

# Responsible AI Guidelines 2026

A PRACTICAL HANDBOOK FOR HUMAN-CENTERED AI



servicenow®

# Table of Contents

## Introduction 3

## About this Handbook 4

- Our AI Vision
- Our AI Mission

## Our Four Principles of Responsible AI 8

## Human-Centered AI Assessments 12

- Empower Personas in the AI Lifecycle
  - Responsible/Human-Centered AI Assessments
- Human Collaboration, Feedback, and Oversight
  - Human Collaboration, Feedback, and Oversight Assessment

## Inclusive 20

- Accessibility and Diversity
  - Evaluation for Inclusivity & Bias Assessment

## Transparent 24

- Prioritize transparency, traceability, and contestability
  - Resources
  - Transparency, Interpretability & Contestability Assessment

## Accountable 28

- Accountability Statement
- Mitigate Undesired Behaviors and Outcomes Checklist
- Understanding the Potential Benefits and Harms of AI
  - Exploring impact, risks, opportunities, and trade-offs behind each AI Solution
- Set Realistic Expectations
  - Realistic Expectations Assessment
- Continuous Risk Monitoring and Mitigation
- Evaluate Early and Often
  - Evaluation Assessment
- Data Considerations
- Social and Environmental Impact
  - Social and Environmental Impact Assessment

## Agentic AI 42

- Human Centered Automation
- About Agentic AI
- Agentic AI Best Practices Checklist
  - Agentic AI Assessment

## AI-native Experiences 47

- About AI-native Experiences
  - Five considerations
  - The delegate-review-own framework
  - Responsible friction in AI-native design
- AI-native Best Practices Checklist
  - AI-native Experiences Assessment

## The ServiceNow Responsibility 64

## Closing Statement 65

# Introduction

In the realm of AI, user experience design and research are not merely steps in the process—they are the keystones of innovation. By grounding AI development in the richness of human diversity and real-world complexity, we ensure technology serves us with relevance and humanity.

Human-centered AI is the development and deployment of AI systems that prioritize human needs, values, and well-being.

**This approach involves designing AI systems that are transparent, responsible, and auditable.**

Ultimately, human-centered AI aims to create technology that serves humanity and aligns with our aspirations, goals, and values.

This guidebook serves as a primer for anyone using or working with ServiceNow's AI technology in their careers or businesses.

Readers will learn practical guidance and valuable information on ways to navigate using AI responsibly. It underscores the need for AI systems that prioritize human needs, values, and well-being, advocating for technology that respects and enhances the human experience.

[ServiceNow's Code of Ethics ↗](#)

[ServiceNow's Commitment to Responsible AI ↗](#)

[ServiceNow's Human Rights Statement ↗](#)

[ServiceNow Trust and Compliance Center ↗](#)

# About this Handbook

## WHAT

Practical guidance and decision-making tools

## FOR

People who use ServiceNow's AI technology

## WHY

To create a better world of work through the appropriate and responsible use of AI

## WHEN

Best practices advocate for regular assessments

**This initiative represents our continued dedication to offer practical guidance and consistent decision-making tools to all individuals involved in building, configuring, and consuming ServiceNow's AI-driven products and features.**

Familiarizing yourself with the content in this handbook is strongly encouraged if you are contributing to an AI-driven product experience or considering incorporating AI into your work.

Resource:

[NIST AI RISK MANAGEMENT FRAMEWORK LINK ↗](#)

**As a note, we are using “AI” as a term that includes various forms of AI including generative AI, LLMs (large language models) and agentic AI.**

While this latest version of the Responsible AI Handbook is designed to provide immediate value, we acknowledge that there is ongoing work to be done.

This is **not** an exhaustive set of guidelines and principles and much of it is aspirational, because there are aspects of technology and building responsible AI that is ongoing, in research and development.

As our products evolve, technology improves, and the market shifts, we will continue to iterate, and release updated revisions of this handbook.

**We have also included worksheets with key questions for each section to help guide meaningful discussions and decision-making within your company and teams.**

# Our AI vision

**As we look toward the future, we envision a world where ServiceNow products leverage AI in meaningful ways to drive better outcomes for our customers.**

We envision a seamless integration of AI across our ServiceNow suite, with tools that not only streamline the creation, configuration, and monitoring of AI services but also democratize AI utility, making it accessible irrespective of expertise in data science.

We are building a future that is conscientious about AI ethics, fostering trust, and embedding our corporate ethos into our AI-driven solutions.

**Our journey ahead is focused with purposeful innovation and a commitment to harness the transformative power of AI to elevate human collaboration and creativity in the workplace and enable individuals to work more effectively.**

# Our AI mission

**At ServiceNow, our mission is to use the transformative power of AI within workflows, driving enhanced performance and world-class experiences.**

We provide our valued customers with powerful and reliable AI tools that empower innovation and unparalleled solutions.

We reimagine workflows and augment existing ones with AI-based efficiency gains.

We focus on the people, use cases, and opportunities that benefit most from AI.

With you in mind, we've established a set of comprehensive pillars that steer our AI technology development.

These pillars are founded on the industry's best practices and ensure the creation of trustworthy and responsible AI for you and your customers.

Each of these pillars are meant to guide you through the AI development cycle, from initial creation by our team to practical application by yours, enhancing the way your end - users and customers interact with our services.

