ATLASSIAN

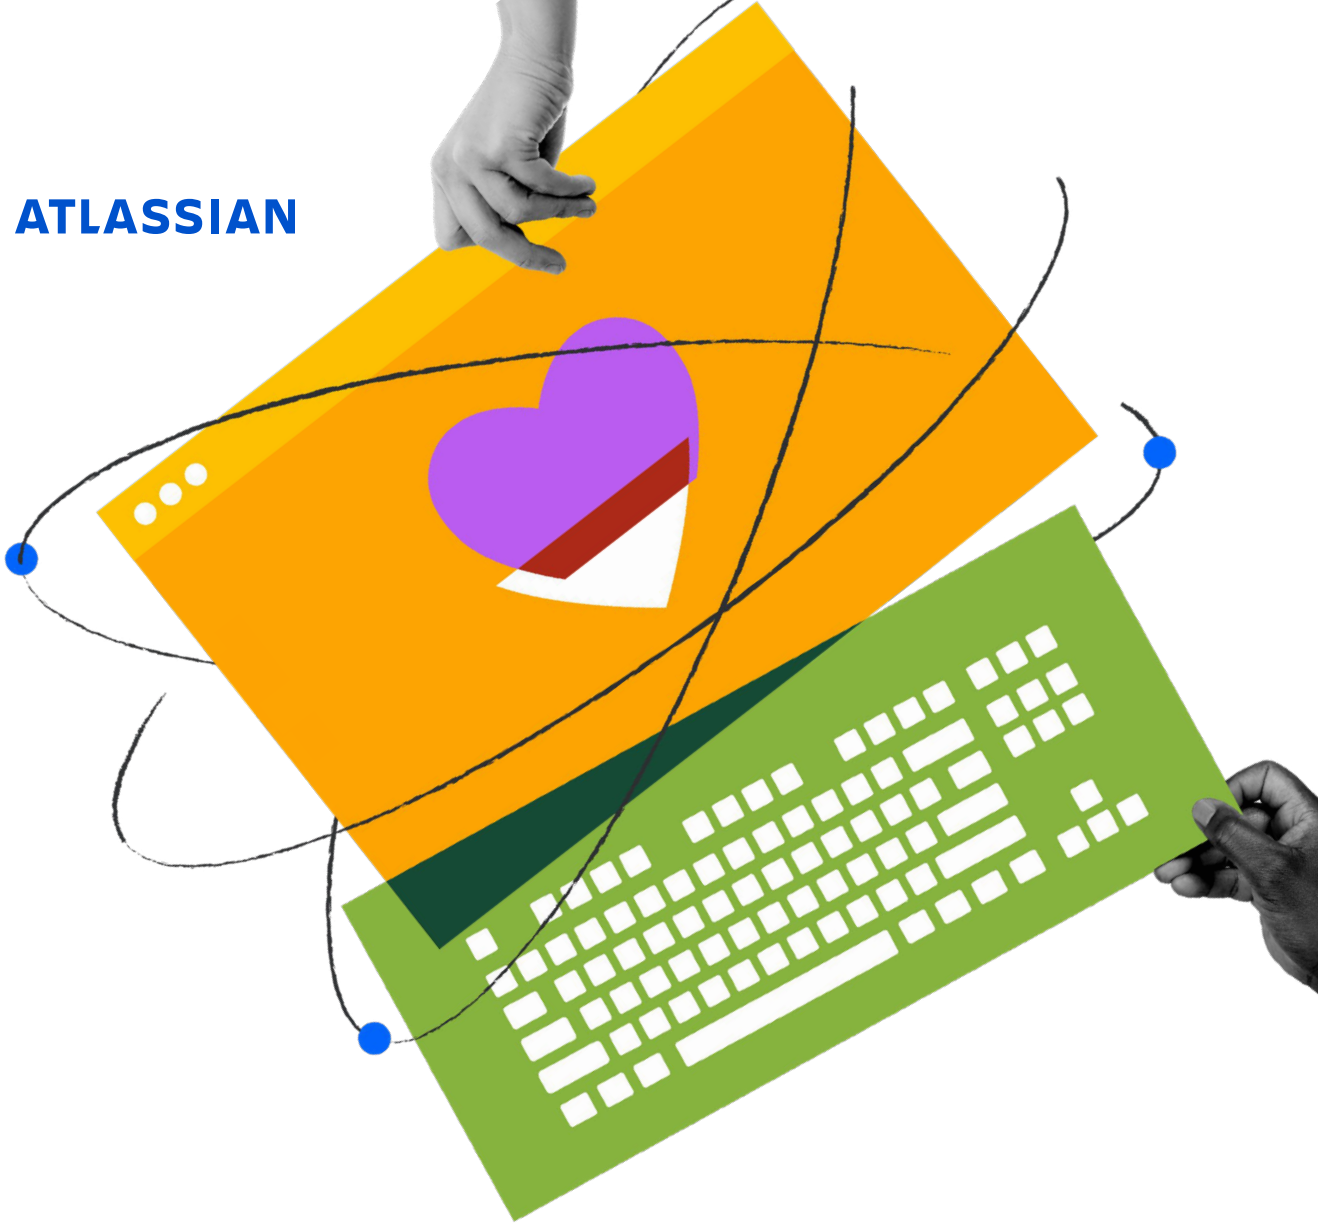

# Our No BS Guide

## to Responsible AI Governance

Playing responsibly, as a team

# Table of contents

# Atlassian's mission is to unleash the potential of every team.

As a company focused on teamwork and collaboration, we strongly believe AI has the power to supercharge teams and lead to better outcomes for our communities. But with this exciting opportunity comes the serious responsibility to develop these technologies thoughtfully and deploy them with care.

We also believe this responsibility is a shared one: responsible technology is impossible alone.

This is partly because the current landscape can be overwhelming. AI tools are increasingly agentic and self-directed. Meanwhile, between existing regulatory frameworks and the wave of new laws, standards, and guidance, it can be hard to navigate the rules of the game.

Atlassian knows this all too well. We've gone through the challenging process of translating the outcomes we're aiming for into the guardrails, processes, and practices needed to achieve them – then embedding them across an entire organization.

That's why we are taking an open, collaborative and, above all, iterative approach to AI governance. Atlassian has a rich tradition of sharing our own practices in the name of unleashing other teams' potential. To further this tradition, we are committed to sharing and updating the tools we use to hold ourselves accountable. These resources are available on our responsible technology principles page.

We know that our ways of working can be unique and our AI use cases will differ from yours. That's why our principles and practices are adaptable and scalable for your organization. These resources can contribute to compliance with technical standards and guidance from recognized authorities and expert bodies, even if they aren't a wholesale replacement for them. They also don't substitute for thoughtful regulation and enforcement, which we believe remains necessary to build trust and help foster the responsible development and use of technology.

