


Leading with Context

Lessons from Atlassian's AI Journey



Executive summary

The context advantage

Our [2026 State of Teams](#) research found that 85% of knowledge workers use AI, but only 6% of executives see organization-wide ROI.

At Atlassian, we've spent the past few years focused on closing the gap between AI activity and impact. As "customer zero," we've used our own products and systems to fundamentally rewire how we plan, decide, and deliver with AI.

Among our most important shifts was changing how we work to give AI the rich context it needs to deliver business impact. Here are our biggest lessons so far.

1

AI needs to be a C-suite initiative, not an IT project

Empower business leaders to drive their own AI outcomes. To support them, establish an internal, forward-deployed engineering model that embeds technical experts directly within their teams.

2

Models are only as useful as the context you give them

Competitive advantage comes from connecting AI to your organization's unique knowledge, decisions, workflows, and teams.

3

Technology's impact is culture-dependent

AI tools give your organization access to intelligence, but culture dictates whether anything changes. Prioritize how your teams think about using AI just as much as the using itself.

4

Start with tasks and skills before evolving roles

Start with practical workflow changes, identify the new skills teams need to develop, and then adapt roles based on what actually works.

5

Evolve what you measure as you mature

You may need to start with adoption, but shift toward concrete business outcomes as you progress. Think: better quality, faster delivery, lower costs, or greater customer impact.

Table of contents

5 The AI-first organization

7 Where we are on our AI journey

8 The five questions every leader is asking, and the answers we've uncovered so far:

10 How do we drive accountability?

12 How do we unlock the ROI of AI?

15 How do we make AI part of our culture?

19 How will teams evolve?


23 How do we measure success?

26 Key takeaways

27 What's ahead



A note from our Chief People & AI Enablement Officer



Avani Prabhakar

At Atlassian, we’re building for human-AI collaboration. But we’re also living it.

Here's a dilemma we hear from leaders every week: [85% of knowledge workers](#) now use AI, but only 6% of executives can point to clear organization-wide ROI. Meanwhile, token usage keeps climbing, and so do costs. The gap between AI activity and AI impact is widening.

We've been there, too. As "customer zero," we've used our own products and systems to fundamentally rewire how we plan, decide, and deliver with AI—and we haven’t always gotten it right. This report shares the five questions we asked ourselves along the way and what we've learned so far.

Our biggest realization: models are not an organization’s key competitive advantage. Your context is.

That's why we are model-agnostic: the moat isn't which model you use, but the institutional memory of every project, goal, decision, and workflow that only your teams understand. The organizations that win won't just bolt on AI. They'll give it access to the full, ever-evolving picture of how their teams actually work.

Technology alone can’t do this. AI transformation at scale requires a cultural shift. Teams need the fluency and psychological safety to experiment. Leaders have to stop treating AI as a personal productivity hack and start embedding it into how teams coordinate and deliver together. When they make that leap, [our research shows](#) their teams are 7.3x more likely to say that AI improves speed without sacrificing quality.

We’re still figuring this out. AI transformation is a continuous evolution, not a one-and-done project. This report is an honest account of what has and hasn’t worked for us so far. We’re sharing it in the hope that it helps you on your own journey.

What is an AI-first organization?

The AI-first organization

AI-first doesn’t just mean “everyone uses AI.” We’ve learned that adoption without intention just results in more activity, not impact.

An AI-first organization is a new operating system for teams, where humans move to the critical frontiers (e.g., high-stakes decisions, creativity, and judgment) while autonomous agents execute at speed, all fueled by institutional context.

[Our research](#) shows that this kind of transformation requires evolving three core parts of how your organization operates.

Context

Ask yourself: Are we effectively capturing our institutional knowledge (e.g., decisions, tradeoffs, and customer interactions) and can agents access it?

What it looks like: Context is documented digitally and structured so agents, not just humans, can read, act on, and extend it.

Without context, AI gives you generic answers. But when it understands your roadmap, why key decisions were made, and how your teams work, AI can deliver actionable, strategic insights.

Our research: The payoff for doing this work is massive. Teams that see higher ROI on AI [are 7.3x more likely](#) to have optimized their data and knowledge foundations for AI.

Culture

Ask yourself: Are we continuously experimenting and evolving?

What it looks like: Teams celebrate effective AI usage, view working with AI as a collaborative exercise, openly share what has and hasn’t worked, and lean into uniquely human skills like creativity and judgment.

Our research: Teams that see higher ROI on AI are [up to 3.2x more likely](#) to say they are encouraged to experiment to find ways AI works for them.

Workflows

Ask yourself: What outcome are we trying to drive, and how can we reimagine our workflows so AI accelerates success?

What it looks like: Leaders and teams don’t just bolt AI onto existing processes. They step back to focus on their biggest problems, the underlying causes, and how AI can help solve them.

Our research: Teams that see higher ROI on AI are [up to 3.5x more likely](#) to have reinvented how they work through new tools and routines.

Where we are on our AI journey

At Atlassian, we believe transformation is an evolution rather than a box you tick once. Our data and teams tell us we’re pointed in the right direction.

“We’re continuously innovating, even when we’re taking a break. We’re working while we’re not. It’s pretty amazing.”

