

Agent Roll Ups in Professional Services Margin Expansion or Hidden Supervision Cost

A Transaction Framework for Workflow Economics Quality Control Client Acceptance and Sustainable Value

Authored by Chennakeshav Adya

Independent Researcher

September 2026

Abstract

An acquisition programme in professional services can appear to combine two familiar sources of value: roll-up economics and artificial-intelligence productivity. The reported case may show lower delivery hours, wider partner leverage, centralised technology and cross-selling. The buyer can still inherit a weaker business if automated production shifts effort into senior review, exception handling, client explanation, data control and remediation. Those costs often sit across partner time, write-offs, support functions and unrecorded intervention rather than in one visible technology budget.

This paper develops a transaction framework for deciding whether an AI-enabled professional-services roll-up creates sustainable margin or hides supervision cost. It reconstructs the economics at workflow level, separating origination, production, review, rework, client acceptance, collection and post-delivery liability. It tests the labour pyramid, supervisory span, utilisation, pricing model, revenue recognition, client consent, professional accountability, technology dependency, partner retention and integration design. It then connects the evidence to sustainable EBITDA, purchase price, earn-outs, warranties, integration funding and the first hundred days.

Four hypothetical targets illustrate the method: an accounting-advisory firm, a legal-services platform, an engineering consultancy and a technology-services business. A common one-hundred-job model shows that a 35 percent reduction in production hours can produce only a 5.8 percent reduction in total delivery hours when review, rework and AI operations expand. A redesigned control model improves the saving to 15.8 percent. A separate hypothetical margin bridge reduces reported EBITDA of USD 9.6 million to sustainable EBITDA of USD 5.0 million after normalising supervision, technology, partner retention, rework and client-remediation costs. These values are management assumptions for method illustration; they are not observations, forecasts or valuation conclusions.

The analysis finds that task automation is not transaction evidence of margin expansion. A buyer needs accepted output, collected revenue, attributable review effort, reproducible quality, bounded professional liability and a delivery model that clients and accountable practitioners will continue to use. Value should be released when those outcomes are demonstrated by workflow cohorts and supported by contracts, time records, quality files, complaints, write-offs, collections and tested integration controls.

JEL Classification: G24, G34, J24, L84, M15, O33

Keywords: professional services M&A, agentic AI, supervision cost, utilisation, revenue quality, roll ups, due diligence, integration, valuation

Introduction

Professional-services businesses convert judgement, expertise, trust and controlled execution into client outcomes. Their economics are commonly described through revenue per professional, utilisation,

recovery, leverage, partner productivity, recurring revenue and operating margin. AI agents can change every measure because they can research, draft, classify, calculate, coordinate, monitor and prepare client work. The apparent productivity gain can be real. Its economic location and durability require transaction evidence.

A roll-up adds another layer. The buyer expects common systems, shared delivery centres, procurement, cross-selling, central functions and stronger market access. Acquired firms may use different engagement terms, time systems, review rules, data environments and definitions of completed work. An agent deployed across the group can amplify standardisation or distribute one weak control across many clients. The acquisition thesis therefore needs a workflow model that joins commercial, professional, technological and financial evidence.

Public standards and regulator materials place continuing responsibility on accountable professionals and firms. The American Bar Association discusses competence, confidentiality, supervision, communication and reasonable fees when lawyers use generative AI [21]. The UK Solicitors Regulation Authority calls for leadership, impact assessment, policies, training, monitoring and protection of client confidentiality [22-23]. The UK Financial Reporting Council and the International Auditing and Assurance Standards Board emphasise professional judgement, evidence and quality management when technology supports audit work [24-30]. These sources do not establish the economics of a target. They identify duties that can create real operating work.

This paper provides a transaction method for strategic buyers, private-equity sponsors, family offices, boards, lenders and management teams. It asks whether the combined platform can deliver more accepted work and cash with a controlled cost base. It treats margin as an outcome of the complete client workflow rather than a claim derived from model speed.

1 Define the acquisition decision

The investment committee should begin with a decision statement. The statement identifies the service lines, jurisdictions, client groups, target entities, technology assets and operating capabilities being acquired. It then identifies the value mechanism: greater delivery capacity, lower unit cost, faster turnaround, improved quality, broader client coverage, recurring managed services, cross-selling or a defensible data and workflow system. Each mechanism requires distinct proof.

The decision statement should define what management means by an AI agent. A research assistant, document classifier, workflow orchestrator, calculation engine and autonomous customer-facing agent create different supervision and liability profiles. The target should identify where the system observes, recommends, prepares, approves, communicates or executes. A high model-capability score does not show that the workflow is commercially accepted or professionally controlled.

The acquisition perimeter should include the legal entities, regulated permissions, engagement contracts, partners, employees, contractors, delivery centres, client data, models, prompts, agent policies, tools, vendor agreements, quality systems, insurance and claims. The buyer should distinguish technology owned by the target from licences, customer-specific configurations and third-party services. Transferability can be limited by consent, professional regulation, independence rules, data rights or vendor terms.

The committee should state the conditions under which the thesis fails. Examples include senior review rising faster than automated production, clients rejecting the delivery model, hourly revenue falling with time saved, partners leaving after consideration is paid, common technology increasing correlated liability, or integration delaying collections. These failure conditions become testable diligence questions and help prevent a broad AI narrative from substituting for an acquisition case.

2 Reconstruct the end to end service workflow

Professional-services margin is created across the whole workflow. A useful map begins with lead qualification and conflict checking, moves through scoping, pricing, staffing, evidence collection, production, review, client discussion, approval, delivery, billing, collection and follow-up. It also includes complaints, claims, corrections, write-offs and knowledge capture. The target should map each material service line at this level.

For each step, diligence should identify the accountable role, system, data source, agent function, review requirement, hand-off, completion evidence and failure response. A task can be fast while the surrounding process remains slow. Automated drafting may reduce production time while incomplete evidence, approval queues or client negotiation become the constraint. A roll-up may inherit multiple definitions of ready, reviewed, delivered and billable.

The map should separate standard and exceptional work. Agents often perform well on repeatable cases and transfer complex, ambiguous or high-risk matters to senior people. If the exception rate rises with volume, the labour pyramid can invert. The buyer should therefore measure the share of work that follows the standard path, the reasons for exception and the senior time consumed by each class.

Evidence should come from workflow logs, time records, matter or project systems, quality files, customer-support records, billing and bank collections. Interviews explain the process but do not measure it. A representative sample should be traced from initial instruction to accepted output and cash. That trace establishes the economic denominator for later productivity claims.

