



eBook

Critical Success Factors to Enterprise AI Adoption



Table of Contents

Introduction: Navigating the AI revolution—from potential to performance	3
1. Prioritizing and Selecting the Right Use Cases	4
2. Determining Business Value to Get Executive Buy-In	5
3. Optimizing Work of Limited AI Resources	6
4. Having High-Quality Data for Accurate and Relevant Results	7
5. Making Sense Out of Agentic AI	8
6. Successfully Passing the Security Evaluation	10
7. Understanding and Mitigating Ethical and Compliance Risks	12
8. Extending and Customizing AI Across the Enterprise	13
9. A Flexible, Scalable Architecture to Innovate Quickly and Efficiently	14
10. Managing Change and Driving End-User Adoption	15
Conclusion: Looking Forward	16


Introduction: Navigating the AI revolution—from potential to performance.

AI, especially Gen AI and more recently agentic AI, is no longer a futuristic concept—it's actively reshaping the business landscape. While the potential for transformation is immense, organizations often stumble when trying to harness AI's power effectively. The path from initial experimentation to realizing tangible business value is paved with significant hurdles.

How do you pinpoint the right AI use cases that deliver measurable results and gain leadership buy-in? How can you ensure your data is of high-enough quality for reliable AI outcomes, navigate the complexities of security and ethical risks, and manage the scarcity of specialized AI talent? How do you make sense of emerging technologies such as agentic AI, build a scalable architecture for rapid innovation, and, crucially, drive user adoption across your workforce?

These aren't just minor roadblocks; they are critical challenges that can stall AI initiatives, erode trust, and prevent companies from capturing the promised benefits. This is where Workday comes in. Built on a foundation of the world's largest and cleanest set of HR and finance data, Workday approaches AI not as a standalone technology but as an integrated capability designed to solve real business problems within the natural flow of work.

Discover how a focus on measurable value and responsible AI helps organizations like yours overcome the top 10 challenges to AI adoption—transforming potential into proven performance and innovating with integrity.



1. Prioritizing and selecting the right use cases.

The challenge.

According to [McKinsey research](#), 78% of large organizations are now using AI in some capacity. This could be developers using a copilot or the whole company using a Gen AI product such as ChatGPT or Google Gemini. One could argue there's more value in organizationwide adoption but the tangibility of that value really depends on the adoption of specific use cases. For example, everyone can use Gen AI to help draft internal emails, but value is more easily measured (and captured) when automating or assisting specific business processes, such as annual budget and resource planning or quarter-end close.

So the challenge is less about how to select the most broad-reaching AI solution and more about selecting the right use case to prove the value to leadership teams that AI can deliver real results. Choosing the wrong use case can result in a mistrust of AI value, leading to a loss of support and funding for AI solutions in the future.



What needs to be addressed.

Starting small in a low-risk environment that supports experimentation and lessons learned is a best practice with any new technology or solution and is particularly important if you want to deploy several AI solutions across the organization:

- **Solve a real business problem:** Is there evidence of a real problem in a business process or user experience? Do you have quotes from employee interviews or supporting data from surveys or systems?
- **Clear and measurable impact to strategic goals:** Is it obvious how it can positively impact top-level goals and can that value be measured with confidence?
- **Aligned to a well-defined process:** Is it automating repetitive tasks or an existing process, or are you creating new behavior or introducing a new user interface?
- **High feasibility:** Are you confident that it can be developed and deployed without significant resources or investment?
- **Fast time-to-value:** Will it take days or months to get into production and then be quickly adopted and utilized by end users?
- **Simplified change management:** Do you need to train more than a single department or team? How much guidance needs to be provided before using this AI feature for the first time? Is it in the natural flow of existing work?

How Workday can help.

Workday collects feedback on AI features early and often directly with our customers. We don't focus on building generic AI capabilities that are disconnected from specific use cases, nor do we release AI features only to keep up with the hype. We believe that de-prioritizing low-impact features enables us to release the right ones that maximize customer value. Workday strives to develop AI that solves real business problems, with a user-centric experience in the flow of work to drive more natural end-user adoption.

