



AI for sales and business development

Making account research, qualification, proposal work, and follow-up more useful.

WORKS WITH THE ASSISTANTS TEAMS ALREADY USE



EDITION 07 / EMBER & INDIGO

EXECUTIVE WORKING EDITION - 08 OCT 2026

07

AT A GLANCE

Making account research, qualification, proposal work, and follow-up more useful.

II

Sales teams lose a surprising share of their week to work that is not selling: researching accounts, writing up calls, assembling proposals, and cleaning the CRM. AI can take much of that load, but only if the team decides which tasks it owns, which facts it may use, and which messages a human must always approve. This playbook shows where AI fits in a sales motion, how to run it with clear ownership, how to protect buyer trust, and how to measure whether pipeline quality actually improves.

WHO IT IS FOR

Heads of sales, revenue operations leaders, business development managers, and sales enablement leads who need to decide where AI belongs in the sales process and how to run it without damaging buyer trust.

TIME TO USE

60 minutes to read; one working session to apply

S.

What you will have at the end

- A ranked list of sales tasks where AI saves real seller time, with the ones to avoid marked clearly
- A review policy that says which outputs a seller must approve before a buyer sees them
- A baseline and a metrics table to judge whether AI improves pipeline quality and not just activity
- A 90-day rollout plan with named owners, a pilot group, and a stop rule

Contents

01	What AI should and should not do in a sales motion	03
02	Use cases, fit, and where to stay away	04
03	How these tools work, in plain language	06
04	Roles, ownership, and decision rights	08
05	Protecting buyers, data, and the brand	09
06	Baseline, metrics, and test sets	11
07	A practical path from pilot to routine	12
08	Common failures, an illustrative scenario, and vendor questions	14
09	A first quarter that produces evidence	15
10	Terms used in this publication	16




EXECUTIVE BRIEF

What AI should and should not do in a sales motion

01

S.

AI is most valuable in sales when it removes preparation and paperwork so sellers spend more time in real conversations. It is least valuable when it is asked to be the relationship.

A typical seller spends a large part of the week on tasks that sit around the sale: reading a prospect's website and filings, preparing for a discovery call, writing call notes, drafting a proposal from old ones, updating fields in the CRM, and chasing follow-ups. Each task is small, but together they crowd out the work that wins deals. Generative AI is good at reading, summarizing, drafting, and sorting. That is a close match for the preparation and paperwork layer of selling.

It is a poor match for the parts of selling that depend on trust, judgment about a specific person, and accountability for promises. A model can draft a sentence that sounds warm, but it does not know that your champion just lost a budget fight, that a competitor's reference call went badly last month, or that legal will never accept a certain payment term. Treat AI as a fast, tireless junior researcher and drafter, with the seller as the accountable owner of everything a buyer receives.

The four jobs worth starting with

- Account intelligence: a one-page brief on a company, its recent news, its likely priorities, and the people involved
- Discovery preparation: tailored questions and hypotheses before a call, and a structured summary after it
- Proposal and response drafting: first drafts assembled from approved content, reviewed by a seller and a pricing owner
- Pipeline hygiene: draft CRM updates, next-step reminders, and flags for stalled or inconsistent opportunities

THE TEST TO APPLY TO EVERY USE CASE

Ask: if this output were wrong and the seller did not notice, who would be harmed and how quickly would we find out? Research notes that are wrong cost minutes. A wrong price or an invented capability in a proposal can cost a deal or a contract dispute.

WORKING QUESTION

Which two tasks, if a seller never had to do them again, would free the most time for conversations with buyers?

WHERE IT FITS

Use cases, fit, and where to stay away

02

Sort sales tasks by how bad a wrong output is and how easily a human can check it. Start where checking is fast and mistakes are cheap.

