

AI Governance
**Lighthouse
Case Study**

The Atlassian logo, consisting of a blue triangle icon followed by the word "ATLASSIAN" in blue capital letters, centered within a white circle.

October 2026

Key Governance Insights

AI governance is iterative, not static

Recognising the rapidly shifting AI landscape, Atlassian continuously reviews whether its governance structures, practices, and enablers are aligned with its strategy and remain fit-for-purpose as the technology, regulatory expectations, and risks evolve.

AI governance structures should fit organisational culture and existing processes

Atlassian's federated approach, with AI oversight managed by existing risk domains, aligns with how the company already operates. At Atlassian, AI governance is not one-size-fits-all. AI governance structures should be designed to match an organisation's culture, operating philosophy, and existing governance architecture.

Governance documents should reflect how staff actually use them


Atlassian has continuously refined its Responsible Technology Review Template in response to staff feedback and practical experience, making it easier to use and more accessible to a broader range of staff.

AI literacy is the foundation of responsible AI

Atlassian treats AI literacy as an ongoing critical capability and much more than just a one-off training exercise. Responsible AI lessons are embedded into broader capability-building around AI.

HTI's AI Governance Operating Model

This case study will explore Atlassian's approach against the four components of HTI's AI Governance Operating Model:





What is Atlassian's strategy for AI?

Atlassian does not have a separate AI strategy

While AI is core to Atlassian's strategy, it is not prioritising AI for its own sake. Instead, Atlassian is embedding AI meaningfully, responsibly, and naturally across its products and workflows to help achieve its mission to "unleash the potential of every team".

Reflecting this mission, AI is increasingly viewed as part of the team, enhancing how people work together, rather than simply automating tasks or delivering individual productivity gains. Human-AI collaboration is at the centre of Atlassian's approach to AI.



Our AI strategy is to respond to change fast vs. focusing too much on building the 'right' thing. This means horizontal investments that enable all AI capabilities in service of solving people problems, not technology problems. We say AI is one of the best things that's happened to us because it lets us do that faster and at a higher quality than ever before. So the question was never 'what's our AI play,' but if the work teams do every day actually gets better because of it.

Sherif Mansour, Head of AI & Product Craft



Atlassian's strategy continues to be shaped by its core values

Atlassian's values are actively applied in day-to-day decision-making and are regularly invoked by staff when working through difficult questions. The value 'be the change you seek' has shaped Atlassian's approach to responsible technology. Teams are encouraged to think through the consequences of their decisions. This has resulted in product teams proactively rethinking features from their own AI use cases where they concluded it was unnecessary or potentially harmful, without being directed to do so.

This values-driven approach has helped foster a culture of safe experimentation, while building AI capability, trust and confidence across the organisation. Atlassian has further reinforced this approach by treating itself as 'customer zero', which focuses the company's efforts on testing and using its own AI products within internal environments before public deployment.



What governance structures does Atlassian have for AI?

Atlassian integrates AI governance

Atlassian takes an integrated approach to AI governance instead of having a dedicated AI governance body. This approach to AI is reflected more broadly in its recent role expansion of the *Chief People and AI Enablement Officer*, bringing People and IT teams together, as well as an internal Forward Deployed Engineering team to help architect and customise AI products across the firm.

Today, AI oversight is integrated into its existing governance architecture, including privacy and security reviews, procurement workflows, product development lifecycles, and enterprise risk management. This is a deliberate choice that reflects the company's culture and broader organisational philosophy: new governance bodies are only implemented where essential.

Rather than creating a standalone AI committee, Atlassian has focused on adapting and strengthening existing structures

The result is a federated model, where every risk domain is responsible for managing AI-related risks within its area of business activity. These risks are reported to the board through standard quarterly audit and risk processes, rather than a separate AI reporting channel.

Consistent with this federated approach, product-facing AI use cases are supported through Atlassian's product counselling model, which complements existing legal review processes. Product counsels advise teams prior to launch, guiding them through key legal, privacy, and compliance considerations. This gives the product team a single point of contact for advice, reducing the need to engage multiple functions concurrently and independently.

While accountability is distributed, key governance practices remain consistent across the organisation. For example, Atlassian's Responsible Tech Review Template provides a common, structured format for assessing new AI use cases across functions.

