September 2026

Insights

Controlling the Agentic Finance Function

Why internal controls must govern the end-to-end workflow — not the AI tool

By Uniqus Consultech

With contributions from Amrita Kapoor, Internal Audit, Snowflake

EXECUTIVE SUMMARY

Scope and action

Internal control for the agentic finance function

Agents are moving from assisting finance to executing work that can affect the financial reporting and book close processes. The control question is no longer whether AI is used in processes; it is what the full workflow can do and how it may affect financial reporting, and how management proves that controls still exist and operate.

Agents now read invoices, prepare accruals, apply cash, route exceptions, and initiate actions. Agents in **procurement, sales, payroll, supply chain, and data platforms** can also affect ICFR given that they affect transactions.

For ICFR, this is more than a technology change. It challenges four long-standing assumptions

- Judgment has a human owner;
- Systems behave deterministically;
- Segregation can be enforced through roles; and
- Change arrives through managed releases.

The Committee of Sponsoring Organizations of the Treadway Commission (COSO) remains a sound framework. But the unit of control is now the full agentic workflow: model, prompts, and configuration, governed data and transformations, tools and permissions, orchestration, independent validation, human review, evidence, and recovery.

Not every AI tool belongs in ICFR. An agent enters the control perimeter when it performs or supports a control, produces information relied upon by a control, initiates, authorizes, or records a transaction, changes relevant master or reference data, or can affect a significant account or disclosure.

Management should first: build the register; identify financially relevant workflows; set clear authority limits; qualify the full workflow before production; capture evidence as the work happens; and monitor material changes, including vendor and data pipeline changes.

The position in brief

An agent that performs any activity that affects financial reporting or disclosures is a component of internal control. Its whole lifecycle, what it may do, how it is qualified, how it is monitored, how changes to it are governed, and how it is retired, belongs inside the Sarbanes-Oxley (SOX) program, evidenced like any other control. The lifecycle built for the first workflow is reused for every one after it. Automating first, without adequate consideration of controls, may result in a substantial increase in the cost of controls.

Sandip Khetan
Co-Founder, Global Head of Accounting & Reporting Consulting

Sharad Chaudhry
Partner, Accounting & Reporting Consulting

Ankur Gupta
Partner, GRC

Shoichi Ohno
Managing Director, Accounting & Reporting Consulting

ACKNOWLEDGEMENTS

Amrita Kapoor
Vice President and Head of Internal Audit, at Snowflake

Amrita Kapoor is Vice President and Head of Internal Audit at Snowflake, with more than 18 years of experience in audit, risk, and compliance across the technology, entertainment, and media industries. Her work spans SEC audits, SOX compliance, technical accounting, and operational audits, with a focus on enabling assurance through data and AI. She was previously a Director in PwC's Risk Assurance practice.

THE SITUATION

Section 1

The control challenge in agentic finance

Agentic finance is not a new control problem. It brings familiar issues into workflows that can act faster, across more systems, and with less direct human involvement. The following examples show where those issues arise in practice.

The accrual journal entry — and the question of judgment

Each month an agent ingests service agreements, purchase orders, receiving data, and prior-period entries; identifies obligations; proposes accruals; and drafts support. Deterministic routines calculate amounts and reconcile the population; cases outside approved criteria route to reviewers, who concentrate on amendments, disputes, and novel arrangements. No material agent-generated amount **reaches the ledger without independently designed validation,** and every entry retains a source-to-approval trail.

The control question this workflow raises is about judgment: **when the agent proposes both the entry and the rationale the reviewer reads, who is actually exercising the judgment the control was designed around?**

Procure-to-pay — and the questions of authority and inputs

In payables, agents read invoices from monitored channels, extract fields, match them to purchase orders and receipts, propose account coding, flag anomalies, and draft discrepancy communications. Independent matching and tolerance checks release routine items to a payment proposal; exceptions route to people. The complexity sits in partial receipts, unit conversions, non-PO invoices, and timing differences — precisely where judgment, fraud, and cut-off risk concentrate. Two control questions follow.

- **Where does the agent's authority stop?** — at proposing a match, correcting reference data, or touching the vendor master where payment fraud concentrates?
- **How are externally supplied documents treated as untrusted inputs?** An instruction embedded in a supplier file must not be able to expand the agent's tools, permissions, or routing.

At period end, open purchase orders and unmatched receipts processed by the payables agent may feed the incurred-but-not-invoiced accrual. A receipt-matching timing error can therefore become a cut-off error in the estimate.