Not every sales task benefits equally. The useful way to sort them is by two questions: how quickly can a human verify the result, and how costly is an unnoticed error. Account research is easy to spot-check against the source pages. A pricing table inside a proposal is also checkable, but an error is expensive, so a named owner must approve it. Anything that makes a commitment on the company's behalf needs a person to read and own it.

USE CASE	FIT	MAIN RISK AND CONTROL
Account and contact research briefs	Strong	Outdated or mixed-up facts. Require source links and a seller skim before use

USE CASE	FIT	MAIN RISK AND CONTROL
Discovery call preparation and question lists	Strong	Generic questions. Feed in your qualification framework and past won-deal patterns
Call summaries and CRM field drafts	Strong	Misheard commitments. Seller confirms next steps and amounts before saving
First-draft proposals and RFP answers	Moderate	Invented capabilities or stale terms. Draw only from an approved answer library
Outbound email personalization	Moderate	Creepy or wrong personalization. Cap volume and require seller review
Pricing, discounts, contract terms	Weak	High cost of error. AI may retrieve approved rules but never decide an exception
Fully automated outreach at scale	Avoid	Brand damage, spam rules, and consent issues. Keep a human in the loop

Where AI does not fit

- Judging how serious a buyer is from tone alone, since models read text, not intent
- Scoring individual people on sensitive traits or inferring attributes about them
- Replacing the first human conversation in a high-value or relationship-led account
- Making forecast commitments, which belong to the sales leader who answers for them
- Writing claims about security, compliance, or performance that product or legal has not approved

A PRACTICAL RULE

If the output goes to a buyer, a seller reads it first. If the output sets a price, a term, or a promise, a named approver reads it too. Put this rule in writing before the first pilot.

WORKING QUESTION

Which sales activity in your team today would be risky to automate even if the AI output were usually good?

HOW IT WORKS

How these tools work, in plain language

03



Most sales AI is a language model plus your data plus a set of instructions. Knowing which of the three is failing is the key to fixing it.

A language model predicts useful text from the context it is given. It does not look anything up unless the system around it does so. In sales tools, that surrounding system usually does three things: it pulls data from the CRM, email, call transcripts, and approved content; it passes that data to the model with instructions; and it returns a draft for a person to use or edit. When the output is poor, the cause is usually one of three things: the data was missing or wrong, the instructions were vague, or the task was too much to ask in one step.



01 Gather the facts

The system retrieves the account record, recent activity, public information, and relevant approved content. Better inputs matter more than a cleverer model.

02 Frame the task

A prompt template tells the model the role, the format, the sources it may use, and what to do when it does not know. This is where your qualification method and tone of voice live.

03 Draft

The model produces a brief, summary, email, or proposal section. It should cite which source supports each claim.

04 Review and act

A seller edits, approves, or rejects the draft. Their edits are your best data on where the system is weak.

Why CRM data quality sets the ceiling

AI that summarizes or scores opportunities inherits every flaw of the CRM. If close dates are optimistic, contacts are stale, and stages mean different things to different managers, AI will confidently restate the confusion. Before investing in tools that reason over pipeline data, agree on stage definitions, required fields, and who owns data accuracy. A short cleanup of five or six fields often improves results more than switching vendors.

PLAIN-LANGUAGE WARNING

Models can state wrong facts fluently. This is sometimes called hallucination. In sales it appears as invented customer names, made-up product features, and wrong dates. Require source citations and test for this explicitly.

WORKING QUESTION

If a seller asked your current CRM for the truth about one deal, how many fields would they trust without checking?



OPERATING MODEL

Roles, ownership, and decision rights

04

Sales AI fails less from bad technology than from nobody owning the prompts, the content library, or the quality checks.

Sales organizations are used to owning quota, territories, and forecasts. AI adds new assets that need owners too: the prompt templates, the approved answer library, the data connections, and the review rules. If these belong to everyone, they belong to no one, and quality drifts. Name an accountable person for each, and keep the group small enough to meet weekly during the pilot.

