA GAP COMES FROM				
	+ 0,3	+ 4,5	+ 1,8	= \$6.6m
A. T&M		Shift to a unit of work	Capitalizing the released capacity	C. Base + outcome- linked + licensing

Note: between A and C, productivity effort and inference cost are identical; the \$6.3m EBITDA gap comes from the unit price and the use of the released capacity. Inference cost rises to 5% because the unit price falls (about 6.7% per month at constant capacity) more slowly than volume grows. Source: Odyssey, 2026; Du, M., Tiered Super-Moore's Law, arXiv, March 2026; Pay-i, 2026. Illustrative simulation.

Deflation is offset by changing the pricing unit, the value-sharing rule and its measurement

Margin is decided at signature, not in annual price reviews. Three clauses are inseparable: the billed unit sets the base, the sharing rule splits it, the protocol establishes it.

“As long as the invoice rests on time spent, a provider's productivity effort funds its own erosion.”

What is AI deflation? It describes a simple mechanism: the same deliverable costs less from one year to the next, because the time needed to produce it falls. It owes nothing to competition or to the cycle: written into the unit of account, it applies at every contract renewal as soon as the client knows the gain.

Change the pricing unit. Billed by time, every productivity gain is subtracted from revenue. The unit must rest on the outcome delivered — volume processed, scope operated, asset made available — not on the means deployed.

Define the sharing rule. A contract silent on ownership of the gain awards it to the client: billed by time, it mechanically passes on any acceleration in production. The sharing clause requires three elements settled before the first invoice — an audited baseline, a scope of attribution and a trigger threshold — without which it cannot be enforced.

Commit to the measurement protocol. The share is enforceable only if the outcome is measured in a way that holds up. The provider owns the instrumentation, since execution traces arise in the systems it operates. Attesting the outcome requires an independent third party.

Three rules define a fair split. The base: the gain is measured against an audited baseline, net of costs incurred. The contribution: the share goes to whoever carries the risk and owns the asset. The limits: a tolerance band below which nothing is triggered, and a cap beyond which the loss is no longer borne by the provider alone.

The unit, the sharing rule and the protocol are defined in the contract. They are delivered through the operating model: cost per engagement, management of the rework rate, capitalization of assets.

The portfolio is reshaping around a dozen new lines. Two criteria decide their three-year yield: the exposure of the service to deflation, and the asset the provider keeps once the engagement ends.

EXHIBIT 4

high

low

Proprietary asset

low

high

Deflation risk

Bubble size: demand momentum

DEFENSIBLE POSITION
Asset retained · low deflation

1, 2, 3

4, 5, 6

TWO COMPLEMENTARY VARIABLES, FROM 1 TO 5

1. Data preparation

2. AgentOps

3. Licensed vertical agents

4. Agent governance

5. Legacy modernization

6. Custom agent deployment

Steel blue: competitive intensity. Light blue: closeness to the provider's current business. Five gradations.

Source: Gartner, Hype Cycle for Agentic AI 2026; Deloitte, 2026; Odyssey analysis. Seven other engines were assessed on the same grid.

On the horizontal axis, the exposure of the service to deflation; on the vertical axis, the proprietary asset the provider retains. Two complementary variables refine the reading: the stronger the competitive intensity, the more entrants contest the line; the closer the line to today's business, the cheaper it is to build the skills.

The portfolio shifts from AI *build* to instrumented *run*, and the window is short

The activities that fund the transition. Custom agents and legacy modernization carry demand today, but rarely leave a reusable asset, and AI automates the work they bill. They remain worth selling while they fund the shift, without the target model resting on them.

The activities that will generate the revenue. Data preparation, AgentOps and licensed vertical agents are billed by unit of work, subscription or license, and leave the provider a reusable asset. They will take over as the revenue base, with a degree of recurrence, once the build phase ends.

Rotating the portfolio is not a catalog extension : it is a shift from build to instrumented operations, and it has a deadline. Building AI systems is a finite wave — once the processes of large accounts, then of mid-market and smaller firms, have been redesigned, demand will contract.

What remains is running those systems — supervision, measurement, continuous improvement — and the assets that make it possible. Positions in operations are taken now, funded by build revenue.