Procure-to-Pay (P2P) — SOX Process Flowchart with AI Touchpoints & Controls

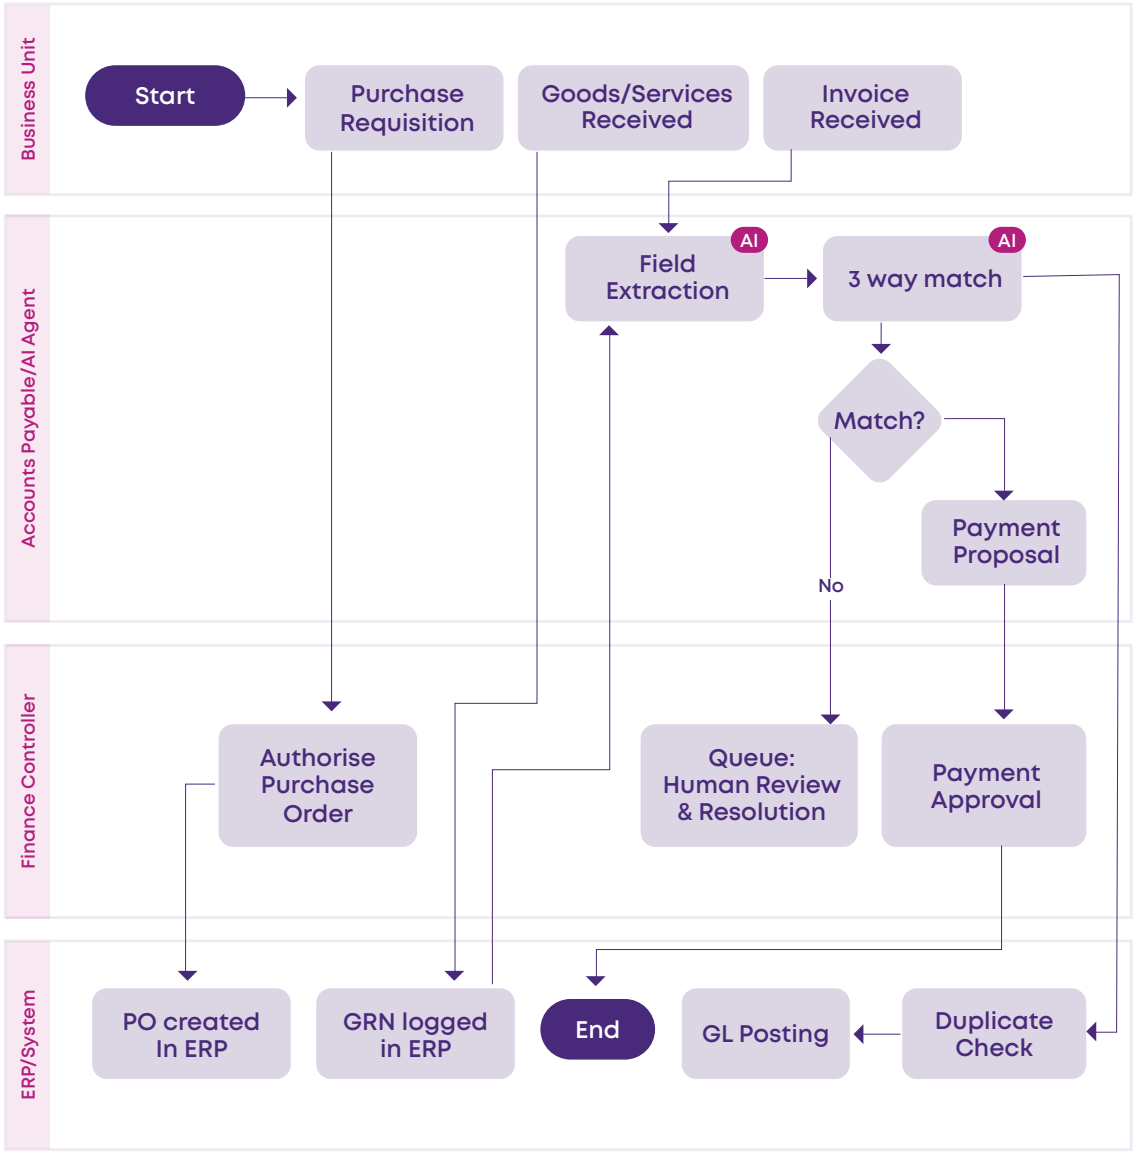

Exhibit — Procure-to-Pay: AI touchpoints mapped to SOX control points (C1–C5).

Order-to-cash — and the question of segregation

In receivables, agents apply cash, prioritize collections, draft outreach, and propose dispute resolutions or credit memos. The messy cases — short payments, bundled remittances, missing references — are where the workflow's authority and review design matter most.

If **one system identity can propose, justify, and post a credit adjustment, it combines duties that the control design should keep separate. Customer-facing messages add immediate** external exposure, so communication authority also belongs inside the control boundary.

