

How to Create Your Responsible AI Charter

A framework for strategy, execution,
and AI innovation.



Table of Contents

Why Responsible AI Matters So Much	3
The Workday Approach to RAI	6
How to Develop an AI Charter	10
Sample Template: a Charter for RAI	12
Conclusion: Responsible AI, Empowered Intelligence	16

Why responsible AI matters so much.

AI has always demanded our attention and care. This is due, in part, to genuine concern over the changes AI will introduce in work and life and the real potential for these changes to include unintended negative consequences. The use of AI is rapidly expanding within organizations and, along with it, the potential for impact on human workers—including many of the people reading this guide right now.

It's unsurprising, then, that the rise of responsible AI (RAI) frameworks and practices have paralleled AI's increased deployment in business. Leaders wish to gain the most efficiency and effectiveness possible from AI systems while also avoiding the risk of unintended consequences. In working to achieve this balance, they are learning that RAI is not just the right thing to do—it's the smart thing to do for a competitive advantage. Like any transformative technology, aligning your AI deployment with organizational priorities early on means you'll be better positioned for faster, safer, and more impactful innovation as AI continues to mature.

RAI is not just the right thing to do—it's the smart thing to do for a competitive advantage.

Today, like AI itself, RAI is at an inflection point: it's time for companies to decide what RAI means in the context of their business, make specific plans that reflect and preserve their values, and test RAI principles in the marketplace of ideas as well as the marketplace of goods and services. At Workday, we define RAI as a sociotechnical approach to developing and deploying AI systems focused on identifying and mitigating risks to peoples' fundamental human rights and safety. It emphasizes the importance of creating AI systems that are fair, accountable, transparent, and explainable. For us, the goal of an effective RAI program is to facilitate the trustworthiness of AI systems, such that they will ultimately have a positive impact on society.

Regardless of the definition you choose to adopt in your organization, the time to act on putting an RAI framework in place is now. Building, growing, and evolving an effective approach to RAI governance is a significant undertaking, but one that is also incredibly rewarding and valuable.

This guide offers you a starting point for your effort, a template for an RAI charter—the foundation to roll out an effective RAI program that will endure, evolve, and serve your entire organization. An RAI charter is meant to formalize your governance framework, ensuring your values and principles are followed as AI products and technologies are built and deployed.

“

There is pressure to act. I do get the sense that there's a bit of an arms race between competing organizations. There are opportunities for organizations to lead and there will be opportunities for organizations to follow. Those who lead will have an innate competitive advantage moving forward.

Dave Mackenzie

Managing Principal, Digital, Aurecon



RAI at a glance.



RAI emphasizes the importance of creating AI systems that are fair, accountable, transparent, and explainable to empower employees, customers, and partners. The ultimate goal of RAI is to facilitate the trustworthiness of AI systems, allowing AI to scale and positively impact businesses and society.

RAI can be thought of as an umbrella term that incorporates many similar and related terms and priorities:

- **Ethical AI** typically focuses on moral principles and societal values
- **Trustworthy AI** elevates the end result of stakeholder trust once ethical and responsible priorities have been met
- **Human-centered AI** refers to the priority of keeping humans in the loop and in control of AI, using it to elevate—rather than automate—their performance and decision-making activities
- **Explainable AI** refers to systems that provide clear, understandable explanations of their decisions and actions to humans, maintaining transparency and accountability throughout any AI process or workflow

RAI is a methodology that encompasses managing risks and building trust while scaling the power of AI to improve business outcomes. An RAI program should start with a commitment from the very top, including key executives from across the company. What's more, this approach is multipronged and requires buy-in and support at all levels of the organization. RAI is not the job of just one person, team, or executive steering committee—it's a company-wide imperative that leadership must recognize and propagate throughout the organization.

An interdisciplinary challenge.

Business leaders recognize the stakes in getting AI right and want to lead effectively in both intellectual and operational contexts. That's going to require open-mindedness, sustained focus, and patience. It's also going to require participation and input from many business functions.

