appfire | accenture

AI-ready, set, scale

A pragmatic pathway to successful enterprise AI adoption

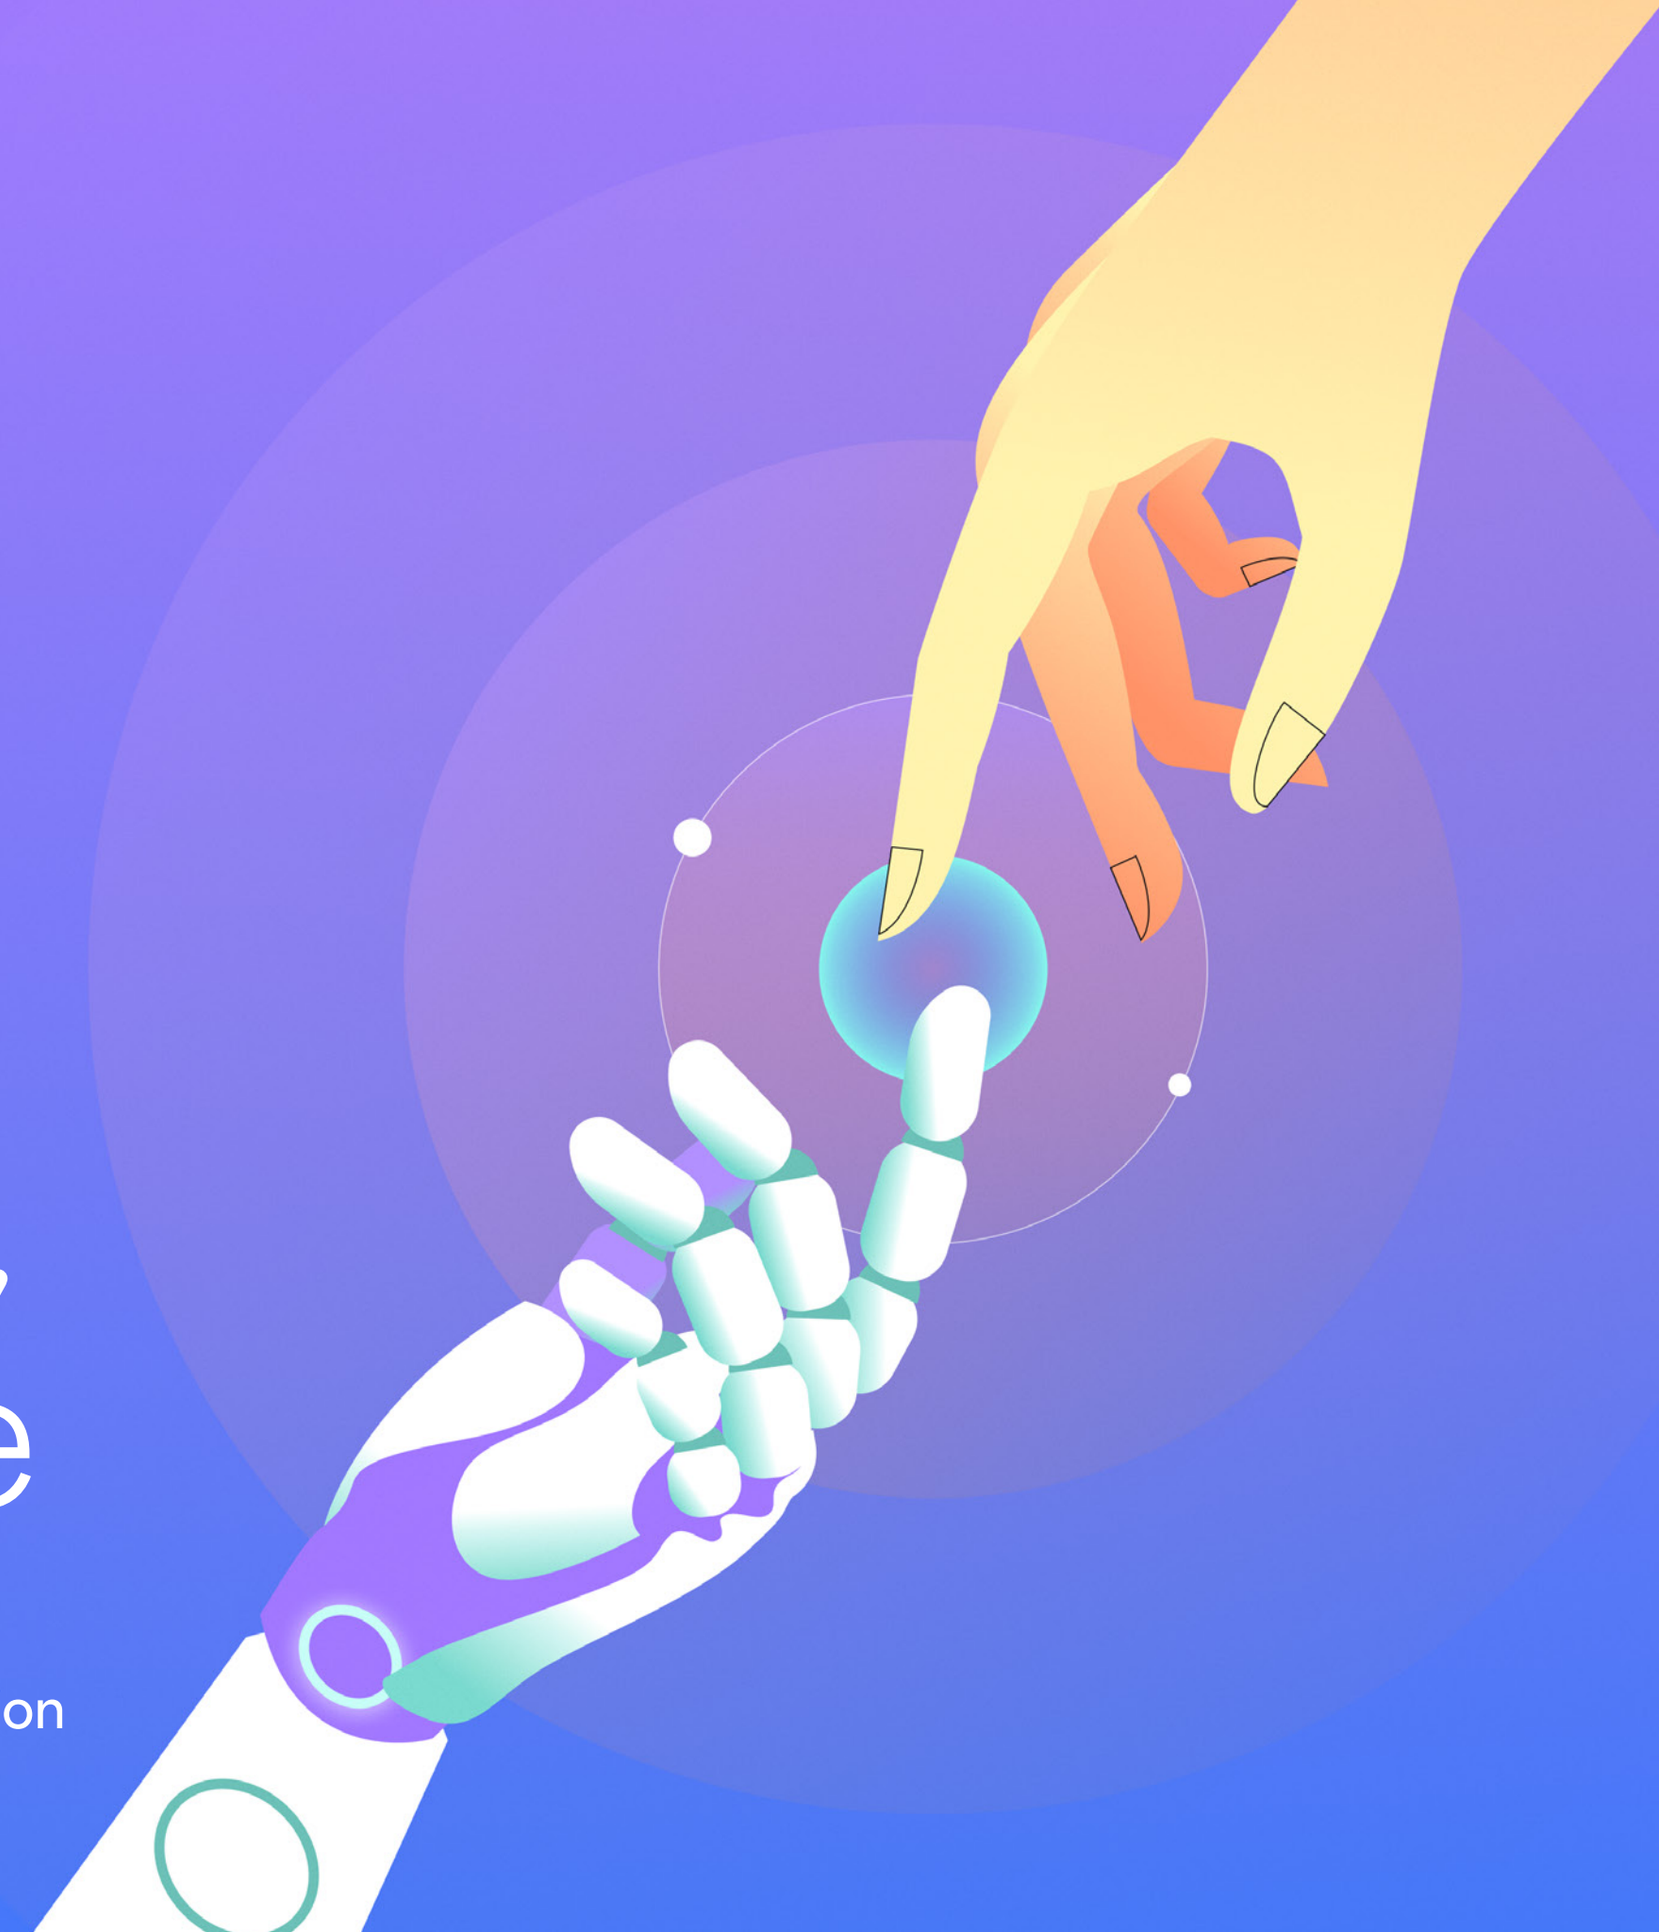

Contents

Executive summary

AI is reshaping how your enterprise creates, finds, and applies knowledge. But tool sprawl, increased information security risks, siloed data, and uneven change management often frustrate implementations and undermine outcomes.

Most organizations are not struggling to experiment with AI. They are struggling to make it accurate and secure at scale.

The gap is not in the technology. It is in how knowledge, systems, and workflows are structured to support it.

Accenture and Appfire provide a pragmatic, ROI-focused pathway to AI adoption that blends Atlassian's knowledge management lessons with Appfire's delivery and measurement capabilities. The result is a human-centered, governed, value-tracked rollout that connects intent to impact.

This path enables CIOs, CDOs, and IT leaders to move beyond pilot projects to scalable, governed, and measurable outcomes. The focus is simple: connecting the full power of your Atlassian investment to an execution and observability layer that ensures knowledge is usable, trusted, and tied to business results.

> AI is less about replacing work and more about removing the friction around thinking, creating, and deciding. The gap now is between organizations that experiment with it and those that actually redesign how they work because of it."
>
> Gabriela Cescatto, Transformation Lead at Husqvarna Group

The critical flaw: knowledge isn't AI-ready

A recent report published by MIT suggests that despite $30-40 billion in investment, 95% of enterprises are getting zero or even negative returns. Most GenAI pilots fail not because the models are flawed, but because they are poorly integrated into the workflows and data that actually run the business.

Organizations layer AI onto fragmented tools, stale and inconsistent data, and workflows that were never designed to support it. These efforts lack the feedback loops, change management, and process redesign needed to make them reliable in practice. The result is "AI theater" instead of measurable impact.

This exposes a critical flaw in the enterprise's underlying cognitive architecture - the very core of how an organization structures, governs, and connects its knowledge, workflows, and decision-making pathways. When this architecture is fragmented or inconsistent, AI systems amplify those weaknesses rather than resolve them.

