

The Business Value Dashboard: Seeing the Real Impact of Your AI Agents

Oracle Fusion Cloud Applications | AI Agent Studio

June 2026, Version 1.0

Copyright © 2026, Oracle and/or its affiliates

Public

The AI accountability problem

Organizations deploying AI agents face a familiar paradox: leadership demands proof of ROI before expanding investment, but proving ROI requires the kind of measurement infrastructure that most teams build last, if they build it at all.

The result is a credibility gap. AI programs that are delivering real value struggle to quantify it. Budget holders approve pilot after pilot without a clear line of sight to cumulative impact. Agent creators iterate on their designs without knowing which improvements move the needle. And when renewal conversations arrive, the case for continued investment rests on anecdotes rather than data.

This isn't because the value isn't there. It's because measuring AI's impact across a portfolio of agents, teams, and business functions requires purpose-built tooling.

Why AI value is hard to measure

Traditional IT ROI is relatively straightforward: a system costs X, and it replaces Y hours of manual work. The math is simple. AI agents, however, introduce complexity that these models don't handle well. For example:

- **Value is distributed, not centralized.** A single agent might save ten minutes per interaction across thousands of users. The savings are real but diffuse because no one person experiences the full impact, so no one person can report it.
- **Value compounds across agents.** An organization running fifty agents across HR, finance, and service accumulates savings that are meaningful in aggregate but hard to attribute to any single deployment decision.
- **Estimates evolve.** When an agent is first deployed, the time and cost savings per interaction are estimates. As usage patterns emerge and the agent matures, those estimates need refinement. Without a feedback loop, initial projections calcify into fiction.
- **Multiple stakeholders need different views.** A CFO wants portfolio-level impact on the bottom line. An AI program lead needs adoption rates and savings by business unit. An agent creator needs per-agent performance to know where to invest iteration effort. The conventional one dashboard, one view construct doesn't serve any of them well.

Introducing the Business Value Dashboard

The Business Value Dashboard is a new capability within the Monitoring and Evaluation module of Oracle Fusion Cloud Applications' AI Agent Studio. It provides a clear, estimated view of ROI—both time savings and cost savings—for your entire AI agent portfolio and each individual agent. And it's updated automatically with real usage data.

A global enterprise running 20 agents across Oracle Cloud ERP and HCM, saving an average of 10 minutes per interaction across 50,000 monthly uses, could unlock more than 100,000 hours annually. That represents millions of dollars in direct productivity gains.

How it works

The system is built on a simple but effective foundation:

- **Every agent has a business metrics template.** When an agent or workflow is created, the author specifies estimated time savings (in minutes) and cost savings (in USD) per agent interaction. Time savings is mandatory; cost savings is optional. These estimates can be seeded with recommended values and refined at any time.

- **ROI calculates automatically from real usage.** The system tracks actual agent interactions—how many times each agent runs, how many events it processes—and multiplies that usage against the author-provided value estimates. The calculation is transparent: $\text{ROI} = \text{number of relevant events} \times \text{estimated value per event}$.
- **Only real work counts.** Debug runs and test executions are excluded from calculations. Only end-user interactions contribute to the ROI figures, ensuring the numbers reflect actual business impact.
- **Estimates are editable, not fixed.** As organizations learn more about how agents perform in practice, they can update the estimated values at any time. The dashboard recalculates within minutes, ensuring that ROI figures stay aligned with operational reality rather than locked to initial guesses.

What the dashboard shows

The dashboard serves multiple levels of the organization with appropriate views:

- **Portfolio-level roll-up:** The default view shows total time savings and total cost savings across all agents and agent teams, aggregated across every Fusion product family. A CFO or budget holder can see the full impact of AI adoption at a glance, without selecting individual agents. This view supports filtering by product family and product, and it displays the top-performing agents and agent teams by both time and cost savings.
- **Product and family views:** An AI program lead overseeing Fusion Cloud HCM can filter to see all HCM agents—recruiting, benefits, compensation, payroll—and compare performance across them. A CHRO can view savings specifically for agents used in Compensation, Recruiting, and Benefits. Multiple products and families can be selected simultaneously for comparative analysis.
- **Per-agent detail:** An agent creator can drill into any individual agent team to see its specific time savings, cost savings, and usage volume over time.
- **Flexible time units:** Time savings can be displayed in minutes, hours, days, weeks, months, or years, making the numbers meaningful regardless of scale. An agent saving ten minutes per interaction might show thousands of hours saved when viewed at the annual level.

