

Data & Infrastructure · Infrastructure

Maintenance Event 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	Data & Infrastructure	≈50% of routine effort	Reviews and approves all output

01 — Executive summary

Maintenance Event Agent is a CAI Score L2 (Draft) agent for data & infrastructure. Auto-populates maintenance events so planned work is logged consistently and is visible to the right teams. It operates with a human review gate: all output is approved before use. Indicatively, it can redirect around 50% of the routine effort this work consumes.

02 — Problem & context

- Planned work gets logged inconsistently, so teams are blindsided by it.
- Maintenance records are filled in late, if at all.

03 — Representative use cases

- Auto-populates maintenance event records from the planned-work description, enforcing all required fields.
- Ensures every maintenance event is logged before work begins, not reconstructed after.
- Notifies the correct teams automatically when a maintenance event is logged, eliminating visibility gaps.
- Enforces consistent record format across all teams and time zones, improving auditability.
- Produces the evidence to track logging completeness, notification accuracy, and maintenance-event visibility.

04 — How to use it

- 1 Connect it in your tenant to the systems this work already touches (infrastructure).
- 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 at the planning step — it fills out maintenance records consistently, improving visibility of planned work across teams.



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

Auto-populates maintenance events so planned work is logged consistently and is visible to the right teams.

Assumptions (illustrative — replace with your own)	Result
3 FTE engaged · 1840 productive h/FTE/yr · 50% addressable · €55/h blended	≈ 2,760 hours/yr redirected · ≈ €152,000/yr indicative value



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

KPIs to track

Improved reliability and awareness of maintenance activity.

09 — Recommendations

- Connect the agent to the maintenance management system before deployment.
- Define required fields and notification routing rules before go-live.
- The owning team confirms each record before it is committed — this is an L2 agent.
- Review logging completeness after the first month; adjust field requirements as needed.

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™.