

Vantage Finance & Tech

Precision Finance. Agile Delivery. Integrated Solutions.



Modern PMO Guide for ERP Cloud Implementations

Project Management in the Age of Agentic AI — SAP S/4HANA & Oracle Fusion ERP

A practical playbook for PMO leaders, program directors and steering committees running cloud ERP transformations that combine traditional program discipline with agentic AI capabilities — covering governance, phase-by-phase activities, AI-specific controls and value realization for both SAP S/4HANA and Oracle Fusion ERP programs.

Prepared by Vantage Finance & Tech · Precision Finance. Agile Delivery. Integrated Solutions.



Table of Contents

1.	Executive Summary
2.	Introduction & Purpose
3.	The Modern PMO: From Governance to Value Enablement
4.	Agentic AI in ERP Programs — What Changes for the PMO
5.	Program Governance & Organizational Structure
6.	Phase 1 — Discover & Prepare
7.	Phase 2 — Explore (Design)
8.	Phase 3 — Realize (Build & Configure)
9.	Phase 4 — Deploy (Test & Cutover)
10.	Phase 5 — Run & Hypercare
11.	The Agentic AI Workstream
12.	RAID Management for AI-Augmented Programs
13.	Quality Assurance & Testing
14.	Organizational Change Management & Training
15.	PMO Tooling, Cadence & Reporting
16.	Value Realization & KPIs
17.	Risk Register — Common Risks & Mitigations
18.	Closing Recommendations



1. Executive Summary

Cloud ERP programs — whether built on SAP S/4HANA or Oracle Fusion ERP — are no longer purely technology migrations. The newest wave of implementations embeds **agentic AI**: autonomous or semi-autonomous software agents that execute multi-step tasks (invoice matching, exception triage, forecasting, narrative reporting) inside the ERP's transactional fabric, often through SAP's Joule and Joule Agents or Oracle's AI Agent Studio and embedded AI agents in Fusion Applications. This shift changes what a Project Management Office (PMO) must govern, test and sustain.

This guide sets out a practical, phase-by-phase model for running such programs. It layers an **agentic AI workstream** on top of the established SAP Activate and Oracle Cloud Success Navigator / OUM-aligned methodologies, so that PMO leaders can plan, staff, govern and report on AI-enabled scope with the same rigor traditionally applied to functional and technical scope.

Key takeaways for PMO and steering committee audiences:

- Agentic AI scope needs its **own backlog, governance gate and risk register entries** — it cannot be folded silently into functional workstreams without losing visibility.
- Human-in-the-loop checkpoints, agent permission boundaries and model/version change control are **new PMO deliverables**, not just IT security concerns.
- Testing must expand beyond functional/UAT/regression to include **agent behavior testing**: prompt regression, hallucination/exception handling, and confidence-threshold tuning.
- Value realization tracking should separate **process efficiency gains from AI-driven gains**, since steering committees increasingly ask for this distinction.
- The core program discipline — RAID logs, RACI, cutover rehearsals, hypercare — remains essential and unchanged in principle, but each artifact needs an AI-aware extension.

This guide assumes a program-level PMO function overseeing one or more cloud ERP workstreams. It is written to be platform-agnostic where possible, with explicit SAP S/4HANA and Oracle Fusion ERP callouts where the methodology, tooling or terminology diverges.



2. Introduction & Purpose

Both major cloud ERP platforms have converged on a similar architectural idea: a transactional core enriched with embedded, conversational and increasingly autonomous AI capabilities. **SAP S/4HANA Cloud** delivers this through Joule (the SAP AI copilot) and Joule Agents built on SAP's agentic AI framework, integrated with SAP Business AI and the Business Technology Platform (BTP) AI Core/AI Foundation services. **Oracle Fusion ERP** delivers this through embedded AI agents across Financials, Procurement and Supply Chain modules, Oracle AI Agent Studio for building custom agents, and the broader Oracle Cloud Infrastructure (OCI) Generative AI services.

