October 2026

Point of View

ENTERPRISE AI SERIES / PAPER 1 OF 3

Architecting AI for the Enterprise



EXECUTIVE SUMMARY

EXECUTIVE SUMMARY

CLASSIFY:
From Idea to Investment Decision

SOLUTION:
Choose the Lightest Fit

GOVERN:
Make Controls Enforceable

OBSERVE:
Manage the Live AI Estate

THE LOOP IN PRACTICE:
Supplier Invoice Exceptions

MAKING IT STICK:
The First 90 Days and Beyond

APPENDIX A:
Qualification Notes

APPENDIX B:
Sources and Disclosure

Classify the demand. Choose the solution. Govern the use. Observe the outcome.

Artificial intelligence (AI) is now spreading across employee productivity, business functions, and enterprise workflows. The leadership challenge is no longer choosing a model. It is deciding which problems deserve AI, what solution fits, where controls sit, and whether each deployment keeps paying its way.

This paper sets out a practical discipline for strategic decisions for business leaders, chief financial officers (CFOs), chief information officers (CIOs), and AI Councils in India and the Middle East.

Central proposition

- **Start with demand, not supply:** Process redesign, enterprise resource planning (ERP) configuration, rules, analytics, and AI are alternative answers to a business problem, not steps on a maturity ladder.
- **Govern the use case, not the technology:** A drafting assistant and an agent that can release payments need different release tracks.
- **Keep architecture under review:** Models, data, usage, costs, and agent behavior change after go-live, and approved designs must change with them.

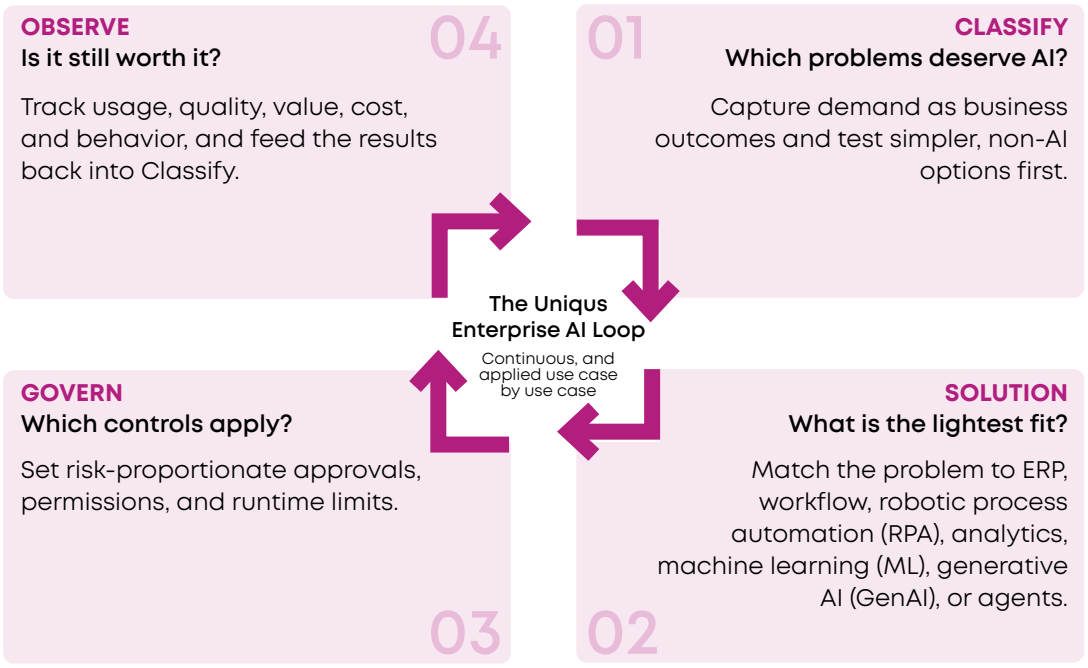


Exhibit 1 | Four connected decisions, repeated across the life of each use case. A finding in a later stage can reopen decisions taken in earlier stages

The aim is a well-run portfolio, not the largest AI estate:

The smallest solution that does the job, on reusable foundations, with proportionate controls and evidence behind every decision to scale, redesign, or stop. The sections below take each decision in turn, apply all four to one finance problem, and end with a 90-day plan and an operating model.

EXECUTIVE SUMMARY

CLASSIFY:
From Idea to Investment Decision

SOLUTION:
Choose the Lightest Fit

GOVERN:
Make Controls Enforceable

OBSERVE:
Manage the Live AI Estate

THE LOOP IN PRACTICE:
Supplier Invoice Exceptions

MAKING IT STICK:
The First 90 Days and Beyond

APPENDIX A:
Qualification Notes

APPENDIX B:
Sources and Disclosure

CLASSIFY

From Idea to Investment Decision

Every request should arrive as a business problem rather than a product: "reduce unresolved supplier invoice exceptions," not "deploy an invoice agent." Classification then involves three separate decisions. Folding them into a single score hides the trade-offs between them.

