ADVISORY BRIEF

Your Staff Are Already Using AI on Corporate Data. This Is What Leaves Your Network.

The Australian Government banned DeepSeek for a reason. But the risk isn't one app. It's every AI tool your staff are using without your knowledge, and every prompt that sends your data somewhere you can't control.

March 2026 Guruswami™ Advisory guruswami.com

ADVISORY BRIEF

Your Staff Are Already Using AI on Corporate Data. This Is What Leaves Your Network.

March 2026 · For CISO, CIO, Board, CEO

EXECUTIVE SUMMARY

The Australian Government banned DeepSeek for a reason. But the risk isn't one app. It's every AI tool your staff are using without your knowledge, and every prompt that sends your data somewhere you can't control.

CONTENTS

Guruswami Advisory

Your Staff Are Already Using AI on Corporate Data. This Is What Leaves Your Network.

Advisory Brief · March 2026 · For CISO, CIO, Board, CEO

Risk briefing - what is leaving your network and why

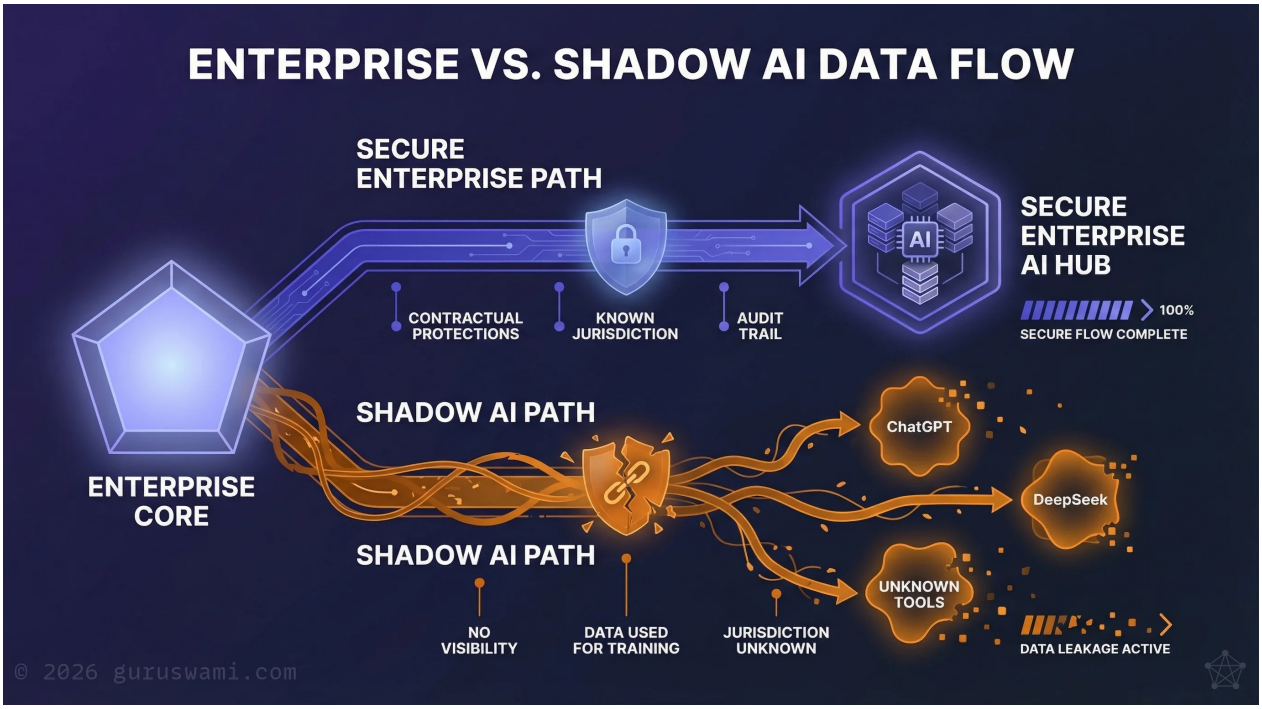

Exhibit 1 Where your data goes when staff use AI - enterprise vs shadow AI data flow

Source: Guruswami Advisory

> **1 in 4**
>
> white-collar workers use generative AI at work without their employer's knowledge.[1] They are not being malicious. They are solving real problems while creating exposures your security team cannot see. **The data is already leaving.**
>
> Jobs and Skills Australia, 2025