For the PMO, this convergence means that an ERP implementation roadmap can no longer treat “AI” as a late-stage enhancement. Agentic capabilities increasingly arrive switched on by default in standard cloud releases, which means governance decisions about which agents are activated, for which processes, with what level of autonomy, must be made **during design**, not after go-live.

2.1 Purpose of This Guide

- Provide PMO leaders with a complete, phase-aligned activity checklist for ERP cloud programs that include agentic AI scope.
- Define the additional governance artifacts, roles and controls that agentic AI introduces.
- Offer a reference structure for risk, quality and change management adapted to AI-augmented processes.
- Remain useful regardless of whether the underlying platform is SAP S/4HANA, Oracle Fusion ERP, or a dual-platform landscape.

2.2 Who Should Use This Guide

- Program/Project Directors and PMO Leads accountable for overall delivery governance.
- Steering committee members who need a shared vocabulary for AI-related decisions.
- Functional and technical workstream leads who must integrate AI activities into their plans.
- Risk, audit and internal control stakeholders evaluating agentic AI exposure.



3. The Modern PMO: From Governance to Value Enablement

The traditional PMO mandate — scope, schedule, budget, risk — remains the backbone of any ERP program. What has changed is the breadth of what the PMO must make visible to leadership. A modern cloud ERP PMO operates as a **value enablement function**: it does not just track delivery, it actively curates the backlog of process, data and AI-enabled improvements against a continuously evolving cloud release cadence.

3.1 Three Shifts Driving the Modern PMO Model

- **From big-bang to continuous delivery** — SAP S/4HANA Cloud and Oracle Fusion ERP both ship quarterly (or more frequent) feature releases. The PMO must run a lightweight, recurring release-adoption process alongside the main implementation timeline.
- **From static functional scope to a living AI backlog** — agentic capabilities are added, deprecated and re-scoped by the vendor between releases; the PMO needs a mechanism to triage these continuously rather than freeze AI scope once at the start.
- **From IT-led governance to cross-functional AI governance** — decisions about agent autonomy, data access and human oversight require Finance, Risk/Audit, Data Privacy and business process owners at the same table as IT.

3.2 Core PMO Functions (Unchanged but Re-scoped)

PMO Function	What Changes With Agentic AI Scope
Scope & backlog management	Adds a parallel AI/agent backlog with its own intake, prioritization and acceptance criteria.
Schedule management	Adds AI-specific milestones: agent design review, guardrail sign-off, agent UAT, autonomy-level go-live gate.
Budget & resource management	Adds AI engineering/prompt-design effort, model governance resourcing, and ongoing AI operating cost tracking.
Risk & issue management	Adds AI-specific risk categories (see Section 12) and a faster escalation path for agent behavior issues.
Quality management	Adds agent behavior testing and confidence/threshold tuning to the standard test strategy.
Stakeholder & change management	Adds AI literacy and trust-building communications, and human-in-the-loop role redesign.



4. Agentic AI in ERP Programs — What Changes for the PMO

Agentic AI differs from earlier generations of embedded analytics or simple chatbots because agents can **plan and execute multi-step actions** across modules — for example, investigating a blocked invoice, proposing a resolution, and posting the correction once a confidence threshold and approval policy are met. This materially changes the PMO's risk and control posture.

4.1 Platform Reference Points

Capability	SAP S/4HANA	Oracle Fusion ERP
AI copilot / assistant	Joule	Embedded AI / Oracle Digital Assistant
Agent framework	Joule Agents, SAP Build / BTP AI Foundation	Oracle AI Agent Studio
Underlying AI platform	SAP AI Core, SAP AI Foundation, partner LLMs	OCI Generative AI Service
Typical agent use cases	Finance close tasks, procurement exception handling, service ticket triage	Supplier risk monitoring, payables exception resolution, demand forecasting narratives

4.2 New PMO Concerns Introduced by Agentic AI

- **Autonomy levels** — each agentic process must be classified (e.g., advisory-only, approve-then-execute, fully autonomous within bounds) and this classification must be a tracked program artifact, not an undocumented configuration setting.
- **Agent inventory / registry** — the PMO should maintain a living register of every active agent, its scope, owning process, data access and current autonomy level.
- **Guardrails and escalation paths** — what happens when an agent is uncertain, when it produces an unexpected result, or when it should hand off to a human must be defined and tested before go-live.
- **Auditability** — every agent action affecting financial or operational data needs a traceable decision log, which becomes a non-negotiable testing and cutover requirement.
- **Continuous re-validation** — because vendors update underlying models and agent behavior between releases, the PMO needs a recurring (not one-time) validation cadence post go-live.

