

The Agent Authority Assessment

Seven questions. About ten minutes. Best answered with your leadership team.

WHY THIS MATTERS

Every AI agent that can act needs two kinds of ownership. The first is ownership of the work: someone configures the agent, grants its access, and keeps it running. That kind almost always has an owner.

The second is ownership of oversight, and it has three parts: someone **notifies** when the agent acts outside its authorized scope, someone has the authority to **stop or contain** it, and someone **decides who needs to be told**. Each part can have an owner while the handoffs between them have none.

This assessment shows where oversight of your AI agents sits today, and where the handoffs between those parts are unowned.

HOW TO SCORE

- 0 No, or we don't know
- 1 Partly, or for some agents only
- 2 Yes, consistently, and we could show the evidence

Answer for the organization as it runs today, not as the policy describes it. Include agents built into software you have purchased, not only the ones your team built. If you are unsure, score it 0. Uncertainty is a finding.

YOUR WORKING AREA

Record what this assessment covers: the whole organization, one department, or one pilot.

SCOPE ASSESSED		DATE	
PARTICIPANTS		EVIDENCE REFERENCED	

The Seven Questions

Select or circle one score for each question.

1	VISIBILITY We can list every AI agent that can take action in our environment, including agents built into software we have purchased.	0 ☐ 1 ☐ 2 ☐
2	SCOPE For each agent, what it is authorized to do is written down, not assumed.	0 ☐ 1 ☐ 2 ☐
3	OWNERSHIP Each agent has a named person accountable for its authorized scope.	0 ☐ 1 ☐ 2 ☐
4	REVIEW A named owner reviews what our agents actually did against their authorized scope, on a schedule that fits the consequences, and we have evidence the review catches deviations.	0 ☐ 1 ☐ 2 ☐
5	VERIFICATION When a request claims special authorization (a test, an exception, an executive request), that claim is confirmed through a trusted approval source independent of the request itself before the agent acts.	0 ☐ 1 ☐ 2 ☐
6	STOP AUTHORITY A named person can pause or contain each agent within a response window set by what is at stake, and we have evidence the team could meet that window.	0 ☐ 1 ☐ 2 ☐
7	HANDOFFS The handoffs between noticing, stopping, and telling each have a named owner, including who decides who gets told.	0 ☐ 1 ☐ 2 ☐

YOUR TOTAL

/ 14

Check questions 4, 6, and 7 first. A 0 on any of them means your team cannot confirm reliable detection, intervention, or escalation. Treat that uncertainty or missing capability as a critical gap, regardless of the total. Individual gaps take priority over the total.

What Your Score Means

The bands guide discussion; they do not certify readiness. A critical gap on question 4, 6, or 7 takes priority over any band.

11 to 14

Oversight Largely Defined

Most parts of oversight have owners, with few partial answers. The total does not show that every agent is covered or that the design works in practice. Start with any question scored below 2, then test it: run a short exercise with your highest-consequence agent and time how long it takes to notice, contain, and notify.

6 to 10

Split Oversight

Oversight exists in places, but some parts or handoffs are missing or partial. The gaps are easy to miss because other parts look covered. Start with your lowest-scoring questions and the highest-consequence agent they affect, and close those gaps first.

0 to 5

Unowned Oversight

Your agents may be acting with no one clearly positioned to notice, stop, or escalate. Nothing may have gone wrong yet; this kind of gap tends to stay invisible until something does. Start by mapping every agent that can act, then address questions 4, 6, and 7 for the highest-consequence one.

WHAT TRANSFORMXPERIENCE BUILDS

TransformXperience is an operating-model and accountability consulting practice, not a software vendor. For AI agents, TransformXperience builds two deliverables: an **agent authority map** showing every agent that can act, including embedded ones, its authorized scope, who reviews what it did, and who can stop or contain it; and an **escalation path** naming who gets notified, who makes the disclosure call, and response windows set by the potential consequence of an action. TransformXperience designs it, then hands it to your team to run.

YOUR NEXT STEP

Book a 90-minute Executive Readiness Session. We will baseline where oversight of your AI agents sits today and name your priority gaps, including decisions no one owns.

We architect it. You own it. You run it.

Michelle McKinney, MBA, CSM | Founder and CEO, TransformXperience