Order-to-Cash (O2C) — SOX Process Flowchart with AI Touchpoints & Controls

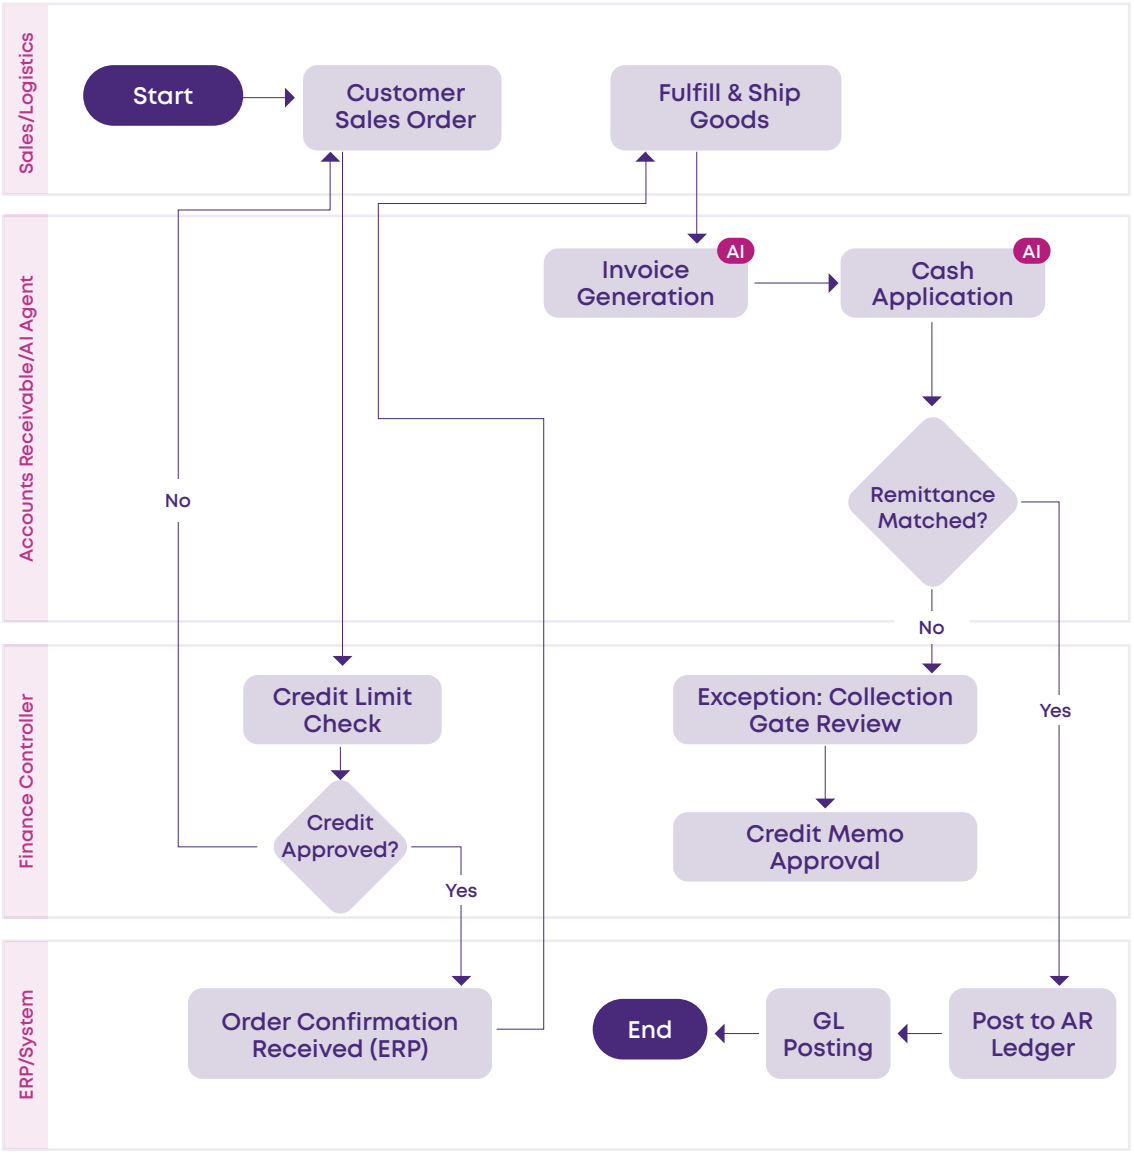

Exhibit — Order-to-Cash: AI touchpoints, control gates, and segregation of duties.

Across the three workflows, the main control questions are the same: who is making the judgment, what authority does the agent have, can it rely on the information it receives and how is segregation of duties preserved?

The same questions arise in close, estimates, disclosures, payroll, inventory, treasury, tax, sales operations, and upstream data services.

They also show why agents cannot be treated simply as faster employees or as another system feature. They challenge four assumptions that traditional control designs rely on:

Judgment has an identifiable human author 01

Control frameworks assume a person interpreted the contract, decided the accrual, assigned the account. When an agent performs the interpretive step and a person approves, judgment is shared. If the approval is only a quick check, the agent has effectively made the judgment.

For each workflow, management should be able to answer: **who is responsible for the judgment, what evidence did they use, and what did the reviewer independently assess?**

Systems behave deterministically 02

Traditional applications are expected to execute set logic consistently. Agentic systems can produce different results from the same or similar inputs. Well-designed workflows bound that variability: rules and reconciliations own calculable attributes, and sufficiently precise independent review addresses judgment. That **containment is a control and must be evidenced.**

Segregation of duties (SoD) can be enforced through roles 03

SoD works by ensuring no one person holds initiation, authorization, recording, and custody together. An orchestrating agent can hold all four unless the architecture prevents it. **User access reviews do not by themselves show what an agent can do through its tools and permissions.**

Change arrives through managed releases 04

Agentic-workflow behavior depends on models, prompts, configuration, tools, data schemas, transformation logic, retrieval sources, and thresholds — several of which may change without a managed release. Treat each **as a governed component,** with approval and sufficient version, input, output, and change evidence.

THE FRAMEWORK

COSO frames the implementation itself as a six-step loop — from establishing governance through to monitoring — that repeats as the agent estate grows. It is the program-build companion to the agent lifecycle above: the framework and lifecycle that follow work through each step in turn.

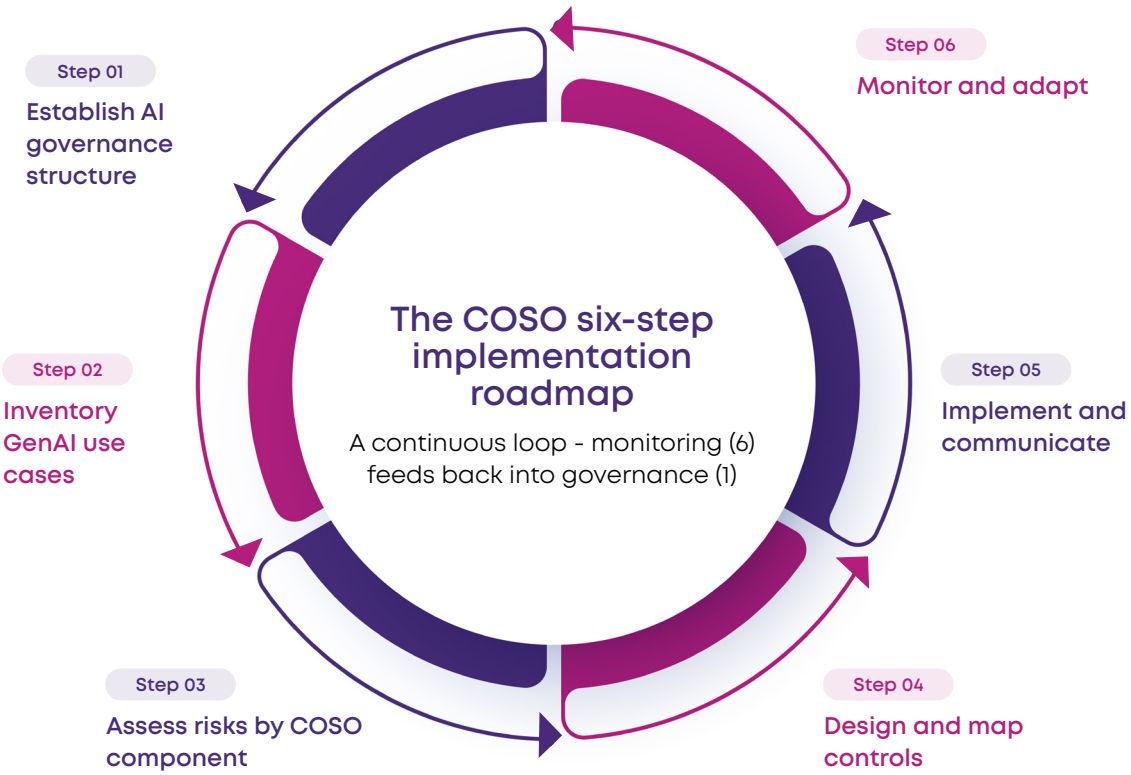

The table below shows where each step is developed in this paper.

