FRONTIER.

FRONTIER / NEW ANALYSIS

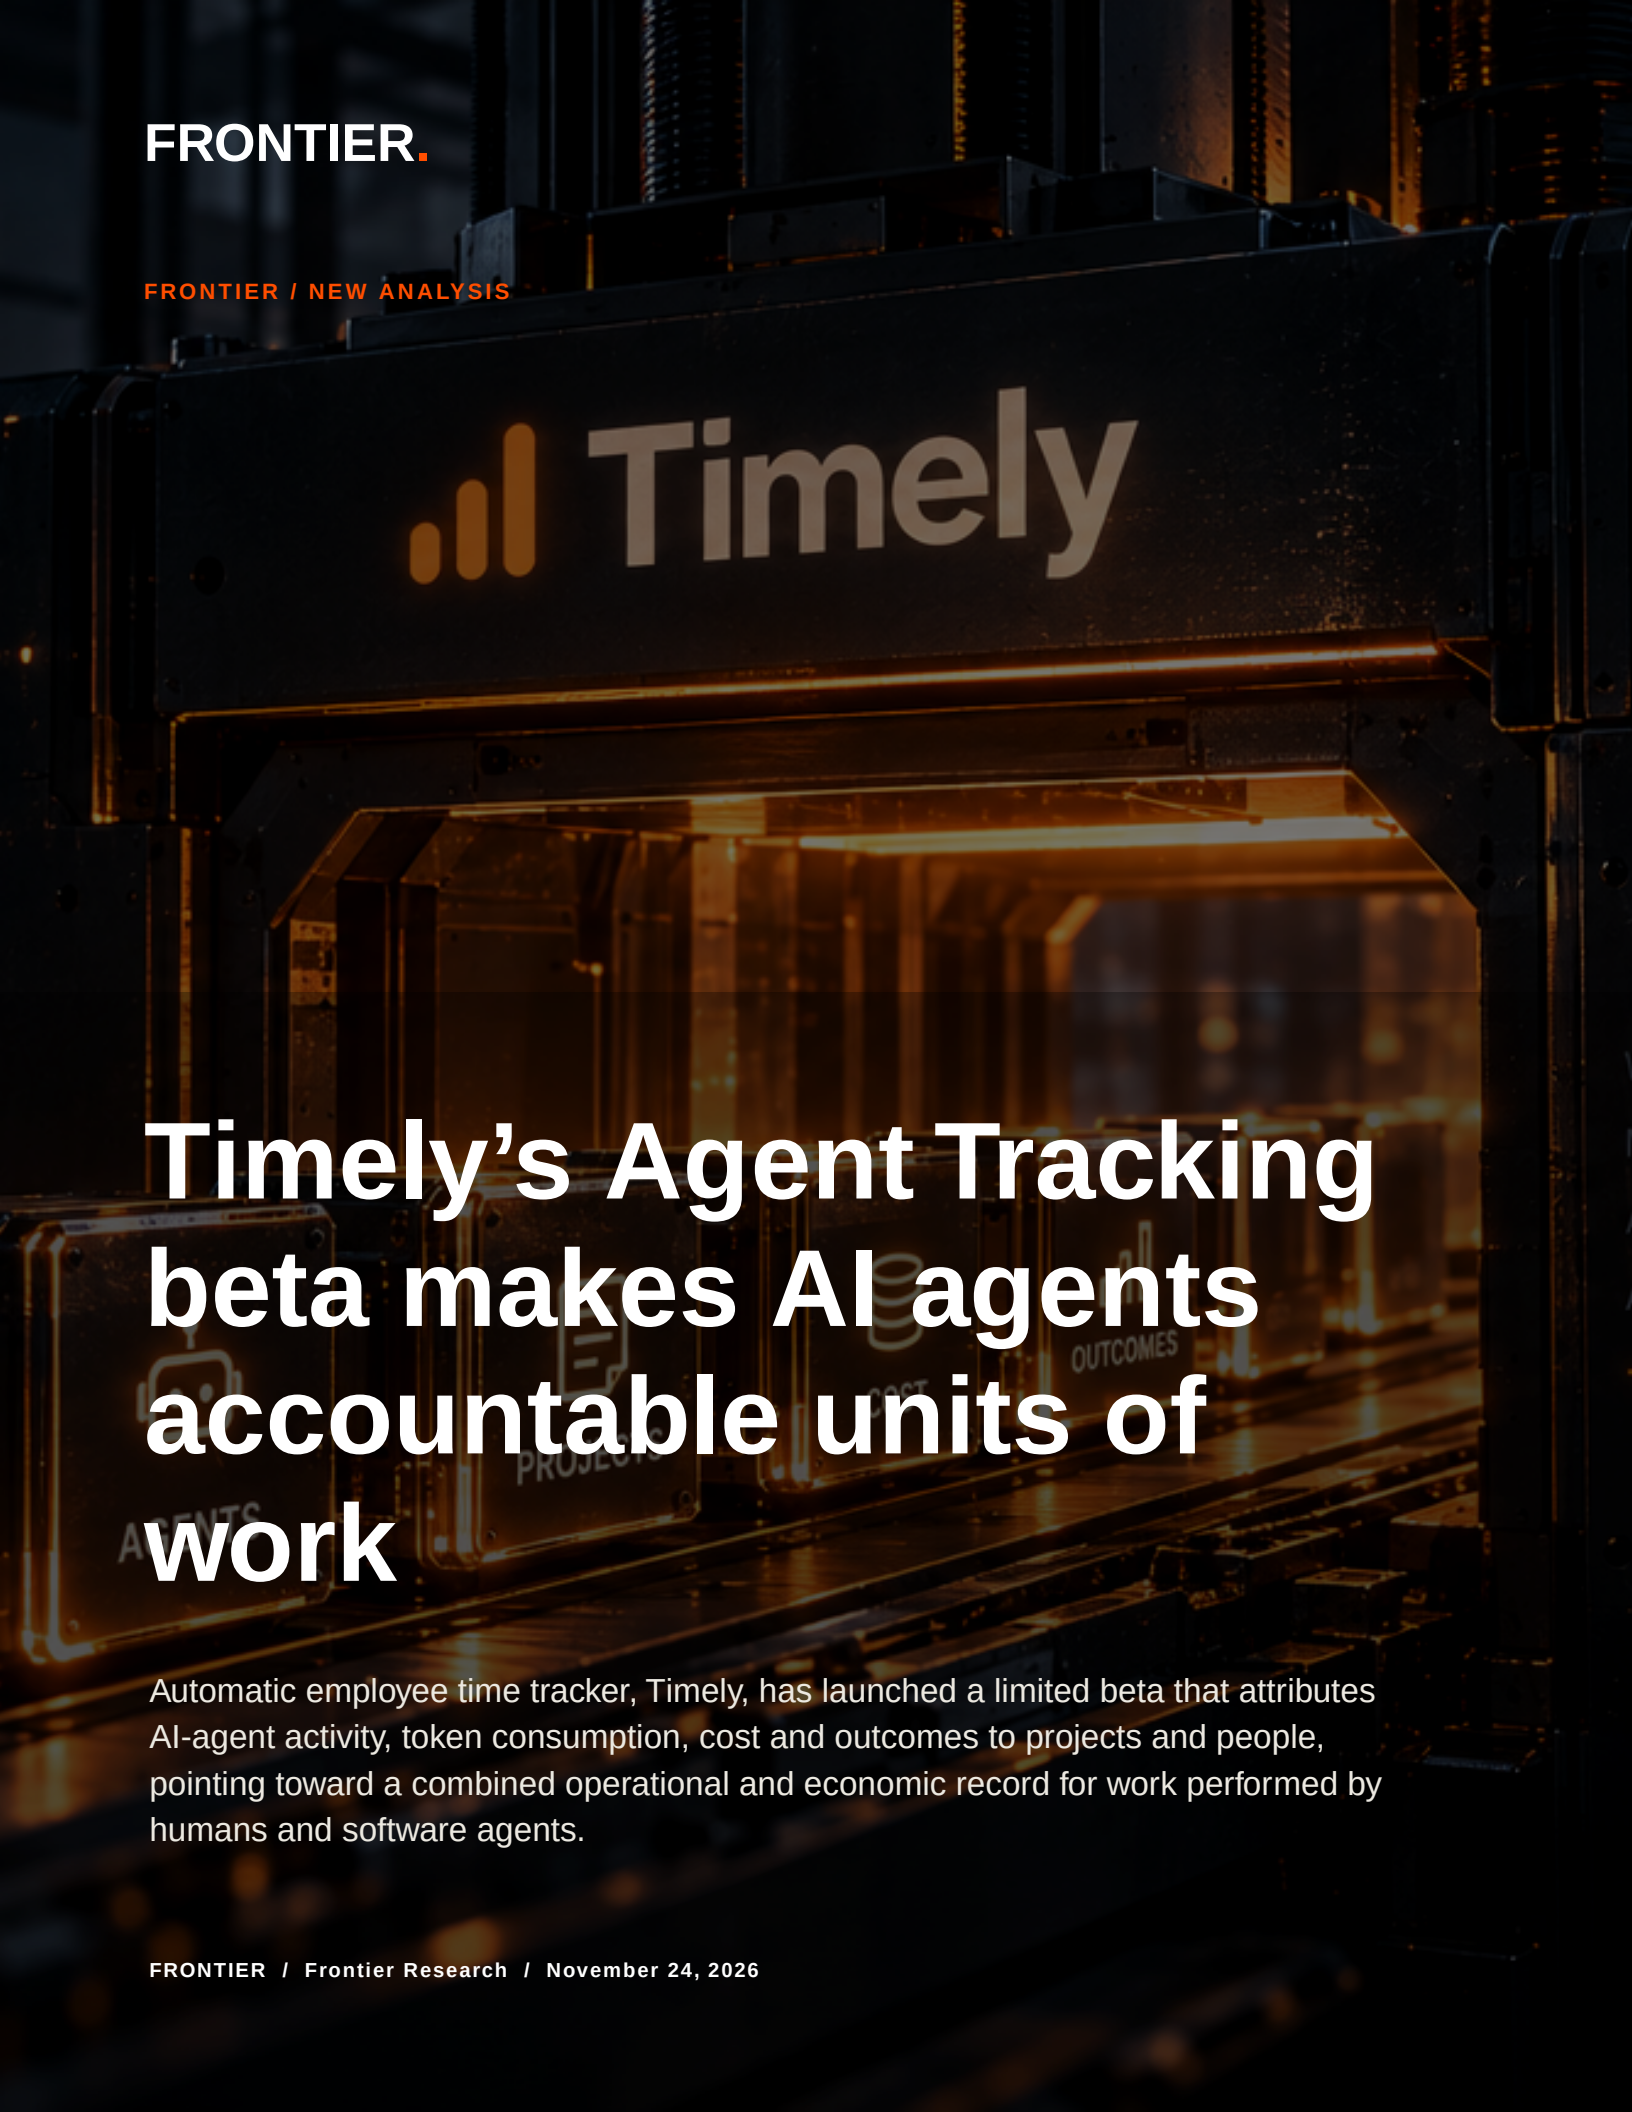

Timely's Agent Tracking beta makes AI agents accountable units of work

Automatic employee time tracker, Timely, has launched a limited beta that attributes AI-agent activity, token consumption, cost and outcomes to projects and people, pointing toward a combined operational and economic record for work performed by humans and software agents.

FRONTIER / Frontier Research / November 24, 2026

NEW ANALYSIS

Built to change a decision.

Automatic employee time tracker, Timely, has launched a limited beta that attributes AI-agent activity, token consumption, cost and outcomes to projects and people, pointing toward a combined operational and economic record for work performed by humans and software agents.

EDITION	**New Analysis**
AUTHOR	**Frontier Research**
PUBLISHED	**November 24, 2026**