# Our Four Principles of Responsible AI



# Our Four Principles of Responsible AI

**Trust is earned in drops but can be lost in buckets.**

ServiceNow recognizes the importance of responsible AI development.

When developing new products and services for our customers, we center our approach around the following principles:

- 1 Human-Centered**
- 2 Inclusive**
- 3 Transparent**
- 4 Accountable**

# 1. Human-Centered

We enable our customers to choose where and when they leverage AI in our products by clearly identifying when AI is being used, keeping humans in the loop.

Our in-product messaging and product documentation clearly identifies when AI is used.

We provide our customers with this guidance, so they understand how to leverage and deploy our AI in a responsible manner.

# 2. Inclusive

We believe in the power of technology to reduce complexity and make the world a more equitable place.

Our AI teams are broadly diverse and test our AI systems with datasets that reflect the diversity of our customers and our end users.

Used responsibly, we believe AI can improve how all humans interact with technology to achieve their goals and aspirations.

## 3. Transparent

We communicate with our customers about AI in a direct manner that informs and provides information on using AI responsibly.

We help provide information to empower AI choices by providing product documentation and model cards that explain our AI systems, in clear understandable terms.

We work to describe the limits and intended usage of AI in our products and err on over-communicating in areas that people care about.


## 4. Accountable

At ServiceNow, **trust** is the cornerstone of our AI operations.

We take our responsibilities regarding AI seriously and have implemented an oversight structure for accountability and governance.

Our team works closely with external experts and the AI community to continually improve our approach, and we have established internal governance bodies to oversee daily product and development operations, so that our customers can utilize ServiceNow's AI capabilities with confidence.

# Human- Centered AI Assessments



# Empower Personas in the AI Lifecycle

**It is important to empower our customers to make informed, intentional decisions by being aware of dilemmas, trade-offs, risks, and opportunities associated with AI.**

We recognize that the AI landscape is continually evolving, and we are committed to growing and solidifying our approach over time.

ServiceNow's role is to lead in establishing, managing, and expanding governance frameworks, educational resources, documentation, and methods that empower you to use AI responsibly and prevent misuse.

**Together, we can navigate the dynamic world of AI with confidence and integrity.**

# Responsible/Human-Centered AI Assessments

**In our commitment to fostering responsible AI innovation, we offer these assessments to assist you throughout the lifecycle of your AI project.**

This tool is crafted to inspire and guide collaborative discussions among your team and stakeholders, especially at the critical early stages of evaluating AI-enabled products.

It's structured to catalyze the kind of deep, inclusive conversations that are essential for navigating the complexities of AI responsibly.

Acknowledging the dynamic landscape of AI applications, we embrace the notion that adaptability is key.

**These assessments are designed as a modular framework, meant to be customized to suit the evolving needs and vision of your projects.**

The questions in each section are designed to spark meaningful conversations and assessments that will empower your decision making.

As technology advances, we will constantly refine our methods, upholding a flexible and ethical AI development process that adapts to change.

**The importance of a diverse team composition cannot be overstated, given the wide range of expertise required to address the multifaceted challenges of AI.**

We advocate for ongoing engagement and end-user feedback with diverse audiences throughout your project's development, to ensure that your strategies remain effective and responsive to new insights and changing circumstances.

**Please consider all following assessments as a starting point for a journey towards more informed and collective decision-making in AI development.**

It is designed with a focus on fostering a comprehensive understanding of AI responsibilities and encourages adapting the questions and considerations it proposes based on how AI is being utilized within your specific context.

[Responsible AI FAQ ↗](#)

# Responsible/Human-Centered AI Assessment 1/2

KEY CONSIDERATION	ANSWER (TEAM'S RESPONSE GOES HERE)	SUPPORTING EVIDENCE OR DECISION PROCESS (SPECIFY HOW TEAM CAME TO THIS DETERMINATION)
Is there a significant risk of unintended outcomes or harm by this application?		
Who will primarily be using this application? Who will benefit?		
Could someone be harmed or disadvantaged by this application and why?		
What unintended outcomes could occur? Which are the biggest risks and who could they possibly impact? E.g., your employees, your customers, your partners?		
Can this AI product or feature be used by everyone? If not, why? And will those who are excluded be impacted?		
Please rank your top three AI values/goals and describe how you'd balance trade-offs among them.		
Who are your stakeholders, and can you list what level of information/documentation each would need to inform their decision-making?		

# Responsible/Human-Centered AI Assessment 2/2

<b>KEY CONSIDERATION</b>	<b>ANSWER (TEAM'S RESPONSE GOES HERE)</b>	<b>SUPPORTING EVIDENCE OR DECISION PROCESS (SPECIFY HOW TEAM CAME TO THIS DETERMINATION)</b>
Do the people monitoring the AI system's outcomes, have awareness of any safeguards that have been put in place and a way to communicate any unintended biases or unintentional outcomes?		
What safeguards have been put in place to help reduce the risk of over-trusting the AI system or becoming complacent when monitoring or auditing the AI system?		
What undesirable biases have been identified? What do you know and what is still unknown?		
To the best of your ability, have you mitigated any known risks or concerns, especially related to aspects of responsible AI (accessibility, inclusion, cultural appropriateness, etc.)?		

# Human Collaboration, Feedback, and Oversight

**Human involvement is crucial in developing responsible AI. Our AI systems are designed to complement human expertise, involving human input and oversight to safeguard subsequent actions and outcomes.**

Continuous improvement in AI requires regular monitoring and updates, relying on user research and human expertise for valuable insights and feedback. This ensures our AI systems can evolve and adapt effectively.

