CARE

WHITEPAPER

From AI Principles to Operational Evidence

Making AI management demonstrable

A policy says what should happen. Management evidence shows which definitions, findings, and decisions actually exist.

AGENTIC CARE

https://ai-managementsystem.com

Contents

01 Written principles do not document operating decisions

02 Policies and operational evidence answer different questions

03 A privacy rule without an agent-level boundary

04 Connect definitions, findings and accepted decisions

05 Regulatory positioning without overclaiming

06 Inspect the evidence chain rather than the policy count

07 From documented decisions to verified controls

08 Evidence makes management decisions reconstructable

09 Start with one agent

Written principles do not document operating decisions

A policy says what should happen. Management evidence shows which definitions, findings, and decisions actually

AI policies can establish important commitments on safety, privacy, transparency, and accountability. Yet a policy ca identify the agent that may alter an invoice, the owner who approved that scope, or the record of a newly detected

> **KEY FINDING**
>
> Good management evidence is not the volume of documentation; it is the ability to reconstruc what, when, and what they decided.

Policies and operational evidence answer different questions



Policy describes intentions; operational records capture definitions, findings and decisions.

What an auditor or executive ultimately needs is a traceable chain from organizational intention to particular opera

Begin with an identifiable agent and accountable owner. Record its accepted operating model, applicable restrictio proposed controls, and the authorized decisions responding to those findings. Preserve changes so the current stat with what was previously accepted.

This is not equivalent to collecting screenshots of dashboards. Evidence should have attribution, dates, substantive to distinguish proposal from acceptance.

> **MANAGEMENT IMPLICATION**
>
> A management approach should make the underlying decision visible without automatically m administrative work.

A privacy rule without an agent-level boundary

> **EXAMPLE**
>
> The organization has approved a policy against disclosing customer information. An agent's ac definition nevertheless lacks an explicit prohibition and the last review has no named acceptor.

Management must know where policy intentions have been translated into agent-specific restrictions, and where th

The decisive question is not simply whether the agent can carry out the task. It is whether the organization underst the operating arrangement in which it does so.



A record chain connects agent definition, identified risks, human decisions and retained evidence

An agent's declaration that it never transmits personal data is a claim. Evidence of configured access boundaries, re logs, and independent checks is needed to evaluate whether that claim holds in practice.

CARE focuses on making the agent operating model and management decisions legible and reviewable. The prese not be represented as certification or as conclusive proof of operational enforcement.