“The wave of AI system building will end. What then?”

The emerging revenue model combines a fixed-fee base, an outcome-linked fee and asset licensing


Six revenue models coexist, to be judged on what is invoiced, what they require and the protection they offer against deflation.

- 1 **Time and materials with a productivity commitment** — *bills time*, while promising to use less of it. It is the most common model, and the one bound to fail: it writes deflation into the contract.
- 2 **Fixed price per deliverable** — *bills a scope*. It protects the gain while the scope holds, but is renegotiated downward as soon as the client sees the time actually spent.
- 3 **Shared unit of work** — *bills a countable object* : a document processed, a ticket resolved. It is the first model that leaves the productivity gain with the provider, on condition that it knows its cost per unit.
- 4 **Outcome-based pricing** — *bills the achievement of a business objective* : a bonus at a threshold, a percentage of the measured gain, or a premium rate while performance holds. Without a signed measurement protocol, attribution is settled by the balance of power.
- 5 **Value sharing on a performance indicator** — *indexed on a line of the P&L*. This model offers providers the highest potential, and is also the most demanding: audited baseline, counterfactual and third-party attestation.
- 6 **Asset licensing** — *bills the use of an asset built by the provider*. It is the only revenue independent of both time spent and the client's outcome, and the only one valued at a software multiple.

Taken alone, none of these models covers everything a provider needs — funding the structure, earning a share of the gain, monetizing assets. Layering them is what makes the difference: the choice is no longer a single revenue model, but a combination of layers built contract by contract.

Layered revenue: a base that funds, a fee that captures, licensing that monetizes

The difficulty is not designing the combination but invoicing it: three layers, a single invoice, and financial control able to track them separately.



The market already pays 2x more for assets and recurring revenue

A provider is not meant to become a software vendor. Within digital services themselves, activities that own proprietary data and generate recurring revenue trade at about 16 times EBITDA, against 10 to 11 times for those selling capacity.

EXHIBIT 5

M&A multiples paid, by sub-sector

Medians, majority transactions 2015 – June 2026		
Sub-sector	EV / EBITDA	EV / revenue
Data and analytics	16,3x	2,2x
Cybersecurity	16,1x	2,1x
Digital engineering	12,1x	1,4x
Software development	11,3x	1,2x
Staffing	11,0x	1,0x
Systems integration	10,7x	1,1x
Managed services (MSP)	10,7x	1,1x
IT consulting	9,8x	1,2x
Value-added reselling	7,8x	0,5x
Median, digital services	10,4x	1,1x
Software, profitable private transactions	22,4x	—

Source: Aventis Advisors and Mergermarket, database of majority transactions January 2015 – June 2026 (n = 520 with a disclosed EBITDA multiple); software and SaaS multiples, Aventis Advisors, August 2026. Odyssey analysis.

The gap separates models, not trades: an asset owned, recurring revenue, a scarce skill — the three layers described above. A provider that shifts toward licensing and subscription climbs that ladder without changing trade.

The evidence is not proof. These medians describe sub-sectors, not models compared on a like-for-like basis: they give the order of magnitude of what is at stake, not what a given provider would be worth after the shift.

2x

the EBITDA multiple gap within digital services: 16.3x for data and analytics, 7.8x for reselling

Aventis Advisors, Mergermarket

Four groups of entrants, and no barrier that scale or relationships can hold

The addressable market is growing while the providers' position within it weakens: AI transformation opens a wider scope that they no longer cover alone.

Model vendors are moving down into consulting and deployment. The services entities they set up in 2026 buy a presence at chief-executive level: framing the use case, then deploying the agent, gives them command of the chain from decision to execution. Audit and consulting firms are acquiring technology capabilities, by buying applied-AI players and by building proprietary assets.

AI-native entrants address, with augmented platforms and small teams, work only a services provider used to be able to run.

Providers have historically been protected by relationships rather than by structure. Their four supports never barred entry to a determined competitor. Two withstand AI — access to clients and knowledge of systems in production. Two give way: mobilizing engineers, once production is delegated to agents; economies of scale, once marginal cost depends on inference rather than headcount.

