October 2026

Before an AI agent touches your ERP, check six things.

SAP announced a gateway that checks and logs outside agents.

Oracle announced ChatGPT plugins that act on Oracle systems.

IBM shipped an identity and an off switch for each agent.

1

The agent has its own login.

Not a shared API key, not the boss's password. One account per agent, with a named person who owns it.

> Test: can you tell, for any change in the system, whether a person or an agent made it?

2

It gets its owner's rights, or fewer.

An agent working for a sales rep sees what that rep sees. It never sees payroll because the integration was easier with an admin account.

> If an agent can do something its owner can't, fix that first.

3

Read-only is the starting point.

Answering "what's in stock?" or "who owes us money?" is useful from day one and can't break anything.

> Give write access one action at a time, after you've watched it read for a few weeks.

4

A person approves money and stock.

The agent can prepare the price change, the supplier payment or the stock adjustment. Someone clicks approve.

> **Prepared by the agent, approved by a person.**
> That's the whole rule.

5

Every action lands in a log your team can read.

Not a debug file only the developer opens. A plain list: what the agent did, when, for whom, and what it changed.

> Monday morning question: "what did the agent do over the weekend?" The answer should take ten seconds.

6

You can switch it off in one click.

Revoke the agent's access without changing anyone's password and without calling the vendor.

> If turning it off needs a support ticket, you don't control it.

You don't need an enterprise contract for this

We build business software where agents follow these six rules by default.

Stock, sales, invoicing and distribution systems, with AI agent access built in. Save this post for the day someone asks to "just connect ChatGPT".

> **sindansolutions.com**
> WhatsApp +213 549 08 91 42