

# Workday's Agentic Framing

## The Software Debate Through Workday's Eyes

February 27, 2026



Periods of confusion can be valuable for the curious. Confusion can cause those with deep knowledge in the area to start talking and sharing more. And the good kind of talking and sharing where they're trying to clarify their thoughts in real time.

We are in one of those moments right now. AI continues to progress. OpenClaw and Claude Cowork marked another ChatGPT-type moment. Agentic advancements led to new levels of confusion for many, particularly in software. When behaviors suggest that nobody fully understands, you sometimes get more candor, deeper explanations, presentations, and interviews where management teams wrestle with the question out loud.

There have been several recent earnings calls where management has attempted to cut through the noise and clarify why they believe their respective software companies are well positioned for the future. But Workday's was the cleanest we've seen in framing the current software debate and boiling it down to what matters most.

It started with Co-Founder Aneel Bhusri who is now back in the CEO seat. He said:

**You've all heard the narrative out there that HR and ERP will be replaced or relegated to the background by AI. I personally just don't see that happening. Our application domains are really, really hard to build. I've been working in the HR and ERP space for over 30 years. These are true systems of record that must process transactions with absolute accuracy and speed, enforce complex security models, and comply with statutory and regulatory requirements all over the world. That kind of complexity is very hard to replicate. No amount of vibe coding is going to produce an HR or an ERP system.**

This framed a couple of key questions within this debate: *Just how hard is it to build systems of record? And how much does this type of domain expertise matter?* In Workday's case, the reality is these systems must process payroll accurately to the penny across many countries with different tax codes, labor laws, and regulatory requirements. They must enforce security models that determine who can see what data (e.g., manager can see their direct reports' compensation but not their peers'). They must comply with statutory requirements that change constantly across jurisdictions. They must handle benefits administration, financial close processes, audit trails, and regulatory reporting with zero tolerance for error. The complexity is in the interaction of thousands of rules, exceptions, and compliance requirements across geographies and industries, accumulated and refined over decades.

The pushback tends to be: *Even if you can't vibe code this type of system because it is too complex for that today, is it just a current limitation?* When Aneel left PeopleSoft to start Workday, incumbents made similar arguments about why cloud-based systems couldn't handle the complexity and security requirements of enterprise HR and finance. They were wrong because the new architecture eventually caught up on the hard parts while delivering advantages the old architecture couldn't match. The question is whether AI follows the same pattern on a compressed timeline or whether the deterministic requirements of these systems represent a harder barrier than the cloud transition did. At the very least, it's probably a multi-year effort.

The next question in this debate is: *well, then what does better software look like?* We all know by now that AI's progress means the software we use will allow us to be far more productive. If agents can't fully take over, at least for a while, what does progress look like from here? Aneel said:

**You can't have probabilistic outcomes in running a payroll. It needs to be 100% accurate and completed 100% of the time. So, what is the future? It's the marriage of deterministic enterprise apps with probabilistic AI that leads to three things: a redefined user experience that is prompt based; a greatly improved business process automation execution platform with work done by agents and humans; and lastly, much deeper AI-generated insights.**

Now he starts getting into the role of agents, which covers the next section of the current debate. It's widely agreed that both the system of record / truth underneath and agents that help improve user experience and business process automation are important. But whose agents can do this and whose will work best? Particularly for agents that are more focused on user experience and will go across apps like Workday HCM and Financials, will it be first-party agents from Workday or third-party agents from a Claude Cowork-type model that win?

Annel said:

**Taken together, these changes will lead to much better business outcomes and deliver far higher ROI for our customers. As you'll hear from Gerrit, this hybrid world and architecture is exactly what Workday is building today, marrying the best system of record for HR and finance with the reasoning capabilities of domain-specific LLMs. That's the future. It's like peanut butter and jelly. They just go together. And it's all rooted in the trust we've built with our customers over the past 20 years from managing their systems and their most critical data.**

**All you have to do is look back at the past year to know that this hybrid world is coming to fruition. We experienced accelerating agentic AI adoption across the solutions we acquired, HiredScore, Evisort, Paradox and most recently Sana. They all had great Q4s, and I'm pleased to say that Sana is now seamlessly integrated into the Workday stack. It's very cool.**

