

AI-Enabled Professional-Services Roll-Ups: Buying Capability without Capitalising Hype

A Transaction Framework for Capability, Data, Talent, Revenue Quality and Sustainable AI Economics

Authored by Chennakeshav Adya, Independent Researcher

September 2026

Abstract

Professional-services consolidators increasingly describe acquisitions through artificial intelligence, proprietary data, specialist talent and cross-selling. These descriptions can identify genuine capability. They can also combine several unlike assets: third-party software licences, reusable delivery methods, customer-specific configurations, unstructured documents, individual expertise and early-stage automation pilots. A buyer that capitalises the label rather than the evidence can pay for benefits that remain dependent on a vendor, a founder, a single client or an untested integration plan.

This paper develops a transaction framework for AI-enabled professional-services roll-ups. It defines the acquisition thesis, maps capability and evidence, reconstructs recurring revenue, measures talent and client dependency, and bridges claimed AI productivity to sustainable cash economics. It also examines data rights, model-provider concentration, cybersecurity, competition review, purchase-price allocation, financing and post-close governance. The framework is designed for advisory, accounting, legal, engineering, staffing, consulting, managed-service and other knowledge-intensive businesses in which human judgement remains part of delivery.

Five original figures and five decision tables turn the analysis into an acquisition scorecard. A worked example shows how apparent EBITDA growth can change after client givebacks, implementation cost, retained human review, model expenditure and capacity release are tested. Every amount, percentage, multiple, score and scenario in the example is a hypothetical analytical assumption. The paper does not provide a valuation, forecast, accounting conclusion, legal opinion or recommendation for any company or transaction.

JEL Classification: G34, G32, L84, M41, O33

Keywords: professional services, mergers and acquisitions, roll-up, artificial intelligence, revenue quality, talent dependency, proprietary data, sustainable EBITDA, purchase price allocation

1. Define the acquisition thesis before valuing the AI label

A professional-services roll-up can seek geographic reach, sector credentials, specialist expertise, client access, recurring managed work, delivery capacity, proprietary workflow assets or a platform for further acquisition. The transaction mandate should state which of these mechanisms is expected to create value and which evidence could disprove it. A buyer seeking a stable cash-generating platform needs a different diligence emphasis from a buyer seeking to transform human-intensive work into a technology-supported service.

The thesis should identify the unit of acquired capability. A named AI product may be a third-party model connected to a document repository. A claimed platform may be a set of scripts used by two teams. A valuable capability may exist without a software product, through curated data, a repeatable method, experienced operators and disciplined quality control. The analysis

should therefore begin with the customer problem, the delivered outcome, the workflow and the evidence of repeatability.

The transaction perimeter should map acquired entities, client contracts, employees, contractors, intellectual property, data, software licences, cloud commitments, subcontractors and shared group services. The financial perimeter should reconcile statutory accounts, project reporting, time records, invoices, receivables, deferred income and cash receipts. The operating perimeter should show where work is sold, designed, performed, reviewed and accepted.

The investment committee should receive falsifiable questions. Does the capability improve price, win rate, delivery time, quality or capacity? Does it remain usable after a change of control? Can it operate across several clients without breaching rights or rebuilding every workflow? Are the people who understand the system likely to remain? Can the buyer fund the integration and governance needed to realise the case? Each answer should connect to price, structure, financing, protection, integration or the decision to stop.

The thesis should specify the counterfactual. The buyer may be able to hire a team, license a comparable tool, build the workflow internally or form a commercial alliance. Acquisition is attractive when control, speed, scarcity, client access or integration value exceeds the cost and risk of those alternatives. A build-buy-partner comparison should use the same capability definition, required investment, time to production, client adoption and risk standard. This prevents the acquisition process from treating scarcity created by marketing as economic scarcity.

Board approval should identify the period over which value is expected to emerge. Immediate value may come from existing client cash flows. Medium-term value can depend on integrating sales, delivery or data. Longer-term value may require productisation and portfolio learning. Separating these horizons clarifies the appropriate source of capital, consideration structure, management incentives and monitoring cadence. It also prevents a long-term option from being used to justify a near-term earnings adjustment.

2. Build a capability taxonomy that separates evidence from branding

AI capability in professional services can be organised into six layers: client problem definition, workflow knowledge, data and rights, models and tools, delivery operations, and commercial distribution. The layers are complementary. A technically strong model can produce limited value when the firm lacks usable data, client consent, integration capacity or a route to market. A strong distribution network can monetise a modest tool when it solves a frequent and costly problem.

The taxonomy should identify whether each component is owned, licensed, customer-provided, employee-held or publicly available. It should state the jurisdictions, contracts and systems that govern use. A method documented in manuals and embedded in training is more transferable than undocumented founder knowledge. A reusable connector with tested controls is different from a demonstration assembled for a transaction process. A governed data corpus with traceable rights is different from documents that clients allowed the target to access for one engagement.


Capability maturity should follow observed production use. Useful stages include concept, prototype, controlled pilot, limited production, repeatable production and portfolio-scale operation. Evidence can include deployment logs, version records, incident history, client acceptance, realised hours, margin cohorts and renewal behaviour. Management presentations and

product demonstrations can explain the proposition. They do not establish production economics on their own.

The capability map should be performed at engagement level before aggregation. A target may use the same brand for several workflows with different data rights, model dependencies and human-review requirements. One workflow may be repeatable across clients while another remains customised. Aggregating the two can overstate transferability and conceal the delivery cost of exceptions. The register should therefore record each material workflow, client population, production version, accountable owner and evidence period.

