

Insights

The System of Work

AI's real disruption to finance isn't a better ERP — it's a new system of work. A practical guide to turning AI's probabilistic intelligence into the deterministic, auditable outcomes finance can sign off on — and to acting now, while the models keep changing.



THE QUIET SHIFT FINANCE HASN'T NAMED YET

For thirty years, finance ran on a clean division of labor: the **ERP was the system of record** — where transactions were stored — and **people armed with spreadsheets were the system of work**, the layer where data was interpreted, reconciled, judged, and turned into a number someone would sign. That second layer was never formally architected. It lived in spreadsheets and in the heads of experienced controllers.

Artificial intelligence is the first technology with a credible claim to operate *inside* that second layer — not as a better record, but as a participant in the work itself. That is the real disruption, and it lands precisely where finance has the least documented process and the highest stakes: judgment, reconciliation, and narrative.

This paper makes a connected argument. The system-of-work shift — not ERP modernization — is what AI changes for finance. The function is still in the infancy of that shift. Finance faces a constraint no other function does as acutely: language models are **probabilistic** by construction, while reporting demands **deterministic, certain, auditable** outcomes — a tension solved by architecture, not by better prompts. The same architecture that delivers auditability also controls cost, answers the “models change every six weeks” paralysis, decides what to build versus buy, informs the ERP you choose, and tells you where on your worklist to start.

The scope is the whole office of the CFO — planning, treasury, spend, and reporting alike — but we go deepest where the constraint is most unforgiving: the controllership and reporting spine. That is deliberate. Every finance workflow decomposes along the same seam —

judgment the model drafts, numbers the artifacts own, sign-off the human keeps. We go deepest on the controllership spine because a pattern that survives an audit survives everything — and Section VI is explicit that deployment starts elsewhere, on the transactional spine, where risk is low and payback is fast.

The conclusion is action-oriented and, we think, freeing: **build the part that does not change — now — and let the models come to you.**

“*The model proposes, the artifact disposes, the human signs.*”



Sandip Khetan

Co-Founder, Global
Head of Accounting &
Reporting Consulting



Shashikant Shenoy

Partner,
Finance Transformation



Gagan Arora

Partner,
Finance Transformation

Written for CXOs of mid-market companies in North America, and for Controllers and CFOs across India and the Middle East. The thesis is global; the “what this means for you” differs by market, addressed directly in Section II.

CONTENTS

A finance leader’s guide, in four movements: the shift, the architecture, the workflows, and how to act. Every entry is a live link.

How to use this paper — and who should read it

This is written for finance leaders, but it is built to be passed down. If four minutes is all you have, the one-page version overleaf is the whole argument. The strategy sections (I–V, IX–X, XI–XIII) are for the CFO, controller, and audit committee deciding whether and how to move. The workflow sections (VI–VIII) are written so that the people who do the work — in payables, the close, FP&A, consolidation, receivables — can find their own function, see what actually changes for them, and absorb where this is heading. If you lead a finance team, read it for the decision; then hand it to your team for the shift.

The One-Page Version	04
THE SYSTEM OF WORK IN FOUR MINUTES — FOR THE BOARD	04
The Shift — And The Constraint	05
FROM SYSTEM OF RECORD TO SYSTEM OF WORK	05
WHERE WE ACTUALLY ARE: THE INFANCY IS THE HEADLINE	07
THE FINANCE-SPECIFIC PROBLEM NO ONE WANTS TO NAME	09
The Architecture That Resolves It	11
THE RESOLUTION: ARTIFACTS PLUS LLMS, BY DESIGN	11
TOKEN CAPITAL: THE AUDIT-GRADE DESIGN IS THE COST-EFFICIENT ONE	15
The Workflows, Reimagined	17
REIMAGINING THE WORKFLOWS: P2P, CLOSE, RECON, FLUX, FP&A	17
THE AI-NATIVE GCC: STANDARDIZATION AS THE OPERATING MODEL	23
BEYOND THE CLOSE: WHERE THE PATTERN GENERALIZES	27
Deciding, Governing, And Acting	29
THE SIX-WEEK QUESTION: ACTING WHEN THE MODELS WON'T SIT STILL	29
EMBEDDED AGENTS, BUILD-VERSUS-BUY, AND THE ERP DECISION	32
GOVERNING IT: THE AI CENTER OF EXCELLENCE	34
THE WORKFORCE SHIFT: FROM PRODUCING TO JUDGING	36
ACT NOW — AND BRING YOUR STAKEHOLDERS WITH YOU	38

THE SYSTEM OF WORK IN FOUR MINUTES

For the board member, the audit committee chair, or the CEO with four minutes. Everything after this page is the evidence.

The shift

Finance runs on two layers: the **system of record** (the ERP — mature, governed) and the **system of work** (judgment, reconciliation, narrative — historically people and spreadsheets, never formally architected). AI is the first technology that operates in the second layer. That layer now has to be deliberately designed and controlled, for the first time.

The constraint

Language models are probabilistic by construction; financial reporting demands deterministic, auditable outcomes. This is not fixed by better models or better prompts. It is fixed by architecture: the model is bounded to judgment; deterministic code owns every number that must tie; a human signs at defined control points. **The model never authors a number that hits the financials** — so model error becomes a caught exception, not a silent misstatement.

Why now

The same architecture that makes AI auditable makes it affordable, and makes model churn an upgrade event rather than a rebuild. Waiting for the models to settle is the expensive choice — they won't, and none of the durable work (controls, plumbing, institutional muscle) arrives with a model release.

The ask

One operated workflow, one quarter, one named owner. Start on the transactional spine (reconciliations, payables), measure four things against the pre-AI baseline — **cycle time, preparer hours reclaimed, error rate, audit-preparation effort** — and scale only if it clears the bar set in advance.

The risk of it being wrong

Failures surface as exceptions at control points, adjudicated by a person, before anything posts. The larger risk is inaction: competitors crossing from pilots to operated workflows while the organization is still evaluating.

The sentence to carry out of the room

Build the part that does not change — now, in one workflow — and let the models come to you.

The One-Page Version

The System of Work in four minutes

The Shift — And The Constraint

The Architecture That Resolves It

The Workflows, Reimagined

Deciding, Governing, And Acting

FROM SYSTEM OF RECORD TO SYSTEM OF WORK

The most consequential thing AI does to finance is not visible on any ERP roadmap.

Every finance function runs on two distinct technology realities. The first is the **system of record** — the ERP, the sub-ledgers, the consolidation tool, the bank feeds — where truth is *stored*. It is governed, access-controlled, and after decades of investment, reasonably mature. When people say a company is “modernizing finance,” they almost always mean modernizing this layer. But records do not close the books, explain why gross margin moved 140 basis points, decide whether a contract contains a distinct performance obligation, or write the variance commentary the audit committee reads. That work happens in a second, quieter layer — the **system of work** — and understanding why AI lands there, rather than on the ERP, starts with seeing what that layer has actually been made of.

The One-Page Version

The Shift — And The Constraint

From system of record to system of work

Where we actually are: The infancy is the headline

The finance-specific problem no one wants to name

The Architecture That Resolves It

The Workflows, Reimagined

Deciding, Governing, And Acting

Excel was never just a calculator. It was the workflow.

For the entire modern era of finance, the system of work has been built out of two things: people and spreadsheets. It is tempting to describe the spreadsheet as a more flexible companion to the ERP. That undersells it. The spreadsheet became dominant because it is where the work actually happened — the modeling, the reconciliations, the allocations, the what-ifs, the schedules that feed the disclosure. The ERP records what occurred; the spreadsheet, wrapped around the judgment of an experienced preparer, decided what it *means* and what to do about it.

This is why “the system of work” was never formally designed. Enterprise performance management systems did make their way in over the last two decades, and they systematized real slices of it — planning, budgeting, consolidation. But the broader layer of judgment and reconciliation was never fully architected: it largely lived in spreadsheets and in the heads of experienced controllers. No one drew an architecture diagram for it. It accreted — one analyst’s workbook at a time — into the invisible operating layer of the function. Its logic lives in cell formulas no one fully documents and in the tacit knowledge of the person who has closed the books for nine straight years. It is, at once, finance’s greatest source of flexibility and its least controlled surface.

Why this matters for control

The system-of-work layer is exactly where spreadsheet errors, key-person dependency, and weak audit trails already live. Auditors and SOX programs have spent two decades trying to impose discipline on it. If AI now enters this layer, it does not enter a controlled environment — it enters the least-governed, highest-judgment part of finance. That is the opportunity and the risk in the same sentence.

AI’s claim is on the work layer — not the record

Every prior wave of finance technology improved the record: better ERPs, faster consolidation, cleaner data. AI is the first technology that credibly operates in the work: reading a contract and proposing the revenue conclusion, drafting the flux narrative, surfacing the reconciliation exception, assembling the disclosure draft. It does not merely store or retrieve — it participates in interpretation and production.

That reframes the conversation. The question is not “which AI features will my ERP vendor ship?” It is: **if AI becomes the system of work, the system of work now has to be deliberately architected and controlled — for the first time.** Finance never had to design that layer before, because people and spreadsheets filled it by default. Now there is a choice to make, and most organizations have not yet recognized that they are making it.




Figure 1 — AI’s disruption lands on the system of work, not the system of record.

“Excel was never just a calculator. It was the workflow.”

The One-Page Version

The Shift — And The Constraint

From system of record to system of work

Where we actually are: The infancy is the headline

The finance-specific problem no one wants to name

The Architecture That Resolves It

The Workflows, Reimagined

Deciding, Governing, And Acting

WHERE WE ACTUALLY ARE: THE INFANCY IS THE HEADLINE

Most writing on AI in finance describes a destination. The more useful truth is how early the journey is.