With Workday AI Foundry, we engage with design partners as new AI features are in early development, prior to our Early Adopter program. Through this, the user experience is designed (and redesigned) to delight and make it easy for diverse groups of users to adopt without needing to become AI experts.

2. Determining business value to get executive buy-in.

The challenge.


It's no secret that AI, particularly with recent innovations in Gen AI and agentic AI, has the potential to revolutionize the way we work and deliver transformational business value. The biggest question on all our minds is: Once I've selected the business challenge to solve, how do I estimate and prove the value of AI to key stakeholders before purchasing or deploying it?

Nascent technologies, such as AI, can make this difficult for a variety of reasons, including lack of real-world evidence, the rapid rate of change and innovation, concerns about end-user adoption, data-quality issues, and time-to-value, just to name a few. While it's necessary to address each of these in order to estimate the expected value of a particular AI solution, navigating this alone within your organization is an unrealistic burden that could either slow approvals down among competing AI priorities or lead to disappointing results amidst inflated expectations.

What needs to be addressed.



It's easy to build a value calculator based on theoretical assumptions, but having real-world customer value metrics to back it up will give you and your leadership team confidence in not only the potential value but also in the expected value that's already been realized by other organizations.

Executive stakeholders will undoubtedly ask for the timeline of that expected value, with clear preference for "ASAP." The shorter the project, the better. Various costs and time-delaying activities include evaluation, development and configuration, testing, data migration or transformation, implementation, and ongoing support and maintenance. The overall value and time-to-value can be eroded depending on the complexity and resources needed to accomplish all of this.





\$306K incremental AI-driven savings per 100,000 invoices.

Impact on cost per invoice.




\$256K incremental AI-driven savings per 20,000 cases reduced.

Impact on total cost of case volume.



\$3M incremental AI-driven savings per 100 employees retained.

Impact on top talent turnover rate.



How Workday can help.

Workday begins calculating value metrics with customers as part of the development process through our Design Partner Group and Early Adopter programs and refining those calculations as AI features are more broadly adopted. We then work with existing customers to confirm tangible value across multiple use cases and processes to help other organizations bolster their evidence-based business cases with real customer statistics across industries. Those value metrics can easily be translated into net present value (NPV) that considers time-to-value without having to consider lengthy implementations or extra complexity of getting AI capabilities up and running.

3. Optimizing work of limited AI resources.


The challenge.

AI, particularly Gen AI and agentic AI, is a new technology with a talent pool that's racing to try to keep up with the rapid pace of innovation. The most cutting-edge, resource-abundant tech companies may not have an issue with this, but for most of us, we do not have the resources to hire or even reskill our technologists to become AI experts to build, test, and deploy new AI solutions across domains.

Even as the majority of organizations plan to purchase off-the-shelf software, trying to evaluate the strategic approach of a vendor, various AI models, and the underlying core technology supporting the AI takes significant experience and knowledge that are unrealistic for many to attain in the short term.

What needs to be addressed.

Assuming that most organizations are not able to obtain AI talent quickly, relying on vendor-developed solutions to reap the benefits from AI is the most likely path. First, evaluating the AI needs to be easy. This means transparency into how the technology works, details into the AI models themselves, and an easy way to get this information into the hands of various evaluators. Then there needs to be limited AI technical expertise required to get the solutions configured and deployed. With both of these angles addressed, your limited pool of AI talent is freed up to take on a larger number of AI projects, particularly ones that are unique to your organization and require in-house development or extensive customization.



How Workday can help.

Workday provides transparent and detailed documentation on every generally available AI feature through our AI fact sheets that customers can find on Workday Community. We also have recorded webinars and sessions that provide “under-the-hood” details and visuals on our AI architecture and how AI models work, making it easy to understand and evaluate how we approach AI.

Additionally, our customers do not need AI expertise to enable AI capabilities in Workday because they are embedded into our core platform and utilized within the natural flow of work across HR and financial processes. With minimal (optional) configuration, admins can simply toggle AI features on and off from a single dashboard. This provides complete customer control over which features are enabled, as well as granular transparency into exactly which data is used for each individual feature.