Capability should also be assessed against client alternatives. A workflow can be operationally effective while remaining easy for a client, software vendor or competitor to reproduce. Defensibility may arise from accumulated outcome data, integration into client processes, regulatory accreditation, specialist judgement, a trusted relationship or distribution. Each claimed barrier should be tested through contracts, observed win-loss evidence, client interviews and replacement analysis. A technical feature without a durable adoption mechanism may support revenue without supporting a valuation premium.

Figure 1. Capability map for an AI-enabled professional-services acquisition



Durable value requires evidence across the connected commercial, operational and control layers.

Table 1. Capability evidence matrix

Capability	Stronger evidence	Weak signal	Transaction question
workflow	documented process, controls and training	founder demonstration	can another team reproduce it?
data	lineage, rights, quality and permitted uses	volume without provenance	can the buyer retain and reuse it?
technology	production logs, code rights and tests	prototype or vendor label	does it work at required scale?
people	role map, retention and succession	generic headcount	who carries irreplaceable knowledge?
distribution	cohort win, price and renewal evidence	pipeline narrative	can capability reach paying clients?

The matrix links each claimed asset to transferability and realised economics.

3. Establish evidence lineage from source to investment decision

Every material claim should retain a path to source evidence. A controlled register can assign stable identifiers to clients, engagements, contracts, statements of work, invoices, employees, code repositories, models, data sets, deployments and incidents. The register should record owner, system, period, extraction date, transformation rule, reviewer and unresolved exception. This allows management, advisers and lenders to reproduce the calculations that support value.

AI can accelerate document classification, clause extraction, code review, invoice matching and anomaly detection. Each output remains an analytical lead until the source and commercial meaning are reviewed. A model may identify a termination clause while missing an amendment. It may group companies with similar names while failing to identify separate budgets. It may detect reusable code while overlooking a licence that prevents transfer.

The decision record should distinguish sourced fact, calculated result, professional judgement and hypothetical scenario. That separation is especially important when management has used an AI-generated classification or forecast. NIST's AI Risk Management Framework and Generative AI Profile emphasise governance, measurement and management of risks across the lifecycle. The practical transaction implication is a traceable chain from claim to evidence, reviewer, limitation and decision.

Population completeness should be proven before samples are interpreted. Contract repositories can omit side letters, local work orders or inherited agreements. Code repositories can omit production scripts maintained by individuals. Customer reports can group several legal buyers. The diligence team should reconcile record counts and values across independent systems, investigate breaks and document residual coverage. A precise analysis of an incomplete population can produce a confident but misleading conclusion.

Evidence access should survive closing. Data rooms are temporary and management teams change. The buyer should preserve the final register, source identifiers, calculation logic, approvals and known limitations in a controlled repository. This package becomes the baseline for purchase-price allocation, integration, covenant reporting and post-close back-testing. It also allows later teams to distinguish an original assumption from a result observed after acquisition.

4. Test whether tools are owned, controlled and transferable

Technology diligence should inventory source code, models, prompts, connectors, data pipelines, deployment scripts, third-party libraries, cloud services and vendor contracts. Ownership records should cover employee and contractor assignments, acquired code, open-source obligations, development partners and client-funded work. Repositories should be reconciled to production deployments and to the people who maintain them.

Third-party models can provide powerful capability without creating a proprietary model asset. The economic question is whether the target owns a differentiated workflow, data advantage, control system, integration or distribution position around the model. Vendor terms should be reviewed for assignment, data use, output rights, service changes, price, availability, security and termination. Technical tests should assess whether another provider can replace the current dependency and at what cost and delay.

The buyer should also identify client-specific configurations. Workflows built inside a client's environment may not transfer. Prompts and templates may incorporate confidential materials. A client may own deliverables or improvements under the engagement agreement. The acquisition case should value only the capability that the buyer can lawfully control and economically deploy after closing.

5. Reconstruct recurring revenue by client and engagement cohort

Professional-services revenue can appear recurring because the same client buys work each year. Contractual commitment, renewal behaviour and project continuity remain distinct. The cohort should identify customer group, contracting entity, engagement, service, start date, term, notice, pricing, delivery team, gross margin, invoice, collection and renewal status. It should separate retainers, managed services, projects, implementation, subscriptions, pass-through and contingent fees.

Revenue measures should use consistent populations. Gross retention should exclude new clients and expansion. Net retention should show price, volume, cross-sell and contraction separately. Backlog should distinguish executed commitments from proposals, framework ceilings, management estimates and work dependent on future approval. Pipeline should be reconciled to historical stage conversion, cycle time, price and delivery capacity.

The cohort should identify whether AI changes the commercial model. An automation capability may support a premium outcome-based fee, protect a fixed fee, reduce billable hours or trigger a contractual productivity giveback. Revenue cannot be credited twice through both retained hours and full productivity savings. Each use case should connect to the relevant engagement and pricing clause.

Renewal analysis should follow the actual decision calendar. Professional-services clients can approve work through annual budgets, panels, tenders, project gates or rolling statements of work. A long relationship may still depend on a near-term rebid. The cohort should record notice dates, procurement events, budget owners, competing providers, satisfaction, unresolved disputes and evidence of future need. Management confidence can then be compared with observable decision progress.

Quality of revenue also includes the cost of serving it. Two engagements with equal revenue can have different economics because of senior oversight, travel, subcontracting, data preparation, custom development, rework or collection delay. The cohort should carry gross profit, contribution, working capital and dedicated capability. This allows the buyer to identify revenue that supports the platform and revenue that consumes scarce people or technology without an adequate return.

