

XAMUN POSITION PAPER NO. 5

The firm that sells hours cannot buy back time.

An AI strategy for the accounting and advisory practice — what to re-derive first, what never to hand to a language model, and what happens to the billable hour.

Arup Maity, Founder and CEO, Xamun Technologies

Version 1.0 · September 2026

Summary. The accounting practice is the purest case of the problem this series has been describing. Its product is the hour. Its leverage model is a pyramid of juniors doing routine work under partner review. Its risk is professional liability for a wrong number. And the technology now being sold to it is a language model — a machine that is fast, fluent, and structurally unreliable about numbers.

This paper applies the method of Paper No. 4 to a practice of 50 to 500 people: bookkeeping, compliance, audit, payroll, advisory. It argues that a language model must never be the last word on any figure a client relies on; that the firm's AI strategy is therefore a question of which operations to rebuild around a deterministic core, with generation placed only where language is the work; that the first operation should almost always be the monthly close for recurring clients, or the document intake that feeds it; and that a firm which re-derives its operations will find its pricing model has been re-derived with them.

It assumes the two-minds architecture set out in Paper No. 1 and the re-derivation method of Paper No. 2, and does not repeat either in full.

1. The firm's particular trap

Every professional services firm sells time, but the accounting practice sells it with unusual precision. The hour is the unit of billing, the unit of staff evaluation, the unit of capacity planning and the unit of profitability analysis. Realisation, utilisation, leverage ratio, recovery rate — the firm's entire management vocabulary is built on the hour.

AI attacks the hour directly. Not the firm, not the client relationship, not the professional judgement — the hour. A task that took a junior four hours takes a machine four minutes. The firm that has priced that task at four hours now faces a choice it has not had to make before: charge the client four hours for four minutes of work, and wait for a competitor to tell them; or charge for four minutes, and watch the leverage model collapse; or find a third unit of sale.

Most firms have responded to this by not responding. They have bought AI-enabled tax research tools and document summarisers, encouraged staff to use them, and left the pricing model untouched. The tools save hours. The hours are quietly not billed, or are billed and eventually challenged. Margin leaks from the bottom of the pyramid, where it was always made, and the partners see a slow decline in recovery that they attribute to fee pressure.

This is the licence strategy from Paper No. 4, and its failure is the same: individual productivity that never aggregates into an operational change, because the operation and its pricing were never redesigned to absorb it.

The firm cannot use AI to do the same operation faster and expect to keep the margin. It must use AI to do a different operation — one the client pays for by the outcome, not by the hour it no longer takes.

2. What an accounting practice actually is

A practice of 50 to 500 people is, operationally, a small number of recurring engines running on a calendar. The calendar is fixed by regulators and clients; the engines are fixed by habit.

The engines

Engine	What it does, and where the friction sits
Client onboarding	Engagement letter, KYC/AML checks, entity setup, chart of accounts, access to client systems. Friction: document chasing, inconsistent risk assessment, the file that is never quite complete.
Document intake	Bank statements, invoices, receipts, contracts, payroll data arriving by email, portal, WhatsApp and shoebox. Friction: chasing clients, manual classification, re-keying, the missing document found at month end.
Bookkeeping and coding	Transactions classified to the chart of accounts, VAT-coded, matched to source documents. Friction: volume, judgement calls on ambiguous items, reviewer corrections that are never fed back.

Engine	What it does, and where the friction sits
Reconciliation and close	Bank, intercompany, control accounts, accruals, prepayments, fixed assets, the trial balance. Friction: waiting on documents, waiting on the client, waiting on review; the close that takes twenty days.
Tax compliance	VAT returns, corporate tax computations and filings, withholding, transfer-pricing documentation where applicable. Friction: deadline concentration, data pulled from a close that is not finished, the rule that changed.
Audit and assurance	Planning, risk assessment, fieldwork, sampling, working papers, review, reporting. Friction: evidence gathering, documentation to standard, review notes cycling between senior and partner.
Payroll	Monthly computation, statutory filings, payslips, end-of-service accruals. Friction: late inputs, rule complexity, the error that reaches an employee.
Client query and advisory	The stream of questions — can I deduct this, what happens if we restructure, why is the VAT figure different. Friction: partner time consumed by routine questions; inconsistent answers from different staff.
Work-in-progress, billing and collection	Timesheets, WIP review, write-offs, invoicing, chasing. Friction: WIP that ages, write-offs that are decided late, collections that take ninety days.

