

# AI Literacy and Compliance at Salesforce

Salesforce Office of Ethical & Humane Use,  
Government Affairs & Privacy Compliance Team





# AI is transforming the way businesses operate.

As the #1 AI CRM, Salesforce is committed to ensuring that everyone—from small business owners to enterprise leaders, from our own frontline employees to our executive leaders—has the knowledge and skills to navigate this shift. With [Agentforce](#), Salesforce enables organizations to deploy autonomous AI agents that act on unified, real-time data across their systems, helping every employee deliver more personalized, efficient, and secure customer experiences. AI is no longer just an emerging technology; it is an essential partner in driving productivity, efficiency, and customer success. As AI adoption accelerates, the need for accessible and comprehensive AI literacy grows.

Understanding AI goes beyond knowing how to use it—it requires recognizing its capabilities, limitations, and ethical considerations. Tools like Agentforce can automate workflows and enhance decision-making, but without AI literacy, users may struggle to make informed decisions regarding AI systems.

Beyond individual proficiency, AI literacy is becoming a regulatory necessity. The [EU Artificial Intelligence Act](#) (EU AI Act) sets expectations for [AI literacy requirements](#) for deployers and providers, with compliance requirements in place since February 2, 2025. Organizations must assess how these regulations apply to their unique contexts, and having the right knowledge is crucial for meeting these obligations. <sup>1</sup> This guide is intended for informational purposes only and it does not provide legal advice. It offers a starting point to help businesses navigate these complexities, offering practical insights to support responsible AI adoption.

At Salesforce, we believe AI should be an enabler, not a barrier. That's why we are investing in AI education programs to ensure our employees, our customers and our communities can harness AI's potential responsibly and effectively. As AI continues to shape the future of work, fostering AI literacy isn't just an opportunity—it's a necessity.

© COPYRIGHT MAY 2025 SALESFORCE, INC.

<sup>1</sup> This document is an overview of AI Literacy and is provided for informational purposes only. It does not provide legal advice. Salesforce urges its customers to consult with their own legal counsel to familiarize themselves with the requirements that govern their specific situations. This information is provided as of the date of document publication, and may not account for changes after the date of publication.

# Introduction to AI Literacy

**Salesforce is dedicated to promoting AI literacy and ensuring the ethical use of AI technologies. This resource outlines our holistic strategy and highlights our ongoing contributions to advancing AI literacy for customers and employees.**

We have established a set of trusted AI principles that emphasizes responsibility, accountability, transparency, empowerment, and inclusion, ensuring that AI technologies are developed and deployed responsibly. Transparency and explainability are key priorities for Salesforce, as the company ensures users understand how AI-driven decisions are made by providing clear explanations and insights into AI models. Salesforce also actively works to identify and mitigate biases in AI systems through rigorous testing and validation processes. Additionally, Salesforce empowers its customers with tools and resources to customize AI models to align with their ethical standards, fostering a responsible and trustworthy AI ecosystem.



## 1

# Trusted AI Principles

At Salesforce, responsible AI has always been a priority—not just to keep up with innovation, but to anticipate it. We developed our first set of [trusted AI principles](#) in 2018. As we entered into the era of [generative AI](#) in early 2023, we augmented our trusted AI principles with a set of [five guiding principles](#) for developing responsible generative AI, as one of the first enterprise companies to publish guidelines in this emerging space. [These principles](#) still hold true for the current, third era of AI – the era of AI agents.