3 Separate task speed from accepted throughput

Task speed measures how quickly a defined activity is completed. Accepted throughput measures completed client work that passes review, is accepted, billed and collected. The distinction is central to an acquisition because software demonstrations usually observe task speed while value depends on accepted throughput. A draft produced in one minute has limited economic value if a partner spends an hour finding and correcting errors.

The buyer should define a common unit of service for each workflow. Units may be completed tax returns, diligence workstreams, contract reviews, engineering packages, audit tests, software releases or client reports. Each unit should have entry criteria, required evidence, review level, acceptance event, price and collection status. The same unit must be used before and after AI adoption.

Throughput should be decomposed into production, review, rework, waiting and external delay. Production time can fall while total cycle time remains unchanged because review capacity is fixed. The target should show cohort data rather than selected examples. Cohorts should control for complexity, client, team, jurisdiction and service type. A comparison without these controls can attribute a favourable mix shift to technology.

The buyer should also measure rejected output and silent repair. Rejected output appears in formal quality records. Silent repair occurs when staff correct work before the system records an exception. Sampling version histories, comments and file activity can reveal this hidden labour. Sustainable margin requires a repeatable reduction in total cost per accepted unit, not merely faster first output.

4 Model the labour pyramid and supervisory span

Professional-services firms traditionally combine junior production, manager review, specialist input and partner accountability. AI can remove junior tasks without removing the need for judgement, explanation or sign-off. The target labour pyramid may therefore become narrower at the base and heavier in review. A buyer should model roles by actual workflow activity rather than job title.

Supervisory span should be measured as the number and complexity of active work units per accountable reviewer. A wider span can improve leverage when work is standardised, evidence is visible and exception routing is reliable. It can weaken control when reviewers receive more outputs, less context and compressed deadlines. Approval rates alone provide weak evidence because routine confirmation can conceal insufficient attention.

The buyer should measure review minutes, rejection, amendment, escalation and subsequent defect by reviewer and workflow. It should also examine waiting time and workload distribution. A small group of highly trusted partners may be absorbing AI exceptions, client concerns and sales support. Their effort can be missing from project codes or charged to business development, administration and integration.

The operating model should specify which decisions remain with licensed or accountable professionals. It should identify when a second review, specialist consultation or independent challenge is required. Professional judgement is not a generic overhead; it is a production input in regulated and high-consequence services. The acquisition case should fund the amount of supervision needed by the accepted control design.

5 Treat supervision as a measurable production input

Supervision cost includes more than formal sign-off. It includes defining the task, selecting evidence, configuring the agent, interpreting output, resolving uncertainty, explaining the result, documenting judgement, escalating exceptions and monitoring performance. It also includes maintaining policies, evaluations and permissions. The buyer should assign these activities to work units and service lines.

Time systems can understate the cost. Senior personnel may review between meetings, correct work without opening a formal matter code, or treat AI oversight as innovation activity. The diligence team should reconcile calendar samples, document histories, approval logs, time entries and project economics. The purpose is to estimate effort, not to monitor individuals beyond legitimate diligence needs.

Review intensity should reflect consequence and evidence quality. A standard internal summary can use sampling and automated checks. A client opinion, regulated filing, financial conclusion or safety-relevant design may require a named professional to inspect underlying evidence and reasoning. The FRC's AI-in-audit materials stress documentation and professional judgement around technology-supported procedures [24-25]. IAASB materials retain accountability within firm and engagement quality systems [26-30].

Supervision should be included in bids, staffing plans, utilisation targets and margin reports. If it remains an unallocated partner burden, reported engagement margins can improve while partner capacity and retention deteriorate. A roll-up that standardises supervision measurement can create value even before it automates more work.

6 Measure review and rework economics

The review model should distinguish first-pass review, substantive correction, technical consultation, client clarification and final approval. Rework should identify whether the cause is incomplete input, agent error, policy mismatch, human error, client change or integration failure. Different causes require different investment and have different recurrence.

A useful metric is total labour hours per accepted unit. Another is cost per accepted unit, which applies role-specific compensation and delivery-centre cost. Review yield measures the share of units approved without substantive correction. Rework recurrence measures whether the same defect returns after remediation. These metrics should be produced for pre-AI, early adoption and current cohorts.

The hypothetical one-hundred-job model in this paper begins with 1,000 production hours, 120 review hours and 80 rework hours. Initial AI deployment reduces production to 650 hours but raises review to 260, rework to 140 and AI operations to 80. Total hours fall from 1,200 to 1,130, a 5.8 percent reduction.

A controlled model with stronger input validation and exception routing reduces total hours to 1,010, a 15.8 percent reduction. These are management assumptions.

The comparison illustrates why the buyer should avoid applying the production-hour reduction directly to payroll. Capacity release depends on review, rework, service demand, employment structure and the ability to redeploy people. Cash savings should be distinguished from avoided hiring, additional capacity and improved turnaround.

7 Reinterpret utilisation and recovery

Utilisation traditionally compares chargeable hours with available hours. AI can reduce the hours needed for a job and therefore lower utilisation even when accepted output rises. A buyer should preserve the traditional measure for workforce planning while adding output, cycle-time, quality and contribution measures. Otherwise management may discourage efficient delivery or report an apparent utilisation decline as commercial weakness.

Recovery compares billed or collected value with recorded time at standard rates. The measure can improve when fewer hours deliver the same fixed fee. It can deteriorate when hourly work becomes faster and clients pay only for actual time. The buyer should segment revenue by hourly, fixed, subscription, contingent, managed-service and outcome-linked pricing. Each model transmits productivity differently.

Available capacity is not automatically saleable. The target needs demand, sales coverage, suitable skills and client permission to use the delivery model. Released junior hours may not match demand for senior advisory work. The acquisition case should show how capacity becomes new revenue, avoided hiring or cash cost reduction. Timing matters because payroll and partner drawings can remain fixed during transition.

The combined group should define a contribution measure by accepted work unit. It should include direct production, review, rework, AI platform, quality, data, client support and expected claims cost. This measure is harder to manipulate through time classification and gives the board a clearer view of whether the operating model creates value.

8 Test pricing and billing-model resilience

AI productivity can create a conflict between cost reduction and revenue preservation. Hourly billing may transmit time savings to the client. Fixed fees can retain part of the efficiency gain but expose the firm to rework and scope risk. Subscriptions can improve visibility while requiring ongoing service capacity. Outcome pricing can align value and increase contingent risk.