Software vendors are moving the other way. Services become their growth engine and give them access to the processes their agents need. Software-as-a-Service is matched by Service-as-a-Software. Providers stand at the crossing point.

£740m

paid by Accenture to acquire
Faculty

05 · THE RESHAPING OF THE MARKET

Palantir set the template: a software company that deploys its own engineers inside the client, capturing the framing a provider would have billed. Mistral AI follows the same path in Europe.

The proprietary asset decides positions. Execution corpus, proven agents and an accepted measurement protocol are slow to reproduce: whoever owns them keeps control of their price.

These assets must remain portable as models evolve. Exclusive alignment with one model vendor buys privileged access and turns the provider into a distribution channel. Full neutrality preserves independence without creating an asset. The middle course is a multi-model strategy: a shortlist of vendors selected on performance and cost, an orchestration layer that makes switching possible, and a decision revisited each cycle — buyers who reassessed models every six months now do so monthly¹.


> €6.5bn

raised by Mistral AI in equity and debt since 2023

Mistral AI; Reuters

EXHIBIT 6

Model vendors move down into services, consulting moves up into technology (2026)



1. Digital Applied, Frontier Model Release Velocity Index, Q2 2026. 2. OpenAI is a shareholder of Thrive Holdings. Source: OpenAI, May 11, 2026; TechCrunch, May 4 and August 12, 2026; Aventis Advisors, May 2026; Blackstone, December 16, 2024.

The defensible position runs through an alliance the provider leads

The architecture of the alliance. Building strategic framing and financial measurement skills in-house would take years; acquiring them requires a shareholder willing to fund deals while the ratios weaken. Alliances are the workable path, provided they rest on a contractual commitment.

- **The operator, the provider, holds the contract and remains the client's single point of contact.** To do so, it must lead the alliance. Otherwise the architect or the certifier takes control of the client relationship.
- **The architect, a strategy consulting partner, brings the framing and the value targets,** white-labeled or co-signed. Three conditions: independence from model vendors, financial expertise, and no competing execution business.
- **The certifier, an independent third party, attests the outcome when measurement sets the price** — no one certifies their own work.

EXHIBIT 7

Execution is held alone; framing and measurement need a third party

1. Frame	2. Execute	3. Measure
SKILL 	SKILL 	SKILL 
Management, organization, finance — to be built	In place — this is the trade; to be strengthened on some AI scopes	Financial control per unit of work; ROI measurement
STANDING 	STANDING 	STANDING 
Held by consultancies, moving down into technology	Established with technical and business leaders	Denied to whoever measures their own work
ECONOMIC UNIT Fixed fee	ECONOMIC UNIT Unit of work, subscription, license	ECONOMIC UNIT Instrumentation subscription; attestation by a third party
Out of reach alone: alliance	To defend and instrument	Instrument yes, attest no

Enforceability of an outcome-based contract rests on the measurement protocol and on the form of the alliance

Providers have historically had limited practice of alliances. Beyond subcontracting, structured partnerships remain rare — at the very moment the model requires them. The gain-sharing contracts of the 2000s failed because attribution proved too difficult.

A system in production now traces every unit produced, every rework and every outcome. Six clauses are enough to make the share enforceable — without them, the provider carries the risk without controlling its measurement.

EXHIBIT 8

The six clauses without which an outcome-based contract stays one-sided

1	2	3	4	5	6
Audited baseline	Counterfactual	Scope of attribution	Sharing formula	Inference cost review	Asset ownership and arbitration

Source: Odyssey analysis, 2026.

The protocol says how the outcome is measured. The contractual form says who answers for it. Three forms of alliance are possible. The choice shapes liability toward the client, the strength of the attestation and the complexity of the split.

— **The bilateral contract held by a joint entity.** The provider and its partners form a project company that contracts alone: the client keeps a single point of contact and the split is settled among the shareholders. In return, the entity is justified for a scope and a duration, rarely for a single engagement.

— **The multi-party contract.** Each party contracts with the client on its own scope: maximum legal clarity and independent attestation — the most robust form when measurement sets the price. The client, however, takes back the coordination it wanted to delegate.

