

Sales & Marketing · Sales

Sales Opportunity Agent

CAI Score L2 · Draft — Human reviews & approves all output

▲ this agent

L1 Assist	L2 Draft	L3 Operate	L4 Decide	L5 Autonomous
TIER	PILLAR	AUTOMATION	HUMAN ROLE	
L2 · Draft	Sales & Marketing	≈45% of routine effort	Reviews and approves all output	

01 — Executive summary

Sales Opportunity Agent is a CAI Score L2 (Draft) agent for sales & marketing. Researches live opportunities to give sellers and managers visibility of deal velocity and risk, using AI research across available sources. It operates with a human review gate: all output is approved before use. Indicatively, it can redirect around 45% of the routine effort this work consumes.

02 — Problem & context

- Deal risk only becomes visible when a forecast misses.
- Reps research accounts manually, so the pipeline view is always behind.

03 — Representative use cases

- Researches each live opportunity continuously, surfacing velocity and risk signals for the seller and manager.
- Replaces manual account research: the pipeline view is current, not built on information from the last call.
- Flags stalled deals and at-risk opportunities before they appear as forecast misses.
- Provides deal-level intelligence summaries for pipeline review meetings without manual preparation.
- Produces the evidence to track pipeline accuracy, deal-risk identification rate, and forecast variance.

04 — How to use it

- 1 Connect it in your tenant to the systems this work already touches (sales).
- 2 Confirm scope, the named reviewer, and the approval the tier requires.
- 3 Receives the task trigger or request
- 4 Produces the draft or analysis output
- 5 Human reviews and approves before use
- 6 Logs every output for audit
- 7 Review quality and acceptance rate after each cycle; tune scope as you learn.

05 — Where it fits in the process


Sits across the active pipeline — surfacing velocity and risk signals on each opportunity to sharpen forecasting decisions.



Drafting loop for a L2 agent — human approves all output.

06 — Risk assessment

#	RISK	LIKELIHOOD	IMPACT	MITIGATION / CONTROL
1	Incorrect or incomplete draft	Low	Medium	Draft is reviewed and approved by a human before any use; errors are caught at the review gate.
2	Over-reliance / reduced diligence	Low	Medium	Tier design keeps a human in the approval seat; acceptance-rate KPIs are tracked each cycle.
3	Exposure of sensitive data	Medium	Medium	Runs inside your own tenant; Colleague AI processes none of your data. Access stays under your existing controls.
4	Scope creep or misuse	Low	Low	Scope is bounded and documented; the agent produces output only — it takes no autonomous action.
5	Quality drift over time	Low	Low	Periodic sampling and KPI review detects drift; prompts are version-controlled and change-managed.



Residual likelihood × impact after the tier's built-in controls. Numbers map to the table above.

07 — Regulatory & compliance impact

FRAMEWORK	RELEVANCE	HOW THIS AGENT ALIGNS
EU AI Act	Operational use under human oversight; logging and risk-tiering apply.	Documented risk tier, human oversight and event logging per agent — designed to map to the obligations.
DORA	Relevant where the function is in scope for financial-sector operational resilience.	Actions are logged and attributable, supporting resilience and traceability requirements — designed to map to.
GDPR	Applies wherever personal data is in scope of the workflow.	Data stays in your tenant; Colleague AI processes none of it. Purpose limitation and minimisation are preserved.
ISO/IEC 42001	AI management-system controls: roles, monitoring, review.	The agent operates within a documented management system — defined tier, named roles, logging and review. Designed to meet; not certified.

“Designed to map to” / “designed to meet” describe alignment with the CAI Score framework, not external certification. Confirm obligations with your compliance function.

08 — Value & illustrative ROI

Researches live opportunities to give sellers and managers visibility of deal velocity and risk, using AI research across available sources.

Assumptions (illustrative — replace with your own)	Result
3 FTE engaged · 1840 productive h/FTE/yr · 45% addressable · €60/h blended	≈ 2,484 hours/yr redirected · ≈ €149,000/yr indicative value



Figures are an illustrative model to be validated against your own volumes and rates – not a guarantee.

KPIs to track

Improved decision quality · better deal-risk visibility · more accurate forecasting.

09 — Recommendations

- Connect the agent to all data sources used in current account research.
- Sellers and managers make all forecasting and deal calls — the agent informs, it does not decide.
- Review signal accuracy against actual deal outcomes after the first quarter.
- Baseline pipeline forecast accuracy and deal-research effort before go-live.

10 — Governance & method

This agent runs in your own Microsoft/Azure tenant; Colleague AI hosts only the control plane (score, policy, audit) and processes none of your data. A person reviews and approves all output before it is used. Every action is time-stamped and attributable. The CAI Score classification is documented and open to challenge.

This dossier is an evaluation document. ROI and automation figures are illustrative models, not commitments. Compliance statements describe the CAI Score framework and are not a certification of your compliance with any law or standard. © Colleague AI — proprietary. CAI Score™.