Step	Focus in this paper	Reference within this document
01 **Establish AI governance structure**	A governance body with authority to approve deployments; risk appetite and non-negotiable authority limits; second-line ownership and accountability; oversight reporting.	§10 Management governance; §4 COSO principles 3 and 5; §11 Audit committee oversight
02 **Inventory GenAI use cases**	The register — one row per capability per workflow, with documented in- and out-of-scope decisions; the eight capability types as the durable inventory unit.	§5 Inventory and risk assessment; §2 Capability taxonomy
03 **Assess risks by COSO component**	Assertion-level risk first, then capability as the failure-mode lens; the seventeen principles reinterpreted for agentic work; deficiency severity in §7.	§2 Capability taxonomy; §5 Risk assessment; §4 COSO principles
04 **Design and map controls**	The six recurring control patterns; how AI controls relate to ITGCs and application controls; a worked risk-and-control matrix.	§3 Six control building blocks; §2 Rebuilding SoD on capability boundaries
05 **Implement and communicate**	End-to-end qualification, the evaluation library, shadow operation and cutover; communication upward to the audit committee and outward to service organizations.	§6 Qualification before deployment; §11 Audit committee oversight
06 **Monitor and adapt**	Balanced monitoring and self-assembling evidence; the change channels and version testing; retirement and the record; vendor-side change; correlated failure, recovery, and deficiency evaluation — the loop then returns to Step 1.	§7 Operating the control; §8 Change; §9 Retirement

Section 2

Capability, not product: applying the COSO taxonomy

The February 2026 guidance classifies GenAI by what it does, in eight capability types. For finance, this is the working unit of both risk assessment and segregation.

COSO classifies GenAI by what it does. Products change; capabilities persist. Using shorter finance labels for the eight official categories, the sequence is ingestion, transformation, transaction processing and reconciliation, orchestration, judgment, monitoring, knowledge retrieval, and human-AI collaboration. This capability view is precise enough to describe and control any agentic workflow.

Capability	What it is?	Illustrative use across finance workflows	Illustrative ICFR risk drivers
Ingestion	Reading and extracting from documents and feeds	Invoice capture; contract reading; remittance parsing	Incomplete or unauthentic population; manipulated inputs
Transformation	Restructuring and matching data	PO matching; cash application; data mapping	Lost, duplicated, or misclassified records; opaque lineage
Judgment	Proposing conclusions under ambiguity	Accrual identification; account assignment; collection prioritization	Plausible error; bias; displaced independent review
Transaction processing	Creating or changing records in financial systems	Payment proposals; credit memo entries; journal posting	Unauthorized or inaccurate records; asset movement
Orchestration	Sequencing tasks and other agents	End-to-end P2P and close coordination	Segregation collapse; correlated error; uncontrolled tool use
Monitoring	Watching processes and flagging anomalies	Duplicate detection; drift alerts	Self-review; missed drift; alert or escalation failure
Knowledge retrieval	Retrieving and summarizing information	Policy, standards, and regulatory lookup	Stale, incomplete, or unauthorized sources; unsupported citations
Human-AI collaboration	The interface between agent and reviewer	Every approval gate; drafted rationales and outreach	Automation bias; review fatigue; unclear accountability

Start with significant accounts, disclosures, relevant assertions, and likely sources of misstatement; then identify which agent impacts this workflow. Capability is a design lens, not a substitute for assertion-level assessment. “AI in payables” is too broad. “A transaction-processing workflow affecting occurrence, accuracy, and cut-off of cash disbursements, with independent release authorization” is specific enough to design and test. In practice this classification is applied one workflow at a time in the inventory that opens the lifecycle (Section 5): each registered workflow is tagged by capability, and that tag drives its risk assessment and control design.

Rebuild segregation of duties on capability boundaries. The agent that interprets is separated from the one that posts; posting cannot alter reference data; and monitoring is sufficiently independent. In payables, the invoice-processing identity has no vendor-master write access, and payment release sits behind independent authorization and deterministic controls the proposing agent cannot satisfy or alter. In O2C, the proposing agent cannot authorize its own credit memo; posting occurs only after independently controlled approval. Separation is enforced across identities, tools, data, and configuration — not merely stated in policy.

Agentic components can change through more — and less visible — channels than traditional applications, increasing the evidence burden **for continuity: immutable version and configuration records,** change detection, requalification triggers, and operating telemetry.

The fraud lens

The capability map doubles as an update to the fraud risk assessment under COSO Principle 8. Three vectors deserve explicit entries: the manipulated input — an instruction embedded in a document the agent will read and act on; the reference-data path — any agent route, however indirect, to the vendor master or customer master, where a change to banking details is the precise move business-email-compromise fraud depends on; and the adjustment channel — credit memos, write-offs, and unapplied cash proposed at machine speed. Each existed before agents. What changed is the speed, the volume, and the absence of a person in the middle who might notice.

The self-review problem

If the same agent prepares a reconciliation and drafts the review memo, it is producing evidence about its own work. Use evidence from sufficiently independent components and sources: deterministic reconciliation or recalculation where feasible, and sufficiently precise independent review where judgment cannot be reduced to rules

Section 3

The six control building blocks

The same six control patterns appear repeatedly across agentic finance workflows.

01

Access and acceptable-use boundaries

Each production agent or orchestration **layer acts through attributable non-human identities with least privilege; shared human credentials are prohibited. Where one technical identity serves several capabilities,** record-level logs attribute each action to the workflow, capability, version, and approver. Permitted tools and actions are enumerated and controlled.

02

Data, lineage, and input controls

The relevant data perimeter does not stop at the ERP. It extends through external feeds, ingestion and change-data-capture processes, cloud warehouses or lakehouses, ETL/ELT pipelines, semantic-layer transformations, retrieval corpora and indexes, and interfaces back to the ledger. For ICFR-relevant flows, management identifies approved sources and owners; establishes refresh and cut-off rules; reconciles source-to-target counts and control totals; detects schema drift; controls transformation logic; handles data-quality exceptions; restricts access; and retains versioned end-to-end lineage. A lineage diagram is useful evidence only when its own completeness and change controls have been established.

03

Prompt and configuration governance

Prompts, schemas, tolerances, routing thresholds, and retrieval sources are a **configuration, and configuration in a control is change-managed:** versioned, segregated between author and approver, promoted through environments, and logged. An untracked prompt edit is an unauthorized program change that happens to be written in English.