**So, as we move into fiscal year 2027, what should you expect from us? Our core enterprise apps in HR, finance and student will continue to thrive. Recognizing that in some areas like HR, we have high levels of market share, so the growth will come from new apps sold to the customer base.**

**The accelerating growth will come from the agents we built on top of Workday, both inorganic and organic. If FY**

And Gerrit Kazmaier, Workday's President of Product & Technology, said:

**Today, we have 12 new organically developed role-based agents that are now starting to move into general availability for all our customers. And more than 400 customers are already using them and get real ROI. With our Self-Service Agent, early access customers have reduced HR case volume by 25% and increased their employee productivity by 20%.**

**And with our R1 release in March, in addition to our Self-Service Agent, Planning Agent, Deployment Agent, Payroll Agent, and Business Process Optimize Agent are all entering into general availability for all our customers.**

Clearly Workday believes first-party agents will play a major role and is all-in on winning the agents that can both automate work within its apps and go across multiple apps. And Workday believes that executing here will allow for *accelerating* growth.

Being able to articulate a plan for *accelerating* growth is central to the debate on whether the current pacing around the ~\$200bn (and growing) in annual capex from Amazon, ~\$185bn from Google, etc. still makes sense. We likely need to see widespread and effective consumer *and* enterprise adoption of agentic workflows (due to the higher

levels of productivity it can unlock) within the next 12 months to avoid valid “plateau” concerns. This will require more companies like Workday (or disruptors) to demonstrate that enterprise software spend can accelerate due to winning AI agent product releases.

Next, Gerrit added:

And here's another real highlight for us, the release of Sana Core and Sana Enterprise into general availability. Sana Core brings conversational AI directly into Workday, so employees can get answers and complete tasks using Workday agents. Sana Enterprise extends this same experience beyond Workday by connecting enterprise tools like Outlook and Google Drive and so many more, so employees can find information, create agents, and automate workflows all in one place. That will create an entirely new way for our 75 million users to engage with Workday.

Sana Core and Sana Enterprise went into GA on February 15, which we executed with world-class speed from project start to finish in just three months. We used AI to dramatically accelerate our product innovation cycle

Workday is hoping its customers adopt its agents not just to find information and automate workflows within its own apps but also across third-party apps as well. Why do Workday agents have a right to win workflows that involve third-party apps? Why don't Claude and other frontier labs instead have the right to win both the workflows that go across apps and within apps like Workday? This is important.

Gerrit said:

finance processes. So, think about a compliance autopilot, which continuously tests and verifies ledger and transactions, always-on analysis of all transactions that are coming in, and then even optimizing them with our document and contract intelligence, for instance, for real-time bookings, making sure they are set at the right value at the right discount rates, if you're a supply chain.

Workday should have the right to win these types of agentic workflows if only its agents have access to the underlying data. But could Claude (and others) do the same or do it better and get there quicker IF they had access to the same underlying data? Workday is incentivized not to give third party agents this type of access, but are new entrants that start with a clean sheet of paper incentivized to provide the frontier labs with full access? If so, how long would it take new entrants to build up the required domain expertise and functionality at the system of record level?

Frontier models like Claude are already exceptionally good at reasoning, language understanding, summarization, and working through complex logic. If you gave it a well-structured prompt with the right context, it could draft a performance review, analyze a financial close checklist, or help plan a workforce restructuring. It could reason about HR and finance problems at a high level. The gap Workday refers to throughout this earnings call is about **access, context, and authority**.

**Access:** Workday's agents operate inside a system that already contains the actual employee records, compensation data, benefits elections, org charts, financial transactions, journal entries, and payroll runs for a given company. Claude sitting outside that system would need to be given all of that data every time it is asked to do something, which raises security, compliance, and integration challenges. Workday's agents don't need to be given the data.

**Context:** Gerrit described 70,000 core business process archetypes instantiated across thousands of customer variations. That means Workday has both the data and the encoded logic of how companies run their HR and finance operations. When a Workday agent handles a financial close process, it understands the specific sequence of

approvals, the compliance rules, the journal entry structures, and the exception handling procedures for that particular customer's configuration. A general-purpose model would need all of that context reconstructed from scratch. It would be reasoning about HR and finance in the abstract rather than operating within the specific reality of how a given company works.

