AI Vendor Risk Checklist

Assess AI vendors before and after approval.

AI vendor risk does not stop at onboarding. Vendors can introduce new AI features, change model behavior, update data policies, expand integrations, or alter default settings after the initial review.

Use this checklist to evaluate AI-enabled vendors before approval and monitor them over time.

CONTENTS

1. Vendor AI Profile

What AI does this vendor provide?

Use this section to understand whether the vendor uses AI, where it appears in the product, and how it may affect your organization.

CHECKLIST

- [ ] Vendor has disclosed whether the product includes AI, machine learning, GenAI, copilots, agents, automation, or model-driven functionality
- [ ] AI features are clearly identified by name, function, and product area
- [ ] The vendor has explained whether AI is core to the product or an optional feature
- [ ] Admins can enable, disable, or restrict AI features
- [ ] The vendor has disclosed whether AI features are enabled by default
- [ ] The vendor has provided documentation on AI feature behavior and limitations
- [ ] The vendor has disclosed whether AI functionality depends on third-party models or AI providers

KEY QUESTIONS TO ASK

- Which AI features are available today?
- Which AI features are planned or in beta?
- Are any AI features enabled automatically?
- Can our organization control which users can access AI features?
- Does the vendor use proprietary, third-party, open-source, or customer-configured models?

EVIDENCE TO COLLECT

AI product documentation | Feature release notes | Admin control documentation | AI transparency documentation | Vendor responses to AI questionnaire

2. Use Case and Business Context

How will the vendor's AI be used?

Use this section to connect vendor AI capabilities to your actual business use case, not just the vendor's general product description.

CHECKLIST

- [ ] Intended business use case is documented
- [ ] Business owner is assigned
- [ ] Technical or system owner is assigned
- [ ] AI use is classified as internal, customer-facing, employee-facing, or operational
- [ ] Potential customer, borrower, employee, or end-user impact is documented
- [ ] Use case has been reviewed for regulated workflow exposure
- [ ] Approved and prohibited uses are documented
- [ ] Any usage limits or approval conditions are captured

KEY QUESTIONS TO ASK

- What business process will this vendor support?
- Will AI affect customers, employees, applicants, borrowers, accounts, funds, or regulated decisions?
- Will users rely on AI-generated outputs to make decisions?
- Will AI outputs be reviewed by a human before use?
- Could the AI feature expand into additional teams or workflows over time?

EVIDENCE TO COLLECT

Use case description | Business owner approval | Workflow documentation | Review notes | Approved-use policy or conditions

3. Data Privacy and Data Use

What data will the vendor's AI access, process, retain, or use?

Use this section to assess whether sensitive, regulated, customer, employee, financial, or confidential data may flow through the vendor's AI capabilities.

CHECKLIST

- [ ] Data categories are documented
- [ ] Sensitive data exposure has been assessed
- [ ] Customer or borrower data exposure has been assessed
- [ ] Employee data exposure has been assessed
- [ ] Confidential business data exposure has been assessed
- [ ] Vendor has disclosed whether customer data is used for training or model improvement
- [ ] Vendor has disclosed retention practices for prompts, outputs, files, logs, and metadata
- [ ] Vendor has disclosed whether data is shared with third-party AI providers
- [ ] Opt-out, deletion, and data control options are documented
- [ ] Privacy review is complete or routed for review

KEY QUESTIONS TO ASK

- What data will the AI feature process?
- Are prompts, outputs, attachments, files, or logs retained?
- Is customer data used to train, fine-tune, evaluate, or improve models?
- Are third-party AI providers involved?
- Can we restrict the data users submit into the AI feature?
- Can we delete AI-related data if needed?

EVIDENCE TO COLLECT

Data processing agreement | Privacy policy | AI-specific data use terms | Retention documentation | Subprocessor list | Data flow diagram, if available

4. Security and Access Controls

How is access to the vendor's AI controlled?

Use this section to assess whether the vendor provides sufficient security, access, and administrative controls for AI-enabled functionality.

CHECKLIST

- [ ] Role-based access controls are available
- [ ] Admin controls for AI features are documented
- [ ] SSO and identity controls are supported, where applicable
- [ ] Logging or audit history is available for AI-related activity
- [ ] Data access permissions are understood
- [ ] AI feature access can be limited by role, team, or user group
- [ ] Security documentation has been reviewed
- [ ] Incident response process includes AI-related events, where applicable
- [ ] Vendor has disclosed relevant certifications or security reports

KEY QUESTIONS TO ASK

- Who can access the AI feature?
- Can admins restrict usage by user, team, role, or workspace?
- Are AI prompts, outputs, or actions logged?
- Can we audit who used the feature and when?
- What security controls protect AI-related data?
- How does the vendor handle AI-related incidents or misuse?

EVIDENCE TO COLLECT

SOC 2 report or equivalent | Security whitepaper | Access control documentation | Audit log documentation | Incident response documentation | SSO and admin settings documentation

5. Model and AI System Risk

What model or AI system sits behind the vendor's feature?

Use this section to understand the model dependency, transparency, limitations, and oversight needs behind the AI capability.

CHECKLIST

- [ ] Model type or AI provider is documented
- [ ] Vendor has disclosed whether the model is proprietary, third-party, open-source, or customer-configured
- [ ] Model purpose and limitations are documented
- [ ] Known risks, failure modes, or accuracy limitations are disclosed
- [ ] Explainability or transparency information is available, where relevant
- [ ] Human review requirements are defined
- [ ] Model change notification process is understood
- [ ] Ongoing monitoring requirements are documented

KEY QUESTIONS TO ASK

- What model or AI provider powers the feature?
- Can the vendor change models without notice?
- How does the vendor evaluate model performance?
- What are the known limitations or failure modes?
- Does the AI produce recommendations, classifications, summaries, decisions, or actions?
- Is human review required before relying on outputs?