Mark, Sales

“Two non-technical teammates and I were able to build an onboarding agent by starting with simple instructions and playing around to get it right. After launch, 96% of new hires became weekly active AI users and 45% became superusers, outpacing the company-wide average.”

Corinne, People Operations

☰ Context

Connecting all our AI tools to the Teamwork Graph (p. 12) yielded better, cheaper, and faster results when put to the test.

44%

Agents grounded in Teamwork Graph (TWG) context deliver **44% more accurate answers**

48%

With TWG, agents also use **48% fewer tokens**, without sacrificing speed*

150+ billion

TWG tracks 150 billion+ objects and relationships across development work, documents, messages, and more

12 billion

TWG updates 12 billion times a day, keeping context current

** We tested AI models across 10 complex scenarios with and without Teamwork Graph context.*

👥 Culture

When we give Atlassians room to experiment with new tools, they run with it.

Over 99%

of Atlassians use Rovo in their day-to-day work without a mandate

~80%

of our Product Management and Design teams participated in our AI Builders Week

Model-agnostic

We’ve adopted a range of AI tools including Rovo, Claude Code, Codex, Cursor, Replit, and Gemini. By plugging each into TWG, we’re able to instantly make them smarter and more cost-effective

🔄 Workflows

We didn't just add AI to how we work; we’re rethinking how the work itself happens.

15k+

Over 15,000 agents actively worked alongside Atlassian teams in a single month

135%

At Atlassian, agent-powered tasks are up 135% month over month, with each agent now handling about 40 tasks a day

30.8%

faster PR cycle time, meaning code moves through review faster



The five questions every leader is asking

We structured this report around the five questions we kept asking ourselves—and hear most often from enterprise leaders.

Our search for answers transformed how we approach context, culture, and workflows, helping our teams make context-led AI transformation a daily practice.

- 1 How do we drive accountability?
- 2 How do we unlock the ROI of AI?
- 3 How do we make AI part of our culture?
- 4 How will teams evolve?
- 5 How do we measure success?

 Question 1

How do we drive accountability?

Organizations often over-invest in tools and under-invest in people: our research shows executives are [84% more likely](#) to invest in new technologies than in developing teams' AI skills.

Real transformation requires both. At Atlassian, we started by aligning our C-suite on the why, the how, and exactly what success looks like.

Here are the operating principles we embraced:

- 1 Make AI transformation a company-level OKR
- 2 Bring technology and people transformations together
- 3 Establish guidelines so AI use is responsible
- 4 Set the entire business up for AI success

1. Make AI transformation a company-level OKR

Define company-level goals: At the start of FY26, we made AI transformation a shared goal across our CTO, CISO, and CPO. This was split into three workstreams:

- 1. Growing systematic usage of AI tools (p. 13)
- 2. Building an AI-ready workforce and culture (p. 15)
- 3. Delivering AI hero use cases by function (p. 19)

Create department-level goals: We shifted AI enablement from Information Technology (IT) directly to department leaders. By embedding AI tools into real work and hosting structured enablement events, every function that set a goal to grow its AI superusers (see next column) more than doubled that rate in just six months.

Evolve success metrics: As our adoption matured, we refined how we measure success (p. 23). We moved from simply tracking weekly active users to identifying the top 10% of users company-wide to our current metric: craft-specific superusers. These are the top 10% of people within their specific function who are actively redesigning how work gets done.

What’s next: In FY27, departments are driving specific, high-impact goals including: increasing customer support satisfaction by leveraging AI to reduce time to resolution, freeing up 30% of recruiter capacity, and launching five new revenue-driving AI use cases in Sales.

Our research: Superusers change what’s possible

A superuser is an employee who ranks in the top 10% of highest AI usage within their specific function, ensuring the metric is tailored to the unique tools and workflows of their craft rather than a flat company-wide number.

Where regular employees use AI to search, sustained superusers* set goals and let AI execute. 89% of these superusers report that AI lets them take on work they previously couldn't do.

This translates to larger, proven productivity gains. Even after controlling for tenure and past performance, superusers close 4.4x more Jira issues and merge 19.5x more pull requests.

Most importantly, this impact scales. Our internal research shows that just one superuser boosts their entire team’s performance by 56%, likely because they uplevel their peers and build infrastructure that lets everyone work smarter.

**Sustained superusers are employees who exceeded their functional AI goals in at least half the weeks between January and June 2026.*

2. Bring technology and people transformations together

AI transformation stalls when technology and culture initiatives run on separate tracks. Only [37% of knowledge workers](#) think their organization effectively balances the people and tech sides of AI. The fix? Put one accountable owner over both tooling and people. If you can't merge functions, create a shared OKR and a joint operating cadence.

At Atlassian, we expanded our Chief People Officer role to Chief People and AI Enablement Officer, taking Avani Prabhakar from running a 700-person Human Resources (HR) team to leading 3,500 people across a newly integrated People and Transformation function spanning IT, Data Science, Customer Tech Support, and Customer Engineering.

Our logic: HR has an enterprise-wide view of how work actually gets done. As Prabhakar put it, “The biggest barrier is not access to models or tools; it’s mindset, trust, behaviors, and confidence.”