The buyer should inspect engagement letters, pricing approvals, invoices, write-offs, disputes and renewals. It should compare realised price and contribution before and after AI adoption. Sales presentations should be reconciled with contractual promises. A firm that markets expert accountability while relying on unreviewed automation may create legal, reputational and retention risk.

The ABA's Formal Opinion 512 discusses reasonable fees and actual time in the context of lawyers using generative AI [21]. Applicability depends on the engagement and jurisdiction. The transaction lesson is broader: the revenue model must match the service actually delivered and the obligations retained by the firm. Technology cost or efficiency does not automatically create a right to preserve historical billing.

The acquisition model should therefore use pricing cohorts. It should identify which clients accepted revised terms, which renewed, which demanded lower fees and which increased scope. Claimed price resilience should be supported by collected cash and retention. Pipeline assumptions require signed or historically evidenced conversion rather than management expectation alone.

9 Verify client acceptance and consent

Client acceptance has commercial, professional and data dimensions. A client may accept the final work while objecting to the use of a particular model, data location, subcontractor or automated decision path. The buyer should identify contractual notice and consent requirements, confidentiality obligations, audit rights and restrictions on model training or reuse.

Evidence should include signed engagement terms, amendments, client communications, procurement questionnaires, security reviews, complaints and renewal decisions. Generic disclosure can be insufficient where a material use affects the means of service, confidential information or accountable judgement. Legal advice is required for the relevant profession and jurisdiction.

The buyer should also test user behaviour. Professionals may use unapproved tools when official systems are slow or incomplete. The SRA's 2026 warning notice identifies inaccurate information and confidentiality as concerns in legal services and describes risks from entering client-sensitive data into public AI tools [23]. A target's stated policy should be reconciled with technical controls, training, logs and incident evidence.

Client acceptance should be modelled as a value driver. Approved use can shorten cycle time and improve consistency. Restricted use can preserve manual cost or require separate environments. A roll-up can create value by building a credible control platform that smaller acquired firms could not fund alone. The cost and customer benefit of that platform belong in the acquisition case.

10 Preserve professional accountability and judgement

Professional services often involve duties that remain with individuals or regulated firms. The system can support judgement without owning professional responsibility. Diligence should identify who signs, certifies, approves or advises; which evidence they must review; how disagreements are handled; and whether the person has authority, competence, independence and time.

The quality file should reproduce the path from source evidence to conclusion. It should identify agent versions, tools, policies, prompts where relevant, human amendments, consultations and final approval. Audit and assurance standards provide useful examples of quality management, engagement responsibility and professional scepticism [26-30]. Legal and engineering services have different rules; the same evidence principle still assists transaction diligence.

Automation bias can weaken review when a plausible output receives less challenge. Controls should require reviewers to inspect critical evidence, uncertainty and contradictory information. The EU AI Act's human-oversight provisions for certain high-risk systems describe the need to understand limitations, avoid over-reliance and retain the ability to disregard or stop an output [18-20]. Applicability requires legal analysis.

The buyer should treat accountable review capacity as a constraint in the operating plan. Growth that outruns partner or specialist supervision can reduce quality and increase claims. A platform that measures and routes judgement effectively can support a wider span; the evidence must show that quality remains stable as volume grows.

11 Examine data confidentiality and privilege

Client work can contain personal data, trade secrets, financial information, legal advice, designs and regulated records. The diligence map should identify what data enters each agent, where it is processed, how long it is retained, whether it is used for training, who can access traces and how deletion operates. It should distinguish firm data from client-owned information and public sources.

Confidentiality risk can arise from prompts, retrieved documents, tool results, logs, support access and cross-client memory. The target should demonstrate tenant separation, access control, encryption, redaction, retention, incident response and vendor terms. The buyer should test representative workflows and verify that controls apply to integrations and secondary agents.

Professional privilege and confidentiality require specific legal analysis. The SRA warns that placing client information into unsuitable public tools can jeopardise confidentiality and privilege [23]. The ABA also identifies protection of client information as an obligation when lawyers use generative AI [21]. These materials are jurisdiction-specific; they show why client data treatment belongs in the transaction perimeter.

The operating model should price privacy and security work. Dedicated environments, private models, regional hosting and client-specific controls can increase cost while supporting retention and premium services. A target that excludes these costs from service-line margins overstates sustainable contribution.

12 Assess quality evidence and claims history

Quality should be measured through defects, corrections, write-offs, complaints, service credits, claims, regulator contact and client loss. The buyer should reconcile registers across risk, legal, professional standards, customer success, insurance and finance. An AI-related problem may be classified as ordinary rework and never appear in a central incident register.

Sampling should compare agent-supported and conventional work at similar complexity. Reviewers should record the nature and severity of changes, not only pass or fail. The target should show how a defect changes prompts, retrieval, policies, training, staffing or client instructions, and whether remediation is tested before closure.

Claims history can understate future exposure when deployment is recent, limitation periods are long or clients have not detected errors. Scenario analysis should therefore supplement observed loss. The model should separate known claims, possible obligations, future operating risk and insurance recoverability. IAS 37 and business-combination requirements inform accounting analysis; transaction advisers determine the treatment for the deal [1-7].

The acquisition committee should ask whether scale changes correlation. A common agent, model, policy or dataset can introduce the same defect across many client matters. Portfolio-level quality testing and staged release can reduce this exposure. Correlated failure deserves a specific downside case rather than an average historical error rate.

13 Standardise without erasing local requirements

Roll-ups often seek one delivery model. Standardisation can reduce duplicated technology, improve data, support training and create a common client experience. It can also erase local professional rules, engagement terms, language, precedent and client expectations. The integration design should identify what can be common and what must remain jurisdictional or service-specific.

A useful architecture has common control components with configurable policies. Identity, logging, evaluation, model inventory and incident management can be shared. Approval thresholds, data locations, professional sign-off and client disclosures can vary by service and jurisdiction. Exceptions should be explicit, owned and tested rather than embedded in undocumented local workarounds.

The buyer should inspect taxonomy alignment. Acquired firms may use different codes for service, complexity, risk, client, stage and completion. An AI system trained or measured on inconsistent labels can produce misleading comparisons. Data harmonisation should precede group-wide productivity targets.