But Atlassian acknowledges that its approach will change over time and according to its needs

As a company built on agile principles, Atlassian takes an iterative approach to governance. Given the pace of change across the technology, regulatory, and risk landscape, it continues to assess whether its current structures remain fit-for-purpose and responsive to its evolving governance needs.

Atlassian recognises that its approach may not be a template for others. It works for them because it reflects how the company already operates. Ultimately, AI governance is not one-size-fits-all and must be tailored to the organisation.



How have Atlassian's AI governance practices evolved?

Guardrails, not gatekeeping

While Atlassian's federated governance approach allows for domain-specific practices, the Responsible Technology Review Template is a key exception. The template must be completed by all business units prior to the development of AI use cases.

Atlassian developed the template early in its AI governance journey, drawing on industry best practice and tailoring it to the firm's context. Reflecting Atlassian's iterative approach to governance, the template has undergone several updates.



We translated the questions into language that we thought our teams would more easily work with and be able to pick up without us having to walk them through it.

Anna Jaffe, Head of Regulatory Affairs & Responsible Technology





How have Atlassian's AI governance practices evolved?

Four important lessons emerged from Atlassian's updates to the Responsible Technology Review Template:

1

Listen to staff and learn from how they answer the questions

Atlassian refined the template based on feedback from staff across different roles and business units. It also examined how staff were completing the template, identifying which questions were answered well and which were not, and refining the wording and structure accordingly.

For example, UX designers answered stakeholder impact questions using 'personas' and outlining the benefits and harms for each. This was such a helpful framing that the template was subsequently adjusted to explicitly ask staff to identify the different stakeholders and consider the potential benefits and harms for each.

2

Keep it simple and accessible

Atlassian refined its template to ensure the questions used clear, simple language. As the use of AI expanded across the organisation, responsibility for completing these assessments extended to a wider audience that may not have a deep familiarity with the underlying AI technology.

3

Direct questions to the right level

Questions about third-party LLM providers were often unanswerable by individual product teams because provider relationships were managed centrally. As staff started building individual agents, questions about the model being used were relevant for the teams working with the providers, rather than the individual product owners. The revised template made it clear that these questions could be redirected and were not necessary for the individual team to answer. This reduced staff friction without losing oversight.

4

Risk tiering is essential as AI use cases scale

Not all AI use cases require the same depth of review. Below a defined low-risk threshold, teams go through a simpler, fast-track process. Above that threshold, teams complete the full template themselves. Even with tiering in place, the volume of full reviews has continued to increase rapidly, indicating that while tiering is necessary, it is not a complete solution on its own.



How does Atlassian enable responsible AI use through AI literacy?

Atlassian sees AI literacy as the foundation of responsible AI

For Atlassian, AI literacy is not simply knowing how to use an AI feature. It is understanding how a system operates, where it performs well, where it doesn't, and the opportunities and risks this creates. Atlassian views the 'trust layer' as key to connecting AI literacy with responsible practice and translating AI knowledge and skills into practical, real-world capability.

Responsible AI is part of AI capability building

Atlassian embeds responsible AI into existing training and capability-building wherever possible. When the company ran AI Builders Week for its product group, a week set aside for staff to experiment hands-on with AI, the Responsible Technology team built a responsible AI element into the programme. The same principle applies across Atlassian's literacy programmes: responsible AI concepts are introduced within broader content, alongside optional pathways to further resources. Atlassian offers a short-form responsible AI module to all staff but has deliberately chosen not to make it mandatory, preferring genuine engagement over compulsory completion.



AI literacy at Atlassian isn't delegated — it's led from the top. Our leadership don't just sponsor AI Builders Week, they open it, participate in build sessions, and showcase work. When leadership visibly experiments, fails, and ships alongside their teams, it gives everyone permission to do the same. That's how we went from 'AI is interesting' to 77% of our Product Management org and 80% of our Design org attending a deeply technical, hands-on curriculum.

Carry Gray, AI Transformation, Office of the Product & AI org





Atlassian's five key takeaways from [AI Product Builders Week](#):



Immersion beats theory

Hands-on experimentation drives more meaningful learning than passive approaches.



Solving for real, lived problems creates momentum

Atlassian employees related their experiments to the work they did day-to-day, resulting in 115 new ways to use AI, which was then captured in an expanding playbook.