EVIDENCE TO COLLECT

Model documentation | AI transparency documentation | Evaluation or testing summary | Known limitations disclosure | Model change policy | Human oversight requirements

6. Agentic AI and Automation

Can the vendor's AI take action or trigger workflows?

Use this section for copilots, assistants, AI agents, automated workflows, and features that can act across systems.

CHECKLIST

- [ ] Agentic or automated capabilities are identified
- [ ] Action scope is documented
- [ ] System permissions are documented
- [ ] Human-in-the-loop requirements are defined
- [ ] Escalation paths are documented
- [ ] Stop, override, or rollback controls are available
- [ ] External communications or customer-impacting actions are reviewed
- [ ] Agent or automation changes are monitored over time

KEY QUESTIONS TO ASK

- Can the AI only suggest, or can it take action?
- Can it create, modify, approve, delete, send, escalate, or trigger workflows?
- What systems can it access?
- What data can it retrieve or use?
- Where is human approval required?
- Can the organization pause, restrict, or revoke agent access?

EVIDENCE TO COLLECT

Agent or automation documentation | Permission model | Workflow diagrams | Human oversight policy | Admin control documentation | Audit log examples

7. Compliance and Governance Fit

Does the vendor support your AI governance requirements?

Use this section to assess whether the vendor can support your internal policies, regulatory expectations, and oversight model.

CHECKLIST

- [ ] Vendor AI use aligns with internal AI policy
- [ ] Risk tier has been assigned
- [ ] Required reviewers are identified
- [ ] Legal, compliance, privacy, security, vendor risk, or model risk review has been routed where needed
- [ ] Approval conditions are documented
- [ ] Exceptions are documented
- [ ] Reassessment cadence is defined
- [ ] Vendor is added to AI inventory or AI register
- [ ] Evidence is stored for future audit, exam, customer, or board review

KEY QUESTIONS TO ASK

- Does this vendor require AI governance committee review?
- Does the use case involve regulated decisions, sensitive data, or customer impact?
- Which internal policies apply?
- Which teams need to approve this vendor before use?
- What controls or conditions are required?
- When should the vendor be reassessed?

EVIDENCE TO COLLECT

AI risk assessment | Internal review record | Approval decision | Conditions or mitigations | Exception documentation | AI inventory record | Reassessment schedule

8. Vendor Change Monitoring

How will you know when the vendor's AI risk changes?

Use this section to define ongoing monitoring after approval.

CHECKLIST

- [ ] Vendor has a process for notifying customers of AI feature changes
- [ ] Release notes are monitored
- [ ] AI feature changes are reviewed after approval
- [ ] Model changes are monitored
- [ ] Data policy changes are monitored
- [ ] Subprocessor changes are monitored
- [ ] Default setting changes are monitored
- [ ] Material changes trigger reassessment
- [ ] Monitoring activity is captured as evidence

KEY QUESTIONS TO ASK

- How will we know when the vendor adds or changes AI features?
- Will we be notified before AI capabilities are enabled?
- Can the vendor change models, subprocessors, data terms, or default settings without approval?
- What changes require reassessment?
- Who owns ongoing vendor monitoring?

EVIDENCE TO COLLECT

Change notification policy | Release notes | Vendor monitoring record | Alert history | Reassessment record | Follow-up actions

9. Approval Decision

What decision should be made?

Use this section to document the final vendor AI review outcome.

DECISION OPTIONS

- [ ] Approved
- [ ] Approved with conditions
- [ ] Requires additional review
- [ ] Restricted use only
- [ ] Not approved
- [ ] Reassessment required before broader rollout

APPROVAL CONDITIONS

Document any limits, controls, or follow-up requirements:

Approved users or teams | **Approved use cases** | **Prohibited use cases** | **Data restrictions** | **Human review requirements** | **Admin control requirements** | **Monitoring requirements** | **Reassessment date** | **Required vendor follow-up**

DECISION RECORD

Vendor	AI feature or product
Business owner	Reviewer or approving team
Risk tier	Decision
Reassessment date	Evidence location
Conditions	

10. Quick Scoring Guide

Use this lightweight guide to determine whether a vendor requires deeper review.

Lower risk

- Internal productivity use only
- No sensitive, customer, borrower, employee, or regulated data
- Human review required before use
- AI feature can be disabled or restricted
- No decisioning, eligibility, prioritization, or automated action
- Vendor provides clear documentation and controls

Higher risk

- Customer, borrower, employee, financial, health, or regulated data involved
- AI affects eligibility, access, pricing, prioritization, recommendations, or decisions
- Vendor uses customer data for training or model improvement
- AI feature is enabled by default
- Vendor relies on third-party AI providers with limited transparency
- AI can trigger actions, modify records, communicate externally, or operate autonomously
- Model behavior, data policy, or default settings can change without clear notice

Final Review

Before approving an AI vendor, confirm:

- [ ] The vendor's AI capabilities are documented
- [ ] The business use case is approved
- [ ] Data use and privacy terms are reviewed
- [ ] Security and access controls are understood
- [ ] Model or AI system risk is assessed
- [ ] Agentic or automated behavior is reviewed, if applicable
- [ ] Required stakeholders have reviewed the vendor
- [ ] Approval conditions are documented
- [ ] Monitoring requirements are assigned
- [ ] Evidence is stored in the AI governance record

LUCIDTRUST

This Checklist Only Gets You So Far

LucidTrust automates AI vendor discovery, continuous risk monitoring, and audit-ready evidence collection — so reviews like this one run themselves.

Book a Demo