Table 2. Recurring-revenue cohort for professional services

Field	Evidence	Test	Deal relevance
commitment	executed contract and work order	termination, scope and minimum	forecast eligibility
delivery	time, workflow and acceptance records	completion and quality	earned economics
billing	invoice and ledger	cut-off, credits and disputes	working capital
collection	receipt and allocation	ageing and deductions	cash conversion
renewal	client budget, notice and competition	decision path and alternatives	retention scenario

Commitment, delivery, billing and collection require separate evidence.

6. Measure client concentration through economic dependency

Top-client percentages are a useful starting point. They can understate dependency when several contracts share an ultimate parent, executive sponsor, procurement framework, sector shock or technology platform. The analysis should include revenue, gross margin, receivables, unbilled work, dedicated staff, references, data access and cross-sell pipeline. A modest client may remain critical when its logo enables entry into a regulated market.

Relationship ownership should be mapped across partners, account leaders, delivery managers and subject specialists. The buyer should test whether client trust belongs to the firm, a team or one individual. Reference calls and client evidence should cover service criticality, satisfaction, switching cost, pricing, automation expectations, data restrictions and future procurement.

The model should run single-client and correlated-cluster losses. It should estimate notice receipts, contribution margin, avoidable labour, retained technology cost, receivable recovery, replacement-sales timing and reputational effects. A concentration statistic becomes decision-useful when translated into cash, covenant and integration consequences.

Concentration should be assessed over time. A stable top-client percentage can conceal churn when one large client replaces another. A declining percentage can reflect rapid low-quality growth elsewhere. The buyer should review monthly or quarterly cohorts, new-client maturation, expansion, contraction and loss. The analysis should explain whether concentration is falling through deliberate diversification, acquisition accounting or temporary project revenue.

Client dependency may also constrain integration. A material customer may require consent before personnel, systems, locations, subcontractors or models change. Cross-selling can trigger conflict, independence or procurement concerns. The integration plan should therefore map each major client's contractual and commercial change path. Expected synergies should enter the model only when the required client action, timing and accountable relationship owner are

identified.

7. Analyse recurring revenue quality alongside delivery mix

A target may combine recurring managed work with discretionary advisory projects, implementation, recruitment, licences and subcontracted delivery. Each stream has different duration, margin, working-capital and capacity characteristics. The cohort should reconstruct revenue and direct cost using the same classification and show historical transitions from project to managed work.

Subcontracted work can extend capability and geography. It can also reduce control, margin and transferability. Diligence should identify contractual rights, key subcontractors, mark-ups, client consent, data access, quality responsibility and termination. A target should not be treated as asset-light when it depends on minimum software commitments, dedicated facilities, specialist contractors or prepaid data.

The exit mix matters. A roll-up may acquire project-led specialists to feed recurring services across a broader platform. The case should state the conversion process, historic evidence, sales ownership, required investment and timing. A portfolio multiple cannot be justified by relabelling all revenue as technology-enabled.

8. Build an AI value score from realised commercial evidence

An AI value score should measure economic contribution across five dimensions: client outcome, commercial adoption, delivery productivity, transferability and risk control. Each dimension should use observable evidence. Client outcome may include cycle time, error, recovery, compliance or decision quality. Commercial adoption may include paid deployments, realised price, renewal and expansion. Productivity should reflect net hours and cost after review and rework.


Transferability asks whether the capability can move across teams, clients, jurisdictions and technology providers. Control asks whether data, security, validation, monitoring, incident response and human oversight match the consequence of the work. A high technical score with weak commercial evidence should remain a capability option rather than a capitalised earnings benefit.

Scoring should be calibrated by service consequence. A research-assistance tool used under expert review has a different risk profile from a system that influences regulated advice, employment decisions, safety or client funds. Higher-consequence use cases require stronger validation, oversight, auditability and incident response before they receive an equivalent control score. The rubric should record the applicable consequence category and prevent a portfolio average from masking one material high-risk deployment.

The score should be updated with longitudinal evidence. Initial diligence may rely on a limited production period. Post-close governance should compare expected and realised adoption, economics, quality and incidents by cohort. Repeated over several acquisitions, the data can improve the buyer's weighting and thresholds. A scorecard becomes strategically valuable when it learns from outcomes and changes capital allocation, rather than remaining a static transaction appendix.

Scores can create false precision. The investment committee should see the underlying evidence, weighting and limitations. A score is useful when it creates a consistent comparison and an explicit gate. It should not replace contract review, technical testing, client evidence or financial reconstruction.

Figure 2. Illustrative AI value score across five evidence dimensions



Scores are hypothetical analytical assumptions and demonstrate the method only.

Table 3. AI value scoring rubric

Dimension	0-25	26-50	51-75	76-100
client outcome	assertion	pilot measure	repeated improvement	independently evidenced outcome
commercial adoption	demonstration	unpaid pilot	paid use by several clients	renewal, expansion and pricing evidence
productivity	gross estimate	limited time study	net production cohort	cash and capacity release evidenced
transferability	individual use	one client	several teams or clients	repeatable across portfolio
control	informal review	partial controls	documented governance	tested monitoring and incident response

The score supports comparison; the underlying evidence remains authoritative.

9. Bridge gross productivity to sustainable AI economics

AI benefit should be calculated from a controlled baseline. The register should identify task, volume, handling time, quality, labour grade, exception rate, review effort, technology cost and commercial treatment before and after deployment. Changes in client mix, process scope and staffing should be isolated. A pilot performed by an expert team may not represent portfolio delivery.