Adoption statistics mislead. “Seventy percent of finance teams use AI” is technically true and practically empty, because it counts an analyst pasting a paragraph into a chatbot the same as a governed pipeline that closes a sub-process end to end. To see clearly, separate AI use into three stages of maturity — and be honest about where the mass of the function sits today.

The One-Page Version

The Shift — And The Constraint

From system of record
to system of work

Where we actually
are: The infancy is the
headline

The finance-specific
problem no one wants
to name

The Architecture That
Resolves It

The Workflows,
Reimagined

Deciding, Governing, And
Acting

ASSISTIVE

*individuals, one task at a
time*

A person uses a general AI tool to draft a memo, summarize a standard, or explain a variance. Useful, widely adopted, and almost entirely outside the workflow. Nothing about how work flows changes; one human got a faster assistant. **This is where most mid-market finance sits today.**

EMBEDDED

*AI features inside existing
tools*

The close platform, the FP&A tool, or the ERP ships AI capabilities — auto-reconciliation suggestions, anomaly flags, narrative drafts. Enterprise LLM plug-ins and skills, and the ERP-native copilots, sit here too. Real and valuable, but bounded by the vendor’s product surface and someone else’s design choices. You consume it; you do not compose it. (Section X returns to this as the “embedded-agent” wave.)

OPERATED

*governed, end-to-end, at
scale*

AI executes multi-step finance workflows from input to reviewable output — humans at defined control points, deterministic checks where numbers must tie, and a full audit trail. The workflow is designed, not improvised. **Almost no one is here yet.** This is the frontier, and where durable advantage will be built.

The honest assessment

Stage 3 barely exists in production finance today. We say so plainly because it is both true and clarifying: the winners will not be those with the most AI pilots, but those who cross from assistive experimentation to *operated* workflows — with the controls, determinism, and economics that make output trustworthy. Getting from 1 to 3 is an engineering and governance problem, not a licensing decision.

The implication is uncomfortable but freeing: if almost everyone is at Stage 1, then being early to a well-built Stage 3 in even one or two high-value workflows is a real, defensible head start. The window is open precisely because the work is hard.

We will be candid about that difficulty, because it is where most of the value and most of the failure live. In our own work building these systems, the first version of an operated workflow is rarely the one that ships. The early attempts tend to over-trust the model, under-build the deterministic checks, and discover only in review that the numbers do not tie the same way twice. The teams that succeed are the ones that treat the first build as a draft, expect two or three iterations before the controls are right, and resist the temptation to declare a demo a solution. None of what follows is easy. It is, however, learnable — and the learning is the durable asset.

The same shift, three different starting lines

The thesis is globally true, but the pressure that makes it urgent — and the constraint that makes it hard — differs across the markets we serve. A CFO should read the row that describes their reality.

The One-Page Version

The Shift — And The Constraint

From system of record to system of work

Where we actually are: The infancy is the headline

The finance-specific problem no one wants to name

The Architecture That Resolves It

The Workflows, Reimagined

Deciding, Governing, And Acting

	North America (mid-market)	India	Middle East
Driving urgency	Lean teams under cost and talent pressure; the mandate is “do more without adding headcount.”	Large captive / GCC finance centers where scale and standardization are the advantage — now contestable.	Fast build-out of finance functions alongside rapid diversification and listing ambitions.
Reporting complexity	US GAAP, SEC reporting, SOX / ICFR discipline; auditability is non-negotiable.	Ind AS plus group reporting into US GAAP / IFRS parents; heavy reconciliation and bridge work.	IFRS adoption, newer regulatory regimes, growing assurance and governance expectations.
Structural edge	Closest to the buyers of AI-embedded finance tools; strong vendor ecosystem.	Deep talent and process maturity — the ideal place to <i>operate</i> AI workflows, not just pilot them.	Little legacy debt; the chance to leapfrog directly toward operated, well-governed workflows.
Trap to avoid	Buying embedded features and mistaking Stage 2 for transformation.	Assuming AI threatens the cost-arbitrage model — when, operated well, it compounds it.	Thin senior talent makes ungoverned AI risky; design controls in from day one.

One thesis, three plays

In North America: leverage without headcount, proven to the auditor. In India: turn process depth into operated-AI advantage that *strengthens* the delivery model rather than undercutting it. In the Middle East: leapfrog legacy and build governance in from the start. Same architecture; different first move. What unites all three is a single requirement — the output of an AI workflow must be something a finance leader can *rely on, defend, and trace*. That is where the real engineering begins.

THE FINANCE-SPECIFIC PROBLEM NO ONE WANTS TO NAME

A marketing team can live with an AI that is usually right. A controller signing a 10-K cannot.

This is the sentence that separates finance from almost every other place AI is being deployed. In most functions, a model that is 95% accurate is a triumph — the remaining 5% is noise a human skim absorbs. In finance the standard is not “impressively good.” It is **right, repeatably, and provably so.**

The difficulty is structural, not a matter of model quality that the next release will fix. A large language model is **probabilistic by construction**. It generates the most likely next token given everything before it. Run the same prompt twice and you may get two defensible but different answers. That property is a feature for creativity and for handling ambiguous language — and it is precisely the wrong property for a system that must produce the same trial balance every time and survive an auditor asking, “why is this number what it is?”

What finance actually demands

Accuracy	
What it means in practice The number is correct against the source data and the standard.	Why an LLM alone struggles Plausible-sounding output is not the same as arithmetically and technically correct output.
Determinism	
What it means in practice The same inputs produce the same output, every run.	Why an LLM alone struggles Sampling means identical inputs can yield different wording or values.
Certainty	
What it means in practice The preparer and reviewer can state confidence and basis.	Why an LLM alone struggles A fluent answer carries no inherent measure of how reliable it is.
Auditability	
What it means in practice Every output traces to inputs, logic, and a reviewer.	Why an LLM alone struggles A single generated block hides the path from source to conclusion.

The One-Page Version

The Shift — And The Constraint

From system of record to system of work

Where we actually are: The infancy is the headline

The finance-specific problem no one wants to name

The Architecture That Resolves It

The Workflows, Reimagined

Deciding, Governing, And Acting

Most thought leadership pretends this tension does not exist — it shows a slick demo and moves on. The demo works because someone checked the output by hand; that does not scale and does not survive an audit. The honest position is that the tension is real and *resolvable* — but resolved by how you **architect around the model**, not by hoping the model becomes certain. A probabilistic engine will not become deterministic, so you stop asking it to be.

The reframe that unlocks everything

The output of a well-built finance AI workflow is reliable not because the model grew more certain, but because the certainty was engineered *around* it. The model is bounded to where probabilistic reasoning adds value; everything that must tie, reconcile, or be signed is handled by deterministic components it never gets to improvise. That is Section IV.

“A marketing team can live with an AI that is usually right. A controller signing a 10-K cannot.”

The One-Page Version

The Shift — And The Constraint

From system of record to system of work

Where we actually are: The infancy is the headline

The finance-specific problem no one wants to name

The Architecture That Resolves It

The Workflows, Reimagined

Deciding, Governing, And Acting



Section 4

THE RESOLUTION: ARTIFACTS PLUS LLMS, BY DESIGN

You make a probabilistic engine produce deterministic outcomes by not asking it to do the deterministic parts.

The architecture that resolves Section III's tension is a **decomposition**. Stop treating the LLM as the system that produces the final number. Treat it as one component in a workflow, surrounded by deterministic **artifacts** — validated code, structured calculations, defined schemas, the actual ERP data, and explicit control checkpoints. The model reasons and drafts; the artifacts compute, validate, and lock down anything that must be exact. **The model proposes, the artifact disposes, the human signs.**

The One-Page Version

The Shift — And The Constraint

The Architecture That Resolves It

The resolution:
Artifacts plus LLMs, by design

Token capital: The audit-grade design is the cost-efficient one

The Workflows, Reimagined

Deciding, Governing, And Acting

What the LLM should do

Interpret messy language, read a contract, classify a transaction, spot an anomaly, draft a narrative, propose a treatment — the genuinely *probabilistic* work where “mostly right and fast, for a human to confirm” is a real gain.

What artifacts should do

Perform the calculation, enforce the schema, reconcile to source, apply the rule the same way every time, and record the trace — the *deterministic* work, anything that must tie, foot, or be signed, handled by code that returns the same answer every run.

Concretely: the LLM might read a lease and conclude it contains an embedded derivative — a judgment. But the measurement, the journal entries, and the disclosure figures are produced by deterministic calculation artifacts, not generated as free text. The model proposes which ledger items match; a deterministic routine confirms the totals tie to the penny and flags what does not. **The human sits at the control point** — reviewing judgment, approving the deterministic result — exactly where a reviewer sits today, but with the routine work assembled for them. The output is reliable because of *where the boundaries are drawn*, and the audit trail falls out for free: each artifact has defined inputs, fixed logic, and a signed output.




Figure 2 — The operated pattern: the model is bounded to judgment; artifacts own everything that must tie or be signed.

FOR THE TECHNICAL EVALUATOR

The **decomposition pattern, made concrete**. A naive implementation asks one model call to read, calculate, conclude, and format in a single block of generated text. It demos well and fails audit, because the number and the logic are entangled inside one probabilistic output. The operated pattern separates concerns:

```
# NAIVE - one probabilistic call owns the number (non-deterministic)
result = llm("Read this lease, compute the ROU asset & liability, write disclosure")

# OPERATED - model judges; artifacts compute, validate, trace
terms      = llm_extract(lease, schema=LeaseTerms)      # probabilistic: interpret
terms      = validate(terms, schema=LeaseTerms)         # deterministic: schema gate
measure    = compute_ifrs16(terms)                      # deterministic: the math
checks     = reconcile(measure, source=erp_balances)     # deterministic: must tie
draft      = llm_narrative(measure, checks)             # probabilistic: words only
review     = human_gate(measure, checks, draft)         # control point: sign-off
```

Three properties fall out. **Determinism**: `compute_ifrs16` returns the same figure every run because it is code, not a sample. **Auditability**: each step has typed inputs, fixed logic, and a recorded output — the trace from source to disclosure is the call graph itself. **Bounded model risk**: the LLM never authors a number that hits the financials, so model error becomes *a caught exception*, not a silent misstatement.