ROLE	OWNS	DECIDES
Chief revenue officer or sales VP	Outcomes and risk appetite	Which use cases run, and the stop rule
Revenue operations	CRM data, integrations, metrics	Field definitions, data access, reporting
Sales enablement	Prompt templates, training, answer library	What good output looks like and how sellers learn it
Product marketing and product	Approved claims and positioning	What may be said about capabilities
Legal and security reviewer	Terms, privacy and security answers	Escalation rules and prohibited content
Frontline sellers and managers	Every message a buyer receives	Whether a draft is good enough to send

Keep a single approved answer library

Proposal and questionnaire work goes wrong when the model draws from old decks, outdated security answers, or one rep's improvisation. Build a library of approved answers with an owner, a last-reviewed date, and an expiry. Point the AI at that library and nothing else for factual claims about the product, security, and terms. When an answer is missing, the right behavior is to flag a gap for a human, not to invent a reply.

- Give each library entry an owner and a review date, and retire entries that nobody will stand behind
- Log which entries are reused most, since these deserve the best wording
- Route gaps to product or legal as a queue, and measure how long they sit
- Version prompt templates the way you version pricing, so you can see what changed when quality moved

MANAGER ROLE CHANGES TOO

Managers coach on judgment, not on formatting. If AI handles notes and drafts, use the time saved in one-to-ones to review real calls and deal strategy, not to read more AI-generated summaries.

WORKING QUESTION

Who in your organization would be asked to fix it tomorrow if an AI-drafted proposal went out with a wrong claim?

RISKS AND CONTROLS

Protecting buyers, data, and the brand

05

The main sales risks are wrong promises, careless use of personal data, and outreach that feels automated and cold. Each has a practical control.

Selling involves personal data about buyers, confidential information about their plans, and your own commercial terms. Sending any of it to a tool without checking where it goes is the fastest route to a problem. Beyond data, there are accuracy risks and trust risks. A buyer who receives an email that is clearly mass-produced, or that cites an irrelevant detail from their LinkedIn profile, loses a little respect for your company, and sellers rarely see that loss in their metrics.

RISK	WHAT IT LOOKS LIKE	CONTROL
Wrong claims	Invented features, customers, or certifications in a proposal	Approved answer library, citations, product review of new claims
Data exposure	Call transcripts or customer files pasted into unapproved tools	Approved tools only, clear rules on what data may enter, vendor terms checked
Privacy and consent	Scraping or profiling individuals, outreach without a lawful basis	Check data protection and marketing rules with counsel; limit personal data use
Voice recordings	Recording or transcribing calls without proper notice	Follow local recording and consent requirements; announce clearly
Brand erosion	Generic, over-familiar, or inaccurate messages	Seller approval, volume caps, tone guide, and spot audits
Bias in scoring	Lead scores that favor some groups for irrelevant reasons	Use business signals, review score drivers, avoid sensitive attributes

A short human-review policy

Before any AI-drafted message reaches a buyer

- A named seller has read the whole message, not skimmed the first line
- Every factual claim about the buyer can be traced to a source
- Every claim about our product comes from the approved library
- Prices, discounts, dates, and commitments have been checked by the accountable owner
- The message would still be acceptable if the buyer forwarded it to a colleague or competitor

Regulations on marketing communications, privacy, and recording differ by country and change over time. Treat this section as a prompt to ask your counsel, not as legal advice. If you operate across regions, check current requirements for each, and note that frameworks such as the NIST AI Risk Management Framework and ISO/IEC 42001 can help structure your internal controls without being sales-specific.

WORKING QUESTION

What is the one piece of buyer or deal data your sellers should never paste into an unapproved tool?

MEASURING VALUE

Baseline, metrics, and test sets

06

Activity numbers will rise whenever AI is introduced. The question is whether quality, speed, and win rates move, and you cannot know without a baseline.

