

Enterprise AI Readiness Assessment

Twelve yes / partial / no questions. The score is how many of S·CORP's readiness practices are in place. It is not an industry benchmark.

UPDATED 2026-10-22

Readiness here means a sponsor, an operator, and a security owner can see the same boundary. The questions are the practices we ask for before an agent meets a live book of work. Score them yourself. A "yes" is two points, a "partial" is one, a "no" is zero.

The bands below are our reading of that score. They are not a statistic about other firms, and they are not a prediction. If the GDPR, PCI DSS, ISO 27001, or the EU AI Act already apply to the workflow, those duties sit on top of this list. This page does not discharge them.

Score the workflow you actually intend to run, not the platform in general.

1. A sponsor can name the user and the unit of work in one sentence.

If the sentence needs a metaphor, score this as not in place.

Yes / Partial / No

2. The decisions a person still makes are listed, including anything that spends money or makes a promise.

A blank list means the integration will choose.

Yes / Partial / No

3. There is a data map: systems, fields, and fields that must not enter a prompt or an embedding.

A share drive is not a data map.

Yes / Partial / No

4. Privacy has a written basis and a retention rule for prompts and traces, where a law such as the GDPR applies.

Naming the law is not the same as having the rule.

Yes / Partial / No

5. Tools are inventoried, each with a narrow verb and a non-shared identity.

One connector called "update" is a gap.

Yes / Partial / No

6. Authentication of the customer, broker, or employee does not depend on the model's opinion.

The agent should receive a result, not invent one.

Yes / Partial / No

7. A person who is not a developer can stop the agent the same day.

If the only stop is a redeploy, this is partial at best.

Yes / Partial / No

8. A case identifier ties the tool calls, the output, and the human decision together.

A chat transcript alone is not the record.

Yes / Partial / No

9. Someone has read where the vendor stores prompts and whether customer content changes a model.

A sales answer does not count. The contract does.

Yes / Partial / No

10. Changes to the agent itself go through a review and a way back.

A private script one person runs is not a path.

Yes / Partial / No

11. The people who will accept or reject the output have used it on their own work.

A steering-committee demo does not count.

Yes / Partial / No

12. There is a named owner for sampling output against the written scope.

The owner needs a calendar, not a slogan.

Yes / Partial / No

How to read the points

- 0–8: the boundary is not written. Do not put the agent in front of a customer or a live system of record.
- 9–16: a pilot can run beside the real process, with a person on every decision the scope sentence named.
- 17–24: the operation is narrow enough to run, if the sample review is actually on someone's calendar.

S-CORP — Engineered for outcomes. Built for trust.

s-corp.app · Enterprise AI Readiness Assessment