CARE

WHITEPAPER

Industrial CARE: Governing Agents as Assets

From isolated AI experiments to an accountable, scalable operational workforce.

AI agents are becoming an operational workforce. CARE provides the missing management layer that connects each agent to its operating model, risks, and the humans accountable for accepting how it is supposed to work.

AGENTIC CARE

https://ai-managementsystem.com

Contents

01 The challenge

02 The problem

1 Why conventional approaches struggle to scale

03 The solution

04 How CARE makes industrial agent management practical

1 Establish

2 Control

3 Optimize

05 What changes for the organization

06 From governance overhead to operational capability

07 Start with one agent

The challenge: AI agents are becoming operational assets



A common agent registry connects work across teams and platforms without replacing their execution envi

Companies already know how to manage employees, applications, infrastructure, and business processes. They hav systems for ownership, authorization, accountability, change management, and operational control.

Yet many organizations cannot answer basic questions about their agents:

- Which agents operate in our organization, and what do they actually do?
- Who is accountable for their behavior and business outcomes?
- What are they authorized to do, and what must they never do?
- Which systems, tools, and information can they access?
- How do we know when their capabilities, behavior, or risks change?

> **KEY FINDING**
>
> The challenge is no longer just deploying agents. It is managing them as an operational workfc

The problem: traditional management systems leave a gap

Organizations already have security policies, identity management, software inventories, approval processes, and ri frameworks.

These remain necessary. But they address different aspects of an agent's operation.

An identity management system may establish which credentials an agent uses. That does not explain whether its b are appropriate.

An application inventory may identify the platform on which an agent operates. That does not establish the agent's its escalation boundaries.

A risk assessment may document the intended behavior of an agent at deployment. That does not ensure that its o remains accurate as tools, capabilities, or processes change.

The result is fragmentation.

An organization may know which technology an agent uses without knowing what that agent is responsible for. It r created an agent without knowing who is accountable for its decisions. It may have extensive security controls with process for accepting changes to an agent's operating authority.

> **MANAGEMENT IMPLICATION**
>
> The missing element is a management system that connects the agent, its operating model, its humans responsible for it.

As adoption increases, however, the management effort grows alongside the agent population. Each new agent ad ownership, permissions, reviews, and potential exceptions.

Organizations face an unattractive choice:

SLOW ADOPTION

Slow down adoption to preserve control.

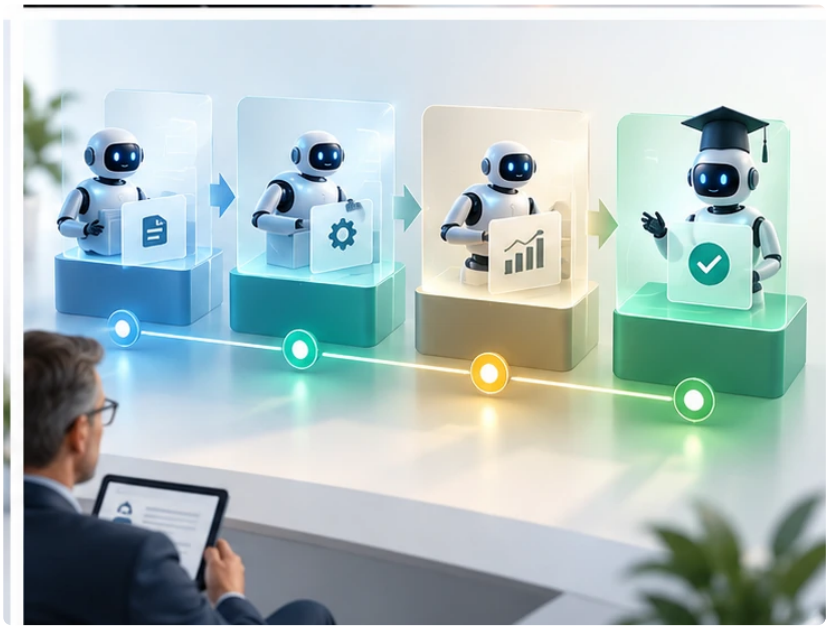
OR

LOSE PACE

Allow agents to proliferate faster th management processes can keep u

Neither is a sustainable industrial operating model.

The solution: industrialize agent management



A defined lifecycle makes agent admission, preparation, review, and accountable acceptance repeata

Industrial management does not mean treating every agent identically. It means establishing consistent manageme allowing each agent to have its own purpose, responsibilities, capabilities, and risk profile.