— **The fully bilateral contract.** The provider contracts alone and subcontracts framing and attestation. This form is easier to sell, but whoever executes becomes whoever attests, and the measurement loses its force.

Outcome

The commitment changes in nature. Buying shifts from an obligation of means to an obligation of result, rarely written down and increasingly demanded.

Margin is hollowed out just as the shareholder prepares to exit

The main friction is a mismatch of horizons: the transition runs two to three years, the capital financing it rarely does. For two to three years, revenue declines at constant activity, working capital rises, and the cost base carries both the legacy structure and the new development.

The difficulty is timing rather than willingness : two to three years of transition sit badly under an EBITDA covenant, a planned exit or an annual bonus. The instruments are what must adapt: a transition measure in incentive plans, a covenant that neutralizes capitalized development, an exit horizon that accommodates the shift.

Shareholders have three options. Extend the holding period; recapitalize with a patient partner; or sell before the ratios weaken, leaving the value of the transition to the buyer.

The transition is funded first from its own resources. Four levers provide for it: (1) the capacity released by AI, directed toward building assets rather than handed back to the client; (2) the contract, where the ownership clause allows reuse of what the client funded; (3) realignment of variable pay, today indexed on revenue and utilization; (4) capitalization of eligible development, which smooths the apparent decline in ratios.

2 to 3

years of weaker ratios before margin rebuilds

Odyssey simulation


€250m

cap on Blackstone's recapitalization of Sia Partners

Blackstone, 2024

EXHIBIT 9

The value of the transition is harvested by the shareholders still present beyond year 3



Source: Odyssey analysis, 2026.

Odyssey © 2026 | After the Billable Day

Capital now prices the execution corpus, an asset most providers cannot protect

In 2026, capital transactions show that investors and buyers are no longer after headcount but after four assets that a services company's accounts do not record — vertical knowledge, methodological tooling, execution corpus, and access to a client portfolio.

The execution corpus gathers execution traces, the exceptions arbitrated and the reasons behind them: the material an AI model needs to work under real operating conditions.

In most providers none of these assets is inventoried, and so none is protected or monetized. The Thrive Holdings transaction, with OpenAI as a shareholder, shows what is at stake: about \$2 billion raised in August 2026 to acquire providers, that is, instrumented access to real processes and business data far more than order books.

Providers need an intellectual property, infrastructure, and balance-sheet strategy.

The stake is sovereignty over data, not over models. On the layer where technology meets real processes, the advantage lies with whoever holds the domain knowledge — healthcare, industry, energy, financial services, and the public sector. In each, data residency and sector rules restrict who may operate the systems. Providers are the natural carriers of that constraint, and it is one of the few advantages an entrant cannot buy.

It shows up on the balance sheet. Inventory the assets, secure ownership and training rights in contracts, capitalize what can be capitalized, choose between licensing and sale. An asset absent from the balance sheet has no value in a transaction.

Infrastructure sets the limit. Running regulated workloads on foreign compute caps the autonomy obtained, and the value of the asset. Investing in one's own capacity remains a differentiation strategy reserved for a few consortiums.

“An asset absent from the balance sheet has no value in a transaction.”

By 2030–2032, leadership will rest on six measures that no dashboard tracks today


The reshaping of the market will produce new leaders. Incumbent providers keep three advantages — their position with client decision-makers, their knowledge of systems in production, and a renewal calendar that regularly opens the chance to change the unit. That is a head start, not a protection: what will separate players is not their size but the timing of the shift.

“Providers that have not completed the rotation by 2030 will not lead the shift — they will follow it.”

Six measures, all of them management decisions, will determine tomorrow's leaders. They are within reach of every provider, so the difference will come from when they are made and how fast. The software industry's move from license to subscription showed as much: the vendors that shifted first opened a valuation gap the others never closed.

EXHIBIT 10

Six measures place a provider, and none appears in today's dashboards



HOW TO READ

- Typical profile of a mid-sized provider, 2026
- Threshold above which the position is defensible

The gap is closed by decisions of the executive committee: change the unit of one contract, open a cost per engagement, sign a measurement protocol.