That said, we've learned *a lot*. We hope the lessons we've taken from this journey so far will spark inspiration for teams everywhere – that's where this guide comes in.

We are profoundly grateful to all the experts, industry peers, and fellow Atlassians who have shaped these practices. In the spirit of our "Open company, no bullshit" value, we have and will continue to evolve them as the AI landscape changes. And that's ok. Unlike a Jira ticket, responsible technology never moves to the "done" column.

Let's dive in.

### How do I use this paper?

In this paper, we outline the key pillars of Atlassian's approach to responsible technology and how we bring them to life – including the lessons we've learned along the way, and how we'll keep improving. We want the pillars, initiatives and lessons in this paper to inspire you as you think through what responsible technology means for your organization.

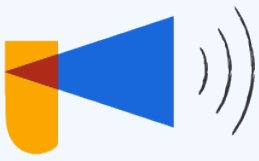

# Embedding responsible technology at Atlassian – and beyond

Our approach to responsible technology is built on three core pillars:

1. Upholding our principles through real-world practices.
2. Learning together, as a team.
3. Delivering transparency and driving dialogue.

These pillars are applicable to every organization – and every person in the organization.

In this guide, we walk through what these pillars look like at Atlassian, our key takeaways on how best to apply them, and what we've learned about each of them over the course of our responsible technology journey.

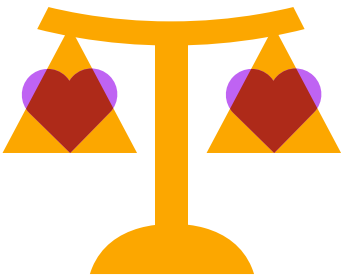

# Upholding our principles through real-world practices

Strong technology governance is more than adhering to the letter of relevant laws, frameworks, and standards. It's built by connecting values, to principles, to practice.