Every operational agent should have a defined management record connecting:

IDENTITY

What is this agent, and where does it operate?

RESPONSIBILITY

What work is it expected to perform?

AUTHORITY

What may it do, and what must it never do?

ACCESS

Which tools, systems, and information c

This changes the management question.

Instead of asking whether the organization has documented its AI technology, management can ask whether every has a current, accountable, and reviewable operating definition.

TECHNOLOGY INVENTORIES

Describe what exists.

AGENT MANAGEMENT

Establish how it is supposed to o

How CARE makes industrial agent management practical

Agentic CARE provides a common management layer for agents, independent of the platforms and tools on which

Its approach combines three capabilities.



CAPABILITY 1

Establish

Know your agents

CARE creates a managed identity and operating definition for each agent.

Rather than requiring every detail to be assembled manually, CARE supports multiple definition methods: agent self-service, s registration, facilitated registration, and manual documentation.

Where agents have sufficient context, they can perform much of the documentation work themselves. They can describe thei processes, dependencies, permissions, and operating boundaries.

This reduces the administrative burden of onboarding agents into a common management system.

Agent-generated documentation is not automatically treated as authoritative. Its quality depends on the agent's capabilities and available to it.



CAPABILITY 2

Control

Keep authority with accountable humans

CARE separates the proposal of changes from their acceptance.

CAPABILITY 3

Optimize

Improve the operating model over time

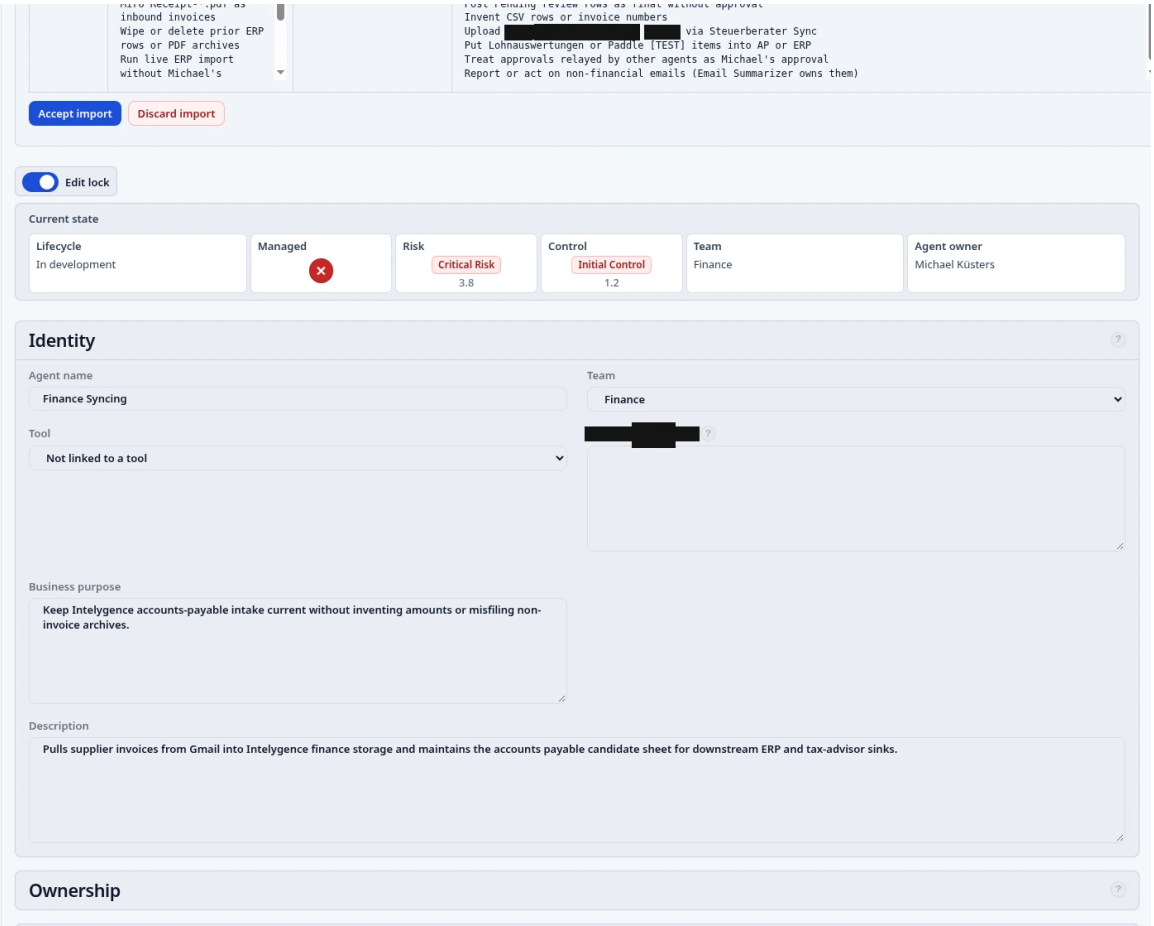
An agent's initial definition is only a starting point.

