




Beyond the Hype: The CTO's Guide to AI Reality_

A Build Circle Whitepaper





Introduction from the Managing Director

At Build Circle, myself and the team speak daily with CEOs, CTOs, and founders who are navigating the immense possibilities and pressures of AI. Right now, every technology leader across our network is asking the same question:

“What is everyone else really doing with AI?”

There is an overwhelming volume of media hype and boardroom expectation, but CTOs are desperate for practical reality and actionable insights. At Build Circle, we act as a “second brain” for tech leaders, convening cross-industry knowledge across fintech, utilities, and scaling tech brands to observe firsthand the real friction between AI expectations and engineering reality.

Our cross-industry analysis shows that while AI is driving massive acceleration, it is not a silver bullet. Real success requires navigating board-level expectations, restructuring into hyper-lean agile pods, establishing strict control over escalating token costs and, above all, ensuring foundational technology is modern and robust enough to handle the speed of AI.

This report cuts through the noise to offer a verified, real-world perspective, actionable strategies, and field-tested insights to inform your strategy.



Rory Seabrook,
MD at Build Circle

The Vantage Point: Strategic Engineering Partner_

This report is compiled from insights and direct experiences across the regulated technology sector, including fintech, utilities, and scaling technology brands. Build Circle works closely with technology leaders to bridge the gap between high-level AI ambitions and modern system architectures. Through flexible engagement models, we deliver trusted engineering excellence to modernise legacy platforms, unblock delivery pipelines, and build lasting in-house capability.

The AI Reality Gap_

BOARDROOM EXPECTATION

What leadership believes AI delivers

Instant headcount reduction

AI replaces engineers immediately

End-to-end speed uplift

10x coding = 10x delivery

Migration is a one-time project

Cloud = done, move on

ROI visible within months

AI pays back fast

Competitors are way ahead

We're falling behind right now

ENGINEERING REALITY

What technical teams are living

AI amplifies great engineers

Quality of talent matters more, not less

Bottlenecks shift upstream

Data and problem definition now constrain

Legacy is the blocker

Bad foundations limit AI's potential

Token costs are escalating

Spend is untracked and growing fast

Most peers are still in pilot

The gap is smaller than it looks



Contents_

Our delivery experience and industry data reveal five key themes driving modern AI adoption. This report breaks down each theme, delivering strategic takeaways and an action plan to help technology leaders assess their operating models and build high-performing engineering teams for the future.

1: The Expectation vs. Reality Gap.....	5
2: Human Critical Thinking in AI-Accelerated Teams.....	8
3: The True Cost and Control of AI	12
4: You Can't Accelerate on Bad Foundations.....	15
5: Navigating Oversold AI Hype and the Agentic Reality.....	19
Conclusion: Engineering Excellence.....	21

The Expectation vs. Reality Gap

There is a profound friction between board-level expectations and the reality of on-the-ground engineering.

The Competitor Paranoia

A frequent observation is the widespread belief among CTOs that their industry rivals are vastly further along in their AI journeys than they actually are. This “competitor paranoia” has triggered a wave of reactive decision-making. Research by the IBM Institute for Business Value reveals that 64% of CEOs¹ admit Fear of Missing Out (FOMO) drives some technology investments before they have a clear understanding of the underlying business value.

Cross-industry data confirms that this competitor anxiety is largely built on an illusion:

- **Performative AI:** A recent Writer enterprise study found that 75% of executives admit their current AI strategy is “more for show” than real operational transformation,² while BCG AI Value Gap study states 60% of companies generate no material value from AI.³
- **The Maturity Vacuum:** According to McKinsey, only 1% of organisations have achieved a “mature” AI operating model where agents are seamlessly integrated into production software pipelines.⁴

While there is a gap between early adopters and cautious followers, modern AI tooling and standardised API ecosystems mean **the gap is far smaller than leaders fear**. Because late-movers do not need to invent underlying foundation models, organisations with clean data and modernised cloud infrastructure can bridge this competitive gap in a matter of months, bypassing the costly trial-and-error phases that tripped up early experimenters.