WHAT THIS BRIEF COVERS

1. **Every prompt is a data transfer.** Your staff are sending corporate data to infrastructure you do not control, in jurisdictions you have not assessed.
2. **The problem is not one app or one country.** DeepSeek is the headline. The risk is every AI tool your organisation has not evaluated.
3. **Inference streams can be intercepted and modified.** Data leakage is the risk you know. Decision manipulation is the one you have not considered.
4. **AI dependency is a business continuity risk.** When AI goes offline, the manual process no longer exists. The organisation does not slow down. It stops.
5. **The compliance position is clear.** Existing obligations already apply. The regulatory gap is widening with every unmanaged prompt.

Every Prompt Is a Data Transfer

When an employee pastes a contract, a financial model, customer records, or a strategy document into an AI tool, that data leaves your network. It is processed on infrastructure you do not control, in a jurisdiction you may not have considered, by an organisation whose interests are not yours.

It is happening now, across Australian organisations, every working day.

The Australian Government understood this when it issued PSPF Direction 001-2025[2], banning DeepSeek from all government systems and devices. The stated reason: *"extensive collection of data and exposure of that data to extrajudicial directions from a foreign government that conflict with Australian law."*

That directive applied to government. The risk applies to everyone.

DeepSeek's Own Privacy Policy

DeepSeek's privacy policy[3] is worth reading. It states plainly that all user data is stored on servers in the People's Republic of China. It states that data is used to train and improve their models.

Under China's National Intelligence Law (2017)[4], Chinese organisations are legally required to cooperate with state intelligence agencies. There is no opt-out. There is no appeal mechanism under Australian law.

DeepSeek is cheaper and in some benchmarks faster than Western alternatives. That is exactly why your staff may already be using it. The attraction is the capability. The cost is your data.

The Problem Is Not One App. Or One Country.

The Australian Government's ban targeted DeepSeek specifically. But the underlying risk is broader than a single application or a single jurisdiction.

Foreign adversary models are the most obvious risk. DeepSeek's data sits in China under laws that compel cooperation with intelligence services. But any offshore AI service creates a version of the same exposure: data leaves your network, is processed on infrastructure you do not control, and may be retained or used in ways you cannot verify.

Western cloud AI is not automatically safe. Microsoft, Google, and OpenAI offer enterprise agreements with stronger protections than consumer tools. Some offer "sovereign" or "on-shore" processing options for Australian customers. But "sovereign" is frequently a checkbox in a configuration panel, not an audited technical reality. It is an optional setting that may get overlooked, assumed, or never enabled in the first place. Processing paths transit international nodes. Shared infrastructure means your data sits alongside other customers' data. And these platforms are high-value targets. When a major cloud AI provider is breached[5], every customer's data is exposed at once.

Some cloud offerings, especially from Microsoft and Google, do have provable sovereign claims. Their technology can be used safely. But configuring them to do so is not trivial. Auditing them is complex. Trusting them requires both verification and faith. Their features are often robust, but they may not be as capable as bleeding-edge tools from Anthropic, Meta, or Grok. Dozens of innovative AI startups emerging from China, France, Russia, the US, and Australia are offering cheaper, smarter alternatives every month. One in four of your staff[1] may be using these tools, at home or on their personal phones, on corporate data.

Many services explicitly state that user inputs may be used to train future models. Every prompt feeds a model that may surface patterns derived from your information to anyone who queries it. Some enterprise agreements exclude training. Many consumer-tier tools do not. Most organisations have not done the work to understand which of their staff are using which tools under which terms.

Shadow AI Is Already Inside Your Organisation

It is already happening.

> *Every prompt is a data transfer.*

The ASD Annual Cyber Threat Report 2024-25[6] documented an 11% increase in cyber incidents and a 111% rise in attacks targeting critical infrastructure. The ASD Cyber Security Priorities for Boards 2025-26[7] specifically expects boards to have visibility of shared responsibilities between their organisation and service providers.

Against that backdrop, Jobs and Skills Australia found[1] that between 21% and 27% of white-collar workers are using generative AI at work without their employer's knowledge. The data leaving through those prompts includes intellectual property, source code, and regulated personal information.