Every one of these is countable — closes per month, returns per period, audits per year, payslips per run, invoices issued. Every one has a cycle time, an error rate and an owner. They are exactly the kind of operation Paper No. 4 says a strategy should be a list of.

The pyramid

Sitting across the engines is the leverage structure: juniors do the instances, seniors review, managers review the reviewers, partners sign. This structure exists for two reasons — it is how the profession trains its people, and it is how a wrong number is caught before it reaches a client or a regulator. Both reasons are real. Neither requires that the first draft of every instance be produced by a human.

The pyramid is where AI lands hardest, because the routine instance work at its base is the most automatable and the most billed. A firm cannot simply remove the base and keep the top; it must redesign what the base does and how the top reviews it. That redesign is the strategy.

Operational gravity in a practice

The firm's unrecorded knowledge lives in its seniors and managers: which client's bookkeeper always mis-codes intercompany, which director's expense claims need a second look, how a particular regulator interprets a particular provision, which working paper format this partner will sign without notes. It is held in people, in review comments, and in a thousand client-specific email threads. Any system that ignores it will be corrected by hand, silently, in the same way the firm's previous practice-management software was — and will then be blamed for not working.

The first job of the firm's AI strategy is to get that knowledge out of heads and into a memory that a process can read. That is not a knowledge-management project. It is a by-product of re-deriving the operations that use the knowledge.

3. Why the chat-first approach fails in a practice

The accounting practice has been an early and enthusiastic adopter of language-model chat tools, and has been early to discover their limits. The limits are worth stating precisely, because they determine where generation can and cannot be placed.

A language model does not know what a number is

A language model predicts text. It does not compute. Asked for a VAT figure, it produces a plausible-looking figure. Asked to reconcile a bank statement, it produces a plausible-looking reconciliation. Asked which depreciation method applies, it produces a plausible-looking answer, which is sometimes right. Plausibility is the product; correctness is incidental. In most industries this is a nuisance. In a firm whose liability insurance is priced on getting numbers right, it is disqualifying.

This is not a criticism of the technology. It is a description of what it is for. Language models are extraordinary at language: reading a contract and extracting the payment terms, classifying an invoice by what it describes, drafting a management letter from a list of findings, explaining a tax position in plain words to a client. None of those is a calculation. All of them are work a practice does every day.

Confidentiality and residency

A client's general ledger is among the most sensitive documents the client possesses. Practices that have allowed staff to paste client data into public chat tools have taken a risk they would not accept in any other form, and most have now issued policies forbidding it. The policy is correct, but it leaves the underlying problem in place: the useful thing the model can do requires seeing the data. The resolution is architectural, not procedural — the model must run where the data is permitted to be, with what it sees and what it retains under the firm's control.

Professional standards do not have a chat exception

Audit documentation standards require that an experienced auditor with no prior connection to the engagement can understand the procedures performed, the evidence obtained, and the conclusions reached. A working paper that says "the model suggested this" does not meet that bar. A firm's quality management system requires that the firm can show how its judgements were made and reviewed. An output from a prompt, reviewed by a junior, does not meet that bar either. The standards are not hostile to automation; they are hostile to unexplained conclusions. Automation that explains itself meets them more reliably than a tired senior at eleven o'clock in March.

The rule for a practice is simple and absolute: a language model is never the last word on a figure, a filing position, or an audit conclusion. It proposes; something deterministic disposes; a named human signs.

4. Two minds in a practice

The architecture this series has described applies to an accounting practice with unusual cleanliness, because the profession has already spent a century separating the things that must be deterministic from the things that require judgement.

The deterministic mind

Everything the profession has already codified belongs here, and runs first:

- The chart of accounts and its mapping rules for each client.
- VAT and corporate tax rules: rates, thresholds, exemptions, the treatment of specific supply types, filing periods and deadlines.
- Materiality thresholds, sampling parameters and risk-assessment logic for audit engagements.
- Reconciliation logic: matching rules, tolerance bands, ageing, the conditions under which an item is an exception.
- Payroll computation, statutory contribution rules, end-of-service formulae.
- Review and sign-off rules: which items require senior review, which require partner sign-off, which may be released unattended.
- The audit trail: who or what did what, when, on the basis of which rule, and what the prior state was.

This mind decides. Every transaction, every computation, every filing figure passes through it, and if a rule is violated the instance stops, regardless of what any other component proposed. Its logic can be printed, tested against a known case, and shown to a regulator. It is the same every time.