64%

of CEOs⁵ admit Fear of Missing Out (FOMO) drives some technology investments before they understand the business value

1 & 5. [IBM Study: CEOs Double Down on AI While Navigating Enterprise Hurdles, May 2025.](#)

2. Writer, [Enterprise Generative AI Adoption Survey 2026](#), 2026.

3. Boston Consulting Group (BCG), [Are You Generating Value from AI? The Widening Gap](#), September 2025.

4. McKinsey & Company, [Superagency in the Workplace: Empowering People to Unlock AI's Full Potential at Work](#), 2025.



Managing the Board

A primary frustration for CTOs today is managing boards and executives who treat AI as a “silver bullet” that will instantly solve all systemic problems and drastically cut costs. Two years ago, boards naively assumed AI simply meant they could instantly cut their engineering teams. Today, technology leaders are tasked with tempering these unrealistic expectations by proving that successful AI adoption depends primarily on workforce readiness, cultural alignment, and workflow redesign.

Gartner research reveals that only 20% of CxOs believe their workforce is AI-ready,⁶ warning that AI implementation requires significantly more change management than any previous technology era. Furthermore, industry data from Boston Consulting Group indicates that up to 70% of AI transformation success hinges on people and process change, rather than the underlying algorithms.⁷ To be successful, CTOs must redirect board conversations away from speculative cost-cutting and onto the pragmatic engineering foundations that make AI delivery work.

only
20%

of CxOs believe their
workforce is AI-ready⁸

Moving AI Beyond Isolated Wins to Systemic Impact

Engineering teams are indeed reporting massive 4x to 10x improvements in specific coding throughput. However, this acceleration is frequently confined to narrow tasks and it does not automatically translate to significant speed increases across the entire software delivery lifecycle.

Instead of accelerating the entire software lifecycle, AI has shifted operational bottlenecks upstream. Because writing code is no longer the primary constraint, high-performing technology teams find their main friction points lie in defining the right business problems and ensuring data readiness before a single line of code is written. Research from Gartner reinforces this delivery reality, with 67% of enterprise IT leaders citing process fragmentation and handoff failures as their primary delivery obstacles, far outweighing developer coding speed.⁹

According to McKinsey,
88% of organisations
have added GenAI to at
least one function, yet
fewer than 20% saw an
improvement in overall
software delivery speed.¹⁰

6 & 8. Gartner, Inc., [Gartner Predicts by 2027, 50% of Enterprises Without a People-Centric AI Strategy Will Lose Their Top AI Talent](#), May 2026.

7. Boston Consulting Group (BCG), [AI Transformation Is a Workforce Transformation](#), February 2026.

9. Gartner, Inc., [Software Engineering Leaders: Strategic Research and Role Guidance](#), 2026.

10. McKinsey & Company, [The AI Revolution in Software Development](#), Research Report.

Strategic Takeaways

- ✦ **The Industry Reality:** Across most sectors, peer organisations are at roughly the same stage of adoption, meaning tech leaders are rarely as far behind as market noise implies.
- ✦ **The Boardroom Gap:** Frustration at executive levels usually stems from viewing AI as an immediate silver bullet rather than an operational capability that requires structured enablement.
- ✦ **Sustainable Acceleration:** Lasting commercial value depends on process redesign, workforce readiness, and measurable output.

Your Action Plan

- ➡ **Audit the Entire SDLC:** Avoid treating AI as a blanket accelerator. Audit your software delivery lifecycle to identify where the bottlenecks are shifting to as coding speed increases.
- ➡ **Differentiate Velocity from Acceleration:** Distinguish between rapid prototyping speed and end-to-end software delivery throughput when presenting metrics to executive boards.
- ➡ **Establish Reality-Based Benchmarks:** Frame AI capability around verified operational metrics, such as a 20% release acceleration, rather than unverified 5x–10x hype claims.
- ➡ **Manage Peer Benchmark FOMO:** Anchor investment decisions in internal architectural readiness and measurable business outcomes rather than competitor adoption anxiety.
- ➡ **Redesign Workflows:** Organisations achieving meaningful returns are nearly 3 times as likely to have fundamentally redesigned their workflows when deploying AI rather than just applying it to existing tasks.