What happens when it gets it wrong

An audit committee's fair question. In a decomposed workflow the realistic failure is not a rogue number in the financials — it is the model proposing a wrong *interpretation* (the wrong classification, a missed clause), which the schema, the calculation, or the human reviewer catches at a control point. The failure surfaces as an **exception to adjudicate**, not a misstatement to discover later. Accountability does not move: **the AI proposes, the human signs**.

The preparer and reviewer own the output exactly as they do today — the difference is that the routine work arrived assembled, with its trail attached. That is a smaller, more familiar risk surface than the black-box demo it replaces, and a very different conversation to have with an auditor.

Mapping the architecture to the audit assertions

Controllers and audit committees do not think in "AI features." They think in **assertions** — the comfort management must have over the numbers before anyone signs. The decomposition is persuasive precisely because it can be read assertion by assertion: for each one, name the risk AI introduces, then the deterministic control that addresses it. The pattern is consistent throughout — *the model proposes, an artifact or a human disposes* — which is what makes this a controls story, not a technology story.

The One-Page Version

The Shift — And The Constraint

The Architecture That Resolves It

The resolution:
Artifacts plus LLMs, by design

Token capital: The audit-grade design is the cost-efficient one

The Workflows, Reimagined

Deciding, Governing, And Acting

Assertion _____ COMPLETENESS

The risk when AI enters the process

The model silently drops part of a population — a subset of contracts, a ledger segment, a class of transactions.

The control in the decomposition

The *artifact* defines and reconciles the population (record counts, control totals) against source; the model never decides what is in scope.

Assertion _____ EXISTENCE / OCCURRENCE

The risk when AI enters the process

The model infers a transaction or match that is not supported — a plausible-looking item with no underlying record.

The control in the decomposition

Every proposed item must trace to a source record before it can post. No source, no entry — enforced deterministically, not by trust.

Assertion _____ ACCURACY

The risk when AI enters the process

Right judgment, wrong number — a fluent output that does not foot or tie.

The control in the decomposition

The arithmetic lives in `compute_*`(), never in generated text, and reconciles to the penny. The figure is the same on every run.

Assertion _____ VALUATION / ALLOCATION

The risk when AI enters the process

A measurement, estimate, or allocation that looks reasonable but is not supportable under the standard.

The control in the decomposition

Deterministic measurement under the relevant standard, plus a human sign-off at the estimate's control point. The basis is recorded, not improvised.

Assertion _____ RIGHTS & OBLIGATIONS

The risk when AI enters the process

The model misreads who controls or owes — a lease, a guarantee, a side agreement.

The control in the decomposition

The model proposes the interpretation; a *human* adjudicates the obligation at a gate. This is a judgment a person owns and signs, with the model's reasoning attached.

Assertion _____ PRESENTATION & DISCLOSURE

The risk when AI enters the process

Classification or disclosure drift inside AI drafted narrative — the words wander from the numbers.

The control in the decomposition

The model drafts words, never the figures; disclosure numbers come from artifacts, and a reviewer approves the presentation against them.

Why this matters for the sign-off

Read down the right-hand column and a pattern emerges: every assertion is supported either by a deterministic artifact or by a human at a control point — never by trusting the model's output as-is. That is exactly the comfort a controller needs to sign and an audit committee needs to approve. The architecture does not ask anyone to take the AI on faith; it gives each assertion a control with a name.

Where this is genuinely hard

It would be dishonest to present this as a clean build. The decomposition is simple to describe and demanding to get right, and in our experience the difficulty concentrates in a few predictable places. Drawing the boundary — deciding exactly where judgment ends and deterministic calculation begins — is the hardest design decision, and the first attempt is usually wrong; some steps that look like judgment turn out to be rules, and some that look mechanical hide a real estimate. The deterministic artifacts themselves take real engineering: a reconciliation routine that ties to the penny across messy, real-world data is not a weekend project. And the control points require judgment about *how much* to surface to a human — too little and the control is theatre, too much and you have rebuilt the manual process with extra steps. We flag this not to discourage but to set the expectation correctly: the first build teaches you where the boundary really is, and the second is the one that works. Anyone who tells you the first demo is the finished system has not built one.

*“The model proposes, the artifact disposes,
the human signs.”*

The One-Page Version

The Shift — And The
Constraint

The Architecture That
Resolves It

The resolution:
Artifacts plus LLMs, by
design

Token capital: The
audit-grade design is
the cost-efficient one

The Workflows,
Reimagined

Deciding, Governing, And
Acting



TOKEN CAPITAL: THE AUDIT-GRADE DESIGN IS THE COST EFFICIENT ONE

The architecture that makes AI auditable and the architecture that makes it affordable are the same architecture.

Here is the insight most finance leaders have not been handed: **rigor and cost discipline are not a trade-off in AI — they are the same design decision viewed from two angles.** The decomposition that produced determinism and auditability in Section IV is also what controls the economics.

Tokens are capital, not consumption

Every interaction with a model consumes *tokens* — the units of text it reads and writes — and every token has a cost. The instinctive framing is a utility bill: usage up, bill up, use less. That framing quietly destroys value. The better framing is **capital allocation**: a token spent on a task that genuinely requires judgment is an *investment* that compounds; a token spent re-deriving something code could have computed once is waste that recurs on every run.

Compounding spend

Tokens spent on judgment a human did slowly — reading the contract, spotting the exception, drafting the narrative. This *buys back human time* on work that scales with the business. The return grows as volume grows.

Recurring waste

Tokens spent asking the model to add up a column, reapply a fixed rule, or restate context it already had. A deterministic artifact does this once, correctly, at near zero marginal cost — forever.

Now the two threads converge. **You do not burn tokens re-deriving what an artifact can compute deterministically and cache.** The reconciliation routine, the calculation, the schema validation — exactly the parts you moved *out* of the model for auditability — are also the parts you removed from your token bill. The expensive, nondeterministic, hard-to-audit way and the cheap, deterministic, auditable way are the same fork in the road. **One design decision; two payoffs.**

The One-Page Version

The Shift — And The Constraint

The Architecture That Resolves It

The resolution:
Artifacts plus LLMs, by designToken capital: The
audit-grade design is
the cost-efficient oneThe Workflows,
ReimaginedDeciding, Governing, And
Acting

FOR THE TECHNICAL EVALUATOR

Where the token economics actually live. Three levers, all unlocked for free by the decomposition pattern:

1. Don't pay the model to do arithmetic. Moving a calc into `compute_*`() turns a recurring per-run token cost into a one-time engineering cost.
2. Cache the stable context; pay only for the delta. Standards text, policy, chart-of-accounts, examples - the large, unchanging prefix - is cached once. Each run pays full price only for the small new input (this lease, this account), not the whole context again.
3. Route by difficulty. Reserve the largest model for genuine judgment. Classification, extraction, and formatting go to smaller, cheaper models or to code.

An organization that adopts AI *without* this structure pays premium rates to a large model to repeatedly perform deterministic work it should never have been asked to do — and gets an unauditable result for it. One that decomposes pays the model only for judgment, caches the rest, and routes the easy steps cheaply. The second runs **far more workflow per dollar** — affording AI across many processes while the first still justifies one pilot.

So how will you know it worked?

The honest answer is not a vendor ROI slide; it is a measurement discipline you set *before* you start. Pick one workflow, define the success bar up front, and measure four things against the pre-AI baseline: **cycle time** on that workflow (days or hours to complete), **preparer hours reclaimed, error and rework rate**, and **audit-preparation effort** (how much work it takes to evidence the result). A sane envelope is one workflow over a single quarter with a named owner. The cardinal rule: **do not declare victory on a demo**. A demo proves the workflow can run once; only the baseline comparison proves it runs reliably, cheaper, and audit-ready at volume. If it clears the bar you set, you scale; if it does not, you have spent a quarter and learned exactly where the gap is — which is itself a return.

“The audit-grade design is the cost-efficient one. One design decision; two payoffs.”

The One-Page Version

The Shift — And The Constraint

The Architecture That Resolves It

The resolution:
Artifacts plus LLMs, by design

Token capital: The
audit-grade design is
the cost-efficient one

The Workflows,
Reimagined

Deciding, Governing, And
Acting

REIMAGINING THE WORKFLOWS: P2P, CLOSE, RECON, FLUX, FP&A

Sit a controller down and the worklist pours out: payables, close, flux, reconciliations, journal entries, board packs — then forecasting, then revenue. That list is not a mess. It is a map — and a few of those processes, reimagined in detail, teach you how to think about all the rest.

The instinct is to ask “which one can AI do?” The better question is “which should I do first, and what does it actually look like when I do?” The list sorts itself along the one axis this paper turns on — the ratio of **judgment to determinism** — which is also the order of readiness.

The One-Page Version

The Shift — And The Constraint

The Architecture That Resolves It


The Workflows, Reimagined

Reimagining the workflows: P2P, Close, Recon, Flux, FP&A

The AI-Native GCC: Standardization as the operating model

Beyond the close where the pattern generalizes

Deciding, Governing, And Acting



Two schools of thought on where to start

This is the one place where thoughtful practitioners genuinely diverge, and it is worth naming the disagreement rather than pretending it away. There are two credible strategies for sequencing an AI transformation of finance, and they point in opposite directions.

The **first school starts with the transactional spine** — payables, receivables, reconciliations — where volume is high, risk is low, and payback is fast, and takes on controllership and the close *last*, precisely because they are the hardest, most scrutinized, most judgment-laden processes in finance. Why lead your program with the audited close, the argument runs, when you can build capability, momentum, and organizational trust on the easier, high-volume wins first, and only then bring that hardened muscle to the parts where the stakes are highest?