04

Output validation and exception handling

Before an agent-generated amount or conclusion affects a financial record, it passes validation independent of the component that produced it: deterministic reconciliation or recalculation for calculable attributes, and sufficiently precise human or other independent review for judgmental matters. Failures route to owned, aging exceptions with documented disposition.

05

Logging and traceability

Every ICFR-relevant output retains sufficient record-level metadata: **model and version where available, prompt and configuration version, data snapshot and transformation version, tools used, key inputs, actions, validation results, exception disposition, and approvals, subject to appropriate privacy and retention constraints. Management controls the logs' completeness, accuracy, retention, access, time synchronization, and tamper resistance.** This makes a later question — which records did version N or transformation M affect? — answerable.

06

Monitoring for drift, anomalies, and unauthorized use

Monitor a balanced set of control and outcome indicators: value and volume processed, population completeness, validation failures, false negatives found through independent sampling, unresolved-exception aging, unauthorized actions, drift or probe results, incidents, and changes in escalation or override rates. Segment material results by model, prompt or configuration, and data version. Metrics are tolerance-banded, investigated, and linked to response; no single rate proves effectiveness.

How AI controls relate to ITGCs and application controls

AI-specific controls are not a third, mutually exclusive category. Relevant ITGCs govern the enabling environment: access to model endpoints and configuration repositories, authorization and deployment of changes, technology operations, interfaces, incidents, and third-party services. Automated application controls address transaction-level risks through reconciliations, matching, duplicate checks, authorization blocks, and completeness checks.

Controls over model qualification, prompts, retrieval sources, data provenance, drift, autonomous tool use, and human review address AI-specific failure modes. Depending on the objective, they may operate as general controls, automated application controls, or IT-dependent manual or management-review controls. The risk and control objective — not the label — determine how each control is scoped and tested.

Section 4

The key COSO principles for agentic workflows

COSO's seventeen principles remain the framework. The nine below require the greatest reinterpretation for agentic workflows.

Principle	What changes for agentic workflows
3 — Structures, authority, and responsibility	Define and technically enforce authority for non-human actors; **assign** a named accountable owner.
5 — Accountability	Agents cannot be accountable; owners answer for outcomes, evidence, exceptions, and recovery.
8 — Fraud risk	**Add manipulated inputs, indirect paths to master** data, and machine-speed adjustments.
9 — Significant change	Assess model, prompt, data, tool, and vendor changes — including changes the entity did not initiate.
10 — Control activities	Rebuild segregation across capabilities; independently **validate financially relevant outputs in proportion to risk.**
11 — General controls over technology	Extend ITGCs to prompts, policies, models, tools, transformations, retrieval, and non-human identities.
13 — Relevant, quality information	Establish source, completeness, accuracy, precision, and record-level lineage for information used in controls.
16 — Evaluations	Use balanced indicators, controlled probes, independent sampling, and requalification after material change.
17 — Deficiencies	Evaluate crossed boundaries, unqualified change, missed monitoring, and correlated failures individually and in combination.

The remaining principles still apply. The table above identifies where agentic work most changes control design, operation, and evidence.

THE LIFECYCLE

Our view of the agentic workflow follows 5 distinct steps:

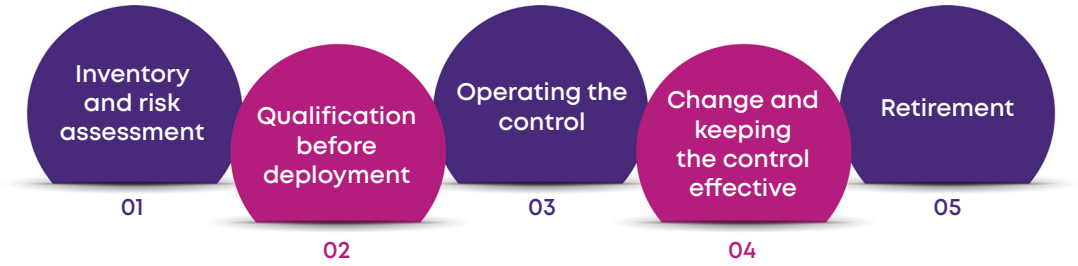

Section 5

Inventory and risk assessment: you cannot control what you have not registered

The lifecycle begins before any design work — with an honest register of what is already running.

This lifecycle operationalizes COSO's six-step roadmap for financially relevant workflows, making change and retirement explicit because control evidence must survive repeated change and decommissioning.

Every agentic control program begins by finding what is already deployed: transformation-team workflows, SaaS features enabled by administrators, departmental experiments that became operational, and embedded vendor capabilities introduced through product updates.

One row per capability per workflow, classified against the eight-capability taxonomy of Section 2, is a practical starting point. Record the purpose; affected accounts, disclosures, risks, and assertions; authority; approved data, transformations, tools, and retrieval sources; model and configuration versions; shared dependencies; control reliance; lifecycle status; and named business, technology, control, and risk-acceptance owners. Use the enterprise register for the full agent population and document both in-scope and out-of-scope ICFR decisions. Periodically test the register for completeness.

Risk assessment then applies that capability lens to each workflow: from the affected accounts, disclosures, assertions, and materiality, it weighs authority, autonomy, reliance, reversibility, fraud susceptibility, volume, observability, and common dependencies. Controls scale accordingly: read-only drafting may need approved-source and review controls; posting, payment, master-data, or external-communication authority requires stronger separation, traceability, validation, and suspension and recovery. Bounded, observable, recoverable workflows generally make better first deployments.

The register is **a living control.** New agents must enter through it and material changes should trigger reassessments. An agent in use but absent from the register is a control gap.

Section 6

Qualification before deployment

Qualification needs to be entity-specific evidence that the complete workflow performs within approved boundaries on the population and conditions on which management will rely.

Familiar ICFR qualification disciplines carry over but must be designed for variability introduced by agents. Before production:

Test the full workflow. Cover data sources and population completeness, transformations, prompts and configuration, tool permissions, independent checks, exception routing, human review, evidence capture, and suspension, rollback, and recovery.

Use a controlled evaluation library. Entity-specific cases with adjudicated answers, stratified to overweight material, rare, adversarial, and historically failure-prone items; version, restrict, and approve the set.

Keep an independent holdout. A separate set, not used to tune the workflow and never exposed to developers, so repeated qualification does not become a test of memorized cases.