Human Critical Thinking in AI-Accelerated Teams

AI is fundamentally restructuring tech departments and operating models, placing an even higher premium on human capability.

Reshaping Team Structures

The traditional software engineering team of 8 to 10 people is rapidly shifting. With AI acting as an enabler, it is now incredibly viable to rely on agile, hyper-lean pods of just 2 to 3 people to deliver business value at high speeds. By 2030, tech hierarchies are expected to be much flatter, with distributed “pods of AI capability” sitting directly alongside business teams to own their outcomes directly.

Traditional Engineering Team Structure (8-10 people)


Product Owner


Scrum Master


QA Leads


4-5 Engineers


DevOps

Modern AI-Enabled Pod Structure (2-3 people)


Human Engineer + AI Agents


Direct Collaboration with Business Units



CASE STUDY:

Accelerated AI Capabilities in Higher Education

The Challenge: A private-equity-backed education provider set out to harness Generative AI to elevate its course offering and lead in EdTech. However, academic teams spent excessive manual hours writing exam prep material, while coaching teams lost critical student contact time to routine session administration.

The Human + AI Execution: Build Circle deployed a lean senior engineering pod directly into the client's product team. Senior engineers defined the system architecture and edge-case validation rules, using LLMs to accelerate routine component building while keeping human academics in the loop for output validation.

The Business Outcomes:

Delivered **multiple millions** in business value within 12 months.

Achieved a **75% reduction** in content review time.

Generated a **300% increase** in overall team productivity.

Automated session documentation, increasing direct student engagement time.

The Shift from Code Generation to System Architecture

With AI capable of generating massive volumes of output, there is a severe risk of engineering teams going down “rabbit holes” and creating immense waste. High-performing organisations will increasingly rely on AI for everyday throughput, but they will need human critical thinkers to drive the models, course-correct, and understand what the business needs. Here, strong human engineering capability acts as the multiplier that supercharges AI performance.

High-quality human engineers are essential to steer models because AI-generated code is frequently unrefined. Left uncurated, AI-generated code introduces compounding structural debt. Recent analysis of over 620 million code changes shows an eightfold increase in duplicated code blocks alongside an unprecedented drop in active refactoring.¹¹ As a result, the developer's role is fundamentally shifting with Gartner predicting that by the end of 2026, 75% of software developers will spend more time orchestrating and architecting systems than writing code directly.¹²

11. GitClear, [Write-Only Mode: AI Code Quality & Maintainability Gap Research](#) (Analysis of 620M+ Code Changes), 2023–2026.

12. Gartner, Inc., [Gartner Identifies the Top Strategic Trends in Software Engineering for 2025 and Beyond](#), July 2025.

The Succession Planning Crisis

A massive, unsolved puzzle in the industry is the future of talent. With AI agents taking over the rote tasks traditionally assigned to entry-level roles, demand for junior developers has collapsed by an estimated 40% in companies that have deployed these tools.¹³ Analysis from the Stanford Digital Economy Lab reveals that U.S. entry-level tech job postings plummeted by 67% as companies shifted toward a “Senior + AI”¹⁴ model, leaving employment for software developers aged 22 to 25 roughly 20% below the late-2022 peak.¹⁵ Closing this gap requires rebuilding the talent entry path around AI rather than freezing it. As business leaders, creating new avenues for early-career talent in an AI-first environment is an operational necessity.

Forward-thinking engineering organisations are adapting their onboarding models to give junior talent immediate exposure to AI-accelerated delivery environments. Rather than assigning traditional, isolated “sandbox” tasks, high-performing teams integrate early-career hires directly into active engineering pods from day one. By pairing structured AI tool training with dedicated mini-project ownership, organisations can fast-track critical thinking and system architecture skills, developing the capabilities needed to transform junior engineers into future senior leaders.

BUILDING THE AI TALENT PIPELINE

“It is our responsibility to create new paths for young people to enter the tech job market. Today, graduates are seeking AI-enabled careers. In this period of flux, we must embrace the challenge of bringing early-career talent into the workforce — integrating them directly into project delivery alongside senior engineers so they can experience AI-driven delivery first-hand.”



