

Data & Infrastructure · Infrastructure

# Infrastructure Wiki Agent

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

▲ this agent

| L1<br>Assist | L2<br>Draft           | L3<br>Operate          | L4<br>Decide                    | L5<br>Autonomous |
|--------------|-----------------------|------------------------|---------------------------------|------------------|
| TIER         | PILLAR                | AUTOMATION             | HUMAN ROLE                      |                  |
| L2 · Draft   | Data & Infrastructure | ≈45% of routine effort | Reviews and approves all output |                  |

## 01 — Executive summary

Infrastructure Wiki Agent is a CAI Score L2 (Draft) agent for data & infrastructure. A dynamic knowledge agent that answers infrastructure questions and guides issue resolution from the infrastructure knowledge base. 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

- Engineers wait on the one person who knows how that system actually works.
- Resolution steps are buried in docs no one searches mid-incident.

## 03 — Representative use cases

- Answers infrastructure questions instantly from the current knowledge base — no wait for the one person who knows.
- Guides engineers through resolution steps during incidents, reducing mean time to resolve.
- Reduces knowledge lock-in: critical system knowledge is accessible to the whole team, not held by individuals.
- Supports new engineers in ramping up on complex infrastructure without needing a dedicated mentor.
- Surfaces knowledge gaps: unanswered queries highlight documentation that needs to be created or updated.

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 beside engineers during incidents and BAU — answering from current docs and guiding resolution, so knowledge is not locked in a few heads.



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 dynamic knowledge agent that answers infrastructure questions and guides issue resolution from the infrastructure knowledge base.

| Assumptions (illustrative — replace with your own)                         | Result                                                       |
|----------------------------------------------------------------------------|--------------------------------------------------------------|
| 4 FTE engaged · 1840 productive h/FTE/yr · 45% addressable · €65/h blended | ≈ 3,312 hours/yr redirected · ≈ €215,000/yr indicative value |



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

**KPIs to track**

Faster infrastructure issue resolution · reduced knowledge lock-in.

## 09 — Recommendations

- Load the current infrastructure knowledge base into the agent before launch; keep it current.
- Assign a knowledge-base owner to review unanswered-query logs monthly and fill gaps.
- Brief engineers clearly: the agent answers from documentation — if the doc is wrong, flag it.
- Review answer accuracy in the first two weeks; calibrate the knowledge base 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™.