4. Having high-quality data for accurate and relevant results.

The challenge.

If you have used AI, you have likely experienced limitations that impact the accuracy and usefulness of its outputs. This is because of low-quality, disparate data sources. The internet is a great data source for the average consumer use of Gen AI, but LLMs that rely solely on this source struggle from data that is completely unstructured, not tailored for a specific organization or use case, and lacking broader business context—especially outside of one’s organizational domain.

Without high-quality, clean enterprise data, AI can fail in generating quality outputs and insights, leading to potentially disastrous outcomes if you are making critical business decisions from these flawed results.

What needs to be addressed.

To achieve the greatly promised AI results, enterprise data needs to be:

- **Structured:** Data that has a consistent structure reduces the “guessing” that AI needs to do to determine whether each individual data point across a large dataset belongs to a specific attribute. Nonstructured data may add in irrelevant data attributes into both initial AI model training and inference as a request is executed.

Example: If content from both employee feedback and development opportunities is not structured and you want to assess existing skills to develop a learning course focused on a specific skill, there’s a possibility that AI may mix these two sets of data points together and mistake which skills already exist across the organization, rather than which skills are currently being developed.

- **Domain-specific:** Particularly with Gen AI, context and intention are critical to produce quality insights and answers. Multiple industries may use the exact same term to mean a similar concept, but guidance, policies, and best practices might be completely different due to context. In this case, AI is using irrelevant data to inform generated outputs, causing confusion and business decisions based on inaccurate insights.

Example: A company may have an internal knowledgebase with content for people managers with best practices on how to prevent and handle employee attrition. However, if AI looked outside the domain of HR (in this case), it may find a myriad of articles and blogs on how to prevent customer attrition from a marketing or sales perspective and use that information to explain (incorrectly) how to optimize marketing funnels or renegotiate sales contracts—both of which are completely irrelevant for the majority of people managers.

- **Organization-relevant:** AI can leverage industrywide data to generate quality domain-specific results but if it does not prioritize your organization’s data in the output (over other organizations), the decisions made from these insights may not actually be relevant to your particular organization.

Example: Category-specific expenses can look very different across organizations within the same industry. Policies, cultures for spending, and even the type of event (revenue generating vs. internal team building) can vary greatly across organizations. So detecting expense anomalies or automatically categorizing an expense needs to be specific to each individual organization in order for AI to truly deliver value.

- **Timely:** Using stale data leads to insights and outcomes that may have worked in the past but may not be effective today. Data does not always need to stream in real time but does need to be refreshed at an interval that makes sense for an AI use case to deliver the expected value.

Example: Generating a shift schedule may require personnel availability data to be refreshed just once per week, while getting insights into payroll information may only require a refresh cadence of every two weeks (or weekly or more often for some companies). But intranet search may only be useful to employees if knowledge articles are indexed each day or even within minutes of being uploaded.

Additionally, there needs to be business context directly related to the data to tell AI models “the whole story” so insights and content are not generated in a silo. And context indirectly related to the data—for example, context around your internal talent skillsets when making external hiring decisions—only increases the quality of AI outputs that leads to better strategic decision-making.

How Workday can help.

Workday has the largest, cleanest set of HR and finance data in the world, with **800B+** annual transactions produced by our **70M+** end users all in a uniform data model—leading to more accurate, relevant, and organization-specific outcomes. Additionally, Workday understands the business context—**72M+** monthly business process events, including past and present context about business processes across the organization, data, people involved, user information, and the current task at hand, along with retaining conversational context from previous dialogues.

5. Making sense out of agentic AI.

The challenge.

The wave of agentic AI and agents is demanding attention, but the topic and technology are so nascent that both technical and business decision-makers aren't sure how to incorporate agents into their AI strategy or buying decisions.

Understanding how agentic AI differs from other forms of AI and the reality of its maturity can help guide your AI strategy and ease concerns about investing in an emerging, largely unproven technology.

What needs to be addressed.