Decisions	What to ask	Possible outcomes
Eligibility	Is the purpose legitimate? Is there an accountable owner, a baseline, acceptable risk, lawful access to the data, and basic readiness?	Proceed, remediate, defer, or reject.
Investment priority	How much incremental value is at stake, and how material is it? What will it cost, and how soon will value arrive? Could the process itself be redesigned?	Now, later, experiment, or decline.
Solution fit	Is the work rule-based or judgment-based? Does it need unstructured context, system integration, or the authority to act?	ERP or rules, workflow, analytics or ML, document AI, copilot, functional AI platform, or agent (see Exhibit 5).

Exhibit 2 | Eligibility is a gate, priority is a ranking and solution fit is a design decision. A high value score does not offset unacceptable risk.

Prioritizing eligible demand

Requests that pass the eligibility gate can be ranked on five dimensions.

Dimensions	What is assessed	What it determines
Value	Operational or financial improvement against an agreed baseline, with credible volumes and attribution.	The size of the benefit and who funds it.
Work type	Known rules, prediction, reading unstructured content, judgment support, or adaptive multi-step work.	Whether AI is needed at all.
Context and connectivity	Reliability of the data, the systems involved, and whether access is read-only or read and write.	How complex the architecture will be.
Authority and exposure	Whether the solution informs, recommends, initiates, or executes, and what happens if it is wrong.	The level of risk and the release track.
Readiness and scale	Process stability, skills, volumes, change effort, testability and operational ownership.	Feasibility and sequencing.

Exhibit 3 | A five-dimension prioritization scorecard.

Decision Rule

High value alone does not justify AI. If an ERP configuration change or a simple rule removes the bottleneck at lower cost and risk, do that first.

EXECUTIVE SUMMARY

CLASSIFY:
From Idea to Investment Decision

SOLUTION:
Choose the Lightest Fit

GOVERN:
Make Controls Enforceable

OBSERVE:
Manage the Live AI Estate

THE LOOP IN PRACTICE:
Supplier Invoice Exceptions

MAKING IT STICK:
The First 90 Days and Beyond

APPENDIX A:
Qualification Notes

APPENDIX B:
Sources and Disclosure

SOLUTION

Choose the Lightest Fit

Most workflows combine fixed rules, selective AI, and human judgment. Choosing a solution means answering three questions in order: which lane the work sits in, which technology fits, and where it should come from.

Three investment lanes

We group AI demand into three lanes, each with its own funding model, owner, and success measure.

	AI for Productivity *Help individuals work faster and better*	**Functional AI** *Improve a defined process within one function*	**AI for the Enterprise** *Coordinate work across systems and functions*
What it covers	Drafting, search, summarizing, and analysis by individual employees.	AI applied to one process, often inside the core application or a domain platform.	Workflows that span systems and teams, increasingly run by agents.
Typical use cases	Email and document drafting; meeting summaries; policy questions; first-pass variance commentary; Arabic-English translation.	Invoice and matching exceptions; reconciliations; contract review; credit-memo drafting; service-case handling.	Order-to-cash dispute resolution; supplier onboarding; close orchestration across ERP and consolidation; periodic know-your-customer (KYC) refresh.
Usual supply pattern	Enterprise productivity software as a service (SaaS).	Application-native AI or a functional AI platform.	Hyperscaler AI platforms or independent orchestration.
Usual release track	**Fast**	**Controlled**	**High assurance, where agents can change records or recommend decisions.**
Funded and measured	Central budget. Individual Adoption and Productivity.	Function budget. Cycle time, errors, backlog and cost to serve. Verified by Finance.	Enterprise AI budget. business process throughput or contribution. Verified by finance

Exhibit 4 | Three investment lanes. Use cases are illustrative. Supply patterns are described in Exhibit 6 and release tracks in Exhibit 9.

The lane describes the kind of work; the release track describes its exposure. They usually match, but a productivity tool that starts sending output to clients needs a stricter track. The lanes are also not a maturity ladder: many productivity use cases deliver their full value as assistants, and a bank can run enterprise AI in collections while still piloting copilots in human resources (HR).