James Parsons,
CEO at Build Circle

13. Second Talent, [How AI Is Changing Engineering Talent Demand](#), 2026.

14 & 15. Stanford Digital Economy Lab & ADP Research Institute, [Canaries in the Coal Mine: Analyzing Junior Developer Employment Shifts in the AI Era](#), November 2025.

Strategic Takeaways


- ✦ **Operating Model Shift:** Software engineering team structures are shifting from large traditional delivery units to hyper-lean 2—3 person pods operating alongside business stakeholders.
- ✦ **Senior Orchestration:** AI tooling speeds up code production, making senior architectural oversight, context verification, and human critical thinking essential to prevent runaway technical waste.
- ✦ **Talent Pipeline Protection:** Protecting long-term engineering capability requires structured mentorship and mini-project ownership, ensuring junior developers continue progressing into senior orchestrators.

Your Action Plan

- ➔ **Restructure into Lean 2—3 Person Pods:** Transition legacy 8—10 person engineering teams into hyper-lean pods where senior engineers and product owners sit side-by-side with AI capabilities.
- ➔ **Enforce Human Architectural Oversight:** Require senior engineers to validate system context and edge cases before deploying AI-generated code to production pipelines.
- ➔ **Shift Developer Roles to Orchestration:** Train engineering teams to act as system architects and orchestrators rather than routine code syntax generators.
- ➔ **Reconstruct Career Pathways:** Begin mapping out alternative career pathways for junior talent to navigate the 40% collapse in entry-level demand.

The True Cost and Control of AI

The era of limitless AI experimentation is ending, replaced by a need for strict financial and operational control.



Token Discipline and Escalating Costs

While initial experiments were cheap, running AI at scale is expensive. The operational costs of AI providers and token consumption are ramping up quickly. According to Gartner's June 2026 report on optimising costs, at least 50% of GenAI initiatives will exceed their planned budgets between now and 2028 due to poor architectural choices and a lack of operational expertise.¹⁶

at least
50%

While model provider token unit costs are decreasing due to chip efficiencies, the move towards agentic AI reverses those savings. Autonomous agentic models require 5 to 30 times more tokens per task than standard GenAI chatbots. Because token consumption scales far faster than unit prices drop, overall inference costs are escalating rapidly. Gartner expects inference to account for at least 70% of an AI model's lifetime operational costs.

will exceed their
planned budgets
between now
and 2028

As vendors pivot from flat seat fees to usage-based pricing, this consumption curve represents a fundamental shift in software economics. **Gartner predicts that by 2028, annual AI coding expenses will surpass the average software developer's salary in organisations without strict token governance.**¹⁷ This dynamic is forcing CTOs to implement rigorous discipline from the outset, requiring engineering teams to prove commercial return before token usage spirals out of control.

16 & 17. Gartner, Inc., [Gartner Predicts AI Coding Costs Will Surpass Average Developer's Salary by 2028 as Token Consumption Surges](#), June 2026.

The Visibility Problem

Tech leaders currently face a critical governance blind spot due to a lack of visibility and control over how individual engineers utilise AI tooling day-to-day. While one developer might use an enterprise AI licence effectively across requirements framing and test coverage, an engineer in an adjacent team may merely be “dabbling”, or using unsanctioned models altogether. Research by Gartner indicates that over 70% of technology organisations lack central visibility into team-level AI usage, output quality, and data exposure.¹⁸ Without standardised governing control, leadership cannot distinguish productivity gains from unmonitored experimentation.

The Tooling Shift: Why Claude is Winning

In the race for reliable enterprise software engineering tooling, cross-industry observations reveal a distinct shift towards Anthropic’s Claude models. Engineering teams are increasingly standardising on Claude due to its superior performance on complex software benchmarks. Claude models lead industry evaluations like SWE-bench (resolving over 49% of real-world software issues autonomously) and its 200,000+ token context window, which enables deep comprehension of massive, legacy codebases.¹⁹ Beyond raw intelligence, tech leaders cite its enterprise-grade security posture, non-training data privacy guarantees, and open protocols like MCP (Model Context Protocol) as key drivers. This malleability allows teams to prototype, test, and integrate models directly into CI/CD pipelines without exposing proprietary logic.