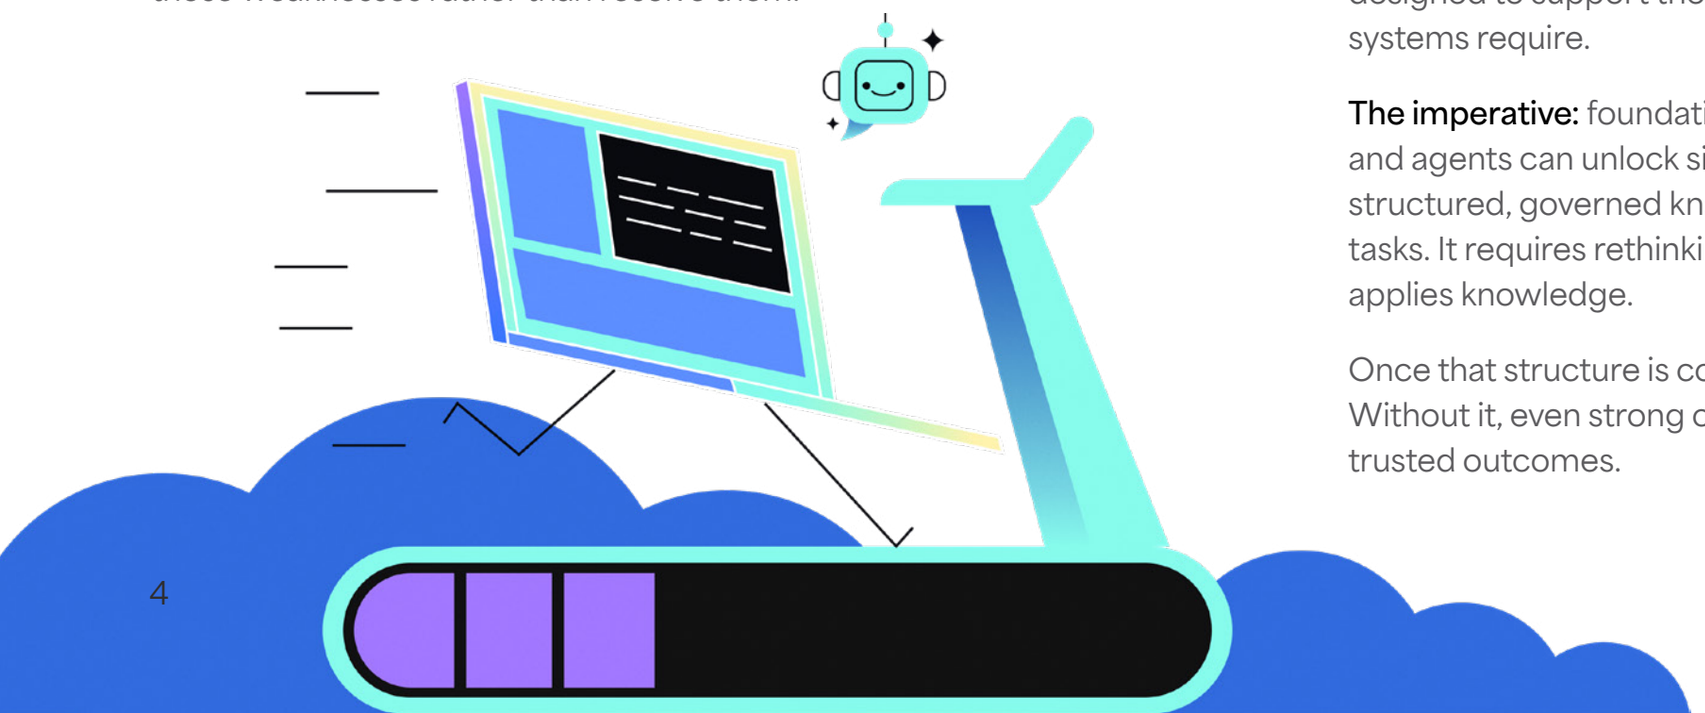

> AI doesn't make disconnected teams more productive. It makes them more expensively disconnected. The organisations winning right now invested in how their people work together before they invested in the tools."
>
> Billy Rollin, Managing Director, Accenture and Atlassian Partnership

The problem: fragmented tools, stale content, permission risks, and low trust in AI output hinder adoption. As one leader put it, most AI initiatives collapse for a simple reason: the organization's cognitive architecture breaks before the technology does.

The failure starts not in the model, but in the structure that was never designed to support the level of context and decision-making intelligent systems require.

The imperative: foundation models are mature, and enterprise search and agents can unlock significant ROI, but only when grounded in structured, governed knowledge. This is not just about automating tasks. It requires rethinking how your enterprise captures, connects, and applies knowledge.

Once that structure is coherent, AI becomes a practical accelerator. Without it, even strong capabilities struggle to produce consistent, trusted outcomes.

Non-negotiable principles for success

A successful AI rollout must be built on clear, non-negotiable principles that act as guardrails for safe, scalable change.

> Most customers come at their AI transformations as a revolution, not an evolution. The ones that land the biggest impact almost immediately are those that look for progress over perfection. These aren't looking for a new shiny thing, but analysing their existing workflows to work out where best to deploy their AI agents."
>
> Sherif Mansour, Head of AI, Atlassian

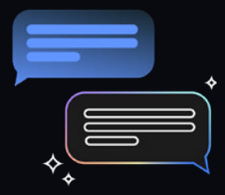

Human in the lead

Augment work, preserve agency, and design for trust with clear feedback loops. AI must serve the user, not the other way around.

Governed by default

Implement permissions-respecting search, data minimization, explainability, and auditability from day one. Governance must be an enabler, not a brake on progress.

Value led

Start small, measure impact continuously, and iterate to scale. Focus on incremental gains in each business unit that build into meaningful enterprise outcomes.

A practical pathway to enterprise AI adoption

To turn principles into practice, we build on a six-step framework grounded in Atlassian's approach to AI readiness. This framework reflects what is required to move from early experimentation to safe, scalable enterprise adoption.

01.
Streamline your technology stack

Conduct a cross-functional application inventory and prioritize systems that can integrate and share context effectively.

If you're already using Confluence as a knowledge hub, tools like Rovo can add an intelligent layer across your ecosystem. By connecting and contextualizing information from integrated systems, Rovo helps turn fragmented data into usable insight, improving decision-making and reducing friction around access and governance.

02.
Standardize the structure and content of knowledge

AI outcomes are probabilistic, but their reliability improves when they are grounded in structured, current, and permission-aware content.

Standardize your information architecture, define "golden sources," and use consistent templates for decision records, runbooks, and product specifications. This reduces ambiguity, improves contextual grounding, and leads to more consistent, high-confidence outputs.

03.
Establish governance and processes

Define acceptable AI use, data handling practices, and ownership roles, and implement controls such as restricted spaces, classification standards, and regular content recertification to ensure AI operates within trusted boundaries.