Fix acceptance criteria in advance. Approve exit criteria before the evidence arrives, or deadline pressure will move them: agreement by risk stratum, with a higher bar for the high-materiality population; no new categories of error regardless of aggregate score; escalation behavior preserved within tolerance; full schema conformance; error weighted by materiality, not count; and cost within a defined envelope.

Run in shadow before cutover. The agent processes live inputs, its outputs compared with the incumbent process, without touching the ledger, over a risk-based period covering representative activity, high-risk exceptions, and period-end conditions.

Make cutover an evidenced approval. Business-owner sign-off supported by the qualification and shadow-operation evidence, followed by hypercare and re-baselined monitoring.

Treat vendor evidence as input, not proof. A benchmark or assurance report may inform the assessment; neither establishes effectiveness for the entity's control.

Regulation will continue to evolve. Companies can act now by applying existing ICFR, technology, vendor, and risk disciplines in proportion to the workflow's **use** and **risk.**

Section 7

Operating the control: monitoring, evidence, and independent review

Once live, the questions change: is the workflow behaving as qualified, is evidence accumulating by design, and is independent review still operating with sufficient precision?

Three disciplines define the operating stage:

01

Monitor a balanced set, not accuracy alone.

Track the control and outcome indicators from Section 4's monitoring block, segmented by model, prompt or configuration, and data version, and compare period-end activity with ordinary operations. Escalation and override rates remain useful boundary signals, not proof on their own.

02

Let the evidence assemble itself.

Agent output used in performing a control is information whose relevance, source, completeness, accuracy, and precision must be established. Where a model contributes a material figure or conclusion, reliability is demonstrated through independently designed validation appropriate to the risk — reconciliation or recalculation where possible, and sufficiently precise review where judgment exists. The evidence package accumulates as the workflow runs: population reconciliation, version identifiers, validation results, exception dispositions, override rationale, and approvals. **If that package exists only when someone compiles it for the auditor, the process is not evidenced by itself.**

03

Defend the quality of human review

Reviewers need competence, capacity, authority, and direct access to underlying evidence. For significant judgments, **they** document an expectation before reading the agent's rationale and in certain cases independently re-perform.

Evaluating exceptions and deficiencies

When an agentic control exception occurs, management should assess more than whether an error reached the financial statements. Consider the control objective, the affected accounts and assertions, the population and period exposed, the risk of fraud, related control failures, and any compensating controls. Evaluate exceptions together when they arise from a shared model, data source, configuration, or workflow component. A detected error may show that a control worked; an undetected or uninvestigated pattern may indicate that the control environment needs further evaluation.

What this looks like in practice

For an ICFR leader, the test of any framework is whether it survives contact with the risk-and-control matrix. The exhibit below applies the six blocks to payables in control-description form, ready to adapt to the entity's facts.

Risk	Control	Type	Evidence
Invoice population incomplete across intake channels	Deterministic reconciliation of received-to-processed counts by source, daily; unmatched items aged and dispositioned	Automated detective	Reconciliation log; exception dispositions
Agent mismatches an invoice or misassigns the account	Three-way match and tolerance check performed by an independent deterministic routine; failures block or route to an exception queue; payment release remains separately authorized	Automated preventive (ITAC)	Match results tied to each entry; queue dispositions
Manipulated invoice content directs agent behavior	Input screening for embedded instructions; agent action set constrained to enumerated, logged tools	Automated preventive	Screening log; agent entitlement listing
Unauthorized path to the vendor master	Agent identity carries no vendor-master write access; changes to banking details require independent verification and a hold before the next payment run; alert on any agent-sourced attempt	Preventive and detective	Access matrix; change log; hold-and-release record; alert history
Duplicate or erroneous disbursement	Deterministic duplicate screen across vendor, amount, reference, and date window before any payment proposal	Automated preventive	Screen results; blocked-proposal log
Agent behavior changes with model, prompt, or configuration	End-to-end qualification against the regression suite and independent holdout before promotion; documented approval; thresholds re-baselined after investigation	AI qualification/ change control (ITGC-supported)	Qualification report; approval; re-baseline record
Human review degrades into concurrence	Balanced control and outcome scorecard against tolerance bands; periodic independent re-performance of sampled approvals	Monitoring / MRC	Scorecard; investigation records; re-performance results; evidence of reviewer challenge
Schema, mapping, or transformation change produces an incomplete or inaccurate agent population	Each production load reconciles source and target counts and control totals; schema or mapping exceptions block or quarantine processing; transformation changes are tested and approved; period-end data snapshot and lineage identifier are retained	Automated application controls supported by ITGCs	Load reconciliation; schema alerts; approved transformation change; automated tests; lineage snapshot; exception disposition

Section 8

Change: keeping the control effective

The model provider announces that the version at its center retires in nine months. What has to happen next?

Everything to this point could be qualified once and operated indefinitely if the workflow held still. It will not — an agentic workflow changes through five principal channels.

- **Prompts, configuration, and thresholds change** at the entity's initiative, under the configuration governance already described. This channel is manageable because it is visible.
- **A model version may change** on a provider-defined timetable. Govern the version change as a program change: qualify against pre-approved criteria, protect the close calendar, recalibrate thresholds, verify rollback, and **retain prior outputs and evidence for the required period.** A withdrawn version may not remain reproducible, so the historical record must stand on its own.
- **The third channel is feedback-driven change. Even when the base model does** not learn in production, reviewer corrections may update memory, retrieval stores, exemplars, or routing. Define what may update automatically; keep thresholds and authority under formal change control; review accumulated changes; and re-run qualification on a defined cadence.
- **Behavior can also change** behind a stable endpoint. Pin immutable versions where available; otherwise use a controlled probe that compares material attributes and structured conclusions with a baseline, using repeated runs where variability requires them. A probe triggers investigation and can narrow the exposure window; it does not prove the exact change date or cause.
- **An upstream schema, mapping, transformation, semantic layer, or retrieval index changes while the model and prompt remain untouched.** The control response is source-to-target reconciliation, schema-drift detection, versioned transformation logic and lineage, impact assessment across downstream consumers, and requalification when the change can alter a financial conclusion.

