

MAYA ENTERPRISE · CASE STUDY FOR TECHNOLOGY INDUSTRY

Improving ROI in Agentic Workflows

Case Study: Novara Technologies

Human Context Infrastructure for Agentic Workflows in Enterprise Technology

Enterprise Technology · United Kingdom · 2026

WORK-SELF©

Executive Summary

Technology industry is rapidly adopting agentic workflows. Human employee and AI agent coordination is the single biggest driver of workflow automation ROI.

Novara Technologies

A 1,300-person UK enterprise SaaS company representing the WORK-SELF enterprise technology cluster: Product & Engineering, Customer Success and Sales & GTM under structural pressure as AI moves inside the operating model.

Takara.ai Independent Testing Agent

As the independent testing agent, Takara.ai ran a standardized test across 50 employees representing a customer-service team, 360 employee-by-task Work Contracts and 42 test cases across seven customer-service task categories.

33%

less rework. The back-and-forth between employee and agent fell from 3 correction rounds* to 2

74%

of the agreed rules followed by the agent, up from 32% without Maya

100%

of the customer-service task types tested improved

83%

biggest gain in Service Requests: agreed rules followed rose from 25% to 83%

Technology companies are entering a decisive transition. Product and engineering, customer success, service operations and sales teams are moving from using AI tools to supervising AI-driven workflows. The strategic question is no longer whether an enterprise can deploy agents. It is whether those agents can produce work that employees can accept, govern and use without doing the work twice.

Most agent business cases focus on model capability, licence cost and automation coverage. The value often disappears later, when employees reread output, rewrite drafts, reroute cases, correct ownership and manually manage escalation. The agent is deployed, but the department does not realise the full operating advantage.

Maya Enterprise (WORK-SELF product) is designed to improve coordination between employees and AI agents by giving each agent the relevant employee and organisational context at the point of work. Maya's Human Context MCP delivers that context to the agent, helping it follow the correct ownership, review, escalation and workflow requirements.

Takara.ai independently tested Maya using the same AI model, tasks and scoring in two standardized testing groups, one with Maya context and one without. The test asked a simple question: does an agent perform better when it understands the operating context of the employee and task?

> With Maya, correction rounds fell from 3 to 2. The agent also followed more of the rules governing ownership, review, escalation and deadlines. Every customer-service task type tested improved, with the largest gains on work requiring the most human judgement and rework.

WHAT THIS MEANS FOR THE TECHNOLOGY BUSINESS

01

Faster time to value

Work moves from initial agent output to approval with one fewer correction round.

02

More capacity

Skilled employees spend less time checking and correcting output, leaving more time for higher-value work.

03

Safer scale

The AI agent handles ownership, reviews, escalations and deadlines reliably, enabling scale with less supervision.

Implementing Maya saves Novara Technologies **£5.38M** annually by **reducing employee rework by 33%**.

**Note: What is a correction round? One pass of an employee taking the agent's output, checking or correcting it, and handing it back so the agent can continue. Three rounds down to two means the work is finished in two passes instead of three.*



From AI Tools to Agentic Workflows of the Future

The next competitive advantage is turning agent output into accepted, governed work with less human correction, as AI moves deeper into the operating model.

TODAY
Isolated AI tools

→

NEXT 12–24 MONTHS
Agentic workflows

→

THE DECIDING LAYER
Accepted, governed work with less human correction

Every enterprise technology company now has to decide which workflows will become agentic, which decisions will remain human-led, and how those two sides will work together. WORK-SELF's market thesis is that this transition will happen over the next 12–24 months across product and engineering, customer success, sales and go-to-market, IT operations, finance and other shared services.

Enterprise technology firms face a structural reckoning. AI is already inside the operating model. Code generation, automated testing, AI-first service, sales assistance and workflow automation are compressing the value of manual execution. Competitive advantage moves to the company that can convert agent output into reliable operating capacity faster than its peers.

The problem is not that agents cannot produce answers. The problem is that a specific employee must still accept responsibility for the outcome. Without human context, the agent may send work to the wrong owner, miss the expected review route, escalate too late, or produce an answer that is plausible but unusable.

Employees then reread, rewrite and manage the agent by hand. The efficiency gain evaporates at the handoff.

The industry has spent three years making agents smarter. The next advantage comes from making them aware of the human responsible for the work.

Where agent ROI is lost?

Automation ROI disappears when employees still have to reread, reroute and correct agent output, turning AI into extra work rather than operating capacity.

When employees must reread, rewrite and reroute agent output, the expected savings are reduced and the full ROI from workflow automation is not realised. In the customer-service workflows tested for Novara, this problem was directly observable. Similar patterns can also emerge across product, sales, operations and other functions within technology companies.

