Konstera

CHECKLIST

AI Agent *Governance Checklist.*

30 commitments the CIO, CFO and risk owner at a mid-market retail, wholesale or distribution business should make before deploying AI agents — grouped into pre-deployment, operational, incident response and compliance posture. Hard gates, not aspirations.

30 items Printable · ungated Konstera · 2026

CHECKLIST

Contents

A Pre-deployment governance
B Operational governance
C Incident response
D Compliance posture
– How to use this

A Pre-deployment governance

Settled before any agent sees production data. Treat these as hard gates — not aspirations.

- [ ] A01
Business owner for the agent is named (not "the AI team") and signed off on the use case.

- [ ] A02
Autonomy policy is defined per agent: what runs autonomously, what requires reviewer sign-off, what requires named-approver approval.

- [ ] A03
Success criteria are written and signed by both sides before any build work.

- [ ] A04
Threshold rules (amounts, counterparties, time-of-day) that tighten or loosen autonomy are documented.

- [ ] A05
Data classification is complete: the agent cannot touch data it isn't cleared for.

- [ ] A06
LLM provider and model are approved by security — specifically for the data it will see.

- [ ] A07
BYOK keys are stored in the credentials MCP service, not in agent code or environment variables.

- [ ] A08
Agent runs in an isolated per-invocation container; no shared state between tenants.

- [ ] A09
DPIA (Data Protection Impact Assessment) is complete where personal data is processed.

- [ ] A10

Vendor's compliance roadmap is in writing: SOC 2 Type II date, GDPR posture, HIPAA path.

B Operational governance

Lives in your platform configuration and your weekly ops rhythm, not a Confluence page.

- [ ] B01
Reviewer inbox is staffed with named reviewers, each with a role-scoped RBAC profile.

- [ ] B02
Reviewer actions (send / edit / redraft / reject) are logged per message for agent tuning.

- [ ] B03
Monthly cost ceilings are enforced in the platform — the agent refuses invocations at the cap.

- [ ] B04
Cost is attributed per agent, team, model, tenant — one ledger across the business.

- [ ] B05
Every invocation produces a durable audit trail: inputs, model, prompt, tool calls, reviewer actions.

- [ ] B06
Replay capability is live: any past invocation can be re-run against a new agent version.

- [ ] B07
Autonomy policy is reviewable + editable by platform admins without redeployment.

- [ ] B08
Fallback LLM chain is configured — the agent survives one provider's outage.

- [ ] B09
Observability (OpenTelemetry traces) exports to the business's existing APM stack.

- [ ] B10
Four KPIs are tracked weekly: reviewer approval rate, time-to-decision, cost per outcome, deflection or capacity shift.

C Incident response

What you do when the agent gets it wrong — and the hour matters.

- [ ] C01
Named on-call owner for every agent in production (name, phone, Slack handle).

- [ ] C02
Runbook per agent: how to disable it, how to pause it, how to redeploy with a corrected policy.

- [ ] C03
Pre-approved customer communication templates for "we got that wrong" scenarios.

- [ ] C04
Legal and PR engagement thresholds (which incident classes escalate immediately).

- [ ] C05
Post-incident replay produces a root-cause document with model, prompt, context, and reviewer decisions.

D Compliance posture

Regulators will ask. The answers should already exist.

- [ ] D01
Data residency is pinned per tenant (EU / UK / US / self-hosted) and documented.

- [ ] D02
  DPA (Data Processing Agreement) is signed with the AI platform vendor.

- [ ] D03
  Subprocessor list is current and surfaced to procurement.

- [ ] D04
  EU AI Act obligations are mapped (high-risk use case? General-purpose AI reliance? Prohibited uses?).

- [ ] D05
  GDPR subject access and erasure requests reach platform data — not just your production stores.

Three ways to run the checklist.

1. **Before signing with any AI platform.** Walk through sections A and D with the vendor. If they can't answer three or more items, that's a procurement red flag — not an onboarding question.
2. **Before promoting any pilot to production.** Section A is the production gate; nothing deploys with open items. Sections B and C come alive once real volume starts flowing.
3. **Quarterly, as a risk-committee audit.** Walk through every item against every production agent. Report green / amber / red to the risk committee. Open items become next-quarter owned actions.

Konstera's honest self-assessment.

Konstera meets every item in sections A, B and C out of the box. On section D, we're ready now on data residency, DPA, subprocessor list and GDPR; SOC 2 Type II is targeted for 2027. See Trust & security for the full picture.

Konstera

One platform for every AI agent in your business.

Governed centrally, costed in one ledger, built on shared context — with people in control of every decision that matters.

konstera.ai · konstera.ai/book-a-call · konstera.ai/resources