**Authority:** A Workday agent can execute. It can change an employee's benefits election. It can post a journal entry. It can move a candidate through a hiring pipeline. It has permissioned write access to the system of record. Claude can tell you what should happen. Workday's agent can make it happen, within the security and approval frameworks the company has already configured.

What is not commoditized is the data substrate, the process logic, and the permissioned execution layer. Those are the pieces that have taken 10+ years to build and that customers cannot easily rip out and replace.

The risk, of course, is Claude and other general-purpose AI platforms ultimately being “pulled in” and connecting to those data and execution layers from the outside. If Claude or a similar model could securely access Workday's data, understand a company's specific process configurations, and execute actions with proper permissions, the value of Workday's proprietary agent layer would erode. Another wrinkle is the legacy platforms that Anthropic, OpenAI, etc. would be disrupting are their customers. The general-purpose AI platforms must be mindful of competing with key customers, particularly prior to having a clear line of sight into how they would obtain the access to the data substrate and process logic they'd need to offer winning agents within these domains. These questions also came up in Workday's Q&A:

**Gabriela Borges**

*Goldman Sachs Group, Inc., Research Division*

Aneel and Gerrit, I wanted to link together a couple of the points you're making. This idea of AI being built on top of cloud and then the potential for Workday to monetize any other vendor or corporate built tool that extracts richer information out of Workday. My question to you is how do you manage the risk that vendors or customers build solutions next to Workday that leverages all the domain experience that you've built, but ultimately, Workday becomes part of this biggest stack such that the incremental value is being accrued in the intelligence way, which would be outside Workday? Maybe just talk to us a little bit about how you manage and mitigate against that type of risk?

**Aneel Bhusri**

*Co-Founder, CEO & Executive Chairman of the Board*

Well, I'd first of all say that the API layer that ultimately will become the same with the agentic layer, whether we have the apps on top or a third-party have apps on the top, it's our intelligence layer. It's our metadata, it's our security model, it's our data model. It's our business process framework, and people underestimate the power of our business process framework. There's 20 years of knowledge into these HR and financial applications that even with the best AI, it's going to take 5, 6, 7 years to replicate, and we're not standing still. So we want to get to a model where we're pretty much indifferent if it's third parties or us building the apps on top, we get compensated either way through the consumption model. Gerrit?

**Gerrit Kazmaier**

*President of Product & Technology*

Gabriela, I think it's a great question. So it allows us to really cut through the chase, I think. Look, the reality is, the question is where this intelligence system is being built and what's driving that? And from what we see, I think what research shows today that the idea that you create just a big model and connect it to a database of data and it just figures **everything out is utterly unviable**, right? It just creates a completely lawless AI which has no understanding of enterprise processes, compliance, determinism, repeatability. So you basically have to build an HR and finance system for AI, right? So that's the intelligence layer, right? It's not bimodal, smart model and dump data. It's actually about an AI-enabled ERP. And that's the intelligence layer that Aneel spoke about that we are building.

Workday's responses highlight the core obstacle for disruptors to win these types of domain specific agentic workflows today. LLMs like Claude are trained in general knowledge. They know broadly how payroll works, what tax withholding is, what FMLA leave entails. But it does not know that Company X's policy is that vacation days roll over up to 40 hours in California but expire entirely for employees in Germany, that the rollover resets on the employee's hire date anniversary rather than the calendar year, that employees who transferred from the London office retain their UK accrual schedule for 18 months under a legacy policy, and that all of this interacts with a union agreement that overrides the standard policy for hourly workers in three specific facilities. That is one company's vacation policy. Multiply it by every HR and finance process across over 10,000 customers, each with their own configurations, exceptions, union agreements, regulatory jurisdictions, and historical policy changes. That is what Workday's business process framework encodes.