We are committed to the ongoing enhancement of our AI systems, recognizing the value of diversity in their training to cultivate robustness and adaptability.

We carry out user experience research and consistently engage with our customers to collect feedback on our AI experiences. This is part of our broader effort to establish review mechanisms that not only gather but also integrate human feedback on both our AI product experiences and AI outputs.

This approach is designed to identify and address unintended outcomes or limitations, to ensure that our AI systems evolve to meet user expectations and embody our ethical commitments.

**Human oversight is key to responsibly deploying AI, reflecting our dedication to align with ethical and societal values, and elevate the human experience with AI.**

# Human Collaboration, Feedback, and Oversight Assessment

This portion of the assessment focuses on how the team will ensure productive and positive human-AI interactions through feedback and oversight.

KEY CONSIDERATION	ANSWER (TEAM'S RESPONSE GOES HERE)	SUPPORTING EVIDENCE OR DECISION PROCESS (SPECIFY HOW TEAM CAME TO THIS DETERMINATION)
Is the scope of what needs to be explained to the user defined?		
Does the user need to have technical or domain knowledge to use the solution?		
What kinds of implicit and explicit feedback are collected from people?		
Is behavioral or passive feedback collected to augment any explicit signals on engagement or feedback with the AI?		
Do people have any control over an AI system they may be interacting with or impacted by? Are they informed how to tailor the AI to meet user needs and adhere to ethical guidelines?		
Are there any occasions where a human is not in control of decision making? Discuss the risks and implications with your team. If AI were to make the wrong decision, what are the consequences?		

# Inclusive



# Accessibility and Diversity

When integrating AI into your enterprise operations, the focus should be on its real-world impact on the people who will use and interface with our technology.

**At ServiceNow, our commitment is to make AI technology practical and accessible for everyone, ensuring it serves a wide range of needs effectively over its entire usage period.**

**We understand the importance of inclusivity, diversity, accessibility, and equality in our AI development process.**

The dimensions of diversity to consider will likely vary by use case and which layer in the application stack that we evaluate (i.e., data governance, AI system training and testing, user experience).

It is important to think about diversity dimensions by interacting with stakeholders and with the broader research community to identify as many relevant aspects as possible. It is important to acknowledge and discuss the trade-offs that come with using AI.

This helps us have meaningful conversations with customers about managing challenges, rather than just trying to wish them away.

**For instance, while we can't completely remove bias, we can take steps to recognize and constrain it.**

### **Examples of dimensions include:**

- Time in role
- Technical ability and understanding
- AI training and understanding
- Cultural background
- Educational and professional backgrounds
- Geographic location
- Gender
- Ethnicity
- Age
- Education
- Language
- Disabilities

[ServiceNow's Accessibility statement](#) ↗

# Evaluation for Inclusivity and Bias Assessment

This part of the assessment is focused on evaluating AI systems for bias with earlier assessment items around inclusivity and diversity.

KEY CONSIDERATION	ANSWER (TEAM'S RESPONSE GOES HERE)	SUPPORTING EVIDENCE OR DECISION PROCESS (SPECIFY HOW TEAM CAME TO THIS DETERMINATION)
Does your team integrate the perspectives of individuals with diverse abilities into your review processes?		
Can your AI features be consumed with equal effectiveness by people with different abilities? If not, which mitigation actions have been taken?		
Should the AI system be assessed for bias, and if so which kinds, and how?		
Does your team continue to monitor for undesired biases in production after deployment?		
Which steps did you take to reduce or compensate for undesired biases?		
What feedback mechanisms are in place for users to report on the inclusivity, biases, and accessibility of their user experience?		

**Transparent**



# Prioritize transparency, traceability, and contestability

We aim to develop AI solutions and provide tools that enhance your control and understanding.

**Our tools and interfaces are built to ensure transparency, traceability, and the ability for you and your customers to challenge and understand the AI's impact on your operations.**

Instead of dictating actions, we present options, maintaining your and your customers' independence in making decisions. Our system provides methods to review outputs, provide feedback, and, if necessary, adjust AI-generated decisions, prioritizing autonomy and informed choice.

**We commit to creating AI that is not just effective but also ethical, prioritizing your and your customers' rights to understand and influence AI decisions, with a strong focus on transparent practices.**

# Resources

Here are a few resources we provide to increase transparency in our AI products.

[Model Cards ↗](#)

[NIST AI RISK MANAGEMENT FRAMEWORK ↗](#)

[Product Documentation ↗](#)

[Responsible AI FAQ ↗](#)

# Transparency, Interpretability, and Contestability Assessment

This portion of the assessment is focused on how to discuss and incorporate transparency, interpretability, and contestability into your AI solution.

KEY CONSIDERATION	ANSWER (TEAM'S RESPONSE GOES HERE)	SUPPORTING EVIDENCE OR DECISION PROCESS (SPECIFY HOW TEAM CAME TO THIS DETERMINATION)
Which personas are involved in installation, use, and monitoring of the AI solution?		
Is every aspect of the AI deployment process documented as appropriate to be referenced by customers and people using the AI system?		
Do all personas fully understand in a demonstrable way that they are working with an AI technology?		
Do individuals whose data is being used by the AI understand how their data is being used?		
Can users contest AI outputs and outcomes as appropriate?		
How will the team measure user trust, confidence, and learning over time?		