The sensible question after an agentic workflow change is whether a new control or a modification to an existing control operation is necessary. Such change does not automatically create a new control but management must assess whether model behavior is integral to how a key control prevents or detects misstatement. Where it is, the pre- and post- change configurations are distinct operating controls; a superseding control must operate for a sufficient period and the external auditor may still test the pre-change control when it matters to the financial-statement audit or control-risk assessment.

Section 9

Retirement and the record

Decommissioning an agent is a controlled event with an evidence obligation, and the lifecycle closes where it opened: at the register.

The COSO six-step roadmap does not name retirement as a discrete step; it sits implicitly within monitor-and-adapt. With models, vendor features, and tools now turning over quickly, decommissioning is **a routine, recurring event rather than a rare one** — which is why this lifecycle makes it an explicit stage with its own evidence obligation.

Agents leave service — workflows are redesigned, vendors change, capabilities consolidate. Retirement done casually leaves two exposures:

Lost evidence: the agent's outputs, version history, exception and override logs, and qualification records support already-issued financial statements – and a decommissioned agent's record often can't be reconstructed, so they must be retained for the full retention period.

Residual access: the agent's credentials, entitlements, and integrations should be revoked with the same discipline as a departing employee's.

GOVERNANCE: MANAGEMENT AND THE AUDIT COMMITTEE

Agentic finance requires two connected layers of governance. Management owns the design and operation of the control environment. The audit committee oversees whether management has set appropriate boundaries, understands the exposure, and responds effectively when risk changes.

Section 10

Management governance

Management should govern agentic workflows as a portfolio, not only one workflow at a time. Individual workflows may have different owners and controls, but they can share models, data pipelines, orchestration tools, identity services, retrieval sources, and monitoring platforms. A failure in one shared component can affect several accounts, controls, and processes at once.

Risk appetite.

Set a short list of non-negotiables and risk-based authority limits. Common starting points include no unregistered agent, no ability to change bank details or master data without independent authorization, and no posting or payment beyond an approved limit. Boundaries are technically enforced and exceptions require enterprise approval. The objective is not a universal ban on autonomy; it is explicit, revocable authority that project teams cannot relax locally.

Ownership of the second line.

Process owners are the first line. Second-line oversight may sit in model risk, technology risk, operational risk, controllership, or another function with the mandate, competence, and independence to challenge the agent portfolio. A smaller company may combine roles, but business ownership, technology operation, control ownership, and independent challenge remain explicit. Internal audit(the third line)then evaluates the lifecycle rather than re-performing the agent's work.

Concentration.

A shared model, data pipeline, orchestrator, identity layer, or monitoring service can affect many workflows at once. Maintain a dependency map, contractual notice, and tested fallback for priority workflows — including current manual close procedures — and consider alternate processing where concentration could make recovery unacceptable.

Section 11

Audit committee oversight

The audit committee does not need to manage the underlying technology. Its role is to challenge whether management's governance is clear, evidence-based, and proportionate to the financial-reporting risk.

Management's reporting should provide six recurring views.

- **The exposure:** built and embedded agents that can affect ICFR, including outside finance, with value and volume processed, autonomy, affected accounts and assertions, owners, and control status.
- **The horizon:** announced model retirements and platform releases, vendor evidence gaps, and the qualification or compensating-control plan for each.
- **The changes:** material model, prompt, configuration, tool, data, transformation, retrieval, and vendor changes, with qualification results and approvers.
- **The incidents:** boundary crossings, detected drift, unauthorized changes, correlated failures, open deficiencies, and remediation status.
- **The assurance:** balanced control and outcome indicators, workflows qualified or requalified, key-control evidence, and tested fallback readiness.
- **The frontier:** embedded agentic features enabled in systems of record, what provider assurance does and does not cover, and how management closes each gap.

These views let the committee ask the right questions: who set the agent's authority; which accounts and controls rely on it; what happens when a shared provider or data pipeline changes; how management knows review remains substantive; and whether suspension and recovery have been tested.

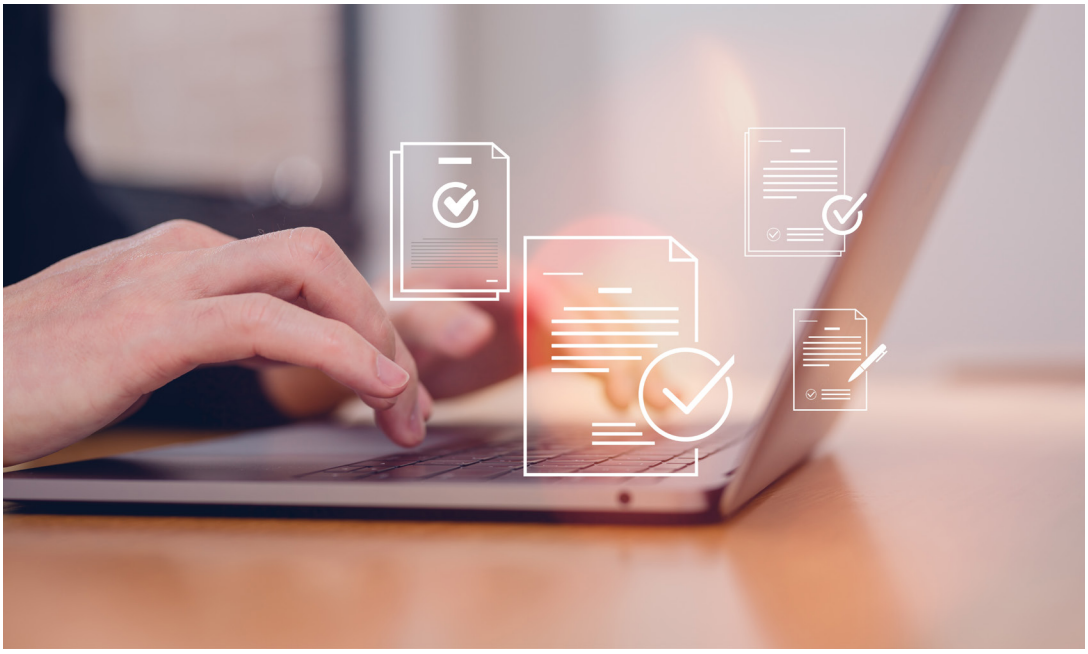

In Closing

QUESTIONS FOR MANAGEMENT

Practice is still evolving in several areas. Management should revisit these questions as technology, control environment, and audit/control expectations develop:

01 Can we rely on a model's explanation of its conclusion? Does the workflow retain the underlying source evidence, independent validation, and a clear record of the action taken? Should this approach change if reasoning traces become verifiable rather than merely persuasive?

