

Who signs for the agent?

A one-page accountability map for your AI agents, the DDQ answers it gives you, and a board summary

Enterprise due diligence questionnaires now ask who is accountable for your AI agents, what they can do without a person, and how fast you can turn them off. The agent can't sign those answers. These templates name the people who do.

	Template	What it's for
1	The one-page accountability map	One row per agent: reach, identity, owner and backup, decision rights, off switch, evidence. An executive signs it.
2	Answers for the AI section of a DDQ	Eight common questions, each answered from one column of the map.
3	Quarterly board summary	What the board sees each quarter, and three questions any director can ask.

How to use this kit

- Fill in the bracketed fields. Gray text in brackets is guidance; delete it when you're done.
- The map has two example rows for Northwind AI, a made-up company. Delete them before you sign.
- Name a person, not a team, as each agent's owner, and add a backup. A team can run an agent; only a person can answer for it.
- Decide what each agent may do by consequence: reading and drafting, changes you can undo, and anything you can't undo, that spends money or that reaches a customer.
- Turn each agent off once and write down how long it took. Buyers ask for that number.
- Have one executive read every row and sign, then put the next review on the calendar.

These templates are general guidance, not legal advice. Liability wording belongs with your counsel. Free to use and adapt.

AI agent accountability map: [Company]

One row per AI agent you build or run, product features and internal tools alike. Reviewed [every quarter]. Signed by the executive who stands behind the DDQ answers drawn from it.

Agent and job	Reaches	Runs as	Owner and backup	Decision rights	Off switch: who, how fast	Pause it when	Evidence	Last reviewed
Example: Support reply drafter Drafts replies to tickets for review.	Customer data, email and messages	Its own identity	Head of Support Backup: Support lead	On its own: read, summarize and draft Approved by the agent on duty: changes it can undo Never: anything it can't undo	Head of Support or on-call engineer Under 5 minutes	A reply sent without review	Logs: agent, who it acted for, approval Tested: [date]	[date]
Example: Release helper Opens pull requests, runs tests.	Code, production systems	Its own identity	Chief Technology Officer Backup: Engineering lead	On its own: read, summarize and draft Approved by the code owner: changes it can undo Never: deploy, delete or spend	On-call engineer Under 5 minutes	Changes outside its repositories	Logs: agent, who it acted for, approval Tested: [date]	[date]
[Agent] [What it does]			[Owner] Backup: [name]	On its own: Approved by [name]: Never:	[Who] [How long]		Logs: Tested:	
[Agent] [What it does]			[Owner] Backup: [name]	On its own: Approved by [name]: Never:	[Who] [How long]		Logs: Tested:	
[Agent] [What it does]			[Owner] Backup: [name]	On its own: Approved by [name]: Never:	[Who] [How long]		Logs: Tested:	
[Agent] [What it does]			[Owner] Backup: [name]	On its own: Approved by [name]: Never:	[Who] [How long]		Logs: Tested:	

Signed	Title	Date	Next review

Answers for the AI section of a DDQ

HOW TO USE THIS TEMPLATE

Each answer comes from one column of the signed map, so update the map first and the answers follow.

Replace every bracket. Have whoever signs the map read every answer before it goes out, because some contracts turn questionnaire answers into warranties.

Send the map itself with the questionnaire when you can. A reviewer who sees the whole page asks a few follow-ups instead of fifty questions.

Who is accountable for the AI agents in your product and operations?

Asked by: Microsoft's supplier data protection requirements; the UK government's AI Management Essentials

[Company] keeps a one-page accountability map of the [N] AI agents we build and run. Each agent has one named owner who answers for it and a backup, and [the executive] signs the map [every quarter].
[Agent: owner, backup.]

From the map: Owner and backup, and the signature

Which actions can your AI agents take on their own, and which need human approval?

Asked by: Reco's AI vendor security questionnaire

We decide by consequence: reading and drafting, changes we can undo, and anything that can't be undone, spends money or reaches a customer. [Agent: may read, summarize and draft on its own; needs approval from [name] to change things it can undo; may not take actions it can't undo.] Approvals are recorded with the approver's name.