Because RAI affects all corners of the organization, and due to its nature as a field of study, any RAI implementation must be an interdisciplinary and collaborative venture. RAI initiatives have grown rapidly in recent months and years, commonly drawing from many areas, including:



Engineering



Legal and compliance



Data science



Public policy



Social science



Equity and social justice



User experience design

It has become clear that single-viewpoint RAI initiatives miss the point and are doomed to stall. The sociotechnical nature of RAI requires the active participation of diverse subject matter experts.



The sociotechnical nature of RAI requires the active participation of diverse subject matter experts.

Context: RAI in business and government.

AI has the potential to change the way we work and the future of work itself. This is why it is so important to balance the benefits of these technologies with their potential risks.

Policymakers and regulators around the world are focused on striking that balance. The EU finalized its landmark [Artificial Intelligence Act](#) in 2024, and the U.S. federal government released the [NIST AI Risk Management Framework](#) and enacted an [Executive Order](#) on safe, secure, and trustworthy AI in 2023. The UK held the first-ever international AI Safety Summit in 2023, and the UN, G7, and others are intensely focused on AI technology and governance.

But RAI is not just about laws, regulations, or public policy. A properly developed and applied RAI framework is not limited to a compliance exercise. Rather, it's about developing and deploying AI for good. Leading organizations will lend thought leadership and participate in public policy efforts to help shape workable and practical regulations but the rubber truly meets the road when organizations work to apply regulations and best-practice frameworks to actual business practices.

Emerging best practices among thought leaders, technological evolutions, and consumer preferences will shape AI solutions and how people use them. RAI helps organizations in the short term by driving user trust to facilitate AI deployment and user adoption. RAI also serves as a continuous risk mitigation mechanism to protect organizations in the long term as the AI landscape rapidly shifts and evolves.

The RAI competitive advantage.

In a business setting, it's important to recognize that without effective RAI governance, AI developers and deployers might waffle and delay innovation for fear of introducing unintended consequences or harmful impacts. On the other hand, without the guidelines and guardrails that robust RAI governance offers, businesses might make decisions that harm the organization's brand and reputation.

Programmatic and proactive risk mitigation efforts clear roadblocks and speed innovation with emerging technologies. They guide businesses to rely on better-quality data and inputs, which leads to better outputs and insights for decision-making. What's more, having a solid RAI framework in place allows organizations to efficiently adapt and respond quickly to technology and regulatory innovations and changes.

We know that users of AI technologies are more likely to adopt and use those technologies with confidence if they trust that the AI was developed and deployed responsibly. Transparency, explainability, and user control help to establish the trust required for widespread deployment and adoption, driving fast time-to-value and tangible business ROI.

This leads us to the following questions: What do businesses do right now to embark on RAI practices that are under their control while also anticipating more forceful influence from regulators and continuously evolving markets? What does "responsible" mean, anyway? And how can this work amplify, rather than limit, the fantastic and transformative benefits that AI is already delivering?

The Workday approach to RAI.

Workday was born in the cloud, and we've benefited from a decade of direct experience with AI. This advantage stems from the true-cloud architecture on which Workday was built, where AI is an integral part of the system, not an afterthought.

From the very start, we have been focused on data privacy and protection, and those safeguards extend to our AI technologies. We are guided by our Workday core values to understand that any perceived dissonance between AI innovation and integrity is a false dichotomy. RAI-by-design enables us to effectively map, measure, and mitigate risk, so we can confidently and consistently innovate.

“

Embracing Workday AI to drive our HR strategies has empowered us to prioritize upskilling and internal mobility, underscoring our dedication to responsible AI and advancing technology in a human-centered, inclusive, and trust-building manner.

Sadie Bell

VP, Innovation & Deployment, People Systems and Automation, HPE



According to Workday research,

94%

of enterprise companies are investing in AI technology ¹

Forrester reports that AI platform budgets will triple, driving a

36%

compound annual growth rate from 2023 to 2030 ²

Accenture reports that

96%

of organizations support some level of AI government regulation but only 2% have self-identified as having fully operationalized RAI across their organization ³

Stand on your pillars.

