

AI's breakthrough moment is being held back by fragmentation



AI is no longer confined to pilots. Among high-performing organisations, around half are already using it with the intent to reshape how their businesses operate and most are redesigning workflows.¹



At the employee level, AI is changing the pace of work itself, with enterprise users recovering 40–60 minutes a day by integrating it into everyday tasks.²



Organisations are reporting revenue uplift, stronger customer experiences, and faster product development cycles as AI moves closer to the core of the enterprise.²

Yet as AI workloads become heavier, more real-time, and increasingly autonomous, they are still being deployed across fragmented compute, storage, networking, and security environments. That fragmentation is now the limiting factor — constraining scale, increasing risk, and slowing the very transformation AI is meant to deliver.

The rise of Agentic AI

A few years ago, generative AI changed content creation. Agentic AI is changing how operations run. AI systems are beginning to act, coordinate, and execute with limited autonomy.




This shift is already taking shape.



62% of organisations are experimenting with AI agents¹



Over 1 million business customers now use OpenAI's tools as part of daily operations¹



ChatGPT enterprise message volume grew 8 times and API reasoning token consumption increased 320 times year-over-year²

The impact extends well beyond IT. Agentic AI is influencing **national security, healthcare, finance, supply chains, and corporate automation**, while accelerating progress in robotics, edge intelligence, and next-generation computing.³

But autonomy raises the bar. Agentic systems rely on real-time data, consistent connectivity, and full accountability. Small inconsistencies in networks or security can now create outsized risk. **As AI systems begin making decisions independently, enterprises need control and visibility built into the foundation.**

AI success now depends on the foundation beneath it

Breakthroughs at the model layer are not slowing down. New AI models will continue to arrive, faster and in greater variety. But as organisations adopt multiple models across teams, regions, and use cases, competitive advantage is shifting to something more durable: the foundation that allows AI to run consistently, securely, and at scale.

This becomes sharper as we enter 2026, often described as the **Year of the Defender⁴** when AI-driven security begins to tip the balance back toward enterprises. The most serious emerging threat is **moving beyond overt breaches to prompt injection⁵**, where corrupted inputs quietly shape untrustworthy models over time. Fragmented infrastructure makes this difficult to detect and even harder to contain. In a world where models will keep changing, a unified foundation is what protects data integrity, preserves trust, and allows intelligence to scale without losing control.



Singtel AI-Ready Infrastructure replaces fragmentation with a pre-integrated foundation.

Singtel AI-Ready Infrastructure is designed for enterprises running AI in production. By integrating performance, connectivity, and security into one foundation, it allows AI workloads to scale reliably today while remaining adaptable to what comes next.

Pre-validated AI architecture



Purpose-built AI PODs with integrated compute, networking and storage.



Plug-and-play stacks that cut deployment time and reduce integration risk.

High-performance compute & scalable storage



Supports full AI lifecycle: training, fine-tuning, RAG and inference.



Flexible modular options or turnkey integrated stacks.

Ultra-low latency, AI-ready networks



Critical low-latency, high-bandwidth connectivity for demanding AI tasks.



Predictable performance for real-time and autonomous systems.

Security engineered into every layer



AI access security to prevent leakage, tampering and unauthorised access.



Air-gapped environments for sensitive and regulated workloads.



SIM-based identity and SASE integration for end-to-end protection.

Centralised governance and observability with CUBE



Unified control across compute, storage, networking and AI runtime.



Real-time insights, orchestration and multi-vendor management.

Operational simplicity



Deployment, optimisation and ongoing management via Singtel Managed Services.



Embedded GenAI assistance for performance, security and digital-experience optimisation.

Ready for what AI becomes next

AI is evolving faster than the infrastructure decisions made to support it. Enterprises need a foundation that is designed for change, one that can absorb new models, new architectures, and new risk profiles without constant reinvention. A unified, resilient infrastructure creates the conditions for AI to scale responsibly, with performance, security, and governance moving in step.

Singtel AI-Ready Infrastructure provides that foundation, supporting AI workloads in production today while remaining resilient enough to endure what lies ahead.

Accelerate your path to enterprise-scale AI.

Contact us

References

- McKinsey, The state of AI in 2025: Agents, innovation, and transformation, 2025
- OpenAI, The state of enterprise AI, 2025
- Forbes, From Generative To Agentic: The New Era Of AI Autonomy In 2026, 2025
- Thomson Reuters, Safeguarding agentic AI: Why autonomy demands governance and security, 2025
- IBM, What is a prompt injection attack?, Year: N/A