November 2026

Point of View

ENTERPRISE AI SERIES / PAPER 3 OF 3

Building the Agentic Enterprise



EXECUTIVE SUMMARY

EXECUTIVE SUMMARY

WHAT MAKES WORK AGENTIC

FROM SYSTEMS OF RECORD TO A SYSTEM OF OUTCOMES

HOW AGENTS FAIL

BUILDING TRUST INTO EXECUTION

WORKED EXAMPLE: Supplier Onboarding

WATCHING AGENTS IN PRODUCTION

OPERATING MODEL AND ASSURANCE

REGIONAL CONTEXT: India, UAE, KSA and Qatar

ROADMAP: From the First Workflow to Agentic Operations

APPENDIX

From systems of record to a system of outcomes.

An artificial intelligence (AI) agent is software that can pursue a goal, decide its next step, and act inside enterprise systems, within limits that people set. That last part changes the stakes. Assistants help people work. Agents do some of the work themselves, at machine speed, in the systems that hold the company's money, customers, and records.

The agentic enterprise is not the one with the most agents. It is one that has redesigned how work gets done: people set goals and authority, agents carry out bounded tasks, and systems of record keep the truth. This paper sets out where agents belong, how to make them safe to trust, and how to start.

At a glance

- **Agentic means bounded, adaptive execution, not a chat interface.** Where the steps can be written down in advance, a conventional workflow is usually the better tool.
- **Autonomy is earned one workflow at a time:** agents move up the autonomy ladder only on evidence, and can move down again.
- **Agents fail differently:** errors compound across steps, and agents can be steered by what they read. Controls have to work at the moment of execution.
- **Systems of record keep the truth:** agents act through approved interfaces and leave evidence as they go.
- **Accountability does not pass to the vendor:** the enterprise that deploys an agent remains responsible for what it does.

Central thesis

Agentic AI becomes consequential when software is trusted to execute work across enterprise boundaries. That requires explicit goals, identity, authority, verification, and human accountability. It is an operating-model change, not a technology upgrade.



EXECUTIVE SUMMARY
WHAT MAKES WORK AGENTIC
FROM SYSTEMS OF RECORD TO A SYSTEM OF OUTCOMES
HOW AGENTS FAIL
BUILDING TRUST INTO EXECUTION
WORKED EXAMPLE: Supplier Onboarding
WATCHING AGENTS IN PRODUCTION
OPERATING MODEL AND ASSURANCE
REGIONAL CONTEXT: India, UAE, KSA and Qatar
ROADMAP: From the First Workflow to Agentic Operations
APPENDIX

WHAT MAKES WORK AGENTIC

An assistant supports a person completing a task. An agentic workflow has a goal, keeps track of where things stand, chooses among permitted next steps, checks the result, and escalates when it cannot proceed safely. Calling tools is not enough on its own: a scripted sequence can call the same systems without ever adapting its plan.

Dimension	Assistant	Agent
Objective	Responds to a request.	Pursues a defined outcome within policies and success criteria.
Context	Uses the information supplied for one answer.	Tracks the state of the work and retrieves permitted context as it changes.
Action	Suggests what the user should do.	Uses authorized tools and systems, within checkpoints and limits.
Feedback	The user judges the answer.	Checks the result, retries within limits or escalates.
Accountability	The user owns review and action.	A named process owner governs the authority delegated; every action is attributable.

Exhibit 1 | The difference is delegated control, not the interface.

The agent test

- **Does the path change case by case?** If the steps and exceptions can be written down and validated, use a workflow, enterprise resource planning (ERP) rules, or robotic process automation (RPA) instead.
- **Is reading evidence or choosing the next step the real bottleneck?** If the delay sits elsewhere, for example in approvals, an agent will not fix it.
- **Does adaptive reasoning pay for its extra cost and risk?** Complex branching on its own does not justify an agent.

The autonomy ladder: authority must be earned

Assist	→ Recommend	→ Execute with approval	→ Bounded execution
The person reviews and acts	The person chooses between options	The agent acts only after authorization	The agent acts within limits; people handle exceptions

Exhibit 2 | The autonomy ladder. These are choices about risk and authority, not maturity stages. Many workflows should stay at Assist or Recommend.