After years of creating, testing, and optimizing our internal approach to RAI, Workday has adopted four pillars—we refer to them as “the four Ps” that inform our comprehensive, actionable approach to RAI going forward.



Principles.

Organizations should start with their core values and develop RAI principles to align. Our commitment to ethical AI at Workday is a reflection of our [core values](#), including a focus on our employees, customer service, innovation, and integrity.

In our effort to develop responsible and trustworthy AI, we aspire to achieve the following goals:

- Amplify human potential
- Positively impact society
- Champion transparency and fairness
- Deliver on our commitment to data privacy and protection

Our AI ethics principles serve as our North Star, guiding every AI technology that we build.



Practices.

We’ve created and operationalized industry best practices to embed RAI-by-design into our groundbreaking AI features and applications. Every new AI feature we build undergoes a rigorous evaluation to ensure that it aligns with our RAI framework from ideation to deployment. This includes risk evaluation, mitigation, and transparency to our customers through our AI fact sheet program.

Our guidelines and safeguards are designed to facilitate the development of AI technologies that are:

- Safe and secure
- Human-centered
- Transparent and explainable



Effective organizational frameworks for responsible AI should consider four fundamental pillars: principles, practices, people, and policy. Companies need to also ensure transparent communication about their approach to each of those areas.

Jim Stratton

Chief Technology Officer, Workday





People.

We rely on a diverse set of cross-functional experts to develop and maintain our RAI governance framework. This includes input and participation from engineering and product teams, data science and quality control, legal and compliance, public policy, UX design, and belonging and diversity.

While the RAI program is enterprisewide in scope, direct participation can be categorized as follows:

RAI team: A team of independent subject matter experts led by our chief RAI officer. This team of scientists and AI experts dedicates 100% of their time to building and maturing RAI governance.

RAI advisory board: This board includes diverse C-suite executives from across the organization. The advisory board guides and supports the work of the dedicated RAI team.

RAI champions network: This network consists of experts from across the company who are embedded within key product and technology teams. They are passionate about the development of responsible and ethical AI solutions.



When I think about the future of AI, I think about the different disciplines coming together—from the humanities, the social sciences, engineering, legal, policy, HR, IT—to make sure that we take the time now to learn from past technological revolutions so that this one creates the best possible future, for everyone.

Kathy Pham

VP, Artificial Intelligence and Machine Learning, Workday



Policy.

Workday takes a proactive approach to responsible AI and stays ahead of changes in the AI landscape. By actively engaging with lawmakers, we help shape regulations that build trust, enable innovation, and ensure the safe and ethical development and deployment of AI technologies.

For example, when the EU AI Act was finalized, Workday customers didn't need to react or take action regarding their Workday applications. We have been in close contact with officials regarding the Act's development and have worked to ensure alignment of our RAI practices to the Act's requirements throughout the entire process. This proactive alignment approach translates to direct efficiencies for our developers and customers, avoiding disruptions and allowing focus to remain on core business activities.

The Workday RAI charter was developed with these pillars in mind through a multidisciplinary process of analysis and strategy led by our chief responsible AI officer, in collaboration with our vice president of artificial intelligence. Driven by social and data science perspectives, the charter provides the foundation for governance needed to direct a focused, yet flexible, RAI program.

Getting started with responsible AI.

Embarking on your responsible AI journey requires a thoughtful approach. Familiarize yourself and your advisory board with relevant regulations and best practices and consider their implications for your business. We also recommend designating a dedicated leader to oversee the RAI initiative.

Here are some key steps to help you get started as you begin building your own RAI charter:

Get familiar with developing regulations and industry best practices:

- NIST AI Risk Management Framework
- EU AI Act
- Future of Privacy Forum AI in HR best practices

Consider what this means in the context of your business:

- The risks posed by AI vary based on context and characteristics
- Consider your own business and AI solutions in the context of the resources listed above

Determine who at your organization should be part of the RAI initiative:

Whether it's a dedicated team, an executive committee, a network of volunteers, or all of the above, someone needs to be accountable for the governance and oversight of the RAI initiative. There is no “one-size-fits-all” approach that works for everyone. Again, you must consider best practices in the context of your own organization.