The **second school starts with controllership** — on the logic that proving the architecture where the constraint is most unforgiving de-risks everything downstream. If the decomposition holds for the audited close, where the numbers must tie to the penny and survive an auditor, it holds for everything easier. Establish it at the hardest point and the rest of finance is a series of simpler cases.

Our view separates the two questions these schools are really answering. As a matter of *intellectual proof*, controllership is the right proving ground — which is why this paper goes deepest there, and why the architecture in Sections IV and V is demonstrated on the close. That is also why the assurance a controller and audit committee need is the bar the whole design is built to clear. But as a matter of *deployment* — where an organization should actually begin its program — we side with the first school: **start on the transactional spine, build the muscle and the credibility where risk is low and payback is fast, and move up the judgment axis toward the close as the controls and the confidence mature.** Leading a live AI program with the audited close, before the organization has built the capability elsewhere, bets the most scrutinized process in finance on the least-proven ground. Prove the pattern where it is hardest; deploy it where it is easiest first. Those are not in tension — they are answers to two different questions, and conflating them is what makes the sequencing feel contested.

That is the menu. But a menu does not let a controller see their work transformed. So take the processes in detail — starting with the one that is already working in practice today, then the one that teaches the most.

Payables and P2P — where AI is already working

If you want proof that operated AI is real and not a slide, look at accounts payable. Procure-to-pay is the **most mature AI use case in finance today**, and it is mature precisely because it decomposes so cleanly. The *interpretation* is judgment the model does well: reading an invoice, extracting the fields, matching it to the right purchase order and contract, coding it to the right account, spotting the duplicate or the anomaly. The *control* is deterministic and belongs to code: the three-way match, the tolerance check, the policy and approval-limit enforcement. And the *exception* — the invoice that does not match, the payment that breaches policy — routes to a human at a gate.

Reimagined, AP stops being a queue that people work through and becomes a **stream that mostly clears itself**: invoices interpreted and matched on arrival, the clean ones flowing straight to payment under deterministic controls, and the accounts-payable

The One-Page Version

The Shift — And The Constraint

The Architecture That Resolves It

The Workflows, Reimagined

Reimagining the workflows: P2P, Close, Recon, Flux, FP&A

The AI-Native GCC: Standardization as the operating model

Beyond the close where the pattern generalizes

Deciding, Governing, And Acting

team spending its time only on the exceptions and the judgment calls — vendor disputes, unusual terms, suspected fraud. It is the ideal first operated workflow: high volume, high repetition, a clean judgment versus- determinism split, an immediate and measurable payback, and a control story an auditor already understands. Start where the pattern is proven.

If you sit in the payables seat today, here is what actually changes: the work stops being “process every invoice” and becomes “resolve the ones the system flagged.” The PO-match you did by opening three screens is done before you look; what reaches you is the invoice that did not match and the reason it did not. Your day shifts from keying to deciding — is this a genuine discrepancy, a vendor error, or something that should not be paid at all. That is a more skilled job, not a smaller one, and it is worth being honest that it is also a different one.

We should be candid about the one place P2P bites even though it is the mature use case: the model is only as good as the data it matches against, and real vendor master files, PO systems, and contract repositories are messier than any demo. In our experience the first weeks of a P2P build are spent less on the AI and more on the unglamorous work of connecting to systems and handling the edge cases — the invoice with no PO, the partial delivery, the vendor who changed their name. The AI part is comparatively easy; the integration and the exceptions are the work. That is true of every workflow in this section, and it is better to know it going in.

The close — the canonical orchestration

Then take the process every reader lives — the close — and reimagine it completely. It is the right one to go deep on, because recon, flux, and journal entries all live inside it: reimagine the close and you reimagine most of the list at once.

The close, the way it actually runs today

Strip away the close calendar and what remains is a **sequential relay race**. The sub-ledgers close, and only then can the reconciliations begin. The recons happen in spreadsheets, tied out by hand, each owner waiting on the data upstream. Journal entries are prepared and emailed for approval; flux commentary is written in a month-end scramble against variances no one saw coming until the numbers landed. The board pack is assembled last, from whatever finally tied. And review — the real review — happens at the end, when there is no longer time to change anything.

Through all of it, the controller is the human orchestrator, holding the entire dependency chain in their head, chasing the piece that is late so the next piece can start. It is linear, it is human-gated end to end, and its defining feature is that **errors surface last, exactly when they are most expensive to fix**.

The close, reimaged as a decomposed agentic workflow

Now apply the architecture of Section IV — the LLM bounded to judgment, deterministic artifacts owning anything that must tie, humans at control points — not to a single task, but to the whole orchestration. The relay dissolves into something continuous.

The One-Page Version

The Shift — And The Constraint

The Architecture That Resolves It

The Workflows, Reimagined

Reimagining the workflows: P2P, Close, Recon, Flux, FP&A

The AI-Native GCC: Standardization as the operating model

Beyond the close where the pattern generalizes

Deciding, Governing, And Acting

Traditional Close - a sequential relay, gated by people



Five-day sprint. Linear. Human-gated end to end. Errors surface last, when they are most expensive to fix.

reimagined as a decomposed agentic workflow

Reimagined Close - continuous convergence, human at control points



Continuous. Parallel. The close becomes a state that converges — not a sprint that starts at period-end.

Figure 3 — The close, reimagined: from a linear, human-gated relay to a continuously-converging agentic workflow with deterministic checks and human control points.

Concretely, what changes:

Reconciliations stop being a period-end event.

Agents monitor the sub-ledgers *continuously*, preparing reconciliations as transactions post rather than after the books close. A deterministic routine confirms the control totals tie and the population is complete; the agent surfaces only the *exceptions* for a human. The recon is no longer a scramble — it is a standing state that is almost always green.

Journal entries arrive with their support attached.

The model proposes the entry and its rationale; the supporting calculation is a deterministic artifact, not generated text; the entry routes to a reviewer who approves at a control point. The preparer's time moves from building the entry to judging the few that need judgment.

Flux narrative is drafted as variances emerge.

The model writes the “why” against movement that deterministic code has already quantified as the “what,” throughout the period rather than in a month-end panic. (The next subsection takes flux on its own.)

The One-Page Version

The Shift — And The Constraint

The Architecture That Resolves It

The Workflows, Reimagined

Reimagining the workflows: P2P, Close, Recon, Flux, FP&A

The AI-Native GCC: Standardization as the operating model

Beyond the close where the pattern generalizes

Deciding, Governing, And Acting

The board pack assembles itself

from artifacts that are already tied, reconciled, and signed — because the components were produced under control as the period ran, not stitched together at the end.

The controller's role inverts

from *doing the assembly to adjudicating exceptions and signing at control points*, throughout the period instead of at the end. The close becomes a state that **converges continuously** rather than a sprint that begins when the period ends.

Nothing here is magic, and nothing asks the model to be trusted with a number. The agent orchestrates and drafts; deterministic artifacts compute and reconcile; the human owns judgment and the signature. That is the reimagination: not “AI does the close,” but a close *redesigned* so the work is continuous, the checks are deterministic, and the human is freed for the judgment only they can give.

If you run the close today, the change is not that your job disappears — it is that the month-end crunch flattens. The frantic five days become a steadier rhythm where most of what used to pile up at period-end has already been prepared and checked. What you spend your time on shifts from assembling and tying to reviewing what the system assembled and deciding the calls that need a person. Worth saying plainly: this is a real adjustment. Teams that have run the close a certain way for years do not love being told the rhythm is changing, and the first close on the new model is genuinely harder than the old way before it gets easier — you are running the new process and checking the old one at the same time. In our experience the payoff is real but it arrives in the second or third cycle, not the first. That honesty matters, because a team told it will be effortless from day one will conclude the whole thing failed when the first close is hard.

The recon, the flux, and the forecast

The recon, reimagined

Taken on its own, the reconciliation is the highest-readiness workflow in finance — mostly matching and thresholds. The traditional version is a spreadsheet built fresh each period, tied out by hand, where the preparer examines *everything* to find the few items that do not agree. Reimagined, an agent matches continuously against source, a deterministic routine confirms the totals and completeness, and the human is presented only with the **exceptions and their probable cause** — not the whole population. The work shifts from *examining everything to adjudicating the few*. This is where to start, precisely because the determinism dominates and the payback is immediate.

The One-Page Version

The Shift — And The Constraint

The Architecture That Resolves It

The Workflows, Reimagined

Reimagining the workflows: P2P, Close, Recon, Flux, FP&A

The AI-Native GCC: Standardization as the operating model

Beyond the close where the pattern generalizes

Deciding, Governing, And Acting

The flux, reimagined

Flux is the textbook decomposition, and the cleanest illustration of the whole thesis. The number — what moved, by how much — is purely deterministic and belongs to code. The *explanation* — why it moved, whether it is expected, what it means — is judgment, and belongs to the model drafting for a human. Traditionally these collapse into one frantic exercise at month-end: a person computing variances and chasing explanations under deadline. Reimagined, deterministic code owns the “what moved” the instant the data lands; the model drafts the “why” against that movement continuously; and the reviewer confirms or corrects the narrative rather than authoring it from scratch. The same split that makes flux *auditable* — numbers from code, words from the model — is what makes it fast.

The forecast, reimagined

FP&A forecasting is the flux decomposition pointed *forward* — and it takes the paper past controllership into planning, the CFO’s other half. The traditional forecast is a fortnight of analysts pulling actuals, rebuilding driver models in spreadsheets, chasing assumptions from the business, and writing the narrative under deadline. Reimagined, the split is the same as flux, only prospective: the *numbers* — the driver-based roll-ups, the sensitivities, the scenario mechanics — are deterministic and belong to the model engine as code; the *narrative and the assumptions* — the story of why the forecast bends, which scenario is plausible, what the business is signaling — are judgment the model drafts for a human to challenge. Agents refresh the base forecast continuously as actuals land rather than in a quarterly scramble; deterministic code owns every number in the model; the FP&A leader spends their time on the assumptions and the scenarios, not on rebuilding the mechanics. Planning stops being a periodic rebuild and becomes a living view — and the same decomposition that made the backward-looking close auditable makes the forward-looking forecast fast and defensible.