THE TOOLING SHIFT

“Claude is transforming how quickly we move from discovery to delivery. In client workshops, we can take live discussions, build functional prototypes in real time, and present a working model before the meeting even ends.”



Henry Keen,
CTO at Build Circle

18. Gartner Research, [AI Governance, Observability & TRISM \(AI Trust, Risk and Security Management\)](#), Strategic Research Report, June 2026.

19. Anthropic Research & SWE-bench, [Claude 3.5 Sonnet & Model Context Protocol Briefing; SWE-bench Evaluation Leaderboard](#), 2026.

Strategic Takeaways

- ✦ **The Inference Reality:** Financial friction escalates rapidly as organisations transition from low-cost pilot testing to production-scale inference and recurring token spend.
- ✦ **Active Cost Governance:** Uncontrolled AI usage risks budget overruns unless technology leaders enforce model right-sizing, caching protocols, and real-time team usage tracking.
- ✦ **ROI Discipline:** Commercial control requires evaluating every GenAI deployment against clear unit economics, keeping token consumption aligned with business outputs.

Your Action Plan

STRATEGIC GOVERNANCE

- ➔ **Standardise AI Platforms:** Establish centralised platform governance to gain complete visibility over model usage, data controls, and output quality across teams.
- ➔ **Enforce an ROI Framework:** Track “token cost vs. business value” against the 70% lifetime operational inference threshold to verify that AI investments outperform traditional delivery baselines.

OPERATIONAL EXECUTION

- ➔ **Context Engineering & Task-Model Alignment:** Train developers to streamline context prompts and intelligently route tasks - mapping routine queries to smaller, low-cost models and reserving premium LLMs for critical system design.
- ➔ **Embed Continuous Token Governance & FinOps:** Combine automated budget caps and real-time API monitoring with regular sprint retrospective audits to maintain financial discipline across development and production workflows.

You Can't Accelerate on Bad Foundations_

The massive acceleration promised by AI carries significant risks if an organisation's underlying architecture cannot support the speed and demands of AI tooling.

The Prerequisite for AI

To be a highly effective user of AI, an organisation must possess strong foundational engineering capabilities. Accelerating processes with AI will only cause a business to crash into underlying system bottlenecks if the basics are not securely in place.


Deloitte benchmark data reveals that technical debt consumes between 21% and 40% of IT budgets. With leaders reporting that this debt is actively increasing, teams are forced to divert more than 20% of project budgets simply maintaining legacy complexity rather than driving new delivery.²⁰

Furthermore, applying AI to existing technical debt creates a dangerous feedback loop. Organisations that fail to monitor code quality while using AI code assistants report massive spikes in security vulnerabilities, with studies showing that roughly 45% of AI-generated code fails basic OWASP security tests.²¹

20. Deloitte, [Technical Debt's Penalty on Value and Growth \(2026 Global Technology Leadership Study\)](#), Deloitte Insights, March 2026.

21. Veracode, [Spring 2026 GenAI Code Security & Vulnerability Analysis](#), 2026.

The Cloud Maturity Bottleneck



While 99% of organisations report that AI adoption is accelerating their cloud investments, 88% acknowledge that their current cloud maturity actively puts those initiatives at risk. NTT DATA survey data highlights that only 14% of enterprises have reached a fully "cloud-evolved" state capable of supporting and scaling production AI workloads.²²

Modernisation as the Enabler

Legacy systems act as a direct tax on innovation velocity. Research shows that 50% of IT leaders identify legacy applications and fragmented data platforms as the primary barrier to AI transformation.²³ Before an enterprise can operationalise AI at scale, it must build a repeatable "factory" engine to modernise core applications, resolve architectural debt, and migrate to cloud-native environments, which serve as the execution layer for AI operating models.

