

Workday's Trusted Agent Platform Now Delivers Headless Access to HR and Finance, Letting Customers Build and Run Lawful AI Agents

New Agent Harness Runs Every Agent with Built-in Context, Memory, and Guardrails, So Every Decision Can Be Proven

Agent-Ready Tools, Workday Data Cloud, Developer Agent, Custom Agents, and Agent Harness Are Now Generally Available

LAS VEGAS, Oct. 13, 2026 /PRNewswire/ -- **Workday Rising 2026** – [Workday, Inc.](#) (NASDAQ: WDAY), the trusted agentic platform for HR, finance, and IT, today announced its [Trusted Agent Platform](#) is now headless, enabling customers to build, run, and govern lawful AI agents across HR and finance with the development tools and models they choose. Every agent, whether it's built by Workday, a customer, or partner, works under the same Workday permissions, approvals, and audit trails.

The Trusted Agent Platform is where organizations build and run AI agents that operate within their business rules. It brings together [Workday Data Cloud](#), [Agent-Ready Tools](#), [Developer Agent](#), [Custom Agents](#), and [Agent Harness](#), all generally available today.

Enterprise software is going headless, and AI agents can now do real work in HR and finance, without anyone logging in. That is the risk. Point an agent at an ungoverned copy of HR data in a data lake and it can read compensation it should never see. Connect it through a system-level account and it can act without approval. Either way, no one can prove what it did or why.

"Headless can't mean lawless," said Gabe Monroy, chief technology officer, Workday. "The rules for HR and finance already live in Workday, so customers don't have to rebuild them for every agent they create. This is what lets companies move fast with agents without losing control."

Bring Workday Rules to Headless

Workday is giving customers a safer way to let AI agents and apps work with Workday data wherever work happens, not just inside Workday. With Agent-Ready Tools and [Live Data Query](#), every request carries Workday's rules with it, including who is asking, what they can see and do, which approvals apply, and how the action is recorded.

Agent-Ready Tools are Workday's Model Context Protocol (MCP) tools, giving any agent that supports MCP an approved way to read Workday data and take action.

"MCP gives agents a standard way to connect with the systems where work happens," said Ankur Sinha, head of enterprise and verticals, Anthropic. "Workday's Agent-Ready Tools bring that model to HR and finance, helping customers give agents access to critical workflows while preserving the permissions and approvals their businesses already trust."

Live Data Query lets teams and agents ask questions of live Workday data in SQL through Workday Data Cloud, with no custom reports, fragile integrations, or stale data copies. Workday applies the right security and business rules, so a BI dashboard or an AI agent gets answers from the same trusted source.

Run Every Agent on Guardrails with Agent Harness

Workday [research](#) shows why agents need guardrails. Large language models can misapply HR policies even when the correct policy is in front of them. Better models alone do not solve this. Agents need the context, rules, and business processes that live inside Workday.

Every agent deployed on Workday runs on Agent Harness, which gives each agent the context a task depends on, including who the person is, their role and reporting line, what they can see and do, and which policies and processes apply. That context comes from the system that runs HR and finance, not from a data copy or policy document, so agents work from the facts and rules the business already trusts. Agent Harness also governs each action and routes each task to the right model for cost and performance.

In early 2027, Agent Harness will add Agent Context Service and Grounded Reasoning. Agent Context Service will identify the Workday context an agent needs and provide it before the agent starts reasoning, so agents work faster and handle less sensitive data. Grounded Reasoning will then check requests against company policies, business processes, and tenant data before an agent acts.

For example, if a manager asks an agent to move someone into a role the manager has no authority to fill, the agent will be able to stop, explain why, and point to the right next step instead of making a change that violates policy.

Build Enterprise-Ready Agents on Workday

Developer Agent is now generally available and lets teams turn plain language requests into custom agents and apps using the

domain expertise and AI skills built into Workday.

Customers describe the agent they need using their preferred development tool, including Claude Code, Codex, Cursor, and the Workday Developer Site. A team that wants to match employees to open internal roles, for example, can ask Developer Agent to build an agent that uses trusted job data through Workday Data Cloud, picks the right Agent-Ready Tools, and follows existing Workday security and business processes.

