

Your AI Strategy Is Already Behind

MXA Consulting — Technology & Transformation Insights | February 2026

By Dr Rishni Ratnam, CEO, MXA Consulting

The technology underpinning your current AI pilots has moved a generation forward. Strategies defined even twelve months ago were defined in a different era. Organisations that settle into productionising their limited pilots risk optimising for a capability baseline that no longer exists.

In November, we published *Path to Value With AI*¹, which argued that the real barriers to AI impact are the same ones we've faced with digital transformation for twenty years: weak architecture, fragmented data, underskilled people, and leadership that treats technology as a cost centre. That argument still holds. But the urgency has increased significantly.

What happened

Anthropic released Claude Opus 4.6 with a one-million-token context window and a capability called Agent Teams, where multiple AI agents coordinate autonomously on shared work.² To demonstrate, an Anthropic researcher set 16 agents to build a C compiler from scratch. No human wrote code. The agents coordinated through git, resolved their own conflicts, and produced 100,000 lines of working code that compiles the Linux kernel.³

On the same day, OpenAI released GPT-5.3-Codex. Two weeks later, Google released Gemini 3.1 Pro. Three competing models, each optimised for different things: depth, speed, and cost efficiency respectively. The emerging pattern among high-performing teams is to route different work to different models — not to pick one vendor.

Why this matters for your organisation

The thing you are trying to productionise was designed against a capability baseline that no longer exists.

Deloitte's 2026 *State of AI in the Enterprise* found that only 28% of Australian organisations have moved even 40% of their AI pilots into production.⁸ Most remain in experimentation mode. Their conclusion was blunt: Australian organisations need to accelerate strategic, enterprise-level decisions to keep up with global peers.

We see the same pattern in our work. Organisations spent 2024–2025 carefully selecting use cases, running

In the first three weeks of February 2026, three frontier AI models shipped within sixteen days of each other. Each represents a step change in capability, not an incremental update. And critically, the infrastructure being built on top of them — agent coordination, governance tooling, CI/CD integration — is what makes this moment different from the last twelve months of hype.

This matters for government and regulated private sector organisations, because the gap between "running a pilot" and "operating at scale" just got wider — not because your pilots were wrong, but because what's possible has moved.

But the model releases are not the story. The story is what's being built around them.

GitHub released Agentic Workflows — AI agents that run inside CI/CD pipelines, authored in plain English, with sandboxed permissions and audit trails.⁴ Apple integrated agentic coding into Xcode.⁵ GitHub's Agent HQ now lets teams run Claude, Codex, and Copilot side-by-side, switching agents per task.⁶ Spotify merges over 650 AI-generated pull requests per month; their co-CEO stated that senior engineers haven't written code since December.⁷

This is not experimentation. This is production infrastructure.

pilots, and building business cases to productionise. That was the right approach at the time. But the technology didn't wait. The chatbot, document summariser, or classification tool you're preparing to scale was designed against models and tooling that are already a generation behind.

This isn't a criticism. It's a structural problem. The rate of technology change has outpaced the rate of organisational change management. And the response to that should not be to slow down, it should be to widen your field of view and increase your investment in the enablers we outlined in *Path to Value*.

Three shifts that matter for governance-conscious organisations

First, AI agents can now coordinate, and the orchestration patterns are maturing.

The February releases introduced multi-agent orchestration as a production-grade capability. A lead agent decomposes a problem into tasks, assigns them to specialist sub-agents in isolated environments, and synthesises results without human intervention for hours at a time.

Frameworks like GSD⁹ (13,000+ GitHub stars, used by engineers at Amazon, Google, and Shopify) illustrate the emerging pattern: the orchestrating agent analyses task dependencies and groups them into sequential "waves." Independent tasks within each wave execute in parallel via sub-agents, each with a fresh context window. Dependent waves run sequentially. The orchestrator never does the heavy lifting; it spawns, waits, integrates, and tracks. This prevents the quality degradation that occurs when AI sessions accumulate too much context, and it keeps the human's oversight clean.

This matters because multi-agent coordination was, until recently, theoretical. It is now how production teams ship software. Your AI strategy needs to account for it.

Second, governance is becoming a first-class design concern, not an afterthought.