At Atlassian, we first established our **Responsible Technology Principles** as our North Star. These principles guide all of our work and help us to take accountability for using technologies like AI responsibly and in line with our company values.

Our Responsible Technology Principles were heavily informed by, and designed to align with, a number of similar principles embedded in frameworks globally. But they also feel uniquely Atlassian. We drew on our company values to articulate our perspective and our corresponding commitments to our customers, employees, and community stakeholders.

These Responsible Technology Principles are also accompanied by other foundational commitments that are critical to AI governance. These include existing policies that govern the development, use and deployment of technologies like AI across our organization, laws that we must abide by, and best practices and industry standards that provide directional guidance to organizations like ours.

Of course, commitments are nothing without action.

To connect our principles to our work, we focused on the practices we use to embed responsible technology into everyday operations. At Atlassian, our core practice is what we call a responsible technology review, with our Responsible Technology Review Template at its heart.

A responsible technology review is a structured practice that helps teams think through, record and explain the key decisions and considerations behind each AI use case. This is similar to a privacy or data protection impact assessment, and works in concert with it. Together with these other core processes like privacy and legal reviews, a responsible technology review can provide a holistic view of that use case. It's our way of making sure we

understand and document the potential impacts, risks, and benefits of using AI – so we can act responsibly and be transparent about our choices.

A responsible technology review follows a consistent format. This helps us ensure a consistent, baseline level of review across all use cases. The Responsible Technology Review Template breaks down our Responsible Technology Principles one by one, then asks teams to consider and respond to a series of open-ended questions that relate to different aspects of that principle.

This is why we're using responsible technology reviews across all our teams that develop or make decisions about technology. But while the review format is consistent, we've learned that achieving scale across all teams and use cases means figuring out the right approach for each of them. Over time, we've adapted the scope of our responsible technology reviews to focus on the most important use cases and critical questions to ask of our teams – and will keep doing this on a regular basis.

We designed this practice to be suitable for any team and to support consistent, efficient, and scalable practices. It's our hope that anyone can pick it up and run with it.

Importantly, our foundational commitments apply across all aspects of our organization and our teams' work. Not just in our customer-facing products, but across all of our internal and external activities that relate to building, deploying, and using AI.

# Our key takeaways for connecting principles to practices

Here's what we've learned about putting our foundational commitments into practice:

### 1 Your principles should be more than an expression of norms and standards.

Responsible technology should feel like an intrinsic part of your organization's goals. When identifying how to put your principles into practice, always come back to your organization's mission and values first. Think of your principles as a bridge between those values and the practices you implement to bring responsible technology to life.

### 2 Don't reinvent the wheel.

Your technology governance programs, like privacy and security, may already cover some aspects of AI development, deployment, and use. You can build on these programs by identifying gaps that would need to be filled by AI-specific governance programs – and then act only to fill those identified gaps.

### 3 Perfection is the enemy of progress.

When it comes to responsible technology, the journey *is* the destination: there's always more to be done. By shipping early and iterating often, we hit the ground running while also welcoming and integrating feedback along the way. In the past two years alone:

a. ***We designed for real-world use, not checklists.*** We continued to revise our Responsible Technology Review Template so that we could meet teams where they are. In asking questions of our teams, we want to encourage curiosity and openness and minimize check-the-box responses. If a question isn't resonating with our teams, we continue to experiment until it does.

b. ***We implemented more support to make our self-service tools successful.*** We released our Responsible Technology Review Template as a self-service tool. Teams are asked to think through how they feel about the alignment between their work and our principles, and take accountability for their answers. This is an inherently subjective task. To support our teams, we are continuing to invest in guidance on risk classifications, cross-functional partnerships, and governance structures to help shepherd teams through these processes.

c. ***We revisited key trade-offs as circumstances changed.*** We first made a conscious choice to approach responsible technology reviews as a practice, with a focus on learning. As a result, our initial efforts to embed responsible technology reviews cast a wide net across use cases, regardless of risk. Over time, we've recalibrated our approach to focus on higher risk use cases.

# Learning together, as a team

Building a strong responsible technology culture is critical to AI governance.