02 Should the organization rely on continuous controls monitoring instead of periodic test of controls

03 Is exception monitoring providing enough evidence?

04 Does the available assurance cover the full workflow on which management relies? Which parts of the end-to-end workflow are covered by provider reports, and which parts remain management's responsibility?

05 Is the evaluation library large and representative enough? What test population is needed to support the control objective and the intended level of reliance?

06 When does an observed error rate become an ICFR concern? What is the difference between a model error, a control deviation, and a financial-statement misstatement in this workflow?

07 What happens when the system of record itself becomes agentic? What is the line between the system of record and the activity performer in that scenario and will the entity's register capture agents it did not build? How will assurance (SOC-1/ SOC-2 reports) cover such risk and how will management fill in the gaps?

Eight questions establish whether an organization is governing agentic workflows or merely running them.

- Can we produce, today, a register of every built or embedded agent that can affect ICFR — including agents outside finance — with capability, authority, ownership, approved data and transformations, and shared dependencies for each?
- For each agent that can post, pay, change master data, or communicate externally, can we state its authority limits and show the technical enforcement and independent authorization?
- Are end-to-end qualification criteria approved before testing, supported by an independent holdout, and repeated after relevant change?
- Could we trace any financially relevant output to its model, prompt, configuration, data snapshot, transformation, retrieval source, tools, validation, exceptions, and approval?
- Would we detect an unauthorized change or correlated failure in a model, prompt, shared component, schema, transformation, retrieval source, or embedded platform — and could we identify affected workflows, records, and recovery actions?
- Does a balanced scorecard reveal displaced human judgment, false negatives, aging exceptions, unauthorized actions, and control drift — and did management investigate material movement?
- If the reviewer of an agentic workflow were challenged to evidence their challenge — not their approval — could they?
- What feature-specific assurance do vendors provide, where are the evidence or continuity gaps, and how does management compensate or limit authority?

Organizations that can answer these eight with current evidence have turned agent adoption into a governed capability.

SOURCES, METHOD, AND ATTRIBUTION

This publication is a practitioner synthesis prepared from primary materials available through 17 August 2026. The three finance workflows are illustrative composites informed by advisory experience; they do not describe a single client or establish prevalence. The control patterns, interpretations, and recommendations are Uniqus Consultech's own.

COSO's Achieving Effective Internal Control Over Generative AI (GenAI) is the principal organizing source. Its five components, seventeen principles, eight capability types, and six-step roadmap are applied to financial reporting and extended here to end-to-end agentic workflows, authority boundaries, evidence, change, and recovery. The SEC's ICFR rule and PCAOB AS 2201 anchor the discussion of ICFR scope, significant accounts and disclosures, relevant assertions, control selection, and deficiency assessment.

Regulatory-horizon references are deliberately limited. Federal Reserve SR 26-2 is banking supervisory guidance; its attached interagency guidance expressly excludes generative and agentic AI from scope while noting that existing risk-management and governance practices should inform controls for out-of-scope tools. The FSB's June 2026 document is a consultation report proposing twelve sound practices, not a final standard. Both are used as directional evidence, not generally applicable requirements.

- COSO, Achieving Effective Internal Control Over Generative AI (GenAI), released February 23, 2026.
- PCAOB, AS 2201: An Audit of Internal Control Over Financial Reporting That Is Integrated with An Audit of Financial Statements, as in effect at August 17, 2026.
- SEC, Management's Report on Internal Control Over Financial Reporting and Certification of Disclosure in Exchange Act Periodic Reports, issued June 5, 2003.
- Federal Reserve, SR 26-2: Revised Guidance on Model Risk Management, April 17, 2026.
- Financial Stability Board, Sound Practices for Responsible Adoption of Artificial Intelligence (AI): Consultation Report, June 10, 2026.

Examples, control designs, and policy positions are the authors' interpretations. Organizations should adapt them to their facts, reporting obligations, risk profile, and applicable law.

Amrita Kapoor, Internal Audit, Snowflake, contributed perspectives on controls over the data layer, correlated failures and recovery, testing across model versions, and the relationship between AI controls, ITGCs, and application controls. Her affiliation is listed for identification only; the views expressed do not necessarily represent Snowflake Inc.

Uniqus Consultech is an AI-enabled global consulting firm specializing in Accounting & Reporting Consulting, Governance, Risk & Compliance, Sustainability, Tech Consulting, and Valuations. Uniqus provides neither audit nor tax services. This publication offers general information, not advice on specific facts.

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 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.

About Our Governance, Risk & Compliance Services:

Our Governance, Risk & Compliance (GRC) practice partners with organizations as they continue to deal with the dynamic governance, risk, and compliance landscape and navigate today's complex business environment. We seek to prepare business leaders and their teams to identify and manage risks, strengthen governance across levels, drive end-to-end compliance, design and test internal controls, anticipate and assess any disruptions, and help build business resilience. We help to deploy platforms aimed at digitization, automation, and AI-enablement of an organization's GRC functions.

A Team That You Can Trust To Deliver

Sandip Khetan
Co-Founder, Global Head of Accounting & Reporting Consulting

Anthony Vitale
Partner, Accounting & Reporting Consulting

Arda Kaya
Partner, Accounting & Reporting Consulting

Ankur Gupta
Partner, GRC

Avinash Ramkumar
Partner, Accounting & Reporting Consulting

Nagaraj Uchil
Partner, Global Head of GRC

Sharad Chaudhry
Partner, Accounting & Reporting Consulting

Vijay Sampathkumar
Partner, Tech Consulting

Harsimran Singh Sangha
Managing Director, Accounting & Reporting Consulting

Marshall Crockett
Managing Director, Tech Consulting

Matt Solomon
Managing Director, GRC

Shival Sehgal
Managing Director, Accounting & Reporting Consulting

Shoichi Ohno
Managing Director, Accounting & Reporting Consulting

Scan this code to find more content like this

Visit us at www.uniqus.com
or follow us on