


The Atlassian Forward Deployed Engineering Playbook

How to turn enterprise AI from
pilots into production





Executive summary

Across 100+ enterprise customers, Atlassian's Forward Deployed Engineers (FDEs) have seen what gets AI into production and what gets in the way.

Four lessons from the field:

- 01 **Pick the right problem.** Focus on issues that are frequent, painful, and measurable, then work backward to figure out what AI needs to solve them.
- 02 **Treat governance as part of the architecture.** Build permissions, security, and auditability in from day one, increasing autonomy as trust grows.
- 03 **Design AI to become part of how work gets done.** Curiosity drives initial adoption. Value drives sustained usage. The most effective, widely used AI is embedded in the tools and workflows where people already collaborate.
- 04 **Optimize the system, not just the model.** At enterprise scale, economics matter. Reuse context, match compute to the task, and measure AI against the work it replaces.

How Atlassian FDEs help

FDEs combine software engineering and applied AI, embedding with customer teams to build AI use cases that drive sustained business value.

Our approach delivers measurable results:

\$20M+ saved

Deployed AI that autonomously resolves tier-1 and tier-2 issues, cutting mean time to resolution from 18 days to 10 minutes while maintaining customer satisfaction.

~50% less triage effort

Built a custom triage agent while navigating complex enterprise governance requirements, unblocking production AI rollout across 40,000 employees.

6.7k+ hours saved/year

Co-built two production Rovo agents, one that archived ~600,000 pages and one for work-item triage, estimated to save 6,700+ hours a year across 20+ teams.


THE GOAL

Move AI from experimentation to production, prove business value, and give customers a repeatable path to scale what works.



Table of contents

Introduction	4
Front-line learnings	9
01 Pick the right problem	10
02 Treat governance as part of the architecture	12
03 Design AI to become part of how work gets done	14
04 Optimize the system, not just the model	16
What we're still figuring out	19
Closing thoughts	21



Introduction



PUBLISHED OCTOBER 2026

Most companies already have AI. The challenge is engineering it into how the business actually works.

Atlassian's Forward Deployed Engineers (FDEs) partner with customers to take AI from a technology to a production-deployed solution.

Across 100+ enterprise customers and 80+ production AI agents, we've seen what gets AI into production and what gets in the way.

FDEs bring those lessons directly to your organization, helping you move faster from AI potential to production value. Here's what that looks like in practice:

 **A technology company saved \$20M+ in direct support costs.** FDEs deployed AI that autonomously triages and resolves tier-1 and tier-2 issues end-to-end, cutting mean time to resolution from 18 days to 10 minutes while maintaining customer satisfaction.

 **A semiconductor manufacturer cut engineering-services triage effort by approximately 50%.** FDEs built a custom triage agent while navigating complex enterprise governance requirements, unblocking the rollout of production AI across 40,000 employees.

 **A travel company is saving an estimated 6,700+ hours annually across more than 20 teams.** FDEs co-built two production Rovo agents: one that archived approximately 600,000 pages and another that automates work-item triage.

Every engagement looks different because every enterprise works differently. But the underlying approach is the same: AI creates value when it understands the context, systems, and workflows that make your organization unique and is engineered into the work people do every day.

That's what the Atlassian FDE program is designed to do.



The Atlassian FDE program

Atlassian FDEs are senior engineers, bringing AI engineering and applied AI together in a single role and building alongside your teams in your environment.

Atlassian FDEs at-a-glance

- **AI-enabled engineers who build with you:** FDEs deploy real-world AI solutions, design new agents, enrich context with your custom domain data, and build custom agentic workflows and MCP integrations.
- **Dual-domain experience:** FDEs bring lessons from a wealth of enterprise deployments to solve both data architecture and organizational bottlenecks.
- **Focused on self-sufficiency:** In ~12 weeks, FDEs leave your teams with clear business value and repeatable patterns to keep scaling wins on their own.
- **Empowered to improve the product:** Rather than applying temporary workarounds, FDEs remediate blockers directly within the Atlassian platform, hardening our products for enterprise use at scale.

