AI Agent Governance Checklist for SMEs

12 Questions to Ask Before Connecting AI to Your Business Systems

Clear Gate Systems | cleargatesystems.com *Eileen Weadick, PhD | AI Governance for Irish SMEs*

How to use this checklist

This checklist is for business owners and directors who are introducing AI tools that can take actions in their systems: accessing files, sending communications, modifying records, querying databases, or triggering processes in connected software.

It is not a technical document. It is a governance document. Every question on it can be answered by a decision-maker without a technical background. If you cannot answer a question confidently, that is useful information. It tells you where to focus.

Work through the checklist before any AI agent goes live in a business system. Come back to it whenever the scope of an AI deployment changes.

Section 1: Authority and Access

Question 1: Is there a written description of exactly what this AI agent is permitted to do, and what it is not permitted to do?

A working instruction is not the same as a governed one. "Summarise customer emails" is an instruction. "Summarise customer emails, flag anything that requires a response within 24 hours, and do not take any further action without approval" is a governed instruction. If the boundaries are not written down, they do not exist in any meaningful sense.

Answer: Yes / No / Partially

Question 2: Does the agent have the minimum access needed for this specific task, and nothing beyond that?

An agent given access to your full file system to complete a task in one folder has more authority than it needs. An agent that can read and write and delete when the task only requires reading has more authority than it needs. Excess access is a governance gap, not a convenience.

Answer: Yes / No / Partially

Question 3: Have we reviewed exactly which credentials, tokens, and system connections the agent can reach?

A credential is a digital key that grants access to a system or service. An access token is a specific type of credential that software systems use to communicate with each other. In several documented AI incidents, agents found and used credentials stored in unrelated files or folders, with access to systems far beyond the intended task. Know what your agent can reach.

Answer: Yes / No / Partially

Section 2: Human Oversight

Question 4: Is there a defined list of action types that require a human to approve before the agent can proceed?

At minimum, this list should include: deleting or modifying records, sending external communications, altering contracts or financial records, triggering regulatory submissions, and modifying live infrastructure. The list should be documented, not assumed.

Answer: Yes / No / Partially

Question 5: Is there a technical control, not just an instruction to the model, that prevents destructive or irreversible actions without human confirmation?

This is the most important distinction in AI agent governance. There is a meaningful difference between telling the agent not to do something and designing the system so

that it cannot. In documented incidents, agents disregarded explicit safety instructions. Governance that depends entirely on the model's behaviour is not governance. The surrounding system architecture should make the unauthorised action structurally impossible.

Answer: Yes / No / Partially

Question 6: Is there a named person in our organisation who is accountable for AI governance?

Someone in your organisation needs to own this. Not as an IT responsibility, but as a leadership one. If an AI agent makes a consequential error, there should be a named individual who can explain what governance was in place, who authorised the deployment, and what the review process was. Accountability is not bureaucracy. It is protection.

Answer: Yes / No / Partially

Section 3: Visibility and Auditability

Question 7: Does the agent log every action it takes, in a format a human can read and audit?

A log is a record of what happened, when, and using which tools. Without logs, you cannot reconstruct the sequence of events during an incident, demonstrate to a regulator or client what the AI did or did not do, or identify the specific governance gap that allowed a problem to occur. Logging is also a legal requirement for certain AI systems under Article 12 of the EU AI Act.

Answer: Yes / No / Partially

Question 8: Can we reconstruct a full record of what the agent did during any given time period?

Logging is the first step. Retrievability is the second. Logs that exist but cannot be searched, organised, or presented in a readable sequence are not much use during an incident investigation. Before deployment, test whether you can actually pull a complete and readable action history.

Answer: Yes / No / Partially

Section 4: Testing and Resilience

Question 9: Have we tested this agent against realistic failure scenarios before connecting it to live systems?

A successful demonstration is not a governed deployment. Testing should include: ambiguous or incomplete instructions, conflicting inputs, attempts to access out-of-scope systems, sensitive data scenarios, and recovery from partial failures. If the agent has only been tested in ideal conditions, it has not been adequately tested.

Answer: Yes / No / Partially

Question 10: Can we pause or shut down this agent immediately if something goes wrong?

Before going live, you should be able to answer this question: if we needed to stop this agent right now, what exactly would we do, who would do it, and how long would it take? If the answer is not clear, the deployment is not ready.

Answer: Yes / No / Partially

Question 11: Can we revoke the agent's access to all connected systems quickly?

Stopping the agent is one thing. Revoking its access to email, databases, cloud infrastructure, and external services is another. Both should be possible, and the process for doing so should be documented and tested before deployment, not figured out during an incident.

Answer: Yes / No / Partially

Section 5: Regulatory Alignment

Question 12: Do we know which provisions of the EU AI Act apply to this AI system and when?

The EU AI Act is directly applicable in Ireland without the need for separate national legislation. There is no SME exemption. Prohibited practices under Article 5 applied from 2 February 2025. Transparency obligations under Article 50 apply from 2 August 2026.

Full obligations for high-risk AI systems apply from 2 December 2027. You do not need to be a legal expert to understand which category your AI system falls into, but you do need to have asked the question.

Answer: Yes / No / Partially

Reading your results

All twelve answered Yes: Your governance architecture is in good shape for this deployment. Review it again whenever the scope of the deployment changes, the agent is given access to new systems, or the underlying AI model is updated.

One or more answered Partially: Partial answers indicate governance that exists on paper but may not be fully implemented or tested in practice. Address each partial answer with a concrete action: who is responsible, what specifically needs to be done, and by when.

One or more answered No: Each No is a documented governance gap. This does not mean the deployment must stop, but it does mean a decision needs to be made and recorded: either the gap is addressed before going live, or a conscious decision is taken to accept the risk with the reasons noted. Undocumented risk is the problem, not risk itself.

A note on the EU AI Act

Ireland's Department of Enterprise confirmed in May 2026 that a provisional agreement has been reached on the Digital Omnibus on AI, which introduces targeted amendments to the EU AI Act intended to increase legal certainty and reduce compliance costs for businesses. These changes do not remove obligations. They make implementation more workable.

Fines for the most serious breaches of the EU AI Act reach up to €35 million or 7% of global annual turnover. For smaller violations, fines reach up to €15 million or 3% of annual turnover.

About Clear Gate Systems

Clear Gate Systems designs and implements auditable AI governance workflows and operating systems for SMEs in Ireland. Our work produces practical, working governance: not policies that describe what controls should exist, but documented systems that demonstrate how AI is used and governed in practice, and that can be shown to regulators, boards, and clients.

If this checklist has raised questions about your own governance architecture, a Clarity Call is a good starting point. It costs nothing and carries no obligation.

Book a Clarity Call at cleargatesystems.com

This document is for informational purposes only and does not constitute legal advice.