Gross hours saved are not cash. The bridge should deduct implementation, integration, testing, model usage, data preparation, security, monitoring, rework, human review, client givebacks and retained capacity. It should then identify whether released people can be removed, redeployed to


paid growth or left on the bench. Timing matters because costs often precede commercial benefit.

The analysis should avoid counting the same benefit in revenue growth, margin expansion and capacity release. If AI enables a fixed-fee team to deliver more work without hiring, the value may appear through avoided cost or additional volume. If the client contract shares savings, some benefit leaves through price. If the service is billed by the hour, productivity can reduce revenue unless the commercial model changes.

Baseline selection requires care. Handling time can fall because the work mix becomes easier, an experienced team replaces trainees, demand falls or quality checks are deferred. The study should use comparable tasks, volumes, grades, quality and periods. It should record failed runs, exceptions and the time needed to prepare data. A controlled parallel run or phased deployment can provide stronger evidence than a before-and-after average exposed to several simultaneous changes.

Capacity release should be traced to a management action. Saved hours create value when the firm removes cost, avoids planned hiring, increases paid volume or improves an outcome that supports price and retention. Time dispersed across many employees may remain economically inaccessible. The workforce plan should show which roles, teams and months are affected, how service resilience is maintained, and whether released people have credible funded demand. The cash model should follow this implementation schedule.

Figure 3. Illustrative bridge from gross AI productivity to sustainable annual value



All monetary values are hypothetical analytical assumptions.

10. Test pricing and who captures the productivity benefit

Professional-services pricing may use time and materials, retainers, fixed fees, subscriptions, transactions, outcomes or hybrids. Each mechanism allocates automation benefit differently. Time-based billing can create a conflict between hours saved and revenue. Fixed fees can allow the provider to retain benefit until repricing. Outcome fees require credible baselines and acceptance. Subscription labels can still hide high service intensity.

The contract review should identify benchmarking, efficiency commitments, rate cards, gainshare, service credits, most-favoured terms and change control. Realised price should include discounts, free work, write-offs, credits and delayed acceptance. AI-related promises in proposals should be compared with contractual obligations and delivery records.

The deal model should link each use case to pricing, adoption, cost, client approval and timing. A productivity case without a commercial route is an operating hypothesis. A premium price requires evidence that clients value a differentiated outcome and have paid or renewed on that basis.

11. Measure talent dependency and knowledge transfer

Professional-services assets often reside in people. The workforce map should identify partners, sellers, subject experts, product owners, data stewards, engineers, reviewers and client leaders. Each role should be linked to revenue, client trust, workflow knowledge, intellectual property and succession. Headcount totals cannot show whether three individuals carry most of the acquisition thesis.


Retention probability should be assessed through contracts, incentives, career position, compensation, workload, geography and stated plans. Transaction bonuses can bridge closing while failing to secure long-term commitment. Non-compete and non-solicitation provisions vary by jurisdiction and circumstance. The operating plan should assume that enforceability and retention are separate questions.

Knowledge transfer should use repositories, operating manuals, model cards, code ownership, training, shadowing and tested succession. A capability is more durable when another team can reproduce the outcome with controlled quality. The buyer should budget retention, replacement, training and productivity effects explicitly.

Compensation analysis should distinguish market pay, sales credit, carried economics, transaction awards and deferred incentives. A founder or partner may accept below-market salary because value is realised through ownership. After closing, replacement economics can be materially higher. The sustainable cost bridge should normalise these roles and include the time needed to rebuild relationships or specialist knowledge.

Culture and professional identity can influence retention and client trust. Specialists may value autonomy, technical standards, publication, partnership prospects or a particular brand. Integration should identify which platform controls are essential and which local practices preserve capability. Retention interviews, organisation design and leadership appointments should occur early enough to inform price and closing conditions. A financial retention pool without a credible professional proposition may defer departure rather than resolve dependency.

Figure 4. Illustrative talent-dependency map



Revenue and capability percentages are hypothetical analytical assumptions.

Table 4. Talent-dependency and retention register

Role	Dependency	Evidence	Protection	Integration action
client leader	renewal and cross-sell	relationship map and calls	retention and handover	joint account plan
domain expert	method and quality	workflow and review record	incentive and succession	codify and train
product owner	roadmap and deployment	backlog and release history	retention and governance	portfolio product council
data steward	rights and lineage	register and approvals	covenant and access	unify governance
engineer	integration and operations	repository and incident record	team retention	architecture migration

Each critical role should have evidence, mitigation and an accountable owner.

12. Verify proprietary data, consent and permitted use

Data value depends on rights, quality, coverage, refresh, governance and the use case. A large document collection may have limited transferable value if it is customer confidential, poorly labelled, duplicated or restricted to delivering one engagement. Diligence should build a data inventory covering source, owner, subject, jurisdiction, purpose, consent, contract, retention, security, model use and deletion.

The target should demonstrate how it separates client environments, controls access, handles personal or sensitive data and responds to deletion or correction. Training, retrieval, evaluation and analytics can have different legal and contractual bases. The buyer should identify whether a model or provider receives data, retains it or uses it for improvement.

Data quality should be tested against the claimed outcome. Coverage, label accuracy, drift, bias, missingness and exception handling matter more than raw volume. A governed narrow corpus can be more valuable than a large unstructured archive. The transaction model should avoid assigning

value to data that cannot be lawfully retained, reused or monetised.

The data-room review should reconcile policy with technical reality. Access lists, storage locations, model logs, vendor settings, deletion records and incident tickets can show whether documented controls operate in practice. Samples should include active, completed and terminated clients. Legacy data deserves particular attention because its original agreements and consents may differ from current standards.