3. Establish guidelines so AI use is responsible

Governance should provide guardrails, not gatekeeping. We published our value-driven [AI principles](#) externally to ensure accountability and treat governance as an iterative process.

To help teams put our principles into practice, we created a [Responsible Technology Review Template](#). Example questions include, “Could use or misuse result in unintended consequences?” and “Will this tech interact with data in a way that is aligned with the expectations we set with customers and employees?”

4. Set the entire business up for AI success

AI transformation succeeds when every function has the technical support they need to rethink how work gets done. We created an internal Forward Deployed Engineering team that embeds within business functions to identify high-value workflows, build AI solutions alongside them, and develop lasting AI capabilities. See p. 21 for more.

Key takeaway

Context-led AI transformation is a cultural shift, not just an IT deployment.

To drive accountability and results, merge the tool side with the people side of transformation and align the C-suite on how work will change.

Question 2

How do we unlock the ROI of AI?

In 2026, your biggest differentiator is context: the institutional memory of every project, decision, and workflow that only your teams understand. Without this context, AI can only deliver generic outputs.

But [69% of knowledge workers](#) say their data and knowledge foundations aren't optimized for AI, meaning most organizations have agents operating on a fraction of what the business knows.

Here’s how we approach it:

- 1 Adopt a model-agnostic strategy
- 2 Realize the full value of your existing data
- 3 Embrace an "open by default" culture
- 4 Focus on forward-looking norms
- 5 Tie every project to a goal
- 6 Capture context by default
- 7 Protect against AI slop with guardrails

At Atlassian, we follow a simple formula:

Intelligence x Context = Acceleration

Intelligence is the engine, but context is the fuel. If your AI doesn't know what choice you made in 2024, it can't help you win in 2026.

You can’t drop AI into disconnected systems and expect it to effectively drive work forward. Building context starts with capturing institutional knowledge, integration, and making data findable and structured. We call this agent-ready data.

1. Adopt a model-agnostic strategy

Between market shifts and unpredictable political interference, relying on one model provider can be a vulnerability. Our bet is on flexibility: invest heavily in your own context layer with an open ecosystem.

At Atlassian, we use a wide range of AI tools (e.g., Claude Code, Codex, and Gemini). We connect each one to the Teamwork Graph, which allows them to understand how work and people are connected, take actions in the right systems, and do all of it within existing permissions.

⌘ The Teamwork Graph

The foundation of Atlassian's context architecture is [the Teamwork Graph](#), a unified data layer that maps the relationships between work, people, decisions, and knowledge across every Atlassian product and connected tool. It connects to 100+ popular SaaS applications out of the box, including Slack, Salesforce, GitHub, and Figma, and supports custom Forge connectors for proprietary and industry-specific systems.

The Teamwork Graph gives any AI model instant access to your company's unique context, making every model smarter without locking you into a single provider.

At Atlassian, we depend on the Teamwork Graph to bring together all our context to make our AI faster and more intelligent.

2. Realize the full value of your existing data

Within most enterprises, context is scattered across dozens of tools where AI can't see the connections. At Atlassian, we integrate every app with the Teamwork Graph so we can map relationships between projects, people, and knowledge. This lets us feed graph context directly into any AI tool, including coding agents like Claude Code and Codex.

Our research: More context makes AI a smarter teammate

In an analysis of 450+ Atlassian teams that participated in ShipIt, our biannual hackathon, those whose Rovo results had access to more apps were rated 2x more innovative. This was true for both technical and non-technical teams across a range of job tenures.

3. Embrace an "open by default" culture

By making work transparent and accessible across the organization where possible, you build richer context that models need to deliver relevant results. We've long made openness a core operating principle. In practice, that means we:

- Work in shared spaces, not private channels or personal drives
- Document decisions where others can find them: the why, not just the what
- Restrict only when necessary (e.g., confidential financials)

4. Focus on forward-looking norms

Resist the urge to clean up every legacy artifact. We focused on establishing clear norms for how knowledge should be captured going forward. Over time, these practices helped our graph compound. A few examples:

- **Link objects** (e.g., connect a Confluence page to other pages or Jira tickets) so AI can more accurately understand the relationships between work.
- **Add AI notetakers to every group meeting**, so decisions and discussions become a searchable record and part of the Teamwork Graph. Our Loom notetakers also publish notes and action items to relevant apps, combatting meeting sprawl and allowing Atlassians who do attend to be fully present.
- **Tag the status of pages** (i.e., we mark each page as “Rough draft” or “Ready for review” or “Verified”) so that AI can pull from the right work.


Our research: Clear statuses lead to better knowledge quality

In a study of 1,000 knowledge workers, we found that reviewers are 22% less likely to catch fundamental flaws in a proposal when AI turns it into a clean, polished-looking document. However, adding a simple “Early Draft” label to the proposal nearly erased this effect, essentially resetting willingness to critique.

5. Tie every project to a goal

AI can’t help you achieve your goals if it doesn’t know what they are. At Atlassian, all teams document their goals in a centralized knowledge base called the [Goals app](#). Team goals are rank ordered and include clearly outlined success metrics. All OKRs are linked so it’s easy to see how every team’s OKRs ladder up to our single, big company goal. Finally, we add weekly OKR updates and monthly scoring to the Goals app.