**Accountable**



# Accountability Statement

**ServiceNow's AI-enabled products are designed to perform within the company's technical standards while maintaining individual accountability.**

Individual accountability includes understanding the role of AI and the role of the human across decisions in the AI product lifecycle.

ServiceNow recognizes the need for safeguards to mitigate and prevent adverse or unjust impacts of our AI technology, ensuring compliance with local laws and guidelines.

We encourage our valued customers to diligently follow our product and service guidance and procedures when implementing ServiceNow technology.

AI systems hold the potential to greatly enhance the efficiency of our customers' operations. We place a high premium on integrating feedback from our customers and their users of our AI services.

To foster continuous improvement and ensure an adequate understanding of AI limitations among our customers, we are committed to:

- 1 Gaining a comprehensive understanding of the AI system's intended purpose and expected outcomes.**
- 2 Designing systems and user experiences that encourage users to leverage AI appropriately, ensuring the right balance between human and AI collaboration.**

**In this way, we strive to provide AI solutions that not only optimize efficiency but also uphold accountability, transparency, and a positive user experience for you and your customers.**

[Human Rights Statement ↗](#)

[Responsible AI FAQ ↗](#)

# Mitigate Undesired Behaviors and Outcomes Checklist

Here are ways that you can help in the mitigation of undesirable outcomes.

- Customization and Personalization**  
Tailor AI systems to your preferences and requirements and ensure data readiness.
- Feedback and Reporting**  
Monitor and review AI outputs, provide feedback, and report issues.
- Regular Updates**  
Keep AI systems up to date with the latest software and security updates.
- Human Oversight**  
Maintain human review of AI outputs and decision-making, especially in critical areas.  
Ensure that the humans creating and monitoring AI systems represent the diverse user population that will use them.
- Data Privacy**  
Ensure appropriate and robust data privacy and security measures are in place.
- Continuous Learning** Now Learning ↗  
Review product documentation and model card if available thoroughly prior to using an AI system. Stay informed about advancements and best practices.
- Monitoring and Auditing**  
Implement processes to detect and address undesired AI behaviors.
- Collaborate with AI Providers**  
Engage in ServiceNow's product discussion and research programs to share insights and provide feedback. E.g. PAC, UX research lab.
- User Training**  
Educate users and stakeholders on effective AI interaction.
- Proper use**  
Ensure that the AI system is being used as intended and not being exploited in harmful ways.

# Understanding the Potential Benefits and Harms of AI

Exploring impact, risks, opportunities, and trade-offs behind each AI Solution

The first question in any responsible approach to AI is to ask ourselves (and broader team/stakeholders) who may benefit and who may be harmed by this AI solution.

**We believe that the humans who use and are impacted by our technology should take precedence over the innovation of technology.**

We should understand the problems we are trying to solve, what the processes are, and provide a baseline for the performance of the intended automation efforts.


We should not seek to use AI for every customer problem or opportunity, but rather have a mature understanding of why AI is the right mechanism for the outcomes we and our customers need.

# Set Realistic Expectations

**As a responsible AI organization, we are committed to maintaining a secure approach to data and careful development of our AI systems.**

To ensure transparency and build trust with our customers, we strive to communicate realistic performance standards and capabilities of our AI systems.

We will also be transparent about the limitations of our AI systems and data and ensure that the degree of uncertainty in our AI systems is always communicated clearly.



# Realistic Expectations Assessment

This portion of the assessment helps teams consider how they will communicate realistic expectations to customers/people using the technology as well as how to grow user knowledge over time.

KEY CONSIDERATION	ANSWER (TEAM'S RESPONSE GOES HERE)	SUPPORTING EVIDENCE OR DECISION PROCESS (SPECIFY HOW TEAM CAME TO THIS DETERMINATION)
What are your expectations from this tool and which metrics are a priority?		
How "bad" are errors, and which errors are particularly bad for the people using the AI system? What is the worst-case scenario?		
Could you offer AI system owners metrics to track usage, accuracy of AI outputs, actions taken, and outcomes?		
Does the AI solution allow a phased approach to help people build confidence? (e.g., enabling assignment group predictions made by an AI to show up as recommendations first, rather than auto filling the fields.)		

# Continuous Risk Monitoring and Mitigation

In addition to proper governance and visibility, we are working to implement tools and controls that mitigate the risk of biased or otherwise undesired, hallucinated, and/or toxic outputs.

We are working to provide the opportunity in the user experience and infrastructure to:

- 1 Detect and allow people to report inappropriate or unintended behaviors**
- 2 Support response processes to mitigate the risk.**
- 3 Ensure adherence to applicable regulations, including high-risk considerations outlined in frameworks such as the EU AI Act.**


Moreover, we invite you to be a part of the solution by actively participating in the monitoring process. Collaborate with other business units, share insights, and help educate your peers on where they can help review AI outcomes.

**We emphasize the importance of involving diverse stakeholders to review AI outputs and provide valuable feedback. Together, we can ensure that our AI solutions meet the highest standards of reliability, and human satisfaction.**

# Evaluate Early and Often

**Consistent with foundational best practices, begin with simpler use cases to facilitate frequent and early assessment of AI performance.**

Early evaluation enables the identification of potential issues and the implementation of necessary safeguards to ensure the responsible use of our technology, including addressing edge cases and mitigating risks.



