Luminote

Playbook · 2026 Edition

2026 Sales Playbook.
AI without the hype.

What artificial intelligence actually changed in B2B sales, where it returns something you can measure, and where it costs more than it delivers. An honest balance of upsides, downsides and the seven traps most teams walked into after starting in 2024. **No slides about revolution.**

20 pages
upsides, downsides, traps

Principle
the machine drafts, the human approves

For
sales directors, RevOps, founders

What's inside

INTRODUCTION

The end of the excitement phase

2024 was the year of demos. 2025 was the year of pilots. 2026 is the first year in which somebody asks for the invoice — and that question is healthy.

Most sales teams have had their first contact with AI: an assistant for writing emails, a sequence generator, a meeting transcription plugin. Some of those tools stayed. Some were uninstalled after three weeks. Very few were ever held to account for their impact on revenue.

This playbook is neither a manifesto nor a warning. It is a division: **where AI in sales gives an advantage large enough to show up in the quarter, and where it manufactures the appearance of work**. In a demo, the two look identical.

We start from one assumption worth stating plainly: a good use of AI in sales is work a rep does not want to do but which has to be done accurately. The note. The qualification. The CRM update. The recap of what was agreed. A bad use is usually anything that requires judgement and taking responsibility for it.

Inside you'll find a map of use cases, seven traps described by their symptoms, an honest list of downsides — including the ones that apply to us — plus a set of metrics and a 90-day rollout plan.

And one principle we repeat in every piece we publish: **the machine drafts, the human approves**. Nothing reaches the CRM or the customer without the rep's decision.

How to read this playbook

Chapters 1 and 3 are diagnostic — read once. Keep chapter 2 open whenever you evaluate a new tool. Chapter 4 and the plan on page 19 are to be done, not considered.

CHAPTER 1

Where AI actually works

Not everywhere it gets deployed, and not where it makes the biggest impression in a demo. **The value is concentrated in a handful of very boring places in the process.**

1.1 · THE MAP

Works, limps, harms

A simple split that saves a quarter of experiments: the more a task is about faithfully recording what already happened, the better it turns out. The more it is about taking responsibility, the worse.

USE CASE	VERDICT AFTER TWO YEARS OF PRACTICE
Notes and CRM entry	Works. Closed task, verifiable, immediate saving
Qualification from a call	Works, if it rests on quotes rather than the model's judgement
Recap and follow-up	Works with human approval. Without it, a reputational risk
Pre-meeting research	Limps. Saves time, but every fact needs checking
Lead scoring	Limps at low volume — the model learns on too short a history
Model-generated forecast	Limps. Only improves once CRM data is complete
Mass generated cold outreach	Harms. Short-term lift in sends, long-term cost to the brand
Autonomous replies to customers	Harms. Responsibility for the relationship stays with you

A test before buying any AI tool

Ask whether the output can be checked in thirty seconds. If yes, the task suits automation. If verification takes as long as doing the job by hand, the tool moves work around rather than removing it.

1.2 · DIAGNOSIS

Four levels of team maturity

Before planning the next step, establish where the team stands today. Most organisations that "have AI" sit on level one and believe they are on level three.

01 **Private tools.** Some reps use a chatbot to write emails, on their own initiative and their own accounts. Nobody measures it, company data leaves the building without a contract. Effect: erratic. Risk: real.

02 **Isolated plugins.** The team has meeting transcripts or a sequence generator. A record exists, but it doesn't reach the place where decisions are made. A human still fills in the CRM.

03 **Wired into the process.** The call automatically feeds the deal: note, qualification, commitments, follow-up ready to approve. The rep stops re-typing. This is where the first effect shows up in the calendar.

04 **Decisions on call data.** The manager runs pipeline reviews on facts, coaching draws on patterns across the whole team, and the forecast rests on what customers actually said. The level at which AI stops being a tool and becomes a data source.

The jump that pays best

Moving from level 2 to level 3 returns fastest and takes the least work — because you are not changing the sales process, only removing the manual re-typing from it. The jump from 3 to 4 requires managerial discipline, not technology.

1.3 · CONTEXT

Three things that genuinely changed

Behind the noise, three shifts actually moved the line of what is worth automating in sales.

1. Transcription stopped being the problem

Not long ago, speech recognition with industry jargon and a poor microphone was a barrier to entry. Today the quality of the record is high enough that the discussion has moved from "does it understand" to "what do we do with it".

2. Unit cost fell to the point of irrelevance

Processing an hour-long call costs a fraction of what it did two years ago. That is why it now pays to analyse every call rather than a sample — and that is a qualitative change, not a quantitative one.

3. The customer's benefit of the doubt ran out

Recipients have learned to spot a generated message. Personalisation that stood out in 2023 can now read as a warning sign. Advantage comes from relevance, not volume.

An inconvenient conclusion

Taken together, these three shifts mean AI in sales pays off today mainly *inward* — organising your own data and taking load off the team. Outward-facing use cases, aimed at the customer, now carry a worse ratio of return to risk than two years ago.