Most of this usage is not malicious. It is the opposite. Your most capable, most innovative staff are the ones most likely to use AI tools to solve problems faster. They are summarising documents, drafting emails, analysing data, writing code. They are doing exactly what AI is good at.

They are also doing it on their phones, on their home computers, on free-tier consumer tools, and on whichever new model is getting attention that week. They are bypassing your network controls not out of negligence but because the approved tools are too slow, too restricted, or do not exist. They are solving real business problems while creating compliance exposures your security team cannot see.

This is the same shadow IT pattern that every CISO has dealt with for two decades, accelerated by AI tools that are free, powerful, and one browser tab away. The difference is that shadow IT historically meant an unapproved SaaS subscription. Shadow AI means your corporate data has left your network and is sitting on someone else's servers. Possibly being used to train their next model.

By the time you discover it, the data is already gone. You cannot recall a prompt. You cannot delete what a model has already ingested. The exposure is permanent.

What ASIC Found

ASIC's REP 798[8] reviewed AI governance across 23 financial services firms and found exactly this gap. Nearly half lacked policies addressing fairness and bias in AI. Fewer than half had policies about disclosing AI use to consumers. None had implemented AI-specific contestability arrangements.

These are regulated firms with compliance obligations. The governance gap in unregulated sectors is likely wider.

ASIC's finding was blunt: governance arrangements are lagging behind AI adoption, and that gap is widening as adoption accelerates. The firms with the widest gap are those adopting fastest with the least oversight.

The Risks We Have Not Encountered Before

Data leakage is the risk organisations are starting to understand. There are others that most have not yet considered.

When AI inference happens outside your network, the data travels to an external service and a result travels back. That is a stream of communication, and like any stream of communication, it can be intercepted and it can be modified.