EXECUTIVE SUMMARY

WHAT MAKES WORK AGENTIC

FROM SYSTEMS OF RECORD TO A SYSTEM OF OUTCOMES

HOW AGENTS FAIL

BUILDING TRUST INTO EXECUTION

WORKED EXAMPLE: Supplier Onboarding

WATCHING AGENTS IN PRODUCTION

OPERATING MODEL AND ASSURANCE

REGIONAL CONTEXT: India, UAE, KSA and Qatar

ROADMAP: From the First Workflow to Agentic Operations

APPENDIX

FROM SYSTEMS OF RECORD TO A SYSTEM OF OUTCOMES

Systems of record keep authoritative transactions. **Systems of intelligence** analyze and reason over enterprise information. **Systems of action** make permitted changes through approved interfaces. A **system of outcomes** is the way these work together: people set goals and limits, agents do bounded work, and results are checked against the business outcome. It is an operating model, not a product to buy. [14]

Record	Intelligence	Action	Outcomes
ERP, customer relationship management (CRM), human resources (HR) and core systems	Analytics, retrieval and reasoning	Agent workflows and governed application programming interfaces (APIs)	Business results and human ownership
Truth: owns records, ledgers, and transactions.	**Context:** explains the current state and the possible next step.	**Execution:** coordinates tools and acts within limits.	**Ownership:** outcomes, decisions, and accountability.

Exhibit 3 | Four capabilities working together. The system of outcomes is a way of operating, not a new application.

Two consequences follow. Agents should never duplicate the ledger or bypass approval, identity, and data controls. And connection standards, such as the Model Context Protocol, can make it easier to plug agents into tools, but they do not decide what an agent is allowed to do. That still has to come from the enterprise's identity, permissions, policy checks, and logging. [15]



EXECUTIVE SUMMARY

WHAT MAKES WORK AGENTIC

FROM SYSTEMS OF RECORD TO A SYSTEM OF OUTCOMES

HOW AGENTS FAIL

BUILDING TRUST INTO EXECUTION

WORKED EXAMPLE: Supplier Onboarding

WATCHING AGENTS IN PRODUCTION

OPERATING MODEL AND ASSURANCE

REGIONAL CONTEXT: India, UAE, KSA and Qatar

ROADMAP: From the First Workflow to Agentic Operations

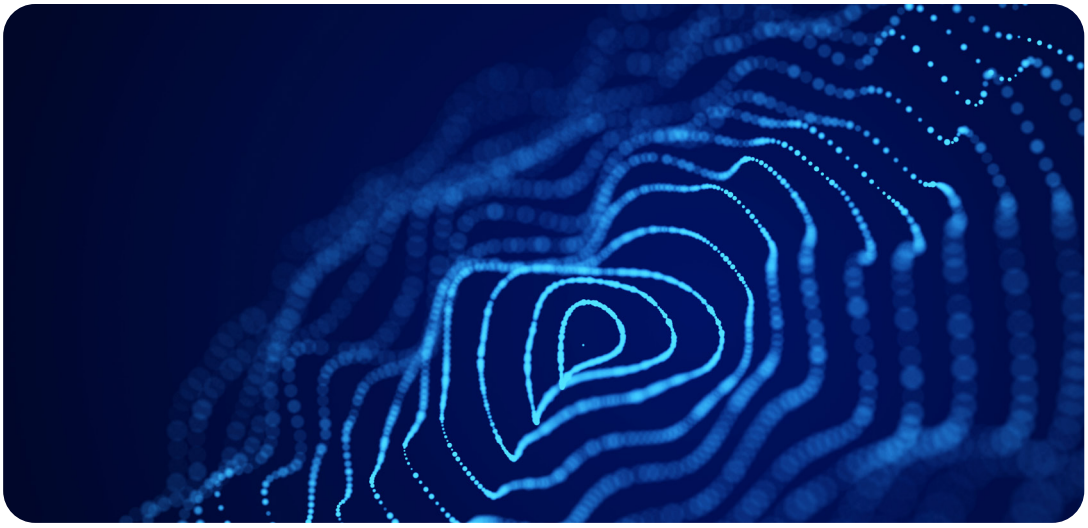
APPENDIX

HOW AGENTS FAIL

Agents inherit the weaknesses of language models and add execution risk of their own. Five failure modes matter most, and none is solved simply by choosing a more capable model. [17]

