SOCIAL ENTERPRISE ADVISORS

Practical guide | September 6, 2026

Before you approve an AI vendor

A working set of questions for evaluating a proposed tool, its place in your workflow, and the evidence needed for a decision.

> A general operational guide. It does not establish clinical safety, legal compliance, or suitability for a particular organization. Involve the appropriate clinical, privacy, security, and legal reviewers for your use case.

1. Define the decision the tool will affect

Write down the actual workflow before reviewing features. Describe the input, proposed output, person using it, and action that might follow. Identify where a mistake would create harm, delay, or rework.

- What specific job should the tool help someone complete?
- Who is accountable for the resulting action?
- What remains a human decision, and when must the workflow stop for review?
- What is explicitly outside the intended use?

2. Establish the data boundary

Ask for a data-flow description that follows information through collection, processing, storage, and deletion. Have the relevant specialists assess the answers against your requirements.

- What information enters the tool, and what can be minimized?
- Which vendors and subprocessors can access it, and where is it processed?
- Can your information be used for training, product improvement, or other secondary purposes?
- What happens to records and access when the agreement ends?

3. Request evidence for your use case

A demonstration shows that a tool can work in selected conditions. An evaluation should examine the tasks, users, and failure modes relevant to your organization. Use approved synthetic examples or appropriately authorized data.

- Which claims are supported by evidence, and what are the limitations?
- Do the tests represent the languages, populations, and workflows in scope?
- How does the system behave when information is missing or conflicting?
- What independent specialist review is needed before relying on the output?

4. Account for review and exceptions

Include the work created by checking, correcting, and escalating outputs. Compare the whole process with the baseline, not just the time spent generating a response.

- Who reviews outputs, and what capacity will that require?
- What happens when the reviewer disagrees or the tool is unavailable?
- Can users identify and report failures?
- Who can pause the tool, and what is the fallback process?

5. Plan for changes after purchase

Record the version and conditions you evaluated. Decide what changes should trigger another review, and identify the person responsible for ongoing monitoring.

- How will you learn about changes to models, data use, subprocessors, and functionality?
- What records can you export to investigate an incident or performance concern?
- Who maintains the inventory entry, review schedule, and unresolved conditions?
- What would justify stopping use or changing vendors?

6. Write a decision memo

Conclude with an explicit recommendation: proceed, proceed with conditions, or do not proceed. Separate known facts from vendor claims and unresolved questions.

- Intended use, exclusions, and accountable owner.
- Evidence reviewed, tests performed, and limitations.
- Required human controls and unresolved conditions, each with an owner.
- Approval, monitoring, and re-evaluation dates.
- Scope and cost of the next step, including review and support.

Further reading

This SEA operational guide draws on the voluntary NIST AI Risk Management Framework. It is not a certification or reproduction of the framework.

- NIST AI Risk Management Framework
- NIST AI RMF 1.0: Govern, Map, Measure, Manage

Apply this to your organization

Explore Clinical AI Vendor Review or contact Social Enterprise Advisors to discuss your priorities.