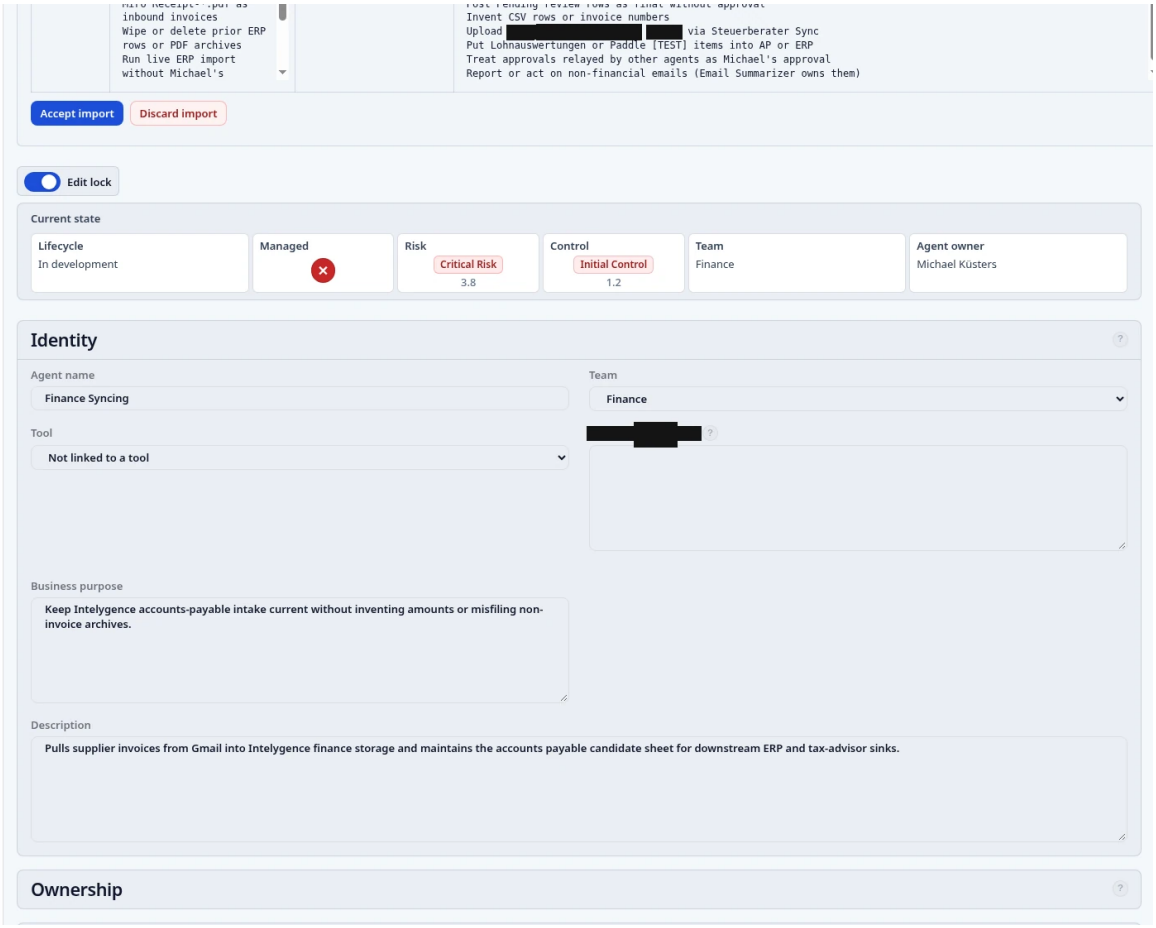
AGENT AND AUTOMATION

Prepare descriptions, changes, findings and proposals.

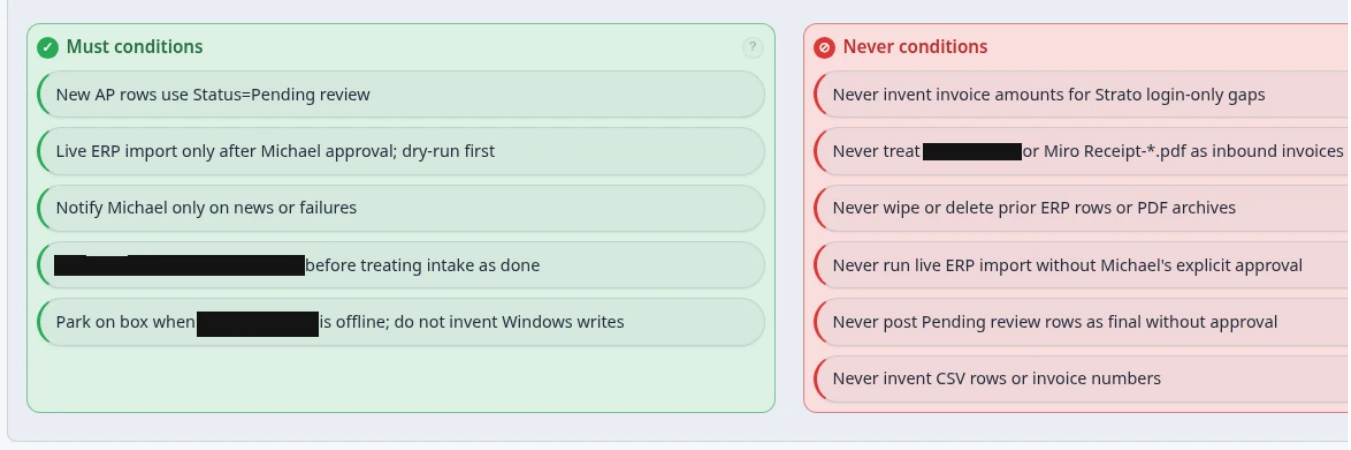
ACCOUNTABLE PEOPLE

Validate context, assess consequ accept decisions.