GitHub's Agentic Workflows run with read-only permissions by default. Write operations require

explicit, auditable approval. Each workflow runs in an isolated sandbox. The AGENTS.md standard, now stewarded by the Linux Foundation, defines how agent behaviour, permissions, and constraints are specified per-repository.¹⁰

This matters for regulated organisations because the tooling is being designed around the same principles you already care about: least-privilege access, separation of duties, audit trails, and human approval for consequential actions. The per-repository permission model means different parts of your codebase can have different agent configurations; different models, different autonomy levels, different approval workflows.

The implication: governance infrastructure is available now. Not as a policy document, but as executable controls that can be embedded in your delivery pipelines.

Third, single-vendor AI strategies are already a constraint.

No single model is best at everything. The emerging pattern is to route different types of work to different models: one for complex reasoning, another for fast execution, a third for cost-sensitive volume work. Locking into a single provider is not just a commercial risk, it's a capability constraint that will compound as the gap between model providers' respective strengths widens.

What executives should focus on

We are not suggesting you abandon current initiatives. We are suggesting you recalibrate them against the current state of the technology, and act on three things.

1. Expand your technical horizon.

If your AI strategy was defined twelve months ago, it was defined in a different era. Your architects and technical leaders need dedicated time; not vendor presentations, but hands-on experimentation to understand what is now possible. This is the single highest-ROI investment available to you right now, and it directly addresses the technology literacy and AI skills we identified in Path to Value as critical enablers.

2. Invest in governance infrastructure, not just governance policy.

As we argued in November, organisations serious about AI must prioritise ICT investment. That argument now extends specifically to agent governance tooling: agent definition files that encode your standards and

boundaries, permission models that enforce least-privilege access, CI/CD pipelines that validate agent output before it reaches production, and audit trails that capture not just what an agent did but why. These are not things to evaluate in 2027. They are available now.

3. Increase your change management investment.

This is the counterintuitive part. When technology moves this fast, the natural organisational response is to pause and wait for things to settle. That instinct is wrong. The organisations that will be best positioned in twelve months are the ones building the internal capability (technical, procedural, and cultural) to absorb rapid change as a continuous process.

As we noted in Path to Value, unlike technology investments, there are rate-limits on how fast people can change, so early and consistent investment is required. Organisations must not put this off. The cost of delay is falling further behind a frontier that is accelerating.

The developer's role has already shifted from writing code to orchestrating agents. The same shift is coming for every knowledge-intensive function. The question is not whether to engage, but whether you engage from a position of understanding or a position of catch-up.

References

¹ MXA Consulting, *Path to Value With AI*, November 2025.

² Anthropic, "Introducing Claude Opus 4.6," February 5, 2026. <https://www.anthropic.com/news/claude-opus-4-6>

³ Nicholas Carlini, "Building a C compiler with a team of parallel Claudes," Anthropic Engineering Blog, February 2026. <https://www.anthropic.com/engineering/building-c-compiler>

⁴ GitHub Blog, "Automate repository tasks with GitHub Agentic Workflows," February 13, 2026. <https://github.blog/changelog/2025-02-04-github-agentic-workflows-technical-preview/>

⁵ Apple Newsroom, "Xcode 26.3 Release Candidate," February 24, 2026. <https://www.apple.com/newsroom>

⁶ GitHub Blog, "GitHub Agent HQ — Public Preview," February 4, 2026. <https://github.blog>

⁷ Gustav Söderström, Spotify co-CEO, public remarks, February 2026. Reported via multiple sources including The Register and InfoQ.

⁸ Deloitte, *State of AI in the Enterprise 2026: The Untapped Edge*. <https://www.deloitte.com/au/en/issues/generative-ai/state-of-ai-in-enterprise.html>

⁹ GSD (Get-Shit-Done), open-source spec-driven development framework. <https://github.com/gsd-build/get-shit-done>

¹⁰ AGENTS.md specification, Linux Foundation stewardship. Per GitHub Agentic Workflows documentation.

MXA Consulting is an AI-native management and technology consultancy serving Australian government and regulated private sector organisations. For a confidential discussion about how these developments affect your organisation's technology strategy, contact hello@mx.com.au.