- Some organizations, including Workday, have a dedicated, independent person in charge of the responsible AI governance framework, a chief RAI officer (CRAIO).
- Other organizations have expanded the scope and skill set for the chief data officer (CDO) or chief privacy officer.
- RAI can sit under the office of the chief legal officer, chief technology officer, chief product officer, or another group.
- Some organizations establish groups of teams (or pillars) for an RAI foundation, each of whom is part of the larger RAI team or advisory board.

Stakeholders, including board members, customers, partners, and regulators, will have questions about your company's use of AI and data, including the RAI strategy, how it was developed, how it is being operationalized and governed, and which KPIs are being measured, to name a few.




How to develop an RAI charter.

At Workday, we created an RAI charter for our internal purposes to formalize the program. We're sharing it externally now to help other organizations "get started on," benchmark, or refine their approach to responsible AI in alignment with this ever-changing landscape of AI innovation and risk.

Charters are documents of conviction and direction. They have been used throughout history to establish rights and authority, such as with the Magna Carta, or "Great Charter," and to formalize important agreements. They also stand out among the various plans and publications that might be circulating within your company, which is appropriate given the importance and reach of AI and related technologies.

Strategically, adopting an RAI charter takes companies from merely talking about RAI to actually codifying it. It can be used to document the mission and vision of the effort, as well as the goals, strategies, roles and responsibilities, and execution of all agreed-upon principles and standards in a way that clearly supports innovation. This is a critical step for any company embracing AI—as most are—as well as for any company that wants to enter tomorrow's working world with momentum and clarity.



Adopting an RAI charter takes companies from merely talking about RAI to actually codifying it.

Notably, an RAI charter is a living document that should be updated over time. The Workday charter includes mechanisms for doing so, and is meant to be a starting point to consider for ideation in your business.



Pework: questions to guide your RAI charter.

The Workday RAI charter (sample structure provided in the next chapter) was developed by our dedicated RAI team, with guidance from our RAI champions and approved by our RAI advisory board. The primary purpose of the charter is to codify the foundational aspects of our RAI program, including its mission, vision, and purpose. Identifying roles, responsibilities, and decision-making authorities for the dedicated team, champions, and advisory board is an important function of the charter. We offer it to you in hopes that it will support or facilitate your RAI journey.

These starter questions will likely get you to a place where you can begin the work of crafting a charter. They will clarify your company's ethical standpoint and practical purposes for doing this work.

✓ The basics.

- Why do we need a responsible AI framework?
- How will we define responsible AI?
- Are we building frameworks for AI development, deployment, partnerships, procurement, or all of the above?

✓ Governance and oversight.

- Who needs to be involved in the creation and maintenance of our AI governance?
- Who has decision-making authority to identify appropriate participants for AI governance efforts?
- How will we address stakeholder concerns and incorporate their feedback?
- What are the escalation procedures for addressing AI-related issues?

✓ Ethical and legal considerations.

- When our AI interacts with a human, what could go wrong? What practices should we put in place to avoid unintended consequences?
- How will we ensure compliance with existing laws and regulations—and readiness for future regulatory frameworks?
- Will we limit RAI efforts to legal and compliance? How will we separate efforts to ensure ethical and responsible development and use from legally compliant development and use?

✓ Fairness, inclusivity, and impact.

- How will we measure and mitigate any unforeseen biases in AI systems?
- What steps will we take to include diverse perspectives in AI development?
- What potential unintended social and environmental impacts could arise from our AI systems?

✓ Transparency and accountability.

- How will we maintain transparency in our AI development and deployment activities?
- What mechanisms are in place for reviewing and evaluating AI performance and impact?

✓ Charter improvement and evaluation.

- What processes are in place for continuous learning and improvement with regard to our RAI practices?
- How will we stay updated on advancements and best practices in RAI?
- What steps will we take to ensure our AI charter remains relevant and useful over time?
- Who has decision-making authority to publish and update the charter?



Sample template: a charter for RAI.