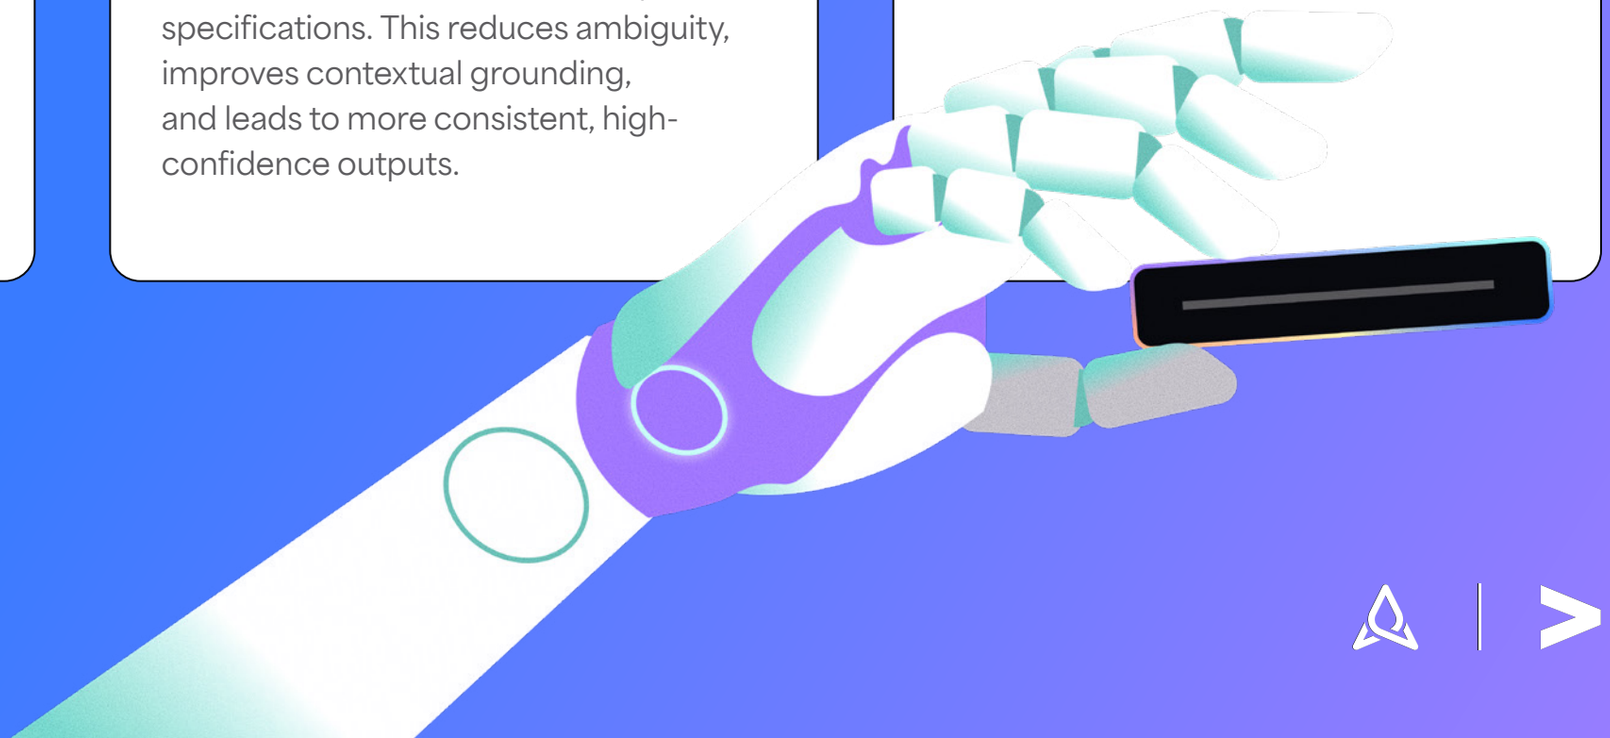

04.

Build excitement and buy-in

Secure executive sponsorship, start with low-risk use cases, and expand as teams gain confidence through visible wins and real adoption.

05.

Foster continuous engagement

Establish feedback loops and track adoption over time. Share progress alongside metrics to reinforce momentum and ensure AI becomes part of how teams work, not an isolated initiative.

06.

Develop advanced capabilities

Publish reusable prompts, enable role-based upskilling, and evolve from assistive use cases to more advanced automation as maturity increases.

Taken together, these steps define what it takes to move from AI pilots to sustained, enterprise-wide impact, and highlight where an execution and observability layer becomes critical in connecting structured knowledge to real workflows.

Governance is the foundation for scalable AI

Adopting enterprise AI means moving beyond basic tool implementation and establishing “governed by default” policies from day one. This is a non-negotiable condition for success, particularly because the race to adopt AI often exposes critical flaws such as permission risks and increased information security challenges that undermine outcomes.

Policy must be designed not as a brake on progress, but as an enabler of safe, scalable change.

In practice, this means defining acceptable AI use, establishing clear rules and taxonomies for data handling and classification, and implementing controls such as restricted spaces, content recertification, and label standards. Without this, AI systems operate on inconsistent or unverified information.

Only by enforcing a governed content lifecycle can organizations ensure that AI agents ground their responses in the most current and approved version of enterprise knowledge.

To translate governance into measurable impact, organizations should look to frameworks that guide system-level adaptation. The DORA AI Capabilities Model identifies foundational practices, including a clear AI policy and a healthy data ecosystem, that amplify the positive impact of AI.