PMO principle: treat every agentic AI capability as a process change with a software component, not a software feature with a process footnote. Scope, test, and govern it accordingly.



5. Program Governance & Organizational Structure

A cloud ERP program with agentic AI scope benefits from a governance structure that adds one dedicated forum and one dedicated role to the conventional steering committee / PMO / workstream model.

5.1 Governance Bodies

Body	Mandate	Cadence
Steering Committee	Overall program direction, budget, major scope and go-live decisions	Monthly / at stage gates
PMO Core Team	Day-to-day schedule, risk, budget, cross-workstream coordination	Weekly
AI Governance Board	Reviews and approves agent autonomy levels, guardrails, data access and ethical/compliance considerations before each agent goes live	Bi-weekly during build, monthly post go-live
Workstream Leads Forum	Functional/technical alignment, dependency management, backlog grooming	Weekly
Change Advisory Board (CAB)	Reviews configuration, integration and agent changes post go-live	Weekly post go-live

5.2 Key Roles

- **Program/Project Director** — overall accountability for scope, schedule, budget and benefits realization.
- **PMO Lead** — owns the governance cadence, RAID log, reporting and methodology compliance.
- **AI/Agentic Solution Owner** (new role) — accountable for the agent backlog, autonomy classifications and coordination with the AI Governance Board; typically sits within the PMO or reports directly to the Program Director.
- **Functional/Process Workstream Leads** — own end-to-end process design including which steps are AI-assisted or AI-executed.
- **Data & Integration Lead** — owns master/transactional data quality and integration architecture, which underpins agent reliability.
- **Risk, Audit & Compliance Representative** — participates in the AI Governance Board and validates control design for AI-executed transactions.
- **Organizational Change Manager (OCM) Lead** — owns communications, training and adoption, including AI-specific trust-building activities.



6. Discover & Prepare

SAP Activate: Discover/Prepare · Oracle: Initiate/Plan (Cloud Success Navigator)

Overview

This phase establishes the business case, program charter, governance structure and high-level scope, including an explicit first pass at which processes are realistic candidates for agentic AI based on data readiness and process maturity.

Core PMO & Delivery Activities

- Develop and approve the program business case and value case, including baseline KPIs.
- Establish program charter, governance structure (Section 5) and RACI.
- Define rollout strategy (big-bang, phased, template-first) and target system landscape (Greenfield, Brownfield, Selective Data Transition for SAP; or Implement/Adopt approach for Oracle Fusion).
- Mobilize the core team; procure environments/tenants and licenses.
- Conduct initial data quality and integration landscape assessment.
- Build the master program plan, RAID log, and communication plan.

Agentic AI Workstream Activities

- Conduct an AI opportunity assessment: inventory candidate processes for agentic AI (e.g., invoice exception handling, supplier risk triage, close-task automation) and score them by data readiness, risk and value.
- Stand up the AI Governance Board and agree its charter, including escalation thresholds.
- Confirm platform AI licensing/entitlements (e.g., SAP Joule/BTP AI Foundation credits, Oracle OCI Generative AI Service capacity) as part of procurement.
- Draft initial AI guardrail principles (data privacy, human-in-the-loop defaults, explainability requirements) for steering committee endorsement.

Key Deliverables

- Approved business case and value case
- Program charter and governance model
- Master program plan and RAID log baseline
- AI opportunity heat map and AI Governance Board charter

Stage gate: Steering Committee approval of business case, program charter and AI governance charter before Explore/Design begins.