# Evaluation Assessment

This portion of the assessment is focused on ensuring the team has a plan to regularly review their technology against changing expectations and landscapes.

KEY CONSIDERATION	ANSWER (TEAM'S RESPONSE GOES HERE)	SUPPORTING EVIDENCE OR DECISION PROCESS (SPECIFY HOW TEAM CAME TO THIS DETERMINATION)
What are the limitations of your evaluation protocol? E.g., what are you not measuring?		
How often are you evaluating the performance?		
Do you have a baseline for your performance? (E.g., is it measured against the existing processes, or if it's a new process, is it an improvement over time?)		
Do the outcomes of the AI system align with your prioritized business values?		

# Data Considerations

**We place a strong emphasis on prioritizing your control, data privacy, and agency.**

We understand the importance of transparency in how your data is used, and that's why we provide you with full access to your raw data for your review.

**When it comes to harnessing the full potential of AI within your organization, ensuring data readiness and adherence to best practices is of utmost importance.**

To garner the most of value from your ServiceNow AI-enabled product, consider the following data practices:

### **DATA READINESS**

Begin by assessing the quality, completeness, and relevance of your data. Ensure that it aligns with the objectives and scope of your AI implementation.

### **DATA PREPARATION**

Prepare your data by cleaning, transforming, and normalizing it as needed. This step is vital for enhancing the accuracy and performance of AI systems.

### **DATA LABELING (when applicable)**

At times, labeling and annotating data may be required. It is important to limit biases during this process and to be meticulous, to provide correction and needed information to better inform the AI system. Define clear labeling guidelines to ensure consistency.

### **DATA USAGE**


Understand what data an AI product may need and ensure that the correct data applies to the intended use cases. Ensure that data usage complies with data privacy and regulatory requirements.

# Social and Environmental Impact

**Considering societal and environmental impacts is important at every stage of AI product development.**

It's essential to evaluate the spectrum of effects on both direct users and those indirectly exposed to the technology.

The systems should be beneficial for themselves and any other people these systems impact.




# Social and Environmental Impact Assessment

This portion of the assessment focuses on considering our AI solutions’ social and environmental impact.

KEY CONSIDERATION	ANSWER (TEAM’S RESPONSE GOES HERE)	SUPPORTING EVIDENCE OR DECISION PROCESS (SPECIFY HOW TEAM CAME TO THIS DETERMINATION)
Does the AI product have the potential to impact how someone feels or thinks about themselves?		
Does the AI product compromise the safety of any individual?		
Does the AI product have the potential to impact how someone may feel or think about someone else?		
Does the AI product require the use of computing resources that may impact the environment?		
How will the AI product impact company sustainability targets?		
Could the AI be used as part of any system that is purposefully lethal to individuals?		
Could the AI be used as part of a system that is purposefully harmful to individuals?		
Does the AI product create possible negative outcomes for the environment or society?		
How can we predict the impact of these systems on job displacement and other affected areas?		

# Agentic AI



# Human-Centered Automation

As we embrace the power of intelligent systems, we recognize that some tasks are repeatable and, in many cases, can be automated to improve efficiency and reduce complexity.

However, automation must be designed responsibly, with clear guardrails, meaningful safeguards, and thoughtful human oversight. Agentic AI represents a shift toward systems that take action in more proactive and goal-directed ways.

**Our role is to ensure these AI agents operate in service of people by reinforcing human intent, aligning with business goals, and maintaining trust across the workflows they support.**

As agentic workflows become more common and complex, we increasingly recognize the need to elevate human-centered considerations, reinforcing transparency, oversight, and alignment with human values.

# About Agentic AI

As artificial intelligence evolves, the term **agentic AI** refers to systems that go beyond generating content or offering suggestions.

**Agentic AI systems are designed to take purposeful action on behalf of a person or organization.**

**These AI agents operate with a defined role, objective, and behavior pattern—interacting with people, data, and systems to complete tasks or drive outcomes within a given context.**

Unlike generative AI, which focuses on producing outputs like text, images, or code, agentic AI is built to **autonomously initiate and execute upon assigned actions.**

For example, while generative AI might help write an email draft, an agentic AI system could detect the need for outreach, generate the content, send the email, and update the relevant record in a workflow—potentially without requiring repeated human approval or rejection, depending on the guardrails or rules established.

**Whether AI agents are resolving tasks or managing cross-functional approvals, ServiceNow’s agentic AI approach is to operate within human-defined boundaries, with transparency, oversight, and alignment at the core.**

# Agentic AI Best Practices Checklist

Here are questions to help guide the responsible design, deployment, and oversight of AI agents.

**These practices support human-centered autonomy by prioritizing clarity, control, and accountability.**

## **Have you clearly defined the AI agent's role and purpose?**

- Is the AI agent's role, purpose, and intended behavior clearly defined and written in plain language that is easy to understand?
- Is there a clear business justification supporting a defined workflow need for the AI agent?
- Is the AI agent's objective aligned with user and business goals?

## **Are you starting with assistive capabilities to build trust?**

- Can people review, configure, or understand the AI agent's planned actions before they are executed when appropriate?

## **Can designated individuals maintain oversight and intervene when needed?**

- Are people able to undo, pause, redirect, halt, or safely terminate the AI agent's actions if necessary?
- Are there appropriate checkpoints or human-in-the-loop interventions established for sensitive, complex, or ethically ambiguous decision points?
- Are the AI agent's actions and decisions understandable at an appropriate level for oversight?