1. Understanding the technology.

There is a wide spectrum of sophistication and capabilities when it comes to AI agents. Some agent solutions offered are simply rebranded AI-driven experiences that combine techniques such as Gen AI and robotic process automation (RPA), offering a great front-end user experience to kick off the automation of a well-defined process. Other agents may have limited or no user interaction, proactively running in the background to surface insights, provide context to other agents, or automate repetitive user tasks until it reasons that human oversight is needed.

While the true definition of agentic AI requires capabilities such as planning and reasoning, autonomous decision-making, role and goal definition, and iterative learning, an AI agent is defined differently across various solution providers. To prevent analysis paralysis caused by trying to deeply understand the unique architecture of each individual agent across every vendor, organizations should prioritize business value, quality, and safety. This is not to say that AI architecture is not important; in fact, it is critical to achieving significant value across agentic solutions. Yet the architecture of a single agent without context of its underlying AI foundation is not as valuable to make decisions when it comes to agentic AI.

2. Consider the AI foundation.

Agents won't deliver the promised value without a strong AI foundation, including high-quality data, broad business context, intuitive user experiences, extensibility, data privacy and security, and responsible AI. Without this foundation, even the most sophisticated agentic orchestration could return hallucinations, expose sensitive information to the wrong people, exhibit unacceptable bias, or reduce employee motivation to adopt the use of AI.

This AI foundation is critical to the success of other forms of AI such as machine learning and Gen AI, with either the business value or potential risks being amplified when using agentic AI.

3. Managing the eventuality of agent sprawl.

Every vendor in the AI space is actively building agents for every use case imaginable. This will inevitably lead to agent sprawl and a new form of shadow IT: managing hundreds, or even thousands, of agents. This includes data and tool access controls, goals and roles, skills or capabilities, onboarding and offboarding, cross-agent orchestration, and user-group and regional restrictions.

Additionally, when considering agents as part of your workforce that's helping the organization achieve its goals, transparency into performance and goal achievement in a single place is critical to answer strategic questions such as:

- Where are there skills surpluses or deficits across my workforce?
- Precisely where should I invest in AI vs. reskilling more people?
- Are investments in agents performing as expected and delivering business value?
- How can we increase the amount of engaging work vs. mundane/repetitive tasks for people?

How Workday can help.

At Workday, our agentic strategy is to transcend simple task completion to deliver transformative business value, such as increasing recruiter capacity by **54%** or decreasing outside legal spend by **70%** through role-based agents.

Role-based agents are designed to augment humans and transform how work gets done by understanding the breadth and nuances of a specific role within a complex organizational model to replace work, not people. Acknowledging that agents will have autonomy, we are committing to building agents responsibly, ensuring humans stay in the loop for key business decisions through transparent human-agent collaboration.

Workday Agent System of Record is the single command center for managing agents, including those extended by customers and built by third parties, with transparency into the business impact of agents to optimize workforce investments.



6. Successfully passing the security evaluation.

The challenge.

When it comes to enterprise AI, everyone is worried about data leakage. We can't let sensitive information fall into the wrong hands, whether it be to a public-facing LLM or even internally to people who don't have the appropriate privileges and access.

But evaluating AI from a security perspective is no easy task. Like the cloud 15 years ago, AI is a new technology where there may not be existing standards and evaluation tools. Creating even more complexity is the difficulty and obscurity in understanding individual AI models, including those powered by third parties. There is no one-size-fits-all security approach to AI, even within the same vendor platform.

What needs to be addressed.

Understanding the sensitivity of both the use case and data can largely determine many of the security requirements when it comes to evaluating the security of AI models. Outside of factors that are needed for every use case, such as data encryption and regional data residency, there are three primary security approaches that will help simplify the evaluation of a particular use case or feature.

Fine-tuning on shared large models: Nonsensitive public data.

Comingling data with other customers for training AI models may sound risky, but this is not the case when it involves only data that is published publicly, such as using Gen AI to create job descriptions. A large model that is shared across tenants is acceptable in this scenario, with fine-tuning to generate high-quality job descriptions tailored to your specific organization. Of course, you want to ensure that the user prompt, which could contain sensitive information, is not used to retrain the model and that the output (the job description in this case) is not used to train that model until it is published publicly.