CHAPTER 2

Traps of the rollout

Each looks like a good decision at the start and each costs a quarter. We describe them by symptom so they can be spotted **before they show up in the numbers**.

2.1 · CATALOGUE

Seven traps of rollout

TRAP	HOW YOU RECOGNISE IT
Tool instead of problem	The rollout began with a product name, not with a number you want to improve
Verification costlier than the work	The rep reads the generated text longer than it would take to write their own
Garbage in	The model gets an incomplete CRM and produces convincing untruths
Automating a bad process	You sped up a step that should have been deleted instead
Recordings used as control	The team speaks more cautiously in meetings, data quality drops
Purchase without an owner	Licences bought, nobody owns adoption, 20% log in after a quarter
Building your own agent	Prototype in a week, maintenance forever, sales ends up maintaining software

The first six can be reversed within a month. The seventh usually costs a year, because nobody wants to be the person who shuts down a project announced at an all-hands. **Agree the exit criteria before the project starts** — a date, a metric, and consent that no result means the end, not another iteration.

The unnamed eighth trap

Executive enthusiasm. If the tool was chosen at a conference rather than in a conversation with reps, adoption will be performative — the team will learn to report usage, not to use it.

2.2 · ANATOMY OF A FAILURE

The pilot that never ends

The most common way an AI rollout in sales dies: not by rejection, but by never being decided.

Pilot start	March, three people, "let's see how it works"
Defined success metric	none
Baseline measured before start	none
Owner on the sales side	"everyone a bit"
Decision in June	extend the pilot
Status in October	two users left, topic returns at budget time

None of these lines is about technology. All of them are about a decision not taken at the start. A pilot without a baseline can neither succeed nor fail — it can only drag on, and dragging on looks safer inside an organisation than closing something down.

Three conditions for a pilot that gives an answer

01 **A number before the start.** Admin hours per person, time to follow-up, field completeness on deals. Measured the week before, not reconstructed later from memory.

02 **A decision date.** In the calendar, with the name of the person who decides. Four to six weeks is enough.

03 **Consent to a negative result.** Said out loud at the start. Without it every pilot ends in success and none changes the number.

2.3 · THE COSTLIEST TRAP

Prospecting on steroids, i.e. spam

Being able to send ten times more personalised messages sounds like leverage. In practice it is the fastest way to burn a domain and a brand at the same time.

The mechanism is simple: once the cost of writing a message drops to nearly zero, everyone sends them. The recipient's inbox fills with content that is correct, polite and empty. They stop reading the whole category, not just your message.

DIMENSION	GENERATED VOLUME	HUMAN RELEVANCE
Effect in month one	More sends, more meetings	No change
Effect after a quarter	Reply rate falls, deliverability suffers	Steady rise in meeting quality
Hidden cost	Domain reputation, brand, market fatigue	Rep time spent on research
Who detects it	Filters and recipients — both quickly	Not applicable
Reversibility	**Months of rebuilding the domain**	**Immediate**

This is not an argument against AI in prospecting. It is an argument for a different use of it: **to prepare the rep before a conversation, not to send messages instead of them**. The same model that generates hundreds of emails can gather in a minute the context a rep would have spent fifteen minutes finding — and that use has no side effects.

Rule of thumb

If the message could go to two hundred other companies after swapping the name, it isn't personalised — it's formatted. Every recipient sees the difference in the first sentence.

CHAPTER 3

An honest balance

Upsides, downsides and limits — including those that apply to tools like ours.

A vendor who cannot name their own limitations simply hasn't found them yet.

3.1 · BALANCE

Upsides and downsides in one table

AREA	THE DOWNSIDE NOBODY MENTIONS	THE UPSIDE YOU CAN MEASURE
Team time	A new tool is a new habit to build	Several hours a week less admin per person
Data quality	The model also records what the customer said imprecisely	Field completeness rises, because it no longer depends on evening discipline
Forecast	Better data won't fix a bad qualification method	A forecast built on quotes, not impressions
Coaching	Requires a manager who actually listens to calls	Patterns across the team instead of a few anecdotes
Onboarding	A recording library needs curating	New hires learn on real conversations from week one
Team trust	Badly communicated recording damages culture for a long time	Well communicated, it takes the paperwork off the team
Cost	**Subscription plus the time to build the habit**	**Comparable to a fraction of the headcount you recover**

The balance comes out positive in one specific scenario: when you automate reproductive work and leave the decisions to people. It comes out negative when you do the opposite — **and that is the whole difference between a rollout that holds up and one that quietly disappears from the budget after a year**.

3.2 · LIMITS

What AI won't do for a rep

A short list, unchanged for two years despite successive generations of models. Every item shares one denominator: it requires taking responsibility, and responsibility cannot be delegated to software.