FDEs start with a business outcome, then trace the work required to achieve it: who is involved, which systems hold the necessary knowledge, where handoffs break down, what permissions apply, and what an agent needs to know and do. Then they build. FDEs write production code, create agents, connect systems, and redesign workflows around AI.

Engagements typically last ~12 weeks and follow four phases:

PHASE	WHAT HAPPENS
Discover	Identify high-value business workflows and map system context and permissions
Build	Co-engineer, test, and deploy production-ready workflow and AI solutions grounded in your organization's context
Adopt	Drive sustained usage, establish operational governance, and equip internal engineering teams with repeatable patterns to continue building autonomously
Value	Measure business impact, demonstrate ROI, and optimize for the metrics that matter: cost, speed, revenue, and reach

The goal is to leave you with working AI in production and a repeatable approach your teams can carry forward.

A head start on enterprise AI

FDEs build on [Atlassian's Teamwork Graph](#): an enterprise-ready context layer connecting the people, knowledge, code, communications, assets, and work involved in deciding, shipping, and delivering results.

For many enterprises, much of that context already lives in Jira, Confluence, and Jira Service Management (JSM). But Atlassian is only one part of your technology stack. A customer relationship management (CRM) system understands your customers. An enterprise resource planning (ERP) system understands your transactions and assets.

Teamwork Graph adds a connected understanding of how work gets done, bringing in context from systems including GitHub, Slack, Google Workspace, Microsoft 365, and Figma while preserving the relationships and permissions that give that information meaning.

That means FDEs can spend less time assembling the context layer and more time identifying how AI can help solve important business problems.

88 The Teamwork Graph in action

44% higher-quality answers

Grounding RoVo in Teamwork Graph context produces **44% higher-quality answers** while using **48% fewer tokens**, with no reduction in speed.*

64% more work shipped

Our research also shows that teams whose AI draws most heavily on the Teamwork Graph **ship 64% more work**.

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



Built for production from day one

FDE solutions are built for your technology stack. FDEs use Atlassian's enterprise platform as a production-ready foundation, then connect it to the tools, data, models, and systems your teams already use. Rovo provides secure, scalable infrastructure, while Teamwork Graph brings together the context AI needs to work across your enterprise.

- ✓ **Permission-aware by design:** Rovo validates permissions at every step, so agents can't access or infer restricted information across connected systems.
- ✓ **Governed at enterprise scale:** Centralized controls manage agents, connectors, data sources, actions, and AI consumption as deployments grow.
- ✓ **Enterprise-grade security and compliance:** SOC 2 Type II, ISO 27001, CSA STAR Level 2, FedRAMP Moderate (Atlassian Government Cloud), GDPR and HIPAA support, with multi-region data residency.
- ✓ **Protected by default:** Customer data isn't used to train third-party foundation models, with user-level permissions and enterprise-grade encryption.
- ✓ **Connected to your existing stack:** Rovo, Forge, APIs, and 100+ connectors bring third-party context into Teamwork Graph and allow FDEs to securely connect proprietary systems and internal tools.
- ✓ **Built for efficient scale:** Shared context and centralized consumption controls help manage AI usage and economics as deployments grow.

THE RESULT

Because Atlassian FDEs start with an enterprise-ready foundation, they can immediately focus on making AI work for your business.



Front-line learnings

Across our FDE engagements, the same four reasons enterprise AI stalls show up again and again. Here is what we've learned from solving them in the field.



01

Pick the right problem

Challenge: “We have AI. Now what?”

For many enterprises, the biggest challenge is deciding where AI can make enough of a difference to change how people work. Hundreds of use cases are possible. Far fewer solve a problem that is frequent, painful, and measurable.

FDE APPROACH

Start with the business problem, not AI.