The most common measurement error in sales AI is counting more emails sent or more meetings booked as proof of success. AI makes volume cheap, so volume will rise. What matters is whether the right accounts are reached, whether conversations are better, and whether deals move faster without a drop in win rate or deal quality. Measure before you start, using two to three months of historical data or a matched control group of sellers who do not use the tool.

METRIC	HOW TO MEASURE	WATCH FOR
Seller time on prep and admin	Time study or survey before and after, same sellers	Self-reported savings that never become selling time
Stage conversion rates	CRM stage-to-stage conversion, pilot vs control	Gaming through looser stage definitions
Sales cycle length	Median days by segment and deal size	Mix changes that hide the effect size
Proposal turnaround and edit rate	Hours to first draft; share of draft text kept after review	High edit rate signals weak inputs
Win rate and average deal value	Closed-won over closed deals, by segment	Short pilots that are too small to tell
Buyer experience	Reply quality, complaints, opt-out/unsubscribe rate	Rising opt-outs after personalization campaigns
Data quality	Share of opportunities with complete, current key fields	Improvement that sellers feel is extra work

Build a small test set before you launch

For each use case, collect 25 to 50 real examples with known good answers: accounts where you know the facts, past calls with agreed summaries, and old RFP questions with approved replies. Run the tool on them, and have two experienced people grade the output against a simple rubric: accurate, complete, useful, and safe to send. Set a pass mark before you look at results, such as no factual errors in any brief and at least four in five drafts acceptable with light edits. These figures are illustrative, so choose yours with your team. Rerun the set whenever you change the model, the prompt, or the data source.

ILLUSTRATIVE EXAMPLE

A pilot group saves roughly an hour per seller per week on call notes, but their stage conversion drops. Investigation shows summaries were auto-advancing deals on vague buyer language. The time saved was real, the quality cost was too, and only both metrics together revealed it.

WORKING QUESTION

What would you need to see in the data after 90 days to say with confidence that AI improved selling and not just activity?

IMPLEMENTATION GUIDE

A practical path from pilot to routine

07

Start with one team, two use cases, and a written review rule. Expand only when the numbers and the sellers agree it is working.

01 Pick the pilot

Choose one segment with a motivated manager, a clean enough CRM, and enough deal flow to measure. Eight to fifteen sellers is usually enough to learn from.

02 Choose two use cases

Combine one research task and one drafting task, for example account briefs and call summaries. Avoid launching pricing or outbound automation first.

03 Prepare the inputs

Fix key CRM fields, assemble the approved answer library, and write the tone guide. Decide what data may and may not be used.



04 Train on real work

Run sessions using the sellers' own accounts. Teach them to give context, check sources, and edit rather than accept. Short habits beat long courses.

05 Run weekly reviews

Review a sample of outputs, edits, and complaints each week. Fix prompts, library gaps, and data problems in the same cycle.

06 Decide with evidence

At the end of the pilot, compare against the baseline and choose: scale, adjust, or stop. Write the decision down with the reasoning.

**Launch readiness**

- Use cases and non-use cases are written down and shared with sellers
- Approved tools list and data rules are published
- Answer library has owners and review dates
- Baseline metrics are captured before go-live
- A named person can pause the tool the same day if something goes wrong
- Sellers know where to report a bad output and how fast it will be fixed

**Sequence for expansion**

- Extend to a second segment only after weekly review findings fall to a manageable level
- Add outbound personalization last, with volume caps and opt-out monitoring
- Integrate into the CRM workflow so sellers do not have to switch screens
- Keep a human accountable for every external message, even when volumes grow
- Revisit the tool list quarterly, since capabilities and vendor terms change



WORKING QUESTION

Which segment has the cleanest data, the most deal volume, and a manager who will run the pilot properly?

PITFALLS AND BUYING

Common failures, an illustrative scenario, and vendor questions

08

Most sales AI disappointments come from rushing outbound volume, ignoring data quality, or buying a tool before defining the job.