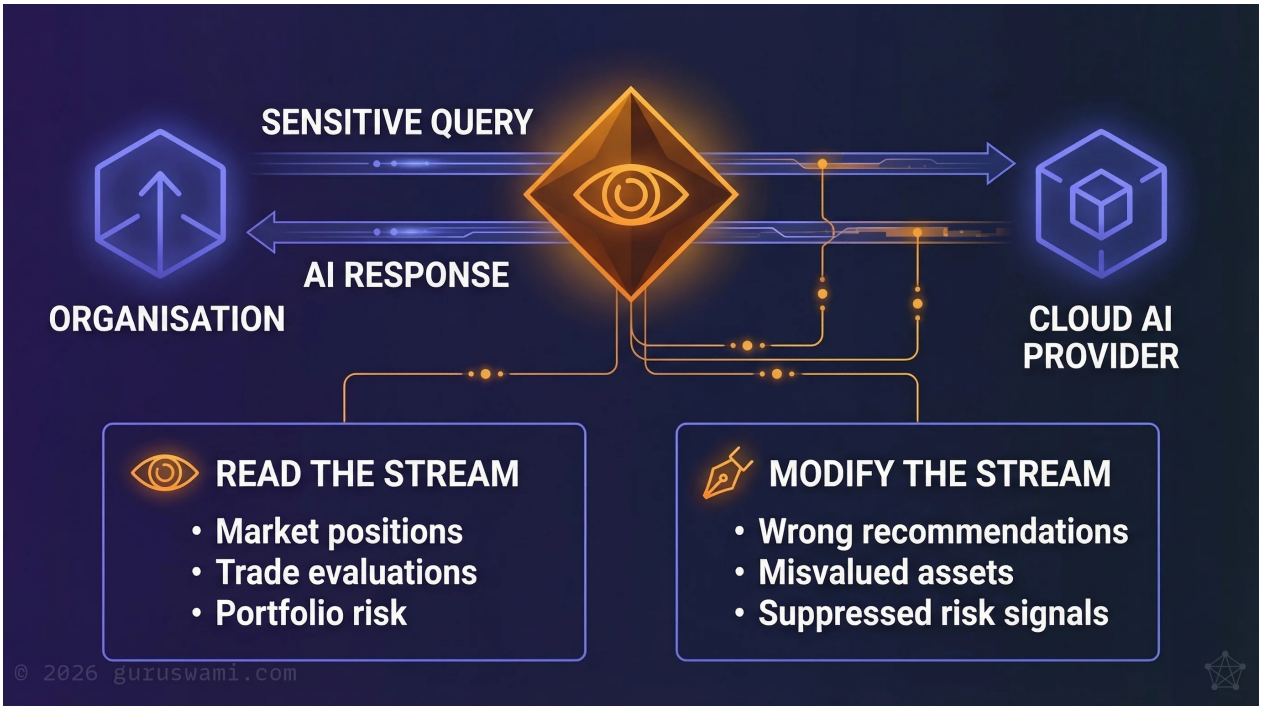

Exhibit 2 Inference interception attack: attacker reading and modifying AI query and response streams in transit

Source: Guruswami Advisory

Consider a stock broking firm using a cloud AI service to analyse market positions, evaluate trades, or model portfolio risk. An attacker who can tap that inference stream has real-time intelligence on what the firm is looking to buy, why, and how much it expects to move the market. That is insider trading delivered as a network exploit.

Now consider what happens if the attacker does not just read the stream but alters it. A modified inference result could recommend the wrong position, misvalue an asset, or suppress a risk signal. The firm acts on what it believes is its own AI's analysis. The manipulation is invisible because the result looks normal. It is just wrong.

This is not limited to financial services. A law firm's AI summarising case precedents. A defence contractor's model assessing supply chain risk. A health insurer's system triaging claims. In every case, the inference stream carries the organisation's most sensitive questions and receives answers it will act on. Intercepting that stream gives an attacker visibility into the organisation's thinking. Modifying it gives them influence over the organisation's decisions.

These are not theoretical attacks. They are extensions of techniques that already exist for other communication channels. The difference is that organisations have spent decades securing email, voice, and data in transit. Almost none have considered that their AI inference traffic is an equally valuable, and equally vulnerable, communication channel.

For the most sensitive workloads, the only way to eliminate this risk is to keep the entire inference pipeline on infrastructure you control. If the data never leaves your network, there is no stream to intercept.