This approach lets AI reason about tradeoffs, prioritize effectively, and spot duplicated or orphaned work. Most importantly, it anchors AI to our business strategy, giving it guardrails instead of letting it guess priorities based on chat volume.



6. Capture context by default

“Write it down” has always been good advice but a chore to action. AI changes that. Capturing decisions and tradeoffs, creating tickets, and documenting meeting summaries can now happen by default, not by discipline. We do this through:

- ✓ **Whiteboard → work, automatically.** We brainstorm in a whiteboard tool, then have AI read the sticky notes, turn them into tasks, and assign them to people or agents.
- ✓ **Capturing decisions where the work lives.** Every project-level call is documented on the work item, page, or goal to which it relates.
- ✓ **Voice-first drafting.** Talking through the full context of an idea into a voice memo or AI prompt produces a richer first draft, and keeps a human in the loop on what gets shipped.

■ **Our research:** Voice-first drafting makes work easier

When we asked Atlassians to test this workflow, we found that [84% felt less resistance to starting](#) and 78% got to a usable draft faster.

7. Protect against AI slop with guardrails

Before AI touches anything, align on what “good” looks like. One way we do this is through AI Working Agreements, a structured team exercise where members identify where and how to use AI. The exercise takes about 60 minutes and is available as [a Play in the Atlassian Team Playbook](#).

■ **Our research:** Clear AI norms drive team success

After piloting AI Working Agreements with nearly 300 Atlassians, [82% said the practice](#) helped their team align on how to use AI to drive progress and 75% learned a new AI use case.

We encourage teams to run [AI retrospectives](#) at the end of projects to assess what worked and what didn't.

A cross-functional group of Atlassians also set explicit [guidelines for how to write with AI](#) and created an agent based on those guidelines. Because model capabilities are rapidly evolving, we committed to revisiting and evolving the guidelines each quarter.

Key takeaway

To unlock the full potential of AI, build agent-ready data by integrating disparate tools, establishing forward-looking knowledge documentation norms, and clearly linking every project to centralized business goals.

Question 3

How do we make AI part of our culture?

[85% of knowledge workers](#) use AI, but only 29% have changed how they work.

The answer to driving meaningful, context-led AI transformation? Don’t issue mandates. Instead, create a culture in which teams feel safe experimenting with AI, identify their wins, and scale their lessons.

Here are three principles to help you get started:

- 1 Embed AI where work already happens
- 2 Create a strong culture of experimentation, modeled by leadership
- 3 Host structured enablement events

1. Embed AI where work already happens

Vet and enable the AI features that exist within your approved tools, then drive awareness so your teams actually use them.

We found that integrating AI into existing workflows accelerates adoption. AI superuser rates spiked as soon as Sales added custom Rovo agents to their revenue workflows, Finance turned on Claude in Excel, and Design started prototyping with Replit.

SpotOn unlocked quick wins with Rovo

SpotOn turned on Rovo, which was embedded in Jira Service Management and Confluence, and then connected it to Slack, too. The results were immediate and measurable:

- ✓ **Faster product requirements documents:** Creation time dropped from 1-2 days to just a few hours.
- ✓ **Scaled support:** A tailored sales agent in Slack answered 95% of questions in-channel.
- ✓ **Higher resolution:** The overall percentage of issues resolved by AI jumped from 40% to 55%.

2. Create a strong culture of experimentation, modeled by leadership

Atlassian has a strong “no mandates” culture. A mandate might get people to log in, but it won’t change how they work.

Instead, we focus on creating a culture in which teams *want* to learn. And it works: we drove near company-wide adoption of [Dia](#), our AI-powered browser, in just six months, entirely through grassroots education and enablement.

To get people comfortable sharing their learnings, we encourage leaders to live demo how they use AI by screensharing in meetings.

Our research: Leader demos spark AI usage

[We find](#) that when a leader uses AI live (e.g., by sharing their screen during a meeting) their team is 4x more likely to make AI part of their daily work.

We also foster a culture where employees write internal blogs outlining new ways of working. For example, the Confluence engineering team shipped four production features with zero manually written code, accelerating their pace by 7.5-15x. Their internal blog post detailing the process spread organically and was read by over 1,500 Atlassians, highlighting a company-wide eagerness to learn from other teams' AI wins.

3. Host structured enablement events

We’ve also expanded and set up programs to give people dedicated time to play with AI. For example:

- **AI Builders Week:** Our Research & Development functions hold a quarterly [AI Builders Week](#) in which teams are asked to ship a tangible change to one of their workflows. Our most recent event led to a 16% increase in Rovo superusers, and 99% of survey respondents said it helped them apply AI in a practical way. We put together [a playbook for how you can run your own](#).
- **Shiplt:** In our series of quarterly [internal hackathons](#), employees from across the business are given 24 hours to turn any idea into reality. Fun fact: [Jira Service Management](#) was originally a Shiplt project! During our most recent Shiplt, Rovo usage among the roughly 2,000 participants was 11% higher than the rest of Atlassian.
- **AI Innovation Days:** [These days](#) explicitly give teams time to experiment with AI. The only success metric is to learn something new. Sessions often kick off with an icebreaker where participants share a successful AI use case, an untapped opportunity they want to explore, and a recent frustration.