The governed language model

Everything that is language belongs here, and is never final:

- Reading a document — bank statement, invoice, contract, board minute — and extracting the fields the deterministic mind needs.
- Proposing a classification for a transaction the rules cannot resolve, with a stated reason.
- Drafting: management letters, audit findings narrative, tax position memoranda, client explanations, engagement letters from templates.

- Summarising a client's question and routing it, or drafting an answer from the firm's recorded positions for a human to check.
- Explaining a figure in plain language: why this quarter's VAT is higher, what changed in the computation.

Every output of this mind is a proposal. The deterministic mind checks it — does the extracted total match the document total, does the proposed classification comply with the client's rules, does the drafted letter reference figures that exist in the ledger — before anything is written to a system of record or sent to a client. Where the check cannot be made deterministically, a human is called.

The graph memory

Between the two minds sits the firm's operational memory, scoped to the operation and extended as operations are added:

- Clients, entities, group structures, directors and signatories.
- Obligations: which filings are due for which entity, when, under which rule set, and their current status.
- Precedents: how this client's ambiguous items were classified last time, and by whom; which filing positions the firm has taken and on what basis.
- Lineage: which source document supports which ledger entry supports which figure in which filing.
- Review history: what the reviewer changed, so that the change becomes a rule rather than a correction.

This memory is what converts the firm's operational gravity into something a process can read. The senior who knows that a particular client always mis-codes intercompany does not have to remember it; the graph does, and the deterministic mind applies it. The manager's review note is not lost in a working paper; it becomes a precedent the next instance consults.

The human

A named professional signs every filing, every audit opinion and every client deliverable, as now. What changes is what they are signing: not a document assembled by a junior over three weeks, but an instance that passed every deterministic check, whose exceptions were resolved by a named person with a recorded reason, and whose lineage from source document to final figure can be walked in seconds. The signature is more defensible, not less.

5. The constraint inventory for a firm

Paper No. 4 described the constraint inventory as the first act of strategy. In a practice it can be run in two weeks by the managers who run the engines, with a partner present, and it should

produce a table like the following. The numbers below are illustrative of the kind of finding an inventory produces, not benchmarks; the firm must measure its own.

Operation	Volume	Cycle time	Constraint and its cost
Monthly close (recurring clients)	Per month, per client	Days from period end to reviewed TB	Waiting on client documents; reviewer corrections; managers absorbing junior rework. Cost: unbilled hours, late VAT data, partner time in March.
Document intake	Documents per month	Days from period end to complete file	Chasing by email; manual classification; re-keying. Cost: the close cannot start until this finishes.
VAT returns	Returns per period	Days before deadline that the return is ready	Depends on an unfinished close; deadline concentration; late-found errors. Cost: penalties, overtime, client trust.
Corporate tax computation	Computations per year	Weeks from year-end accounts to filing	Manual adjustment schedules; rule interpretation held in one person's head. Cost: partner review time, inconsistency across clients.
Audit fieldwork and working papers	Engagements per year	Fieldwork days; review cycles	Evidence requests cycling with the client; documentation to standard; review notes. Cost: over-run against budget, write-offs.
Payroll	Payslips per run	Hours from input to release	Late inputs; manual rule application. Cost: errors reaching employees; rework.
Client queries	Queries per week	Hours to answer	Partner and manager interruption; inconsistent answers. Cost: senior time consumed by routine.
WIP and billing	Invoices per month	Days from work done to cash	Late WIP review; write-off decisions deferred; slow collection. Cost: working capital, realisation.

The inventory will show a dependency the firm already knows and has never written down: almost every downstream operation waits on the close, and the close waits on intake. That dependency is why the choice of first operation is less open than it appears.

6. The first operation

Applying the four tests from Paper No. 4 — countable, constrained, reversible, owned — to the inventory produces, for most practices, one of two answers.

The monthly close for recurring clients

The close passes every test. It is countable: closes per month, days per close. It is structurally constrained: the delay is in waiting and rework, not in any individual. It is reversible: if the re-derived close fails on a client, the firm does the close the old way for that client that month, and nothing has reached a regulator. It has a natural owner: the manager who runs the bookkeeping and compliance engine.

It is also the operation with the highest leverage, because everything downstream consumes it. A close that finishes in five days instead of fifteen means VAT returns prepared with a margin, corporate tax computations starting earlier, advisory conversations held while the numbers are still relevant, and billing raised while the work is fresh. The firm does not have to re-derive VAT to improve VAT; it has to fix what VAT waits for.