EXECUTIVE SUMMARY

CLASSIFY:
From Idea to Investment Decision

SOLUTION:
Choose the Lightest Fit

GOVERN:
Make Controls Enforceable

OBSERVE:
Manage the Live AI Estate

THE LOOP IN PRACTICE:
Supplier Invoice Exceptions

MAKING IT STICK:
The First 90 Days and Beyond

APPENDIX A:
Qualification Notes

APPENDIX B:
Sources and Disclosure

Matching work to technology

Within each lane, start with the simplest pattern that removes the bottleneck.

	Start with	Move beyond it when
Structured, stable decisions	ERP configuration, rules, or conventional workflow; RPA only where no application programming interface (API) exists.	Exceptions require reading unstructured inputs or changing the sequence of steps.
Forecasts and scores from historical data	Statistical methods, analytics or conventional machine learning.	Narrative or unstructured evidence would materially improve the result.
Reading contracts, emails or invoices	Document AI and extraction, with validation.	Extraction alone does not resolve the business case.
Summarizing or drafting	An approved copilot or domain assistant.	A defined process outcome and integration justify more.
A domain-specific business process	A functional AI platform, or AI embedded in the function's core application.	Work has to be orchestrated across domains, portability is essential, or the domain capability is missing.
Adaptive work across several systems	Workflow orchestration with agents, plus permissions, checkpoints, and recovery.	This is already the most complex option. Allow agents to act only where a simpler workflow cannot do the job.

Exhibit 5 | Starting points, not fixed rules. A process may combine several solutions, each tested against the actual problem/ bottleneck.

Five supply patterns

The choice is wider than hyperscaler versus application vendor. Functional AI platforms are a distinct category: they package domain workflows that run across enterprise systems and can shorten time to value without replacing the ERP. Most large enterprises will use several patterns.

Enterprise productivity SaaS

01

Examples

Microsoft 365 Copilot; Gemini for Workspace; enterprise editions of ChatGPT and Claude.

Natural fit and main trade-off

Personal productivity and knowledge work. Little integration effort, but time savings need careful testing for quality and value.

Application-native AI

02

Examples

Oracle Fusion Agentic Apps, Fusion Claw and Fusion AI studio; SAP Joule and Business AI; Salesforce Agentforce; ServiceNow Now Assist.

Natural fit and main trade-off

Deep process context inside a core platform. Check external integration, licensing, controls and dependence on one vendor.

EXECUTIVE SUMMARY

CLASSIFY:
From Idea to Investment Decision

SOLUTION:
Choose the Lightest Fit

GOVERN:
Make Controls Enforceable

OBSERVE:
Manage the Live AI Estate

THE LOOP IN PRACTICE:
Supplier Invoice Exceptions

MAKING IT STICK:
The First 90 Days and Beyond

APPENDIX A:
Qualification Notes

APPENDIX B:
Sources and Disclosure

Functional AI platforms

03

Examples

Numero for the Office of the CFO [3]; BlackLine and other finance automation specialists; HR, procurement and legal AI platforms.

Natural fit and main trade-off

Ready-made domain rules, agents, controls and workflows. Test real domain depth, integration with your ERP, audit evidence, and vendor maturity.

Hyperscaler AI platform services

04

Examples

Oracle OCI Enterprise AI; Microsoft Foundry; Amazon Bedrock AgentCore; Google Gemini Enterprise Agent Platform.

Natural fit and main trade-off

Flexible models and engineering foundations. Requires in-house capability for design, integration, evaluation, and agent operations.

Independent orchestration and custom builds

05

Examples

UiPath and other workflow and integration platforms; agent frameworks working through governed APIs.

Natural fit and main trade-off

Process control across platforms and bespoke orchestration. The enterprise owns more of the quality, controls, operations and support.

Exhibit 6 | Examples illustrate categories. They are not endorsements and do not confirm regional availability. Vendors increasingly span categories.

In our view, the functional AI platform option should be tested. For accounts payable (AP), close, reconciliation, revenue, or procurement, compare a specialized functional platform and your core application's AI before building on a hyperscaler. A functional AI platform can cut configuration effort, but it adds a vendor, data copies, and exit risk. Prove the case on real exceptions, not in a demonstration.

Seven selection tests

Put every candidate through the same seven tests. Regulatory and control tests are pass or fail; convenience and price are scored only after a candidate passes them.

	What to validate in a pilot on real exceptions	Minimum evidence
01 Problem and domain fit	Does it resolve real exceptions end-to-end using the right business rules?	Scenario pass rate, exception coverage, and a domain workflow map.
02 Context, model and quality	What is retrieved, sent to the model, and kept? How are errors caught?	Data-flow diagram; evaluation set including local languages.
03 Enterprise integration	Does it connect to your ERP, customer relationship management (CRM) systems, and data through supported APIs, keeping roles and segregation of duties?	Identity design, a live API demonstration, and read and write mapping.