Scale from 0 to 5. Rework rate: share of agent output corrected by a human before delivery. To be applied to your own scope: the exercise places the company before it places the market.

Five of the eight measures that steer the new model do not exist in providers today

A provider's dashboard measures what no longer makes the margin: the utilization of its headcount.

Eight measures take over from the average daily rate, the utilization rate and margin per client, across four areas: revenue, cost, production, and contract. Five have to be built from scratch.

These measures are not for reporting but for decisions. Fully loaded cost per engagement decides whether a contract continues or stops; the rework rate decides when an agent is recalibrated; the gap between promised and attested outcome decides when to renegotiate before renegotiation is imposed. Each is read at engagement level, monthly, not at a quarterly management committee.

EXHIBIT 11

The dashboard for revenue detached from time: eight measures, five of them to be built

REVENUE	COST	PRODUCTION	CONTRACT
Share of revenue billed outside time spent Unit of work, subscription, outcome and licensing, as a share of the total replaces average daily rate	Full cost per engagement and per task completed Personnel, subcontracting, inference, supervision, and rework replaces utilization rate	Human rework rate Share of agent output corrected before delivery; profitable below 30% no equivalent today	Contracts under a signed measurement protocol Share of revenue covered: baseline, counterfactual, and sharing no equivalent today
Licensing revenue and capitalized assets Recurring revenue from licensed assets; assets carried on the balance sheet replaces margin per client	Inference cost per unit of work Tokens and compute consumed per unit delivered, by model no equivalent today	Asset reuse rate Share of engagements delivered from existing assets no equivalent today	Gap between promised and attested outcome Difference, by contract, between the commitment and the third party's measurement no equivalent today

Source: Odyssey analysis, 2026.

Four decisions the executive committee should take within 90 days

These four decisions require a trade-off within the executive committee, with no investment budget. Each produces a dated deliverable.

1 Measure the share of revenue detached from time spent

By client and by service line, over the last twelve months. Owner: finance. Deliverable: a single dated figure, presented to the executive committee.

2 Establish a cost per engagement, inference included

On three representative pilot engagements, covering personnel, subcontracting, inference, and supervision. Owner: financial control. Deliverable: a margin per engagement that holds up internally.

3 Change the pricing unit on one contract

Pick a renewal in the next six months, replace the unit with a unit of work processed or a subscription per scope operated, and sign the measurement protocol before the first invoice. Owner: the chief executive. Contributors: *delivery* to define the unit, financial control to compute the cost, sales to negotiate. Deliverable: a signed contract.

4 Obtain a five-year mandate from the shareholder

Present the trajectory, the margin trough of the first two years and the four levers that fund the transition. Owner: the chief executive. Deliverable: a written agreement on the horizon over which performance is judged.

Secure the mandate before starting the other three. Those three decisions weaken the ratios before they restore them. Without the shareholder's explicit agreement on that horizon — and without adjusting executive incentives beyond the single year — they will not be carried through. Owner: the chief executive and the board. Deliverable: a mandate letter setting the horizon and the measures used to track the transition.

When the outcome cannot be measured: cybersecurity contracts on exposure

In cybersecurity, billing has already detached from time spent — subscription, event, scope covered. The outcome itself remains impossible to contract for: no one guarantees the absence of an incident.

Three forces set this segment apart. Regulation creates demand that is largely insensitive to the budget cycle — SEC incident disclosure and CMMC in the United States, the NIS Regulations in the United Kingdom, NIS 2 and DORA in the European Union; the skills shortage holds prices up; data residency rules narrow the field of eligible operators.

Deflation is present nonetheless. It bears on the layer that justified volume — first-level analysis, triage, false positives. A SOC equipped with agents covers the same scope with fewer analysts, and a subscription sized on headcount is renegotiated.