With Custom Agents also now generally available, teams can build their own agents and deploy them on Agent Harness. Teams that prefer their own framework can build with tools like Google ADK and LangChain frameworks. They can also choose which AI model powers each agent, including Anthropic's Claude, Google's Gemini, Meta's Llama, Mistral models, OpenAI's GPT, and now, NVIDIA's Nemotron open models.

"Efficient agents require a system of both open and proprietary models," said Pat Lee, vice president, strategic enterprise partnerships, NVIDIA. "Workday's Trusted Agent Platform gives developers the ability to work with the right model for the job – including NVIDIA Nemotron open models – so they can build optimized workflows they can adapt, govern, and trust."

Availability

- Workday Agent Harness is generally available, powering all agents built by Workday, customers, and partners.
- Agent-Ready Tools, Custom Agents, and Developer Agent are generally available in Workday Extend AI.
- Workday Data Cloud, including Live Data Query and Agent Lake, are generally available.
- Support for LangChain frameworks and NVIDIA Nemotron is now generally available. Support for Google ADK is accepting design partners.
- Agent Context Service is expected to be delivered within Agent Harness in early access in late 2026. Grounded Reasoning is expected to be delivered within Agent Harness in early access in early calendar year 2027.

For More Information

- Read the blog, [Inside Workday's Trusted Agent Platform](#).
- Discover [how to build and run AI agents lawfully](#) with Workday Extend AI's Developer Agent, Custom Agents, and Agent-Ready Tools.
- Learn how [Workday Data Cloud](#) opens the system of record directly to AI agents, so teams no longer have to extract data, duplicate it, or rebuild security.
- Register and attend [Workday Rising](#) sessions in the digital program:
 - Office of the CIO Keynote: The Agent-Ready Platform [26-KEY6], Oct. 13 at 11:00am PT.
 - Workday Data Cloud: Governed Access and trusted contextualized data for Autonomous Agents [26-PSVIT2743], Oct. 13 at 4:30pm PT.

About Workday

[Workday](#) operates at the heart of the enterprise – HR, finance, and IT – where the margin for error is effectively zero. By tightly coupling AI with the context, guardrails, and trusted processes that run the business, Workday goes beyond AI that assists work to agents that are capable of driving measurable outcomes. More than 11,500 organizations worldwide, including more than 65% of the Fortune 500, trust Workday to deliver. For more information about Workday, visit [workday.com](#).

© 2026 Workday, Inc. All rights reserved. Workday and the Workday logo are trademarks of Workday, Inc. All other brand and product names are trademarks or registered trademarks of their respective holders.

Forward-Looking Statements

This press release contains forward-looking statements including, among other things, statements regarding Workday's plans, beliefs, and expectations. These forward-looking statements are based only on currently available information and our current beliefs, expectations, and assumptions. Because forward-looking statements relate to the future, they are subject to inherent risks, uncertainties, assumptions, and changes in circumstances that are difficult to predict and many of which are outside of our control. If the risks materialize, assumptions prove incorrect, or we experience unexpected changes in circumstances, actual results could differ materially from the results implied by these forward-looking statements, and therefore you should not rely on any forward-looking statements. Risks include, but are not limited to, risks described in our filings with the Securities and Exchange Commission ("SEC"), including our most recent report on Form 10-Q or Form 10-K and other reports that we have filed and will file with the SEC from time to time, which could cause actual results to vary from expectations. Workday assumes no obligation to, and does not currently intend to, update any such forward-looking statements after the date of this release, except as required by law.

Any unreleased services, features, or functions referenced in this document, our website, or other press releases or public statements that are not currently available are subject to change at Workday's discretion and may not be delivered as planned or at all. Customers who purchase Workday services should make their purchase decisions based upon services, features, and

functions that are currently available.

SOURCE Workday Inc.

For further information: Investor Relations: ir@workday.com, Media Inquiries: media@workday.com

<https://en-au.newsroom.workday.com/2026-10-13-Workdays-Trusted-Agent-Platform-Now-Delivers-Headless-Access-to-HR-and-Finance.-Letting-Customers-Build-and-Run-Lawful-AI-Agents>