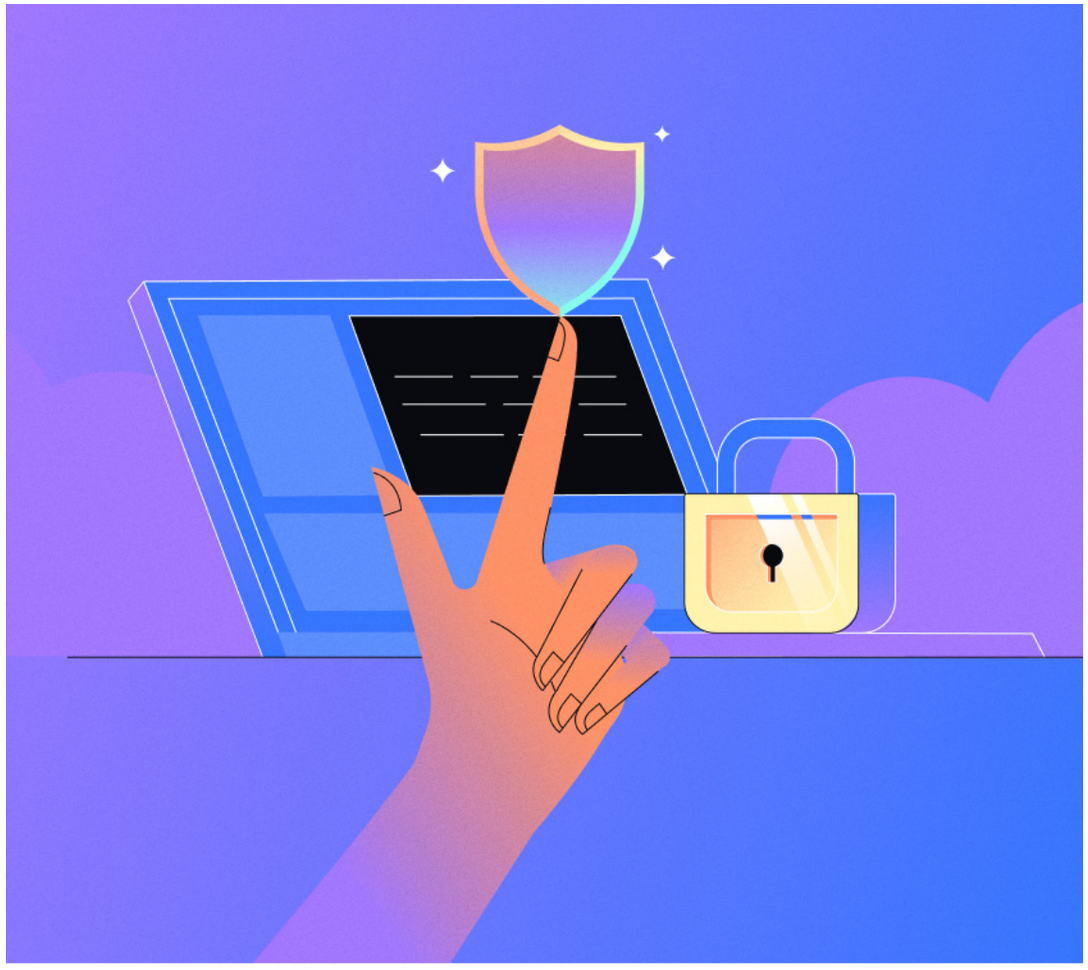

A clearly communicated AI stance provides psychological safety and boundaries for experimentation. This reduces resistance within workflows and strengthens AI’s impact on both individual effectiveness and organizational performance.

By prioritizing these governance capabilities, from defining data owner roles to ensuring AI-accessible internal data, organizations move past failed pilot projects and toward scalable, governed, and measurable enterprise success.

Where AI meets execution

While Atlassian Cloud provides a powerful knowledge and experience layer, companies like Appfire provide the acceleration and observability layers that connect AI-ready knowledge to business outcomes.

This is where your enterprise's execution data becomes a secure, governed knowledge source for intelligent systems. Instead of relying on static or incomplete context, AI can ground its responses in real-time, permission-aware operational data.

Without this layer, even well-structured knowledge remains disconnected from how work actually gets done. This is what turns AI from a capability into something teams can rely on in day-to-day operations.

With this foundation in place, the focus shifts to making knowledge usable, applying AI in real workflows, and measuring impact.