New capabilities, additional tools, changing business processes, and emerging risks can make an existing operating model in inappropriate.

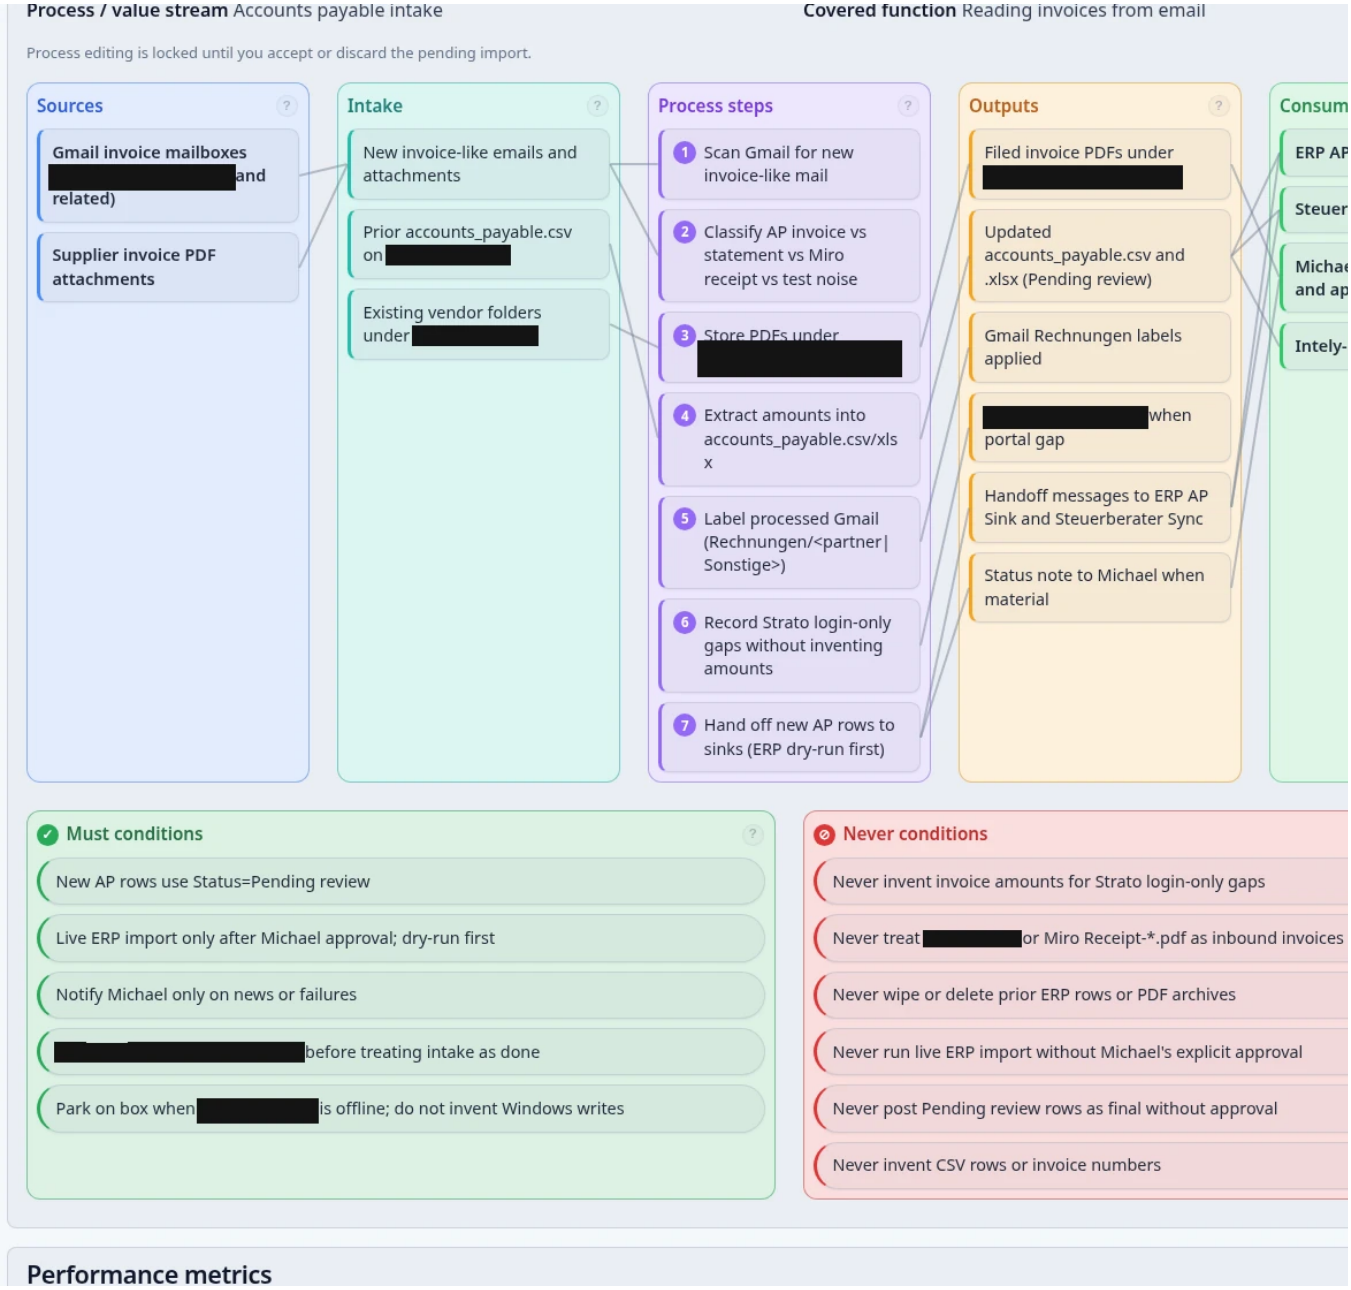
CARE provides a structured way to identify gaps, propose mitigations, record decisions, and maintain a history of changes.