EXECUTIVE SUMMARY

CLASSIFY:
From Idea to Investment Decision

SOLUTION:
Choose the Lightest Fit

GOVERN:
Make Controls Enforceable

OBSERVE:
Manage the Live AI Estate

THE LOOP IN PRACTICE:
Supplier Invoice Exceptions

MAKING IT STICK:
The First 90 Days and Beyond

APPENDIX A:
Qualification Notes

APPENDIX B:
Sources and Disclosure

	What to validate in a pilot on real exceptions	Minimum evidence
04 Regulatory and data fit	Where are prompts, context, outputs, memory, logs, and support data processed, and under which entity?	Applicability matrix; vendor subprocessor and region evidence.
05 Runtime controls and assurance	Can we limit actions, require approvals, detect loops, suspend the system, audit it, and test updates?	Demonstrated policy enforcement, agent traces and an incident runbook.
06 Full economics and scale	What does each completed task cost, including tokens, failed runs, support and human review?	Pilot total cost of ownership, cost caps and volume sensitivity.
07 Portability, exit and vendor risk	Can data, workflows and evidence be exported? Is the provider sound?	Exit drill, export formats and vendor risk assessment.

Exhibit 7 | The same tests apply to functional vendors and hyperscalers. Claims of ERP compatibility should be shown on your own ERP configuration.

Two common evaluation mistakes

- **Treating vendor labels as architectures:** Application vendors now support cross-system agents, and hyperscalers host packaged functional solutions. Decide the architecture and controls first, then the vendor. [4][5]
- **Treating a cloud region as proof of local processing:** The model, agent runtime, memory, search, logs, and support can each run in different regions. Check each provider's service level documentation and include the locations in the contract. [6][7][8]

Two principles for every lane

Every solution draws on the same stack: interfaces, orchestration, models, enterprise context, integration, and infrastructure, with identity, security, monitoring, and cost control operating across all of them. Whatever the lane, two principles hold:

- **Keep systems of record authoritative:** AI can assist or orchestrate, but the ERP, CRM, and HR systems remain the definitive record.
- **Use governed interfaces:** APIs and least-privilege application identities, not screen scraping or shared credentials.



EXECUTIVE SUMMARY

CLASSIFY:
From Idea to Investment Decision

SOLUTION:
Choose the Lightest Fit

GOVERN:
Make Controls Enforceable

OBSERVE:
Manage the Live AI Estate

THE LOOP IN PRACTICE:
Supplier Invoice Exceptions

MAKING IT STICK:
The First 90 Days and Beyond

APPENDIX A:
Qualification Notes

APPENDIX B:
Sources and Disclosure

GOVERN

Make Controls Enforceable

A policy document cannot stop an agent from posting a journal entry. Controls have to be enforced through the systems the enterprise already runs: identity, applications, integration, models and cloud. We call this a common control plane: shared rules and evidence, implemented across several products.

Control area	Enterprise rule	How it is enforced and evidenced
Portfolio and policy	Approved purposes, a model and vendor catalog, exposure levels and permitted jurisdictions.	Use-case register, approved deployment patterns and a named owner for exceptions.
Identity and authorization	Every action traces to an accountable person or agent identity with least-privilege access.	Single sign-on and identity management, delegated scopes, segregation of duties and credential rotation.
Data, context and inference	Source-system permissions carry through to retrieval, and only approved endpoints receive approved data.	Retrieval access controls, redaction and data-loss prevention, subprocessor and location checks, and output handling rules.
Action and workflow	Agents act only within defined tools, limits and approval thresholds, and people can intervene.	API gateway policy checks, approval steps, transaction caps, rollback, and a kill switch.
Assurance and lifecycle	Models, prompts and tools are evaluated, versioned and change-controlled.	Evaluations, adversarial testing, telemetry, audit trails, and incident and decommissioning procedures.
Economics	Consumption is budgeted per use case and tied to outcomes.	Cost allocation, rate and step limits, model routing, alerts and suspension.

Exhibit 8 | A common control plane, consistent with the lifecycle approach of the National Institute of Standards and Technology (NIST) AI Risk Management Framework (RMF) [1]. Open Worldwide Application Security Project (OWASP) guidance on agentic applications is a useful reference for adversarial testing [2].

EXECUTIVE SUMMARY

CLASSIFY:
From Idea to Investment Decision

SOLUTION:
Choose the Lightest Fit

GOVERN:
Make Controls Enforceable

OBSERVE:
Manage the Live AI Estate

THE LOOP IN PRACTICE:
Supplier Invoice Exceptions

MAKING IT STICK:
The First 90 Days and Beyond

APPENDIX A:
Qualification Notes

APPENDIX B:
Sources and Disclosure

Three release tracks

Each use case is assigned to a release track that sets its minimum controls.

