

Operations & Service Delivery · Operations · Procedures

Operations Procedures Assistant

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	Operations & Service Delivery	≈45% of routine effort	Reviews and approves all output

01 — Executive summary

Operations Procedures Assistant is a CAI Score L2 (Draft) agent for operations & service delivery. A conversational assistant that lets staff find and follow internal operational procedures on demand instead of searching documents. 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

- Staff ping a senior colleague every time they hit a how-do-I moment.
- Procedures live in documents no one can find when they actually need them.

03 — Representative use cases

- Answers how-do-I questions instantly from the approved procedure set — no document search required.
- Reduces dependence on tribal knowledge: the right step is surfaced for every member of the team, not just experienced staff.
- Supports new-joiner onboarding by providing on-demand procedure guidance from day one.
- Covers any operational procedure domain — finance, operations, IT, HR — from a single interface.
- Produces an interaction log that shows which procedures are queried most, driving targeted documentation improvement.

04 — How to use it

- 1 Connect it in your tenant to the systems this work already touches (procedures).
- 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


Lives wherever staff get stuck mid-task — answering how-do-I questions from the approved procedure set, so the right step is one question away.



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

A conversational assistant that lets staff find and follow internal operational procedures on demand instead of searching documents.

Assumptions (illustrative — replace with your own)	Result
6 FTE engaged · 1840 productive h/FTE/yr · 45% addressable · €50/h blended	≈ 4,968 hours/yr redirected · ≈ €248,000/yr indicative value



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

KPIs to track

Faster task completion · fewer how-do-I escalations · reduced senior-staff interruption.

09 — Recommendations

- Load all approved procedure documents into the knowledge base before launch; keep them current.
- Assign a procedure owner per domain to review the interaction log and flag gaps monthly.
- Brief staff clearly: the agent answers from approved documents only — it does not advise on undocumented exceptions.
- Review acceptance and escalation rates after the first four weeks; update content where the agent falls short.

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™.