- **Responsibility:** We start by asking not just whether we can build AI features, but whether we should. We work with external human rights experts to strengthen protections and continuously improve. We empower customers, Salesforce employees, and partners to use AI responsibly by offering tools and resources that mitigate bias, enhance transparency, and promote safe, secure AI practices.
- **Accountability:** We believe in accountability and ensure all parties understand their responsibilities by inviting customer feedback, collaborating with industry peers and civil society groups, enabling employees to raise concerns, and continuously reviewing our AI models and training data.
- **Transparency:** We provide customers with disclaimers and disclosures whenever AI is used, and we ensure our AI systems are understandable and explainable. This includes offering detailed documentation and insights into how the models function, the data they rely on, and the reasoning behind their outputs. By prioritizing transparency, Salesforce enables users and stakeholders to see and understand how AI decisions are made—building trust in the technology. Transparency also supports regulatory compliance by making it easier to audit and validate AI processes.
- **Empowering:** We're making AI development accessible—enabling users of all skill levels to build AI applications without writing code. Through Trailhead, our free learning platform, we equip users with the skills they need to succeed in an AI-driven world. In addition, accessibility is core to this mission: we design AI solutions that empower all individuals, including people with disabilities, by enhancing independence, increasing productivity, and expanding opportunity.
- **Inclusive:** We strive to reduce bias and ensure fairness in AI outcomes by implementing rigorous testing and validation processes to identify and mitigate biases in AI models. Salesforce uses diverse datasets and employs techniques such as fairness-aware machine learning to create more equitable AI systems. By prioritizing fairness, Salesforce aims to deliver AI solutions that provide equal outcomes and benefits to all users, regardless of their background or characteristics.






## 2

# AI Council

The formation of Salesforce's AI Council—comprising members from the Office of Ethical and Humane Use, Privacy, Product Legal, Employment Legal, Enterprise Architecture, and Security—reflects a proactive and comprehensive approach to responsible AI. This cross-functional group plays a crucial role in cultivating a strong internal culture of ethical AI use, supporting compliance with evolving regulatory frameworks such as the EU AI Act, ensuring comprehensive risk management, and providing consistent policy guidance. Let's explore the impact of each function:

- **Employment Legal:** Our employment lawyers work with the teams developing and implementing AI to help them understand the specific risks to employment-related processes, the requirements of emerging laws at the intersection of AI and labor laws, and develop appropriate controls that allow for the responsible use of AI where people's employment may be affected.
- **Enterprise Architecture:** Enterprise Architecture provides a technical perspective to the AI Council, ensuring AI technologies are securely, scalably, and efficiently integrated into Salesforce's infrastructure. They enforce risk controls, mitigations, and design standards, establishing a secure foundation for AI systems. Their expertise is essential for translating ethical and legal requirements into actionable guidelines for AI development and deployment.
- **Office of Ethical and Human Use:** This team keeps ethical considerations at the core of AI development and deployment, championing Salesforce's guidelines for responsible AI. Their leadership in AI Council reviews shapes the Council's decisions and ensures AI aligns with Salesforce's values—fostering customer trust and minimizing potential risks by recommending controls, mitigations, testing, and user-facing guidance.
- **Privacy:** Privacy ensures that AI initiatives comply with data protection laws like GDPR and the EU AI Act, focusing on principles such as data minimization, purpose limitation, and security. Their involvement in privacy reviews ensures AI systems respect user privacy, safeguarding customer trust and mitigating legal and reputational risks.
- **Product Legal:** Product Legal brings expertise in AI regulations, ensuring Salesforce's AI features comply with regulatory frameworks such as the EU AI Act, focusing on safety, transparency, and user rights. They review AI projects for compliance and advise on legal risks, providing proactive oversight that helps prevent legal challenges and ensures responsible, compliant development and deployment of AI solutions.



- 
- **Security:** The Security team focuses on safeguarding AI systems and the data processed. They identify and mitigate security risks, ensuring the confidentiality, integrity, and availability of AI services. Their involvement in security reviews and risk controls is crucial for protecting sensitive data and maintaining customer trust.

The cross-functional AI Council plays a pivotal role in advancing responsible AI enablement at Salesforce. By bringing together experts from ethics, legal (product, employment, privacy), security, and architecture, the Council ensures comprehensive risk management, consistent policy guidance, and proactive legal compliance. Its work not only reduces ambiguity for teams developing AI but also promotes company-wide AI literacy through clear guidelines and training. This integrated oversight strengthens security, protects privacy, and reinforces ethical standards, empowering employees to engage with AI responsibly. The AI Council's framework and effectiveness are periodically reviewed to ensure ongoing alignment with Salesforce's commitment to responsible AI innovation and governance, contributing to company-wide AI literacy and trust.

### 3 Transparency and Explainability Initiatives

Salesforce prioritizes making AI systems transparent and explainable through several key initiatives. One such initiative is the use of [Model Cards](#), which provide detailed documentation for AI models. These Model Cards include essential information about the model's purpose, performance metrics, and limitations. By offering this level of detail, Salesforce ensures that users have a clear understanding of what each AI model is designed to do, how well it performs, and any potential constraints or biases it may have. This level of transparency helps users make informed decisions about deploying AI models in their specific contexts.