Standardisation should be released by workflow cohort. The first cohort should have clear inputs, representative volume, stable client terms and available reviewers. Results should be measured through accepted throughput, quality, contribution and collections. Expansion should follow evidence. A group-wide launch based on demonstration performance can distribute one control gap across the estate.

14 Diligence the technology and agent architecture

The technology review should connect architecture to service economics. It should identify models, retrieval sources, orchestration, agent identities, tools, permissions, memory, evaluation, observability, deployment, security and vendor dependencies. The buyer should determine which components are proprietary, configurable, transferable and necessary for current revenue.

Agent identity and authorisation matter when software can access client systems or change work product. NIST's 2026 concept paper highlights identification, authorisation, audit and non-repudiation for software and AI agents [10-13]. OWASP and MITRE provide threat and control taxonomies for excessive agency, prompt injection and other AI-system risks [14-16]. The target's actual implementation remains the evidence.

The architecture should support attribution. A reviewer needs to know which agent version, evidence, policy and tool produced the work. The system should prevent an agent from exceeding the originating user's authority and should bind approval to the exact action. Shared privileged accounts and uncontrolled connectors can increase both security exposure and professional uncertainty.

Portability should be tested. Replacing a model or vendor can change output, latency, cost, data terms and evaluation. A multi-model diagram does not establish operational portability. The buyer should observe a controlled substitution and measure the effort, quality and customer consequences. Vendor concentration belongs in valuation and integration funding.

15 Test talent retention and partner incentives

Professional-services acquisitions depend on people whose client relationships, judgement and reputation may not be contractually transferable. AI can increase concern about role reduction, quality and compensation. The buyer should identify key partners, specialists, reviewers, sales leaders, data owners and technology operators, then map their economic and decision rights.

Retention analysis should distinguish contractual restrictions from willingness to remain. The current US federal noncompete rule is not in effect, according to the Federal Trade Commission's status page [45]. State and other jurisdictions vary. The buyer should obtain current legal advice and avoid assuming that a restriction guarantees retention.

Incentives should align accepted work, client value, quality and cash. A bonus based only on utilisation can discourage automation. A bonus based only on automated volume can weaken review. Partner economics may need to recognise origination, accountable judgement, knowledge contribution and platform adoption. The transition should avoid shifting uncompensated supervision to a small group.

The buyer should model departures and capacity loss. A senior professional can take clients, delay sign-off and remove tacit knowledge. Earn-outs, rollover equity, deferred consideration, retention awards and governance roles can support continuity when designed carefully. The deal model should include their cash cost and behavioural consequences.

16 Normalise reported EBITDA

Reported EBITDA should be reconciled from general ledger to service-line and workflow economics. Adjustments should cover owner compensation, partner drawings, related parties, capitalised

development, non-recurring items and acquisition costs. The AI operating model adds review, rework, platform, evaluation, security, data, insurance, client-remediation and integration costs.

Hidden supervision often appears in partner administration, business development, innovation or unrecorded time. The buyer should estimate a sustainable charge based on observed activity and compensation. It should also normalise technology cost where pilots, vendor credits or buyer-funded services reduce the historical expense. Current cost may not reflect the licences and infrastructure required at combined-group scale.

The hypothetical bridge in this paper begins with USD 9.6 million of reported EBITDA on USD 48 million of revenue. It deducts USD 2.4 million of unrecorded supervision, USD 1.2 million of steady-state AI and control cost, USD 1.0 million of partner-retention economics, USD 0.6 million of recurring rework and USD 0.4 million of client-remediation cost. Sustainable EBITDA is USD 5.0 million, or 10.4 percent of revenue. These are management assumptions.

The bridge should not automatically classify every deduction as permanent. A validated control redesign can reduce review and rework. Client acceptance can support better pricing. The investment case should show timing, owners, expenditure and evidence for each improvement. Purchase price should reflect the state at closing and the risk of delivery.

17 Build the synergy case from operating evidence

Synergies should be defined as cash, capacity or risk outcomes with a baseline, action, owner, cost, timing and evidence. Technology procurement, common delivery, cross-selling, shared functions and reduced rework can create value. Integration, retention, quality, data migration and client concessions can consume it.

The hypothetical gross annual synergy case totals USD 7.8 million: USD 2.4 million from delivery capacity, USD 1.6 million from shared technology and data, USD 1.3 million from central functions and USD 2.5 million from cross-selling contribution. It deducts USD 2.1 million of continuing supervision, USD 1.2 million of AI operations and assurance, USD 0.9 million of retention cost and USD 0.9 million of client and integration leakage. Net annual synergy is USD 2.7 million. These amounts are assumptions.

Capacity synergy should not be treated as immediate cash unless roles leave or hiring is avoided. Cross-selling should be supported by client overlap, consent, sales capacity and historical conversion. Procurement savings should reflect termination fees and volume tiers. Quality improvements should be linked to observed rework and claims.

The synergy register should distinguish run-rate, cumulative cash and present value. It should also show downside if partner departures, client restrictions or technology migration delay implementation. Evidence gates can release management incentives or contingent consideration as accepted throughput and cash are demonstrated.

18 Value the business on sustainable economics

Valuation should triangulate discounted cash flow, market approaches and precedent transactions while adjusting for the target's evidence state [1-7]. Revenue growth should be linked to client retention, realised price, accepted capacity and pipeline conversion. Margin should include the complete delivery and control model. Working capital should reflect billing, disputes and collection.

The assembled workforce has economic importance even though IFRS 3 explains that it is not recognised separately from goodwill [7]. Customer relationships, contracts, technology and other identifiable assets require accounting assessment [1-4]. The transaction team should keep purchase-price allocation distinct from commercial valuation and legal transfer.

The hypothetical valuation illustration applies a 9.0 times multiple to USD 5.0 million of sustainable EBITDA, producing USD 45 million. It adds USD 12 million for evidence-weighted synergy present value, deducts USD 15 million for integration and control investment, and deducts USD 8 million for retention, client and execution risk. Illustrative enterprise value is USD 34 million. The figures are not a market benchmark or recommendation.

Sensitivity should vary accepted-work growth, review intensity, partner retention, realised price, claims, technology cost and integration delay. The committee should see which assumption drives value and which evidence can change it. A high headline multiple can be justified only by durable cash economics and a controlled risk profile.

19 Select deal protections that fit the findings

Transaction documents can allocate identified risks when definitions match the business. Representations may address client authority, professional compliance, data use, technology ownership, agent controls, incidents, complaints, claims, insurance, key people and financial information. Disclosure schedules should identify exceptions by client, service and workflow where material.