Failure mode	What happens	How to design against it
Compounding error	Small errors at each step multiply. Ten steps, each 95% reliable, succeed together only about 60% of the time if the steps are independent and nothing is corrected.	Shorten chains, check business state after critical steps, and send uncertain cases to people.
Hijacked instructions	Content the agent reads, such as an email, a supplier document, or a web page, contains instructions that the agent then follows. This is often called prompt injection.	Treat retrieved content as data, never as commands; limit tool rights; require approval for sensitive actions.
Tool and permission misuse	An agent with broad access takes a valid but unintended action.	Separate agent identities, least privilege, transaction limits, and segregation of duties.
Runaway consumption	Loops and retries burn tokens and API calls without producing an outcome.	Set time, step, and cost budgets per run, with automatic suspension.
Poisoned memory	False, stale, or malicious information sits in the agent's memory or knowledge base and steers later decisions.	Log where memory comes from, let it expire, isolate untrusted sources, and be able to roll it back.

Exhibit 4 | Five ways agents fail. The ten-step figure is simple arithmetic, not a measure of any product's reliability.



EXECUTIVE SUMMARY

WHAT MAKES WORK AGENTIC

FROM SYSTEMS OF RECORD TO A SYSTEM OF OUTCOMES

HOW AGENTS FAIL

BUILDING TRUST INTO EXECUTION

WORKED EXAMPLE: Supplier Onboarding

WATCHING AGENTS IN PRODUCTION

OPERATING MODEL AND ASSURANCE

REGIONAL CONTEXT: India, UAE, KSA and Qatar

ROADMAP: From the First Workflow to Agentic Operations

APPENDIX

BUILDING TRUST INTO EXECUTION

"Human in the loop" is not a control design. Asking for approval on every low-risk step recreates the bottleneck. Letting agents take material, irreversible, or regulated actions on their own may exceed the authority the business has granted. The answer is a governed workflow in which each stage pairs what the agent does with a control the enterprise owns.

Stage	What the agent does	Control the enterprise owns
Goal and plan	Turns the outcome into bounded steps.	Approved purpose, owner, stop conditions, and a time and cost budget.
Retrieve context	Uses the relevant policies, records, and knowledge.	Access filtering, data classification, provenance, and transfer controls.
Decide and act	Chooses permitted actions and calls approved systems.	Least privilege, policy checks, approval thresholds, and segregation of duties.
Verify and recover	Checks the result, retries safe steps or raises an exception.	Actions that can be safely repeated or reversed, rate limits, and anomaly detection.
Record and learn	Records what happened and how success was measured.	Traceability, evaluation, change control, and an accountable person.

Exhibit 5 | A governed agentic workflow. The National Institute of Standards and Technology (NIST) AI Risk Management Framework (AI RMF), its Generative AI Profile, and International Organization for Standardization and International Electrotechnical Commission (ISO/IEC) 42001 provide the governance frame; they do not certify agent security. [1][2]

Controls apply before, during, and after execution:

- **Before execution:**
  Document the goal, owner, permitted tools and data, the agent's identity, its delegated authority, and approval thresholds. Test with ordinary and adversarial inputs.
- **During execution:**
  Enforce policy before each material action, cap time, cost, and tool use, validate what the agent reads, log its decisions, and stop or escalate on anomalies.
- **After execution:**
  Reconcile transactions to the system of record, verify the business outcome, keep the evidence, investigate incidents, and review delegated authority periodically.

Design principle: controls at machine speed

A well-governed agent refuses a forbidden action because its identity and policy engine do not allow it, not because someone notices weeks later. Human judgment belongs at material decisions, ambiguities, and exceptions. Evidence should be produced as the work happens, not reconstructed afterward.

EXECUTIVE SUMMARY

WHAT MAKES WORK AGENTIC

FROM SYSTEMS OF RECORD TO A SYSTEM OF OUTCOMES

HOW AGENTS FAIL

BUILDING TRUST INTO EXECUTION

WORKED EXAMPLE: Supplier Onboarding

WATCHING AGENTS IN PRODUCTION

OPERATING MODEL AND ASSURANCE

REGIONAL CONTEXT: India, UAE, KSA and Qatar

ROADMAP: From the First Workflow to Agentic Operations

APPENDIX

Earning more autonomy

An agent moves up the autonomy ladder only when it passes a promotion gate.