Company-specific models: Sensitive data, relevant to only your organization.

Other use cases, such as those involving sensitive, nonpublic data (for example, detecting anomalies in financial data) require complete isolation from other customers and the public. Additionally, there is no value in having a shared model, as this type of data and resulting insights are completely unique to only a single organization—even similar-sized companies in the same industry. Every organization must have its own smaller, fine-tuned model to achieve both the security and the business value of these particular use cases. These organization-specific models are only trained on data within that same tenant, due to the complete isolation across models and tenants.

Shared large models without full retraining (e.g., RAG): Sensitive data that benefits from a fine-tuned model.

For AI use cases with sensitive data that produces more accurate and high-quality results using an LLM fine-tuned for a specific purpose—such as answering questions about HR policies specific to a particular individual—a technique called retrieval-augmented generation (RAG) can navigate the complexity of multiple unstructured datasets to produce relevant, quality outputs.

In most cases, an open LLM is fine-tuned with public, nonproprietary data for a specific purpose to create a shared model. The data fields used for training of this model are both minimized and de-identified. Then a vector database is created with only tenant-specific data to augment the fine-tuned LLM. At runtime, when a user enters a prompt, it goes through that vector database, then through the fine-tuned LLM to produce a high-quality generated output that is organization-specific and even department- or user-specific, depending on the prompt and data source of the vector database. The shared model does not combine company-specific data from vector databases nor the original data source, and the LLM also never “learns” from either the prompt or the generated output.

How Workday can help.

Workday provides significant transparency into how our AI models work at a granular level for each AI feature. We ensure that the security approach we take is right-sized for every use case, balancing the appropriate level of security for the sensitivity of the use case and the associated data while delivering as much business value as possible to both our customers and end users. This is determined as part of our AI risk evaluation framework prior to any AI development.

Additionally, we adhere to a multitude of critical AI-related security practices:

- Data is never shared to train third-party public models.
- Data is always encrypted in transit and at rest.
- Regional data compliance (data residency in training and inference) can be met per region.
- AI features can be configured to exclude specific locations or user groups.



7. Understanding and mitigating ethical and compliance risks.

The challenge.

While mitigating AI security risks is more straightforward, understanding the nuanced nature of ethical and compliance concerns goes beyond a technology conversation. The term “responsible AI” (RAI) has become the industry standard but can mean very different things depending on the use case, level of perceived risk, AI technique, region, and impact to individual users or stakeholders.

So determining potential exposure to ethical and compliance risks requires a more comprehensive approach than a technical evaluation and relies on a level of trust between AI solution vendors and customers (software developers* and software deployers** in this case). Additionally, AI policy and regulations are at a very early stage and are constantly evolving, so that trust must extend through the long term to mitigate future uncertainties.

What needs to be addressed.

Principles.

The most foundational level of trust is derived from a mutual agreement and alignment of organizational values, such as using AI to amplify human potential rather than to replace humans. While stated principles do not prove any risk mitigation, they are critical as the compass for not only what an AI developer chooses to create, but more importantly, how they design AI and provide transparency.

Practices.

- **Risk-based approach.** Much like a one-size-fits-all security approach isn’t useful or realistic, the same goes for responsible AI. A robust risk framework that is standardized in tiers can address high volumes and varieties of innovation in a repeatable manner that ensures that all AI innovation meets the highest levels of safety and fairness, particularly when it comes to not introducing bias into AI.

- **Transparency.** Transparency into both the AI technology itself and how that is exposed to end users is also critical to maintain AI compliance now and in the future; for example, with the European Union AI Act.
- **Human-centered.** AI should be designed with the goal to amplify the unique skills of humans, such as strategic decision-making, creativity, and nurturing relationships. It should also be designed for ease of use by a variety of users.

People.

Responsible AI is not limited to being either technical, compliant, secure, or ethical. It is all of those things, requiring a multidisciplinary team to develop the AI capabilities themselves, as well as to make decisions on best practices, governance, and user experiences.