Specific indemnities can address known claims or defined liabilities. Escrow, holdback and warranty insurance can support recovery subject to limits, exclusions and collectibility. Professional indemnity, cyber and technology errors-and-omissions policies should be reviewed for insured entities, retroactive dates, AI-related exclusions, change of control and notice.

Earn-outs can link price to collected revenue, client retention, contribution after supervision cost, quality and key-person continuity. A revenue-only earn-out can reward low-margin or poorly controlled work. An EBITDA earn-out can be affected by integration allocations. Measures require clear accounting policies, access rights and dispute mechanisms.

Closing conditions may be appropriate for material transfer consents, professional permissions, data rights, critical control gaps or key-person commitments. Lower-severity findings can be managed through covenants and funded integration plans. Transaction counsel, tax advisers, accountants and regulators determine the actual structure.

20 Design the integration around client work

Integration should preserve client service and accountable authority. The buyer should avoid changing identity, data location, model, engagement terms and responsible professional at the same time for high-risk work. A cohort sequence allows the combined group to test controls and maintain rollback.

The first phase preserves engagement files, permissions, logs, quality records, time data, billing and collections. The second phase reconciles taxonomies, client restrictions, reviewer roles and vendor dependencies. The third phase tests workflow, approval, exception routing, recovery and reporting. The fourth phase migrates suitable cohorts and measures accepted output, quality, contribution and cash.

Client communication should be planned by contract and relationship. Some changes can be operational. Others require notice, consent or a revised engagement. Sales teams should receive accurate descriptions of the delivery model. A claim of autonomous service can create expectations and liability that the operating design does not support.

Integration governance should include service leaders, professional-risk owners, technology, security, data, finance, HR, legal and client teams. The committee should review a common evidence pack rather than separate functional dashboards. Value is released when the complete workflow operates in the combined environment.

21 Build a transaction evidence room

The evidence room should contain service-line economics, workflow maps, time and activity data, pricing, contracts, invoices, collections, write-offs, quality records, complaints, claims, insurance, people data, technology architecture, model and agent inventories, permissions, vendor terms, data controls and integration plans. Files should be current, attributable and reconciled.

The buyer should request machine-readable data where available. Samples should retain identifiers needed to connect workflow, review, invoice and collection while respecting privacy and privilege. The diligence protocol should define who can access sensitive client and personnel information. Unsupported management estimates should remain labelled as estimates.

Evidence should be tested for completeness. Time data should reconcile to payroll and project systems. Revenue should reconcile to the ledger and bank receipts. Agent activity should reconcile to workflow units. Quality issues should reconcile across professional-risk, support, legal and insurance records. Differences can reveal hidden work or incomplete control.

The final transaction file should record each material finding, evidence, responsible adviser, value effect, document response and integration owner. This supports investment-committee review and reduces the risk that a technical observation is lost before pricing or closing.

22 Compare four hypothetical target types

The accounting-advisory target uses agents to assemble evidence, perform checks and draft workpapers. Its value depends on consistent documentation, reviewer judgement, independence and client acceptance. The legal-services platform uses agents for research, document review and drafting. Its material issues include confidentiality, competence, supervision, fee treatment and candour.

The engineering consultancy uses agents to retrieve standards, prepare calculations and coordinate design packages. Professional sign-off, version control, safety consequence and evidence trace are central. The technology-services business uses agents for code, testing, support and cloud operations. Security, customer authority, service levels and rollback become material.

The hypothetical cases have different initial production savings and review burdens. Accounting shows 25 percent production saving with 12 percent additional review and rework. Legal shows 32 percent saving with 20 percent additional control effort. Engineering shows 22 percent saving with 18 percent additional review. Technology services shows 38 percent saving with 17 percent additional operations and quality work. These are management assumptions expressed as shares of baseline labour cost.

The comparison shows why one group-wide margin target can misallocate capital. Each service requires its own accepted-work unit, reviewer model, pricing pathway and risk case. The roll-up creates value when the common platform supports those differences efficiently.

23 Analyse the illustrative unit economics

The one-hundred-job model provides a controlled comparison. Baseline delivery uses 1,000 production hours, 120 review hours and 80 rework hours. Initial AI deployment uses 650 production, 260 review, 140 rework and 80 AI-operations hours. Controlled AI uses 650 production, 200 review, 90 rework and 70 AI-operations hours.

The initial AI case saves 350 production hours and gives back 280 through additional review, rework and operations. Total saving is 70 hours. The controlled case gives back 160 hours and saves 190. The difference comes from input validation, bounded workflows, clearer reviewer information, defect prevention and exception routing. No additional model capability is assumed.

The same hours can have different cost because senior review is more expensive than junior production. The buyer should apply actual role cost and delivery location. It should also account for fixed payroll, severance, hiring avoidance and redeployment. A percentage labour saving is not equivalent to the same percentage cash saving.

Revenue should be added only after testing pricing. Fixed-fee work may preserve value; hourly work may transmit efficiency to clients. Increased capacity creates revenue only if demand and sales conversion exist. The board should therefore monitor cost per accepted unit and contribution per accepted unit alongside hours.

24 Translate evidence into the valuation bridge

The valuation bridge should begin with reported results and move through identifiable evidence. Supervision normalisation is supported by time, calendar, approval and version data. Technology normalisation is supported by licences, cloud use, staffing and security. Retention cost is supported by compensation, rollover and key-person plans. Rework and remediation are supported by quality, write-off and complaint records.

Synergy should then be added only where the buyer has a funded action and capacity to execute. Gross savings should be reduced for control operations, client concessions, partner retention and integration leakage. Present value should reflect timing, tax, capital expenditure and probability. A one-year run-rate should not be added to value without duration and cost.

The committee should maintain separate central, downside and severe cases. The downside can combine slower client acceptance, higher reviewer intensity and delayed technology migration. The severe case can include partner departures, correlated quality failure and a material claim. These are scenario tools, not predictions.

The resulting price should be compared with the transaction structure. A lower upfront price with evidence-based contingent consideration may align uncertainty. A higher fixed price requires stronger verified economics and protection. The board should record which evidence justified each release of value.


25 Review the figures and decision tables

The figures and tables below convert the framework into an acquisition control pack. They use hypothetical data and proposed processes. Their purpose is to make the relationships between task automation, supervision, margin, synergy and transaction evidence explicit.