That history supports operational reviews, accountability, and the production of management evidence.

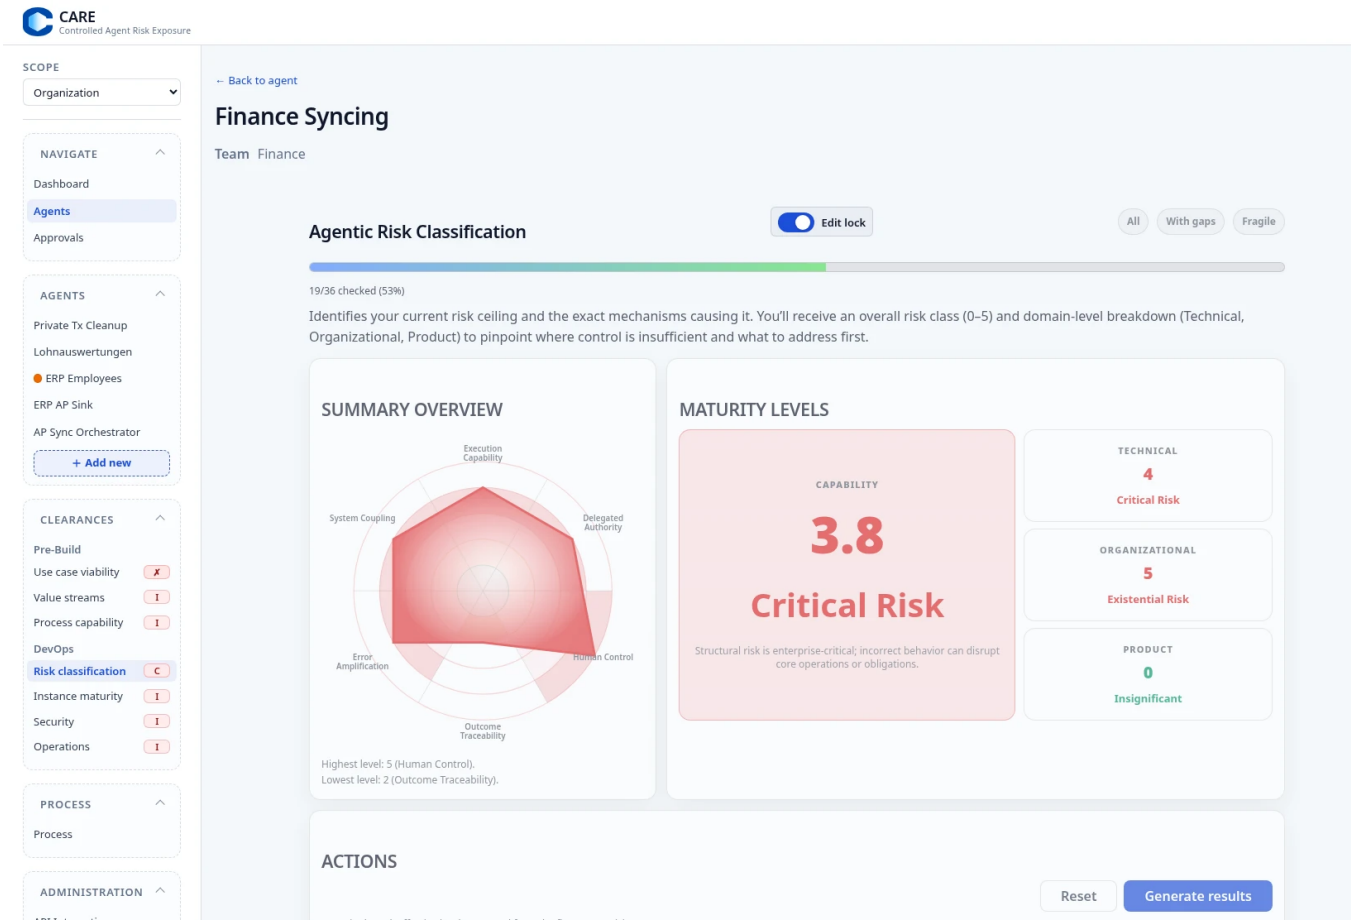
The objective is not to prevent agents from changing. It is to make change visible, deliberate, and manageable.



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 SIPOC-style value stream for Accounts payable intake: Gmail sources, seven process steps, outputs inc review AP rows, and consumers including ERP AP Sink and Steuerberater Sync.



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).

What changes for the organization?

Consider a company where several departments independently introduce AI agents.

Without a common management system, each department may rely on its own documentation, credentials, approv understanding of acceptable behavior.

As the number of agents grows, management has to reconcile these fragmented approaches.

With CARE, departments can continue using their chosen agent platforms while adopting a common management

Management concern	Fragmented approach	Industrial CARE approach
Visibility	Separate platform inventories and informal knowledge	Common agent registry
Accountability	Depends on local practices	Explicit ownership and hum
Operating boundaries	Scattered across prompts, policies, and implementations	Managed agent definitions
Risk management	Periodic or project-specific assessments	Agent-specific risks, propos decisions
Change control	Platform- or team-specific	Staged definition changes v acceptance
Documentation	Manual effort across multiple systems	Agent-assisted preparation standardized records

The benefit is not simply better documentation.

It is the ability to expand agent adoption without requiring management effort to grow proportionally with the nur

CARE does not replace identity management, platform security, or enterprise execution controls. It provides the mis layer connecting agents to their organizational responsibilities and accountable decision-makers.

From governance overhead to operational capability

For many organizations, AI governance is perceived as something added after implementation: a set of documenta assessments, and approval obligations.

Industrial agent management reverses that relationship.

An agent enters a defined management lifecycle. Its responsibilities and risks are documented as part of onboardin established route to acceptance. Management information becomes an ongoing operational output rather than a s documentation project.

This provides a foundation for organizational oversight and supports evidence collection relevant to frameworks su and applicable requirements under the EU AI Act.

It does not automatically establish regulatory compliance. It gives organizations a practical system in which respons assessments, controls, and decisions can be documented and reviewed.

The goal is not to produce more governance paperwork. It is to make accountable agent repeatable

Start with one agent.

1. Point your agent to https://ai-managementsystem.com/under-care.
2. When it's finished, log into your CARE account and review its documentation.

AGENTIC CARE

Instant AI Management System for your Agents, as a Managed Service.

https://ai-managementsystem.com