Where data value is part of the acquisition thesis, the buyer should define the economic unit. Value may arise from improving a client's workflow, training an internal classifier, evaluating models, benchmarking performance or identifying commercial patterns. Each purpose has different rights, quality and refresh requirements. The valuation should model the permitted use and necessary governance cost rather than a generic price per record.

13. Assess model-provider, cloud and software dependency

An acquisition can be operationally dependent on model APIs, cloud platforms, workflow software, data vendors and cybersecurity services. The dependency register should show spend, term, assignment, service levels, regions, data treatment, price change, termination, portability and alternatives. Minimum commitments and prepaid credits should be reconciled to usage.

Technical testing should assess abstraction, fallback, monitoring, version control and the effort required to change provider. A capability that operates only with one model may remain commercially sound. Its economics should reflect concentration, repricing and migration risk. The buyer should also test whether model updates can alter output quality or regulatory classification.

Provider dependency can become a covenant or liquidity issue when usage grows faster than client receipts. The deal case should model unit cost, peak use, currency, payment terms and service interruption. Integration should assign an owner for provider strategy and contract renewal.

14. Reconstruct pipeline, bookings and distribution evidence

Roll-up cases frequently rely on cross-selling acquired capability through a larger client base. The analysis should reconstruct opportunities by source, stage, service, client, expected value, probability, cycle time, pricing and delivery need. Historical conversion should use frozen snapshots rather than a current pipeline that has been repeatedly restaged.

Bookings definitions should distinguish executed contracts, work orders, framework ceilings, cancellable commitments and management estimates. Revenue timing should reflect mobilisation, client dependencies, staffing and acceptance. A signed master agreement can generate no revenue without awarded work.

Cross-sell evidence should show named client needs, responsible sellers, access, proposition, price, technical fit and delivery capacity. A capability may sell well through specialists and poorly through generalists. Integration should preserve the acquired sales motion until the broader channel demonstrates competence.

15. Bridge reported EBITDA to a sustainable base

Reported EBITDA should be reconciled from audited or statutory results to management reporting and monthly ledgers. Adjustments should identify owner compensation, transaction cost, acquisition effects, project overruns, capitalised development, subcontracting, disputed revenue, credits, bad debt and deferred hiring. Recurring costs should not be removed merely because management describes them as investment.

The sustainable bridge should incorporate client retention, pricing, utilisation, wage inflation, role replacement, technology spend, governance, insurance, security, integration and stranded capacity. AI productivity enters only after the commercial and cash bridge described earlier. Synergies should be separated from standalone sustainability and assigned timing, cost and owner.

Cash conversion should reconcile EBITDA to tax, capital expenditure, capitalised development, working capital, deferred income and contingent payments. Professional-services firms can finance growth through staff effort before billing and collection. WIP, unbilled revenue and ageing require engagement-level evidence.

16. Analyse working capital, WIP and cash conversion

The buyer should reconstruct monthly billing and collection by client and engagement. Time records, milestones, acceptance, invoices, credit notes, receivables and receipts should reconcile. Unbilled work can represent normal timing, disputed scope, delayed approval or unrecoverable effort. Deferred income can support liquidity while creating future delivery obligations.

Project businesses may experience cash pressure when hiring and subcontracting precede milestones. Managed services may provide steadier billing while requiring transition investment. AI implementation can add cloud and specialist cost before client pricing changes. The forecast should model these lags and payment terms.

Quality of earnings and net debt definitions should address accrued bonuses, deferred compensation, employee leave, client credits, litigation, software commitments, tax, earn-outs and change-of-control payments. The acquisition agreement should align the completion mechanism with the working-capital pattern of the business.

17. Design integration around capability preservation

Integration should protect the capability before standardising systems. The first 100 days should identify client communications, key-person retention, access continuity, code and data controls, provider contracts, financial reporting and incident escalation. A forced migration can disrupt delivery or remove the evidence needed to measure value.

The operating model should assign decision rights for product roadmap, client acceptance, model validation, data use, pricing and resource allocation. Acquired specialists require a credible position in the combined organisation. Cross-selling should use joint account planning, qualified demand and delivery gates.

Architecture integration should sequence identity, repositories, cloud, data, monitoring, billing and reporting. Each migration needs rollback, client consent where required and accountable testing. Integration expenditure belongs in the deal model before synergies are credited.

Integration governance should use measurable release gates. A capability can move to a new platform after ownership, security, data, validation, client consent, performance and rollback requirements are satisfied. The gate should identify evidence, approver and residual risk. Schedule pressure should not convert an unresolved control into an assumed completion.

Commercial integration should be sequenced with equal discipline. Sellers need qualification criteria, client examples, pricing guidance, subject-expert access and rules for claims. Delivery capacity should be reserved before broad campaigns begin. Early cross-sell can focus on clients with a documented problem and compatible data environment. Results from those cohorts should inform wider rollout and the acquisition scorecard.

18. Examine cybersecurity, privacy and AI governance

Professional-services firms can hold commercially sensitive, personal, regulated and privileged information. Diligence should review security governance, access, identity, encryption, logging, vulnerability management, incident response, backups, business continuity, suppliers and insurance. Penetration tests and certifications should be reconciled to scope, exceptions and remediation.

AI governance should identify approved uses, prohibited data, model evaluation, human oversight, change control, monitoring and incident escalation. NIST's AI RMF and Generative AI Profile provide a useful risk-management structure. The buyer should map controls to actual deployments and client obligations.

Regulatory classification can change by jurisdiction and use case. The integration plan should maintain a register of systems, decisions, consequence, users and affected data. Contractual warranties cannot substitute for tested operating controls and a funded remediation plan.