The margin bridge separates reported EBITDA from sustainable delivery economics. The labour stack shows why production-hour reduction can be absorbed by review, rework and AI operations. The supervisory curve illustrates that exception volume can rise non-linearly as a reviewer carries more concurrent work. The synergy bridge separates gross initiatives from continuing control and integration costs. The first-hundred-day sequence links evidence preservation to cohort migration.


The tables define the diligence perimeter, workflow evidence, one-hundred-job economics, target-type differences, margin normalisation, transaction structure and evidence gates. A real transaction should replace every assumption with source data, adviser conclusions and approved management estimates. The board should preserve the version used for the acquisition decision.

Figure 1 Hypothetical reported to sustainable EBITDA bridge




Management assumptions in USD millions; the bridge is not an observation forecast or valuation conclusion.

Figure 2 Hypothetical labour hours per one hundred accepted jobs




Management assumptions; total hours are 1200 baseline 1130 initial AI and 1010 controlled AI.

Figure 3 Hypothetical supervisory span and exception burden



Management assumptions; the curve illustrates a possible capacity constraint and is not a benchmark.

Figure 4 Hypothetical gross to net annual synergy bridge



Management assumptions in USD millions; net annual synergy is USD 2.7 million before tax financing and present value.

Figure 5 First hundred days workflow economics sequence

Release value when accepted work quality contribution and cash are demonstrated



Proposed sequence; actual timing should reflect client professional regulatory people and systems requirements.

Table 1 Professional services roll up diligence perimeter

Component	Required evidence	Decision question	Hidden cost risk
Client promise	engagement terms scope price and acceptance	what outcome is sold and accepted	automation changes service without client alignment
Workflow	task review exception delivery and collection records	where capacity and delay actually occur	task speed hides end-to-end work
People	roles time compensation retention and authority	who performs and owns judgement	partner effort is unrecorded or unstable
Technology	agents models tools identity data and vendors	what is owned transferable and controlled	platform cost or dependency is omitted
Quality and liability	corrections complaints claims insurance and regulation	what failure costs and who bears it	sparse history understates tail exposure
Economics	revenue price labour write-off and cash by cohort	what contribution is sustainable	reported margin excludes supervision and rework

Proposed structure; target-specific legal accounting professional technology and commercial review is required.

Table 2 Workflow evidence from instruction to cash

Stage	Evidence	Metric	Failure indicator
Scope	instruction conflict check engagement and price	conversion and scope change	work begins without authority or economic baseline
Produce	source data task logs and versions	production hours and first-pass time	demonstration data differs from production
Review	comments approval consultation and exceptions	review intensity yield and delay	silent partner repair or routine approval
Deliver	signed output transmission and client response	acceptance and correction	output sent before accountable approval
Bill	invoice support write-off and dispute	realised price and recovery	hours fall and revenue falls faster

Stage	Evidence	Metric	Failure indicator
Collect	bank receipt ageing and credit	cash conversion and days	reported revenue is not collected

Proposed evidence chain; every material workflow should be sampled and reconciled.

Table 3 Hypothetical hours per one hundred accepted jobs

Case	Production hours	Review hours	Rework hours	AI operations hours	Total hours	Saving versus baseline
Baseline	1000	120	80	0	1200	0.0%
Initial AI	650	260	140	80	1130	5.8%
Controlled AI	650	200	90	70	1010	15.8%

Management assumptions; the illustration separates task automation from total accepted-work economics.

Table 4 Hypothetical AI economics across four target types

Target type	Production saving	Added review rework and operations	Principal value driver	Principal risk
Accounting advisory	25%	12%	evidence assembly and consistent workpapers	insufficient professional judgement and documentation
Legal services	32%	20%	research review and drafting capacity	confidentiality accuracy fees and supervision
Engineering consultancy	22%	18%	retrieval calculation and package coordination	safety consequence versioning and sign-off
Technology services	38%	17%	coding testing support and operations	security authority service levels and rollback

Management assumptions expressed as shares of baseline labour cost; actual economics require cohort data.

Table 5 Hypothetical sustainable EBITDA normalisation

Item	Amount	Evidence required	Potential treatment
Reported EBITDA	9.6	ledger management accounts and quality of earnings	starting point only
Hidden supervision	-2.4	partner activity approval and version history	sustainable operating cost
AI platform and control	-1.2	licences cloud people security and assurance	sustainable operating cost
Partner retention economics	-1.0	compensation rollover and succession	operating or transaction allocation
Recurring rework	-0.6	defects write-offs and correction hours	sustainable operating cost
Client remediation	-0.4	concessions support and complaints	sustainable or contingent cost
Sustainable EBITDA	5.0	reconciled workflow economics	valuation input subject to diligence

Table 6 Transaction response by finding

Finding	Value effect	Potential deal response	Integration measure
Verified accepted productivity	supports capacity and contribution	base value or evidence-weighted synergy	cost and contribution per accepted unit
Unrecorded supervision	reduces sustainable margin	price adjustment earn-out metric or funded redesign	review time exception and reviewer span
Client restriction	limits deployment or pricing	consent condition disclosure or perimeter change	approved client cohort and retention
Quality or claim exposure	creates cash and reputation risk	indemnity escrow insurance or holdback	defects claims recovery and closure
Key-person dependency	threatens revenue and authority	rollover retention governance or deferred value	client retention and accountable capacity
Transfer blocker	prevents lawful or controlled operation	closing condition delayed scope or no-go	verified permission consent and operating acceptance

Proposed framework; actual instruments require current legal tax accounting regulatory insurance and financial advice.

Table 7 Evidence gates for releasing transaction value

Gate	Minimum evidence	Decision	Measure after release
Workflow truth	sampled instruction-to-cash trace and reconciled hours	accept or revise baseline	accepted units cycle and cash
Client acceptance	contract authority consent and renewal evidence	include eligible revenue	retention realised price and disputes
Quality control	reviewer authority complete records and tested exceptions	approve operating design	review yield defects and claims
Sustainable economics	full supervision technology rework and retention cost	set valuation earnings	contribution per accepted unit
Integration readiness	people data systems vendor and client dependencies	approve cohort migration	service continuity and control exceptions
Synergy release	implemented action accepted output and cash evidence	recognise or defer value	recurring cash and residual risk

Proposed governance; board-approved thresholds should be specific to the service and transaction.