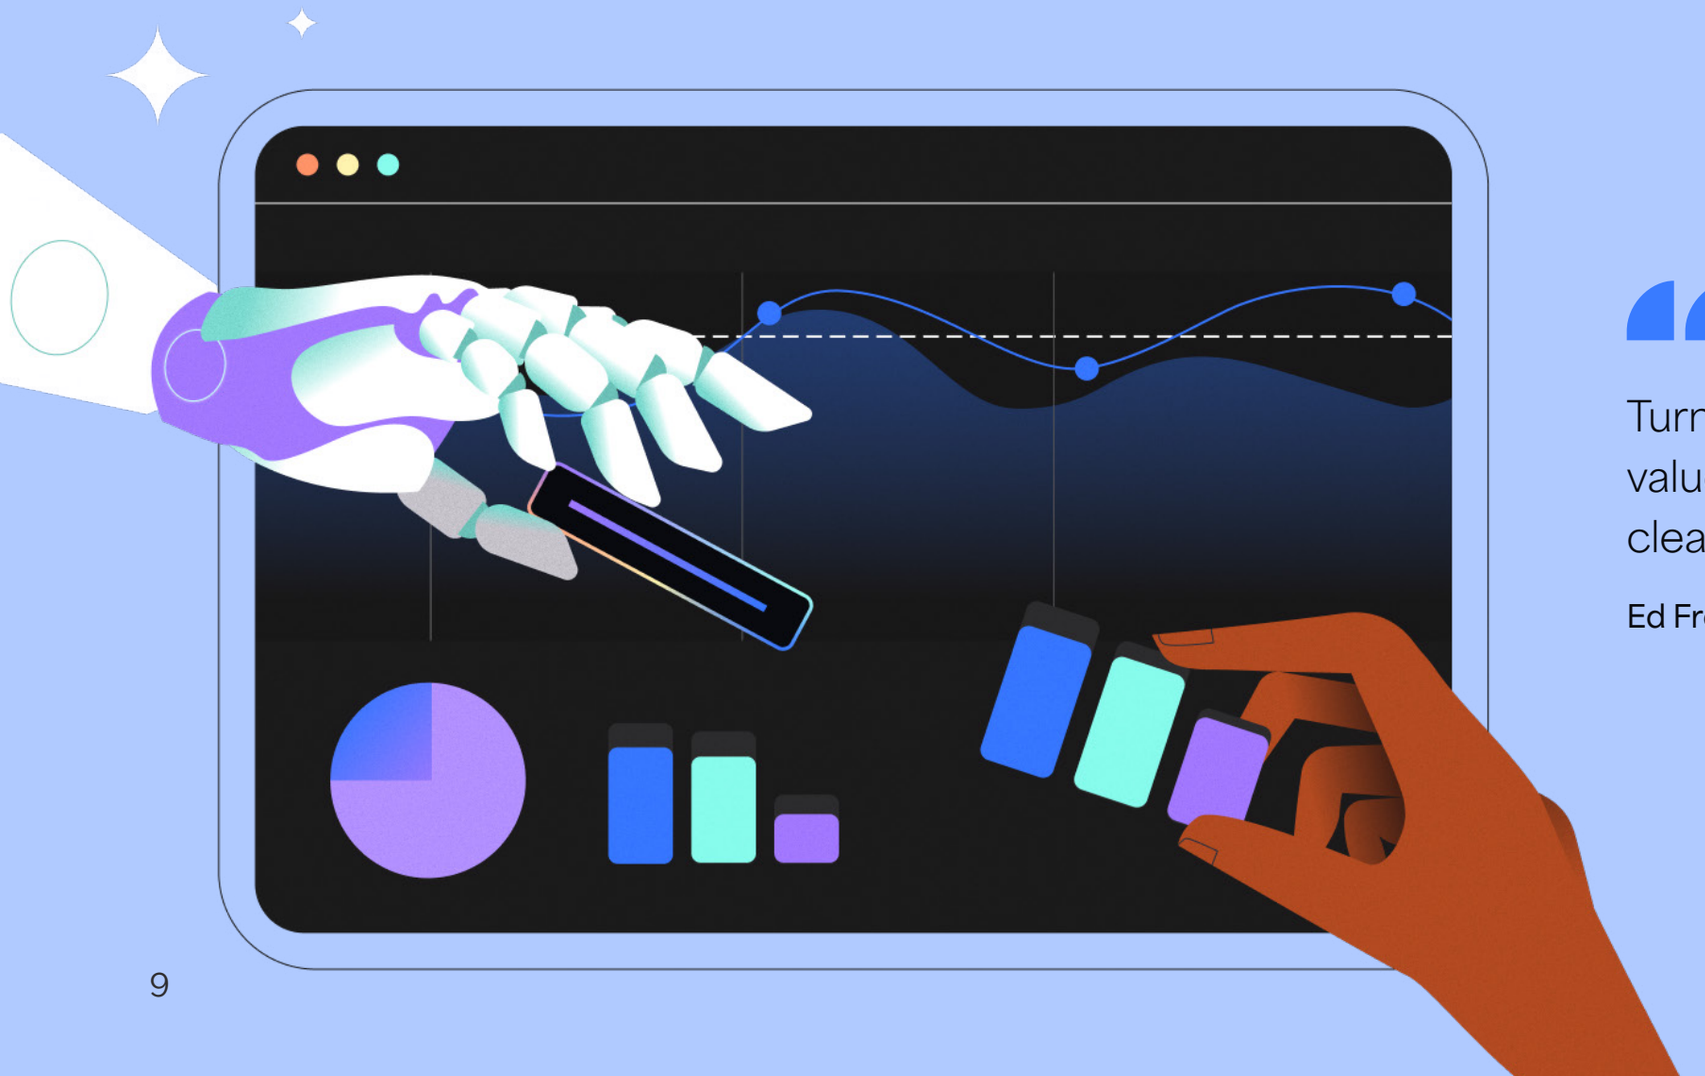

> Turning AI investment into real business value requires grounding adoption in clear priorities."
>
> Ed Frederici, Chief Technology Officer at Appfire

Make knowledge AI-ready

Governed content lifecycle and publishing

Implement auditable review and approval workflows so AI agents ground responses in the most current and approved version of knowledge, and control how content is published so only verified material is surfaced.

At Helmut Fischer, teams needed strict control over how documentation was created and approved to support a formal quality management system. By implementing structured approval workflows within Confluence, they ensured that only validated and approved content was accessible across the organization.

This reduced the risk of outdated or unverified information being used in day-to-day operations and created a more reliable foundation for downstream AI use cases.

Apps like Comala Document Management support this by enforcing staged review processes and controlling how content is published.

> Every time I see a customer who is amazed by Rovo, it's not just the tech. It's their behaviour. They've already built a culture of documentation and openness."
>
> Sherif Mansour, Head of AI, Atlassian

Data classification and governance

Standardize and enforce metadata, ownership, and content completeness so AI systems consume clean, permission-aware data.

In practice, this means applying consistent labels, ownership fields, and classification standards across both documentation and work items. For example, tools like JXL can be used to apply mandatory sensitivity labels and ownership fields across Jira work items at scale. Workflow automation apps such as Jira Miscellaneous Workflow Extensions (JMWE) can enforce conditions so work items are only surfaced to AI systems once they reach defined statuses and meet required criteria.

This ensures that AI outputs are grounded in accurate, validated information, improving both consistency and trust.

Apply AI in real workflows

Have teams start small and combine existing tools and platforms with AI agents to improve how work gets done.

Easy: AI-powered customer self-service and agent assist

AI can respond to customer requests when grounded in approved knowledge sources, such as Confluence knowledge bases. For more complex tickets, the same agent can generate a private internal comment, providing suggested solutions, relevant runbooks, and links to authoritative sources. This reduces search time and improves response accuracy.

Medium: AI-driven work assignment and routing

AI can match incoming work to the most suitable resource based on skills, roles, and availability.

By integrating with systems such as HRIS platforms, organizations can access up-to-date data on employee skills and capacity. AI can then analyze incoming work and cross-reference it with real-time team availability and expertise. This improves first-time assignment accuracy, reduces delays, and balances workload more effectively. Tools like Appfire's BigPicture can then be used to visualize and optimize resource allocation across projects.

Advanced: Operationalize your meeting notetakers

The opportunity is to move beyond transcripts and treat meeting content as a high-value input for downstream AI use cases.

AI can identify commitments, risks, and decisions from conversations and convert them into structured work items in Confluence or Jira. For example, sprint retrospectives can generate follow-up tasks, and risks discussed in meetings can be captured earlier and tracked within delivery workflows.

This creates a tighter feedback loop between discussion, planning, and execution. It depends on consistent tagging and classification so these items can be surfaced, tracked, and reused across systems. When structured properly, meeting outputs shift from one-time records to reusable enterprise knowledge.

Measure impact in execution

Track how AI-driven work performs inside real workflows and how it impacts outcomes.

Closed-loop efficiency measurement

Issues assigned through AI-supported workflows can be tagged, for example by assignment method. Appfire's BigPicture can then track throughput and resource utilization across portfolios and programs.

Reporting tools such as Dashboard Hub Pro provide visibility into metrics like cycle time, SLA adherence, and reassignments, comparing AI-assisted workflows with manual ones. This helps quantify improvements in efficiency, utilization, and quality.

AI-driven effort and cost forecasting

Ground forecasting in actual, time-spent data rather than estimates.

By integrating real-time work data, organizations can compare similar projects and forecast budget variance for new AI-enabled initiatives using historical execution metrics. This improves the reliability of cost projections and ties AI adoption directly to measurable outcomes.

Unified insights

Aggregate operational metrics across systems such as Jira and Confluence to create a connected reporting layer.

At Claritas, teams struggled to consolidate reporting across tools, limiting visibility into performance and making it harder to align priorities. By implementing centralized dashboards, they created a unified reporting layer that improved visibility and enabled more informed decision-making.

Apps like Dashboard Hub Pro support this by aggregating data into a single view and enabling real-time insight across systems.

Measurement and ROI: proving your impact

Successful AI adoption shows up in how work changes: how quickly teams find information, how consistently decisions are made, and how reliably outcomes improve.

Without clear measurement, it becomes difficult to separate real progress from isolated wins. Many organizations invest heavily in AI but lack visibility into where it is creating value and where it is failing to scale.

The previous section focused on execution-level signals. This section steps back to the leadership view.