The Dependency You Are Not Planning For

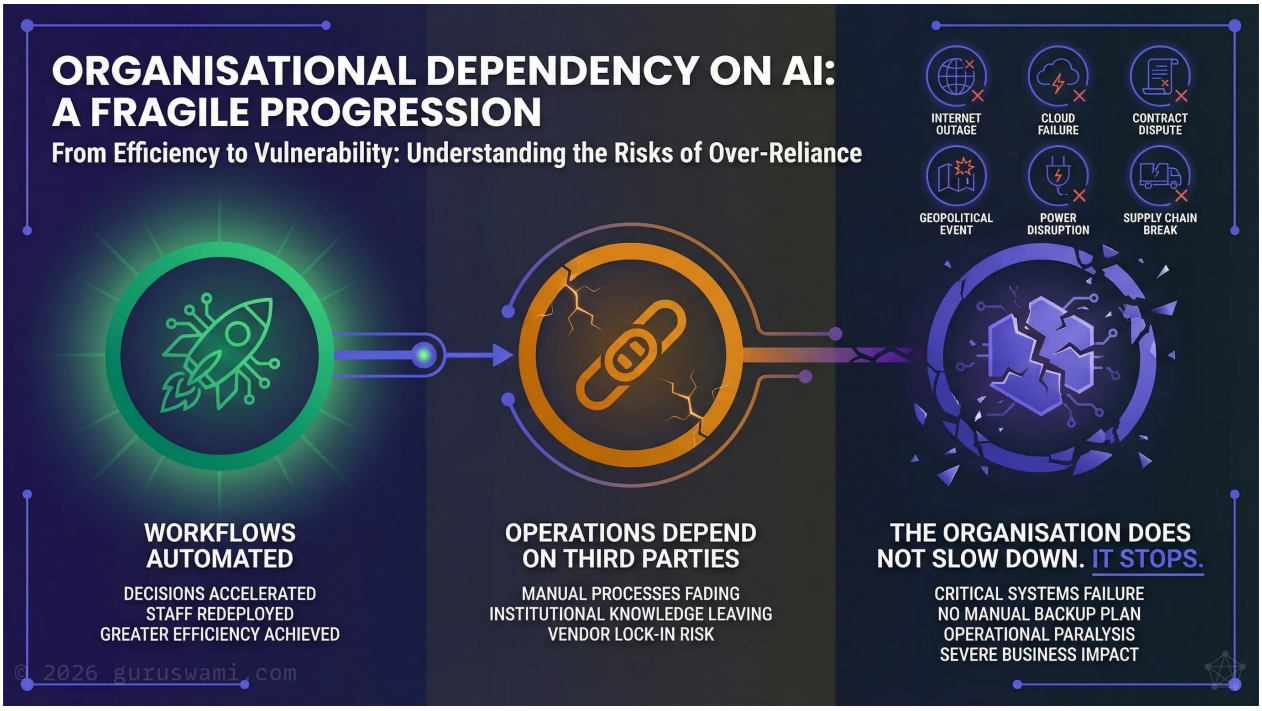

Exhibit 3 AI dependency risk: from adoption success through embedded dependency to disruption, the organisation does not slow down - it stops

Source: Guruswami Advisory

Organisations that adopt AI successfully will become dependent on it. That is the point. AI will automate workflows, accelerate decisions, and replace manual processes that were slower but self-contained. Staff who previously performed those functions will move on or be redeployed. The institutional knowledge of how things worked before AI will leave with them.

That is the natural outcome of doing it well. The efficiency gains are real, and no organisation will maintain parallel manual processes just in case.

But it creates a dependency that most organisations are not planning for. When AI inference is running on a cloud provider's infrastructure, your operations depend on an internet connection, a third-party service, and a commercial relationship, all of which can be disrupted.

An internet outage. A cloud provider's service failure. A contract dispute. A geopolitical event that restricts access to foreign infrastructure. A conflict that disrupts undersea cables or satellite links. A sustained power disruption. Any of these can take your inference capability offline.

When that happens, the workflows that AI automated do not revert to manual. The manual process no longer exists. The people who understood it are gone. The documentation, if it ever existed, is out of date. The organisation does not slow down. It stops.

It is a business continuity problem that AI creates. And almost no organisation is building it into their disaster recovery planning.

The question is straightforward: if your cloud AI provider went offline tomorrow and stayed offline for a week, which of your operations would stop? Do you have an alternative inference capability, on-premises, with a different provider, or both, that could keep those operations running? Do you have staff who could operate without AI assistance at all?

For most organisations today, the honest answer is that they have not thought about it. The time to think about it is before the dependency is embedded, not after.

The Compliance Position Is Clear

The Australian Government's position is documented across multiple frameworks:

Exhibit 4 Australian AI Compliance Framework: 5 regulatory sources including PSPF, APRA CPS 230, DTA AI Policy, DTA Impact Assessments, and ASIC REP 798 with status and deadlines

Source: Guruswami Advisory

- **PSPF Direction 001-2025**[2]: DeepSeek specifically banned from government systems. The precedent is set for foreign AI services that expose data to extrajudicial foreign government access.

- **APRA CPS 230**[9]**:** Operational risk management obligations effective since 1 July 2025. AI vendor oversight is explicitly in scope for APRA-regulated entities.
- **DTA AI Policy**[10]**:** Mandatory AI Impact Assessments required for all government AI deployments by 15 December 2026.
- **ASD Cyber Security Priorities for Boards 2025-26**[7]**:** Boards are expected to have visibility of shared responsibilities between their organisation and service providers, including cloud and AI vendors.
- **ASIC REP 798**[8]**:** ASIC has stated it will take enforcement action where AI use results in breaches of existing consumer protection obligations. The regulatory framework is technology-neutral. Existing obligations apply now.

The direction of travel is unmistakable. Organisations that have not assessed their AI exposure are building a compliance liability that grows with every unmanaged prompt.

These protections are necessary. They stem the immediate damage from unmanaged AI use. But compliance is not the same as competence. Organisations that treat regulation as the ceiling, that restrict AI use without building the capability to use it well, will find themselves globally uncompetitive when peers in less regulated jurisdictions have been learning and innovating for years.

> *The opportunity is for those who comply and learn at the same time.*

Governance Response

The response is not to ban AI. That approach fails because staff will use it anyway, and because AI capability is a genuine competitive advantage when governed properly.

The response is to know what is happening in your organisation and make deliberate decisions about it.