Modernising the underlying tech stack liberates proprietary data, turning necessary system maintenance into a long-term strategic advantage. While commodity LLMs and off-the-shelf AI tools are accessible to every competitor, an organisation's unique, historical data foundation is the one asset that cannot be replicated. Modernising the technology stack cleanses data pipelines, eliminates operational silos, and establishes automated governance. This transforms static, fragmented data into real-time, AI-ready context and turns legacy technical burden into a defensible commercial advantage.

22. NTT DATA, [Cloud-Led Innovation in the Era of AI: Global Enterprise Maturity Study](#), 2026.

23. NTT Global, [Cloud Modernization: How AI Is Exposing the Limits of Legacy Systems](#), 2026.



SAFETY AT SCALE

“You cannot layer 2026 AI capabilities onto a 2010 legacy foundation and expect speed. Modernisation is about removing system friction so your proprietary data can feed the AI engine safely and at scale.”



Henry Keen,
CTO at Build Circle

CASE STUDY:

Enterprise Payment & Cloud Modernisation

The Challenge: A major retail payments provider supported critical settlement infrastructure across 400+ server applications. The estate was bottlenecked by decades of technical debt, legacy message queues, and ageing codebases that choked delivery and innovation speed.

The Human + AI Execution: Build Circle embedded a senior engineering pod to establish a “modernise where necessary, migrate where possible” refactoring model. Senior engineers resolved architectural bottlenecks and defined migration guardrails, using AI tooling to accelerate routine code refactoring and test coverage.

The Business Outcomes:

Modernised 400+ application workloads without committing to costly multi-year rewrites.

Enabled **rapid, continuous delivery** at a fraction of the cost and risk of traditional transformations.

Systematically **retired ageing software dependencies** and infrastructure risk.

Established a **repeatable migration path** to transition workloads smoothly to the cloud.

Strategic Takeaways

- ✦ **The Legacy Speed Cap:** Deploying AI onto tightly coupled, legacy infrastructure creates a false sense of speed that collapses when code hits production bottlenecks.
- ✦ **Pragmatic Modernisation:** Accelerating delivery requires resolving monolithic architectural debt and refactoring core backend systems first.
- ✦ **Standardised Protocols:** Connecting foundational AI capabilities cleanly to backend data requires robust API integration layers and standardised Model Context Protocols.

Your Action Plan

- ➡ **Do Not Lay AI on Debt:** Resist the urge to layer automated AI coding tools or agentic workflows on top of fragile, undocumented legacy codebases without senior oversight.
- ➡ **Prioritise the Cloud Execution Layer:** Accelerate cloud migration and containerisation to establish a resilient, elastic environment where AI agents can execute safely.
- ➡ **Turn Proprietary Data into a Strategic Advantage:** Treat application modernisation as a data enablement exercise - structuring, securing, and cleansing enterprise data pipelines to feed LLMs with high-fidelity context.
- ➡ **Establish Repeatable Modernisation Engines:** Deploy structured "modernisation factories" to systematically reduce technical debt in parallel with AI adoption, ensuring legacy friction does not erode velocity.

Navigating Oversold AI Hype and the Agentic Reality_

As the initial frenzy of generative AI settles, the industry is facing a reality check regarding what is coming next.

The “Oversold” Reality

While industry figures like Mark Zuckerberg have suggested the market may currently be “oversold” on the immediate capabilities of AI, this echoes what technology leaders are experiencing on the ground. Boards and executives initially viewed AI as an instant, magical fix that would immediately slash costs and cut team sizes. However, as the market enters a more mature phase, organisations are learning that successful AI adoption ultimately hinges on people, cultural alignment, and disciplined change management.

The Shift to Agentic Operating Models

Despite the market being oversold on instant ROI, the future of high-performing tech organisations lies in agentic execution. Agentic AI capabilities are growing at a staggering 119% CAGR.²⁴ Gartner predicts that 35% of traditional point-SaaS tools will be entirely replaced by autonomous AI agents by 2030, putting \$234 billion in enterprise application software spend at risk.²⁵

Replacing point-SaaS tools with autonomous agents requires tech leaders to rethink integration architecture. Moving beyond basic prompt engineering, enterprise teams are deploying Model Context Protocols (MCPs) and API connection layers to execute complex, end-to-end operational workflows.