		Typical use case	Minimum controls
01	Track **Fast**	Approved, low-exposure drafting and search, with no transactions.	Approved tools, user guidance, review by the user and usage telemetry.
02	Track **Controlled**	Functional workflows, sensitive data or restricted write actions.	A named owner, a test suite, data mapping, approval of actions and a security review.
03	Track **High assurance**	Material financial or customer outcomes, or autonomous or regulated actions.	Formal impact assessment, least privilege, independent review where needed, retained evidence, intervention rights and regulatory approval where required.

Exhibit 9 | Release tracks set minimum controls according to exposure.

Tracks are not fixed. Reclassify a use case when its data, scope, autonomy, transaction limits, market or model changes. Human review is not required for every client-facing output; oversight should be designed in proportion to the law, the consequences of an error and the residual risk.

Regional overlay: India, United Arab Emirates (UAE), Saudi Arabia and Qatar

The control plane is shared across the group; each market adds its own overlay.

Market	Where things stand	What it means for architecture
India	The Digital Personal Data Protection (DPDP) Act and Rules 2025 are being implemented in phases, with the main obligations expected to apply by May 2027. Reserve Bank of India (RBI) and Securities and Exchange Board of India (SEBI) add sector rules. No general AI statute. [9]	Map data, purposes, transfers, sector localization, and effective dates. Keep testing records and clear decision accountability.
UAE	Federal regime onshore; Dubai International Financial Centre (DIFC) and Abu Dhabi Global Market (ADGM) have their own rules. DIFC Regulation 10 covers autonomous and semi-autonomous systems; 2026 amendments are proposals only. [10][11]	Identify each entity's jurisdiction first, then check processing locations, notices, oversight, and transfers.
Saudi Arabia	Personal Data Protection Law (PDPL) and transfer rules apply to personal data; National Cybersecurity Authority (NCA) controls apply within scope. NCA's 2026 AI Cybersecurity Guidelines are a consultation draft. [12]	Test each data flow for location, lawful basis, and applicable controls. Local hosting is not mandatory for every workload.
Qatar	Law No. 13 of 2016 onshore; the Qatar Financial Centre (QFC) has its own regime. The Qatar Central Bank (QCB) AI Guideline sets approval, oversight, and register duties for QCB-regulated entities. [13][14]	Confirm whether the entity launches its own AI or buys high-risk AI, as approvals differ, and set release tracks to match.

Exhibit 10 | Position as of October 14, 2026. A summary, not legal advice; confirm status and applicability for each entity, sector and use case. See Appendix A.

EXECUTIVE SUMMARY

CLASSIFY:
From Idea to Investment Decision

SOLUTION:
Choose the Lightest Fit

GOVERN:
Make Controls Enforceable

OBSERVE:
Manage the Live AI Estate

THE LOOP IN PRACTICE:
Supplier Invoice Exceptions

MAKING IT STICK:
The First 90 Days and Beyond

APPENDIX A:
Qualification Notes

APPENDIX B:
Sources and Disclosure

OBSERVE

Manage the Live AI Estate

Accountability for the architecture starts at go-live. Monitoring has to tell a heavily used tool apart from a useful one, a good model score from a safe business action, and a low price per token from a low cost per outcome.

Lens	Signals to track (illustrative)	Owner and resulting decision
Lens: Usage	Eligible users completing intended tasks; repeat use; use of unapproved tools instead.	Adoption lead in the function: retrain, simplify, or retire.
Lens: Quality	Accuracy on the use case's test set; grounding in sources; Arabic and Indic-language cases; override and rework rates.	Product or process owner: improve context, retrain users or re-evaluate the model.
Lens: Value	Cycle time, capacity captured, errors or backlog against baseline, with attributable changes.	Business owner with Finance: scale, redesign or stop funding.
Lens: Cost	Full cost per successful completion, including tokens, failed calls, hosting and human review, against budget.	AI Center of Excellence (CoE): route, cache, cap, throttle, or suspend.
Lens: Risk	Data loss and policy events, privacy incidents, model changes, review failures, and complaints.	Risk or security owner: constrain, remediate, or suspend.
Lens: Agent behavior	Action traces, tool calls, approvals, retries, unexpected loops and recovery.	Agent or process owner: adjust authority, workflow and fallback.

Exhibit 11 | Each signal needs a baseline, a trigger threshold, a named owner, and a predefined action. Set thresholds against your own risk appetite and volumes.