We want to ensure that all teams have the skills and awareness to understand and responsibly embrace AI. Achieving this aim will look different in every organization: depending on the knowledge, experience, and training within and across teams, and the ways in which they will use AI systems.

Concepts like AI literacy are a core part of this pillar. AI literacy refers to the idea that the humans developing, deploying, and using AI systems should be informed about what that means, and aware of the opportunities and limitations of AI – as well as the harm it could cause.

As the Human Technology Institute explained in their report on the **State of AI Governance in Australia**, harms from AI can arise from a number of scenarios. This includes when AI fails to operate as intended, is used maliciously or misleadingly, or when it is overused or used in circumstances where it should not have been relied upon.

But taking a truly human-centric approach means keeping the best case in mind too. Understanding the capabilities and opportunities of AI systems is what allows us to maximize the benefit and value that we all derive from AI.

Because of this, we take an expansive view of AI literacy.

At Atlassian, we have implemented many different initiatives to promote AI literacy. We have developed traditional training programs designed to provide all Atlassians with a basic understanding of responsible technology. This includes a dedicated responsible technology training program as well as covering responsible AI concepts in privacy and security trainings.

Our literacy and awareness raising initiatives are not always this formal. We rely on a wide variety of initiatives and mechanisms: from providing one-to-one expertise and advice to teams working with AI, to providing channels for sharing information and learnings, to supporting more formal, role-based training programs across the organization.

We also lean into existing cultural practices – or “plays” as we call them. Atlassian teams use dozens of plays as templates for addressing common challenges, starting important conversations, or just getting sh!t done day-to-day. Many of these plays can help build the feedback and learning loops that are critical to keep building AI literacy over time.

### Interested in learning more about plays?

We open-sourced the Atlassian Team Playbook! Check it out here or try some of these plays for your teams:

- AI Training Workshop - Atlassian Plays
- AI Teammate
- AI Innovation Day
- How to Run Pre-Mortem Exercises
- Ultimate Guide to Empathy Mapping
- Complete Guide to the 5 Whys Exercise
- Disruptive Brainstorming

# Our key takeaways on building a responsible technology culture through learning

Here's what we've learned through our literacy initiatives:

### 1 Build a literacy program that meets your organization's unique needs.

Building understanding and awareness of AI will look different for every organization. Staff building AI products will have different needs than staff in teams or at organizations that are less exposed to AI in their day-to-day work. To get started:

a. ***Know your teams.*** Identify where your organization sits, and what your staff need to know to use AI responsibly in their work. For example, our AI teams have benefited from structured programs covering highly technical information, whereas less technical teams have often preferred to start from the basics.

b. ***Take what works, and leave what doesn't.*** Take some time as early as possible to think about how you will identify and curate the resources that best serve your organization's goals. Explore any existing programs, initiatives and information-sharing channels that you can use, as well as guidance and resources made available by other organizations and technology providers.

### 2 Experiment first, and iterate over time.

We knew it would be difficult to design a short responsible technology training program that would strike a chord with each of our 13,000+ Atlassians. So we released a scrappy MVP version first, using Loom videos and polls. We sent this version to as many teams as possible, gathered feedback and engagement metrics, and iterated before releasing our full version.

### 3 Sometimes awareness is the first step.

When you have a lot of programs, initiatives, and internal expertise, discoverability can be your biggest challenge. It was important for us to dedicate time to making teams aware of our internal responsible technology resources – even before those resources were perfectly refined.

# Delivering transparency and driving dialogue

Open communication is at the heart of our approach to responsible technology. It's one of our values, after all – and it allows us to share what works and learn from what doesn't.

We want to equip our customers, stakeholders, and community with the right information to realize the benefits of new technologies like AI, and to keep learning and improving as we do so. We also see this as a shared task: accountability is a team sport, and we all have a role to play.

To achieve this, we focus not only on *what* and *why* we want to communicate, but *how.*

We ask ourselves three key questions:

- *How* do we communicate to our customers and users?
- *How* do we partner with others to understand the AI landscape we're operating within?
- *How* do we use feedback and partnership dialogues to learn and improve?

In each of these areas, we think about making sure that our communication and dialogue is purposeful and meaningful.