26 Establish board governance and reporting

The board should receive a compact decision report that connects commercial performance, workflow, quality, people, technology and cash. Core measures include accepted units, cycle time, realised price, total hours, review intensity, rework, defects, claims, client retention, key-person capacity, AI operating cost, contribution and collections. Measures should be segmented by service and cohort.

Governance should define accountable owners for professional quality, technology, data, security, finance, people and client outcomes. The investment committee approves value and risk. Operating leaders manage delivery. Independent risk and assurance functions challenge evidence where appropriate. Escalation thresholds should cover quality deterioration, permission change, client objection, key-person loss and correlated incidents.

Management should distinguish observations from estimates and forecasts. Scenario values should retain assumptions and approval dates. The board should receive contradictory evidence and unresolved items, not only the central case. This supports professional scepticism and a defensible transaction record.

Compensation and contingent consideration should use measures that management can influence without weakening quality. Accepted contribution and cash can balance growth and control better than automated volume. Any measure can be gamed; definitions, audit rights and board judgement remain necessary.

27 Limitations and conclusion

The framework does not provide legal, accounting, tax, regulatory, insurance, engineering or investment advice. Professional obligations, AI regulation, employment restrictions and transaction treatment vary by service and jurisdiction. The hypothetical figures illustrate the method and do not describe a company, market benchmark or expected result. Public disclosures and standards cannot replace access to target contracts, systems, records and people.

The analysis is also limited by measurement. Time records can be incomplete, quality problems can appear after delivery and client behaviour can change. AI systems, vendor terms and regulation continue to evolve. Buyers should update evidence through signing, closing and integration. A strong historical cohort does not guarantee future performance when scale, workflow or model changes.

The transaction conclusion is operational. An AI-enabled professional-services roll-up creates sustainable value when the combined group delivers more accepted work and cash with reproducible quality, accountable judgement and a fully costed control model. Production-hour reduction is one input. Review, rework, pricing, client permission, people, data, liability and integration determine whether that input becomes margin.

The buyer should therefore underwrite workflow truth rather than a technology narrative. It should normalise sustainable EBITDA, design evidence-based synergies, protect identified risks and migrate client work by controlled cohorts. This approach allows investment committees to reward genuine productivity while recognising the senior judgement and trust on which professional-services value depends.

Frequently asked questions

Why can AI reduce production hours without expanding margin

Review, rework, AI operations, client explanation and professional accountability can absorb the production saving. The buyer should measure total cost per accepted and collected work unit using observed workflow cohorts.

Which metric should replace utilisation

Utilisation remains useful for workforce planning. It should be supplemented by accepted units, cycle time, review intensity, rework, realised price, contribution per accepted unit and cash collection.

How should hidden supervision be identified

Reconcile time records with calendars, approval logs, version histories, consultations, quality files and representative workflow traces. Senior effort recorded as administration, sales or innovation should be assigned where it supports delivery.

Can fixed fees preserve the AI productivity benefit

Fixed fees can preserve part of the benefit when scope, quality and client acceptance are stable. They can also transfer rework and outcome risk to the firm. Contract terms, renewals, disputes and collected cash provide the evidence.

How should an earn out address AI economics

Use clearly defined measures such as collected revenue, client retention and contribution after supervision and control cost. Include quality and key-person protections. Transaction counsel and accountants should design the terms.

What is the most important integration control

Migrate by workflow cohort while preserving accountable authority, client terms, evidence and rollback. Release wider deployment only after accepted throughput, quality, contribution and cash are demonstrated.

Does one common AI platform suit every professional service

Common identity, logging, evaluation and incident components can help. Approval, data, client and professional rules often require configurable service and jurisdiction controls.

Which Matchpoint practice is relevant to this work

This research relates to Matchpoint Partners mergers and acquisitions, commercial due diligence, valuation, transaction structuring, post-merger integration, strategy and execution, and technology-enabled operating-model work for buyers, investors and professional-services firms.

References

- [1] IFRS Foundation. IFRS 3 Business Combinations. <https://www.ifrs.org/issued-standards/list-of-standards/ifrs-3-business-combinations/>
- [2] IFRS Foundation. IAS 36 Impairment of Assets. <https://www.ifrs.org/issued-standards/list-of-standards/ias-36-impairment-of-assets/>
- [3] IFRS Foundation. IAS 38 Intangible Assets. <https://www.ifrs.org/issued-standards/list-of-standards/ias-38-intangible-assets/>
- [4] IFRS Foundation. IFRS 13 Fair Value Measurement. <https://www.ifrs.org/issued-standards/list-of-standards/ifrs-13-fair-value-measurement/>
- [5] IFRS Foundation. IAS 37 Provisions Contingent Liabilities and Contingent Assets. <https://www.ifrs.org/issued-standards/list-of-standards/ias-37-provisions-contingent-liabilities-and-contingent-assets/>
- [6] IFRS Foundation. IASB consults on proposals to improve reporting of acquisitions. <https://www.ifrs.org/news-and-events/news/2024/03/iasb-consults-on-proposals-to-improve-reporting-of-acquisitions/>
- [7] 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>
- [8] National Institute of Standards and Technology. Artificial Intelligence Risk Management Framework AI RMF 1.0. <https://nvlpubs.nist.gov/nistpubs/ai/NIST.AI.100-1.pdf>
- [9] National Institute of Standards and Technology. Artificial Intelligence Risk Management Framework Generative Artificial Intelligence Profile. <https://nvlpubs.nist.gov/nistpubs/ai/NIST.AI.600-1.pdf>
- [10] National Institute of Standards and Technology. AI Agent Standards Initiative. <https://www.nist.gov/artificial-intelligence/ai-agent-standards-initiative>
- [11] National Cybersecurity Center of Excellence. Accelerating the Adoption of Software and AI Agent Identity and Authorization. <https://www.nccoe.nist.gov/publications/other/accelerating-adoption-software-and-ai-agent-identity-and-authorization-concept>
- [12] National Institute of Standards and Technology. Cybersecurity Framework 2.0. <https://nvlpubs.nist.gov/nistpubs/CSWP/NIST.CSWP.29.pdf>
- [13] National Institute of Standards and Technology. Security and Privacy Controls for Information Systems and Organizations SP 800 53 Revision 5. <https://csrc.nist.gov/pubs/sp/800/53/r5/upd1/final>
- [14] OWASP Foundation. Excessive Agency. <https://genai.owasp.org/llmrisk/llm06-excessive-agency/>
- [15] OWASP Foundation. Agentic AI Threats and Mitigations. <https://genai.owasp.org/resource/agentic-ai-threats-and-mitigations/>
- [16] MITRE. Adversarial Threat Landscape for Artificial Intelligence Systems. <https://atlas.mitre.org/>