Accelerating the Catch-Up Window

Because modern multi-agent frameworks build on standardised foundation models and open protocols, the time required to achieve delivery maturity has compressed significantly. Organisations with clean data pipelines and cloud infrastructure can bypass early experimenters’ trial-and-error phases, shortening requirements-to-deployment cycles by 35% to 50% within months.²⁶

CLOSING THE GAP

“By bringing in the right strategic partner to cut through the hype and implement an agentic framework, organisations can close the competitor gap incredibly quickly.”



Rory Seabrook,
MD at Build Circle

24. Software Strategies, [Gartner Forecasts Agentic AI Overtakes Chatbot Spending by 2027](#), February 2026.

25. CIO, [Agentic AI Puts \\$234B in Enterprise SaaS Spending at Risk](#), Gartner Says, 2026.

26. Gartner, Inc., *AI Delivery Orchestration: Multi-Agent Software Delivery Frameworks*, Gartner Software Engineering & Delivery Practice, Research Note, January 2026.

Strategic Takeaways

- ✦ **Agentic Execution:** Enterprise software is moving beyond chat-based copilots toward API-driven multi-agent workflows executing end-to-end tasks directly across backend systems.
- ✦ **Point-SaaS Disruption:** Autonomous AI agents are bypassing seat-based software interfaces, placing traditional point-SaaS expenditures at direct risk of replacement.
- ✦ **Fast-Follower Advantage:** Technology leaders can bypass early adopter trial-and-error by leveraging open protocols and foundation models to compress delivery cycles efficiently.

Your Action Plan

- ➡ **Architect for Multi-Agent Workflows:** Transition from simple prompt interfaces to API-driven multi-agent systems that execute end-to-end tasks directly across backend platforms.
- ➡ **Audit Point-SaaS Expenditures:** Review software portfolios to identify point-SaaS tools vulnerable to agentic replacement, reallocating seat-based spend to internal agentic workflows.
- ➡ **Capitalise on Fast-Follower Velocity:** Bypass early adopters' trial-and-error phases by leveraging standardised foundation models and open protocols to compress delivery cycles.

Conclusion: Engineering Excellence_

If navigating this landscape feels overwhelming, take a breath. You are doing just fine. Beyond the hype, AI capability grows through solid architectural foundations and steady, disciplined delivery. You do not need to pivot your strategy every time a new model drops or an industry headline predicts immediate disruption. Knee-jerk technology shifts driven by market paranoia only generate waste, inflate token budgets, and distract from your core delivery goals.


While this paper has focused extensively on AI, the fundamental laws of business and software delivery have not changed. AI is a powerful enabler that can accelerate how teams build, scale, and iterate, but technology generates value only when it serves broader strategic objectives.

When the market tells you to “just use AI,” the pragmatic counter-question must always be:

“What business outcome are we trying to achieve, and how does this serve our customers?”

To answer that question, organisations must build and maintain high standards of **Engineering Excellence**. In today’s market, engineering excellence means deploying teams who do more than just cut code - they must “look up from the keyboard” to hyper-focus on wider commercial priorities.

Whether it is reigning in runaway token costs, modernising legacy data centres, or restructuring into lean, high-performing pods, technology must serve a tangible business purpose. AI will accelerate your delivery, but only when layered on top of solid technical foundations and clear customer focus. And ultimately, successful AI adoption requires deploying pragmatic problem-solvers who take full ownership of the final result.





About Build Circle_

Build Circle operates at the intersection of engineering rigour and commercial reality. By pairing AI-accelerated delivery with senior technical expertise, we help tech leaders expand their engineering capabilities and drive measurable business outcomes on time.

Our experienced teams own delivery problems beyond the codebase to unlock productivity, modernise core systems, and maximise technology ROI.

Ready to cut through the hype and scale your engineering reality?

Visit buildcircle.co.uk or contact us at hello@buildcircle.co.uk to start a conversation.