For example, we know that transparency alone isn't enough for meaningful communication. With new technologies like AI, just telling users that a feature is AI-powered won't mean much unless we are also able to equip the user with an understanding of what it means when the AI tool is in use, how best to interact with it, and how to respond to its outputs. Thinking through these elements also help us to ensure that our humans in the loop bring true oversight of, and effective collaboration with, AI.

Similarly, we know that we are often only one link in the supply chain that makes up the AI ecosystem. We need to own and take accountability for the technology we put out into the world. But we also need to work with our vendors, partners, customers, and peers to make sure that our advocacy and dialogue is informed by genuine feedback and a real-world understanding of AI adoption.

This approach helps us bring our Responsible Technology Principles to life outside of the four walls of our organization – and gives us the information we need to keep improving inside of them.

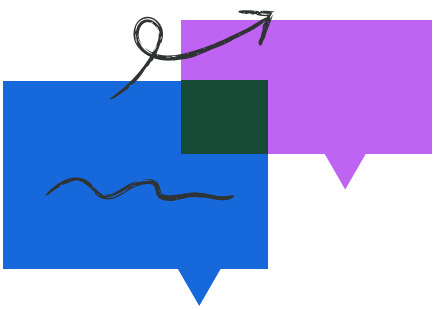

# Our key takeaways for establishing transparency and dialogue

Here's what we've learned as we've brought transparency and dialogue to life:

### 1 Go beyond legal requirements to meet your stakeholders' expectations.

Some AI laws are setting out specific requirements for providing transparency to end users about AI. Be thoughtful about not just what you *can* do within the bounds of these laws, but what you *should* do in light of what your users and stakeholders will expect.

### 2 Transparency is more than just disclosures and documentation.

Communication isn't limited to structured documents like AI transparency disclosures. Over time, we have made sure that our understanding of transparency and accountability is informed by more prominent design and UX features like our consistent use of logos, actionable footer disclosures, and easily accessible features like our thumbs-up/thumbs-down feedback collector.

### 3 Tailor your transparency efforts to your place in the AI ecosystem.

Every organization has a role to play in the AI landscape. Your organization might develop AI models or systems, integrate them into products, deploy ready-made systems to your end users, or include end users within your staff. In fact, you could hold any or all of these roles. Knowing where your responsibilities start and end can help you craft more meaningful communications.

## 4 Work with like-minded partners to drive dialogue.

Once you know where you sit in the ecosystem, you can more easily identify like-minded organizations and associations that you can share and learn from. These might include industry associations of peer organizations, smaller coalitions of companies, or even partnerships with governmental or academic institutions. At Atlassian:

a. ***Community partnerships.*** Our longstanding partners like the **Human Technology Institute** have provided us with incredible opportunities for two-way dialogue with the broader community.

b. ***Dialogue with regulators and policymakers***. More recently, our commitment to the **EU AI Pact** meant that we publicly undertook commitments towards compliance with the EU AI Act – and set ourselves the important task of working towards and learning critical lessons from meeting those commitments.

# Where to next?

We are committed to using our principles alongside the practices, plays, and processes that underlie them to guide our work, decision-making, and communications on the use of responsible technology.

We are also committed to continuing to learn, iterate, improve – then share. We have learned not to let perfection be the enemy of progress, and we also believe that no one company can solve this challenge alone. We will keep being open about our journey, and inviting your feedback.

Find your Responsible Technology Review template and other resources here.

