

DIGITAL ADVENTURES / SALES × AI × LEADERSHIP

When AI acts in sales, who is accountable?

Mario Pufahl · 25 September 2026

Decision rights, approvals and KPIs for sales agents

The previous issue explored the new AI sales stack. CRM remains the system where customers, opportunities and activities are recorded. AI adds a system of action: information is interpreted, work is prepared and tasks are increasingly executed independently.

This sounds like technological progress. In practice, it quickly raises a leadership question: who decides when a sales agent acts?

An agent can analyse a buying group, recommend the next step, prepare a proposal or draft a customer message. But may it also send that message? Change a price? Reassess an opportunity? Adjust a forecast?

If you wait until the first error to answer these questions, you have started managing the system too late.

The real issue: work is delegated, accountability remains

In traditional sales, responsibilities are usually clear. The account manager owns the customer relationship. The sales lead prioritises resources and reviews critical deals. Finance and legal approve certain terms.

AI creates new handovers. A person gives an instruction, one agent accesses data, another assesses outputs, and a workflow executes an action. Work is distributed, but accountability cannot be handed over to the technology.

Sales therefore needs a simple rule: every automated action has a business owner, a defined scope of action and a traceable audit trail. Without these three elements, relief can quickly become loss of control.

Not every action needs the same approval

Reviewing every AI response manually yields little speed. Letting every action run automatically puts trust, margin and customer relationships at risk. Differentiate by impact and risk. I would distinguish four levels:

1. Inform. The agent gathers information, summarises signals or prepares a briefing. People use it to decide. For low-impact work, sampled quality checks may be sufficient.

2. Recommend. The agent prioritises accounts, identifies risks or suggests next actions. It should expose its data, assumptions and uncertainty. The responsible person decides.
3. Prepare. The agent drafts a proposal, email or forecast. A person checks and approves it before it reaches the customer or takes effect in an authoritative system.
4. Execute. The agent performs tightly defined steps, such as creating an internal task or triggering an approved standard process. This requires boundaries, logging and a way to stop or correct actions.

The right level depends on the consequence of the action, not the agent's technical capability. Misclassifying an internal task is not the same as sending incorrect pricing to a strategic customer.

The key question: what may the agent change?

Read access alone is not enough. Leaders must also define what an agent may write, send or trigger. Start with five questions:

- Data: Which customer, contract and employee information may it access?
- Action: May it only draft, or also change CRM data and contact customers?
- Boundary: Which amounts, customer segments and exceptions rule out automatic execution?
- Approval: Who approves higher-risk steps, and where?
- Evidence: Where can inputs, recommendations, approvals and changes be reviewed later?

This makes governance concrete. "The human stays in control" is not enough; define which person exercises which control at which point.

An example from the proposal process

Imagine an agent preparing a proposal for an existing customer. It reads the opportunity, uses current price lists, identifies an upcoming renewal and drafts a proposal. It may initially create the draft and flag missing information. The account manager checks context, value and contacts. Finance approves any non-standard terms. Only then is the proposal sent and the final version stored in CRM.

Part of a standard renewal process might later be automated. A complex tender or unusual discount still requires human approval. The aim is a decision framework for each transaction, not a single automation level for all sales.

Leadership needs different metrics

Counting prompts or agents says little about sales impact. Assess performance in four dimensions:

1. Quality: Are outputs accurate, complete and useful for the customer case? How many need major rework?

2. Process: How long does it take to reach an approved action? Where do questions, handovers and waiting occur?
3. Risk: How often are boundaries breached, incorrect data used or approvals bypassed? Which errors are caught in time?
4. Business impact: Are proposals faster and better? Does pipeline quality improve? Is more time available for relevant customer conversations? Do conversion and contribution margin improve in comparable cases?

Not every change can be attributed causally to an agent. Establish a baseline and compare similar cases rather than crediting AI with every positive development. Ask which specific process has demonstrably improved.

The sales leader designs decision rights

Leaders continue to manage people, while also designing collaboration between people, data and digital agents. They need to make responsibilities and conflicts explicit:

- Who decides on escalation when an agent flags an opportunity?
- What happens when its recommendation conflicts with an account manager's judgement?
- Who owns a message that is technically correct but culturally inappropriate?
- Which priority wins if a metric improves but the customer relationship suffers?

These are not edge cases. They describe a hybrid sales team's everyday work. Technical permissions alone miss the point: AI shifts decision work, and that requires leadership.

How I would start in the next 30 days

1. Week 1: Select one process. Choose a frequent, bounded workflow such as meeting preparation or standard proposals. Describe current effort and typical errors.
2. Week 2: Define decision rights. Specify where AI may inform, recommend, prepare or execute. Name business owners and approval points.
3. Week 3: Test real cases. Review a manageable sample. Document corrections, exceptions and situations where the agent should deliberately do nothing.
4. Week 4: Assess impact and boundaries. Measure processing time, quality and interventions. Ask users whether the support helps. Only then consider expanding its scope.

