KAIZEN AI

KAIZEN AI · READINESS PLAYBOOK

The Kaizen AI Readiness Playbook

Eight Diagnostics Before Scaling Agentic AI in Oracle-Heavy Environments — Scored, Interpreted, and Actioned.

Method. Measure. Momentum.

Governed agentic AI for Oracle-powered enterprises.

kaizen-orbit.lovable.app · info@kaizenai.ai

EXECUTIVE SUMMARY

Why Readiness Is the Difference Between a Pilot and a Program

The most common failure mode in enterprise agentic AI is not a model failure. It is not a technology failure. It is a readiness failure. Pilots succeed — the agent performs the target workflow, the cycle-time signal is real — and then the program stalls because the operating model was never built. The approval authority is ambiguous. The rollback plan was never tested. The Oracle entitlement model was never formalized. The audit team sees the agent for the first time in a field audit.

This playbook is the Kaizen AI diagnostic framework for enterprise readiness. Score each of the eight dimensions before you build anything. Below 10: not ready to scale. Below 6: a pilot will fail. Use the scoring to invest in the right gaps before committing to production deployment.

8	0–16	<10	<6
Diagnostic dimensions	Total possible score	Not ready to scale	Pilot will fail

The Business Problem: Why Readiness Is Skipped

The readiness gap is not ignorance. It is optimism. Executive sponsors see an AI pilot succeed in a controlled environment and authorize production deployment before the operating model exists to support it. The gap between 'the pilot worked' and 'the program scales' is almost always an operating model gap — not a technology gap.

In Oracle environments, the operating model gap has a specific anatomy:

- **Process gaps:** the workflow map covers the happy path but not exceptions. The agent hits an unhandled case and fails silently — or worse, takes an unintended action.
- **Entitlement gaps:** the agent runs under a shared admin account, creating SOD violations and ambiguous audit trails that immediately become SOX findings.
- **Policy gaps:** approval authority exists as a Slack agreement or tribal knowledge, not as policy-as-code. The agent cannot reliably enforce it.
- **Ownership gaps:** no one owns the agent after go-live. Monitoring lapses. Policy drifts. The agent begins exhibiting behavior that no one authorized.
- **Audit gaps:** the internal audit team encounters the agent for the first time during a field audit and treats it as an untested IT control — an automatic finding.

Why Oracle-Heavy Enterprises Are Uniquely at Risk

Oracle-heavy manufacturers face a specific readiness challenge that cloud-native companies do not. Oracle EBS, JDE, and PeopleSoft environments carry years of accumulated customization, SOX control documentation, and entitlement architecture that was never designed to accommodate agentic actors. An agent operating in these environments without a formal entitlement model and policy-as-code framework is an uncontrolled change to a tested IT control — which is a finding regardless of how well the agent performs.

Maturity Model: Five Readiness Levels

Level	Name	Score Range	Characteristics	Recommended Action
1	Fragmented	0–5	No workflow maps; shared Oracle accounts; no policy documentation; audit posture undefined	Do not pilot. Invest 8–12 weeks in operating model foundations before any agent development.
2	Piloting	6–9	Target workflow partially mapped; some policy informal; no rollback testing; audit not engaged	Invest in specific gaps (policy-as-code, rollback architecture, audit engagement) before go-live. Kaizen Assessment recommended.
3	Governed	10–12	Workflow mapped; Oracle entitlement defined; policy-as-code drafted; HITL designed; audit pre-engaged	Proceed to first agent development. Expect 10–12 week deployment to production.
4	Scaling	13–14	First agent in production; governance scaffold reusable; operating model active; KPIs instrumented	Begin second agent design. Expect 4–6 week scaffold reuse deployment.
5	Compounding	15–16	Multiple agents in production; scaffold reuse >70%; audit evidence automated; spreadsheet dependency declining	Expand use-case portfolio. Focus on finance-supply chain connected workflows.

The Eight Diagnostics

Diagnostic 1: Method Clarity

Category: PROCESS

Standard: The target workflow is mapped end-to-end with named owners and documented exception paths.

Scoring guidance: A workflow map that covers the happy path and stops is a 1. A complete specification — trigger, data sources, decision logic, exception paths, completion definition in Oracle — is a 2.

The test question: What does the agent do when required Oracle data is missing or stale? Who owns the exception? What is the SLA? If you cannot answer these without looking something up, score 1.

How to reach a 2: Conduct a structured walkthrough with the two or three people who currently do this workflow manually. Map every branch against Oracle data schema. Validate that every data field the agent needs is populated in the Oracle environment.