FDEs work backward from a business bottleneck to identify where AI can remove meaningful friction. They look for workflows that are:

- ✓ **Frequent:** The work happens often enough for small improvements to compound
- ✓ **Painful:** People genuinely want the work to become easier or disappear
- ✓ **Measurable:** Someone already owns an outcome that can prove whether the agent is working

What that first workflow looks like depends on the organization, though we’ve identified patterns across industries.

Industry	Common starting points
Software & technology	Developer productivity / DX: Jira + GitHub workflows such as sprint productivity, pull requests, and development workflows
Financial services	Governance-first AI: trusted, lower-risk experiences such as Rovo Search and Chat, then expand into agents and automated workflows
Healthcare	Data hygiene & compliance: identify and sanitize personally identifiable information (PII) so enterprise data can safely be used by AI, then build agents on top
IT / service management	Ticket triage: automate high-volume classification, routing, and resolution using JSM and organizational context



FIELD PROOF

What we've built with customers

-  **Telecommunications company: \$1M saved in custom development costs.** A large telecom needed to translate complex planning discussions into structured engineering work. FDEs built an agent directly into its planning process that turns requirements into Jira stories, sub-stories, and work hierarchies. Leadership shared that in addition to saving costs, the engagement significantly accelerated throughput.
-  **Semiconductor manufacturer: Automating engineering service triage.** FDEs targeted a high-volume operational problem: routing requests through an engineering services desk. The resulting triage agent reduced triage effort by roughly 50%.
-  **Telecommunications company: Improving service ticket triage.** A major telecommunications company was using an AI assistant to triage technical service tickets but experienced frequent incorrect routing. FDEs built a Rovo triage agent grounded in JSM and Teamwork Graph context. The customer found the results more accurate and replaced its previous solution for the workflow.
-  **Financial institution: From uncertainty to a working agent in less than three weeks.** FDEs identified a high-value workflow and quickly delivered a working agent in a secure sandbox, giving the organization something concrete to evaluate and scale.

Common AI solutions we're building with customers

SOLUTION	HOW IT WORKS
AI first responder	Automatically categorizes incoming requests, verifies initial diagnostics, and routes issues to the right team
Subject matter expert routing	Uses ticket context and ownership data to find the right expert for complex issues
Product lifecycle	Connects the path from idea → Product Requirements Document → Jira work → code, reducing manual handoffs

FIELD LESSON

Focus on solving for high-frequency, measurable friction over novel use cases.

02

Treat governance as part of the architecture

Challenge: Enterprise governance wasn't designed for nondeterministic software

Traditional software is relatively predictable: given a particular input, you usually know what the system will do.

AI introduces nondeterministic behavior, which leads to important questions: What data can an agent see? Which permissions does it inherit? What actions can it take?

We've seen technically successful pilots stall because those questions were addressed after the build instead of before it. We've also seen off-the-shelf agents fail when dropped into enterprise environments because they don't have the necessary context, integrations, and guardrails.

FDE APPROACH

Treat governance as part of the architecture.

FDEs build for security, permissions, and risk controls from day one, so solutions can safely and quickly move into production. Agents only access approved data and operate within clear guardrails that meet InfoSec and legal requirements. As trust grows, those guardrails can evolve to give agents greater autonomy.

READ → DRAFT → ACT

- 01 **Read:** Start with search, retrieval, and summarization
- 02 **Draft:** Let AI recommend an action, with a person approving it
- 03 **Act:** Automate well-understood, bounded actions once the organization is comfortable with the controls

Teamwork Graph provides a permission-aware foundation underneath those experiences, validating access as agents traverse connected organizational context.



FIELD PROOF

What we've built with customers

</> **Technology company: Building governance in from day one to save ~11.8k hours annually.**

FDEs worked with a finance team to deploy a commercial operations agent that drafts orders with human review and a full Jira audit trail built into the workflow. The approach cut order creation from hours to minutes and achieved 95% accuracy before human review, increasing automation without sacrificing oversight or accountability.

 **Bank: Governance-first AI for tens of thousands of employees.** A major bank enabled Rovo Chat and Search for approximately 35,000 users, starting with read-oriented AI while establishing the auditability, security evidence, and governance processes needed to expand into custom agents.