## **Is there a clear process for feedback and continuous improvement?**

- Can people easily report errors or misalignments in real time or near real time?
- Are feedback channels regularly reviewed and acted upon?

## **Are the AI agent's limitations and boundaries clearly communicated?**

- Do people understand what the AI agent can and cannot do, including what tasks it performs well and its known limitations, before they begin using it?
- Are there visible cues, documentation, or onboarding guidance?
- Can people contest or correct AI agent-driven outcomes?

# Agentic AI Assessment

This worksheet is intended to help teams evaluate agentic AI features or systems before and during deployment. It is designed to prompt critical thinking, facilitate cross-functional dialogue, and ensure AI agents are developed with intentionality, oversight, and user alignment.

KEY CONSIDERATION	ANSWER (TEAM'S RESPONSE GOES HERE)	SUPPORTING EVIDENCE OR DECISION PROCESS (SPECIFY HOW TEAM CAME TO THIS DETERMINATION)
What is the intended role or purpose of this AI agent? (Include any areas the AI agent should not engage with or take action on.)		
Who is the AI agent supporting, and what value does it create?		
What level of autonomy is the AI agent designed to operate with (collaborative with human oversight, or autonomous within defined boundaries)?		
Can designated people override, pause, or redirect the AI agent's actions?		
What mechanisms are in place to gather feedback and report misalignments?		
How transparent are the AI agent's operations, and who has access to this information for review or oversight?		

# AI-native Experiences



# What is AI-native?

## Definition of AI-native

A product or experience fundamentally designed around AI capabilities from the ground up. The experience reimagines workflows based on what AI makes possible and would not exist in its current form without AI.

*Example:* A traditional IT service desk routes tickets through predefined rules. An AI-native service desk might instead interpret the human's intent through conversation, gather relevant context automatically, resolve common issues without routing, and escalate with full context when human judgment is needed.

## The distinction from AI augmented

AI augmentation adds AI features to existing workflows (the core experience follows conventional patterns). AI native rethinks the experience entirely with AI as a foundational capability.

*Example:* An AI-augmented experience might surface recommended knowledge articles alongside a support case. An AI-native experience might diagnose the issue, draft a resolution, and present it for the agent's review before the agent opens a single article.

## Why this matters to people

Understanding the difference between AI augmented and AI native helps organizations set appropriate expectations, plan adoption strategies, and evaluate whether products are delivering on the intended value or the appropriate use of AI capabilities. In practice, AI native is a spectrum rather than a binary state, and a single product may sit at different points on that spectrum across different flows or tasks.

*Example:* Adopting an AI-native experience may require rethinking team workflows, roles, and oversight practices, while an AI-augmented tool can often be adopted within existing processes. Knowing which you are working with helps set realistic expectations for both the technology and the people using it.

## Connection to agentic AI

The agentic AI section covers responsible practices for autonomous AI agents. AI-native experiences represent the broader design paradigm within which agentic systems operate, encompassing how people interact with, oversee, and collaborate alongside AI across entire workflows.

*Example:* An AI agent that autonomously resolves IT incidents operates within an AI-native service experience that determines how people are notified, how they oversee the agent's work, and how they intervene when needed.

# Five Considerations for AI-Native Experiences

AI-native experiences introduce fundamentally different dynamics for the people using them. Traditional software requires people to initiate every action, with systems responding in turn. AI-native systems, in contrast, can initiate, anticipate, and adapt across those networks. This inversion creates new considerations that should be addressed throughout the design and deployment process.

The following five considerations offer a lens for evaluating AI experiences. Teams can use them to assess whether their implementation reflects AI-native thinking or whether AI capabilities have been layered onto existing patterns. This distinction is not binary, and the right approach will vary by context, but the considerations can help guide the conversation.

As with the rest of this handbook, these are intended as practical guidance to support responsible design, not as exhaustive requirements.

These considerations and the framework that follows are not set-once decisions. The appropriate balance of autonomy, friction, and human oversight may shift as systems learn, as organizational context changes, and as regulatory environments and implementation maturity evolve.

**Ongoing reassessment is part of responsible AI-native design.**

# 1. Attention Awareness

It should be clear to people when AI requires their input versus when it is working autonomously, through signals appropriate to the context and their needs.

AI-native systems can invert traditional monitoring, with AI requesting attention when needed rather than requiring constant human checking. How escalation is determined, whether by the AI itself, by rules defined outside the system, or by a combination, is itself a design decision with implications for reliability and oversight.

A guiding principle is to match the urgency of a notification to the importance of the decision, so that significant moments stand out and routine activity stays in the background.

*Example:* An AI-native system might silently auto-resolve routine requests, surface a summary of unusual patterns for awareness, and interrupt only when a critical decision requires human input. A password reset is handled automatically, whereas an unusual access request from an unfamiliar location could prompt a person to review.

## 2. Status Transparency

It is helpful for people to have appropriate visibility into what AI is doing and why, including context about the reasoning or data informing AI suggestions or actions, where feasible.

Traditional software has predictable, visible steps. AI-native systems may need to surface reasoning that is not inherently visible.

This is a design challenge: the mechanisms used to surface reasoning, such as chain-of-thought traces, are not always reliable representations of why a system produced a given output and may not be intelligible to all users without thoughtful design.