The One-Page Version

The Shift — And The Constraint

The Architecture That Resolves It

The Workflows, Reimagined

Reimagining the workflows: P2P, Close, Recon, Flux, FP&A

The AI-Native GCC: Standardization as the operating model

Beyond the close where the pattern generalizes

Deciding, Governing, And Acting

The data-readiness reality check

One precondition the demos underplay: every workflow above assumes the source data and access exist and are usable. AI has changed this picture — agents can now help cleanse, structure, and map data, so the plumbing is genuinely easier than it was. But easier is not solved. What does not go away is the *governance* of that foundation — lineage, access, ownership, and control over the cleansing agents themselves. That governance is durable, model independent work, and (as Section IX argues) no model upgrade will waste it. Let AI help build the foundation; do not let it absolve you of owning it.

Notice the method. Payables showed the pattern already working; the close showed how one reimagined workflow pulls recon, flux, and journal entries along with it; the forecast carried the same split forward into planning. Take the process you know best, redesign it around the decomposition, and the rest of your worklist becomes legible. **Depth in a few workflows beats breadth across ten pilots — every time.**

THE AI-NATIVE GCC: STANDARDIZATION AS THE OPERATING MODEL

A large share of the world’s finance now runs inside Global Capability Centers and shared-services organizations — a great deal of it from India. AI does not just speed up the work these centers do. It changes what a GCC is.

For a generation, the GCC was built on one idea: labor arbitrage. Move headcount to a lower-cost location and run the same manual processes there. That model delivered real savings, but it also exported the same spreadsheet-and heroics way of working, just to a different geography. AI changes the premise. The organizing principle of the modern GCC is no longer *where the work sits* but **how standardized it is and how well it is augmented** — a shift from a cost center running manual processes to an AI-native capability center where a deterministic standardization layer and an AI augmentation layer do the routine work, and people supervise the judgment. This is not a close story or a reporting story. It is an operating-model story that applies across every finance process a GCC runs.

The One-Page Version

The Shift — And The Constraint

The Architecture That Resolves It

The Workflows, Reimagined

Reimagining the workflows: P2P, Close, Recon, Flux, FP&A

The AI-Native GCC: Standardization as the operating model

Beyond the close where the pattern generalizes

Deciding, Governing, And Acting




Figure 4 — The AI-native GCC operating model: a standardization foundation, an AI augmentation layer, and human control — sequenced so the standardized spine comes first.

The structure is the same decomposition this paper has argued throughout, now applied to a whole operation rather than a single task: **a deterministic standardization engine** that runs each process one controlled way across every entity; an **AI augmentation layer** that interprets, drafts, and monitors continuously to lift productivity and outcome; and **human judgment at the control points**, where one team supervises a converging portfolio instead of running parallel queues. Crucially, it is *sequenced*. The standardized

transactional and operational spine — payables, receivables, record-to-report, intercompany, reconciliations — comes first, because it is the highest-volume, cleanest-to-standardize work with the fastest payback. Individual entity-specific reporting, which carries the most local judgment, is a deliberate later stage, not the entry point. Standardize and augment the spine; prove it; then move up into the reporting judgment.

Two situations call for this thinking, and they are genuinely different decisions.

Transforming an existing GCC

If you already run a GCC, the question is not “how does AI speed up our close.” It is “how do we standardize our processes across the whole center and augment them with AI, so productivity and outcomes lift across everything we deliver.” The starting point is honest standardization. Most established centers carry years of accumulated local variation — the same process run subtly differently for each geography or client, much of it undocumented and living in people’s heads. The deterministic layer forces and encodes one controlled way to run each process; that standardization is the precondition for everything else, and it is where the hard, unglamorous work sits. Once a process is standardized, the AI augmentation layer lifts it: agents interpret inputs, draft outputs, and monitor continuously, so the same team delivers more, faster, and with a cleaner control trail. The productivity gain is real, but it compounds from standardization — augmenting a non-standard, undocumented process just accelerates the chaos.

The sequence matters as much as the ambition. In our experience the centers that succeed standardize and augment the transactional spine first — where volume is high and variation is low — build the muscle and the credibility there, and only then take on the multi-entity close and, later still, the entity-specific reporting judgment. Trying to transform everything at once, or leading with the hardest reporting work, is how these programs stall.

Designing a new GCC with AI as a first principle

For anyone standing up a GCC now, there is a rarer and more valuable opportunity: **design the operating model with AI as a first principle, rather than retrofitting it later.** The last generation of centers was architected around headcount — lift a manual process, shift it to a lower-cost location, staff it, and years later try to bolt AI onto a way of working that was never designed for it. That force-fit is expensive and it under delivers, because the process was built for people doing it by hand.

The center designed today can start from the opposite premise. Architect each process as a standardized, deterministic workflow from day one; build the AI augmentation in rather than on; place human control points by design; and staff the center for judgment and supervision rather than manual production. Instead of asking “how many people do we need to run this process,” the question becomes “what is the standardized, AI-augmented design of this process, and where do humans exercise judgment.” A GCC built this way is not a cost center that might later become efficient; it is an AI-native capability center from the first day — leaner, more productive, more governable, and more scalable than any lift-and-shift could become. The people setting up GCCs in this window have a structural advantage the last generation did not, and it is available only to those who design for it rather than force-fit it after the fact.

The One-Page Version

The Shift — And The Constraint

The Architecture That Resolves It

The Workflows, Reimagined

Reimagining the workflows: P2P, Close, Recon, Flux, FP&A

The AI-Native GCC: Standardization as the operating model

Beyond the close where the pattern generalizes

Deciding, Governing, And Acting

How it runs: standardized, multi-framework, follow-the-sun

Whichever situation applies, the operating model has the same properties in motion — and this is where the single entity architecture visibly changes shape at scale.




Figure 5 — Inside the GCC operating model: one team runs many entities and frameworks through an entity-parameterized deterministic engine, with follow-the-sun human control.

The One-Page Version

The Shift — And The Constraint

The Architecture That Resolves It

The Workflows, Reimagined

Reimagining the workflows: P2P, Close, Recon, Flux, FP&A

The AI-Native GCC: Standardization as the operating model

Beyond the close where the pattern generalizes

Deciding, Governing, And Acting

What changes at GCC scale

Standardization engine. Forty closes run one controlled way, not forty spreadsheet dialects — the deterministic layer is what makes multi-entity governable at all.

Multi-framework by configuration. Ind AS / IFRS / US GAAP differences live in the artifact layer as config; the model reads local nuance, the engine applies the rules. **The close never sleeps.** Agents converge continuously; human control points hand off between hubs across time zones. The close is a state, not a nightly stop.

The controller supervises a portfolio — from firefighting forty parallel closes to adjudicating exceptions surfaced across all of them, in one view. At scale, the architecture does not just help — it changes shape: the deterministic layer becomes the standardization engine.

Multi-framework becomes configuration, not re-work

Ind AS, IFRS, and US GAAP differences — the lease nuance, the revenue timing, the impairment trigger — live in the *artifact layer as configuration* per entity, not as dozens of hand-worked variations. The model reads the local nuance and proposes the treatment; the deterministic engine applies the right rule set for that entity, identically every period. A framework change becomes a configuration change, applied once and inherited everywhere it is relevant, rather than a scramble across every affected entity.

The work follows the sun

With agents converging continuously and human judgment at defined control points, the work stops being a nightly event that halts when a hub goes home. It converges around the clock; control points hand off between hubs — an India team adjudicates and hands exceptions to a Middle East hub, which hands to a US hub — and each process is a single converging state that never fully stops. Time-zone spread turns from a coordination tax into an advantage.

The center supervises a portfolio, not a queue

The GCC controller's role inverts hardest of all: from firefighting dozens of parallel processes to **supervising a converging portfolio** — one view across all entities, exceptions surfaced and ranked wherever they arise, judgment applied where it is needed rather than everywhere at once. The span that was unmanageable by headcount becomes tractable because the routine, deterministic work is handled uniformly underneath and only the genuine exceptions rise to a person.

Why this matters for the India and shared-services market

For the GCCs and shared-services centers that run a large share of global finance — and for the enterprises deciding where and how to build them — this is the difference between competing on cost and competing on capability. The center that standardizes and augments, or that is designed AI-native from the start, does not just run finance more cheaply; it runs it better, more governably, and at a scale headcount could never reach. For India, where the GCC is often the buyer and the builder both, this is the strategic frame: AI is not a productivity tool bolted onto the center — it is the operating model the next generation of centers will be built on.

“AI is not a productivity tool bolted onto the center — it is the operating model the next generation of centers will be built on.”

A necessary caution, from doing this at scale: the standardization is where the effort concentrates, and it is easy to underestimate. Encoding one controlled way to run a process across dozens of entities — the genuine local differences, the statutory quirks, the subsidiary that does everything differently for a real reason — is painstaking work, and it determines whether everything above it holds. Prove the model on the transactional spine and a few representative entities first, then extend. The prize is earned process by process, entity by entity — not switched on.

The One-Page Version

The Shift — And The Constraint

The Architecture That Resolves It

The Workflows, Reimagined

Reimagining the workflows: P2P, Close, Recon, Flux, FP&A

The AI-Native GCC: Standardization as the operating model

Beyond the close where the pattern generalizes

Deciding, Governing, And Acting

BEYOND THE CLOSE: WHERE THE PATTERN GENERALIZES

The close showed the method. The same decomposition runs through the whole of finance, reshapes the control framework itself, and reaches into the most judgment-heavy corners of accounting.

