Risk, Security & Compliance · Risk

Risk Event Capture Agent

CAI Score L3 · Operate — Bounded autonomy

▲ this agent

L1 Assist	L2 Draft	**L3 Operate**	L4 Decide	L5 Autonomous

TIER	PILLAR	AUTOMATION	HUMAN ROLE
L3 · Operate	**Risk, Security & Compliance**	**≈60% of routine effort**	**Owns exceptions and oversight**

01 — Executive summary

Risk Event Capture Agent is a CAI Score L3 (Operate) agent for risk, security & compliance. Walks users through creating and logging risk events in the GRC tooling, enforcing consistent classification and complete required fields. It operates with bounded autonomy: a person owns exceptions and retains workflow oversight. Indicatively, it can redirect around 60% of the routine effort this work consumes.

02 — Problem & context

— Risk events go under-reported because logging them is slow and fiddly.

— Half the events are missing fields, so the risk picture is never complete.

03 — Representative use cases

— Guides users through risk event logging step by step, enforcing all required fields before submission.

— Increases reporting completeness: events that previously went unlogged due to friction are now captured.

— Enforces consistent risk classification across all teams, improving the reliability of the risk register.

— Reduces the time to log a risk event from minutes of frustration to a guided two-minute process.

— Produces the evidence to track event-capture completeness, classification consistency, and reporting frequency.

04 — How to use it

1. Connect it in your tenant to the systems this work already touches (risk).
2. Confirm scope, the named owner, and the oversight the tier requires.
3. Detects the trigger in the workflow
4. Handles routine in-policy cases end to end
5. Escalates exceptions to the owner
6. Logs every action for audit
7. Review the audit log and KPI movement after each cycle; tune scope as you learn.

05 — Where it fits in the process

Sits at the point of capture — the moment someone needs to log a risk event — guiding them so events are recorded consistently and nothing gets under-reported.

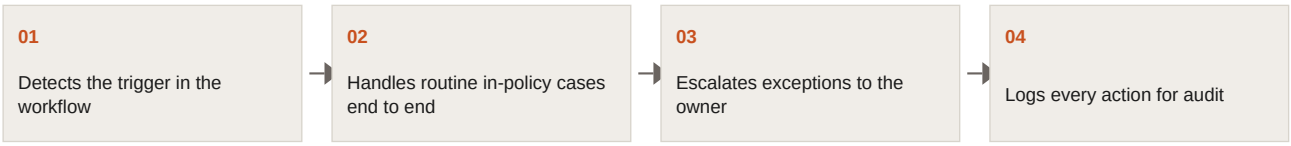

Operating loop for a L3 agent — bounded autonomy.

06 — Risk assessment

#	RISK	LIKELIHOOD	IMPACT	MITIGATION / CONTROL
1	**Incorrect or incomplete output**	Medium	Medium	Routine cases only; exceptions and anything material are escalated to the named owner, with a full audit trail.
2	**Over-reliance / reduced human scrutiny**	Medium	Medium	Tier oversight keeps a human in the loop; KPIs monitor quality so reliance never outruns control.
3	**Exposure of sensitive data**	Medium	Medium	Runs inside your own tenant; Colleague AI processes none of your data, and access stays under your existing controls.
4	**Behavioural drift over time**	Medium	Low	Periodic review against KPIs and sampling detects drift; scope and prompts are version-controlled.
5	**Out-of-policy or edge-case handling**	Medium	Medium	Scope is bounded and documented; anything outside it is escalated, not improvised.

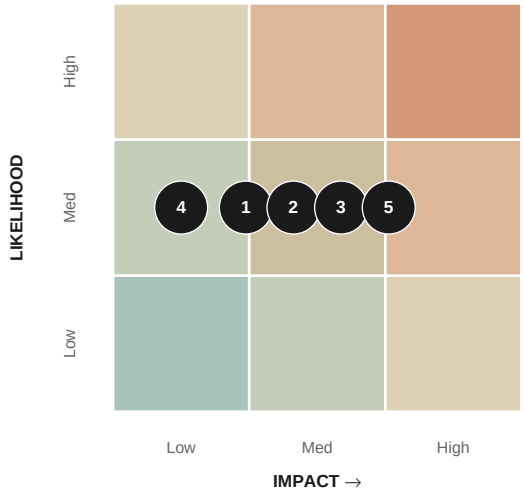

Residual likelihood × impact after the tier's built-in controls. Numbers map to the table above.

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

Walks users through creating and logging risk events in the GRC tooling, enforcing consistent classification and complete required fields.

Assumptions (illustrative — replace with your own)	Result
3 FTE engaged · 1840 productive h/FTE/yr · 60% addressable · €60/h blended	≈ 3,312 hours/yr redirected · ≈ €199,000/yr indicative value

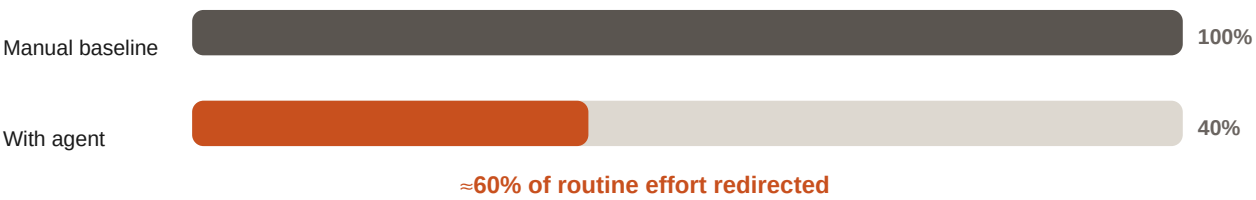

Figures are an illustrative model to be validated against your own volumes and rates — not a guarantee.

KPIs to track

Time saved on event logging · improved reporting completeness · consistent classification.

09 — Recommendations

- Load the risk taxonomy and required-field rules into the agent configuration before launch.
- Brief all risk-event reporters on the agent; position it as removing friction, not adding process.
- Review the escalation log weekly for the first month to identify classification gaps.
- Measure reporting completeness before and after go-live; use the delta as the primary KPI.

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 owns exceptions and retains workflow oversight. Every action is time-stamped and attributable. The CAI Score classification is documented and open to challenge.