Experimentation and iteration are essential

Consistent with Atlassian's organisational philosophy and governance style, employees were encouraged to try, fail, and try again.



Cross-craft collaboration accelerates learning

Collaboration between disciplines and business areas improved collective understandings across the wider firm.



Upskilling across crafts means AI literacy for all

Through working together on practical scenarios, individuals gained greater confidence for how AI can be applied within their teams.



AI Governance Enablers



AI literacy is a continuous loop, not a one-off exercise

Atlassian's approach to AI literacy is a four-step loop: **learn**, **explore**, **play**, and **share**. Atlassian's AI School illustrates each stage of this approach in practice.

	What each step involves	How Atlassian applies it
Learn	Build your teams' AI foundations. Give teams a clear understanding of how AI systems work in practice: not just how to use them, but how they operate, what data they rely on, and where they perform better or worse. Tailor learning to different roles and treat it as continuous, rather than a one-off, and encourage teams to document lessons learned, including from failures.	AI School provides practical training for employees working on AI features, explaining how AI systems actually work. It blends self-paced modules, workshops, and guest lectures.
Explore	Curate your AI resources. Consolidate AI resources into an accessible hub. Draw on internal experience, external guidance, and vendor materials, adapting them to reflect your organisation's values, policies, and risk appetite. Encourage team members to contribute regularly as team practices, tools, and regulations evolve.	AI School has evolved from an engineering-focused program that relied on external resources into a set of tracks and modules that are developed and shared by in-house experts and customised to Atlassian's unique products and context. Each wave of the program iterates on previous waves, and inspires the future program and the broader organisation.
Play	Turn understanding into real skills. Give teams opportunities to experiment with AI on real work in a safe environment. Start with routine, low-risk tasks before progressing to more complex use cases. Encourage teams to test assumptions, identify limitations and edge cases, and openly share both successes and failures. Hands-on experimentation is one of the fastest ways to build capability and confidence.	Each wave of AI School builds on the last, and includes hands-on practical activities, workshops, and a final capstone hackathon that gives participants the opportunity to apply what they've learned.
Share	Build a culture of AI learning. Building a culture of continuous AI learning across your organisation depends on your ability to draw upon early adopters and champions. Enable them to share practical examples, lessons learned, and effective practices across the organisation. Integrate AI discussions (e.g. where has AI been useful for me?) into existing team meetings, workshops, and retrospectives so that individual experimentation becomes shared organisational capability.	Teams share what they've learned via internal channels and Confluence pages, so individual experience becomes shared knowledge.

For more on Atlassian's AI literacy guide, see this resource: www.atlassian.com/whitepapers/ai-literacy-guide

What's next for Atlassian?

Continuous improvement aligned with values

Atlassian's governance structures, practices, and enablers are continuously refined and improved in line with its values. Consistent with its values of *being the change you seek*, *build with heart and balance*, and *don't 'F' the customer*, Atlassian is focused on being both a trusted user of AI and a trusted provider of AI products.

Building AI capability

Atlassian is expanding its internal literacy, fluency, and awareness efforts. Scaling this capability is central to Atlassian's view that responsible AI starts with AI literacy.

Anticipation and alignment with future trends

Atlassian is also considering how future technological developments and governance demands will shape its approach. Like many organisations, it is navigating how AI agents will be used and governed, and how they will reshape the future of work, particularly in software development. Atlassian is positioning itself to engage with these issues responsibly as they arise and how they will reorient certain roles, practices, workflows, job families, and in software development.

“

We're asking teams to let AI into their most important work, and trust isn't something you assume. You have to design for it. The model we keep coming back to is human expert, agent execution: hand the repetitive work to an agent, keep the human as the reviewer and the tastemaker. Build it that way, and governance isn't a separate thing you bolt on afterward. It's already there in how we decide what to build, and what not to.

Sherif Mansour, Head of AI & Product Craft

”

 Human Technology
Institute

For more information on HTI's AI Corporate Governance Program,
or to join our AI Governance network, please contact:

Llewellyn Spink, AI Corporate Governance Lead
llewellyn.spink@uts.edu.au

Jack Goldsmith, AI Corporate Governance Specialist
Jack.goldsmith@uts.edu.au