This PDF is a publication rendering of the canonical Frontier News article. The web article remains the live source of truth; analysis, figures, and source citations travel with this edition.

What changed

Timely launched Agent Tracking on November 18, 2026, as a limited beta for recording what AI-agent sessions accomplish, how much they cost and which projects they serve.[1][2] The original event was actually the November 18 launch, with Dealroom's subsequent November 23 coverage making the news more visible to business leaders.[3] Timely is admitting companies a few at a time rather than offering the product as generally available enterprise software.[2]

Timely is a B2B SaaS platform that uses AI to automatically track work hours, generate timesheets, and manage project budgets without manual timers or invasive employee monitoring.

The news is vital to organization adopting agentic AI because product adds an agent-oriented accounting layer to Timely's existing focus on human work. Its interface can aggregate sessions and display activity by time, tokens or cost, organized by model, project or person. Timely also presents agent sessions alongside human-work timelines and links them to tasks in Jira, Linear, monday.com or Asana.[2] That combination is the important development: AI use is no longer represented only as a software subscription or a line on a cloud invoice. It can be associated with a task, an employee and a stated result.

Shifter reported before the formal launch that Timely had developed a background desktop application capable of observing an agent session's duration, token consumption, actual or estimated cost and subsequent work summary.[4] The distinction between actual and estimated cost matters. A measured token count, an inferred monetary value and a reconciled vendor charge are not interchangeable records, particularly when enterprises use multiple models, agent clients or subscription arrangements.

