

# AI AGENT READINESS CHECKLIST

7 questions before you build

A practical scoring tool for deciding whether a workflow is ready for an AI agent, Copilot agent, or custom workflow assistant.

**Most AI agent projects do not fail because the model is weak. They fail because the business process underneath it is unclear.**

Use this checklist before you build. If the process is not ready, the agent will only expose the mess faster.

| 0 = Not clear                                      | 1 = Partially clear                                  | 2 = Ready                                |
|--|--|--|
| Undocumented, inconsistent, or dependent on memory | Known by the team, but not fully documented or owned | Clear, documented, measurable, and owned |

# 1. Who owns the process?

An AI agent needs an accountable human owner. If nobody owns the workflow, nobody owns the agent output.

| Check yourself  |                             |
|---|-----------------------------|
| ■ Who is responsible for this process today?<br>■ Who updates the process when it changes?<br>■ Who decides whether the agent is working correctly?<br>■ Who reviews failures, exceptions, and escalations? |                             |
| <b>Score</b>  | <b>Your score: ____ / 2</b> |
| 0: No clear owner<br>1: Informal owner, but not documented<br>2: Named owner with decision authority  | Notes: _____<br>_____       |

# 2. Where does the data live?

Agents are only useful when they can access the right information at the right time with the right permissions.

| Check yourself   |                             |
|--|-----------------------------|
| ■ What systems does the agent need to read from?<br>■ What systems does the agent need to write to?<br>■ Is the data structured, searchable, and current?<br>■ Are there privacy, compliance, or permission constraints?<br>■ Does the data live in email, spreadsheets, shared drives, CRMs, ERPs, or someone's head? |                             |
| <b>Score</b>   | <b>Your score: ____ / 2</b> |
| 0: Data location is unclear or scattered<br>1: Data sources are known, but messy or inconsistent<br>2: Data sources are clear, accessible, and permissioned  | Notes: _____<br>_____       |

# 3. What counts as success?

“Save time” is not a success metric. A useful agent needs a measurable business outcome.

| Check yourself   |                             |
|--|-----------------------------|
| ■ Response time<br>■ Error rate<br>■ Cost per transaction<br>■ Number of manual touches<br>■ Follow-up completion rate<br>■ Revenue recovered<br>■ Customer satisfaction<br>■ Employee hours saved |                             |
| <b>Score</b>   | <b>Your score: ____ / 2</b> |
| 0: Success is vague<br>1: Success is directionally clear but not measured<br>2: Success is specific, measurable, and tied to business value  | Notes: _____<br>_____       |

# 4. What exceptions happen often?

The happy path is easy. Exceptions determine whether the agent is safe to use.

| Check yourself |  |
|----------------|--|
|----------------|--|

■ What edge cases happen regularly? ■ Which exceptions stop the process? ■ Which exceptions require judgment? ■ Which exceptions can be resolved with rules? ■ Which exceptions should always go to a human?

|   |                             |
|---|-----------------------------|
| <b>Score</b>  | <b>Your score: ____ / 2</b> |
| 0: Exceptions are not documented<br>1: Common exceptions are known but not mapped<br>2: Exceptions are documented with clear handling rules | Notes: _____<br>_____       |

## 5. Who approves final decisions?

Not every workflow should be fully autonomous. Many of the best first agents prepare, recommend, draft, summarize, or route work while a human approves the final action.

|   |                             |
|---|-----------------------------|
| <b>Check yourself</b>   |                             |
| <p>■ What can the agent do without approval? ■ What can the agent draft but not send? ■ What can the agent recommend but not decide? ■ What dollar amounts, risks, or customer situations require approval? ■ Who gives the approval?</p> |                             |
| <b>Score</b>  | <b>Your score: ____ / 2</b> |
| 0: Approval authority is unclear<br>1: Approval is understood informally<br>2: Approval rules are explicit and risk-based   | Notes: _____<br>_____       |

## 6. When should a human step in?

Good agents do not pretend to know everything. They know when to stop.

|   |                             |
|---|-----------------------------|
| <b>Check yourself</b>   |                             |
| <p>■ What should trigger escalation? ■ Who receives the escalation? ■ How quickly should they respond? ■ What context should the agent include? ■ Where should the escalation happen — email, Teams, CRM, ticketing system, or dashboard?</p> |                             |
| <b>Score</b>  | <b>Your score: ____ / 2</b> |
| 0: Escalation rules are undefined<br>1: Escalation paths exist but are inconsistent<br>2: Escalation triggers, owners, and context are documented   | Notes: _____<br>_____       |

## 7. What should the agent never do?

Guardrails are not optional. The fastest way to create risk is to give an agent authority before defining its limits.

|  |  |
|--|--|
| <b>Check yourself</b>  |  |
| <p>■ Send customer-facing messages without review<br/>■ Change financial records directly<br/>■ Approve refunds, discounts, or credits<br/>■ Share sensitive data<br/>■ Delete records<br/>■ Override employee judgment<br/>■ Make legal, medical, HR, or compliance decisions<br/>■ Act outside documented policies</p> |  |

|   |                             |
|---|-----------------------------|
| <b>Score</b>  | <b>Your score: ____ / 2</b> |
| 0: No explicit guardrails<br>1: Some limits are understood, but not documented<br>2: Clear prohibited actions are documented and enforced | Notes: _____<br>_____       |

# Your Readiness Score

Add your scores from all 7 sections. Total score: \_\_\_\_ / 14

|       |                              |   |
|-------|------------------------------|---|
| 0–5   | <b>Do not build yet</b>      | Map the workflow, assign ownership, document exceptions, and define success.        |
| 6–10  | <b>Cleanup needed</b>        | Close documentation, data access, approval, and escalation gaps before building.    |
| 11–14 | <b>Ready with guardrails</b> | Start with a narrow pilot, human review, measurable outcomes, and clear escalation. |

## Best First AI Agent Candidates

Start with workflows that are repetitive, rules-based, high-volume, low to medium risk, easy to measure, and clear about inputs, outputs, owners, and escalation points.

| Good first candidates   | Bad first candidates  |
|---|---|
| <ul style="list-style-type: none"> <li>• Customer inquiry triage</li> <li>• Invoice follow-up drafts</li> <li>• Meeting summary and action-item routing</li> <li>• Internal policy Q&amp;A with citations</li> <li>• Lead intake qualification</li> <li>• Proposal draft assembly</li> <li>• Field service appointment preparation</li> <li>• Report generation from structured data</li> </ul> | <ul style="list-style-type: none"> <li>• Final hiring decisions</li> <li>• Legal or compliance determinations</li> <li>• Medical, financial, or HR advice</li> <li>• Customer refunds without review</li> <li>• Autonomous accounting changes</li> <li>• Sales commitments or contract changes</li> <li>• Anything nobody currently owns</li> </ul> |

## The Simple Rule

**If a human cannot explain the workflow clearly, an AI agent cannot run it safely.**

Map the process first. Automate the repeatable parts. Keep humans in the loop where judgment matters. Measure the outcome in business terms.

**AIA Copilot helps small and midsize businesses design, build, and manage practical AI solutions with measurable ROI. Visit [aiacopilot.com](https://aiacopilot.com) or message Scott Hay on LinkedIn with AGENT.**