 **Healthcare company: Making sensitive data ready for AI.** A healthcare organization had PII in source systems that it did not want exposed to AI, and that its desired use cases didn't need. FDEs built a solution to identify and remove sensitive information before the data entered Teamwork Graph, then built agents on top of the sanitized data.

 **Technology company: 135,000+ issues resolved autonomously.** FDEs deployed AI that triages and resolves tier-1 and tier-2 service issues end-to-end, with guardrails determining when AI can act and when humans need to step in. The result: mean time to resolution fell from 18 days to 10 minutes.

Common AI solutions we're building with customers

SOLUTION	HOW IT WORKS
Content governor	Detects outdated Confluence content and proposes updates
Risk agent	Reviews epics and initiatives against risk standards
Sign-off agent	Coordinates approvals, records sign-offs, and nudges reviewers in Slack

FIELD LESSON

Design governance into agentic workflows from the start.

03

Design AI to become part of how work gets done

Challenge: Initial adoption doesn't translate into sustained usage

An agent launches, people try it, and usage rises. Then the novelty wears off.

We've seen that pattern enough to distinguish two very different signals: a successful launch tells you people are curious. Sustained usage tells you that you've built something valuable.

Our data shows that agents triggered automatically by real-world events are 12x more likely to become part of sustained workflows than agents that rely on users to prompt them manually.

FDE APPROACH

Turn AI from something people try into something they rely on.

FDEs design AI to become part of how work already happens:

- **Trigger it automatically:** Agents run in response to work events: when a ticket enters a queue, a pull request opens, a sprint ends, or an incident begins
- **Meet people where they work:** AI lives inside Jira, JSM, Confluence, Slack, Bitbucket, and other existing workflows rather than another AI destination
- **Measure what happens after launch:** FDEs track AI-enabled users, workflow outcomes, and customer-defined business metrics

The best AI experience may eventually be one people barely notice: AI routes the ticket, updates the documentation, and flags the dependency. The work simply happens with less effort.



FIELD PROOF

What we've built with customers

-  **Technology company: Cutting onboarding plan creation time by 40%.** FDEs partnered with a People team to build an onboarding agent that uses company, team, and role context to automatically draft tailored 30/60/90-day plans and support new hires from day one. The agent is estimated to save hiring managers 2,500 hours annually
-
-  **Travel company: From 120 experiments to one agent people actually used.** An enterprise had built 120+ AI agents, but adoption remained low. FDEs focused on a recurring workflow, embedding an automated ticket-triage agent directly into live work. Adoption far exceeded the company's previous agents.
-
-  **Motorsports organization: Embedding AI into daily operations.** FDEs helped an engineering-intensive organization go from zero AI use cases to six production agents in under five months. They automated reliability reporting and replaced constant alerts with AI-generated digests, helping engineers spot critical changes without alert fatigue.
-
-  **Travel company: Building data hygiene into the flow of work.** Missing data across ~1,200 Jira projects was blocking reliable executive reporting, while automated reminders failed to change behavior. FDEs built an AI Slack assistant that suggests missing start dates and acceptance criteria and writes approved updates to Jira in one click. The solution restored reliable roll-ups without asking developers to leave Slack.

Common AI solutions we're building with customers

SOLUTION	HOW IT WORKS
Daily scrum facilitator	Runs async daily standups and surfaces blockers
Weekly status reporter	Generates weekly status summaries from Jira activity
Similar incident resolution helper	Uses historical incidents to propose solutions

FIELD LESSON

Don't ask people to remember to use AI; put AI inside work they already do.

04

Optimize the system, not just the model

Challenge: Scaling AI can get expensive fast

An AI workflow that makes sense for 100 people may look very different at 50,000.

As usage grows, enterprises need to understand not only what AI costs, but what they're getting for that spend. We're still early in measuring the long-term economics of enterprise agents. What we can already optimize for are three major cost drivers: context, compute, and the cost of the work AI replaces.