Your charter should include crucial RAI program objectives, governance structure, and related roles and responsibilities to maintain alignment across all key stakeholders over time. By necessity, the language and sections of a charter should evolve as the program matures.

Cross-functional stakeholders.

Your RAI program should engage a diverse range of teams across the enterprise, typically including:

- Product and technology
- Legal, compliance, and corporate affairs
- UX design
- Belonging and diversity
- Marketing and communications

This diversity extends beyond departmental boundaries, incorporating varied perspectives and backgrounds, such as social science, engineering, data science, privacy, legal, and security. As the program evolves, additional teams should be considered.

“

It can be difficult to get a responsible AI effort off the ground. Formalizing the mission, vision, purpose, goals, and team responsibilities in a formal charter is a crucial early step toward success.

Kelly Trindel, PhD

Chief Responsible AI Officer, Workday

Program objectives.

The development and execution of an RAI program has numerous benefits, including the shared objectives identified below, which should be agreed upon by all teams referenced in the charter.

OBJECTIVE	OUTCOMES
Governance	<ul style="list-style-type: none">• Develop a clear governance framework and culture that operationalizes AI ethics principles• Provide the appropriate policies, procedures, guidelines, standards, and tools for responsible AI development
Proactive AI risk measurement and management	<ul style="list-style-type: none">• Develop and maintain a scalable risk-evaluation approach. Align to existing and developing global, federal, state, and local regulations• Develop a risk management methodology to identify RAI-relevant risks, assess impact, and connect to existing risk management structures and priorities
Enhanced decision-making	<ul style="list-style-type: none">• Develop frameworks that empower teams to make effective and efficient decisions as they build AI technologies. This includes providing the resources to:<ul style="list-style-type: none">• Clarify and communicate risk tolerances• Document agreed-upon cross-functional roles and responsibilities for building or procuring AI products and services• Develop clear escalation and reporting paths to the RAI program's senior leadership as necessary
Responsible innovation	<ul style="list-style-type: none">• Foster a governance program to empower, rather than discourage, AI innovation. This means:<ul style="list-style-type: none">• Develop streamlined review processes with fewer and consolidated touchpoints• As much as possible, develop self-service tools and resources for internal teams to uptake• Create a culture that understands and respects the diversity of risk types—for example, legal, business, reputational, integrity—and in certain instances, consider how risks stack up against benefits that innovative technologies can bring
Cross-functional collaboration	<ul style="list-style-type: none">• Promote collaboration on RAI program activities to drive a culture of co-ownership of the governance framework and deliverables
Communication	<ul style="list-style-type: none">• Develop clear and consistent messaging on company perspectives, success stories, and best practices on RAI for relevant internal and external audiences

Governance and decision-making.

The next two pages outline the groups and organizations that should be involved in further developing the RAI program and ensuring program objectives are met.

COMMITTEE	EXECUTIVE PROGRAM SPONSOR
Overall responsibilities	<ul style="list-style-type: none">• Provide top-down support and guidance for the RAI program• Ensure the RAI program and its objectives are accurately represented in executive meetings and discussions, including company board meetings• Allocate necessary budget and other resources to support RAI
Decision-making framework	<ul style="list-style-type: none">• Pilot in charge of escalating issues, decisions, and significant concerns to executive leadership• Advisor on program priorities and objectives• Informed on program status• Decision-maker for program budget and resourcing
COMMITTEE	RAI ADVISORY BOARD
Overall responsibilities	<ul style="list-style-type: none">• Promote and prioritize the RAI vision, mission, objectives, and scope throughout the company and within their individual organizations• Provide expert guidance to the RAI team on ideas and proposals for program enhancements, as well as needs, gaps, and challenges• Review and approve RAI program deliverables that meet program objectives• Review novel or emerging issues about AI technologies that are not currently in scope of existing governance framework and make decisions that reflect company values and business priorities• Provide recommendations on resourcing and prioritization to support program objectives
Decision-making framework	<ul style="list-style-type: none">• Pilot for elevating program awareness across the company• Advisor as it relates to providing subject matter expertise; addressing program needs, gaps, and challenges; and resourcing and prioritization• Informed on RAI program objectives and prioritization of governance framework deliverables• Decision-maker for (1) approving new AI governance policies, procedures, assets, and collateral, and (2) novel or emerging RAI issues or escalations