- [17] International Organization for Standardization. ISO IEC 42001 Artificial Intelligence Management System. <https://www.iso.org/standard/81230.html>
- [18] European Commission. EU AI Act Article 14 Human Oversight. <https://ai-act-service-desk.ec.europa.eu/en/ai-act/article-14>
- [19] European Commission. EU AI Act Article 4 AI Literacy. <https://ai-act-service-desk.ec.europa.eu/en/ai-act/article-4>
- [20] European Commission. EU AI Act Article 26 Obligations of Deployers. <https://ai-act-service-desk.ec.europa.eu/en/ai-act/article-26>
- [21] American Bar Association. Formal Opinion 512 Generative Artificial Intelligence Tools. https://www.americanbar.org/content/dam/aba/administrative/professional_responsibility/ethics-opinions/aba-formal-opinion-512.pdf
- [22] Solicitors Regulation Authority. Compliance tips for solicitors regarding the use of AI and technology. <https://www.sra.org.uk/solicitors/resources/innovate/compliance-tips-for-solicitors/>
- [23] Solicitors Regulation Authority. Misuse of AI warning notice. <https://www.sra.org.uk/solicitors/guidance/misuse-ai/>
- [24] Financial Reporting Council. AI in Audit. <https://www.frc.org.uk/library/standards-codes-policy/audit-assurance-and-ethics/guidance/ai-in-audit/>
- [25] Financial Reporting Council. AI in Audit illustrative example and documentation guidance. https://media.frc.org.uk/documents/AI_in_Audit.pdf
- [26] International Auditing and Assurance Standards Board. Technology Position. <https://www.iaasb.org/focus-areas/technology>
- [27] International Auditing and Assurance Standards Board. Technology Quality Management Workstream. <https://www.iaasb.org/consultations-projects/technology-quality-management-workstream>
- [28] International Auditing and Assurance Standards Board. Technology Quality Management Roundtables Outcomes and Next Steps. <https://www.iaasb.org/publications/technology-quality-management-roundtables-outcomes-and-next-steps>
- [29] International Auditing and Assurance Standards Board. International Standard on Quality Management 1. <https://www.iaasb.org/publications/international-standard-quality-management-isqm-1-quality-management-firms-perform-audits-or-reviews>
- [30] International Auditing and Assurance Standards Board. ISA 220 Revised Quality Management for an Audit of Financial Statements. <https://www.iaasb.org/publications/international-standard-auditing-220-revised-quality-management-audit-financial-statements>
- [31] International Ethics Standards Board for Accountants. Technology related revisions to the Code. <https://www.ethicsboard.org/publications/technology-related-revisions-code>
- [32] Public Company Accounting Oversight Board. AS 1201 Supervision of the Audit Engagement. <https://pcaobus.org/oversight/standards/auditing-standards/details/AS1201>
- [33] Public Company Accounting Oversight Board. AS 1105 Audit Evidence. <https://pcaobus.org/oversight/standards/auditing-standards/details/AS1105>
- [34] Public Company Accounting Oversight Board. Staff Update on Outreach Activities Related to the Integration of Generative Artificial Intelligence in Audits. <https://pcaobus.org/resources/staff-publications/staff-update-on-outreach-activities-related-to-the-integration-of-generative-artificial-intelligence-in-audits>
- [35] Accenture plc. Fiscal 2025 Form 10 K. <https://www.sec.gov/Archives/edgar/data/1467373/000146737325000217/acn-20250831.htm>
- [36] Accenture plc. Fourth Quarter and Full Year Fiscal 2025 Results. <https://www.sec.gov/Archives/edgar/data/1467373/000146737325000213/q4fy25earnings8-kexhibit.htm>
- [37] CBIZ Inc. 2025 Form 10 K. <https://www.sec.gov/Archives/edgar/data/944148/000094414826000038/cbiz-20251231.htm>
- [38] CBIZ Inc. 2024 Annual Report. <https://www.sec.gov/Archives/edgar/data/0000944148/000119312525072379/d855075dars.pdf>
- [39] Thomson Reuters Corporation. Annual Reports. <https://ir.thomsonreuters.com/financial-information/annual-reports>
- [40] UK Office for National Statistics. Management consultancy activities industry data. <https://www.ons.gov.uk/businessindustryandtrade/business/activitysizeandlocation>
- [41] US Bureau of Labor Statistics. Professional scientific and technical services productivity. <https://www.bls.gov/productivity/highlights/professional-services.htm>
- [42] Organisation for Economic Co-operation and Development. OECD Employment Outlook 2023 Artificial Intelligence and the Labour Market. <https://www.oecd.org/employment-outlook/2023/>
- [43] UK Competition and Markets Authority. Merger Assessment Guidelines. <https://www.gov.uk/government/publications/merger-assessment-guidelines>
- [44] US Department of Justice and Federal Trade Commission. 2023 Merger Guidelines. <https://www.justice.gov/atr/2023-merger-guidelines>
- [45] Federal Trade Commission. Noncompete Rule status. <https://www.ftc.gov/legal-library/browse/rules/noncompete-rule>
- [46] UK Information Commissioner's Office. AI and Data Protection Risk Toolkit. <https://ico.org.uk/for-organisations/uk-gdpr-guidance-and-resources/artificial-intelligence/guidance-on-ai-and-data-protection/ai-and-data-protection-risk-toolkit/>
- [47] UK National Cyber Security Centre. Guidelines for Secure AI System Development. <https://www.ncsc.gov.uk/collection/guidelines-secure-ai-system-development>
- [48] Cybersecurity and Infrastructure Security Agency. Secure by Design. <https://www.cisa.gov/securebydesign>
- [49] OpenID Foundation. Identity Management for Agentic AI. <https://openid.net/identity-management-for-agentic-ai/>
- [50] National Institute of Standards and Technology. Privacy Framework. <https://www.nist.gov/privacy-framework>

About the Author

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.

<https://www.linkedin.com/in/ckadya/>

<https://www.matchpoint-partners.com/team/ck-adya.html>

This paper is part of a continuing series on the structure of private and alternative markets. The views expressed are the author's own. The paper is for information only, describes market structure in general terms, and does not constitute investment, legal, tax or regulatory advice or a recommendation in respect of any security, vehicle or counterparty.