01 Agent produces output → **02** Employee rereads and rewrites → **03** Ownership and route corrected → **04** Escalation managed by hand → **RESULT** Automated in name, not in economics

The value is realised only when **agent output reaches the right owner**, **follows the right controls** and **requires less human correction**.

The technology functions under pressure

FUNCTION UNDER PRESSURE	WHAT BECOMES AGENTIC	WHERE ROI IS LOST WITHOUT MAYA
Product & Engineering	Code assistance, testing, incident and release workflows.	Senior staff review generic output, correct ownership and take difficult exceptions back.
Customer Success & Support	Triage, response drafting, routing, escalation and service recovery.	Work sent to the wrong person, repeated corrections, escalations missed and deadlines handled inconsistently.
Sales & GTM	Account research, opportunity support, proposal and handoff workflows.	Generic output must be rewritten for the account, owner and approval path.
Operations & Shared Services	Case handling, onboarding, routing for approval, and field dispatch.	Humans remain the manual control layer because decision rights are not explicit.

MARKET EVIDENCE

1%

of companies believe they have reached AI maturity; the binding constraint is leadership steering, not employee resistance.

McKinsey, Superagency in the Workplace.

2030

the horizon by which AI is expected to structurally reshape the labor market, with technology among the most affected sectors.

WEF, Future of Jobs 2025.

2 in 5

workers will require reskilling due to AI within three years; reskilling is the #1 organizational challenge globally.

IBM, Augmented Workforce 2025.

How Was Maya Performance Tested?

Takara.ai independently tested the same customer-service tasks with and without Maya context. The AI model, tasks, test cases and scoring remained the same, making employee and task context the only variable.

Takara.ai designed, ran and scored an independent standardized test to assess whether Maya's employee and task context improves agent performance. Two testing groups completed the same customer-service tasks, one with Maya context and one without. Both groups used the same AI model, tasks, test cases and scoring methodology. The only variable was whether the agent received permissioned context for the specific employee and task from Maya.

Each answer was assessed against the same four operating rules: correct ownership, review route, escalation and deadline. Takara.ai held the correct answers and independently scored both groups using the same criteria, creating a like-for-like comparison of agent performance with Maya on and Maya off.

A Like-for-Like Comparison

Held constant in both groups
Same AI model, tasks, test cases and scoring.

ARM A
Maya off
Agent receives no employee- or task-specific human context.

ARM B
Maya on
Agent receives a permissioned Context Capsule for this employee and this task.

Single variable: whether the agent received Maya context

Source: Takara.ai standardized testing methodology

WHAT WE TESTED

Does giving an AI agent the right employee and task context reduce rework and improve rule-following?

ELEMENT	WHAT WAS TESTED
Evaluation population	50 employees, representing a customer-service team.
Operating rules	360 employee-by-task Work Contracts.
Test set	42 test cases across the 7 customer-service task types (service request, complaint, billing, technical support, account access, feedback and other); 6 cases per category.
Controls	The same AI model, tasks and scoring method in both arms.
Single variable	Whether the agent received Maya context for the specific employee and task.
Scoring	Each answer scored on four rules: correct owner, correct review route, correct escalation, agreed deadline met.







How Maya Employee-Agent Coordination Works?

Maya gives each AI agent the right context and each employee clear decision rules, reducing rework and strengthening control.

Technology companies need better coordination between employees and AI agents to reduce rework, eliminate inefficiencies and improve ROI from automation. WORK-SELF Maya Enterprise provides this coordination by capturing relevant employee preferences and context, then serving that context to AI agents at the point of work via MCP. This improves agent output, supports employees as their roles change and gives leadership an auditable view of adoption, risk and performance.



MAYA CAPABILITY	WHAT IT DOES IN PLAIN LANGUAGE	VALUE TO NOVARA
01 Understand the work and the person	Combines authorized employee, role, task, policy and workflow information.	A usable view of who is responsible, what is changing and where support is required.
02 Set clear operating rules	Creates a Work Contract for one employee doing one task: ownership, agent autonomy, review, escalation and deadline.	Less ambiguity at the handoff; more consistent work and supervision.
03 Give the agent only what it needs	Sends a small, permissioned Context Capsule to the approved agent at runtime.	Better output without exposing unnecessary personal data or rebuilding the agent stack.
04 Help the employee supervise AI	Uses Review Maps, autonomy settings and stop/escalate signals to show what to read, decide, delegate or stop.	Less manual quality control and faster confidence in the new workflow.
05 Improve and govern the system	Records accepted, edited, escalated and rejected outputs and updates the operating rules under governance.	Leadership can see rework, adoption, control issues and improvement over time.
06 Support workforce transition	Coaches employees through role change and gives leaders readiness, redeployment and transition visibility.	The technology deployment and the workforce change are managed as one programme.

Maya is Human Context Infrastructure for enterprise AI. It is the governed layer that coordinates the agent and the employee responsible for the outcome.