In addition to Model Cards, Salesforce offers a suite of Explainable AI Tools designed to help users understand how AI models make decisions. These tools include features like feature importance, which highlights the most influential factors in a model's decision-making process, and decision trees, which visually map out the steps an AI model takes to arrive at a particular outcome. By providing these transparency tools, Salesforce enables users to gain insights into the inner workings of AI models, fostering trust and confidence in the technology. This commitment to explainability ensures that AI systems are not just powerful but also understandable and accountable, aligning with Salesforce's broader ethical principles.

## 4

## Customer Education and Training

Salesforce invests significantly in AI education to ensure responsible and effective use of AI technologies. [Trailhead](#), Salesforce's free online learning platform, has expanded courses to offer 400 AI-specific skills modules, including AI fundamentals, ethical AI use and prompting. Salesforce [offers](#) its existing premium AI courses and AI certifications free of charge and available to anyone on Trailhead through the end of 2025.

Specifically for agentic AI, Salesforce introduced [Agentblazer Status](#) which provides guided, step-by-step support ranging from fundamentals of Agentforce to managing complex multi-agent systems.

To support customers on their AI journey, the Office of Ethical and Humane Use has curated a comprehensive list of relevant resources. Organized into [Beginner](#), [Intermediate](#), and [Advanced/Technical](#) categories, these resources make it easy to identify the most suitable learning paths. Whether they're just starting out or seeking to deepen their expertise, customers can rely on this guide to build the knowledge required to make informed and responsible decisions in an AI-powered world.



## Beginner - AI 101

Designed for individuals with little to no experience with AI, this level covers fundamental concepts, common AI applications, and basic responsible AI principles.



### TRAILHEAD

<a href="#"><u>EU's Privacy Law Basics</u></a>	Develop foundational knowledge about EU privacy laws, crucial for understanding the regulatory landscape.
<a href="#"><u>Trusted Agentic AI</u></a>	Learn how Agentforce uses safeguards and responsible AI principles to create ethical AI.
<a href="#"><u>Artificial Intelligence Fundamentals</u></a>	Learn the basics of AI and the technology that drives it.
<a href="#"><u>Get Ready for Agentforce</u></a>	This trail will walk you through understanding how AI and agents can transform your jobs and increase productivity. Then, launch your very first agent!
<a href="#"><u>Explore Generative AI Tools</u></a>	Discover the fundamental tools for your generative AI journey including Natural Language Processing Basics, LLMs, Prompt Fundamentals, and Model Fine-tuning.
<a href="#"><u>Drive AI Within your Organization</u></a>	Learn how to develop an AI strategy and successfully implement AI projects. Courses include Generate AI for Organizations, AI + data: Project Planning, and Change Management for AI Implementation.





## ADDITIONAL RESOURCES

### [Managing the Risks of Generative AI \(Harvard Business Review\)](#)

Organizations need to be mindful of the ethical implications and take necessary steps to reduce risks. Specifically, they need to: use zero or first party data, keep data fresh and well labeled, ensure there's a human in the loop, test and re-test, and get feedback.

### [Trusted AI Resource Page](#)

Directory of essential resources for building trustworthy and responsible AI.

### [The Agentic AI Era: After the Dawn, Here's What to Expect](#)

The era of agentic AI is here—what comes next? Explore how autonomous AI agents are reshaping industries, the challenges they bring, and what businesses need to do to stay ahead.

### [From Adaptability to AI Literacy: Critical Skills Needed to Succeed in the Digital Labor Revolution](#)

As AI transforms the workforce, reskilling is more critical than ever. Discover how businesses can empower employees with the skills needed to thrive in the era of digital labor.



## Intermediate – Build with AI

Aimed at users with some familiarity with AI, this level delves into how AI models work, practical use cases within Salesforce tools like Agentforce, and key ethical considerations.