Cost control happens at runtime. Budget by application, owner, and agent, and measure the full cost of each successful outcome. Route work to the cheapest model that passes the task's evaluation set: a small model can classify invoices, while a frontier model handles the few disputes that need reasoning across contracts and emails. Trim context, cache repeated inputs, cap retries and steps, and throttle or suspend automatically at the approved spending limit. Retest model choices quarterly; a low price per token can hide expensive failures and rework.

Closing the loop

When Observe shows declining value, the response should be a decision, not just a dashboard warning: reduce the scope, revise the business case, change the technology, move the use case to a different release track, or stop it.

EXECUTIVE SUMMARY

CLASSIFY:
From Idea to Investment Decision

SOLUTION:
Choose the Lightest Fit

GOVERN:
Make Controls Enforceable

OBSERVE:
Manage the Live AI Estate

THE LOOP IN PRACTICE:
Supplier Invoice Exceptions

MAKING IT STICK:
The First 90 Days and Beyond

APPENDIX A:
Qualification Notes

APPENDIX B:
Sources and Disclosure

THE LOOP IN PRACTICE

Supplier Invoice Exceptions

Consider a large regional business with a growing backlog of supplier invoice exceptions. The obvious move is to buy a single autonomous accounts-payable agent. A demand-led approach breaks the problem down before choosing any technology.

CLASSIFY

Evidence and action

Separate the exception types: matching tolerance errors, missing goods receipts, supporting documents sent by email, and genuinely disputed charges. Baseline volumes, aging, and handling effort.

Owner and acceptance test

AP owner and Financial Controller agree on the target improvement and which root causes can be prevented.

SOLUTION

Evidence and action

Fix ERP matching rules first. Use document extraction for emailed evidence. Compare the ERP's own AI with specialist finance platforms for reconciliation and case resolution. Keep cross-system agents for cases where the path genuinely varies.

Owner and acceptance test

Architect and AP owner test real exception samples, access rights, APIs and end-to-end cost.

GOVERN

Evidence and action

Let the solution read invoices, purchase orders and receipts, and propose a resolution. Keep vendor-master changes, credit notes, journal postings and payments behind authorized, controlled steps.

Owner and acceptance test

Financial Controller and Risk confirm that approvals, segregation of duties, evidence and recovery match the authority granted.

OBSERVE

Evidence and action

Track backlog aging, straight-through resolution, accuracy, rework, supplier disputes, cost per resolved case, and benefit captured.

Owner and acceptance test

AP owner and Finance scale the solution if the improvement holds, and simplify or stop it if not.

Exhibit 12 | Illustrative example only; not a client implementation or a measured result.

What the procurement evaluation must prove

- **Functional platform or application-native AI:** Which resolves the exception with fewer custom integrations, the right approval controls, and a lower total cost?
- **Functional platform or hyperscaler build:** Are ready-made accounting workflows and evidence worth the vendor dependency, or does a custom build give the control and portability the business needs?
- **The minimum autonomy needed:** Can the outcome be achieved by preparing evidence and seeking human approval, rather than executing payments directly?

Run one pilot on real data, with the same sample cases and guardrails for every candidate. Marketing material, vendor productivity figures, and claims of ERP compatibility are not evidence.

MAKING IT STICK

The First 90 Days and Beyond

The roadmap moves from visibility to evidence in 90 days, and each phase ends with a decision gate.

Phase	Management actions	Gate at the end of the phase
Days 0–30	Inventory AI use cases, agents, spend, and vendors. Name owners and set baselines. Apply the eligibility gate. Set up the AI Council and appoint a CoE lead.	A prioritized register, with duplicate or unsuitable proposals removed.
Days 31–60	Pick two or three workloads across the lanes and run the selection tests. Define the control plane, release tracks and country overlays. Build test cases and cost tracking.	An approved solution pattern, delegated risk decisions and objective pilot criteria.
Days 61–90	Test on real cases, comparing functional, application-native and hyperscaler options. Review the evidence on quality, cost, risk and oversight.	Scale, change approach, or stop, with an evidence pack signed by the business owner.
Beyond day 90	Federate delivery on shared standards. Maintain the model and vendor catalog and review the live estate monthly.	A quarterly AI Council portfolio review with named value and control decisions.

Exhibit 13 | Implementation roadmap. Timings are indicative and depend on the size of the existing AI estate.