Many teams are now successfully running and scaling their own strategic AI experiments.

For example, our Rovo Growth team ran over 40 product experiments in just six months. By using tools like Cursor for rapid prototyping, they shipped 80% of what they tested—more than double the industry average. Crucially, speed didn't compromise rigor; they only shipped features that drove statistically significant lifts in pre-set key metrics like Monthly Active Users.

Similarly, our Product Growth team is testing lightly supervised AI agents to run website experiments. These agents autonomously ideate, code, and analyze tests while humans step in at key checkpoints to steer the process. Early results show positive impacts on both experiment velocity and business metrics like website sign-up rates.

Our research: Structured enablement events build AI habits

March’s AI Builders Week lifted superuser rates in Product Management and Design, with gains holding through April, signaling habit formation over event hype.

Similarly, about 30% of our People team participated in April's Shiplt event, which drove a 7.4 percentage point increase in the function's AI superuser rate.



In progress



Agent

cx-modernization

Q3: How do we drive AI adoption that sticks?

16


Our research: Culture can erase AI stigma

[We've learned](#) a penalty exists when organizations fail to invest in a culture that celebrates AI experimentation. In a study of 1,000 knowledge workers, we found colleagues rated AI users as 10x lazier than their peers who didn't use AI, even when evaluating the same work. They were also 24% less likely to recommend them for visible projects.

But this stigma almost vanishes when companies actively celebrate AI wins and leaders model its use. We also found that when people frame AI as a tool to support the team or client—rather than just a personal productivity hack—they are rated 11 percentage points higher on effort.

Company culture determines whether disclosing AI use is rewarded or penalized

Percent of respondents who perceive colleagues that disclose AI use as lazy, based on company culture



How we think about change management

Our survey of 1,000 knowledge workers shows that skepticism is the most common negative reaction to introducing AI within teams. Here are our recommendations:

- ✓ **Address skepticism head-on.** [45% of workers](#) think leaders see AI as a "magic switch." Avoid hype; be clear about what's changing, why, and what it means for their work. At Atlassian, our executive team regularly gives updates on our transformation via Looms and Confluence pages, on which anyone can comment and ask questions.
- ✓ **Make wins visible.** [Only 13% of workers](#) say their organization celebrates AI usage; in our most recent survey, just 19% say they openly share how they use AI. Surfacing real demos, wins, and failures turns individual experimentation into shared growth opportunities.
- ✓ **Invest in enablement.** Give employees dedicated learning days and encourage managers to participate.

Jolin

Agent

How frontier teams work with AI

To understand what to scale, we often look to our frontier teams, or those furthest along in using AI to change how they work. Here's what they do differently:

- **The goal is speed-to-outcome, not “AI adoption.”** These teams frame AI as a tool to solve a real problem (e.g., beat a competitor to market), not as something to adopt for its own sake.
- **Treat clarity as the new superpower.** AI executes fast but fails on ambiguity. The critical skill has shifted from doing the work to defining exactly what you want upfront.
- **Let roles dissolve around goals.** When implementation is this fast, job descriptions blur and people just pitch in. This only works on teams with high trust, psychological safety, and zero internal competition.

- **Mind the 30% quality gap.** AI gets you most of the way there, but that final 30%—review, refinement, and human judgment—requires deliberate effort.
- **Run time-bound experiments on real work.** Most frontier teams start by identifying specific constraints and running a bounded experiment against real deadlines.

How we identify frontier teams:

Top-down: Leadership asks specific teams to experiment with new ways of working.

Bottom-up: Teams experiment on their own and share what they learn through internal blogs.

Sam

Key takeaway

Create space for teams to experiment together, have leaders model how they work with AI, and study and scale what frontier teams do differently.

?

 Question 4

How will teams evolve?

Knowledge workers are [90% more likely](#) to worry about lacking AI skills than losing their jobs.

To address this anxiety, start small. Our **Talent Evolution Framework** does exactly this: we first examine how specific tasks are changing, then identify the new skills required, and only then rethink broader roles.

Here are the three steps to our Talent Evolution Framework:

- 1

 Start with tasks
- 2

 Skills come next
- 3

 Roles are where we go last

1. Start with tasks

Grounding AI discussions in tasks instead of titles makes change feel tangible rather than existential. At Atlassian, we started by having functional leaders map high-value use cases step-by-step and ask their teams two questions about the tasks within each workflow:

1. Where can AI take the manual work (e.g., high-volume tasks)?
2. Where must humans provide intent and taste (e.g., high-stakes tasks)?

Using this framework, we rapidly scoped 14 use cases, shipped 12 minimum viable products in 11 weeks, and successfully scaled 10 projects based on proven business value.