Before extending authority	Evidence required	If the evidence is weak
Quality in real work	Representative scenarios, including multilingual cases, error rates, and checks on business state.	Improve context, tool design or human review.
Security and recovery	Prompt-injection testing, least privilege, dry runs, rollback and a working kill switch.	Keep the approval step, or reduce permitted actions.
Economics and ownership	Rate of successful outcomes, cost per completed case, exception burden and a named owner.	Redesign or stop. Do not extend autonomy.

Exhibit 6 | The autonomy promotion gate. Authority can be reduced as easily as it is extended.



WORKED EXAMPLE

Supplier Onboarding

In a multi-country group, onboarding a new supplier touches five teams. Procurement collects documents, Legal reviews clauses, Compliance runs screening, Finance checks tax and bank details, and an ERP operator creates the vendor. Each system works as intended; the delay sits in the hand-offs. The real design question is not whether to use an agent, but which steps an agent should do, at what level of autonomy, and against which failure modes.

EXECUTIVE SUMMARY

WHAT MAKES WORK AGENTIC

FROM SYSTEMS OF RECORD TO A SYSTEM OF OUTCOMES

HOW AGENTS FAIL

BUILDING TRUST INTO EXECUTION

WORKED EXAMPLE: Supplier Onboarding

WATCHING AGENTS IN PRODUCTION

OPERATING MODEL AND ASSURANCE

REGIONAL CONTEXT: India, UAE, KSA and Qatar

ROADMAP: From the First Workflow to Agentic Operations

APPENDIX

	What the agent does	Autonomy level	Main risk controlled	What stays with people
Step 01 **Collect documents**	Checks completeness and requests missing items through approved channels.	Bounded execution	Hijacked instructions hidden in supplier documents.	The business owner sets eligibility and communication policy.
Step 02 **Check policy**	Retrieves the current approved policy, flags inconsistencies and prepares an evidence pack.	Recommend	Poisoned or out-of-date policy content.	Compliance and Legal decide interpretations and exceptions.
Step 03 **Country checks**	None. Deterministic rules run tax registration, sanctions screening, and bank validation.	No agent	Compounding error in legal rules.	Maker-checker and screening results stay authoritative.
Step 04 **Create vendor record**	Prepares a proposed ERP record through approved APIs.	Execute with approval	Tool and permission misuse.	A named approver releases the record.
Step 05 **Track and close**	Tracks dependencies, confirms the final ERP state, and records the trace.	Bounded execution	Runaway consumption and silent failure.	The process owner supervises and steps in on anomalies.

Exhibit 7 | Illustrative design, not a client deployment. Autonomy varies step by step, and one step deliberately uses no agent at all.

Step 3 shows what should not be left to an agent. Goods and Services Tax Identification Number (GSTIN) checks in India, value-added tax (VAT) registration where it applies in the United Arab Emirates (UAE) and Kingdom of Saudi Arabia (KSA), and commercial registration checks in Qatar, which has not introduced VAT as of October 2026, belong in fixed rules. [16] The agent earns its place by reading documents, resolving ambiguous discrepancies, and choosing the approved next action.

Measure the outcome, not the agent: elapsed onboarding time, first-time-right master data, manual touches, exception aging, policy breaches, and cost per supplier onboarded. If the approval queues stay the same, faster document reading may leave the overall cycle time unchanged.

EXECUTIVE SUMMARY

WHAT MAKES WORK AGENTIC

FROM SYSTEMS OF RECORD TO A SYSTEM OF OUTCOMES

HOW AGENTS FAIL

BUILDING TRUST INTO EXECUTION

WORKED EXAMPLE: Supplier Onboarding

WATCHING AGENTS IN PRODUCTION

OPERATING MODEL AND ASSURANCE

REGIONAL CONTEXT: India, UAE, KSA and Qatar

ROADMAP: From the First Workflow to Agentic Operations

APPENDIX

WATCHING AGENTS IN PRODUCTION

Agents need closer monitoring than other AI because they act. The aim is to see a problem early and to have agreed in advance what happens next, including taking away authority.