FDE APPROACH

Optimize the system, not just the model.

FDEs focus on three levers:

- 01 **Reuse context instead of rebuilding it:** Teamwork Graph structures connected organizational context upstream, reducing the need for every agent to repeatedly reconstruct the same context
- 02 **Match compute to the task:** FDEs evaluate models based on the complexity, latency, accuracy, and economics of each workflow rather than defaulting to the most powerful model
- 03 **Measure against the status quo:** FDEs compare AI costs with the manual effort and operational friction the workflow eliminates

In Atlassian testing, grounding Rovo in Teamwork Graph context produced 44% higher-quality answers while consuming 48% fewer tokens, with no reduction in speed.

This becomes especially important in large enterprises, where a workflow may ultimately need to serve tens of thousands of people.



FIELD PROOF

What we've built with customers

 **E-commerce company: Cutting the cost of a daily AI brief.** A daily morning brief was consuming more than 1 million tokens per run by searching Slack, Jira, Gmail, and Calendar separately. FDEs connected Teamwork Graph to the company's external LLM workflows, significantly reducing token consumption and latency.

 **Telecommunications company: Getting more value from existing AI infrastructure.** When an existing AI assistant struggled to route service tickets accurately, FDEs built a triage agent grounded in JSM and Teamwork Graph context. The solution cut initial response times from days to minutes while delivering significantly higher accuracy.

 **Technology company: Making enterprise AI economics work at scale.** FDEs extended AI-powered search into third-party sources, giving teams access to more organizational context without adding another costly search layer. Leadership estimated the solution costs a fraction of competing enterprise search offerings.

Common AI solutions we're building with customers

SOLUTION	HOW IT WORKS
Executive insights agent	Creates high-level portfolio briefings from Teamwork Graph context
Chief Product Management Office reporting agent	Automates reporting while cutting costs by reusing structured context
Test summary report generator	Uses context to synthesize readiness summaries and anomaly flags

FIELD LESSON

Manage AI economics across context, compute, and the value of the work.



Key takeaways

These four challenges compound.

- 01 The wrong use case won't sustain adoption.
- 02 An ungoverned agent won't reach production.
- 03 AI that sits outside the workflow won't become habit.
- 04 And scaling usage without understanding the economics can increase cost without increasing value.

KEY LESSON FROM THE FIELD

Enterprise AI often stalls because everything around it isn't ready to put AI capability to work.



What we're still figuring out


Across our FDE engagements, we're seeing clear patterns in what works. The questions now are about how to make those lessons compound. Here's what we're focusing on next.

01 How much of today's custom work can become tomorrow's starting point?

Customers want to know what similar companies are building. We're turning patterns across engagements into reusable agents, connectors, and use cases so each customer doesn't start from scratch, without losing the context that makes each solution valuable.

02 How can we continue to help customers scale what works without us?


The goal isn't long-term dependency. It's to leave customers able to identify new use cases, build agents, and scale successful patterns on their own.



03 What are the next bottlenecks as building gets dramatically faster?

As AI accelerates implementation, we're seeing constraints move elsewhere: planning, review, cross-team dependencies, governance, and operations. Each engagement helps us better understand where new bottlenecks emerge and how the development system needs to evolve around them.

The goal behind all three: make every engagement compound into a better platform, a better starting point for the next customer, and greater AI capability inside each organization.



Closing thoughts



Helping you get more from AI

Enterprise AI rarely stalls because AI can't do the work. The challenge is everything around it: choosing the right use cases, connecting AI to the right context, building in governance, driving sustained adoption, and scaling in a cost-effective way.

Those are the challenges Atlassian FDEs solve. They bring lessons from across enterprise deployments, build alongside your teams, and start with the advantage of an enterprise-ready platform, so they can spend less time assembling the foundation and more time making AI work for your business.

Every FDE's goal is to help you get more from AI, give your teams the patterns to keep building on your own, and make every successful deployment a starting point for the next.

→ If you're ready to move AI from experimentation to production, please reach out at fde-contact@atlassian.com