There's a structural barrier today that prevents general-purpose AI labs from replicating enterprise systems like Workday. It lies in the distinction between raw data and institutional logic. While an LLM can process data points, it lacks the "Business Process Framework" required to interpret them within a stateful, deterministic environment. In addition, replicating Workday's integrated domain security model is a big undertaking; a third-party AI "dumping data" into a model cannot inherently understand the complex web of permissions that prevent a peer from viewing a manager's salary, for example.

For an LLM to handle this well, it needs access to the company-specific rules and configurations. Today these live inside Workday's system. An LLM would need structured access to all of this configuration data for a given customer. It would need the rules engine: which policies apply to which employee populations, what the approval chains look like, what the exception handling logic is, how different jurisdictions interact. This is technically solvable. If a company exported or exposed its full configuration logic through APIs, an LLM could ingest it. The question is whether that configuration logic is currently accessible outside Workday's system in a structured, complete way. For most customers, it probably isn't. It lives inside Workday's proprietary framework. But it's not impossible to extract or reconstruct. And perhaps in a world of agentic "screen-reading" or advanced ETL (Extract, Transform, Load) tools, AI can "watch" how a company handles its German vacation rollovers for six months without needing the formal API to reconstruct the rule-base. But then it would need to do the same for the next customer, and so on, all while these customers already have software that does the job.

Then there's deterministic execution within probabilistic reasoning that was discussed. Payroll cannot be 99.7% accurate. But this is a framing problem as much as a technical one. The solution isn't to have the LLM guess the payroll amounts. The solution is to have the LLM orchestrate deterministic computation. The model reasons about which rules apply, which tax tables to use, which deductions to calculate, and then calls deterministic functions that execute the math with precision. This is already how tool use appears to work in AI systems. The model decides what to do, and structured tools execute with exactness. The model doesn't calculate the tax withholding itself. It calls a tax calculation function with the right parameters. This architecture, which is probabilistic reasoning orchestrating deterministic execution, seems to be advancing rapidly. It's essentially what Workday itself is building with its agents. The question is whether this architecture requires Workday's specific process framework or whether it can be built on top of a well-structured alternative.

Then there's regulatory and compliance knowledge. Payroll in Germany follows different rules than payroll in Brazil. Tax withholding in California differs from Texas. Statutory reporting requirements change constantly. Workday maintains this knowledge through continuous updates from a large team of domain experts who track regulatory changes across dozens of countries and update the system accordingly. For an LLM to replicate this, it would need

either continuously updated training data on regulatory changes across all relevant jurisdictions, or access to structured regulatory databases that are kept current. Neither exists in a comprehensive, machine-readable form today. But companies like Thomson Reuters, Wolters Kluwer, and government agencies maintain much of this information. If that information were structured into machine-readable regulatory APIs, which is a direction the industry could be moving, an LLM could consume it. This could be the layer where the timeline is longest because it requires not just AI capability but an ecosystem of regulatory data providers making their content available in formats AI can reliably consume.

Then there's the testing and auditability requirement. Enterprise systems need to prove they got the right answer. Every payroll run needs an audit trail. Every financial close needs documentation of which rules were applied. Regulators need to examine the logic. This means the AI system can't be a black box. It needs to show its work in a way that satisfies auditors and regulators. This is a real constraint for current LLMs, which don't naturally produce auditable reasoning chains that map to specific regulatory requirements. But it's a solvable engineering problem. If the system is designed so that the LLM's decisions are logged as structured rule applications like "applied German tax table 2026-Q1, line 43, to employee classification B2", auditability becomes possible.

So how would Workday's intelligence layer stop mattering? You'd need all these things to converge. Full access to company-specific configuration logic, or a new system that captures it from scratch. Reliable orchestration of deterministic execution by probabilistic models. Comprehensive, machine-readable regulatory databases that stay current. And auditable AI reasoning that satisfies enterprise compliance requirements. Each of these is individually solvable. None of them is solved today at the scale and reliability required for enterprise deployment. And they all need to work together seamlessly, which is a systems integration challenge on top of the individual technical challenges.

It seems Gerrit is right that this is "utterly unviable" today. The full stack does not exist today. At the same time though, the threat probably doesn't need to be complete to start mattering. It probably doesn't need to handle every edge case across 40 countries on day one. It needs to be good enough for a 200-person company operating in two countries to choose it over Workday at a fraction of the cost. Then good enough for a 1,000-person company in five countries. The disruption would likely start at the bottom of the market and work its way up, which is the pattern Workday itself followed against PeopleSoft and Oracle.