If the close is the proof that an everyday workflow can be reimagined, three further moves show how far the pattern travels. They are described as patterns, not products — the point is the *category* of outcome a finance leader should now consider in scope.

The same seam runs through the whole office of the CFO

Section VI went deep on payables, the close, and the forecast — but the decomposition is not a close-and-reporting technique. It is how *any* finance workflow reorganizes, because every one splits along the same seam this paper has now established. Two more workflows complete the CFO's transactional and reporting span, and the point is less the mechanics — those are by now familiar — than the *range*: the same shape holds whether the work is consolidation or collections.

Consolidation

Group consolidation is a clean and high-value generalization. The mechanical heart — intercompany elimination, currency translation, the roll-up across the group — is purely *deterministic* and belongs to code that ties every period. The judgment sits at the edges: the top-side adjustments, the classification of an unusual intercompany arrangement, the consolidation narrative that explains the group result. Reimagined, deterministic engines run the elimination and translation and prove the group foots; the model drafts the adjustments' rationale and the narrative for a reviewer; the group controller signs at the control point. The consolidation that used to be a multi-day, spreadsheet-linked assembly becomes a continuously-tied group view with judgment applied only where it belongs.

Order-to-cash and receivables

O2C rounds out the transactional span opposite payables. Cash application — matching incoming payments to invoices — is largely deterministic matching; collections prioritization and dispute interpretation are judgment the model does well (which customer to chase, what a short-payment note actually means, how to phrase the follow-up). Reimagined, an agent applies cash and matches continuously, deterministic rules confirm the ledger ties, and the collections team is handed a ranked, reasoned worklist — the accounts most worth chasing and why — instead of a flat aged-receivables report. The same split as payables, on the other side of the ledger.

The One-Page Version

The Shift — And The Constraint

The Architecture That Resolves It

The Workflows, Reimagined

Reimagining the workflows: P2P, Close, Recon, Flux, FP&A

The AI-Native GCC: Standardization as the operating model

Beyond the close where the pattern generalizes

Deciding, Governing, And Acting

The pattern is the point. Prove the architecture on the hardest instance — the audited close — and it radiates outward across the whole span of finance without changing shape: payables and receivables, consolidation and reporting, planning and treasury.

Controls over AI — and AI over controls

As AI enters the work layer, the control framework itself must evolve. What *is* the control when a model proposes an entry? Mapping AI-touched processes to a controls framework — with human checkpoints, input validation, and logged evidence at each gate — is now a first-class design problem, and it is the natural extension of the assertion mapping in Section IV. The same techniques run in reverse: AI assists in testing controls, spotting deficiencies, and drafting remediation, under deterministic evidence rules. The function that designs controls over AI is also the one best placed to point AI at its controls.

Treasury and hedge accounting — the pattern at its hardest

Hedge accounting is judgment-heavy at the edges and rule-rigid at the core — which makes it the most demanding test of the decomposition, and a clean one. The model interprets the hedge documentation and the economic relationship; deterministic engines run the effectiveness testing and the measurement under the relevant standard, identically every time, with the working papers generated as a by-product. Even here — high judgment, high stakes, heavy standards — the split holds: model for interpretation, artifacts for measurement, human for the sign-off. If it holds for hedge accounting, it holds for most of what finance does.

The ecosystem reality — not a single-vendor story

Making AI real for the office of the CFO is not the work of one model or one platform. It is a combination of foundation model partners that supply the reasoning, system-of-record vendors that hold the data, and an operator that stitches them into a governed, auditable workflow that runs in a finance function. Uniquis works across that ecosystem of alliances precisely so the CFO does not have to assemble it alone — the value is in the integration and the controls, not in any one component.

The One-Page Version

The Shift — And The Constraint

The Architecture That Resolves It

The Workflows, Reimagined

Reimagining the workflows: P2P, Close, Recon, Flux, FP&A

The AI-Native GCC: Standardization as the operating model

Beyond the close where the pattern generalizes

Deciding, Governing, And Acting

THE SIX-WEEK QUESTION: ACTING WHEN THE MODELS WON'T SIT STILL

A new model lands every few weeks; open-source is catching up fast. So when is the right time to act — and what exactly should you act on?

This sounds like a timing question. It is really a **regret-minimization** question. Every finance leader is quietly afraid of two opposite mistakes: *act too early* and build on a model that is obsolete in a quarter, or *act too late* and watch competitors operationalize while you are still piloting. Those two fears feel like a trade-off — reduce one, increase the other. The most important thing in this paper is that **they are not a trade-off**, and the thing that dissolves them is the architecture you have already read.

THE VIEW FROM THE CAO'S CHAIR

If I move now, I'm betting on a model that's probably beaten by the one shipping next month. If I wait, I read about a competitor who automated their close while I was still 'evaluating.' My IT team wants to build something custom; my auditor wants to know how I'll evidence it; my board wants to know why we don't have an AI story yet. So I do the safe thing — another pilot, another quarter — and somehow the safe thing is the one that keeps costing me time I can't get back."

The resolution: you are not behind because you haven't picked the right model. You are behind only if you haven't built the part that doesn't change.

The architecture is the hedge

The fear of acting early is really the fear of building something *coupled* to today's model. If your workflow is one giant prompt to this quarter's model, then yes — when the model changes, you rebuild, and you are exposed. But the decomposition pattern that gave you determinism, auditability, and token economics is *explicitly* a design that treats the model as a **swappable component**. The schema, the calculation, the reconciliation routine, the control points, the audit trail — none of that changes when the model changes. You swap the reasoning engine and the deterministic scaffolding stands.

So the same architecture that earns the auditor's trust also makes model churn an **upgrade event, not a rebuild event**. A better model — or a cheaper open-source one that clears your reliability bar — drops in, and the whole workflow gets better or cheaper for free. Model acceleration stops being a threat to your investment and becomes a tailwind to it. That reframes the timing question entirely: **you do not time the model. You build the part that does not change, now, and let the models come to you**. Waiting is the more expensive choice, because the dust will not settle — acceleration is the steady state, not a phase — and while you wait you are not building the durable scaffolding, the controls posture, or the institutional muscle, none of which a model upgrade hands you.

The One-Page Version

The Shift — And The Constraint

The Architecture That Resolves It

The Workflows, Reimagined

Deciding, Governing, And Acting

The six-week question: Acting when the models won't sit still

Embedded agents, build-versus-buy, and the ERP decision

Governing it: The AI Center of Excellence

The workforce shift: From producing to judging

Act now — and bring your stakeholders with you

What it takes, concretely

The fair next question is what this actually costs in time and effort. It varies with the process and the state of the data, so any single number is a simplification — but to be concrete rather than evasive: in our experience a first operated workflow is a matter of **weeks to a few months, not days and not a year**. The early weeks go less to the AI than to the unglamorous work of connecting to source systems, encoding the deterministic checks, and handling the edge cases; the model is the comparatively easy part. The important pattern is that the *second* workflow is meaningfully faster than the first, and the third faster still, because the reusable core — the control-point patterns, the artifact libraries, the integration plumbing — is built once and carried forward. The first workflow is where you pay the learning cost; from there the curve bends in your favor. A team should budget for that first build honestly — expect two or three iterations before the controls are right — and expect the economics to improve with every workflow after it.

The decision rule: build now, stay loose, or wait

The One-Page Version

The Shift — And The Constraint

The Architecture That Resolves It

The Workflows, Reimagined

Deciding, Governing, And Acting

The six-week question: Acting when the models won't sit still

Embedded agents, build-versus-buy, and the ERP decision

Governing it: The AI Center of Excellence

The workforce shift: From producing to judging

Act now — and bring your stakeholders with you

POSTURE	WHAT IT COVERS	WHY
Build now	The decomposition discipline, the workflow itself, control points and audit-trail design, data and access plumbing, token instrumentation, and the muscle of running one operated workflow.	Model-independent, compounding, and slow to acquire. None of it is wasted by a model upgrade; all of it takes longer than a quarter to get good at.
Stay loose	The specific model, the specific vendor, long lock-in contracts, anything priced on today's token rates, custom work a near-term release will commoditize.	Fast-moving and cheap to swap. Architect for substitution; do not marry one provider where the next release changes the math.
Genuinely wait	A capability that is on the roadmap, near, and load-bearing for your use case — e.g. a workflow that only works at a reliability the current generation can't hit.	Rare, and worth naming honestly. Pilot to learn; do not scale prematurely. This is the exception that keeps "act now" a matter of judgment, not hype.

The sentence to keep: **act now on everything that survives a model change; stay liquid on everything the next model will change anyway**

Open source is an argument *for* this, not a reason to wait

Open-source models maturing is not a complication to wave away — it is pure upside *if* your workflow is model swappable. It becomes a data-residency option for the Middle East and India, a cost floor on your token bill, a hedge against vendor pricing power, and an option for sensitive workloads that cannot leave the perimeter. The leader who decomposed gets to *use* that optionality; the one who hard-wired a single vendor's API cannot. “Open source is catching up” is therefore one more reason the decoupled architecture is the right bet — not a reason to sit still.

You do not time the model. You build the part that does not change, and let the models come to you.

The One-Page Version

The Shift — And The Constraint

The Architecture That Resolves It

The Workflows, Reimagined

Deciding, Governing, And Acting

The six-week question: Acting when the models won't sit still

Embedded agents, build-versus-buy, and the ERP decision

Governing it: The AI Center of Excellence

The workforce shift: From producing to judging

Act now — and bring your stakeholders with you



EMBEDDED AGENTS AND THE BUILD-VERSUS-BUY TRAP

Your ERP and your SaaS tools are all shipping agents. Your IT team wants to build. Both roads have a trap — and the same architecture marks the safe path through.