This is less spectacular than fully autonomous sales, but it is how a demonstration becomes a dependable process.

A decision register, not just a general AI policy

Many companies already have principles covering confidentiality, privacy and approved tools. An acting sales agent needs rules for specific situations. A decision register should capture at least seven fields for each use case: trigger, data used, possible action, permitted scope, accountable business owner, approval point and evidence. Add an explicit exception rule.

For meeting follow-up, the trigger might be an approved meeting record. The permitted action could be drafting internal tasks and updating specified CRM fields. Customer statements must remain identifiable. The agent must not invent commitments, change prices or send external messages without approval. The account manager reviews the draft before the record changes.

Creating a standard internal task may justify automatic execution; making a statement to a customer is a higher-risk action. The difference lies in the consequence, not the model. A register also supports scaling: another team can inspect and adapt the permitted action instead of copying a successful pilot into a different risk context.

Accountability needs names and an escalation path

Each agent needs a named business owner. That person does not personally generate every technical response, but decides the agent's purpose, acceptable quality and when its use must change or stop. Data, technical operations and oversight also need owners. Sales leaders own how outputs inform customer decisions; legal, security or compliance may need to participate in critical processes.

- Business owner: Why does the agent exist, and what outcome should it improve?
- Process owner: Where does it intervene, and who takes over when exceptions occur?
- Data owner: Are sources accurate, current and permitted for this purpose?
- Technical operations: Is it running as intended, and can faults be isolated quickly?
- Approver: Who decides on pricing, commitments and externally consequential steps?

What happens if nobody approves in time? The system should wait or hand over to a person. Silence must not count as permission.

What effective approval actually checks

"Human in the loop" sounds reassuring, but approval can become a routine click when someone must review dozens of suggestions quickly. Effective approval shows the basis for the decision, not just the finished text: data used, assumptions, deviations from the standard case and the exact consequence of approval.

For a proposal, show existing customer terms, suggested price, contribution margin, discount justification, non-standard clauses and recipients. Approval should sit within the employee's normal workflow. A generic "please approve" email can become a formality.

Corrections must feed back into better data, rules or agent behaviour. Otherwise the organisation repeatedly pays for the same error through technology costs and staff rework.

When an agent acts incorrectly

Errors will occur. The key is detecting them quickly and limiting their impact, not promising an error-free system. Before production use, test three scenarios: the agent uses a wrong source, misinterprets correct information, or carries out a permitted action in the wrong customer context.

For each scenario, define who is notified, which actions stop, which records or customers are affected, who corrects the information and when work may resume. A stop mechanism must be accessible in everyday operations. Requiring a developer to intervene while a sales leader watches an error unfold creates unnecessary delay.

Record completed and pending steps. External messages cannot always be recalled, so safeguards should be stronger where reversal is harder. An internal summary can be corrected; an incorrect customer commitment must be actively explained.

The business case must include rework

Time saved is only part of the calculation. Suppose automated proposal preparation saves 20 minutes in a standard case, but every second draft needs 25 minutes of correction. The expected saving is small. An additional approval step could even make the process slower.

Track four elements separately: original task time, review and correction time, technical operating costs, and impact on quality and business outcomes. Only the combined picture reveals value.

A faster process may still be valuable without saving much labour if proposals become more consistent, generate fewer questions and give customers a reliable answer earlier. Conversely, large time savings may have no value if they damage the relationship. Compare like with like: standard proposals with standard proposals, complex tenders with complex tenders.

From metrics to leadership decisions

A dashboard helps only when its signals trigger decisions. Falling quality needs a defined response. Increasing manual interventions should prompt a review of data, rules or the use case itself.

For a pilot, start with a few clear measures: usable first drafts, average rework time, customer-impacting errors, time to approval and the share completed without further questions. Add a business outcome such as proposal quality or time to the jointly agreed next step.

For each metric, name an owner and a response. Quality below a threshold may trigger extra review. Repeated source errors may justify blocking that source. Repeated approval of a

low-risk standard action may support a controlled expansion of scope. AI performance management becomes a learning cycle: observe, understand the cause, change the rule and measure again. The number of agents is not a success metric.

Scaling is a management decision

After a successful pilot, do not assume the same agent belongs everywhere. Customer segments, pricing rules, contracts and data quality differ. A standard renewal is not a strategic new-customer tender.

Before expanding, check data comparability, process similarity, approval rights, employee understanding and whether escalation works at higher volume. If these remain uncertain, the next step may be using the agent only for advice or drafting. Scaling means transferring a reliable mechanism, not simply increasing automated actions.

My recommendation

The next step is not the agent with the most features. It is an organisation that designs digital work so decisions improve, processes become faster and results more dependable. That requires clear decision rights, proportionate approvals and metrics showing quality and business impact.

AI can take on tasks. Accountability remains a leadership responsibility.

How does your sales organisation decide which decisions an AI agent may prepare or execute?

First published in Digital Adventures. English translation by Sales.Decoded.

[Read the original on LinkedIn](#) · sales-decoded.com/en/