19. Apply purchase-price allocation discipline to intangible assets

IFRS 3 requires an acquirer to recognise identifiable assets and liabilities separately from goodwill when the recognition criteria are met. Customer relationships, contractual relationships and technology may be identifiable. The value of an assembled workforce is generally subsumed into goodwill. Potential contracts without contractual or other legal rights do not become identifiable assets merely because management expects them.

Transaction diligence and accounting valuation serve different purposes. The buyer should still understand which economic benefits depend on customer relationships, technology, data, brands, contracts and people. Useful life, attrition, obsolescence, contributory assets and replacement cost require evidence. A purchase-price allocation cannot validate an unsupported acquisition thesis.

Useful-life assumptions should reflect contract attrition, technology change, employee mobility, provider dependency and the pace at which methods become standard. A customer relationship can remain valuable while the associated technology changes. A technology asset can become obsolete while the firm retains client trust. Separating these mechanisms improves accounting analysis and management's understanding of where reinvestment is required.

Goodwill monitoring should connect to the acquisition case. Revenue retention, margin, key-person departure, product adoption, technology cost and integration delay can act as early indicators. Reporting-unit aggregation should not prevent management from tracking the

performance of the acquired capability. The same evidence used to approve a premium can support timely challenge when expected benefits do not emerge.

Public acquisition accounting illustrates the scale of judgement. Marsh McLennan has attributed acquisition value to customer relationships and goodwill associated with synergies and assembled workforce. Kelly Services has reported goodwill impairments in a professional-services reporting unit. These examples do not determine another transaction. They demonstrate that customer durability, expected synergies and post-close performance remain material after closing.

20. Value scenarios rather than applying a technology label

Valuation should begin with standalone cash economics, then add explicit capability and integration scenarios. A base case can include evidenced retention, price, wage, technology and cash conversion. An execution case can add qualified cross-sell, productivity and integration benefits. A downside can test client loss, slower adoption, provider repricing, talent departure and delayed capacity release.

Comparable-company and precedent multiples require consistent revenue, growth, margin, capitalisation and business mix. A target with a small AI-enabled service should not receive a technology multiple across unrelated revenue without evidence that capability changes economics at scale. The valuation bridge should show which cash flows support any premium.

Sensitivity should focus on mechanisms: retention, realised price, gross-to-net productivity, key-person departure, technology cost and integration delay. Probability weights should be visible and decision-makers should see the unweighted cases. A valuation range is strongest when it can be traced to operating evidence.

Terminal assumptions deserve particular scrutiny. AI-enabled services can grow quickly while their tools, prices and competitive advantages change. A long forecast should include continued product, security and talent investment. Margin expansion should not continue indefinitely if clients capture productivity, rivals adopt similar tools or scarce expertise becomes more expensive. The terminal case should describe the durable client problem and capability system, rather than assuming that current AI terminology remains differentiated.

The board should also see the value of waiting. Additional production cohorts, a client renewal or a resolved ownership issue can materially improve evidence. Waiting can carry competitive and execution cost, but that cost should be compared with the value of uncertainty reduction. Staged investment, a minority position or a commercial partnership can preserve access while evidence matures.

21. Use deal structure to allocate capability and execution risk

Price, deferred consideration, earn-outs, rollover equity, retention and warranties can allocate different risks. Earn-outs should use measures that remain observable after integration and avoid incentives that damage long-term value. Revenue can be influenced by pricing and allocation. EBITDA can be influenced by shared cost. Product milestones can be achieved without commercial adoption.

Useful conditions may address assignment of material contracts and licences, key-person arrangements, intellectual-property ownership, data remediation and client consent. Warranties

can cover accounts, contracts, technology, data, security, employment, tax and disputes. Specific indemnities can address identified exposures where appropriate.

The buyer should retain operational flexibility while providing a fair measurement mechanism. Governance, accounting policies, investment commitments, dispute resolution and access to records should be defined before signing. Transaction protection works best when it follows a quantified diligence finding.

22. Assess serial-acquisition competition and regulatory exposure

A roll-up should be evaluated as a strategy, including cumulative market effects. The US 2023 Merger Guidelines state that agencies may examine a series of acquisitions as part of an industry trend or strategy. Competition analysis should consider the relevant service, geography, client segment, labour market, data, interoperability and the ability of rivals to enter or expand.

Professional services can involve local licensing, independence, conflicts, ownership restrictions or procurement rules. The European Commission's review of merger guidelines includes digitalisation, innovation and dynamic effects. UK merger review can apply where jurisdictional tests are met. The transaction plan should identify filing, information, timing, remedy and integration constraints early.

Small individual acquisitions can still create strategic concentration in scarce talent, specialised data, a narrow customer group or a local service market. The acquisition register should therefore track cumulative share, common dependencies and prior commitments across the platform.

23. Align acquisition financing and covenants with cash durability

Debt capacity should follow recurring cash, client concentration, working capital, technology investment and integration risk. Lenders may focus on revenue retention, adjusted EBITDA, cash conversion, acquisition baskets, earn-outs, permitted debt, distributions and reporting. A rapidly growing roll-up can consume liquidity through purchase consideration, retention, integration and delayed billing.

The financing model should include downside covenant headroom and cash needs by month. Earn-outs and deferred consideration may behave like future claims on liquidity. Capitalised development and adjusted EBITDA definitions require clear treatment. The platform should maintain acquisition-level and consolidated reporting so lenders can distinguish organic performance from purchased growth.