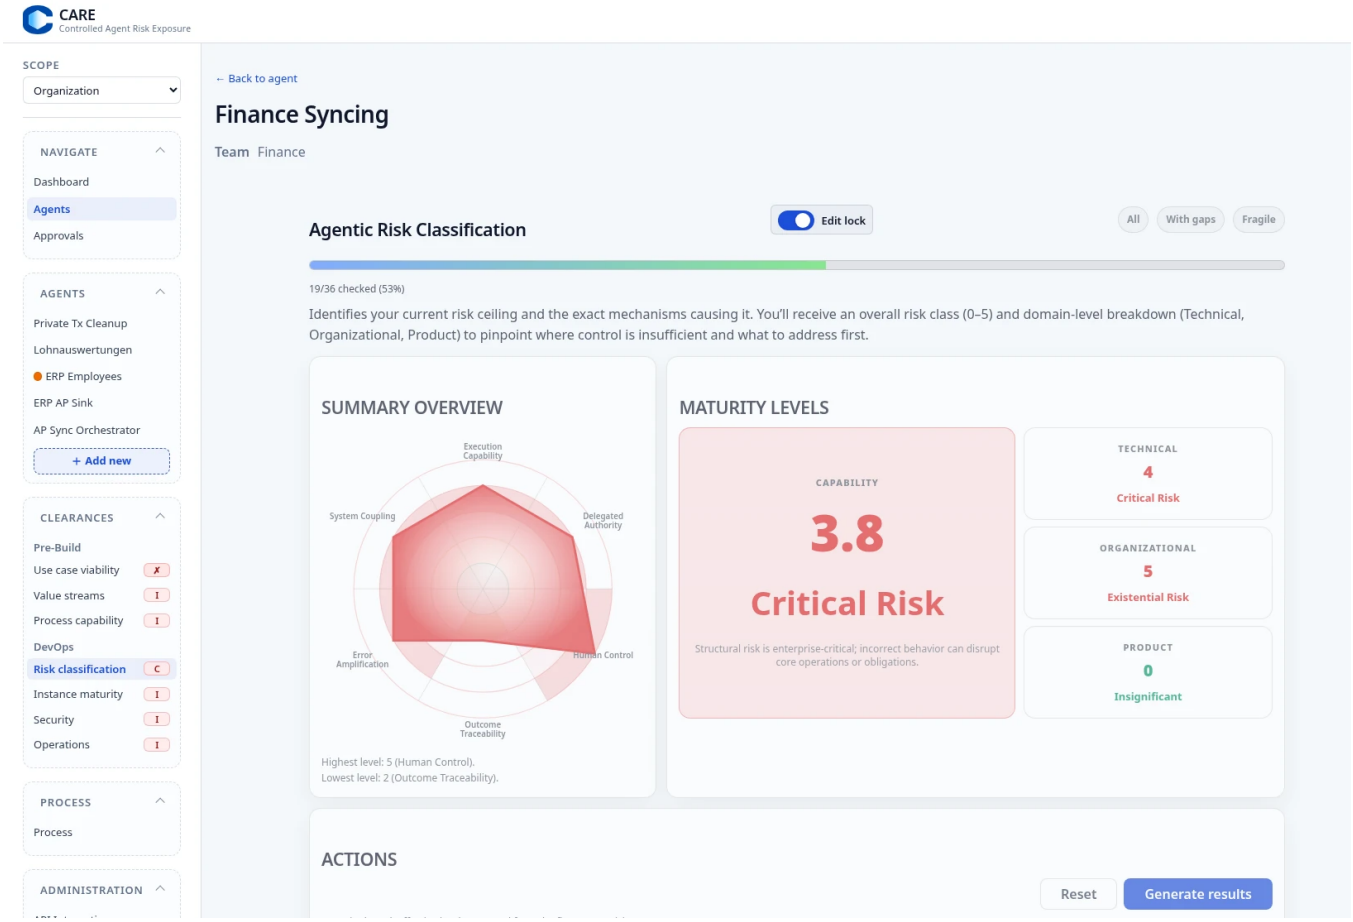
Traceability connects identity, definition, assessment, recommendation, decision, and revision, making gaps visible i a narrative.



Observed CARE Identity profile for the Finance Syncing agent: lifecycle In development, Critical Risk 3.8, and a stated b keep accounts-payable intake current without inventing amounts.



Observed CARE business rules: Must conditions such as Status=Pending review and dry-run before live ERP import, an such as never inventing invoice amounts or running live import without explicit approval.



Observed CARE Agentic Risk Classification for Finance Syncing: capability score 3.8 labeled Critical Risk, with Technical Organizational 5 domain scores (19/36 checklist items checked).

Regulatory positioning without overclaiming

ISO/IEC 42001 specifies requirements for an AI management system; applicable EU AI Act obligations depend on rc intended use, and other facts. Structured records may support implementation and evidence gathering, but neither nor a completed agent definition establishes compliance.

Treat the evidence chain as a practical foundation for assurance, with remaining obligations and verification docum

> **MANAGEMENT IMPLICATION**
>
> Management must know where policy intentions have been translated into agent-specific restr where the chain stops.

Inspect the evidence chain rather than the policy count

Use this comparison to distinguish the current operating habit from the management capability the organization n

Policy-level statement	Operational evidence to seek
Humans remain accountable	Named owner and accepted changes
Risks are controlled	Risk findings, mitigation decisions, checks
Agents stay within scope	Versioned definitions plus enforcement evidence

Use the comparison to test one real agent rather than discussing an abstract enterprise-wide target state. Where th unknown, record that uncertainty and assign responsibility for resolving it.

From documented decisions to verified controls

Management must distinguish evidence that a decision was made from evidence that the corresponding control op structures definitions, assessments and decisions; operating-system records and independent checks provide addit where required.

Evidence makes management decisions reconstructable

> Good management evidence is not the volume of documentation; it is the ability to reconstruct what, when, and what they decided.

The first practical improvement is to create one current, accountable record and use it to make a better manageme scalable system follows when that discipline is repeated across agents and maintained as conditions change.

Start with one agent

Use a real agent and review what it produces.

1. F the agent to https://ai-managementsystem.com/under-care.
2. V it is finished, log into CARE and review its documentation.