

Your AI Agent Has an Identity. Who Owns It?

By Junior Williams

Draft for review · 14 October 2026

When an AI agent reads a document, updates a record or calls another system, it needs a way to obtain access. That may involve a person's delegated permissions, a service account, an application identity or an API credential.

The model can propose an action. The surrounding software and its credentials determine what the agent can actually reach and do.

That makes a familiar question worth asking: who owns this access?

Giving someone a building pass involves more than printing their name on it. Someone decides which doors it opens, when access expires and who can cancel it. The same decisions belong in an AI workflow, even when the access is used by software.

Start with what the agent can do

Consider a hypothetical assistant that prepares responses to customer support requests. It reads a case, searches approved guidance and drafts a reply for a member of staff to review.

The team later connects it to a customer records system so it can gather more context. Then someone proposes letting it update the case status and send the reply after approval.

Each addition changes the access required. Reading guidance, reading customer information, changing a record and sending a message are separate capabilities. An approval to build the first version does not settle the authority needed for every later version.

I would map those capabilities before deciding how much autonomy to give the assistant. The useful starting point is a list of actions, the systems involved and the identity used at each step.

One agent may use several identities. Several agents may also share access through the same integration. A friendly name in a chat window will not necessarily tell you which account performed an action.

Make ownership specific

“The AI team owns it” leaves several decisions open. A workable ownership arrangement should identify who can answer these questions:

- Who approves the business purpose and the actions the agent may perform?
- Who maintains its integrations, credentials and configuration?
- Who reviews its access when its role changes?
- Who can suspend it, and who handles the resulting interruption?

These responsibilities can sit with different people. The important point is that they are assigned and understood. The person who can fix an integration may have no authority to approve a new use of customer information.

Access should fit the workflow's current purpose. If the assistant prepares a draft, consider whether it needs the ability to send it. If it searches approved guidance, establish which repositories belong in that scope. Where an action needs approval, design the workflow so the approval applies to the specific action and relevant content.

Instructions help describe the intended behaviour. Permissions and application controls determine the access available when the workflow runs. Review both.

Keep an identity record you can use

A short record for each agent or workflow can make this manageable:

Field	What to capture
Accountable owner	The person responsible for the workflow and its approved purpose.
Identity and credentials	The identities used at each connection and where credentials are managed; do not copy secrets into the record.
Access scope	The information it may read and the actions it may perform.
Approval points	Which actions require a person, and who is authorized to approve them.
Lifecycle	Access review dates, credential handling and events that require a change.
Suspension and evidence	How access can be withdrawn and which records support investigation.

Keep the record connected to the operating process. Review it when a new connector is added, the agent takes on a different task or the accountable owner changes. Retire access when the workflow is retired.

Check whether stopping it actually works

An operating review should establish what happens when the team suspends the workflow. Are existing sessions or queued actions still active? Do other integrations depend on the same credential? What work needs to be completed manually while access is unavailable?

These are useful questions for a controlled review with the system owners. The outcome should be an understood stop-and-recovery procedure, with evidence of its limits.

I would expect a leader to be able to identify who authorized an agent's access, what it covers and who can withdraw it. That clarity helps the technical team operate the workflow and helps the business understand the responsibility it has accepted.

JWC's AI Exposure and Control Assessment includes reviewing ownership and access within agreed AI workflows. To discuss the scope, [book an introductory call](#).

Further reading

The [NIST AI RMF Playbook](#) calls for clear responsibilities (Govern 2.1) and assigned responsibility for disengaging systems (Manage 2.4). The identity record above is a practical way to apply those ideas to an agent workflow.