Or: document intake

For a firm whose close is bottlenecked entirely on receiving client documents, intake may come first — or, more often, the two are re-derived as one operation, because the close cannot be inverted without inverting its trigger. The re-derived intake watches for documents to arrive, classifies them on arrival, matches them to expected items, and chases the client for what is missing — automatically, on a schedule, with the firm's tone. The close begins on the first document, not the last.

What not to start with

Audit is the operation partners most often propose first, because it carries the highest fees and the most visible pain. It should not be first. It is the least reversible — an audit opinion, once issued, cannot be quietly redone — and its governance requirements are the most demanding. A firm that has re-derived its close and its intake, and has six months of exception logs showing the two-minds shape working on client data, is a firm that can approach audit fieldwork with earned confidence. A firm that starts there is a firm that will run a pilot.

Client-facing chat — a portal where clients ask questions of a model — should also wait. The firm's recorded positions must exist in the graph before a model can answer from them; until then, the model answers from the internet, and the firm signs its name to it.

7. Baseline before build

Before anything is designed, the close is measured over two or three cycles, per client, and the numbers are written down:

- **Cycle time:** period end to complete file; complete file to reviewed trial balance; reviewed trial balance to client pack. As distributions. The clients in the long tail share a cause, and the cause is the design brief.

- **Error rate:** the proportion of junior-prepared entries changed at review; the proportion of closes reopened after delivery; the proportion of VAT returns amended.
- **Throughput:** closes completed per preparer per month; review hours per close; chase emails sent per client per month.
- **Realisation:** hours recorded against the close versus hours billed. This is the number the partners actually care about, and it should be on the page from the start.

No projection is made. The firm does not estimate that AI will save forty percent of bookkeeping time, because no such estimate can be held to account. It records what the close costs today, builds the re-derived close, and reports the same numbers next quarter. Paper No. 1 sets out why this stance is non-negotiable; a practice, of all businesses, should need no persuading that a number without a basis is not a number.

8. Re-deriving the close

Paper No. 2 describes three design moves. Applied to the close, they look like this.

Invert the trigger

Today the close starts when the period ends and a preparer picks up the file. In the re-derived operation the close starts when data arrives: a bank feed line, an uploaded invoice, a payroll file. Each item is classified, matched and posted on arrival, against the client's rules, with exceptions queued. By period end most of the close has already happened. The preparer's job is no longer to do the close; it is to clear the exception queue and confirm the result.

Exception-first

Nobody looks at a transaction that matched its document, complied with the client's coding rules, and fell within tolerance. Humans see only the exceptions: the unmatched item, the ambiguous classification, the intercompany balance that does not agree, the invoice whose total does not equal its extracted lines. The exception queue is ordered by materiality and by the confidence of the proposal. The reviewer's decisions on exceptions are recorded with a reason and become precedents, so that the same exception does not recur.

Place every decision

Every decision inside the close is assigned to exactly one of three places: the deterministic mind (rules, matching, computation, control totals), the governed language model (extraction, classification proposals, narrative), or a named human (materiality judgements, unusual items, sign-off). No decision is left to "whoever is doing it." The placement is written down and can be shown to a reviewer, a regulator, or the firm's insurer.

Kill the parallel process

On a named date, for a named set of clients, the old close is no longer permitted. This is the step firms flinch at and the step without which nothing changes. A re-derived close that runs alongside the old one is a pilot; the firm pays for both and trusts neither. The reversibility test guarantees the firm can go back for a client if it must. It should expect not to.

Staged autonomy

Decision classes are promoted separately. Bank-feed matching against an existing invoice, once it has been right for a few hundred instances, runs unattended. Classification of a new supplier's first invoice stays proposed-and-reviewed. Journal entries above a materiality threshold stay with a human indefinitely. The manager who owns the close sets the pace, reading the exception log weekly, and the pace is visible to the partner who signs.

9. Governance, standards and liability

A practice's AI strategy is subject to a governance test that most industries can defer: a professional body, a regulator and an insurer will each, eventually, ask how the firm's automated work is controlled. The two-minds shape answers each of them from the same construction.

The professional body and the audit regulator

Documentation standards ask that procedures, evidence and conclusions be recorded so that an experienced reviewer can reconstruct them. A re-derived operation records more than a human one does: every rule applied, every proposal made, every exception and its resolution, every reviewer's change and reason. The lineage from source document to filed figure is walkable. The firm's quality management system gains a control it could not previously implement — a complete, contemporaneous record of every judgement inside every engagement.