Maya gives each AI agent the right context and each employee clear decision rules, reducing rework and strengthening control.



Maya is interoperable with Novara's Existing Technology Stack

Maya gives approved AI agents the right human context without requiring Novara to replace its existing systems, models or workflows. WORK-SELF interoperability layer ensures connectivity with any existing system and data already in the enterprise.

Novara does not need to replace its existing systems, models or workflows. Maya reads authorised information from the existing technology stack, turns it into governed human operating context and provides each approved AI agent with only the context it needs for the task. Existing systems continue to provide the enterprise record, while Maya adds the employee, task and decision context needed for the agent and employee to work together effectively.

DATA ALREADY IN THE ENTERPRISE	WHAT MAYA CREATES	BUSINESS VALUE
HRIS: role, team, manager chain, skills, tenure	Employee and role context	Correct ownership, role boundaries and appropriate support.
CRM and ticketing: case history, task type, customer interaction	Employee-by-task Work Contract	Better routing and review, escalation at the right moment, and deadlines met more often.
ERP and workflow: process steps, throughput, cost and exceptions	Workflow context and decision boundaries	Clearer division of work between agent and employee.
Knowledge base: policies, SOPs, style guides and service rules	Permissioned Context Capsule	The agent receives the minimum context required for the task.
Agent and human feedback: accepted, edited, escalated or rejected output	Audit record and governed learning loop	Traceability, measured improvement and evidence of human oversight.

REPRESENTATIVE TECHNOLOGY STACK

Workday / BambooHR · Salesforce / HubSpot · Jira / Linear · Zendesk-class service platforms
Confluence / Notion · Slack / Teams · CockroachDB / PostgreSQL · Workato / UiPath · any approved MCP-enabled agent