Policy.

Finally, being involved with and abreast of developing regulations and best practices ensures organizations can adapt quickly as laws are inevitably enacted to avoid massive amounts of rework to reimplement or even rethink AI.

***Developer/Provider:** Entity that develops an AI system or a general-purpose AI model, or that has an AI system or a general-purpose AI model developed and places it on the market or puts the AI system into service under its own name or trademark, whether for payment or free of charge

****Deployer:** Entity using an AI system under its authority except where the AI system is used in the course of a personal nonprofessional activity

[Source: EU AI Act](#)

How Workday can help.

Since 2019, Workday has grown an industry-leading responsible AI program that promotes confident and continuous innovation through an ever-evolving AI landscape. Our robust risk framework helps identify and mitigate risks very early in development with a focus on fairness and bias mitigation. We champion transparency so our customers can make informed decisions based on how our AI features and models work and design AI with “human-in-the-loop” to ensure AI is not replacing human decision-making. Lastly, Workday is heavily involved with policies such as the EU AI Act and NST Risk Management Framework so customers don’t need to waste time reacting to new legislation to stay compliant.

8. Extending and customizing AI across the enterprise.

The challenge.

While the aforementioned AI talent pool remains a shortage, the pool of app developers does not—yet they are still tasked with integrating AI into custom applications without the resources to successfully do so. Navigating the complexity of data standards, selecting the right AI platform for each individual app, and ensuring that security and data privileges extend into AI capabilities make the task of leveraging AI (particularly Gen AI) increasingly difficult and resource-intensive.

Integrating with external AI partners is another way to build tailor-fit solutions for your organization if app developers are tied up with other projects. But there is still no guarantee that these solutions will persist the same granular security roles and permissions, nor address inevitable ethical and compliance concerns related to responsible AI.

What needs to be addressed.

An open and extensible developer platform that makes it easy to integrate AI into new and existing custom applications negates the need for expensive AI talent and frees up developers to focus on building solutions rather than spending months transitioning into fully ramped AI engineers. AI APIs provide flexibility to scale AI functionality across apps and use cases, leveraging an already standardized and structured dataset, as well as a potentially complex security framework.

And when this same standardization exists across data, security, and responsible AI requirements for partner solutions, it decreases the time, complexity, and risk of leveraging outside vendors to build industry-specific AI solutions.

How Workday can help.

Developers can quickly build custom Workday applications using developer copilot and power them with prebuilt AI APIs from Workday and AWS, taking advantage of the robust AI foundation of data, business context, and security. This same foundation makes it easy to integrate with Workday partners with ready-made AI apps that operate seamlessly on the Workday platform and curated AI solutions that align to our responsible AI governance framework.



9. A flexible, scalable architecture to innovate quickly and efficiently.

The challenge.

The speed with which AI advances can be an issue for customers as they deploy it. The lead time on validating and deploying new capabilities means that the market can move past you even before you are ready to go live with a new feature.

Second, as AI features are adopted and used, the capacity of compute can become an issue. Without flexible and dynamically expandable infrastructure, AI features can quickly overwhelm the systems they are running on.

What needs to be addressed.

Building decentralized AI infrastructure in multiple places across applications creates unmanageable complexity and maintenance efforts that exponentially slows down the pace of innovation as the number of AI capabilities increases.

Centralizing the testing, building, and deploying of AI features provides a more robust and future-proofed approach to AI that supports scale when considering the sheer amount of AI capabilities and various AI techniques that exist now or will exist in the future. Models need to be shared across teams as they are costly to customize, secure, test, and deploy. Replicating this work overcomplicates an already costly process.

A centralized infrastructure facilitates the consolidation and auto-scaling of the vast amount of compute resources needed for AI in a cost-conscious manner. It also permits the sharing of costly external tools (via API access) across teams and features, preventing oversubscription or illicit access to such third-party tools.



How Workday can help.

Workday has built the AI Infrastructure from the ground up to meet these challenges. Workday provides a single point of access for customers and developers. The service exposes not just the tools to build and deploy AI features but also the tools to test and validate the effectiveness and accuracy of the output. The platform auto-scales each feature as needed to ensure acceptable execution times without being overbuilt for less-active features or during times of low end-user traffic.