Liquidity planning should include acquisition cadence. A serial acquirer can appear covenant-compliant after each deal while accumulating integration cost, deferred consideration and underperforming working capital. The model should show committed and discretionary uses, minimum cash, undrawn facilities and downside headroom across the full pipeline. Acquisition approval should be conditional on the combined platform retaining capacity to fund remediation and client service.

Reporting definitions should be agreed before financing. Organic growth, recurring revenue, adjusted EBITDA, capitalised development, synergies and earn-outs can be interpreted differently. The lender model, investment committee model and management reporting should reconcile through a controlled bridge. This reduces later disputes and makes deterioration visible

before a covenant test date.

Equity should fund uncertainty that cannot responsibly be placed on fixed claims. The capital structure should allow the buyer to pause acquisitions, invest in controls and absorb client or talent shocks without forcing value-destructive action.

24. Worked hypothetical example


Consider a professional-services platform with hypothetical annual revenue of 100 monetary units and reported EBITDA of 20. Management identifies gross AI productivity of 10 through faster research, document review and workflow automation. The target also expects 8 of cross-sell revenue through the buyer's distribution.

Diligence finds that 2.2 of the gross productivity is contractually shared with clients. Implementation and integration require 1.1, model and data services require 0.8, and retained human review requires 1.4. Capacity release contributes 1.7 because only part of the saved time can be removed or redeployed within the period. Sustainable annual AI value is therefore 6.2 before tax and financing in this hypothetical bridge.

Client cohorts show that 14 of revenue is linked to one partner-led relationship and 9 is subject to annual rebid. The downside case removes one client, delays cross-sell by twelve months and assumes the product lead departs after closing. The execution case includes only named cross-sell opportunities with qualified demand and delivery capacity.

The buyer responds with a lower upfront value, rollover equity, targeted retention, a deferred component linked to independently measurable recurring gross profit, and specific pre-close remediation for code ownership and data lineage. Integration funding is committed before any productivity synergy enters covenant EBITDA. Every figure in this example is a hypothetical analytical assumption and does not describe an actual company.

Figure 5. Acquisition decision workflow for an AI-enabled professional-services roll-up



The workflow requires evidence gates before price and integration commitments.

25. Final acquisition scorecard and governance cadence

The final scorecard should combine capability, commercial evidence, delivery, talent, data, technology, controls, cash, valuation, structure and integration. Each item needs an evidence owner, status, limitation, decision effect and post-close action. Red, amber and green labels should follow defined criteria. An unresolved high-impact matter should not become green through an average score.

The investment committee should receive the standalone case, execution case and downside separately. The approval should state the value attributed to AI capability, the evidence that supports it, the conditions required before closing and the funding required after closing. Any benefit excluded from price can remain an upside option.

Post-close governance should back-test the thesis quarterly. Measures can include client retention, realised price, paid AI adoption, gross-to-net productivity, incident rates, key-person retention, cross-sell conversion, integration cost, cash conversion and covenant headroom. Variance should be traced to original evidence and assumptions. This creates institutional learning across a serial-acquisition programme.

Governance should assign escalation thresholds. A material client notice, critical-person resignation, significant security event, provider price increase, failed control test or repeated model-quality issue may require an immediate review rather than waiting for the quarterly cycle. The review should reassess forecasts, liquidity, covenants, integration and external obligations. Actions and owners should be recorded in the same evidence system used during diligence.

Portfolio reporting should preserve acquisition cohorts. Consolidated growth can obscure whether value comes from organic performance, new acquisitions, price, currency or accounting changes. Cohort reporting allows the board to compare each acquisition with its approved case and to identify which capabilities, sectors, structures and integration choices produce durable returns. That evidence should determine future acquisition criteria and capital allocation.

Table 5. Final acquisition scorecard

Domain	Core decision	Minimum evidence	Transaction response
capability	does it solve a paid problem repeatedly?	production, outcome and client evidence	value attribution
rights	can capability transfer and operate?	ownership, licence and data lineage	condition or indemnity
economics	does benefit reach cash?	cohort and gross-to-net bridge	sustainable EBITDA
dependency	who or what can break the case?	client, talent and provider maps	retention and downside
integration	can the buyer realise value safely?	funded plan, owner and gates	timing and structure

Material exceptions remain visible even when the total score appears acceptable.

Conclusion

AI-enabled professional-services roll-ups can create defensible value by combining expertise, workflow, governed data, technology, delivery and distribution. The acquisition process should examine those elements as a connected capability system. It should value realised client and cash outcomes, identify transfer and dependency risk, and fund the controls required to scale.

The central discipline is evidence lineage. Production logs, client cohorts, contracts, time, invoices, cash, talent, repositories, data rights and incidents should connect to the transaction model. Gross productivity should be bridged to sustainable value after commercial givebacks, implementation, model cost, human review and capacity release. Valuation and structure should follow the resulting scenarios.

A buyer that preserves optionality can pay for verified standalone economics and share future value when capability proves transferable at portfolio scale. The scorecard and quarterly back-test turn each acquisition into evidence for the next decision.