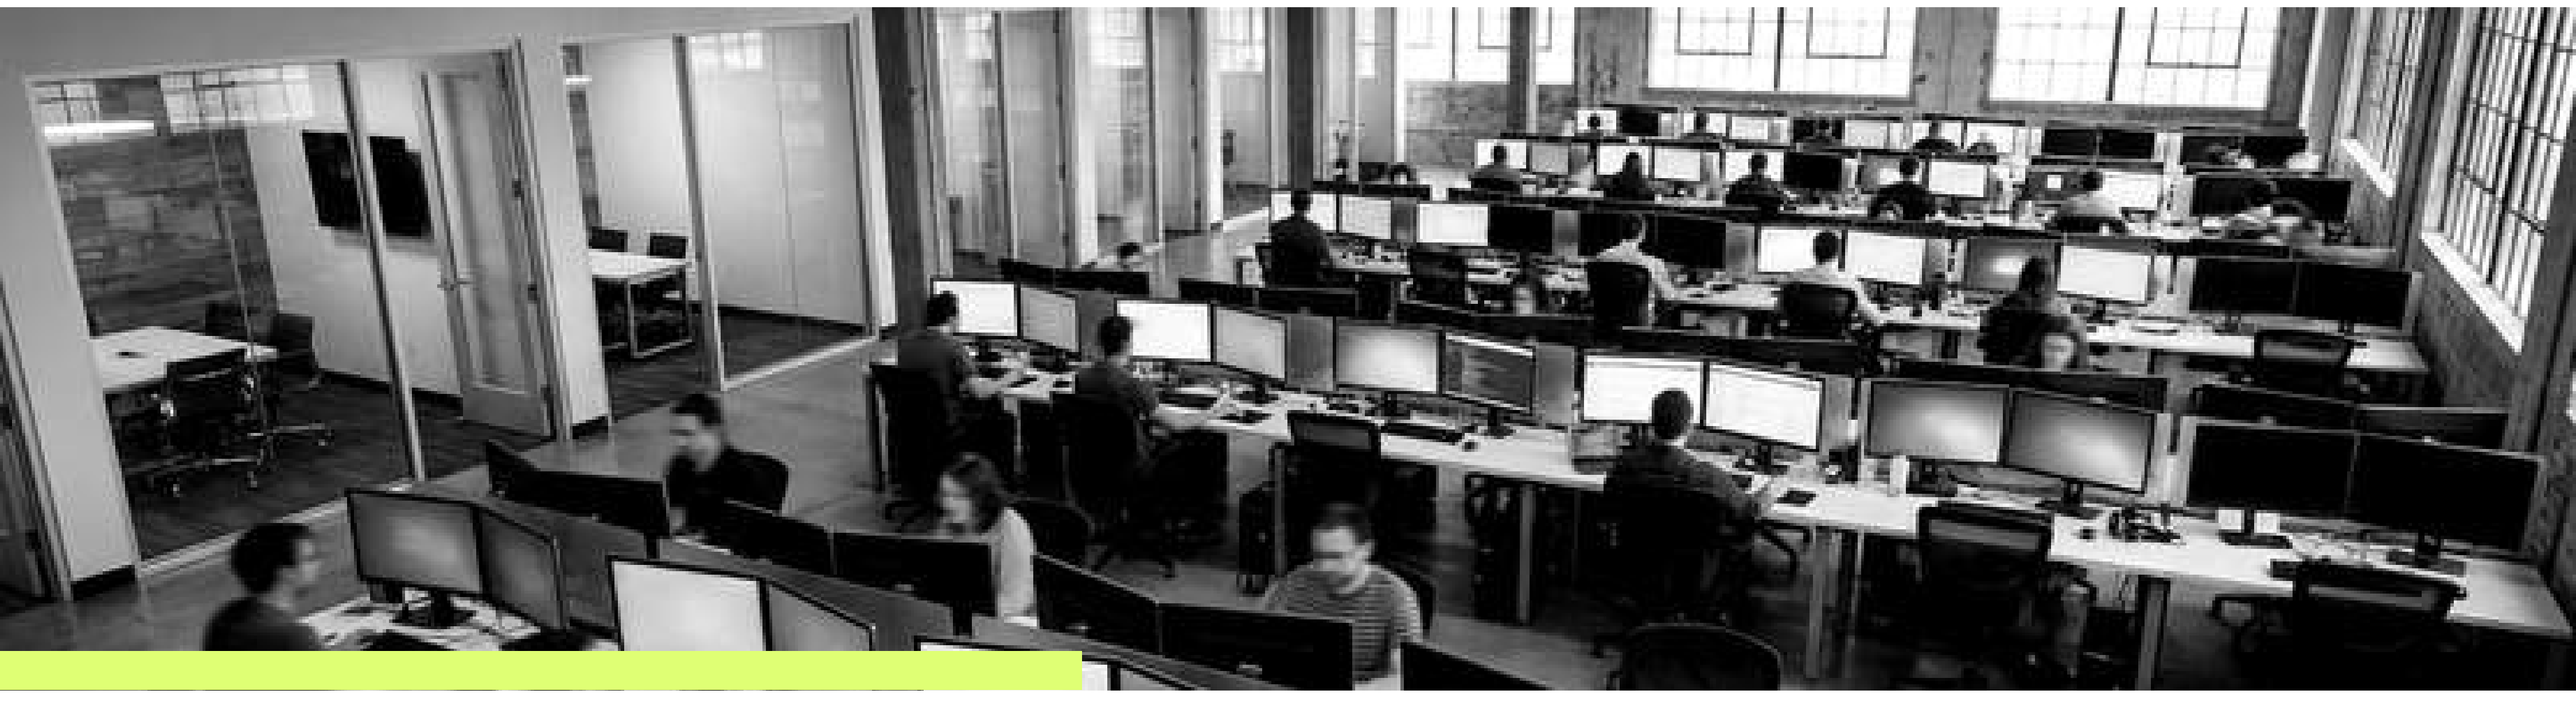
MCP

MCP is the open standard Maya uses to pass governed context to approved AI agents.
This enables Novara to **add human context to its existing agent stack without rebuilding its models, tools or workflows**. Maya is **built to work with any model;** this standardized testing methodology used one AI model, applied identically in both arms.

Maximum interoperability

The systems of record tell Maya what the organization looks like. Maya adds how this person and this agent should work together on this task.