Signal	What it tells you	Illustrative trigger	Action
Escalations rise	The agent is meeting cases it cannot handle.	Escalation rate doubles against the pilot baseline.	Review scenarios and narrow the agent's scope.
Denied or unusual tool calls	Possible manipulation or a design flaw.	Any attempt at a restricted action.	Investigate; suspend the agent if unexplained.
Business-state mismatch	Actions are not producing the intended result in the system of record.	Any unreconciled transaction.	Stop the affected actions, reconcile, and find the root cause.
Loops and retries	Wasted consumption, often from a broken tool or prompt.	Steps per case 50% above design.	Fix the workflow and enforce step caps.
Cost per completed case rises	The economics are eroding.	20% above the approved value per case.	Re-route models, cap usage, or redesign, using the loaded-cost approach in Paper 2.
Model or tool version changes	Behavior can change without warning.	Any version change.	Re-run evaluations before accepting the change.

Exhibit 8 | Agent monitoring signals. Triggers are illustrative and should be set for each workflow.

Autonomy is reversible

A monitoring finding should be able to move an agent down the autonomy ladder within hours, not after the next governance meeting. Design the kill switch and the fallback to manual work before go-live, and test both.

EXECUTIVE SUMMARY

WHAT MAKES WORK AGENTIC

FROM SYSTEMS OF RECORD TO A SYSTEM OF OUTCOMES

HOW AGENTS FAIL

BUILDING TRUST INTO EXECUTION

WORKED EXAMPLE: Supplier Onboarding

WATCHING AGENTS IN PRODUCTION

OPERATING MODEL AND ASSURANCE

REGIONAL CONTEXT: India, UAE, KSA and Qatar

ROADMAP: From the First Workflow to Agentic Operations

APPENDIX

OPERATING MODEL AND ASSURANCE

Agentic AI changes how work is allocated more than it removes whole roles. Coordination and checking shrink, while domain judgment, exception handling, assurance, and workflow engineering grow. How far teams change shape depends on sector, regulation, process design, and adoption. This is not a forecast of job losses.

Outcome or process owner
Owns the business outcome and delegates authority within approved policy. Residual risk is accepted only by an authorized risk owner.

Domain professional
Encodes business logic, improves context and workflow, and resolves judgment-heavy exceptions.

Agent or workflow owner
Maintains the agent, its tools, operating limits, evaluations, and change control.

AI assurance, cyber and privacy
Assesses quality, security, data permissions, evidence, and regulatory requirements.

Agent operations (AgentOps) and platform engineering
Runs availability, monitoring, cost controls, incident response, and recovery.

Exhibit 9 | Roles in a human and agent operating model. In smaller organizations, one person may hold several roles.

A person should not be held accountable for an automated action without the visibility, authority, and means to intervene. For each workflow, set out who sets policy, who approves material decisions, who handles exceptions, and who can suspend the agent. Reward people for challenging agent decisions, not for accepting them.

Three lines of assurance for agents

Agents should sit inside the organization's existing three lines rather than on a separate governance track. The Institute of Internal Auditors' (IIA) Three Lines Model provides the frame. [18] The Committee of Sponsoring Organizations of the Treadway Commission (COSO) published 2026 guidance on generative AI treats orchestration and posting to records as capability types that need their own controls. [19]

Line	Role for agents	Typical activities
First line: outcome or process owners, and agent or workflow owners	Run agents within delegated authority.	Keep each agent's inventory entry, test scenarios, limits and evidence up to date; handle exceptions.
Second line: Risk, Compliance, Security and Privacy	Set and challenge the control envelope.	Approve autonomy promotions for high-risk workflows, test security and data controls, track regulatory approvals.
Third line: Internal Audit	Independent assurance.	Audit the agent inventory, delegated authority, segregation of duties, and evidence trails; test whether promotions followed the gate; report to the Audit Committee. Advise without owning controls. [20][21]
Board or Audit Committee	Oversight.	Set risk appetite for autonomous execution; review material incidents and assurance results.

Exhibit 10 | Assurance over agents across the three lines.