What the standards do not accept is an unexplained conclusion, and this is why the placement of decisions matters. A conclusion the deterministic mind reached is explained by its rule. A conclusion a human reached is explained by their recorded reason. A conclusion the language model reached, unreviewed, would be unexplained — and so the design never lets it be a conclusion.

Data protection and residency

Data-protection law in most jurisdictions where the firm operates governs automated decision-making about individuals, requires impact assessment for higher-risk processing, and assigns accountability for how personal data is handled. Payroll and any engagement touching individual directors or employees is inside that scope. The firm should be able to say, for every operation, where the data is processed, which model component sees it, what is retained, and on what legal basis. The architecture makes this a configuration decision rather than an audit finding: the

deterministic mind and the graph run where the firm's data must reside; language-model tasks are routed to in-region, open-weight models by default, with external frontier models used only for tasks that are both non-sensitive and demonstrably improved by them.

The firm should also be precise about what regulation actually requires. There is no binding AI-specific statute in the UAE, and no AI self-assessment deadline; a vendor asserting one should be asked for the citation. The obligations that apply are those of data protection, professional standards, and sector regulators — and they apply equally to a spreadsheet macro and a language model.

The insurer

Professional indemnity underwriters are beginning to ask about AI use, and the firm that answers "we have a policy" is in a weaker position than the firm that answers "no figure a client relies on is produced by a language model; every computation is rule-based and logged; every exception is resolved by a named person; here is the log." The second answer describes a firm with lower error rates and better records than it had before. It should be priced accordingly, and the firm should make the case.

10. What happens to the hour

A firm that re-derives its close will find, within two quarters, that the close for a typical recurring client consumes a fraction of the hours it did. The firm then faces the question this paper opened with, and it should face it deliberately rather than by drift.

The three responses

The first response is to keep billing the old hours. This works until a client asks, or a competitor prices the same close at a third of the fee. It is a short-term margin gain and a long-term relationship risk, and it is what most firms will do by default.

The second response is to bill the new hours. This is honest and ruinous; it hands the entire gain to the client and dismantles the leverage model.

The third response is to stop selling the hour for that operation. The close becomes a product: a fixed monthly fee per entity, at a defined service level — trial balance reviewed and delivered by the fifth working day, VAT return prepared by the tenth — with the firm keeping the efficiency and the client getting the certainty. The unit of sale is the closed month, not the hours inside it. This is the outcome model, and it is the same shift the technology industry is making from the man-day to the delivered result.

Why the third response is available to a firm now

Fixed-fee bookkeeping has existed for years, and firms have been wary of it because it transfers the risk of a difficult client onto the firm. A re-derived close changes the risk. The firm can see, per

client, how many exceptions a close generates and how long they take to clear; it can price the difficult client differently, on evidence, and it can show the client why. The exception log is a pricing tool. A firm that has it can price by outcome with confidence; a firm that does not is guessing.

The pyramid after

The leverage model does not disappear; it changes shape. Fewer juniors do instance work; more of them clear exceptions and maintain rules — which is, incidentally, a better way to learn what the numbers mean than re-keying invoices ever was. Seniors review exception decisions rather than every entry. Managers own operations and read logs. Partners sign more engagements with more confidence in less time, and spend the difference on advisory, which is where the hour still has a defensible value because judgement is the product.

Headcount is not the design objective and should not be. The firm that re-derives its close will, in most cases, redeploy rather than reduce, because the constraint that was absorbing junior time was also suppressing work the firm could have sold. Growth without proportional hiring is the outcome; it arrives around the third operation, and it arrives as a discovery rather than a target.

11. Sequence: four quarters

For a practice of this size, one operation per quarter is achievable and sufficient. A representative sequence:

1. **Quarter one — intake and close for recurring clients.** Constraint inventory, baseline over two cycles, re-derived intake and close switched on for a first cohort at proposed-and-reviewed autonomy. Parallel process killed for that cohort at the end of the quarter. Graph seeded with client entities, coding rules and precedents.
2. **Quarter two — VAT and periodic compliance.** The return is now generated from a close that finishes early. Deterministic mind holds the rule set and deadlines; graph holds obligations per entity; language model drafts the client cover note. Close rolled out to remaining recurring clients.
3. **Quarter three — corporate tax and payroll.** Adjustment schedules rule-based and lineage-linked to the close; the person whose head held the interpretations writes them down as rules and precedents. Payroll re-derived on the same shape, with the strictest human sign-off on any change to an individual's pay.
4. **Quarter four — audit fieldwork and client query.** With three operations of exception logs behind it, the firm approaches evidence gathering, sampling and working-paper drafting on the same shape, with a partner owning autonomy levels. Client queries routed to a governed model answering only from the firm's recorded positions, with every answer reviewed before release until the log earns otherwise.