The scenario Workday should be most worried about probably isn't a frontal assault on their largest, most complex enterprise customers. It's new entrants like Darwinbox with MCP integrations, or an AI-native startup we don't know yet, winning SMBs with a simpler, cheaper, AI-first system that handles 80%+ of use cases, then gradually expands the scope of what it can do.

A lot of this comes back to some of the most important pillars of business sustainability like internal talent, agility, and culture of innovation. How fast and effectively can Workday (and its peers) deliver the types of agents that work within specific apps and that can go across its apps in a way that allows thousands of customers to get to quick yeses? Incumbents also face innovator's dilemma issues that could lead to suboptimal product and pricing decisions. That said, Workday's customers would be happy for it to win the agentic layer within its areas of domain expertise because this would be more convenient than having to identify and work with additional third-party software providers. But with all signs pointing to agents being ready to add huge productivity soon, its customers likely won't have much patience for anything short of highly effective agents at affordable prices.

## Legal Disclaimer

DENMARK Capital Management LLC (the “**General Partner**”) is not registered as an investment adviser with the Securities and Exchange Commission (“**SEC**”) or any state securities authorities. The limited partnership interests (the “**Interests**”) in DENMARK Capital Partners LP (the “**Fund**”), are offered under a separate private offering memorandum (the “**Offering Memorandum**”), have not been registered under the Securities Act of 1933, as amended (the “**Securities Act**”), nor any state's securities laws, and are sold for investment only pursuant to an exemption from registration with the SEC and in compliance with any applicable state or other securities laws. Interests are subject to restrictions on transferability and resale and may not be transferred or resold except as permitted under the Securities Act and applicable state securities laws. Investors should be aware that they could be required to bear the financial risks of this investment for an indefinite period of time.

All information contained in or derived from this letter is proprietary to the Fund and the General Partner.

A prospective investor should only commit to an investment in the Fund if such prospective investor understands the nature of the investment and can bear the economic risk of such investment. **THE FUND IS SPECULATIVE AND INVOLVES A HIGH DEGREE OF RISK. AS A RESULT, AN INVESTOR COULD LOSE ALL OR A SUBSTANTIAL AMOUNT OF ITS INVESTMENT.** In addition, the Fund's fees and expenses may offset its profits. There are restrictions on withdrawing and transferring interests from the Fund. In making an investment decision, you must rely on your own examination of the Fund and the terms of the Offering Memorandum and such other information provided by the General Partner to you and your tax, legal, accounting or other advisors. The information herein is not intended to provide, and should not be relied upon for, accounting, legal, or tax advice or investment recommendations. You should consult your tax, legal, accounting or other advisors about the matters discussed herein.

**PAST PERFORMANCE IS NOT INDICATIVE OR A GUARANTEE OF FUTURE RESULTS.** This document may present past performance data regarding prior/other investments, funds, and/or trading accounts managed by the General Partner and/or the Principal. This is presented solely for explanatory purposes. The Fund may face risks not previously experienced or anticipated by the General Partner and/or Principal, and therefore, prospective investors should evaluate the Fund on their own merits.

Certain information contained in this document constitutes “forward-looking statements” which can be identified by use of forward-looking terminology such as “may,” “will,” “target,” “should,” “expect,” “attempt,” “anticipate,” “project,” “estimate,” “intend,” “seek,” “continue,” or “believe” or the negatives thereof or other variations thereon or comparable terminology. Due to the various risks and uncertainties, actual events or results in the actual performance of the Fund may differ materially from those reflected or contemplated in such forward-looking statements. The General Partner is the source for all graphs and charts, unless otherwise noted.

This document may also present “sample holdings” or “case studies” of a type of asset(s) the Fund may invest in or are expected to be typical of its holdings. Such “sample holdings” are not currently holdings of the Fund and are presented solely for explanatory purposes. Prospective investors should not assume that such “sample holdings” will actually be purchased by the Fund when determining whether to make an investment in the Fund.