Two pressures arrive together. From outside, every platform — the ERP, the close tool, the FP&A suite — is now pitching “our product has agents.” From inside, an eager IT team wants to build something custom. Treated separately, each leads a finance leader into a different expensive mistake. Treated through the decomposition lens, they resolve into one clear principle.

The embedded-agent wave is real — and it is Stage 2

When SAP, Oracle NetSuite, Workday, or your close platform ships agents, that is genuinely useful for the workflows *inside that system's four walls*. Take it. But see its three structural limits clearly:

It is bounded by the vendor's perimeter.

Your real workflows cross systems — ERP plus the close tool plus spreadsheets plus the data warehouse plus the contract repository. No single vendor's agent governs the *cross system* workflow, which is where most of the bottleneck pain and the orchestration value actually live.

You consume its controls posture; you do not author it.

For a SOX or audit bar, “trust the vendor's agent” is a different risk conversation than “here are our control points and our audit trail.” You may be fine with it — but it is *your* sign-off on *their* black box.

It can deepen lock-in just as the model layer commoditizes.

If reasoning is becoming a swappable commodity, betting your whole agentic future on one ERP's embedded agents couples you to a vendor exactly where you want to stay liquid.

The resolution is not build-versus-buy. It is **buy the embedded agents for in-system work; own the orchestration and controls for the cross-system workflow**. The decomposition architecture is what lets you do both — the vendor's agent becomes one more swappable component you orchestrate and govern, not the thing you surrender the workflow to.

The One-Page Version

The Shift — And The Constraint

The Architecture That Resolves It

The Workflows, Reimagined

Deciding, Governing, And Acting

The six-week question: Acting when the models won't sit still

Embedded agents, build-versus-buy, and the ERP decision

Governing it: The AI Center of Excellence

The workforce shift: From producing to judging

Act now — and bring your stakeholders with you

Own your core, rent the rest

Which exposes the build trap. We are seeing finance leaders, encouraged by in-house IT, build *the wrong layer* — the model wrapper, the orchestration plumbing, the UI: the parts commoditizing fastest and costing the most to maintain forever. Meanwhile they underinvest in the part they should own: *their* decomposition logic, *their* control points, *their* audit trail, the judgment-versus-determinism map of *their* close. The principle is neither “always build” nor “always buy”:

The rule: own your core, rent the rest

Build the thin, durable, finance-specific layer that encodes your judgment and your controls. Rent or buy everything generic — the models, the orchestration frameworks, the embedded SaaS agents — that someone else will maintain and upgrade better and cheaper than you can. Do not let IT talk you into building the part that is about to be free.

The word that hides the real cost is **maintain**. A custom build is not a one-time cost; it is a permanent staffing liability — and the six-week model cycle makes it worse, because now you maintain *and* re-integrate on every release. The build looks cheaper at quarter one (one engineer, a prototype) and is brutally more expensive by year two: maintenance, upgrades, re-integration on every model change, key-person risk when the engineer leaves, and your own ownership of custom code for security and SOX. Over any multi-year horizon, the expensive-looking “partner and leverage” path is usually the cheaper one — the same logic as token capital: do not spend scarce capital, engineering or token, on the parts that are commoditizing. Stick to your core; leverage what already exists for the best possible outcome.

Traditional ERP with embedded AI, or an AI-native ERP?

A live question for emerging-growth companies: when the system of record is up for selection, do you choose a mature ERP with embedded AI, or a newer “AI-native” platform? The market is framing this decision the wrong way — as a bet on whose AI is smartest today. But the model layer is the fastest-commoditizing part of the stack. **Choosing your system of record on the strength of its embedded AI is optimizing on the variable that will change most and matter least in eighteen months.**

The durable criteria for a system of record have not changed: data-model integrity, controls, auditability, scalability, and fit to your business. Evaluate the record on those. Then treat the AI layer as something you *architect around* it and keep swappable — exactly the decomposition this paper argues for. The sharp warning is the one that follows from the spine: a tightly-coupled AI-native ERP can mean trading lock-in on the record layer for lock-in on the *work* layer — and that is the worse trade, because the work layer is where both the differentiation and the commoditization are moving fastest.

The POV for the ERP decision

Choose the system of record on system-of-record fundamentals, not on this quarter’s AI demo. Keep the work layer swappable and owned. Done that way, an emerging-growth company can have both — strong embedded AI from a mature platform and ownership of the cross-system orchestration *and* controls — rather than betting the finance function on one vendor’s proprietary native stack. We are deliberately neutral on the vendor and opinionated on the criteria, because that is the position that holds regardless of how the ERP market shakes out.

The One-Page Version

The Shift — And The Constraint

The Architecture That Resolves It

The Workflows, Reimagined

Deciding, Governing, And Acting

The six-week question: Acting when the models won’t sit still

Embedded agents, build-versus-buy, and the ERP decision

Governing it: The AI Center of Excellence

The workforce shift: From producing to judging

Act now — and bring your stakeholders with you

GOVERNING IT: THE AI CENTER OF EXCELLENCE

This paper tells you to build. But when everyone starts building, a new problem appears — and it is an organizational one, not a technical one.

Give a capable finance team the decomposition and the tools, and adoption spreads on its own. That is the goal — and it is also the risk. Left ungoverned, it produces exactly what uncontrolled enthusiasm always produces: a payables team builds one agent, FP&A builds another, a controller in a subsidiary spins up a third, and within a year no one can say who is running what. The symptoms are predictable — **shadow AI, duplicated builds, sprawling and ungoverned data, and uncontrolled token and infrastructure spend** — and they are expensive in precisely the ways this paper has warned against.

The answer — the one we have introduced with clients, to a consistently strong response — is an **AI Center of Excellence**: a central function that owns the reusable core so the rest of the organization can move fast without fragmenting. It is how “own your core, rent the rest” stops being a slogan and becomes an operating model.




Figure 6 — The AI Center of Excellence: a central core of reusable patterns, controls, model decisions, and economics that every finance workflow draws on.

The One-Page Version

The Shift — And The Constraint

The Architecture That Resolves It

The Workflows, Reimagined

Deciding, Governing, And Acting

The six-week question: Acting when the models won't sit still

Embedded agents, build-versus-buy, and the ERP decision

Governing it: The AI Center of Excellence

The workforce shift: From producing to judging

Act now — and bring your stakeholders with you

What the CoE owns

The CoE is deliberately thin and central. It does not build every workflow — the workflow teams do that. It owns the *durable, reusable assets* that every workflow should draw on rather than reinvent:

-  **The reusable decomposition patterns.** The judgment-versus-determinism templates, the control-point designs, the artifact libraries — built once, reused across payables, close, consolidation, and the rest, so each team is not re-deriving the architecture from scratch.
-  **The control and audit standards.** One consistent way that AI-touched processes are controlled and evidenced, so every workflow meets the same bar and the auditor sees one framework, not a dozen improvisations.
-  **The model and vendor decisions.** Which models, which providers, on what terms — decided centrally and kept swappable, so the organization negotiates and governs from one position instead of many.
-  **The token and infrastructure economics.** Central visibility and control over spend, so token capital is allocated deliberately rather than leaking across a dozen uncoordinated experiments.
-  **The AI-touched-process inventory.** A single register of what AI is doing where, who owns it, and how it is controlled — the thing that makes the whole estate governable and auditable at all.

The One-Page Version

The Shift — And The Constraint

The Architecture That Resolves It

The Workflows, Reimagined

Deciding, Governing, And Acting

The six-week question: Acting when the models won't sit still

Embedded agents, build-versus-buy, and the ERP decision

Governing it: The AI Center of Excellence

The workforce shift: From producing to judging

Act now — and bring your stakeholders with you

Why the CoE comes before the sprawl, not after

It is far easier to establish the core before a dozen teams have built a dozen incompatible solutions than to consolidate afterward. The CoE is the governance twin of the token-capital argument: do not spend scarce capital — tokens, engineering, data, or trust — on duplicated, ungoverned work. Stand up the center early, let it own the reusable core, and every workflow team compounds the same set of durable assets instead of each paying to rebuild them. This is the control that keeps an AI-enabled finance function coherent as it scales.

We will state our position plainly, because this is where we see organizations make the most expensive mistake: **most finance functions will stand up their AI capability too late and too fragmented, and will pay to consolidate what they should have centralized from the start.** The common path — let every team experiment, see what sticks, formalize later — feels prudent and is in fact the costly one. By the time “later” arrives, there are incompatible tools, duplicated data pipelines, no common control standard, and a token bill no one owns. Our view is that the CoE is not a maturity milestone you reach after success; it is a precondition for scaling without waste.

That said, we will not pretend the CoE is easy to get right either. Stood up wrong, it becomes a bottleneck — a central team every workflow has to queue behind, which kills the speed that made AI attractive. The version that works is deliberately thin: it owns standards and reusable assets and gets out of the way on execution. A CoE that tries to build everything centrally fails as surely as no CoE at all. The balance — central enough to govern, light enough not to obstruct — is the real design challenge, and it takes iteration to find.

THE WORKFORCE SHIFT: FROM PRODUCING TO JUDGING

Every reimagined workflow in this paper moves the human from producing the work to judging it. That is the most important people change in finance in a generation — and it is not automatic.

Read back through the workflows and one pattern repeats: the person who used to build the reconciliation now adjudicates exceptions; the analyst who used to assemble the flux now reviews the drafted narrative; the associate who used to key the invoice now handles the judgment calls the machine escalated. The center of gravity of finance work moves **up the judgment stack** — from mechanical production to review, skepticism, and analysis. This is the honest answer to the question every team is quietly asking.

THE QUESTION EVERY TEAM IS ASKING

“Does this replace me?”

The honest answer: no — but it does change what you are good for. The work moves from producing the numbers to judging them, from doing the assembly to deciding whether the machine got it right. That is a promotion in the nature of the work. But it is a different job, and pretending the transition is automatic is how organizations fail their people through this shift.