> **Why This Matters**
>
> Ambiguous exception paths cause silent failures. In a SOX environment, a silent failure is a control gap.

Diagnostic 2: Measure Discipline

Category: METRICS

Standard: KPI thresholds, audit metrics, and adoption metrics are defined, instrumented, and baselined before go-live.

Scoring guidance: Vague success criteria ('faster cycle time') are a 1. Named, quantified, pre-baselined metrics ('AP exception resolution from 4.2 days to under 8 hours; 80% agent-handled volume within 90 days') are a 2.

The test question: What is the specific cycle-time target? What percentage of exceptions will the agent handle at full deployment? How will you measure spreadsheet reduction? If you are still working this out, score 1.

How to reach a 2: Define three KPI tiers: operational (cycle time, resolution rate), governance (approval latency, policy adherence), adoption (agent vs. human volume). Capture the baseline before go-live.

> **Why This Matters**
>
> Without a pre-defined measurement framework, executive sponsors cannot evaluate program health — and programs die in the measurement gap.

Diagnostic 3: Oracle Entitlement Model

Category: SECURITY

Standard: The agent operates under a named service account with scoped Oracle responsibilities — never a shared admin.

Scoring guidance: An informal plan to 'figure out the permissions during development' is a 0. A dedicated service account with documented Oracle responsibilities, provisioned in non-production, is a 2.

The test question: What Oracle responsibilities will the service account hold? Is it provisioned yet in non-production? Does its activity appear distinctly in Oracle audit trails? Does it create any SOD violations?

How to reach a 2: Engage your Oracle DBA and security team before any agent development. Define and document the scoped responsibility set. Test in Oracle sandbox. Validate the audit trail shows the service account distinctly.

> **Why This Matters**
>
> A shared admin account makes the audit trail ambiguous and creates SOD violations. Both are immediate SOX findings that will require remediation regardless of agent performance.

Diagnostic 4: Policy-as-Code

Category: GOVERNANCE

Standard: Approval thresholds, SoD rules, and rollback windows are written as version-controlled policy — not Slack agreements.

Scoring guidance: Policy existing as email threads or tribal knowledge is a 0. Approval authority matrix and SoD constraints in a version-controlled, machine-readable format is a 2.

The test question: Can you show the auditor the policy document today? Is it machine-readable? Is there a change control process for policy updates? If not, score 1.

How to reach a 2: Work with internal audit to formalize the approval authority matrix and SoD constraints for the target workflow. Express in JSON schema or YAML. Store in version control. Require change control for updates.

> **Why This Matters**
>
> Policy drift is the primary source of agentic AI compliance failures. Policy-as-code makes the agent's behavior directly testable against the policy.

Diagnostic 5: Human-in-the-Loop Design

Category: OPERATIONS

Standard: Approvers, escalation paths, and SLAs are explicitly defined — not assumed to be 'whoever is available.'

Scoring guidance: Knowing 'someone in finance will approve' is a 0. Named approver roles (not individuals — roles), SLAs, escalation paths, and backup approver chains are a 2.

The test question: Who are the named approver roles? What is the SLA per approval tier? What happens when the SLA is breached? What happens when the approver is unavailable?

How to reach a 2: Design and test the approval workflow before the agent is developed. Integrate with Fusion BPM Worklist, Slack, or email. Define escalation paths. Run approval workflow tests with actual approver roles.

> **Why This Matters**
>
> Ill-defined HITL creates either approval latency that negates cycle-time benefit, or approval rubber-stamping that creates the illusion of oversight without the reality.

Diagnostic 6: Reversibility

Category: RESILIENCE

Standard: Every action the agent can take has a documented, tested undo path with a stated rollback window.

Scoring guidance: Knowing 'we can probably reverse it in Oracle' is a 1. Documented rollback transaction, rollback window, and named Oracle responsibility for execution — tested in non-prod — is a 2.

The test question: What is the standard Oracle reversal transaction for each agent write-back? What is the rollback window before irreversibility? Who has the Oracle responsibility to execute it?

How to reach a 2: Document rollback procedures for every write-back action before go-live. Test rollback in non-production. Include rollback window in the policy-as-code.

> **Why This Matters**
>
> An agent that cannot be reversed cannot be trusted. Inability to undo is the single largest barrier to executive authorization of production deployment.

Diagnostic 7: Audit Posture

Category: COMPLIANCE

Standard: The SOX/ITIL/regulatory evidence model is agreed with internal audit before the first agent action ships.