COMMITTEE	RAI TEAM
<p>Overall responsibilities</p>	<ul style="list-style-type: none"> • Develop and implement the program vision, mission, objectives, and scope • Develop RAI governance frameworks, collateral, policies, and tools • Provide RAI subject matter expertise throughout the organization and in external settings when appropriate • Organize regular meetings with the RAI advisory board to report on the progress of program objectives, solicit expert advice, and facilitate the review of new governance collateral and decision-making for novel and emerging issues • Keep the executive program sponsor apprised of the successes and challenges of the RAI program, and escalate issues, questions, and blockers as needed • Document and track historical and ongoing decisions and approvals as they relate to the RAI governance framework • Develop RAI training and communication initiatives • Maintain, grow, and enable the RAI champions network
<p>Decision-making framework</p>	<ul style="list-style-type: none"> • Pilot for developing the RAI program and associated governance structure, including policies, practices, tools, and training • Advisor on RAI subject matter expertise across the company • Informed on (1) AI development and relevant product and technology roadmaps and strategy, (2) legal and compliance guidance for AI development, (3) company efforts in public policy, and (4) appropriate company-wide communication channels • Decision-maker for (1) RAI program objectives, (2) prioritization of RAI governance framework deliverables, (3) RAI training and development, (4) internal and external RAI communications and talking points, and (5) management of the RAI champions program

COMMITTEE		RAI CHAMPIONS PROGRAM	
Overall responsibilities		<ul style="list-style-type: none"> • Contribute to RAI governance framework deliverables. • Coordinate with relevant stakeholders in their individual organizations to understand team priorities and provide feedback on RAI objectives and activities. • Promote adoption and effective utilization of the RAI governance framework. • Amplify RAI messaging at their team level. • Keep the RAI team informed as to the need for new or improved governance collateral, training, or tooling. 	
Decision-making framework		<ul style="list-style-type: none"> • Pilot for organizing input and feedback from their various teams to the RAI program • Advisor for development of new governance tools and training, program challenges, concerns, and prioritization • Informed on new aspects of the program approved by the RAI advisory board • Decision-maker for local workflows and methodologies at their team level to execute program objectives 	
DOCUMENTED BY	VERSION	APPROVED BY	DATE
[name]	[version number]	Responsible AI advisory board	[release date]

Conclusion: responsible AI, empowered intelligence.

If you're reading this guide, odds are you're probably already knee-deep in AI and RAI. If so, keep it up. The work you've done is meaningful and valuable.

Workday has always believed that AI is best used to amplify human potential and thereby advance business goals and organizational missions. As demonstrated throughout this guide, effective RAI is a powerful way to establish what needs to be done, by whom, and in what ways in order to unleash the potential of AI and related technologies in business. The charter template we've shared gives you an important, reliable resource to help guide the effort as you refine your bespoke organizational approach.

For further inspiration and to see how leading organizations have crafted their responsible AI charters, you can explore the published guidelines from several industry leaders:

[Microsoft Responsible AI Standard](#)

[PwC Responsible AI](#)

[IBM AI Ethics Guidelines](#)

[Google AI Principles](#)

To learn more about our RAI principles, methodologies, and the impact of AI initiatives at Workday, read our [whitepaper here](#).

To learn more about how to get started with your own RAI initiative, watch our AI masterclass [here](#).

You can also access further education, strategy, and global insights on the Workday RAI web page [here](#).



Sources

¹ Workday, "AI IQ: Insights on Artificial Intelligence in the Enterprise"; 2023.

² Forrester, "Global AI Software Forecast, 2022"; Michael O'Grady and Mike Gualtieri; 2022.

³ Accenture, "Accenture's blueprint for responsible AI"; 2024.



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

© 2024 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.
20240805-getting-started-rai-charter-roadmap-guide-enus