By the end of the year the firm has a graph that holds its clients, their obligations, its precedents and its lineage; four operations running with visible autonomy levels; and a pricing conversation it is having on its own terms. It has not built a data platform, hired a data scientist, or run a pilot.

12. Vendors

The firm will be approached by three kinds of vendor, and should recognise each. Its practice-management and ledger software will add AI features; these are worth taking where the operation sits entirely inside that system, and insufficient where it does not — which is the case for intake, close and anything touching more than one client system. Consultancies and development firms will offer to build, by the hour, against a specification the firm must write; the firm that sells hours should be the first to see the problem with buying them. And outcome-based partners will offer to re-derive an operation and be paid on the countable unit — closes, returns, engagements — with the baseline as the contract.

Whichever it chooses, the firm should read every proposal against the six questions in Paper No. 4, and add one: *show me the decision placement for this operation*. A vendor who cannot say, for the close, which decisions are rules, which are proposals, and which are a named human's, is selling a chat window with a ledger attached.

13. What the managing partner should ask

Before approving any AI initiative in the firm, the managing partner or executive committee should require answers to the following. General answers mean the initiative is not ready.

1. Which operation, and for which cohort of clients?
2. What is its baseline — cycle time, error rate, throughput, realisation — and over how many cycles was it measured?
3. Which manager owns it, and is their evaluation tied to the baseline moving?
4. For every decision in the operation: rule, proposal, or human? Show the placement.
5. Which figures can a language model touch, and what checks them before they reach a ledger, a filing, or a client?
6. Where is client data processed, which components see it, and what is retained?
7. How is every judgement recorded, and can a reviewer reconstruct an engagement from the log alone?
8. On what date does the old process stop for this cohort?
9. What does the fee for this operation become, and when do we tell clients?
10. If we switch it off in month two, what is the cost, and how fast is the return?
11. What is the second operation, and what did the first one's exception log tell us about it?

14. Closing

The accounting practice is not a laggard in AI adoption; it is a bellwether. Its product is the hour, its risk is the wrong number, and its governance is a matter of professional survival — which means the failures of chat-first, pilot-driven, licence-led AI show up in a practice sooner and more expensively than anywhere else. The firms that have already been burned by a confident model producing a plausible figure understand something the rest of the market is still learning.

What follows from that understanding is not caution. It is construction discipline. Put the rules first and let them have the final say. Let the model read, classify and draft, and never let it conclude. Keep one memory that holds what the seniors know. Surface a human at every genuine judgement, and record every one. Start with the close, because everything waits on it. Measure before building, and report the same numbers after. Stop the old process on a date. And when the hours fall, stop selling them — sell the closed month.

A practice that does this has not adopted AI. It has re-derived itself, and it will find that its pricing, its leverage model and its risk position have been re-derived with it. That is what a strategy is for.

The firm that sells hours cannot buy back time. The firm that sells outcomes never needed to.

About this series

This is the fifth paper in a series from Xamun Technologies on the conditions under which AI shows up in a company's P&L. Paper No. 1, *You have AI. It isn't in your P&L.*, sets out the two-minds architecture and the five-level maturity ladder. Paper No. 2, *Don't renovate the work. Re-derive it.*, describes how an operation becomes AI-native. Paper No. 4, *Your AI strategy is a list of tools. It should be a list of operations.*, sets out the method for the 50-to-500-person company that this paper applies to the accounting practice.

Xamun Technologies, licensed in the Dubai International Financial Centre, builds purpose-built operational systems for mid-market and government-related enterprises in the Gulf and Southeast Asia, and is paid on outcomes rather than effort. The author has spent twenty-six years building enterprise software, most of it sold by the man-day, and writes about the end of the hour as someone who used to sell it.

Sources

- MIT NANDA, *The GenAI Divide: State of AI in Business 2025*; McKinsey & Company, *The State of AI* (2025); Gartner agentic-AI forecast (2025) — as cited in Papers No. 1 and No. 4.
- International Standards on Auditing (documentation and quality management requirements), as applicable in the firm's jurisdiction.

- Applicable data-protection legislation in the firm's operating jurisdictions; UAE Federal Tax Authority requirements for VAT and corporate tax filing.
- Xamun Position Papers No. 1, No. 2 and No. 4 (August–September 2026).