7. Explore (Design)

SAP Activate: Explore · Oracle: Implement — Define & Design

Overview

Fit-to-standard (SAP) or Conference Room Pilot / Fit-Gap (Oracle) workshops confirm how standard cloud functionality, including AI agents, will support target processes. This is the phase where autonomy levels and guardrails are formally designed, not just discussed.

Core PMO & Delivery Activities

- Run fit-to-standard / fit-gap workshops by process area; log gaps and configuration decisions.
- Finalize organizational structure (company codes/business units, cost centers, profit centers, ledgers).
- Confirm integration architecture (middleware, APIs, event-driven interfaces).
- Define data migration strategy and object scope.
- Build detailed solution design documents per process area.

Agentic AI Workstream Activities

- For each candidate AI process, document the agent's scope, trigger events, data inputs, decision logic boundaries and autonomy level (advisory / approve-then-execute / autonomous-within-bounds).
- Design human-in-the-loop checkpoints and exception-handling/escalation paths for each agent.
- Define the agent inventory/register structure (owner, process, data scope, autonomy level, review cadence).
- Design audit logging and traceability requirements for agent-executed transactions, in coordination with Risk/Audit.
- Obtain AI Governance Board sign-off on the design for each in-scope agent before build begins.

Key Deliverables

- Solution design documents (functional and technical)
- Agent design specifications and autonomy classification per agent
- Data migration strategy and object list
- Updated RAID log and refined program plan

Stage gate: AI Governance Board sign-off on every agent's design and autonomy level before that agent enters the Realize backlog.



8. Realize (Build & Configure)

SAP Activate: Realize · Oracle: Implement — Configure & Test

Overview

Iterative configuration, integration build and data migration development proceed alongside agent build and unit testing. PMO discipline here focuses on sprint/iteration governance and keeping the AI backlog synchronized with functional build.

Core PMO & Delivery Activities

- Execute iterative configuration and build sprints per process area, with regular demo/playback to business process owners.
- Develop integrations, extensions and custom developments (e.g., SAP BTP extensions, Oracle Visual Builder/PaaS extensions).
- Build and unit-test data migration programs; execute mock data loads.
- Maintain the integrated backlog (functional, technical, AI) with sprint-level burn-down reporting to the PMO.
- Conduct iterative internal quality reviews / playback sessions.

Agentic AI Workstream Activities

- Build and configure agents per approved design, including prompt/decision logic, guardrails and escalation routing.
- Unit test each agent against defined scenarios, including edge cases and intentionally ambiguous inputs.
- Establish the agent inventory/register in its operational form and populate it as agents are built.
- Begin drafting agent-specific runbooks (what the agent does, what it hands off, and how to disable it if needed).
- Validate data quality thresholds required for reliable agent operation (agents are materially more sensitive to master data quality than static reports).

Key Deliverables

- Configured and unit-tested system (functional, technical, AI agents)
- Updated integrated backlog with sprint metrics
- Agent inventory/register (build-stage version)
- Draft agent runbooks and data quality assessment results

Stage gate: Build complete and unit/integration tested for both functional scope and agent scope before entering formal test cycles.



9. Deploy (Test & Cutover)

SAP Activate: Deploy · Oracle: Implement — Test, Train, Transition

Overview

Formal test cycles (SIT, UAT, regression, performance) run alongside dedicated agent behavior testing. Cutover planning must explicitly sequence when each agent is switched on relative to go-live, since not every agent needs to be live on day one.

Core PMO & Delivery Activities

- Execute System Integration Testing (SIT), User Acceptance Testing (UAT), regression and performance/load testing.
- Conduct cutover rehearsals (mock cutover 1, 2, ...) including full data migration dry-runs.
- Finalize training materials and deliver end-user training.
- Build and test the hypercare support model and command-center structure.
- Secure go/no-go decision at each rehearsal and the final cutover gate.

Agentic AI Workstream Activities

- Execute dedicated agent behavior testing: prompt/decision regression testing, confidence-threshold tuning, and adversarial/edge-case testing.
- Test human-in-the-loop handoffs end-to-end (agent escalates → human reviews → outcome logged).
- Validate audit log completeness and traceability for every agent-executed transaction type in scope.
- Decide and document the go-live autonomy posture per agent (e.g., start advisory-only even if designed for higher autonomy, then step up post-hypercare).
- Brief the hypercare team specifically on agent monitoring, override and rollback procedures.