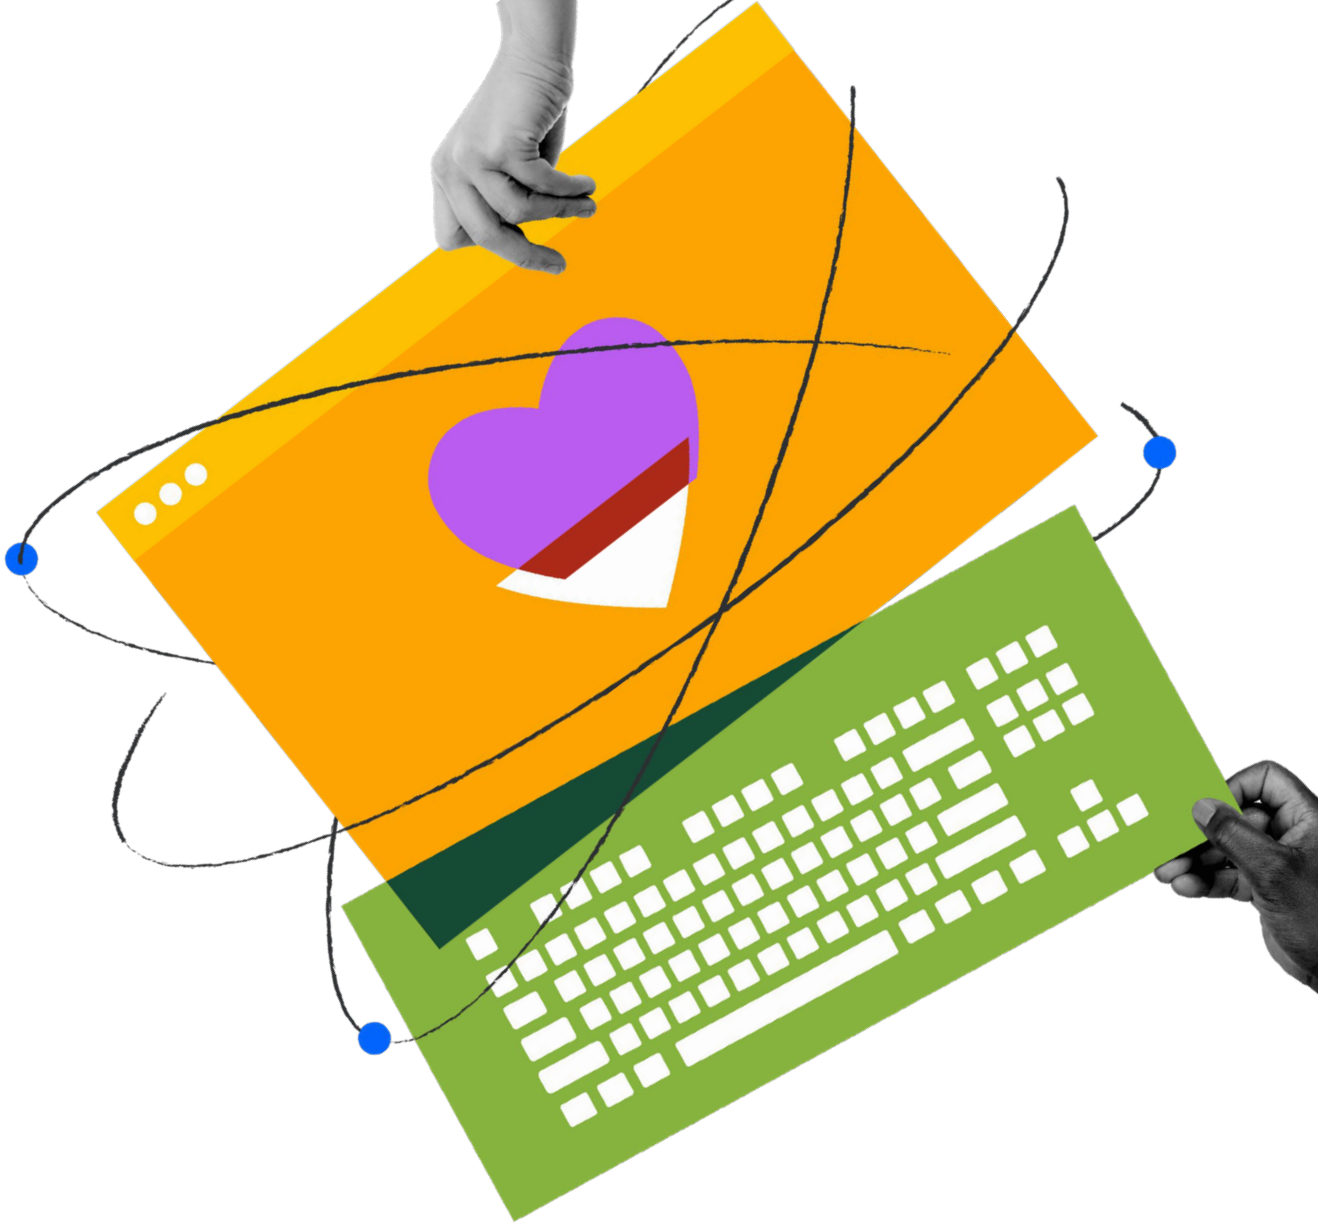

ATLASSIAN