Immediate steps:

1. **Find out what your staff are using.** Most organisations do not have visibility of AI tool usage. Treat AI service endpoints like any other high-risk SaaS in your CASB, proxy, or firewall rules. Endpoint management and an honest conversation with your teams will surface the rest. Assume it is more than you think.
2. **Establish an acceptable-use policy.** Define which AI tools are approved, what data can and cannot be entered, and what the consequences are. Make it specific enough to be useful and short enough to be read.

3. **Assess your data exposure.** For any cloud AI service in use, read the terms. Understand where the data goes, whether it is used for training, and what jurisdiction governs it. If you cannot get clear answers, that is your answer.
4. **Consider local inference for sensitive workloads.** Running AI models on your own infrastructure means the data never leaves your network. The technology exists. The cost is manageable. For regulated, sensitive, or competitively valuable data, it may be the only responsible option.
5. **Get an independent assessment.** If you are unsure of your exposure, bring in someone who can assess it objectively. Not a vendor with a product to sell. An independent practitioner who will tell you what is actually happening.

If you have not assessed your AI exposure, start now. The data does not wait for a governance review.

If you are a Board member: Ask your CISO to report, within 30 days, which AI tools staff are actually using and what data categories are leaving your network. Ensure your disaster recovery plan addresses the scenario where your primary AI provider goes offline for a week.

If you are a CISO: Instrument your network for AI service traffic now. Establish an acceptable-use policy that names approved tools, prohibited data categories, and consequences, then communicate it before enforcement begins.

If you are a CIO: Evaluate local inference for your most sensitive workloads. For any cloud AI service already in use, verify the processing path, data residency, and training-exclusion terms yourself. Do not rely on vendor assurances.

KEY TAKEAWAYS

- Every prompt is a data transfer. When staff paste corporate data into AI tools, it leaves your network, processed in jurisdictions you may not have considered, under terms you may not have read.
- Shadow AI is not malicious. Your most capable staff are the ones most likely using unapproved tools. They are solving real problems while creating exposures your security team cannot see.
- "Sovereign AI" is frequently a sales claim, not an audited technical reality. Unless you have verified the entire processing path, you are trusting a vendor's word.
- AI dependency is a business continuity risk. When AI-automated workflows fail, the manual process no longer exists and the people who understood it are gone.
- The response is not to ban AI. It is to know what is happening and make deliberate decisions. Find out what your staff are using, assess your exposure, and consider local inference for sensitive workloads.

Guruswami Advisory assesses AI governance, data exposure, and shadow AI risk for Australian government and enterprise. Every engagement is led personally by the Principal Advisor.

References

1. https://www.jobsandskills.gov.au/publications/generative-ai-capacity-study-report
2. https://www.protectivesecurity.gov.au/publications-library/direction-001-2025-deepseek-products-applications-and-web-services
3. https://cdn.deepseek.com/policies/en-US/deepseek-privacy-policy.html
4. https://www.chinalawtranslate.com/en/national-intelligence-law-of-the-p-r-c-2017/
5. https://www.wiz.io/blog/wiz-research-uncovers-exposed-deepseek-database-leak
6. https://www.cyber.gov.au/about-us/view-all-content/reports-and-statistics/annual-cyber-threat-report-2024-2025
7. https://www.cyber.gov.au/business-government/protecting-business-leaders/cyber-security-for-business-leaders/cyber-security-priorities-for-boards-of-directors-2025-26
8. https://www.asic.gov.au/regulatory-resources/find-a-document/reports/rep-798-beware-the-gap-governance-arrangements-in-the-face-of-ai-innovation/
9. https://www.apra.gov.au/operational-risk-management
10. https://www.digital.gov.au/ai/ai-in-government-policy

About Guruswami Advisory

Independent AI security and strategy advisory for Australian boards, leadership teams, and regulated organisations. No vendor ties. No platform allegiance. Every recommendation tested on our own infrastructure.

Paul Nevin, Principal Advisor. 28 years in cybersecurity and cyber-intelligence. Six years of applied AI research. Every engagement led personally.

Contact

info@guruswami.com
guruswami.com
linkedin.com/in/paul-nevin

Guruswami™ Pty Ltd | ABN 11 695 354 020 | Canberra, ACT, Australia