Key Deliverables

- Test summary reports (SIT/UAT/regression/performance)
- Agent behavior test report and go-live autonomy decision log
- Cutover plan and rehearsal sign-offs
- Hypercare support model and agent monitoring runbook

Stage gate: Final go/no-go decision by the Steering Committee, with the AI Governance Board separately confirming go-live autonomy levels for each agent.



10. Run & Hypercare

SAP Activate: Run · Oracle: Operate / Continuous Innovation

Overview

Post go-live stabilization extends beyond traditional functional hypercare to include a structured period of agent monitoring before any autonomy level is increased. The PMO transitions toward a continuous-improvement operating model aligned to the vendor's quarterly release cadence.

Core PMO & Delivery Activities

- Run hypercare command center with defined severity levels and escalation paths.
- Track and resolve post-go-live incidents and defects; monitor system performance.
- Conduct a formal hypercare exit review against pre-agreed stability criteria.
- Transition to steady-state support model (internal team, AMS partner, or hybrid).
- Establish a recurring release-adoption process for quarterly cloud updates.

Agentic AI Workstream Activities

- Monitor agent performance against defined KPIs (accuracy, escalation rate, time saved, false-positive rate) during a defined observation window before increasing autonomy.
- Hold a formal AI Governance Board review before any agent's autonomy level is stepped up from its go-live posture.
- Establish a recurring (e.g., monthly) agent re-validation cycle to catch behavior drift after vendor model/release updates.
- Maintain and continuously update the agent inventory/register as part of steady-state operations, including decommissioning agents that underperform.
- Feed agent performance data into the value realization report (Section 16).

Key Deliverables

- Hypercare exit report
- Steady-state support model and SLAs
- Agent performance monitoring dashboard
- Updated agent inventory/register (operational version)



11. The Agentic AI Workstream

Rather than scattering AI-related activities across functional workstreams only, leading PMOs stand up a horizontal **Agentic AI Workstream** that owns the artifacts below across every phase of the program.

11.1 Workstream Charter

- Owns the AI/agent backlog from intake through retirement.
- Maintains the agent inventory/register as the single source of truth.
- Coordinates agent design, build and test activities across functional teams (an agent rarely belongs to a single functional silo — e.g., an invoice-exception agent touches Procurement, Finance and Master Data).
- Acts as secretariat to the AI Governance Board.
- Owns AI-specific risk register entries and escalates to the PMO core team as needed.

11.2 Agent Inventory / Register — Recommended Fields

Field	Purpose
Agent ID & name	Unique identifier for tracking across design, build and operations
Owning process & business owner	Accountability for outcomes
Platform / framework	e.g., SAP Joule Agent, Oracle AI Agent Studio agent
Data scope & access	Which objects/tables/APIs the agent can read or write
Autonomy level	Advisory / approve-then-execute / autonomous-within-bounds
Guardrails & escalation rule	What triggers human handoff
Audit log location	Where agent decisions/actions are traceable
Performance KPIs	Accuracy, escalation rate, cycle-time reduction
Review cadence & next review date	Recurring validation schedule
Status	Designed / built / testing / live / suspended / retired

11.3 Autonomy Classification Model

- **Advisory** — the agent recommends an action; a human makes the decision and executes it.
- **Approve-then-execute** — the agent prepares the action and a human approves before the system executes it.
- **Autonomous-within-bounds** — the agent executes directly within pre-approved thresholds (e.g., value limits, vendor whitelists) and only escalates exceptions.

Recommended PMO default: launch every new agent at the most conservative autonomy level feasible, and require a documented Governance Board decision — supported by hypercare performance data — to step up autonomy. Stepping down should always be a same-day operational decision, not a governance-board agenda item.



12. RAID Management for AI-Augmented Programs

The standard RAID (Risks, Assumptions, Issues, Decisions) log remains the PMO's central control tool. Agentic AI scope introduces categories that deserve explicit tagging so they are visible to the AI Governance Board as well as the PMO core team.