This extends the Transparent principle into the specific context of AI-native workflows.

*Example:* When an AI system takes several steps to resolve a request, it might show what it is currently doing ("reviewing similar past cases"), what it found ("3 matching cases from the last 90 days"), and where it needs input ("the recommended resolution differs from the standard process; how would you like to proceed?").

This gives a person enough context to stay informed without requiring them to monitor every step, with the ability to go deeper when they need to understand more.

# 3. Capability Discovery

People should have the opportunity to discover what AI can do for them, beyond immediate tasks.

AI capabilities are contextual and expansive. People may not be able to discover them through traditional menu exploration alone.

Contextual, progressive revelation of capabilities during active use can support better adoption and more effective use.

*Example:* An employee who manually reassigns cases to the same team each week might receive a suggestion: "You reassign these cases regularly. Would you like AI to handle this automatically, based on the criteria you've been using?"

The capability is surfaced when it is relevant, rather than explained during onboarding and forgotten.

# 4. Trust Calibration

Where feasible, people should have access to information that helps them assess AI output reliability, such as data sources, known limitations, or the basis for a recommendation, so they can make informed decisions about when to verify, contest, or override AI output.

Teams can support this by designing and testing how reliability and uncertainty are communicated, rather than relying on technical metrics alone to drive user trust.

Communicating the basis for AI outputs in ways people can interpret and act on helps prevent both over- and under-reliance. Confidence indicators, where used, should be tested for comprehension and should not be assumed to calibrate trust on their own.

*Example:* Rather than displaying a raw confidence score like "87%," an AI-native experience might explain the basis for its recommendation: "This resolution is based on a strong match to 12 similar cases resolved in the last 60 days."

For a less certain suggestion, it might say: "This situation has limited precedent. You may want to review before proceeding." A person can then decide how much scrutiny to apply based on context they can understand.

# 5. Human-AI Collaboration

People should feel in control when benefitting from AI-initiated actions.

In many cases, an effective experience is a collaborative partnership with appropriate human oversight.

This includes the ability to stop or pause, redirect, adjust, and formally challenge AI decisions.

Preferences and autonomy settings should remain adjustable over time, so that people can recalibrate the collaboration as their needs, context, or comfort level change.

*Example:* An AI system that drafts customer responses might present the draft for review, allow the agent to edit or override it, and remember the agent's preferences over time ("you always adjust the tone for this customer segment, would you like me to apply that going forward?").

If the agent disagrees with a routing decision, there is a clear path to redirect the case and provide feedback on why.

# The Delegate-Review-Own Framework

## The Framework

A practical model for determining what AI should handle autonomously, what requires human review, and what remains a fundamentally human responsibility.

## Delegate

Tasks AI handles autonomously. Typically well-specified, repeatable actions within established guardrails. What qualifies as low-risk or reversible will vary by organization, industry, and regulatory context.

## Review

Tasks where AI proposes and humans verify. Outputs that require judgment, moderate-stakes decisions, or domain expertise validation.

## Own

Tasks where human judgment and responsibility are central. These often include strategic decisions, novel problems, high-stakes or irreversible actions, and ethical judgments.

## **The Autonomy Spectrum**

A sliding scale from AI-assisted to AI-autonomous, influenced by reversibility, stakes, novelty, AI confidence level, and regulatory or compliance requirements.

## **The five interaction patterns across the autonomy spectrum**

These describe different ways humans and AI can share responsibility. They are not sequential. A single product or workflow may use different patterns for different tasks, and the most appropriate pattern will depend on the context, stakes, and maturity of the implementation.

**Operator:** A human controls all actions; AI executes specific commands.

**Collaborator:** A human and AI work together interactively.

**Consultant:** AI provides advice and recommendations, a human decides.

**Approver:** AI proposes complete actions, a human approves or rejects.

**Observer:** AI acts autonomously, a human monitors and intervenes on exception.

## **Applying the Framework**

Questions to ask when designing or configuring AI features:

- What is the worst outcome if AI gets this wrong?
- How reversible is the action?
- Does the human have context that the AI lacks?
- What is the cost of requiring human approval?
- Can the AI's output be made visible to the human approver or observer at the right time and in a format they can act on?

# Responsible Friction in AI-native Design

**As AI systems take on more responsibility within workflows, the moments where people pause, review, and engage become more important, not less.**

**Intentional friction, when applied thoughtfully, is a way to keep people informed and in control at the moments that matter most.**

## The case for intentional friction

Most AI design optimizes for speed and seamlessness. While the intention is to avoid unnecessary friction, "frictionless" as an unexamined goal can create problems. When AI handles everything seamlessly, people can lose situational awareness and struggle to recover when things go wrong. This is especially relevant in enterprise contexts where AI-driven decisions can affect employee experiences, customer outcomes, or compliance obligations.

*Example:* A frictionless AI system that automatically reclassifies and closes support cases without any review step may improve speed metrics, but if a case is misclassified, the affected person may not know it happened, the error may not surface until downstream consequences appear, and the team loses the opportunity to catch patterns of misclassification early.

## When friction adds value

Friction can strengthen decisions in situations such as:

- High-stakes or irreversible decisions
- Situations where human judgment matters
- Moments where people are still building or adjusting their trust in the AI
- Transitions between AI autonomy levels
- Novel or unusual situations where AI reliability is lower

*Example:* An AI system processing a standard equipment request might proceed automatically. The same system processing an accommodation request that involves personal employee information might pause and present the relevant details for a human to review before acting, because the stakes, sensitivity, and potential for error are meaningfully different.