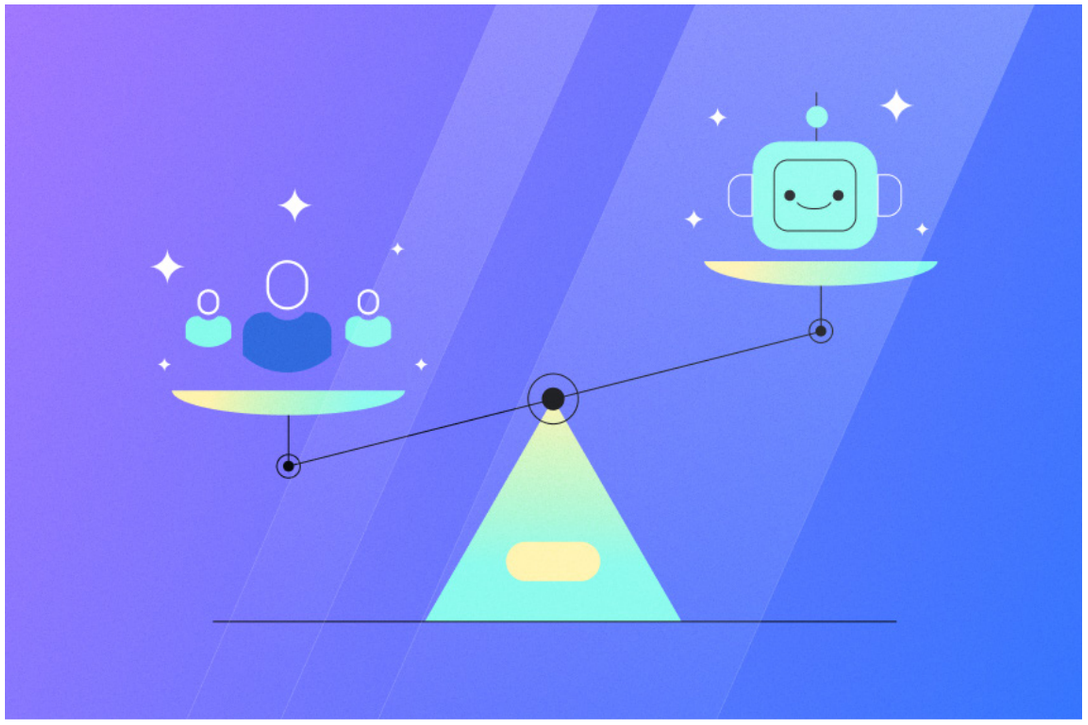

Metric	Definition	Target signal
Time-to-knowledge	Median time for a user to find authoritative info to complete a task	Decrease vs. baseline after rollout
Content freshness and approval	Share of pages within review interval, recertified, and approved via workflow	Sustained increase in compliant, approved content
Delivery velocity quality	Reviewed pull request rate, time to merge, iteration rate (DORA metrics)	Improves as knowledge standardizes
Value realization	Budget/effort allocation vs. hypothesized value delivered	Positive correlation between AI investment and outcomes
Effort and cost accuracy	Variance between planned effort and actual time spent on work	Variance decreases, cost forecasting reliability increases
AI data hygiene compliance	Share of content that meets mandatory classification and metadata standards	Sustained increase in content completeness, freshness
Unified insight coverage	Number of complex, cross-platform, executive-level questions answerable by AI	Increase in sophisticated report generation and coverage
AI answer quality	Show AI usefulness by when the user confirms results as “Solved” or gets answers with valid source references	Sustained increase in correctly answered, source-backed AI responses

Together, these signals show where AI is improving speed, trust, and quality, and where it still isn’t ready to scale.

AI outcomes depend on the system around them

AI does not create value on its own. It amplifies the systems it operates within.

If knowledge is fragmented and workflows are disconnected, AI will reflect those same limitations. If knowledge is structured, governed, and connected to execution, AI becomes a practical accelerator.

The difference is not the technology. It is how the organization is designed to use it.

This is where Accenture and Appfire excel. Working together to help organizations connect strategy, knowledge, and execution in practice.