12.1 AI-Specific Risk Categories to Track

Category	Example
Data quality risk	Agent relies on master data fields with known quality gaps, leading to incorrect autonomous actions
Model drift risk	Vendor updates the underlying model/release, changing agent behavior without explicit program testing
Over-reliance risk	Process owners stop reviewing agent recommendations, eroding the human-in-the-loop control
Audit/traceability risk	Agent-executed transactions lack sufficient logging for internal or external audit
Scope creep risk	Business users request expanded agent autonomy informally, bypassing the AI Governance Board
Vendor dependency risk	Agent functionality depends on a specific vendor AI service tier or roadmap commitment
Change fatigue risk	Parallel functional and AI change overwhelms end users' capacity to adopt new ways of working

12.2 Decision Log Discipline

- Every autonomy-level decision (initial assignment, step-up, step-down, retirement) should be logged as a formal Decision with date, approver and supporting rationale/evidence.
- Decisions should be cross-referenced to the agent inventory/register entry they affect.
- A lightweight monthly digest of AI-related decisions should be included in steering committee reporting, even outside formal stage gates.



13. Quality Assurance & Testing

The test strategy for an agentic AI-enabled ERP program extends the conventional test pyramid (unit → integration → system → UAT → regression) with a dedicated agent behavior testing track that runs in parallel from Realize through Deploy.

13.1 Standard Test Cycles (Unchanged)

- Unit testing of configuration, extensions and integrations.
- System Integration Testing (SIT) across end-to-end business processes.
- User Acceptance Testing (UAT) with business process owners.
- Performance/load testing and regression testing ahead of go-live and each subsequent release.

13.2 Agent Behavior Testing (New Track)

- **Scenario testing** — validate agent decisions against a representative library of historical transactions with known correct outcomes.
- **Edge-case and adversarial testing** — deliberately feed ambiguous, incomplete or conflicting data to confirm the agent escalates rather than guesses.
- **Confidence-threshold tuning** — calibrate the point at which the agent should escalate to a human versus act autonomously, and document the chosen thresholds.
- **Regression testing after model/release updates** — re-run the scenario library whenever the vendor updates the underlying AI service or quarterly release.
- **Bias and fairness checks** — for agents touching vendor selection, credit decisions or workforce-related processes, validate outcomes do not systematically disadvantage any group.

Build the scenario test library once and treat it as a reusable regression asset — it is the single most valuable artifact for keeping agent behavior testing affordable across every subsequent quarterly release cycle.



14. Organizational Change Management & Training

AI-augmented processes change not just *how* work is done but *who* is accountable for outcomes when a machine has taken the first action. Change management must address this directly rather than treating AI as an invisible efficiency upgrade.

14.1 Core OCM Activities

- Stakeholder impact assessment and change readiness diagnostics, run at least twice (Explore and pre-Deploy).
- Role and job-design updates reflecting new human-in-the-loop responsibilities.
- Communication plan covering both functional changes and AI-specific changes.
- End-user training curriculum and delivery, including role-based training paths.
- Super-user network establishment per business unit/plant/region.

14.2 AI-Specific Change Management Activities

- **Trust-building communication** — explain what each agent does, what it does not do, and how to override or escalate, in plain language for affected process owners.
- **“Why this recommendation” training** — train users to interpret and challenge agent outputs rather than rubber-stamp them, particularly during the early autonomy-stepped-down period.
- **Escalation drills** — run simulated exception scenarios during training so users practice the human-in-the-loop handoff before go-live.
- **Feedback channel** — give end users a simple, visible way to flag agent recommendations they disagree with, feeding directly into the agent inventory/register and the AI Governance Board's review cadence.