EXECUTIVE SUMMARY

WHAT MAKES WORK AGENTIC

FROM SYSTEMS OF RECORD TO A SYSTEM OF OUTCOMES

HOW AGENTS FAIL

BUILDING TRUST INTO EXECUTION

WORKED EXAMPLE: Supplier Onboarding

WATCHING AGENTS IN PRODUCTION

OPERATING MODEL AND ASSURANCE

REGIONAL CONTEXT: India, UAE, KSA and Qatar

ROADMAP: From the First Workflow to Agentic Operations

APPENDIX

REGIONAL CONTEXT

India, UAE, KSA and Qatar

Agentic execution raises questions that data-residency analysis does not answer: who is accountable for an automated action, whether a regulator must approve it, and how evidence is kept. Each market is set out on its own terms below; the broader regulatory snapshot is in Paper 1.

India

In our view, accountability sits firmly with the regulated entity, and data obligations are phasing in. Treat every agent, including a vendor's, as your own.

What matters for agents

- The Securities and Exchange Board of India (SEBI) makes regulated entities responsible for their use of AI and machine learning tools, including third-party tools. [10]
- The Reserve Bank of India (RBI) Framework for Responsible and Ethical Enablement of Artificial Intelligence (FREE-AI) report is guidance; proposed model-risk measures are not yet binding rules. [11]
- The Digital Personal Data Protection (DPDP) Act and Rules are commencing in phases. [13]

Design implications

- Keep vendor and in-house agents in one inventory.
- Map data rights, human supervision and evidence to the regulated entity.
- Build consent and purpose checks into agent retrieval as DPDP provisions commence.

Watch-outs

- Assuming a vendor contract transfers accountability.
- Treating RBI proposals as settled rules, or ignoring them.
- Testing agents only in English where customers use Indic languages.

United Arab Emirates

In our view, obligations depend on where the entity is licensed. Map that first, then design oversight for customer-affecting actions.

What matters for agents

- Guidance for licensed financial institutions issued by the Central Bank of the UAE (CBUAE) in February 2026 covers consumer protection, governance, and responsible AI. [12]
- Dubai International Financial Centre (DIFC) Regulation 10 governs personal data processed by autonomous or semi-autonomous systems; amendments consulted on in June 2026 are proposals, not law. [3][4]

Design implications

- Assess Regulation 10 applicability for DIFC entities, including notification duties.
- Define human oversight for material customer-affecting decisions at CBUAE-licensed entities.
- Test agents in Arabic and English.

Watch-outs

- Assuming every agent needs a separate autonomous-system registration.
- Treating proposed DIFC amendments as already in force.

EXECUTIVE SUMMARY

WHAT MAKES WORK AGENTIC

FROM SYSTEMS OF RECORD TO A SYSTEM OF OUTCOMES

HOW AGENTS FAIL

BUILDING TRUST INTO EXECUTION

WORKED EXAMPLE: Supplier Onboarding

WATCHING AGENTS IN PRODUCTION

OPERATING MODEL AND ASSURANCE

REGIONAL CONTEXT: India, UAE, KSA and Qatar

ROADMAP: From the First Workflow to Agentic Operations

APPENDIX

Kingdom of Saudi Arabia

In our view, current law already applies to agents, and AI-specific rules are in draft. Design for today's obligations and track the drafts.

What matters for agents

- The Personal Data Protection Law (PDPL) and applicable National Cybersecurity Authority (NCA) and sector controls apply now. [5]
- The Saudi Data and Artificial Intelligence Authority (SDAIA) consulted on a draft Responsible AI Policy in April 2026. [6]
- The NCA consulted on draft AI Cybersecurity Guidelines in July and August 2026. [7]

Design implications

- Apply PDPL and NCA controls to agents and the data they retrieve.
- Keep regulated postings, such as e-invoices and ledger entries, behind fixed, approved interfaces.
- Design so that registration or reporting duties can be added if the drafts are finalized.

Watch-outs

- Presuming that proposed duties are already enforceable.
- Building long-lived designs that ignore them.

Qatar

In our view, Qatar Central Bank (QCB)-licensed entities face specific approval triggers. Find the trigger early, because it sets the timeline.