Scoring guidance: Planning to 'explain it to audit after we go live' is a 0. Written pre-engagement with internal audit, control classification agreed, evidence model documented is a 2.

The test question: Has internal audit reviewed the agent's governance architecture? Have they agreed on the evidence model? Do they know how the agent's actions appear in the SOX control framework?

How to reach a 2: Schedule a pre-deployment review with internal audit — and external auditors if possible. Walk through governance architecture, audit record model, and evidence package generation. Obtain written agreement on control classification.

> **Why This Matters**
>
> A SOX auditor encountering an agent for the first time in a field audit treats it as an untested IT control — an automatic finding, regardless of agent performance.

Diagnostic 8: Operating Model

Category: ORGANIZATION

Standard: The tripartite ownership model — business sponsor, Oracle application owner, AI ops lead — is named and committed.

Scoring guidance: Theoretical ownership ('the project team will maintain it') is a 0. Three named individuals with documented accountabilities, meeting cadence, and escalation paths is a 2.

The test question: Who is the named Business Sponsor? The Oracle Application Owner? The AI Ops Lead? Are their accountabilities documented? Is the operating charter signed?

How to reach a 2: Complete the operating charter before go-live: named owners, documented accountabilities, monthly meeting cadence for the first quarter, escalation path for owner conflicts.

> **Why This Matters**
>
> Without named individual owners, accountability diffuses and agents die in production support — not because they fail technically, but because no one owns the improvement cycle.

Score Interpretation

Score	Level	Readiness Status	Recommended Next Step
14–16	Compounding	Production-ready and scaling	Second+ agent design; connected finance-SCM workflows
10–13	Governed	Ready for first agent	Address remaining gaps; proceed to agent development
6–9	Piloting	Specific gaps to close first	4–8 week operating model investment; Kaizen Assessment recommended
0–5	Fragmented	Not ready	12-week foundation build before any pilot; start with process mapping and Oracle entitlement

Executive Workshop Format

The Kaizen AI Readiness Workshop is a half-day structured session with four required participants: the business process owner, the Oracle application owner, the CISO or IT security lead, and the internal audit or compliance lead. The agenda:

- Hour 1: Target workflow walkthrough — map the exception paths the process owner cannot currently articulate.
- Hour 2: Oracle entitlement and governance architecture — define the service account model and policy-as-code scope.
- Hour 3: Audit and compliance pre-engagement — agree the evidence model and control classification.
- Hour 4: Operating model and ownership — name the three owners and draft the operating charter.

CIO / CFO / COO: Your Specific Readiness Questions

CFO Questions

- Which SOX controls will the agent touch, and is internal audit pre-engaged?
- What is the approval authority matrix, and is it written as policy?
- What is the rollback procedure if an agent action is wrong?
- How will cycle-time improvement be measured against the current baseline?
- What is the agent's evidence model for audit purposes?

CIO / COO Questions

- Does the agent run under a scoped Oracle service account — not a shared admin?
- Is the human-in-the-loop workflow designed with SLAs and escalation paths?
- Has the operating model (business / Oracle / AI ops) been named and committed?
- What is the KPI baseline, and how will adoption be tracked?
- Is the second agent design already in scope to capture scaffold reuse value?

Implementation Roadmap: From Readiness to First Agent

Phase	Duration	If Score 0–5	If Score 6–9	If Score 10+
Foundation	Weeks 1–4	Process mapping; Oracle entitlement design; policy-as-code authoring; audit engagement	Close specific gaps identified in scoring; validate with Kaizen Assessment	Begin agent design brief immediately

Build	Weeks 5–10	Not started yet	First agent development in Oracle sandbox; approval workflow integration	First agent in development
Deploy	Weeks 11–14	Not started yet	Staged production deployment; KPI capture; audit evidence generation	First agent live; scaffold established
Scale	Weeks 15+	Foundation still in progress	Second agent design; scaffold reuse; operating model maturing	Second and third agents deploying

The Kaizen AI 3M+ Assessment

The full Kaizen AI Readiness Assessment is a four-week structured engagement with deliverables: scored readiness map across all eight dimensions; operating model recommendation; governance architecture blueprint; use-case sequencing plan by scaffold reuse; first-agent design brief ready for development.

The Assessment eliminates the readiness gap before it becomes a production failure.

Contact: info@kaizenai.ai · kaizen-orbit.lovable.app/assessment

Method. Measure. Momentum.