15. PMO Tooling, Cadence & Reporting

15.1 Recommended Tooling Stack

- Program/portfolio management tool for schedule, RAID and resource tracking (e.g., Smartsheet, MS Project, Jira/Confluence for agile workstreams).
- Test management tool integrated with the platform's native testing tools (e.g., SAP Cloud ALM Test Management, Oracle Application Test Suite) plus a dedicated repository for the agent scenario test library.
- Agent inventory/register, maintained as a structured, governed dataset (not a static spreadsheet) accessible to the AI Governance Board.
- Dashboarding/BI tool for executive reporting, blending program RAG status with agent performance KPIs.

15.2 Governance Cadence Summary

Forum	Frequency	Primary Output
Steering Committee	Monthly + stage gates	Go/no-go decisions, escalations, budget approval
AI Governance Board	Bi-weekly (build) / Monthly (run)	Agent design sign-off, autonomy decisions
PMO Core Team	Weekly	RAID review, schedule/budget status
Workstream Leads Forum	Weekly	Dependency management, backlog grooming
Change Advisory Board	Weekly (post go-live)	Change/release approval



16. Value Realization & KPIs

Steering committees increasingly expect the PMO to separate traditional process efficiency gains from gains specifically attributable to agentic AI. A clean baseline-to-actual tracking model, established during Discover, makes this possible.

16.1 Representative KPI Set

KPI Category	Example Metrics
Program delivery health	Schedule variance, budget variance, defect density, go-live readiness score
Process efficiency (general)	Cycle time reduction, touchless transaction rate, error rate reduction
AI-specific value	Agent-handled transaction volume, agent accuracy rate, escalation rate, time saved per agent-handled transaction
Adoption & trust	User satisfaction with agent recommendations, override rate, training completion rate
Risk & control	Audit findings related to agent-executed transactions, time-to-detect agent behavior drift

16.2 Reporting Practice

- Baseline every KPI before go-live using pre-implementation process data.
- Report a rolling trend, not just a point-in-time snapshot, since agent accuracy typically improves over the first few months as thresholds are tuned.
- Present AI-attributable value separately from general process redesign value in steering committee packs, with a clear methodology note on how the split is calculated.



17. Risk Register — Common Risks & Mitigations

Risk	Mitigation
Agent scope expands informally without Governance Board review.	Enforce a single intake channel for new agent requests; block configuration changes that increase autonomy without a logged Governance Board decision.
Data quality issues cause agents to take incorrect autonomous actions.	Gate agent go-live on data quality thresholds defined during Explore; monitor data quality continuously, not just at cutover.
Vendor releases change agent behavior unexpectedly.	Maintain the reusable scenario test library (Section 13.2) and run it after every quarterly release before accepting new functionality into production.
End users disengage from reviewing agent recommendations over time.	Track override/review rates as a KPI; refresh training and escalation drills periodically, not only at go-live.
Parallel functional and AI change overwhelms adoption capacity.	Sequence and pace AI capability rollout deliberately; avoid activating multiple new agents in the same cutover wave as major functional changes.
Audit cannot trace agent-executed transactions.	Validate audit log completeness as an explicit UAT exit criterion, not an afterthought.
Dual-platform programs (SAP + Oracle) develop inconsistent AI governance.	Run a single, platform-agnostic AI Governance Board and agent inventory/register across both platforms, even if build teams are separate.



18. Closing Recommendations

- **Stand up the Agentic AI Workstream and AI Governance Board on day one** — retrofitting AI governance after build begins is materially harder than designing it in from Discover.
- **Keep the agent inventory/register as a living, governed artifact**, not a one-time design document — it is the backbone of ongoing AI governance.
- **Default to conservative autonomy at go-live** and earn the right to increase it with evidence from hypercare monitoring.
- **Invest in a reusable agent scenario test library** early — it pays for itself across every subsequent quarterly release cycle.
- **Report AI value separately from general process improvement value** so steering committees can make informed investment decisions about expanding AI scope.
- **Keep core PMO discipline intact** — RAID, RACI, stage gates and cutover rehearsals remain the foundation; agentic AI is an extension of good program management, not a replacement for it.

This guide is a planning reference, not a complete methodology. Validate roles, cadences and tooling against your organization's specific governance standards, and align AI governance design with applicable data privacy and AI regulation in each operating jurisdiction before go-live.