Timely founder and CEO Mathias Mikkelsen told Kampanje that six employees at the company generated NOK 549,000 in AI-agent charges during one month.[1] That figure is a vendor-attributed account of Timely's internal spending, not an independently audited benchmark. It nevertheless illustrates the problem Timely is targeting: agent activity can accumulate material expense across individual work sessions without giving managers a project-level explanation of what the expenditure produced.

The company also says Agent Tracking does not retain prompts or the underlying work and instead records outcomes and usage information.[1] That privacy description has not been independently validated. Enterprises evaluating the beta therefore need to examine the desktop application's capture behavior, data flows and retention controls rather than treating a product-level assurance as a completed security review.

Why it matters

CIOs now face two linked purchasing questions. The first is whether an agent can perform a useful task. The second is whether the enterprise can reliably attribute the agent's activity, expense and output after deployment. Timely's beta addresses the second question by connecting agent sessions to models, people and projects and by allowing the same activity to be viewed through time, token and cost measures.[2] This is an early implementation, not proof that its records are sufficiently accurate for financial allocation or customer billing.

That qualification changes how the product should be evaluated. A team may use estimated costs to spot unexpected consumption or compare patterns between projects. It should not automatically use those estimates for chargeback, client invoicing or investment decisions. Before doing so, the CIO and finance organization need to establish how Timely converts observed activity into cost, whether the amount reflects metered usage or a subscription allocation, and how credits, bundled capacity or other pricing arrangements are represented. Shifter's reporting that the desktop application can present actual or estimated costs makes this validation a central procurement issue rather than a reporting detail.[4]

Project attribution requires similar scrutiny. Linking a session to Jira, Linear, monday.com or Asana can create a useful operational record, especially when agent activity appears beside human work.[2] But a session-task link does not by itself establish the value or quality of the result. An agent may run for a long time, consume many tokens and produce an unusable output; another may complete a valuable task with little recorded activity. CIOs should therefore resist treating token consumption or elapsed time as a proxy for productivity. The economically useful record combines consumption with a defined outcome and an accountable owner.

This creates a build-versus-buy choice for enterprises already collecting model telemetry, cloud costs and work-management data. A packaged product can reduce the work required to join those records and may provide visibility at the employee desktop, where agent use occurs. An internal implementation can preserve control over cost logic, data boundaries and integration design. Timely's limited admission model means buyers should assess it as an early-access observability and project-accounting product, not assume it can replace established financial, security or operational systems today.[2]

The security tradeoff is equally concrete. Timely says it records outcomes and usage rather than prompts or underlying work.[1] Even under that stated design, outcome summaries, employee associations and project links can reveal information about what a team is doing. The appropriate capture boundary will vary by workflow. A company may permit observation of development tasks while excluding sensitive client matters or other restricted work. The decision should follow a technical review of what the application observes and transmits, not merely a broad judgment that agent monitoring is desirable.

The larger structural change is that AI spending can increasingly be assigned to work rather than left as an undifferentiated technology expense. Once activity is attributable, CIOs can ask which projects consume agent capacity, which models generate the cost and which people initiate the sessions. Timely's product page explicitly organizes its reporting around those dimensions.[2] The opportunity is better allocation and governance; the risk is false precision if estimates, task mappings or outcome descriptions are accepted without reconciliation.

Frontier take

Agent observability will become financially useful only when it records an accountable chain from consumption to project to outcome. Timely's beta is notable because it attempts to assemble that chain across agent sessions and human-work timelines, rather than stopping at model-level token monitoring.[2] Its strategic importance lies in the accounting model it proposes, not in the maturity of this particular release.

That model treats an agent as a participant in enterprise work. The agent consumes resources, operates on behalf of a person, contributes to a project and produces a claimed result. Those elements can support budgeting and portfolio decisions when the underlying records are reliable. They can also produce misleading allocations when estimated cost is presented with the apparent authority of a settled financial amount or when an automated summary is treated as evidence that useful work occurred.

CIOs should therefore separate observability from accounting authority. Agent Tracking may provide a valuable operational view during its beta, including session aggregation and reporting by model, project or person.[2] The enterprise should grant that record authority over chargeback or client billing only after reconciling it with provider charges, testing project attribution and defining how outcomes are accepted. This preserves the visibility benefit without allowing an immature measurement layer to become a shadow financial system.

The same separation applies to workforce analysis. Placing human and agent activity on a shared timeline can help managers understand how work is divided.[2] It does not make an hour of human effort equivalent to an hour of agent runtime, nor does it establish that either produced value. The useful comparison is the cost and accepted outcome of a workflow under different operating designs. Time and tokens are inputs to that analysis, not its conclusion.

