

AI Governance Checklist for Retail Back-Office & Real Estate Systems

A twelve-standard guide for IT departments evaluating any AI tool that reads from or writes to lease administration, store development & construction, facilities, accounts payable or portfolio planning systems. Companion to "Governing AI in the Retail Back Office" (Valeri, October 2026).

How to use it. Score each line Y / N / Partial. A read-only tool should clear standards 1–3, 7, 8, 10 and 11 before approval; any tool with write access to a system of record must clear all twelve. Standard 12 is the retailer's own obligation and applies regardless of vendor. This checklist is additive to, not a substitute for, your standard vendor security review (SOC 2, encryption, vulnerability management, penetration testing). Autonomous payments, legal commitments and personnel decisions are out of scope for any configuration.

#	Standard	Check (Y / N / Partial)	Evidence to request
1	Data boundary & no-training	☐ Customer data is never used to train or fine-tune any model (vendor or provider) ☐ Model providers retain prompts/outputs only for the request lifecycle ☐ Data residency is stated and acceptable	<i>Named model providers; no-training / zero-retention contract terms; DPA</i>
2	Identity binding & least privilege	☐ Every inference and tool call runs under the requesting user's permissions ☐ No autonomous "super-identity" for the agent ☐ SSO (SAML/OIDC) and MFA enforced; revocation takes effect next interaction	<i>Identity-propagation architecture description; live revocation demo</i>
3	Tenant & data isolation	☐ Data, credentials, indexes, prompt context and logs isolated per customer ☐ No prompt caching shared across tenants ☐ Single-tenant option available for sensitive workloads	<i>Written isolation model by layer; single-tenant terms</i>
4	Deny-by-default tools & action tiers	☐ No tool or write capability active until the retailer enables it ☐ Each tool declares typed inputs, outputs and scope; out-of-scope calls rejected ☐ Workflows tiered: read-only / suggested / standing-authorization / restricted ☐ Autonomous payments, legal commitments and personnel decisions unsupported	<i>Tool catalog with scopes; tier map; demo of a rejected call</i>
5	Write-back validation	☐ AI changes to a system of record are validated against its current state first ☐ Non-conforming values (bad ID, out-of-term date, out-of-tolerance amount) are rejected and surfaced, not retried ☐ Retry and cost ceilings are set	<i>Write-back path description; validation rules; sample rejection log</i>
6	Human review & completeness disclosure	☐ Outputs that change a financial record, payment, counterparty notice or reported balance require logged human approval ☐ Abstracts/analyses disclose what was found, what was not, and from which source	<i>List of approval-gated workflows; UI screenshots; sample audit entry</i>
7	Untrusted content & prompt-injection controls	☐ External content (landlord/vendor email, uploads, portals) is labeled by source and treated as data, never instructions ☐ System prompts and tool definitions live in code/config, not editable content ☐ Model-generated code or commands are never executed against customer systems	<i>Instruction-content separation design</i>

#	Standard	Check (Y / N / Partial)	Evidence to request
8	Tamper-evident audit trail	☐ Every inference, tool call and data access logged at capture: tenant, user, workflow, tools, records, provider, model name & version, response ☐ Logs append-only; retention period stated; exportable to SIEM / on request ☐ Vendor personnel access appears in the same trail	<i>Sample audit record; retention & export terms; turnaround SLA</i>
9	Customer-controlled kill switch	☐ Retailer admins can pause the environment, a workflow, a tool or a user without vendor involvement ☐ Effect on next interaction	<i>Live demo; admin documentation</i>
10	Third parties, credentials & incident notification	☐ All subprocessors and model providers in the inference path disclosed, with region ☐ Integration credentials encrypted, workflow-scoped, never exposed to the model, never shared across customers ☐ Confirmed-incident notification within a stated period	<i>Subprocessor list + change notice; credential handling; SLA in contract</i>
11	AI-specific assurance	☐ Controls mapped to NIST AI RMF / AI 600-1, ISO/IEC 42001, OWASP LLM Top 10 & Agentic taxonomy, CSA AICM ☐ AI-specific incident playbook: prompt injection, provider-side leakage, erroneous AI-initiated action ☐ Stated cadence of adversarial testing of agent behavior	<i>Framework mapping; incident playbook summary; testing cadence</i>
12	Internal governance (retailer-side)	☐ Register of every AI use case touching in-scope systems: tier, data classes, owner, approver, review date ☐ Sanctioned AI tool provided + published policy on what may go into any other tool ☐ Role-appropriate AI literacy delivered to operators and reliant staff ☐ Each use case mapped to disclosure, privacy and financial-reporting obligations; reviewed annually	<i>The register; the policy; training records; regulatory mapping</i>

Decision rule

- Any N on standards 1, 2, 4, 5 or 9 for a tool with write access: do not approve until remediated.
- Partial on 3, 7, 8, 10 or 11: approve for read-only use with a dated remediation plan.
- Record the outcome in the AI use-case register (standard 12) with a review date no more than 12 months out.

Reference frameworks

NIST AI Risk Management Framework (AI 100-1, Jan 2023) and Generative AI Profile (AI 600-1, Jul 2024) · ISO/IEC 42001:2023 · OWASP Top 10 for LLM Applications 2025 and Agentic AI Threats & Mitigations · Cloud Security Alliance AI Controls Matrix (Jul 2025) · NRF Principles for the Use of AI in the Retail Sector (Nov 2023) · NSA/CISA Deploying AI Systems Securely (Apr 2024) and AI Data Security (May 2025)