The embedded nature of AI tools enables Workday to build new functionalities quickly and expose them to developers and customers earlier in the product development cycle. As new models or tools are released, engineers are able to access them as soon as they are integrated into the service. Additional AI services that are utilized across features, such as content moderation, are easily tested and deployed via the centralized APIs.

10. Managing change and driving end-user adoption.

The challenge.

Even with executive buy-in for the right use case and successful evaluations, end-user adoption is still needed to capture the promised value of AI. While [52% of employees welcome AI in the workplace](#), many others are unsure of the value it brings to them personally, aren't confident in using it effectively, or are even fearful it could negatively impact their job.

Leaders are tasked with executing on effective change management across a diverse workforce that may not have much exposure to AI while also lacking the knowledge and resources to create excitement and establish adoption.

What needs to be addressed.

To increase end-user adoption, three key areas must be addressed.

- **Trust:** Will this work and be accurate? Will using this result in negative impacts to my role?

Fundamentally, end users need to trust that AI is going to actually work and that it will be used for good within their organization. Companies can provide guidelines on intended usage that resonate with specific user personas and share examples of other respected organizations that have used AI successfully and responsibly. Some level of explainability should be incorporated into the user experience, such as listing key data inputs that are most highly weighted in the generated output, so users can make more informed decisions and understand how AI came to specific conclusions or generated certain content.

- **Value:** Will this make my job easier? How will leveraging AI capabilities free up time to perform more value-add activities that are uniquely suited for humans?

Next, users need to have not only perceived business value but also a belief in the value to them personally. Framing the personal value as freeing up time for more strategic, engaging work is a great way to quell employee concerns that “increased productivity” can be interpreted as simply doing more repetitive, mundane work. The challenges, benefits, and solution need to be articulated and customized to the individual role. Real-world value metrics can boost confidence that this AI capability is worth using.

- **Enablement:** Do I understand how to use AI features? Which capabilities are automated versus those that require action on my part?

AI that is naturally integrated into existing workflows requires minimal user enablement. For those features where a user interacts with AI-generated data such as analytics to support critical business decisions, enablement should focus on how that information was aggregated and how to effectively make data-driven decisions.

How Workday can help.

With a uniquely robust AI foundation built on data and business context and strengthened by value-driven and immediately useful AI use cases, end users can quickly experience value in their specific role. Conversational, contextual AI guidance in the natural flow of work simplifies the user experience, so every user can easily take advantage of AI capabilities without memorizing where to navigate or experience unnecessary context switching. Additionally, explainability is built into the user interface to boost confidence and understanding of both the intended usage and the outputs of AI features.

Conclusion: Looking forward.

Looking forward to the future of AI, it is difficult to predict the next major technical innovation or how legislation will impact AI across the globe. But one thing we know for certain is that AI will transform how organizations in every industry operate and that the AI revolution is inevitable. Having a basic understanding of the critical AI success factors will be fundamental to capturing significant tangible value safely and securely on a continuous basis. Workday calls this “innovating with integrity.”

Whether you want to get started with AI for HR and finance or you are looking to automate and transform other parts of your organization, continuing to learn about AI remains a requirement for being successful and competitive now and into the future.

Learn more.

- [Workday AI Masterclass: Become an Enterprise AI Expert](#)
- [Your AI Journey Starts Here Guide](#)
- [Responsible AI: Empowering Innovation with Integrity](#)
- [How Workday Secures and Protects Sensitive Data in Gen AI](#)



+1-925-951-9000 +1-877-WORKDAY (+1-877-967-5329) Fax: +1-925-951-9001 [workday.com](https://www.workday.com)

© 2025 Workday, Inc. All rights reserved. WORKDAY and the Workday logos are trademarks of Workday, Inc. registered in the United States and elsewhere. All other brand and product names are trademarks of their respective holders.

20250530-critical-success-factors-to-enterprise-ai-adoption-ebook-enus