What matters for agents

- The QCB's 2024 AI Guideline requires system registers, risk-based obligations, and approval for certain cases, such as some provider launches, high-risk procurement, and autonomous high-risk deployments. [9]
- The National Cyber Security Agency (NCSA) has published guidelines on secure AI adoption. [8]

Design implications

- Identify the exact QCB approval trigger and its lead time.
- Maintain the AI inventory, human oversight, and lifecycle controls.
- Decommission agents by revoking credentials and keeping required records.

Watch-outs

- Building before confirming whether approval is needed.
- Leaving retired agents with live credentials.

Exhibit 11 | Agent-specific considerations by market, as of October 2026. Several instruments are drafts or consultations; confirm status and applicability before relying on them.

One principle holds in every market: using a third-party agent does not remove the deploying entity's own obligations. The precise duties, and any approvals, depend on the jurisdiction, the licensed entity, the activity, and how the agent is deployed.

EXECUTIVE SUMMARY

WHAT MAKES WORK AGENTIC

FROM SYSTEMS OF RECORD TO A SYSTEM OF OUTCOMES

HOW AGENTS FAIL

BUILDING TRUST INTO EXECUTION

WORKED EXAMPLE: Supplier Onboarding

WATCHING AGENTS IN PRODUCTION

OPERATING MODEL AND ASSURANCE

REGIONAL CONTEXT: India, UAE, KSA and Qatar

ROADMAP: From the First Workflow to Agentic Operations

APPENDIX

ROADMAP

From the First Workflow to Agentic Operations

Start with one material workflow and a full view of the agents already running. Each phase ends with a decision, not just a deliverable.

	Days 0–30 Map	**Days 31–60** Bounded build	**Days 61–90** Test and decide	**Beyond Day 90** Industrialize
Key actions	Build the agent inventory. Pick one material workflow whose path varies case by case. Baseline cycle time, hand-offs, errors, and cost; confirm the business owner.	Separate fixed country rules from agent tasks; set delegated authority and intervention points; connect approved tools; build normal and adversarial test scenarios.	Run representative and adversarial cases; verify business state, quality, cost per completed case, control evidence, and recovery.	Run agents as products with owners, AgentOps support, model-change testing, cost budgets, and periodic control reviews. Internal Audit includes agents in its plan.
Outputs	Agent inventory; selected workflow with baseline.	Bounded workflow design, evaluation set and control envelope.	Autonomy decision: expand, redesign or stop.	A live agent portfolio with reversible autonomy.
Lead	AI Center of Excellence (AI CoE) with the process owner.	Process owner, Architecture, and Risk.	AI Council, with second-line sign-off.	Process owners, AI CoE and Internal Audit.

Exhibit 12 | Agent roadmap. Timings are indicative.



EXECUTIVE SUMMARY

WHAT MAKES WORK AGENTIC

FROM SYSTEMS OF RECORD TO A SYSTEM OF OUTCOMES

HOW AGENTS FAIL

BUILDING TRUST INTO EXECUTION

WORKED EXAMPLE: Supplier Onboarding

WATCHING AGENTS IN PRODUCTION

OPERATING MODEL AND ASSURANCE

REGIONAL CONTEXT: India, UAE, KSA and Qatar

ROADMAP: From the First Workflow to Agentic Operations

APPENDIX

How to build the agent inventory

Many organizations already run agents they have not counted: features switched on inside enterprise applications, bots upgraded with AI steps, and tools built by individual teams. COSO's generative AI guidance makes the inventory an early step, and QCB requires an AI register for its licensed entities. [19][9] Six sources usually find them:

Source	What it reveals	Watch for
Agent-building platforms	Agents created by business and IT teams.	Agents built by individuals without review.
Enterprise application admin settings	Agent features enabled in ERP, CRM and service platforms.	Embedded agents running with broad access.
Identity and access management	Non-human identities, service accounts and API keys used by agents.	Shared or over-privileged credentials.
API gateway and integration logs	Which agents call which systems, and how often.	Calls that cannot be traced to a registered agent.
Automation and RPA	Bots that now include AI decision steps.	Legacy bots that have quietly become agents.
Procurement and contracts	Vendor agents and AI services in use.	Agent capabilities bundled into wider contracts.