EXECUTIVE SUMMARY

CLASSIFY:
From Idea to Investment Decision

SOLUTION:
Choose the Lightest Fit

GOVERN:
Make Controls Enforceable

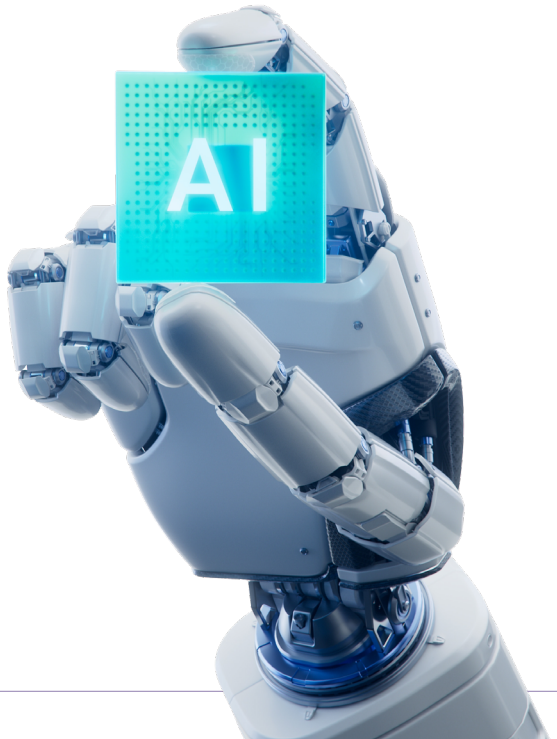
OBSERVE:
Manage the Live AI Estate

THE LOOP IN PRACTICE:
Supplier Invoice Exceptions

MAKING IT STICK:
The First 90 Days and Beyond

APPENDIX A:
Qualification Notes

APPENDIX B:
Sources and Disclosure



EXECUTIVE SUMMARY

CLASSIFY:
From Idea to Investment Decision

SOLUTION:
Choose the Lightest Fit

GOVERN:
Make Controls Enforceable

OBSERVE:
Manage the Live AI Estate

THE LOOP IN PRACTICE:
Supplier Invoice Exceptions

MAKING IT STICK:
The First 90 Days and Beyond

APPENDIX A:
Qualification Notes

APPENDIX B:
Sources and Disclosure

Operating model: centralize standards, federate delivery

Portfolio decisions, formal risk acceptance, and day-to-day accountability are different things, and existing delegations of authority should stay in force. A central AI CoE gives the AI Council one view of usage, cost, model selection and risk across all three lanes, while functions build within its guardrails.

Body	Owns and decides	Should not
AI Council	Priorities and funding; scale, redesign, or stop decisions; high-risk progression within its mandate; realized value and vendor concentration.	Approve every prompt, or accept risk beyond its authority.
Central AI CoE with Enterprise Architecture	Architecture, model and tool standards; approved patterns and catalog; assurance, AI financial operations (FinOps), and monitoring across all three lanes.	Become the bottleneck for every project.
Business and function owners	The outcome, baseline, process redesign and workforce impact. Finance verifies benefits.	Assume that license adoption equals financial value.
Risk, Legal, Privacy and Security	Interpreting requirements, designing controls and formal risk acceptance under existing delegations.	Treat guidance or drafts as binding law.

Exhibit 14 | Decision rights across the AI operating model.

Executive decision

Fund an AI initiative only when the problem, the business owner, a eligible architecture, the control posture, and the measurement plan are all explicit. The best answer to a high-value problem may involve no generative AI at all.

Next in the series: Paper 2, Closing the AI Value Gap, examines how operational improvements become financial outcomes. Paper 3, Building the Agentic Enterprise, looks at how work and accountability change when AI can act across systems.



EXECUTIVE SUMMARY

CLASSIFY:
From Idea to Investment Decision

SOLUTION:
Choose the Lightest Fit

GOVERN:
Make Controls Enforceable

OBSERVE:
Manage the Live AI Estate

THE LOOP IN PRACTICE:
Supplier Invoice Exceptions

MAKING IT STICK:
The First 90 Days and Beyond

APPENDIX A:
Qualification Notes

APPENDIX B:
Sources and Disclosure

APPENDIX A
Qualification Notes

These notes set the boundaries of the regulatory and architectural statements in this paper. They reflect the position as of October 14, 2026; product and regulatory status may change.

India DPDP

Position	Implication
Commencement is phased. Some provisions took effect in November 2025, and the main obligations on data fiduciaries apply by May 2027. [9]	Record which provisions are already in force and plan against the published timetable.

UAE jurisdictions

Position	Implication
The onshore federal regime differs from those of DIFC and ADGM. DIFC Regulation 10 is relevant to AI; the 2026 consultation on amendments is not yet binding. [10][11]	Assess each subsidiary against the jurisdiction in which it is licensed.

Saudi AI cybersecurity guidance

Position	Implication
NCA's AI Cybersecurity Guidelines were under public consultation in July and August 2026. A final, binding version had not been confirmed at the time of writing. [12]	Treat the proposed AI-specific controls as good practice, not legal requirements, until they are issued.