### TRAILHEAD

<a href="#"><u>Agentforce Specialist</u></a>	As part of the <a href="#">AI for All initiative</a> , Agentforce Specialist exams are free through 2025! Level up as an Agentblazer and master the art of driving customer success and productivity with personalized, intelligent solutions powered by Agentforce and generative AI.
<a href="#"><u>Become an Agentblazer Champion</u></a>	Confidently explain Agentforce concepts and its business impact, gain foundational knowledge, and build an agent. Modules include “Autonomous Agents” and “Establish Trust with Autonomous Agents.”
<a href="#"><u>Build an AI Agent with Agentforce</u></a>	This trailhead starts with the AI basics and takes you through the journey to set up and test a conversational AI agent.
<a href="#"><u>Learn Data for AI</u></a>	Explore best practices for making data-driven decisions with AI. Modules include Critical Thinking and Decision-Making with Data and AI, Data Fundamentals for AI, and Data Bias Recognition and Prevention in AI.
<a href="#"><u>Prepare Your Data for AI</u></a>	Successfully implement AI with the Einstein Trust layer, enhanced data quality, and optimized data sets.
<a href="#"><u>Get Started with Prompts and Prompt Builder</u></a>	Get the most out of AI with Salesforce Prompt Builder, ultimately learning how to create prompt templates and use them in your orgs.
<a href="#"><u>Drive Productivity with Salesforce AI</u></a>	Learn to use out-of-the-box Salesforce AI products, get hands-on, and prep for the AI Specialist Certification.

## ADDITIONAL RESOURCES

[Salesforce's Agentforce Is Here: Trusted, Autonomous AI Agents to Scale Your Workforce](#)

Unlock the power of AI-driven automation to enhance customer interactions, boost productivity, and build intelligent, personalized solutions.

[Achieving a Trusted Agentic AI Ecosystem: Salesforce Report Offers a Roadmap](#)

[Salesforce's Trusted AI and Agents Impact Report](#) outlines the company's approach to designing and deploying AI agents responsibly, emphasizing trust, transparency, and ethical standards throughout their lifecycle.







## Advanced/Technical - AI Deep Dive

Geared toward developers, data scientists, and AI practitioners, this level explores in-depth technical concepts, and model optimization.



### TRAILHEAD

[The Next Generation of RAG: How Enriched Index Redefines Information Retrieval for LLMs](#)

Discover how Salesforce Engineering is advancing retrieval-augmented generation (RAG) with Enriched Indexing—enhancing information retrieval for LLMs to deliver more accurate, context-aware AI responses.

[How Generative AI Simplifies Data Analytics Formula Creation for Faster Insights](#)

Generative AI is transforming data analytics by simplifying formula creation—making it easier and faster to uncover insights. Learn how AI-powered automation is streamlining complex data analysis.

[How Slack AI Processes Billions of Messages to Reduce Information Overload with AI-Powered Search and Summarization](#)

Slack AI is transforming the way teams navigate information overload. Discover how it processes billions of messages to deliver smarter search and AI-powered summarization, helping users find what matters—fast.

[Prompt Injection Detection: Securing AI Systems Against Malicious Actors](#)

Prompt injection attacks pose a serious risk to AI security—but detection is evolving. Learn how Salesforce is tackling this challenge with advanced defenses to keep AI interactions safe and trustworthy.

[Salesforce AI Research Publications](#)

Dive into our publication database to explore the latest breakthroughs in state-of-the-art AI research.

[New Salesforce White Paper Tackles LLM Security Risks](#)

[White paper](#) outlines emerging security threats posed by Large Language Models (LLMs), such as prompt injections, training data poisoning, supply chain vulnerabilities, model theft, and unsafe training environments, offering actionable strategies for organizations to mitigate these risks.





## 5

## Customer Zero

As Customer Zero, Salesforce also prioritizes investing in our own employees, recognizing the importance of internal AI training and development.

### Key initiatives include:

- **Agentblazer Ranch:** A dedicated floor in Salesforce's San Francisco headquarters designed to help employees solve business challenges and drive customer success using Agentforce.
- **Annual Compliance Training:** Annual mandatory training underscores the critical importance of privacy in the context of AI development and use. It highlights core privacy tenets like transparency, choice, and privacy by design as essential for building ethical AI systems. By connecting these principles to Salesforce's foundational value of Trust, the training emphasizes responsible AI practices.
- **Career Connect:** An internal AI-powered talent marketplace that provides employees with personalized recommendations for skill development, learning opportunities on Trailhead, and potential new roles within the company.
- **Quarterly Agentforce Learning Days:** Regular events encouraging employees to engage with AI tools and assistants, enhancing their understanding and proficiency in using AI technologies.
- **AI Enterprise Skilling:** AI and AI agent literacy has been identified as one of the top ten Enterprise Skills. In response, Enterprise Learning programs are enhancing their curricula to equip employees with the knowledge and skills needed to interact with, manage, and use AI agents and automated systems both effectively and ethically.
- **Employee Feedback Channels:** We've developed many employee feedback channels including new questions to our company-wide bi-annual Great Insights Survey that matches employees' self-assessment of AI literacy with targeted action and training.