Our initial use cases included:

- **IT Operations:** Incident response teams across IT and Dev [leverage JSM AI alert grouping](#) to consolidate thousands of alerts from different apps and infrastructure layers, helping detect and prevent disruptions to service. In the most recent month measured, alert grouping saved 839 engineering hours—equivalent to 8 full-time engineers who can now focus on solving more complex problems—with a 59% reduction in time spent on alert details.

• **Customer Engineering:** Our internal engineering team built AI-assisted tooling so our support engineers can resolve complex customer issues faster. We cut resolution time from days to minutes, and 38% of customer support issues are now resolved entirely by AI without any human involvement. That's freed up our support engineers to focus on the complex, high-judgment work, and customer satisfaction has increased as a result.

• **HR:** We implemented an AI agent that auto-generates personalized onboarding plans, pulls resources, and completes admin tasks. It drove a 43% reduction in onboarding plan creation cycle time and a 82% customer satisfaction score from the pilot cohort.
- Q4: How will the shape of teams evolve?
- 19

- **Finance and Legal:** We built an Intelligent Document Processing (IDP) platform that autonomously processes invoices and contracts. In FY26, IDP processed 23,000 tax invoices, eliminating 15 manual workflows and accelerating quarterly tax claims by 2 weeks. We also streamlined our previously manual sales order processing time by ~87% (from hours to minutes). At full rollout, we forecast 19,000 hours freed annually.
- **Sales:** We created a Call Preparation agent that pulls context from 4-5 data sources into a concise brief before each call. The initiative drove a 151% increase in call conversions within the first month of the pilot.

A closer look at what didn't work

Of the 14 AI use cases we scoped, we ultimately paused two.

- **Finance and Legal:** We scoped a compliance evidence agent, but it required broad access that violated our security and privacy model. Building a compliant version required engineering resources beyond what we had.
- **Sales:** We built an agent to generate first-call pitch decks, but the tool lost relevance as our selling motion changed. We pivoted the work to persona-specific POV document generation, where the real value sits today.

2. Skills come next

We leverage internal work signals and new job postings across industries to track rising "Sunrise" skills (e.g., prompt design, eval craft, and narrative judgment) and declining "Sunset" skills (e.g., manual data wrangling and first-draft production work).

Through this analysis, we discovered that the most critical, scarce skill is "AI acumen," or knowing exactly when to apply AI and how to rigorously judge its output.

We've also identified two distinct, emerging skill profiles: one for teams building on the frontier and one for AI-first managers.

Frontier team skills:

- **Context creation:** The ability to write precise specs, prompts, and requirements, since AI executes fast but fails on ambiguity.
- **Review, validation & judgment:** Evaluating AI-generated output for correctness, quality, and fit.
- **Cross-functional fluency:** Operating outside your traditional role.
- **Systems thinking:** Building tools, processes, and frameworks that multiply impact across teams.
- **AI orchestration:** Moving beyond occasional prompting to embedding AI as a full operating system – orchestrating agents, building skills, and designing human-AI workflows end to end.

Manager skills:

- **Delegation & direction:** Routing work fluidly between people and AI, breaking big goals into assignable pieces, and coaching models like teammates.
- **Strategic prioritization:** Knowing what matters and pointing the team at outcomes, not output.
- **Quality arbitration:** The taste to pressure-test AI output, catch errors, and differentiate "good enough" from great.
- **Workflow design:** Rebuilding processes around AI rather than bolting it onto how things already work.
- **Talent development:** Deliberately building growth paths for junior talent as AI increasingly handles entry-level tasks.

3. Roles are where we go last

We believe AI will create more jobs than it eliminates – but those jobs will demand very different skills. Our approach to role evolution follows a deliberate process:

- Pick a job family
- Study it in an AI-first environment
- Define the new skills + new job

By applying this method, we are reshaping our teams, starting with the roles below:

Evolving role: Software Engineer

As coding becomes increasingly automated, engineers are spending more time deciding what to build, directing AI, and evaluating its output.

As this role evolves, we are restructuring the entire development lifecycle. We now track "AI Code Quality" to ensure faster output doesn't overwhelm reviewers. Downstream, we use agentic gating and Rovo agents to automate bug triage and catch vulnerabilities, reclaiming the 28% of engineering time currently lost to operational toil.

New role: Director of Capacity Planning, Human & AI, Strategic Modeling

We [believe that HR needs to move](#) from managing headcount to architecting capacity, or the total mix of human talent and agentic capability. This requires a new role, one that owns workforce modeling and treats AI capacity as a real, plannable resource alongside headcount.

Expanded role: Design Technologist

AI enables this role, which sits at the intersection of Design, Engineering, and Product, to turn ideas into prototypes much faster than before. We’re investing more in this role; we’ve seen 35% month over month growth in AI-assisted prototyping output, meaning teams can test a wider ranges of ideas to make decisions based on data instead of instinct.

Hiring new graduates

While much of the industry is pulling back on junior talent, we're leaning in. Our data shows early-career teammates are 19% more likely to be AI superusers internally and across the tech industry, they are [2.1x more likely](#) to use AI to enhance decision-making.

Flatter structures and smaller teams

Newer, AI-first teams run structurally different. Our Forward Deployed Engineering team operates with just 3 layers versus the traditional 5 in a comparable engineering org. We also collapsed Forward Deployed Product Managers and Forward Deployed Engineers into a single job archetype to eliminate "whose job is this?" ambiguity.