“Your people are not being replaced. They are being elevated. You just have to invest in the elevation.”

Why the shift is a real skills discontinuity

Here is the uncomfortable truth most AI narratives skip: **being excellent at producing is not the same as being excellent at judging**. The associate who is brilliant at building a flawless reconciliation is not, by that fact, good at sitting above a machine's output and knowing when to distrust it. Reviewing well demands different muscles — skepticism, judgment under ambiguity, the pattern-sense to catch the exception that *looks* fine, the confidence to override the machine and the humility to know when not to. Those are not the skills that got someone hired into a preparer role, and they do not appear on their own. A finance function that automates production without deliberately building judgment capacity will find its people promoted into a job they were never trained for — and its controls quietly weakened, because a reviewer who defers to the machine is not a control at all.

The One-Page Version

The Shift — And The Constraint

The Architecture That Resolves It

The Workflows, Reimagined

Deciding, Governing, And Acting

The six-week question: Acting when the models won't sit still

Embedded agents, build-versus-buy, and the ERP decision

Governing it: The AI Center of Excellence

The workforce shift: From producing to judging

Act now — and bring your stakeholders with you

What finance leaders have to do about it

The shift is an opportunity — more valuable, more interesting work for people freed from mechanical production — but only if it is invested in. Two moves are non-negotiable:

1

Reskill deliberately for judgment.

Train explicitly for the reviewer's craft: how to interrogate AI output, how to spot the plausible-but-wrong, how to exercise and document judgment, when to escalate and when to override. Assume the gap between a good preparer and a good reviewer is real and must be closed, not that it will close itself.

2

Change the performance criteria.

You cannot measure a reviewer with the metrics you used for a preparer. Throughput and volume no longer capture the value; quality of judgment, exceptions caught, appropriate overrides, and sound documentation do. Performance management has to move with the work, or you will keep rewarding the old job while asking people to do the new one.

The human message underneath the architecture

The decomposition does not remove the human — it keeps the human exactly where judgment lives, at the control point. That is a genuinely optimistic outcome: finance people spend less time producing and more time thinking. But the optimism is conditional. The organizations that come through this well will be the ones that treat the producing-to-judging shift as a deliberate program — reskilling, new performance criteria, and honest conversations — not as something that happens by itself.

The One-Page Version

The Shift — And The Constraint

The Architecture That Resolves It

The Workflows, Reimagined

Deciding, Governing, And Acting

The six-week question:
Acting when the models won't sit still

Embedded agents,
build-versus-buy, and
the ERP decision

Governing it: The AI
Center of Excellence

The workforce shift:
From producing to
judging

Act now — and bring
your stakeholders with
you

ACT NOW — AND BRING YOUR STAKEHOLDERS WITH YOU

The goal is not “adopt AI.” It is to cross from experimentation to one operated workflow you can trust — and to take the room with you when you do.

Two things decide whether this happens: a concrete first move, and the ability to answer what each stakeholder is actually afraid of. Start with the moves.

	Move	Why it matters
The One-Page Version	Pick one high-value, high-repetition workflow — not ten.	Depth beats breadth. One operated workflow (a reconciliation, a close sub-process) proves the model and builds the muscle. Ten pilots prove nothing. Section VI tells you which one.
The Shift — And The Constraint	Map where judgment ends and arithmetic begins.	This single distinction tells you what the LLM touches and what an artifact owns — the decision that delivers determinism, auditability, and cost control at once.
The Architecture That Resolves It	Design the control points and the success bar before the automation.	Decide where the human signs, what evidence is captured, and what “it worked” means — first. Retrofitting controls onto a demo is how AI projects fail audit.
The Workflows, Reimagined	Build the durable core.	Own the decomposition, controls, and trail; leverage models, frameworks, and embedded agents. Stay liquid on the model and the vendor so churn is an upgrade, not a rebuild.
Deciding, Governing, And Acting	Stand up the core — name an owner, and bring the auditor into the design.	Establish the Center of Excellence early (Section XI) to own patterns, controls, model choices, and the AI-touched-process inventory — not stranded in IT. If the trail is a by-product of the architecture, the audit becomes a walkthrough, not a defense.
The six-week question: Acting when the models won’t sit still	Invest in the people shift before you need to.	The work moves from producing to judging (Section XII). Reskill for the reviewer’s craft and change the performance criteria — or you will promote people into a job they were never trained for, and weaken the very controls you built.
Embedded agents, build-versus-buy, and the ERP decision		
Governing it: The AI Center of Excellence		
The workforce shift: From producing to judging		
Act now — and bring your stakeholders with you		

What each stakeholder is really asking

A CFO does not adopt this alone. Each person at the table fears a different thing; here is the question behind their question, and the answer this architecture gives.

CFO

What they're really asking

"Will this spend compound or evaporate?"

The answer the architecture gives

Tokens as capital plus a model-swappable design: spend on judgment, rent the rest, and churn upgrades you for free.

CAO / Controller

What they're really asking

"Can I sign it and defend it to the auditor?"

The answer the architecture gives

Determinism by decomposition and an audit trail that is a by-product of the design — the AI proposes, you sign.

Audit Committee / Board

What they're really asking

"What's our risk if it's wrong — and are we behind?"

The answer the architecture gives

Failures surface as caught exceptions at control points, not misstatements; the larger risk is the cost of not acting.

CIO / CISO

What they're really asking

"Another platform to secure, another vendor to depend on?"

The answer the architecture gives

Swappable components, open-source and residency optionality, controls designed in — less lock-in, not more.

The finance team

What they're really asking

"Does this replace me?"

The answer the architecture gives

No — the work moves from producing to judging, and the human stays at the control point. But the shift is real and must be invested in: reskilling and new performance criteria (Section XII).

The One-Page Version

The Shift — And The Constraint

The Architecture That Resolves It

The Workflows, Reimagined

Deciding, Governing, And Acting

The six-week question: Acting when the models won't sit still

Embedded agents, build-versus-buy, and the ERP decision

Governing it: The AI Center of Excellence

The workforce shift: From producing to judging

Act now — and bring your stakeholders with you

The Uniquis view

Strategy decks will tell you AI matters; you already know that. The harder, more valuable work is making it real — rewiring an actual workflow, engineering the determinism, building the controls around it, and standing it up so it runs every period and survives an audit. That is an operator's job, not an advisor's slide. It is the work Uniquis was built to do.

The one idea to keep

AI's real disruption to finance is not a better record — it is a new **system of work**. Making that system trustworthy means engineering certainty *around* a probabilistic model, not waiting for the model to become certain. The architecture that earns the auditor's trust is the same one that controls the cost, the same one that turns model churn into a tailwind, and the same one that tells you to own your core and rent the rest. So the timing answer is not “wait and see.” It is: **build the part that does not change — now, in one workflow — and let the models come to you**. The teams that internalize that early will be operating at scale while everyone else is still piloting.

The One-Page Version

The Shift — And The Constraint

The Architecture That Resolves It

The Workflows, Reimagined

Deciding, Governing, And Acting

The six-week question:
Acting when the models won't sit still

Embedded agents,
build-versus-buy, and
the ERP decision

Governing it: The AI
Center of Excellence

The workforce shift:
From producing to
judging

Act now — and bring
your stakeholders with
you



ONE WORKFLOW. ONE QUARTER. A RESULT YOUR AUDITOR CAN WALK THROUGH.

The distance between an AI pilot and an operated workflow is one quarter of disciplined work. Uniquis runs it with you: select the workflow, map judgment versus determinism, design the control points and the success bar, build, and measure against your pre-AI baseline — cycle time, hours reclaimed, error rate, audit-preparation effort. If it clears the bar you set, you scale with the reusable core already built. If it does not, you know exactly where the gap is.

Uniquis partners with the office of the CFO to turn AI into operated, auditable workflows — across US GAAP, IFRS, and Ind AS, and across North America, India, and the Middle East.

START THE CONVERSATION



About Uniquus Consultech:

Uniquus Consultech is an AI and global tech-enabled consulting company that specializes in Accounting & Reporting Consulting, Governance, Risk & Compliance, Sustainability & Climate Consulting, Tech Consulting, and Valuations. The Company is co-founded by consulting veterans Jamil Khatri and Sandip Khetan and backed by marquee investors such as Nexus Venture Partners, Sorin Investments, and UST.

Uniquus has a global team of 800+ professionals, led by 100+ Partners and Directors, across 12 offices in the USA, the Middle East, and India. The company serves more than 400+ clients, including marquee names in each of the markets it operates in.

Uniquus is committed to leveraging technology and an integrated global delivery model to provide best-in-class consulting services to its clients.

About Our Accounting & Reporting Consulting Services:

In today's dynamic business environment, organizations are continuously required to adapt to the dynamic shifts in accounting & reporting standards and regulations. Enhanced management reporting requirements and accelerated reporting timelines necessitate increased automation within the finance function. Special events, such as preparing for an IPO/capital market transaction or fulfilling the reporting requirements arising from an M&A transaction, introduce additional complexities and challenges for the finance function.

As a result, organizations need to transform their finance functions from time to time with a focus on people, processes, and technology. Our Accounting & Reporting Consulting (ARC) practice is designed to partner with the finance function through these challenges. Our teams function as an extension of your finance teams to execute your plans seamlessly.

We bring the necessary technical expertise, skills, technology, and bandwidth, enabling you to navigate your plans in real time without adding additional manpower costs within the organization.

A Team That You Can Trust To Deliver



Sandip Khetan

*Co-founder & Global Head
of Accounting & Reporting
Consulting*



Shashikant Shenoy

*Partner, Finance
Transformation*



Gagan Arora

*Partner, Finance
Transformation*

Scan this code to find
more content like this



Visit us at www.uniquus.com

or follow us on