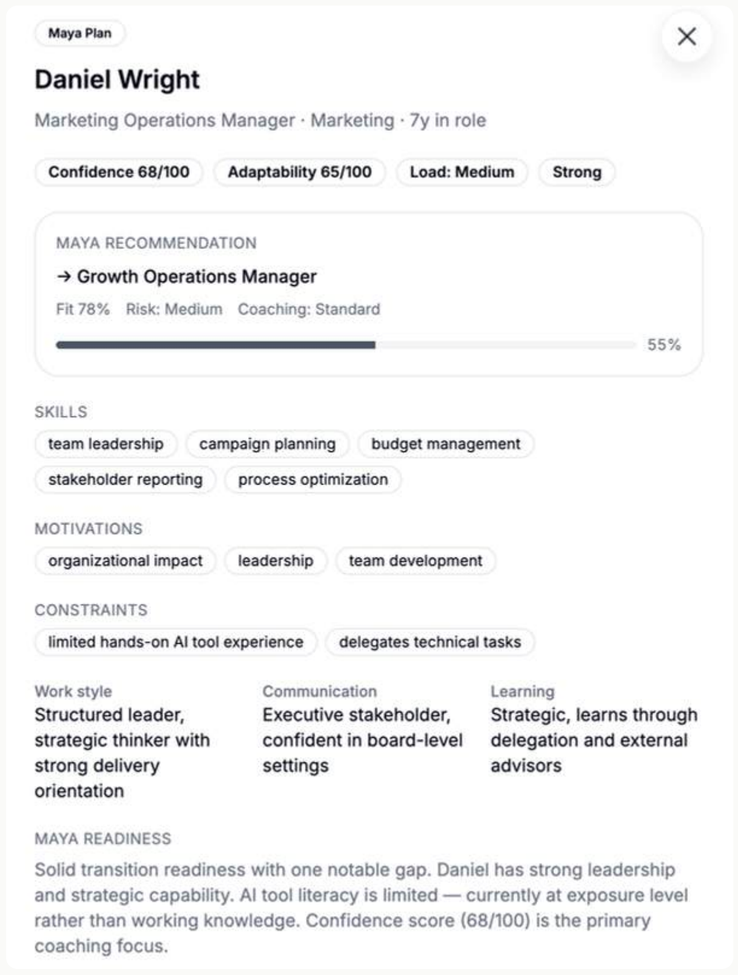


What are core components of Maya Enterprise?

Four operating views give Novara's leadership visibility from individual decision rights through to workforce readiness, risk, agent performance and ROI realization.

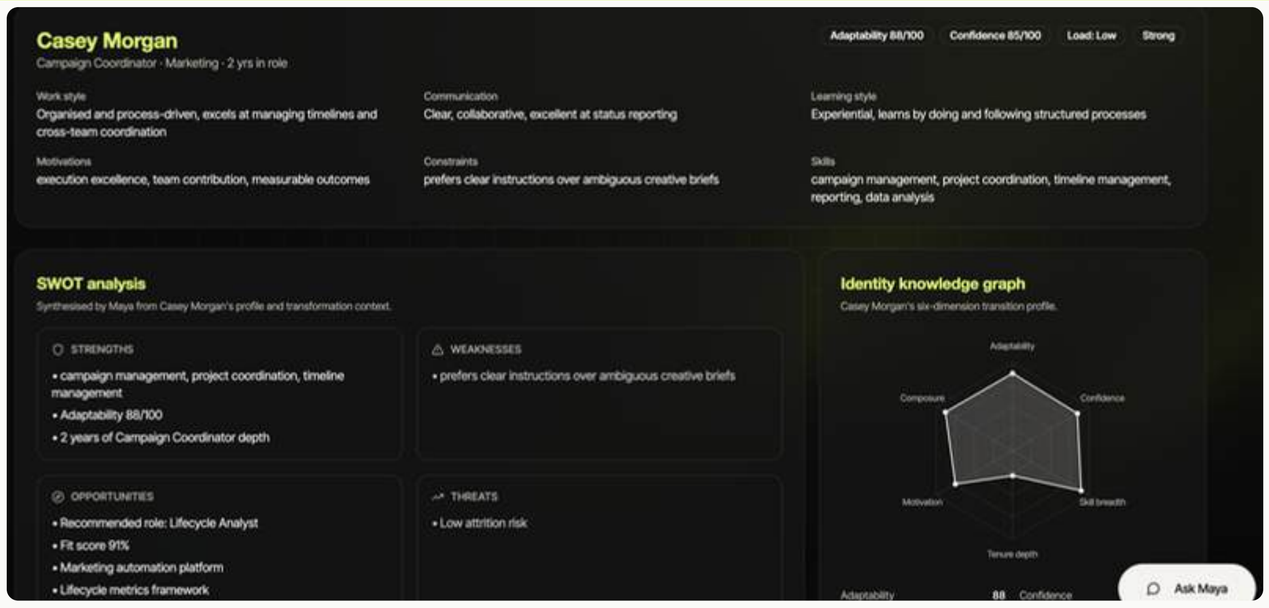
Whether CEO, COO, CFO, CHRO or CAIO, Maya Enterprise gives Novara's leadership continuous visibility into readiness, risk and progress across agentic workflows. Its dashboards and operating views bring together employee transition and adoption, AI-agent performance, compliance and ROI realisation, helping executives see where the organisation is ready to scale, where intervention is required and where automation is delivering value. The four views below illustrate some of the ways Maya provides that visibility.

01 Review Map



Shows the employee where human judgement remains necessary within an agent-assisted workflow, including what to review, decide, delegate or escalate. It helps prevent employees from becoming the manual quality-control layer for AI output and gives Novara clearer visibility into where human intervention is still required.

02 Work Contract card



Defines how an employee and AI agent work together on a specific task, including ownership, decision rights, agent autonomy, review requirements, escalation triggers and deadlines. It gives Novara a governed operating record of who is accountable for the outcome and where the boundary between human and agent responsibility sits.

03 Context Capsule card



Provides an approved AI agent with the permissioned employee, task and organisational context required to perform a specific piece of work. It controls what context the agent receives at the point of work, helping Novara improve output while limiting unnecessary access to employee information and maintaining traceability.