+11%

growth in Q2 2026 at Sopra Steria in defense, space and security

Sopra Steria, H1 2026

EXHIBIT 12

The framework tested against cybersecurity: 3 components in place, 2 to be built

The framework	In cybersecurity	Adaptation
Pricing unit	Often outside time: subscription per asset covered, per qualified event, per scope operated	In place — but AI absorbs the analyst headcount that sized the subscription
Sharing rule	The gain is not shared: it is reinvested in coverage or returned to the client through a lower price per protected asset	To be replaced by wider scope at constant price and by service-level guarantees
Measurement and attestation	No one guarantees the absence of an incident and the counterfactual cannot be observed; the third party already exists — audit, certification, regulator	Measure exposure rather than outcome: time to detect and remediate, estate coverage, false positives, cost per event
Growth engines	Security of agents and models, non-human identities, the AI software supply chain, compliance (SEC, CMMC, NIS 2, DORA)	A new field where cyber providers have strong standing and model vendors are not yet present
Revenue architecture	Coverage subscription + variable fee on events handled + licensing of detection content (rules, response playbooks)	The third layer already exists, as intellectual property still barely inventoried and barely monetized

Source: SEC cybersecurity disclosure rules (Item 1.05); CMMC; UK NIS Regulations 2018; NIS 2, DORA and the Cyber Resilience Act (EU texts 2022/2555, 2022/2554, 2024/2847); Sopra Steria, 2025 annual results; Odyssey analysis.

I APPENDIX · FREQUENTLY ASKED QUESTIONS

01 Will wider demand offset the decline in revenue?

Our reading. No, because the objection is about volume and the thesis is about unit price. New work billed by the day deflates at the same rate as the old: widening the base does not protect margin while the pricing unit remains time.

02 Why would outcome-based pricing succeed now, having failed in the 2000s?

Our reading. Those contracts stumbled on attributing the gain, for lack of a way to measure it. An agentic system in production traces every unit processed and every rework, and a control scope can run alongside the automated one. The condition is contractual: a protocol signed before the first invoice.

03 Does the falling cost of inference make changing the revenue model secondary?

Our reading. The unit price of inference falls, but the cost per completed task rises. One company in five already reports AI use constrained by its costs, tokens included¹. More to the point, that cost is set by model vendors at their discretion: it calls for a review clause and for a choice among vendors.

04 What is the effect of this shift on the balance sheet?

Our reading. A provider's balance sheet is light: goodwill, receivables, almost no assets produced. Layered revenue thickens it — agents and reference libraries are capitalized, and buyers pay for them². Working capital is pulled both ways: the outcome-linked fee lengthens the cycle, licensing shortens it. The trade-off is explicit: better quality of revenue, higher capital intensity, return on capital employed to watch.

05 How can the transition be funded without shareholder money?

Our reading. It is funded by operations themselves. Through the released capacity, directed toward assets; through the ownership clause, which allows reuse of what the client funded; through realignment of variable pay; through capitalization of eligible development.

06 What happens if the promised performance is not delivered?

Our reading. That is what the measurement protocol must settle before signature: trigger threshold, review clause in case of a gap, risk sharing, arbitration of disagreements. An outcome commitment without a review clause transfers to the provider a risk it does not control alone.

1. McKinsey Global Survey, *The state of AI in 2026: On the road to ROI, August 2026* (1,719 respondents, 97 countries): 20% of respondents report AI use constrained by operating costs, tokens included; among the highest performers, the constraint affects 18% of coding-agent use, against 6% elsewhere.

SOURCES AND METHOD

Principal sources

Market and outlook

Accenture, Cognizant and Endava, fiscal 2026 results and guidance; Gartner, worldwide IT spending forecast, 2026; McKinsey Global Survey, The State of AI in 2026, August 2026 (1,719 respondents, 97 countries).

Analysts

Gartner, Hype Cycle for Agentic AI 2026 and 2026–2027 forecasts; Deloitte, agentic AI survey 2026; Forrester, Predictions 2026; Digital Applied, Frontier Model Release Velocity Index, Q2 2026.

Company disclosures

EPAM (Form 10-Q, Q1 2026); Accenture (Form 10-Q, Q2 and Q3 fiscal 2026); CGI (Form 6-K, fiscal 2026); Endava (fiscal 2026); Kyndryl (fiscal 2026); Booz Allen Hamilton (Form 10-K, fiscal 2026); Cognizant (2026); HCLTech, Infosys, Wipro and Coforge (FY26–FY27); Capgemini (2025 Universal Registration Document and Integrated Annual Report); Sopra Steria (2025 results, Q1 and H1 2026); Atos (Q1 and H1 2026); annual reports of TCS, Infosys, Capgemini and Sopra Steria (2010–2025).