References

- [1] National Institute of Standards and Technology. Artificial Intelligence Risk Management Framework: Generative Artificial Intelligence Profile. NIST AI 600-1, 2024, updated 2026. <https://www.nist.gov/publications/artificial-intelligence-risk-management-framework-generative-artificial-intelligence>
- [2] National Institute of Standards and Technology. AI Risk Management Framework resources. <https://airc.nist.gov/>
- [3] UK Department for Science, Innovation and Technology. AI Adoption Plan: Professional and Business Services. 8 June 2026. <https://www.gov.uk/government/publications/ai-champions-ai-adoption-plans/ai-adoption-plan-professional-and-business-services>
- [4] Skills England. Sector skills needs assessment: Professional and Business Services. 2026. <https://www.gov.uk/government/publications/skills-england-annual-skills-report-and-sectoral-skills-needs-assessments-2026/sector-skills-needs-assessment-professional-and-business-services>
- [5] Accenture plc. Annual Report on Form 10-K for the year ended 31 August 2025. <https://www.sec.gov/Archives/edgar/data/1467373/000146737325000217/acn-20250831.htm>
- [6] Accenture plc. Fourth Quarter and Full-Year Fiscal 2025 Results. <https://www.sec.gov/Archives/edgar/data/1467373/000146737325000213/q4fy25earnings8-kexhibit.htm>
- [7] Accenture plc. 2025 Proxy Statement. <https://www.sec.gov/Archives/edgar/data/1467373/000130817925000651/acn-20251210.htm>
- [8] Genpact Limited. Annual Report on Form 10-K for the year ended 31 December 2025. <https://www.sec.gov/Archives/edgar/data/1398659/000139865926000004/g-20251231.htm>
- [9] Paychex, Inc. Annual Report on Form 10-K for the year ended 31 May 2025. <https://www.sec.gov/Archives/edgar/data/723531/000095017025095300/payx-20250531.htm>
- [10] Kelly Services, Inc. Annual Report on Form 10-K for the year ended 29 December 2024. https://www.sec.gov/Archives/edgar/data/40570/000147793225009011/job_10k.htm
- [11] Marsh McLennan. 2025 acquisition accounting disclosure. <https://www.sec.gov/Archives/edgar/data/62709/000006270926000022/R15.htm>
- [12] IFRS Foundation. IFRS 3 Business Combinations. <https://www.ifrs.org/issued-standards/list-of-standards/ifrs-3-business-combinations/>
- [13] IFRS Foundation. IFRS 3 Business Combinations, issued standard. <https://www.ifrs.org/content/dam/ifrs/publications/pdf-standards/english/2022/issued/part-a/ifrs-3-business-combinations.pdf?bypass=on>

- [14] IFRS Interpretations Committee. Customer-related intangible assets, agenda decision, March 2009. <https://www.ifrs.org/content/dam/ifrs/supporting-implementation/agenda-decisions/2009/ifrs-3-ias-38-customer-related-intangible-assets-march-2009.pdf>
- [15] IFRS Foundation. IAS 38 Intangible Assets. <https://www.ifrs.org/content/dam/ifrs/publications/pdf-standards/english/2021/issued/part-a/ias-38-intangible-assets.pdf?bypass=on>
- [16] US Department of Justice and Federal Trade Commission. Merger Guidelines. 18 December 2023. https://www.ftc.gov/system/files/ftc_gov/pdf/2023_merger_guidelines_final_12.18.2023.pdf
- [17] Federal Trade Commission. Guide to Antitrust Laws: Mergers. <https://www.ftc.gov/advice-guidance/competition-guidance/guide-antitrust-laws/mergers>
- [18] European Commission. Review of the Merger Guidelines. 2026. https://competition-policy.ec.europa.eu/mergers/review-merger-guidelines_en
- [19] European Commission. Overview of EU merger control. https://competition-policy.ec.europa.eu/mergers/overview_en
- [20] European Commission. Competition policy for professional services. https://competition-policy.ec.europa.eu/sectors/services_en
- [21] European Commission. Annual Competition Report 2025. https://competition-policy.ec.europa.eu/document/download/a8e84d54-af12-4f60-9361-2b7638ac2de2_en?filename=annual-competition-report_2025_report_part1_EN.pdf
- [22] UK Competition and Markets Authority. Mergers: how they are investigated. Updated 19 June 2025. <https://www.gov.uk/government/publications/mergers-how-they-are-investigated/mergers-how-they-are-investigated>
- [23] UK Competition and Markets Authority. Annual Plan 2025 to 2026. https://assets.publishing.service.gov.uk/media/6784fd5df0528401055d2310/CMA_Annual_Plan_2025_to_2026.pdf
- [24] European Commission. Competition in Artificial Intelligence: keynote speech. 20 March 2024. https://competition-policy.ec.europa.eu/document/download/8d60e1b5-ec5d-4220-9d54-6d3360e863b5_en?filename=20240320_GCR_TMT_Conference_Guillaume-Loriot_keynote_speech.pdf

About the Author

Authored by Chennakeshav Adya, Independent Researcher

Chennakeshav (CK) is a corporate finance and investment banking executive with 25+ years of global experience in deal origination, structuring and execution across M&A, growth capital and corporate strategy. He has led value-creation mandates for founders, corporates and funds - bridging the boardroom view to hands-on execution and close.

His career spans Morgan Stanley, HSBC, Lloyds Banking Group, EWEC, ADQ portfolio companies and Emirates Growth Fund, across TMT, real estate, fintech, deeptech, cleantech, infrastructure and energy. He has partnered with C-suite leaders, private equity and venture funds, sovereign wealth funds and family offices to finance complex fund raises and scale-up ventures, and has led M&A due diligence, post-merger integration and business-transformation initiatives to create value.

At Matchpoint Partners he is Managing Partner, leading the firm's corporate finance, M&A and capital-raising practice. He holds an MBA from London Business School, an engineering degree from VTU and a Master of Laws (LLM, in progress) from UCL London.

An active start-up mentor, CK mentors at Techstars, DIFC FinTech Hive, Startup Grind, Founder Institute and IN5, serves as Entrepreneur Mentor in Residence (EMiR) at London Business School, and judges the Entrepreneurship World Cup.