04 Executive dashboard tile



Brings together workforce readiness, adoption, agent performance, rework, escalations, compliance and risk across agent-assisted workflows. It gives Novara's leadership an operating view of where automation is delivering value, where intervention is required and where the organisation is ready to scale.

How Maya Enterprise Improves ROI in Automated Workflows

In Novara's standardized customer-service test, Maya reduced employee correction rounds from three to two, a 33% reduction in rework, while adherence to operating rules increased from 32% to 74%.

Takara.ai observed significant improvements across the customer-service workflows tested for Novara. The most prominent was a reduction in employee rework, with one correction step removed in every three, while the agent also followed ownership, review, escalation and deadline requirements more consistently. All seven customer-service task types tested improved.

3 to 2

correction rounds per task

33%

reduction in rework

74%

of the agreed rules followed, up from 32%

100%

of the customer-service task types tested improved

As described earlier, Takara.ai constructed two testing groups using the same AI model, tasks, test cases and scoring. The only difference was whether the agent received Maya context for the specific employee and task. Takara.ai held the correct answers and scored every response against four operating rules: correct ownership, review route, escalation and agreed deadline. This created a like-for-like comparison in which Maya context was the only variable.

TAKARA STANDARDIZED TESTING METHODOLOGY

MEASURE	MAYA OFF	MAYA ON	MEASURED CHANGE
Correction rounds to reach the right answer	3	**2**	One correction round removed in three **- a 33% reduction in rework**.
Agreed rules followed by the agent	32%	**74%**	The agent followed **74% of the agreed rules**, up from 32% without Maya.
Customer-service task types improved	–	**100%**	All seven customer-service task types **improved**, in the same direction.

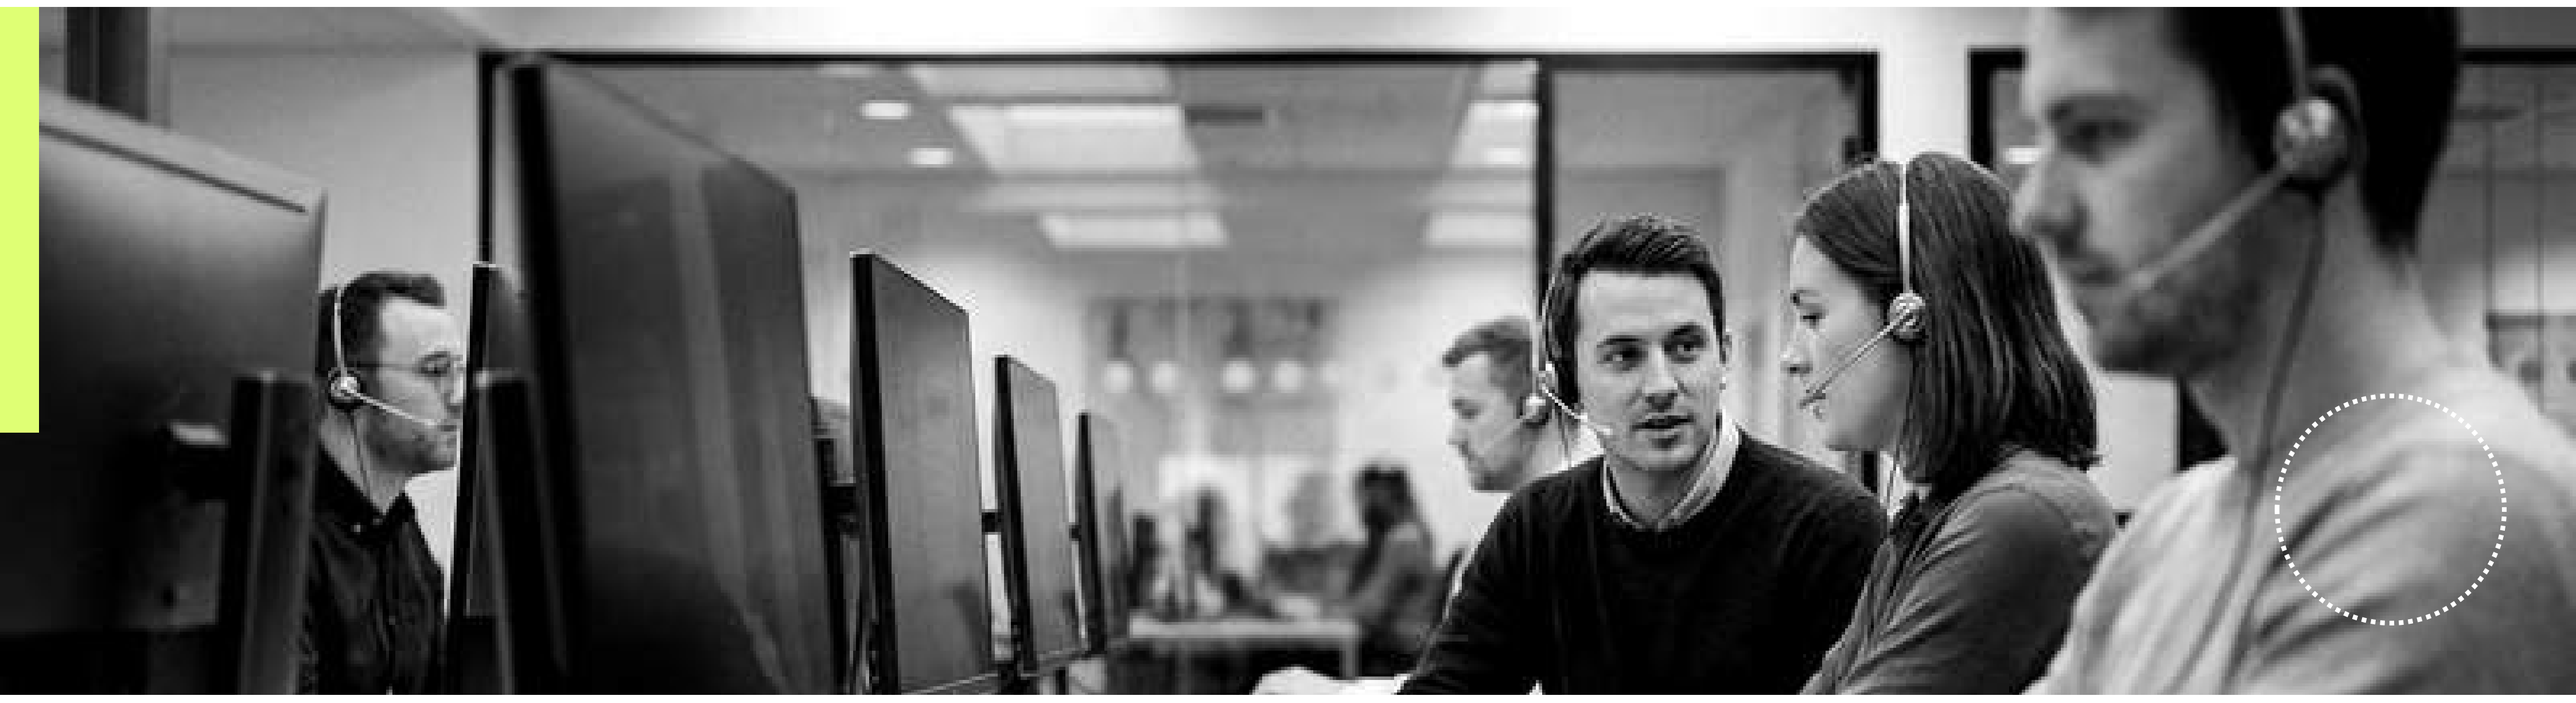
Source: Takara.ai independent standardized test