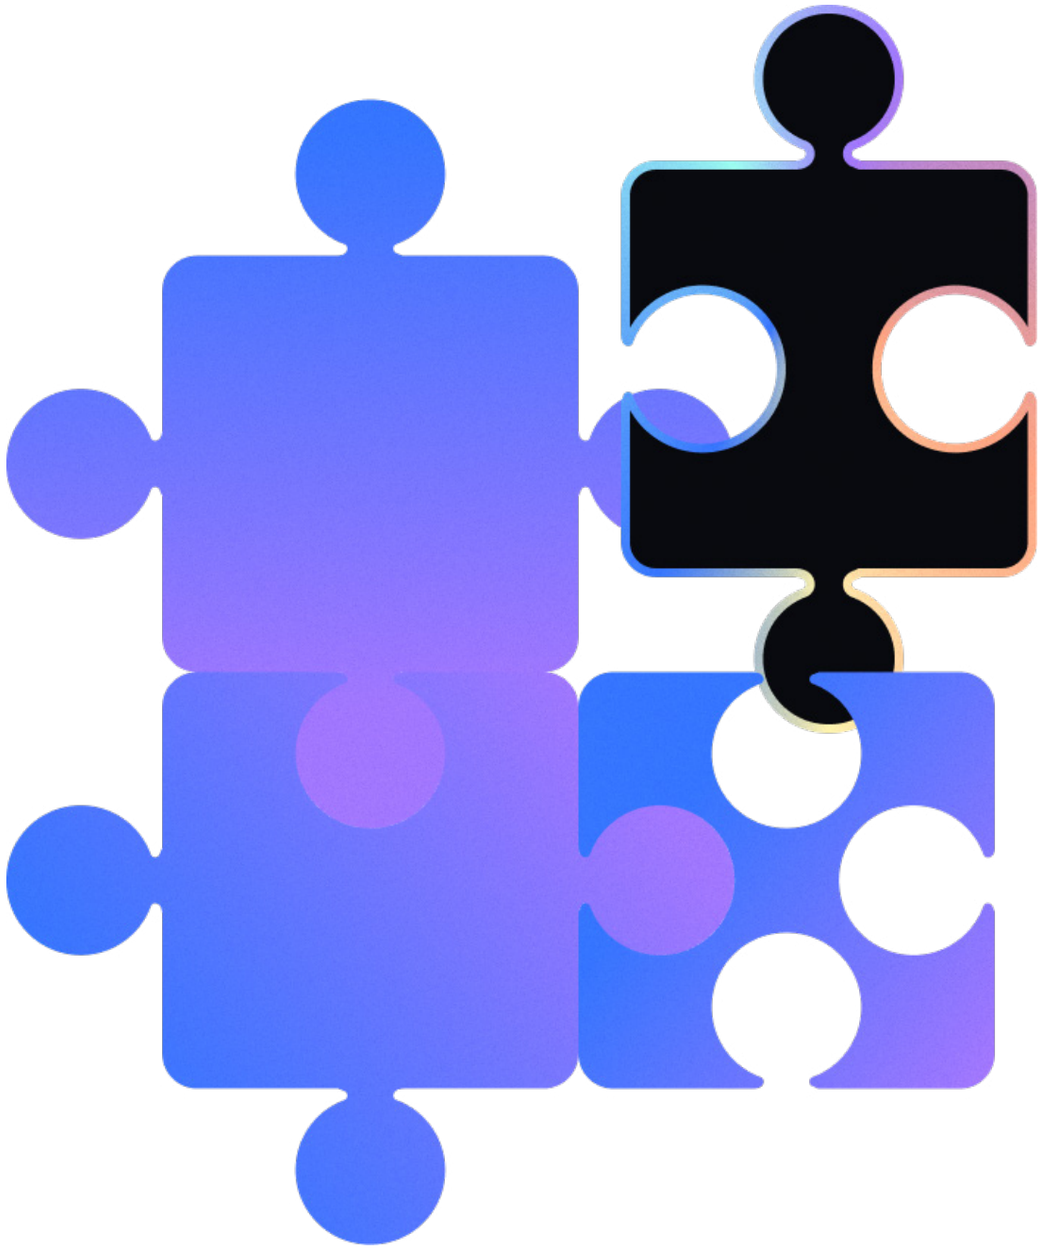

The Accenture and Appfire partnership

Accenture and Appfire together establish a unique and decisive leadership position within the Atlassian ecosystem. Our methodology blends Accenture's industry-leading solutions experience with Appfire's operational, automation, and measurement capabilities.

Our combined approach helps organizations realize the full value of their Atlassian investment by augmenting Atlassian platforms with third-party applications as the execution and observability layer that connects AI-ready knowledge to measurable business outcomes.

This paper reflects that combined perspective, focusing on a core challenge in enterprise AI adoption: moving beyond the "AI theater" of failed, frustrated, and disconnected pilot projects to deliver a pragmatic, ROI-focused pathway to real, scalable success.

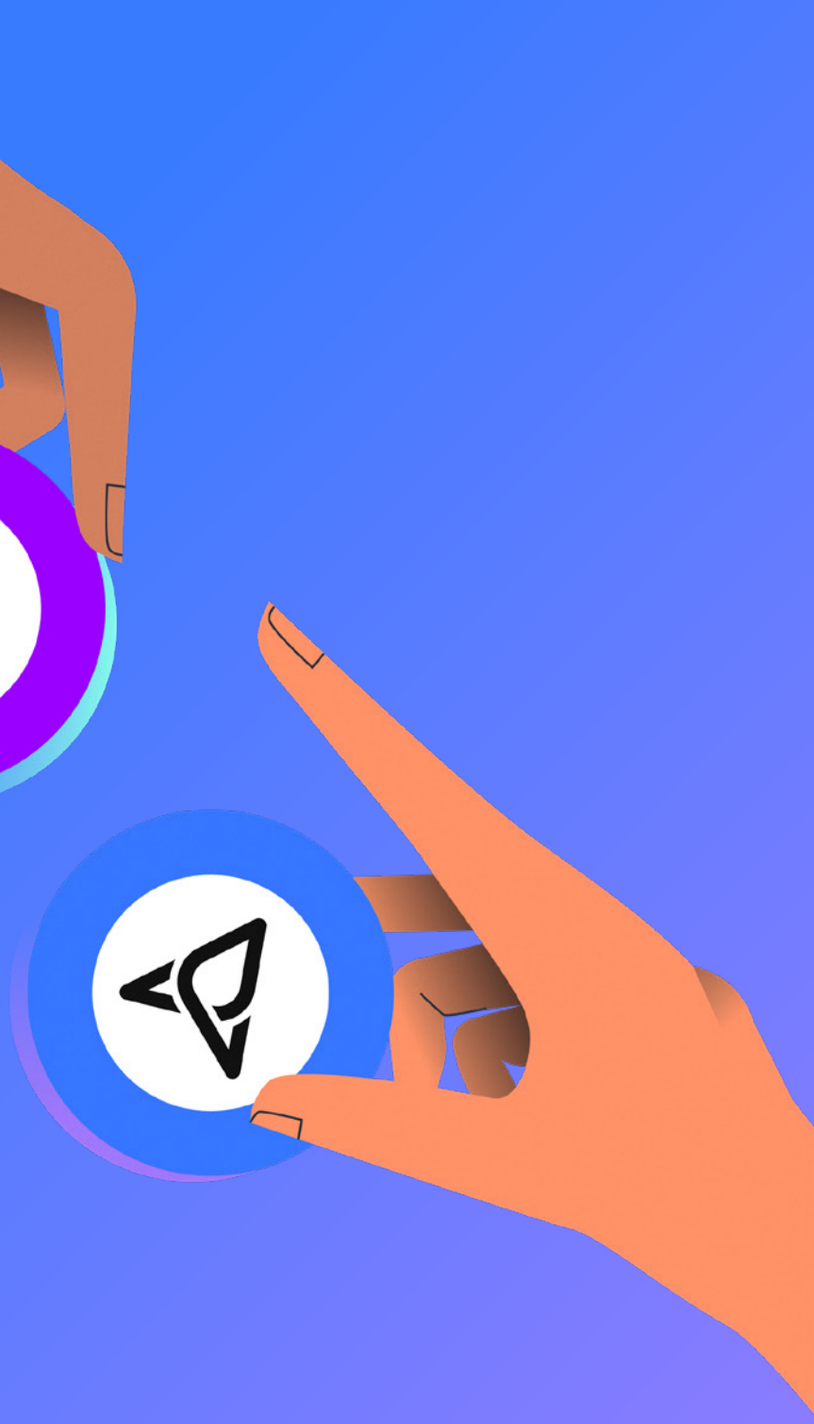

Acknowledgements

We'd like to acknowledge and thank the many contributors who helped shape this whitepaper, including leaders and practitioners across Accenture, Appfire, and Atlassian, as well as the teams who brought it to life.

Their insights and real-world experience help ground this perspective in practice and continue to shape how organizations approach AI adoption and execution in the years ahead.

Appfire authors

Emily Peet-Lukes

Steven Kling

Accenture authors

Arpita Visen

Dan Steigerwalt

Kit Friend

Notice of Liability

The information in this book is distributed on an "As Is" basis, without warranty. The publisher assumes no responsibility for errors or omissions, or for damages resulting from the use of the information contained in this book. Use of the information and instructions contained in this work is at your own risk.

© Copyright 2026 Appfire Technologies, LLC. All rights reserved.