- × **It won't build trust.** A customer buys from someone who bears the consequences if the project fails. A model bears none.
- × **It won't sense what the customer didn't say.** The most important information in a call is often the hesitation before an answer, or the topic that quietly left the agenda. Only the human in the room hears it.
- × **It won't decide on a concession.** Discount, scope, deadline — those are decisions about margin and about the company, not about the wording of an email.
- × **It won't fix bad qualification.** If the team talks to people without budget, automation lets them do it more efficiently and at greater scale.
- × **It won't replace a meeting with the decision-maker.** No amount of deal data substitutes for one conversation with the person who signs.

Hence the principle: the machine drafts, the human approves

This is not marketing caution but a division of labour matched to what each side does well. The model is incomparably better at faithful recording and at spotting patterns. The human is the only one who can judge a situation and answer for it.

3.3 · RISK

Data, compliance and accountability

A sales call is personal data belonging to at least two people, often confidential customer information, sometimes the details of their budget plans. By recording calls you begin processing material more sensitive than the contents of your CRM.

QUESTION	SETTLE IT BEFORE THE FIRST RECORDING
Consent	Who informs the other party, and when. In the calendar invite and at the start of the call
Data location	Country of the servers, sub-processors, processing agreement. The answer should come immediately
Model training	Whether your calls feed the vendor's model. The only acceptable answer is "no"
Retention	How long recordings are kept and who can delete them
Access	Who in the company can see a rep's recording. Manager, team, everyone?
Exit	In what format you get your data back, and how quickly

One non-technical question is worth settling too: who is accountable for a model error that reached the customer. If the recap sends automatically, the company is accountable and nobody in particular. If the rep approves it, the rep is — **and that is the only arrangement in which the team actually reads what it sends**.

How we handle it at Luminote

Data processed and stored on servers in the European Union, we never train models on your calls, and the processing agreement and documentation are ready before the conversation with your lawyer. Nothing reaches the CRM or the customer without the rep's approval.

CHAPTER 4

Rolling out in 2026

No transformation, no steering committee, no AI strategy. One layer of the stack, a few metrics and a date on which the decision is made. **Anything beyond that is a project, not a rollout.**

4.1 · ARCHITECTURE

Layers of the sales stack

Instead of a list of tools — four layers. The order matters: each one works only if the one below is filled with data you can trust.

LAYER 1

The record of the call

Transcript and a structured note from every meeting. Without this layer, everything above rests on the rep's memory.

LAYER 2

System of record

A CRM fed automatically: stage, contacts, qualification, commitments, next step. The rep approves rather than re-types.

LAYER 3

Action

Follow-up, preparation for the next call, reminders about commitments. Always as a proposal awaiting approval.

LAYER 4

Management decisions

Forecast, coaching, onboarding and loss analysis — based on what was actually said in calls, not on impressions from the review.

Where Luminote sits in this

In layers 1–3, exposing data to layer 4. We listen to the call, prepare the CRM entry, MEDDIC qualification, the forecast update and the follow-up — the rep approves it at a glance. 35 integrations, seven-minute setup, data in the European Union.

4.2 · MEASUREMENT

Metrics that measure effect, not noise

The number of generated summaries and the count of active users are the vendor's metrics, not yours. The four below speak to the result, and all can be read after a month.

METRIC	HOW TO READ IT
Admin hours per person	The primary one. A drop the team doesn't feel means the habit hasn't taken hold
Time from meeting to follow-up	The fastest to react. Change visible in the first week
Qualification field completeness	Share of deals with budget, decision-maker and timeline filled — on quotes, not claims
Forecast accuracy	Slow to react but decisive. Compare forecast against result, quarter over quarter

1
primary metric for the whole pilot

4–6
weeks to a decision

0
new headcount in the plan

What not to measure

Volume of generated content, number of model queries, time spent in the tool. All three rise with team engagement and none says anything about revenue.

TO DO

The 90-day plan

Three months, one team, one decision at the end of each. If a stage gives no result, the next one doesn't start.

DAYS 1-30

Measure, and pick one process

Record the four numbers from page 18 as your baseline. Choose one process to automate — ideally the note and the CRM entry, because the effect shows fastest. Name an owner and a decision date.

DAYS 31-60

One team, every call

State the rule plainly: recordings exist to remove paperwork and to learn together, not to control. Nobody updates the CRM in the evening — if something is missing, that's a configuration error, not a discipline problem.

DAYS 61-90

Compare, decide, expand or close

The same four numbers. If the difference is visible to the team without looking at a report — roll it out further. If not — close the project and come back to it in a year knowing why.

What not to do in those 90 days

Don't roll out three tools in parallel, don't build your own agent, and don't start with customer-facing use cases. The first makes attribution impossible, the second costs a year, the third is the slowest to reverse.

FINALLY

AI won't sell for you. But it can stop getting in the way.

Luminote listens to sales calls and prepares the CRM entry, MEDDIC qualification, the forecast update and the follow-up — the rep approves it all at a glance. **35 integrations, seven-minute setup, data on servers in the European Union — and we never train models on your calls.**

+8 h
per rep, per week

12 days
to return on investment

See the demo — luminote.io →