Exhibit 13 | Sources for the agent inventory.

Record each agent's owner, process, systems touched, permitted actions, autonomy level, data classification, and any regulatory touchpoints. An agent that cannot be traced to an owner should be suspended until it can.

The Uniqus point of view

The agentic enterprise is not defined by the number of agents it runs. It is defined by work redesigned around outcomes: agents coordinate and execute permitted activity, systems of record keep the truth, and people keep the judgment, the accountability, and the ability to intervene.

EXECUTIVE SUMMARY

WHAT MAKES WORK AGENTIC

FROM SYSTEMS OF RECORD TO A SYSTEM OF OUTCOMES

HOW AGENTS FAIL

BUILDING TRUST INTO EXECUTION

WORKED EXAMPLE: Supplier Onboarding

WATCHING AGENTS IN PRODUCTION

OPERATING MODEL AND ASSURANCE

REGIONAL CONTEXT: India, UAE, KSA and Qatar

ROADMAP: From the First Workflow to Agentic Operations

APPENDIX

APPENDIX: AGENT READINESS GATES FOR PRODUCTION

Check	Evidence expected
Business goal	A bounded objective, success test, accountable owner, and acceptable failure threshold.
Right tool	A documented reason why a fixed workflow is not enough.
Enterprise context	Approved data sources, permission filtering, and inference and transfer rules for each country in scope.
Agent identity	Separate least-privilege credentials, permitted tools, and segregation of duties.
Workflow safety	Approval thresholds, safely repeatable or reversible actions, and defenses against hijacked instructions.
Runtime assurance	Traces, anomaly detection, cost ceilings, a kill switch and exception routing.
Human role	A named approver, exception owner, override rights, and workforce training.
Autonomy decision	A recorded promotion decision: approved actions, human checkpoints, test results, cost per completed case owner and rollback.
Retirement	A plan to decommission the agent, revoke its credentials, and keep the evidence.

Sources and qualification notes

1. NIST, AI Risk Management Framework 1.0 and Generative AI Profile (2024)
2. ISO/IEC 42001:2023, Artificial intelligence management systems
3. DIFC, Regulation 10 (2023)
4. DIFC, Consultation on amendments to Regulation 10, June 2026 (a proposal)
5. SDAIA, Personal Data Protection Law, Implementing Regulations and guidance
6. Saudi Press Agency, SDAIA public consultation on a Responsible AI Policy, April 2, 2026
7. Saudi NCA, consultation on AI Cybersecurity Guidelines, July 5, 2026
8. Qatar NCSA, Guidelines for Secure Adoption and Usage of AI, version 1.0 (2024)
9. Qatar Central Bank, Artificial Intelligence Guideline (2024)
10. SEBI, Intermediaries (Amendment) Regulations 2025
11. RBI, FREE-AI Committee Report (2025) and 2026 consultation on model risk
12. CBUAE Rulebook, Guidance Note on responsible AI and machine learning for licensed financial institutions, February 11, 2026
13. Ministry of Electronics and Information Technology (MeitY), Digital Personal Data Protection Rules 2025 and phased commencement
14. Uniqus, Enterprise Agentic AI in the Middle East (earlier working paper, internal)
15. Model Context Protocol, specification and security considerations (2026)
16. International Monetary Fund (IMF), Taxation in the Middle East and North Africa (April 2026), including Qatar VAT status
17. Open Worldwide Application Security Project (OWASP) GenAI Security Project, Agentic AI Exploit Roundup, Q3 2026
18. The IIA, Three Lines Model (2020, with a refreshed Statement of Position in 2026)
19. COSO, Achieving Effective Internal Control Over Generative AI, released February 23, 2026; this source is cited through The IIA's Internal Auditor as a secondary summary. Confirm against COSO's full publication before relying on specific capability descriptions
20. The IIA, Global Internal Audit Standards, effective January 9, 2025
21. The IIA, Artificial Intelligence Auditing Framework

The system-of-outcomes construct, the supplier-onboarding design, and the promotion gate are Uniqus conceptual designs, not claims about deployed products or verified benefits. Regulatory references are a summary of public information as of October 2026, and instruments under consultation are not binding. Confirm applicability and controls for each entity and use case with qualified advisors.

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