- Starting with mass outbound because it is easy to measure, then damaging sender reputation and brand
- Buying a platform before agreeing what a good account brief or call summary looks like
- Letting each seller keep private prompts, so quality cannot be reviewed or improved
- Ignoring the review burden, so managers spend the time saved reading summaries
- Measuring seller satisfaction only, without pipeline quality or buyer response

ILLUSTRATIVE SCENARIO, NOT A REAL COMPANY

A regional software firm turns on an email generator for its sales development team. Meetings booked rise quickly, but qualified opportunities do not, and several prospects complain about references to outdated news. The firm cuts volume, adds a source check and a seller review, and rewrites the brief around qualification criteria. Meetings fall, qualified pipeline recovers.

Questions to ask any vendor

Vendor diligence

- Where is our data stored and processed, and is it used to train models for other customers
- Which models and sub-processors are involved, and what happens when they change
- Can every generated claim link to a source, and can we restrict the sources
- How are call recordings, transcripts, and personal data retained and deleted
- What admin controls, audit logs, and role-based permissions exist
- How does the tool connect to our CRM, and what happens to our data if we leave
- What security attestations can you provide, and when were they last issued



Ask vendors to run your test set on your data in a trial, not their demo data. Demos are built to succeed. Your own accounts, messy CRM, and awkward RFP questions will show the real quality, and they will show how well the vendor handles the cases where the right answer is to say it does not know.

WORKING QUESTION

Which of these failures is your team most likely to repeat, and what early signal would tell you?



30 / 60 / 90

A first quarter that produces evidence

09

K

Move from one measured pilot to a routine, reviewed way of working, and expand only when the evidence supports it.

First 30 days: ground it

- Pick the pilot team and two use cases, and write the do and do-not list
- Capture baseline metrics for prep time, conversion, and cycle length
- Fix key CRM fields and name an owner for the answer library
- Publish the approved tools list and the human-review rule



Days 31 to 60: test it

- Run the pilot with weekly output reviews and a graded test set
- Fix prompt, library, and data issues found in review
- Train managers to coach on the use of drafts and summaries
- Compare pilot and control on early indicators

Days 61 to 90: decide

- Review results against the baseline and decide to scale, adjust, or stop
- Document the operating model, owners, and quarterly review
- Plan the second segment and any outbound personalization with caps
- Report buyer-experience signals to leadership alongside efficiency gains

GLOSSARY

Terms used in this publication

10

TERM	MEANING
Account intelligence	A brief that pulls together public and internal information about a target company and its people to prepare for outreach or a meeting.
Agent assist	A tool that suggests information or drafts to a human while they work, leaving the human in control of the final action.
Answer library	A governed collection of approved replies to common questions about product, security, and terms, each with an owner and review date.
CRM hygiene	The accuracy and completeness of customer relationship management records such as stages, close dates, and contacts.
Discovery	The early sales conversation used to understand the buyer's problem, priorities, and buying process.
Hallucination	A fluent but unsupported or false statement produced by a language model.
Human in the loop	A design in which a person reviews and approves AI output before it has an effect.

TERM	MEANING
Prompt template	A reusable set of instructions, context slots, and format rules given to a model for a recurring task.
Qualification	The process of judging whether an opportunity fits your offering, has a real need and budget, and is worth pursuing.
Sales cycle	The time from first meaningful contact with a prospect to a closed deal.
Test set	A fixed group of real examples with known good answers, used to compare tool quality over time.
Stage conversion	The share of opportunities that move from one pipeline stage to the next.

A note on use

Free assessments are available at halden.company/assessments, and you can reach the team at halden.company/contact. Figures in this playbook are illustrative unless stated, and nothing here is legal advice.

TAKE THE NEXT STEP

Take the free Halden assessment at halden.company/assessments, or talk to us at halden.company/contact about turning this into a plan for your organization.