Operations & Service Delivery · Operations · Controls

Four-Eyes Control Assistant

CAI Score L4 · Decide — Senior human accountable for every decision

▲ this agent

L1 Assist	L2 Draft	L3 Operate	L4 Decide	L5 Autonomous

TIER	PILLAR	AUTOMATION	HUMAN ROLE
L4 · Decide	**Operations & Service Delivery**	**≈40% of routine effort**	**Named accountable decision-maker**

01 — Executive summary

Four-Eyes Control Assistant is a CAI Score L4 (Decide (supervised)) agent for operations & service delivery. Supports the second reviewer in a four-eyes process — checks the work, highlights discrepancies, and captures the validation evidence. It operates with senior human accountability: a named person owns every decision it supports. Indicatively, it can redirect around 40% of the routine effort this work consumes.

02 — Problem & context

— Dual-control sign-off stalls waiting for a second reviewer to find time.

— Reviewers eyeball everything manually, so discrepancies slip through under deadline.

03 — Representative use cases

— Performs the structured first-pass review in any dual-control chain — finance, compliance, legal, or operations.

— Highlights discrepancies between the work product and source data so the human reviewer sees a diff, not a blank document.

— Accelerates month-end close dual-control: second-review of adjusted entries before journal posting.

— Supports payment and settlement authorisation: structured check of files against thresholds before release.

— Produces a structured evidence pack for every review, satisfying SOX 404 and equivalent ICFR documentation requirements.

04 — How to use it

1. Connect it in your tenant to the systems this work already touches (controls).
2. Obtain senior sign-off and confirm the named accountable person before go-live.
3. Receives the case or work product
4. Assembles evidence and flags risk or discrepancy
5. Presents structured findings to the named human
6. Named human makes the decision; rationale is logged
7. Review decision outcomes and the audit trail after each cycle; adjust scope as needed.

05 — Where it fits in the process

Slots into the dual-control step of any sign-off chain. It does the first-pass check; the human reviewer confirms and remains the named approver.

01	→	02	→	03	→	04
Receives the case or work product		Assembles evidence and flags risk		Presents findings to named human		Human decides — decision logged

`Decision-support loop for a L4 agent — human owns every call.`

06 — Risk assessment

#	RISK	LIKELIHOOD	IMPACT	MITIGATION / CONTROL
1	**Incorrect evidence assembly**	Medium	High	Named human reviews all agent output before any decision is made; the human owns the conclusion.
2	**Over-reliance on agent assessment**	Low	High	L4 tier structure: agent evidences, human decides. Decision authority cannot be delegated to the agent.
3	**Exposure of sensitive data**	Medium	Medium	Runs inside your own tenant; Colleague AI processes none of your data.
4	**Behavioural drift over time**	Medium	Medium	Periodic review against case outcomes detects drift; scope and prompts are version-controlled.
5	**Out-of-policy or edge-case handling**	Low	High	Hard scope boundary: the agent flags anything outside policy; senior sign-off is required at go-live.

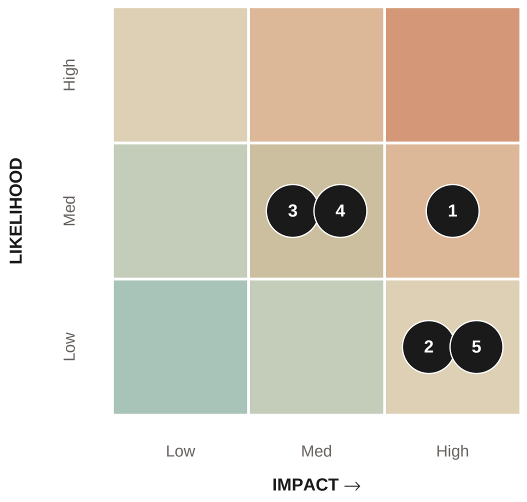

Residual likelihood × impact after the tier's built-in controls. Numbers map to the table above.

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

Supports the second reviewer in a four-eyes process — checks the work, highlights discrepancies, and captures the validation evidence.

Assumptions (illustrative — replace with your own)	Result
3 FTE engaged · 1840 productive h/FTE/yr · 40% addressable · €65/h blended	≈ 2,208 hours/yr redirected · ≈ €144,000/yr indicative value

≈40% of routine effort redirected

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

KPIs to track

Faster control cycles · improved auditability · stronger discrepancy detection.

09 — Recommendations

— Require documented senior sign-off (CFO / CRO / General Counsel) before go-live — mandatory for L4.

— Define the named approver role explicitly; the agent will not operate without it.

— Start with a low-volume, well-documented process before extending to higher-stakes chains.

— Engage internal audit at configuration to confirm discrepancy rules meet their evidence standards.

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 named senior person owns every decision the agent supports. Every action is time-stamped and attributable. The CAI Score classification is documented and open to challenge.