Based on our research into frontier teams, we believe 5–6 people is the optimal team size for greenfield projects. This team size is small enough for rapid alignment but large enough to take on meaningful scope.

Selection and assessment criteria

We've evolved our hiring assessments to evaluate AI acumen directly. For example, we run two coding assessments with technical talent: one that tests raw fundamentals and one that uses HackerRank's AI-Assisted Integrated Development Environment to evaluate AI competency. Our goal is to assess how someone uses AI to solve problems.

Our research: How to support new ways of working

Within frontier teams, rigid roles are dissolving. When outcomes are clear and execution is swift, people just pitch in. But this only works in psychologically safe environments.

To support this new way of working, leaders need to make three structural shifts:

- 1. Reward the team, not the hero.** Performance management should prioritize shared outcomes over individual wins.
- 2. Prioritize time zone overlap.** When project timelines shrink from quarters to weeks, teams need frequent syncs to stay aligned.
- 3. Invert the planning cadence.** Ditch the detailed, long-range roadmaps. Teams still need a clear North Star, but they should focus their planning energy on what happens over the next 1-2 weeks.

Why this sequencing matters

This approach grounds the conversation and builds trust. When we lead with everyday tasks, teams see practical changes they can actually try and benefit from. Following up with skills gives everyone a clear development path. That way, when it's finally time to rethink roles, we're doing it based on actual evidence. This framework allows us to be rigorous instead of reactive.

Key takeaway

Role redesign should be the last step, not the first.

When you start with tasks and skills, you build trust and you're more likely to end up with changes that stick.

Question 5

How do we measure success?

Many organizations are investing heavily in AI without knowing if it's driving mission-critical outcomes. Our research shows that [only 6% of executives](#) can point to clear org-wide ROI from their AI implementation.

As our AI implementation has matured, so has our approach to measuring its business impact.

Here's what we've learned:

- 1 Adapt how you measure as you mature
- 2 Measure the work, not just the user
- 3 Invest in experimentation

1. Adapt how you measure as you mature

Success happens in stages: first, focus on getting people to use AI, then aim for better workflows and clear ROI. Update your metrics accordingly. [One of our key measurement frameworks](#) evolved in three stages:

Phase 1: Measuring Weekly Active Users (WAU) to understand if our employees were using AI. When AI first landed in the enterprise, the only question anyone could answer was the simplest one: did this employee use an AI tool this week?

WAU was a useful starting point. It told us whether our employees were trying the tools. It did not tell us whether anything was changing. So we instead started asking, “Are people using AI heavily enough, and is it improving how they work?” The shift from breadth to depth became the foundation of our AI superuser metric.

Phase 2: Developing a company-wide superuser metric. We defined a superuser as anyone using AI tools more than 40 times a week (the 90th percentile of company-wide usage from the prior half) and aimed to double the percentage of superusers within six months.

We blew through that target but realized that a single bar penalized non-engineers (whose work naturally generates fewer AI interactions) and rewarded volume over actual transformation, so we changed the definition again.

Phase 3: Evolving the superuser metric to be function-specific. Today, each craft has its own 90th-percentile baseline of weekly AI interactions, set from January 2026 usage data. The bar moves with the work. A superuser engineer looks completely different than one in legal: different tools, different cadence, different signature. What they share is not a number. It’s a position: the top ~10% of how AI is actually being used in their function.

To drive context-led AI transformation, we believe you have to identify the small group of employees who are redesigning how work gets done in their function. That’s what our current superuser metric helps us do (though we’re already reevaluating it for the new fiscal year).

These employees *are* your AI strategy. Most companies just don’t know how to find them yet.

Explainer: Why token usage measures the wrong things

Assessing AI success by token usage rewards activity rather than actual capability or return on investment. A high token count doesn't mean your team is building AI fluency or improving how they work. Instead of driving real ROI, chasing tokens just leads to tool sprawl, encouraging teams to use more AI agents and tools rather than the right ones.

2. Measure the work, not just the user

Measuring AI value is harder than measuring AI usage.

We’re developing a measurement framework that will have roughly four stages. We plan to start with adoption and usage, then look at the task level, layer up to workflow outcomes, and ultimately assess business impact. The questions we want to be able to answer are:

- Are people actually using AI?
- Have actions been taken to finish tasks?
- Is each unit of work being done better or faster?
- Is AI delivering ROI at the company level?

The core design principle is to measure the work, not just the user. In a world where agents run overnight and automations close tickets without anyone clicking a button, user-centric metrics miss most of the value. What matters is whether the work is better, faster, and easier.

We also warn against relying on any single metric, e.g., adoption rates or lines of code, because they can easily be gamed and rarely reflect real transformation.

Instead, measure across multiple dimensions, including quality, simultaneously and pair these metrics with cost guardrails.

Within our developer teams, we use the DX AI Measurement Framework to assess whether AI tools are delivering value, not just being used. This framework tracks metrics across three dimensions: utilization, impact, and cost (see table below).