Same model. Different context.

The improvement came from the context, not from a better model. Maya raises the value of the agent stack you already own.

Why the context matters

Generic context helped only so far. When the agent received the wrong employee's context, the improvement disappeared, showing that the gain comes from knowing who is responsible for the task.



Maya Delivered the Biggest Gains on the Hardest Tasks

Novara's customer-service test covered the most common service workflows, and the largest gains appeared where company-specific context and human handoffs matter most.

The standardised test covered seven common customer-service task types for a technology company like Novara, including service requests, complaints, billing, technical support, account access, feedback and other cases. Maya improved performance across every category, with the largest gains in service requests, complaints and technical support, where company-specific context, human judgement and handoffs matter most.

TASK CATEGORY	MAYA OFF	MAYA ON	GAIN	RELATIVE MOVEMENT
Service Request	25%	83%	**+58 pts**	**83%**
Complaint	8%	63%	**+54 pts**	**63%**
Billing	38%	83%	**+46 pts**	**83%**
Technical Support	21%	63%	**+42 pts**	**63%**
Other	29%	67%	**+38 pts**	**67%**
Account Access	50%	83%	**+33 pts**	**83%**
Feedback	50%	75%	**+25 pts**	**75%**

Source: Takara.ai independent standardized test, August 2026. Percentages show how many of the four agreed rules the agent got right. Six test cases per task type. A task counts as difficult when it needs a lot of company-specific knowledge and several handoffs before it can be closed.