Timely's own reported experience underscores why enterprises will pursue this capability. Mikkelsen's claim of NOK 549,000 in one month of agent charges across six employees has not been independently audited, but it describes the kind of spending surprise that project-level attribution is meant to surface.[1] Enterprises do not need to accept that figure as typical to recognize the governance gap. As employees place more work through agents, procurement records alone cannot show which initiatives generated the expense or whether the results justified it.

The winning control layer will not merely count tokens. It will preserve enough context to explain who initiated agent work, where the cost belongs and whether the output met a defined acceptance condition, while collecting no more sensitive content than necessary. Timely has put that proposition into a limited beta. CIOs should test the proposition now, but with financial and security guardrails that reflect the product's early-access status.

Three moves for CIOs

1. — Run a shadow agent ledger before authorizing chargeback Use the beta's project, person, model, time, token and cost views as a parallel operational ledger. Reconcile a sample of sessions against provider charges or subscription records, verify the linked project with the responsible employee and require an accepted outcome before the record influences budget allocation. Keep estimated and actual costs visibly separate.
   - Decision trigger: Begin this control when a business unit asks to allocate agent expense to projects, departments or clients, or when recurring agent activity becomes material enough to affect a project budget.
   - Why now: Timely can present actual or estimated costs, but the available evidence does not establish that either will reconcile automatically with every commercial arrangement.[4] A shadow period identifies attribution and pricing errors before an operational dashboard acquires financial authority.
2. — Define capture zones for desktop agent monitoring Map what the application observes, derives, stores and transmits, then approve it by workflow rather than deploying it indiscriminately. Exclude projects whose outcome summaries or task associations would expose restricted information, and document the retention and access rules for employee-linked agent records.
   - Decision trigger: Require this review before installing the background application on endpoints used for sensitive projects or before connecting Agent Tracking to a work-management system containing restricted task data.
   - Why now: Timely says it retains outcomes and usage data rather than prompts or underlying work, but that claim has not been independently technically validated.[1] The limited beta is the appropriate stage to establish capture boundaries before the resulting history becomes embedded in management processes.
3. — Fund agent workflows by accepted outcomes, not token volume Define an acceptance condition for each measured workflow and compare its total cost with the result delivered. Use time and token data to investigate variance, but require project owners to classify whether the agent output was accepted, revised or discarded before increasing capacity or changing models.
   - Decision trigger: Apply this rule when a team requests more agent budget, proposes a different model or cites increased agent activity as evidence of productivity.
   - Why now: Timely's reporting can connect sessions to projects and show time, tokens and cost, but none of those measures alone proves that useful work occurred.[2] Adding an acceptance decision turns observability into an investment signal while limiting the risk that higher consumption is mistaken for higher value.

Sources

- Timely — AI Agent Tracking product and beta-access page [2]
- Kampanje — launch-day interview and disclosure, published November 18, 2026 [1]
- Shifter — independent report and founder interview, published November 17, 2026 [4]
- Dealroom News — candidate coverage published November 23, 2026 [3]

SOURCES / CITATIONS

Trace the argument.

Every source below is linked. Source numbering follows first appearance in the report.

01 **Kampanje — launch-day interview and disclosure, published November 18, 2026**
https://www.kampanje.com/premium/skal-fa-kontroll-pa-ki-agentene-mye-dyrere-enn-folk-tror/1042563

02 **Timely — AI Agent Tracking product and beta-access page**
https://www.timely.com/ai-agents/

03 **Dealroom News — candidate coverage published November 23, 2026**
https://dealroom.co/news/174916-timely-opens-beta-to-track-ai-agents-time-tokens-and-cost/

04 **Shifter — independent report and founder interview, published November 17, 2026**
https://www.shifter.no/nyheter/vil-la-sjefene-se-hva-ai-agentene-egentlig-driver-med/482831

FRONTIER.

TAKE THE NEXT STEP

Push this further.

Research is useful when it changes a decision. Bring us the question, the constraint, or the argument you want tested.

RESEARCH / BRIEFINGS

inquiries@thisisfrontier.ai

Briefings, advisory work, executive discussions, and research inquiries.

REPRINTS / LICENSING

reprints@thisisfrontier.ai

Reprints, enterprise distribution, excerpts, and licensing.

FRONTIER

www.thisisfrontier.ai

Research, Data Insights, News, and the live Frontier corpus.

All rights reserved. Contact Frontier for redistribution or reprint rights.

SEE FURTHER. MOVE SOONER.