From the map: Decision rights

Can you turn an AI agent off, who can do it, and how quickly?

Asked by: Microsoft's supplier data protection requirements; Reco's AI vendor security questionnaire

Yes. Each agent has a named person who can turn it off, and we time it. [Agent: [who], under [time] when last tested; we pause it when [conditions].]

From the map: Off switch and pause conditions

Do your AI agents use their own identities, or users' and shared credentials?

Asked by: Reco's AI vendor security questionnaire

Every agent runs under its own scoped identity, never an employee's account or a shared key, so every action is attributable to the agent that took it. [If not yet: [N] of [N] agents do today; the rest move by [date].]

From the map: Runs as

How are AI agent actions logged and attributed?

Asked by: Kognitos's agentic AI RFP template

Every agent action is logged with the agent's identity, the person it acted for and any approval, in logs the agent can't edit, so we can answer who did what, when and why.

From the map: Evidence: what the logs record

How do you test your AI agents?

Asked by: The UK government's AI Management Essentials, which asks whether you monitor AI systems to check they perform as expected

We test each agent's limits [before each release and after any model or prompt change], including attempts to make it act outside its job, and we record the results. [Agent: last tested [date].]

From the map: Evidence: limits last tested

Who reviews AI risk, and how often does leadership see it?

Asked by: The UK government's AI Management Essentials; NIST's AI Risk Management Framework

[The executive] reviews and signs the accountability map [every quarter], and [the board or its audit committee] receives a one-page summary: the agents and what they reach, their owners, changes since the last review, incidents and near misses, off-switch and limit test results, and any exceptions someone approved.

From the map: Signature and review cadence

Who is responsible when an AI agent's action causes harm?

Asked by: Reco's AI vendor security questionnaire

We are. [Company] takes responsibility for the agents we build and operate, the prompts we write and the permissions we grant them. [How liability is allocated between us, you and our model providers is set out in section [X] of our agreement. Confirm this wording with counsel.]

From the map: Your contract, with the map as the evidence behind it

Quarterly board summary

HOW TO USE THIS TEMPLATE

One page, built from the map, for the board or its audit committee every quarter.

Keep it short and dated. Directors don't need the detail, but they need to see that someone owns it.

Item	This quarter	Notes
Agents and what they reach	[N] agents; [N] reach customer data; [N] reach money	[New reach since last quarter]
Owners	[N] of [N] with a named owner and backup	[Changes in owners]
What changed	[New agents, retired agents]	
Incidents and near misses	[None / summary]	[What changed after]
Off-switch drills	[N] agents turned off; slowest took [time]	[Date]
Limit tests	[N] agents tested; [findings]	[Date]
Exceptions approved	[What, and who approved it]	
Next review	[Date]	

Three questions any director can ask

1. Who answers for this agent, and who is their backup?
2. Pick one action at random. Can we produce the agent, the policy, the authority, the evidence and the exception owner?
3. Who can stop it, how fast, and what would make them?

The second and third questions adapt Manny Sandhu's random-transaction test and Probal Lala's advice-to-action test, both posted on LinkedIn in October 2026.

Five working days to a signed map

1. Day 1: list every agent. Product features and internal tools, including the one a single engineer set up on a Friday.
2. Day 2: owners and identities. One named owner and a backup per agent. Move every agent off borrowed and shared credentials onto its own identity.
3. Day 3: decision rights. Tier each agent's actions by consequence, put a named approver on the irreversible few, and write down what it never does.
4. Day 4: off switches and evidence. Turn each agent off and time it. Write the pause conditions. Check that the logs show the agent, the person it acted for and any approval. Test each agent's limits and record the date.
5. Day 5: sign it. The executive reads every row and signs. Send the board its summary, draft your DDQ answers from the map, and put the next review on the calendar.

Want a second pair of eyes?

We'll read your map the way your buyer's reviewer will and tell you what we'd change first. Book a free 15-minute security strategy call:

cal.com/ysecurity/15min

The full guide, with the map builder and the sources behind every question:

ysecurity.io/blog/who-signs-for-the-agent

YSecurity is the on-demand cybersecurity team. This kit is general guidance, not legal advice. Northwind AI and its agents are made up.