## The design question

Does this friction make the human more informed, safer, or more deliberate? If yes, it is adding value. If no, it is adding frustration. Friction should be proportional to the stakes of the decision and consistent enough that people can learn what to expect.

*Example:* A confirmation step that asks, "Are you sure?" on every action, regardless of consequence, trains people to click through without reading.

A confirmation step that appears only for irreversible actions and surfaces what will happen ("This will permanently reassign 47 open cases to a new team. This cannot be undone.") prompts genuine reflection because it appears selectively and provides specific, relevant information.

## Connection to existing principles

Intentional friction supports the Accountable principle (realistic expectations, continuous monitoring) and the Human-Centered principle (keeping humans meaningfully engaged in decisions rather than passively approving automated actions).

Thoughtful friction at the right moments can strengthen accountability and build user confidence over time.

# AI-native Best Practices Checklist

These questions are intended as a starting point for design discussions, not a compliance checklist. Not every item will apply to every feature, and the right depth of consideration will vary by context.

## Have you established which AI autonomy level is appropriate for the intended feature or workflow?

- Is there a clear rationale for what AI handles autonomously, what it proposes for review, and what remains a human responsibility?
- Have you considered the reversibility, stakes, and novelty of each action?

## Do people have visibility into what AI is doing and why?

- Can people understand AI reasoning at a level appropriate for their role?
- Is status communicated during longer-running AI processes?
- Are people informed both when AI is working autonomously and when it needs their input?

## Can people discover AI capabilities in context?

- Are relevant capabilities surfaced when they are useful, in addition to being explained during onboarding or in documentation?
- Are limitations communicated proactively, before people encounter them through failure?

## Does the experience help people calibrate trust appropriately?

- Can people distinguish between high-confidence and low-confidence AI outputs?
- Is verification encouraged where it adds value, without requiring verification of everything?

## Do people maintain meaningful control and the ability to challenge AI decisions?

- Can people stop, pause, undo, or redirect AI actions?
- Is there a clear process for challenging AI decisions?
- Does human feedback influence future AI behavior?

## Is friction proportional to the stakes of the decision?

- Do high-stakes or irreversible actions include appropriate checkpoints?
- Are confirmation steps meaningful (prompting actual reflection) rather than habitual?
- Can people adjust the level of AI autonomy based on their confidence, context, and needs?

# AI-native Experiences Assessment 1/2

These questions are intended as a starting point for design discussions, not a compliance checklist. Not every item will apply to every feature, and the right depth of consideration will vary by context.

KEY CONSIDERATION	ANSWER (TEAM'S RESPONSE GOES HERE)	SUPPORTING EVIDENCE OR DECISION PROCESS (SPECIFY HOW TEAM CAME TO THIS DETERMINATION)
Is this experience AI native or AI augmented? What is the rationale for the distinction, and how does it shape your design and deployment approach?		
At which autonomy level (Operator, Collaborator, Consultant, Approver, Observer) is each AI-driven feature or workflow designed to operate? What factors informed this decision?		
How does the experience address each of the five user needs: Attention awareness, Status transparency, Capability discovery, Trust calibration, and Human-AI collaboration?		
Where have you introduced intentional friction, and where have you prioritized seamlessness? What criteria guided those decisions?		

# AI-native Experiences Assessment 2/2

These questions are intended as a starting point for design discussions, not a compliance checklist. Not every item will apply to every feature, and the right depth of consideration will vary by context.

**KEY  
CONSIDERATION**

**ANSWER  
(TEAM'S RESPONSE GOES HERE)**

**SUPPORTING EVIDENCE OR DECISION PROCESS  
(SPECIFY HOW TEAM CAME TO THIS DETERMINATION)**

How do people discover what the AI can and cannot do? Are expectations set through demonstration and contextual guidance rather than documentation alone?

What mechanisms are in place for people to challenge, override, or provide feedback on AI decisions? How is feedback captured, reviewed, and used to inform future system behavior, governance, or product improvements?

# The ServiceNow Responsibility

ServiceNow is committed to designing and developing products that embody our four principles of responsible AI.

- 1 **Human-Centered**
- 2 **Inclusive**
- 3 **Transparent**
- 4 **Accountable**

**Embracing responsible AI is a shared responsibility between ServiceNow and our customers.**

This guidebook, along with its accompanying tools, is designed to assist our customers in navigating the auditing and management of AI systems for enhanced accountability and ethical practices.

As our company, technology and customer landscapes evolve, we will update and adapt how we articulate and embody these principles.

**Our principles are a living reflection of our company's stance on AI and responsibility.**

# Closing Statement

**This guidebook is crafted to be a dynamic and evolving resource, designed to advance in tandem with the ongoing progress and transformative potential of AI.**

Grounded in our core human-centric principles, our intention is to pave a path for AI that not only meets technological expectations but also enriches human potential and respects individual values and diversity.

Journeying forward, enhancing and refining these guidelines will be an iterative process—one that flourishes through active engagement and feedback from you, our valued customers and partners.

Your insights are vital to ensuring our AI technologies remain as responsive and beneficial as possible.

We deeply appreciate the time you have invested in understanding our approach to responsible AI.

**Thank you.**

April 2026