INTERPRETATION

MEASURED RESULT	EXECUTIVE INTERPRETATION	BUSINESS IMPACT
33% reduction in rework	Work reaches an acceptable state with one fewer employee correction round in three.	Work is approved sooner, employees spend less time reviewing, and the agent investment pays back faster.
Agreed rules followed rose from 32% to 74%	The agent gets ownership, review route, escalation and the deadline right far more often.	Fewer avoidable delays, less cost to handle each case, and more control for the business over how the work runs.
All 7 categories improved	The gain held across the whole tested workload, not just on the easy tasks.	If Maya improves every customer-service task type, the same gain is likely on the similar tasks in finance, operations, sales and marketing.
Largest gains on the most difficult tasks	The biggest gains came on complaints, service requests and technical support.	The strongest value sits where risk, customer impact and the cost of human review are highest.

The more context a task required, the more Maya helped, reducing the need for employees to take the work back from the agent.

What 33% Less Rework Means for Novara

For Novara, a 33% reduction in employee rework represents £5.38M in annual labor value and 112,000 skilled hours returned to higher-value work.

MEASURED RESULT	OPERATING EFFECT	FINANCIAL LINE TO MODEL	STRATEGIC OUTCOME
One correction round removed in three	Less human review per completed task.	Hours reclaimed × fully loaded labor cost.	Faster payback and more capacity without linear review headcount.
Agreed rules followed rose from 32% to 74%	Ownership, review route, escalation and deadlines handled correctly more often.	Cost avoided from delays, work sent to the wrong person, missed deadlines and rework.	More reliable service and a stronger control environment.
Most difficult tasks improve most	Less takeover by senior or specialist staff.	Reduced high-cost review and escalation burden.	Complex workflows can move into the agentic operating model sooner.
Auditable context and decisions	Clear evidence of what the agent knew and what the human decided.	Lower assurance and investigation effort; fewer control gaps.	Greater confidence from risk, procurement and leadership.

Capacity and Value Impact

INPUT	VALUE
Employees in agent-assisted workflows [N]	**250**
Agent-assisted tasks per employee per week [T]	**80**
Minutes per correction round [M]	**7**
Working weeks per year	**48**
Fully loaded hourly cost [H]	**£48**

Calculated directly from the Novara case-study operating inputs shown on this page.

Correction rounds removed per year

960,000

employees × tasks per week × 48 weeks × 1 round removed

Skilled hours reclaimed

112,000 hrs

rounds removed × [M] ÷ 60

Annual labor value saved

£5.38M

hours reclaimed × hourly cost.

> For Novara, reducing employee rework by 33% created £5.38M in annual labour value by returning 112,000 skilled hours to productive work. The benefit comes from less manual review, fewer avoidable corrections and escalations, and more capacity from the existing workforce.





The Human Handoff Has a Measurable Solution

The next 12–24 months will separate agent deployment from agentic operating advantage.

Technology companies will not become more competitive simply by buying more agents. They will become more competitive when agent output moves through the organization faster, reaches the correct owner, follows the right controls and allows employees to focus on the work that still requires human judgement.

Takara.ai's independent standardized test shows that task-specific human context changes operating performance. With Maya, one full correction cycle was removed, adherence to ownership, review, escalation and deadline rules improved, and every customer-service task category tested performed better. The largest gains appeared where company-specific context, human judgement and rework matter most.

At an executive level, a 33% reduction in rework gives the existing automation a shorter path to value. The department can handle more work with the same skilled Team, service can become faster and more consistent, and leadership can scale the agentic operating model with stronger control.

Maya is the Human Context Infrastructure for that transition. It joins the agent stack already in place, replaces nothing, helps employees adapt to AI-driven roles, gives agents the human operating context they are missing, and gives leaders an auditable view of performance and risk.

For Novara, the measurable result with Maya Enterprise is less employee rework, stronger adherence to operating rules and governance, amounting to £5.38M in annual value at the minimum.

33%

less rework. The back-and-forth between employee and agent fell from 3 correction rounds to 2

83%

biggest gain in Service Requests: agreed rules followed rose from 25% to 83%

£5.38M

annual labor value realised by reclaiming 112,000 skilled hours through reduced rework.

Request an Enterprise Demo

www.work-self.com/enterprise

Book an Enterprise Consultation

api.leadconnectorhq.com/widget/bookings/ws-enterprise-consultation

About Us

WORK-SELF

WORK-SELF builds Maya Enterprise, the Human Context Infrastructure layer for enterprise AI. Maya helps employees adapt to AI-driven work and helps approved AI agents adapt to the employees responsible for the outcome. The platform combines governed human context, workforce transition support, human-agent operating rules and executive visibility in one enterprise solution.

TAKARA.AI

Takara.ai is an AI engineering firm building production agent systems for enterprise clients. As the independent testing partner, Takara.ai designed, ran and scored the standardized test used in this case study.



Wolf Magdelinic
Co-Founder & CEO, WORK-SELF
Wolf.magdelinic@work-self.com



Abinesh Kanagaraj
Corporate Development, WORK-SELF
Abinesh.kanagaraj@work-self.com

WORK-SELF©