### In addition to these programs, Salesforce has established AI Centers to further support employee upskilling:

- **London AI Center:** Opened in June 2024, this center serves as a hub for AI innovation and collaboration, bringing together industry experts, partners, and customers to advance AI while providing upskilling opportunities for employees.

## 6

## Compliance with Regulations

- **Continuous Monitoring:** Salesforce is committed to staying updated with the ever-evolving landscape of AI regulations. This involves continuously monitoring changes in laws, guidelines, and best practices related to AI. By doing so, Salesforce prioritizes keeping its AI technologies and practices compliant with the latest regulatory requirements. This proactive approach helps in identifying potential compliance issues early and allows for timely updates to policies and procedures.
- **Internal Audits:** To maintain high standards of AI literacy and ethics, Salesforce conducts regular internal audits. These audits are designed to assess the company's adherence to established AI policies and ethical guidelines, with data access limited to functional needs. By systematically reviewing AI systems and processes, Salesforce can identify areas for improvement, ensure that employees are following best practices, and address any compliance gaps. This rigorous auditing process helps in maintaining transparency and accountability within the organization.
- **Collaboration with Legal Experts:** Salesforce recognizes the importance of interpreting and implementing complex regulatory requirements accurately. To achieve this, the company collaborates closely with legal experts who specialize in AI and technology law. These experts provide valuable insights and guidance on how to navigate the evolving regulatory landscape. By working together, Salesforce ensures that its AI initiatives are not only innovative but also legally sound and ethically responsible. This collaboration helps in aligning AI development with legal standards and mitigating potential legal risks.

## Collaboration and Advocacy

Salesforce actively collaborates with industry partners, policymakers, and other stakeholders globally to promote responsible AI practices:

- **Partnerships with Industry and Governments:** Salesforce actively partners with various industry organizations and participates in relevant government-led initiatives to foster a collaborative environment for developing and sharing best practices in AI literacy. By working with these organizations, Salesforce contributes to the creation of guidelines and educational resources that promote responsible AI usage across the industry. These partnerships enable Salesforce to stay at the forefront of AI advancements while ensuring that ethical considerations are integrated into AI development and deployment.

We are proud to be part of the first cohort of companies who signed onto the EU AI Pact with other industry leaders and to reinforce our commitment to AI governance by ensuring that our systems meet the highest standards of trust, transparency, and accountability. The EU AI Pact has helped raise awareness internally and externally about the upcoming regulation. It has also created opportunities for organizations to share best practices on a number of topics, including AI literacy.

- **Policy Advocacy and Thought Leadership:** We actively engage with policymakers across the globe to ensure AI is developed and deployed responsibly, while fostering innovation. This involves participating in discussions, providing expert testimony, and contributing to the development of policy and regulatory frameworks that are both practical and protective. We also publish resources to inform policymakers and other stakeholders on AI advancements, offering recommendations to harness the opportunities of AI.
- **Research and Development:** We prioritize strategic investments in research to advance the understanding and implementation of ethical AI. This includes funding academic studies, collaborating with research institutions, and conducting in-house research to explore new methodologies and technologies that enhance AI ethics. By prioritizing research and development, Salesforce aims to lead the way in creating AI systems that are not only powerful but also aligned with ethical principles and regulatory requirements.

# Conclusion

These initiatives reflect Salesforce's proactive approach to responsible AI adoption by embedding internal governance mechanisms—such as the Trusted AI Principles and AI Council—while ensuring external transparency and alignment with regulations like the EU AI Act. They also demonstrate a strategic investment in upskilling through dedicated training resources, enabling both employees and customers to build the critical capabilities needed to thrive in an AI-driven future.