The DX AI Measurement Framework gives us a precise view of AI's impact on our developer teams. Recently, as AI writes more code, the "Customer focus" and "Production debugging" drives in our Developer Experience Index improved, but scores for "Ease of release" and "Build and test" dropped due to increased reviewer load. Because we can pinpoint exactly what's getting harder or becoming a bottleneck, we can proactively design solutions to unblock our teams.

DX AI Measurement Framework		
Utilization How much are developers adopting and utilizing AI tools?	Impact How is AI impacting engineering productivity?	Cost Is our AI spend and return on investment optimal?
• AI tool usage (DAUs/WAUs) • Percentage of AI-assisted PRs • Percentage of AI-generated committed code • Tasks assigned to agents	• AI-driven time savings (dev hours/week) • Developer satisfaction • DX Core 4 metrics, including: • PR throughput • Perceived rate of delivery • Developer Experience Index (DX) • Code maintainability • Change confidence • Change fail percentage • Human equivalent hours (HEH) of work completed by agents	• AI spend (both total and per developer) • Net time gain per developer (time savings - AI spend) • Agent hourly rate (HEH/AI spend)

3. Invest in experimentation

We've long invested in the [Teamwork Lab](#), a dedicated group of scientists that run safe, rapid experiments to figure out what to change in our business and measure whether those changes actually work. The Lab launches structured pilots and translates the best grassroots discoveries from our employees into scalable processes for the whole organization. Much of the research in this report comes out of the Lab. (Wharton professor Ethan Mollick has also written about [the importance of having a Lab](#).)

How the Teamwork Lab's research evolves our ways of working

The Teamwork Lab doesn't just study how we work with AI; it changes it. For example:

- By [studying our AI superusers](#), the Lab identified the habits and mindsets that turn AI from a novelty into real leverage, informing how we track adoption and develop AI talent.
- After a pilot showed [AI Working Agreements](#) helped 82% of participants align on how to use AI, we scaled the practice across Atlassian.
- Over the last few years, the Lab has taught teams to work in ways that feed context back into the Teamwork Graph, making every tool we connect more useful to the next person who needs it (p. 12).
- Its [creativity research](#) found that when AI creates a first draft, individual work looks sharper but the group's ideas narrow around the model's defaults—reshaping our guidance to think first, then use AI.

Key takeaway

Evolve your metrics beyond basic adoption to measure how the actual work, rather than the user, is improving.

By identifying internal superusers and investing in structured experimentation, leaders can successfully scale grassroots workflow transformations into measurable business impact.

Key takeaways

AI adoption without intention just results in more activity, not impact.

Here's how we've approached unlocking true organizational transformation, leading with context.

1 Drive accountability

AI is a cultural shift, not just an IT deployment. Merge the people and technology sides of transformation into a shared company-level OKR.

2 Prioritize making context accessible

Models are no longer your competitive advantage; your institutional memory is. Connect disparate tools to build "agent-ready data" and anchor your AI to centralized business goals.

3 Model experimentation

Erase AI stigmas by having leaders visibly use AI and encourage teams to treat it as a strategic planning partner.

4 Evolve roles slowly

Sequence change safely to maintain trust. Start by rethinking everyday tasks, identify emerging skills, and redesign roles last.

5 Measure the work

Basic user metrics miss the point. Measure how the actual work improves and empower function-specific superusers to lead the way.



What’s ahead

What we’ll be researching next

This is just the beginning. AI transformation is a constant evolution and we’ll continue to learn, experiment, and share.

Here are the big questions we're focused on answering next.


➔ How do roles and job families need to change across functions?

➔ How should we redesign workflows to take advantage of AI?

➔ How do we need to adapt planning cadences as AI accelerates execution?

➔ How does management need to change to suit new team structures?

To stay updated:



Subscribe to the [Inside Atlassian newsletter](#)
Get our latest research and insights



Follow Atlassian [on X](#)



Follow Atlassian [on LinkedIn](#)



Follow [Avani Prabhakar on X](#) and [on LinkedIn](#)
Chief People & AI Enablement Officer, Atlassian



Follow [Alicia Lenart on LinkedIn](#)
Vice President of Human Resources, Atlassian



Follow [Dr. Molly Sands on LinkedIn](#)
Head of the Teamwork Lab, Atlassian

About us

Atlassian unleashes the potential of every team.

A recognized leader in software development, work management, and enterprise service management software, Atlassian enables enterprises to connect their business and technology teams with an AI-powered system of work that unlocks productivity at scale. Atlassian’s collaboration software powers over 80% of the Fortune 500 and 350,000+ customers worldwide—including NASA, Rivian, Deutsche Bank, United Airlines, and Bosch—who rely on our solutions to drive work forward.

Learn more about our products, including Rovo, Jira, Confluence, and Jira Service Management, at atlassian.com.

Want to learn more?



How Domino’s Pizza Enterprises is cooking up faster, better delivery by unifying 130,000+ team members on Atlassian Cloud Enterprise

[Read more](#)



Delivery Hero achieves efficiency and savings by centralizing 40,000+ employees on one System of Work

[Read more](#)



How Sprout Social saved \$180,000 and hundreds of hours per year with Teamwork Collection plus Rovo

[Read more](#)