Competitive reshaping

OpenAI, announcement of its deployment business, May 11, 2026; TechCrunch, articles of May 4 and August 12, 2026; Aventis Advisors, AI services M&A study, May 2026; Blackstone, investment note, December 16, 2024.

Academic work and cost measurement

Du, M., Tiered Super-Moore's Law, arXiv, March 2026; Pay-i, 2026.

Regulation and public texts

SEC cybersecurity disclosure rules (Item 1.05 of Form 8-K); Cybersecurity Maturity Model Certification (CMMC); UK Network and Information Systems Regulations 2018; Directive (EU) 2022/2555 (NIS 2), Regulation (EU) 2022/2554 (DORA), Regulation (EU) 2024/2847 (Cyber Resilience Act); European Commission, proposed Cloud and AI Development Act, June 3, 2026; Odyssey, Europe's AI strategic autonomy, 2025.

Methodological note

This report draws on Odyssey's sector expertise with digital services providers and vertical software vendors in Europe, and on the engagements conducted with their executives and shareholders. It applies Odyssey's AI strategy framework together with its proprietary methods and its TEV-AI value measurement framework.

The thresholds presented — rework rate, discount levels, timing of the shift — are orders of magnitude drawn from that work, not statistical averages. The simulation in Exhibit 3 is illustrative: it rests on the assumptions stated and is not a forecast. The rankings in Exhibit 4 reflect an Odyssey scoring whose criteria are set out on the corresponding page.

Disclaimer. This report is an independent Odyssey publication: it was neither commissioned nor funded by any company, investor or institution. It rests on public information as of October 2026, which Odyssey has not independently verified and whose accuracy and completeness it does not guarantee. The analyses of the companies mentioned, some of them listed, express the author's opinion only; forward-looking elements are working assumptions, not forecasts. This document is distributed for general information and constitutes neither an investment recommendation nor legal or tax advice. Odyssey undertakes no obligation to update it and accepts no liability for any use made of it.

About the author



Axel Tombereau

Managing Partner, Odyssey

Axel Tombereau advises executive committees and private equity firms on the value created by artificial intelligence. He founded Odyssey in 2020 and has since worked with more than fifty technology companies across Europe.

With over ten years of experience with digital services providers, notably at Alten, Segula Technologies, and Atos, he has led some forty strategic projects and mergers and acquisitions across more than twenty countries, for a combined value of about €4 billion. He now advises the sector's executives and shareholders.

A graduate of ESSEC and INSEAD, he has taught corporate finance at Mines-ParisTech and financial markets at HEC.

axel.tombereau@odyssey.tech

About Odyssey

Odyssey designs companies' artificial intelligence strategies and demonstrates their returns. Our team works with executive committees and investors through two practices: **AI StratOps**, which defines tailored artificial intelligence strategies and carries them through to an executable roadmap, and **AI FinOps**, which measures the return on AI investment to the standard boards, shareholders, and acquirers expect, fully loaded costs included, using the proprietary TEV-AI™ methodology. Odyssey sells no execution services.

OUR TWO PRACTICES

AI StratOps

Tailored AI strategies, prioritized use cases, operating model, governance, and roadmap

AI FinOps

AI value measured by use case, in the P&L, on the balance sheet, and in valuation, hidden costs included

Two ways to engage. Advisory — strategic framing at executive-committee level, ahead of any deployment decision. Skills — ten diagnostics and reports, available separately on a defined scope.

50+

strategic
engagements
delivered

10

diagnostics and
reports in the catalog

4

dimensions of value in
The Odyssey AI ROI
Framework

2020

year Odyssey was
founded

Odyssey Tech SAS — 128 rue La Boétie, 75008 Paris — odyssey.tech · contact@odyssey.tech

To cite this document: Tombereau, A. (2026), "After the Billable Day", Odyssey, Paris — ODY-NR-2026-003 · v6.0