QCB Guideline

Position	Implication
Section 11.1 of the QCB Guideline covers launches and material modifications by providers; Section 11.2 of the Guideline covers agreements to buy, license or outsource high-risk AI. [13]	Establish whether the entity acts as a provider or a user, and whether the AI is high-risk.

Indexes, embeddings and outputs

Position	Implication
Retrieved passages, embeddings, generated outputs and fine-tuned models can carry sensitive information. Their legal classification depends on the facts.	Classify them through a risk-based privacy and security assessment.

Cloud regions

Position	Implication
A cloud region does not prove local processing of model inference, agent runtime, memory, search, logs, or support. [6][7][8]	Confirm the exact service and location in the contract and the service documentation.

The use-case architecture record

For each production use case, keep one record covering:



Ownership and decision: The business owner and baseline, the eligibility decision, and the alternatives considered and rejected.



Data and processing: Data classification; where source data, embeddings, context, outputs, logs, and support data access are processed; the model and its subprocessors.



Authority and controls: Read and write tools, agent permissions, approval thresholds, evaluation evidence, fallback and rollback, and regulatory sign-offs.



Operations: Costs, the service runbook, monitoring thresholds and the date of the next reclassification review.

EXECUTIVE SUMMARY

CLASSIFY:
From Idea to Investment Decision

SOLUTION:
Choose the Lightest Fit

GOVERN:
Make Controls Enforceable

OBSERVE:
Manage the Live AI Estate

THE LOOP IN PRACTICE:
Supplier Invoice Exceptions

MAKING IT STICK:
The First 90 Days and Beyond

APPENDIX A:
Qualification Notes

APPENDIX B:
Sources and Disclosure

APPENDIX B

Sources and Disclosure

References are listed for verification. Laws and regulations, regulator guidance, consultation drafts and vendor statements carry different weight.

1. NIST, AI Risk Management Framework 1.0 and Core (voluntary framework)
2. OWASP, Top 10 for Agentic Applications 2026 (security guidance)
3. Numero, AI agent platform for the Office of the CFO (vendor description)
4. Oracle, AI Agent Studio key capabilities, including external REST and third-party integrations (vendor documentation)
5. SAP, Business AI and Joule capabilities (vendor documentation)
6. Microsoft, Foundry Agent Service supported regions and feature limitations
7. Amazon Web Services (AWS), Amazon Bedrock AgentCore supported regions
8. Google Cloud, Gemini Enterprise Agent Platform locations and data residency
9. Ministry of Electronics and Information Technology (MeitY), Digital Personal Data Protection Rules 2025 and commencement timeline
10. DIFC, Regulation 10 guidance on autonomous and semi-autonomous systems
11. ADGM, Data Protection Regulations 2021, guidance and transfers
12. Saudi NCA, public consultation on AI Cybersecurity Guidelines, July 5, 2026
13. Qatar Central Bank, Artificial Intelligence Guideline (see Section 11 of the Guideline)
14. Qatar National Cyber Security Agency (NCSA), Guidelines for Secure Adoption and Usage of AI, version 1.0

This paper does not constitute legal advice. Applicability by jurisdiction, sector, and product feature must be confirmed with qualified advisors before any production deployment.

About Uniqus Consultech:

Uniqus 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.

Uniqus has a global team of 800+ professionals, led by 100+ Partners and Directors, across 13 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.

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

About Our Tech Consulting Services:

Technology is now a core driver of competitiveness, resilience, and enterprise value.

As organizations digitize operations end-to-end, technology environments are becoming more complex. Fragmented systems, inconsistent data, evolving regulatory expectations, cyber threats, and rapid AI adoption can slow transformation and increase enterprise risk.

Our Tech Consulting practice helps clients cut through this complexity with clarity, discipline, and execution focus. We combine deep technology capabilities with differentiated functional expertise across finance, risk, governance, sustainability, and compliance. Our services help organizations modernize applications, strengthen cyber and technology risk management, build trusted data foundations, automate processes, and adopt AI optimally.

We believe durable transformation requires more than technology implementation. It requires secure architecture, governed data, embedded controls, human oversight, and measurable business outcomes. By integrating these principles from day one, we help clients innovate with confidence and scale with control.

A Team That You Can Trust To Deliver



Anubhav Srivastava
Partner, AI Solutions Architect



Bharat Chadha
Partner, Tech Consulting



Vijay Sampathkumar
Partner, Tech Consulting



Monosij Sinha
Associate Partner, Tech Consulting



Santhanagopalan Varadharajan
Associate Partner, Tech Consulting



Sultan Kreishan
Associate Partner, Tech Consulting

Scan this code to find
more content like this

Visit us at *www.uniqus.com*

or follow us on