The agent team view

Many organizations deploy agents in coordinated teams. A recruiting team might include agents for resume screening, candidate communication, interview scheduling, and offer generation. The dashboard handles this naturally:

- Agent teams display aggregate savings across all member agents.
- Individual agent contributions are visible within the team view.
- Teams with agents measuring different metrics (some tracking time, others tracking cost) show a clear breakdown by metric type.

Editing and refinement

The dashboard acknowledges an important reality: value estimates for AI agents are just that—estimates. Organizations are encouraged to treat initial values as hypotheses and refine them as operational data accumulates.

From the Monitoring and Evaluation interface, authorized users can:

- Update the estimated time or cost savings per agent run
- See the immediate impact of revised estimates on the dashboard (recalculation within five minutes)

This edit capability is essential for maintaining credibility with budget holders. As the AI program matures, the numbers should get *more* accurate, not less, and the dashboard supports that progression.

Who benefits

Here's how the Business Value Dashboard supports the requirements of different users:

- **CFOs and budget holders** get a defensible, data-backed view of AI's impact on the bottom line. The portfolio roll-up provides the executive summary needed for board presentations and investment decisions, grounded in actual usage rather than vendor projections.
- **AI program leads** get the operational intelligence they need to steer rollout strategy. Which agents are delivering the most value? Where should the next wave of deployment focus? Which pilots aren't performing to expectations and need attention? The dashboard answers these questions with data, not intuition.
- **Agent creators** get a direct feedback loop between their design decisions and real-world impact. Understanding which agents drive the most savings and which need improvement can transform agent development from a craft exercise into a measurable discipline.

What's next

The current release delivers the foundational ROI calculation engine and dashboard. The roadmap includes:

- LLM-generated value estimates that analyze an agent's description and automatically suggest time and cost savings benchmarks, reducing the burden on agent creators to estimate from scratch.
- Implicit feedback integration that refines accuracy by distinguishing between successful completions and interactions that required human intervention, producing a more honest picture of actual (not theoretical) savings.
- Multi-currency support beyond the initial USD-only implementation.

Getting started

The Business Value Dashboard is available within AI Agent Studio's Monitoring and Evaluation module. To start seeing value:

1. **Populate business metrics for your existing agents.** Review each active agent to verify time savings estimates are filled in. Cost savings are optional but strengthen the ROI picture.
2. **Calibrate estimates against reality.** For agents that have been running long enough to generate operational data, compare initial estimates against what you're actually seeing. Update values where needed.
3. **Establish a review cadence.** Monthly or quarterly reviews of the dashboard by internal stakeholders and AI leaders create the accountability loop that keeps AI investment decisions grounded.
4. **Use the data in investment conversations.** The dashboard is designed to produce the kind of evidence that budget holders need. Use it.

The hardest part of proving AI's value is making it visible. The Business Value Dashboard solves that problem, turning every agent interaction into a measurable data point in your organization's AI investment story.

The Business Value Dashboard is available within Oracle AI Agent Studio's Monitoring and Evaluation module. For more information, contact your Oracle account team or visit [the Fusion AI documentation](#).

Connect with us

Call +1.800.ORACLE1 or visit **oracle.com**. Outside North America, find your local office at: **oracle.com/contact**.

 blogs.oracle.com

 facebook.com/oracle

 twitter.com/oracle

Copyright © 2026, Oracle and/or its affiliates. This document is provided for information purposes only, and the contents hereof are subject to change without notice. This document is not warranted to be error-free, nor subject to any other warranties or conditions, whether expressed orally or implied in law, including implied warranties and conditions of merchantability or fitness for a particular purpose. We specifically disclaim any liability with respect to this document, and no contractual obligations are formed either directly or indirectly by this document. This document may not be reproduced or transmitted in any form or by any means, electronic or mechanical, for any purpose, without our prior written permission.

Oracle, Java, MySQL, and NetSuite are registered trademarks of Oracle and/or its affiliates. Other names may be trademarks of their respective owners.