FUSION COMPUTING FIELD GUIDE

Shadow AI in Canadian SMBs: How to Find It Before Your Regulator Does

Mike Pearlstein, CISSP, MSc Computer Science (AI)

August 12, 2026

WHAT'S INSIDE

fusioncomputing.ca · 416-566-2845 · CISSP-led managed IT, cybersecurity, and AI for Canadian SMBs · Toronto · Hamilton · Metro Vancouver

TRUSTED BY **Toronto law firms** **Hamilton manufacturers** **Vancouver clinics** **GTA accounting firms** **Ontario non-profits** **British Columbia professional services**

Shadow AI is your staff using AI tools the business never approved. In practice that means a consumer ChatGPT, Claude or Gemini account, on a personal login, with company or client information pasted into it. Start with four places: Microsoft 365 identity and app-consent records, network records showing traffic to AI services, managed-browser extension reports, and expense records. None is complete on its own.

In Canada the stakes changed on May 6, 2026, when the federal Privacy Commissioner and the Quebec, British Columbia and Alberta regulators jointly published findings on OpenAI. On collection of personal information from public and licensed sources they found OpenAI could not rely on implied consent. Training on user interactions was assessed separately, under both the appropriate-purposes and consent issues. If your people pasted client information into a consumer account, that is the specific environment the regulators assessed.

Key takeaways

- Shadow AI is a discovery problem before it is a policy problem. A policy written over an unknown estate governs nothing.
- Start with four useful surfaces: Microsoft 365 and Entra, network or DNS egress, managed-browser reports, and expense records. Each catches a different subset, and coverage depends on identity, device, browser, network and logging configuration.
- In a small firm, expense records are a low-friction place to start because they can expose paid subscriptions without adding telemetry. A charge is only a lead.
- A finding may be an unmanaged personal account signed in on a managed device rather than anyone acting in bad faith.
- The May 2026 joint privacy finding on OpenAI's GPT-3.5 and GPT-4 practices is the reference point for anything staff pasted into those models.
- Turning on AI activity logging can copy the same confidential data you are trying to protect. Decide what may be logged first.

What shadow AI actually is, and how it differs from shadow IT

Shadow IT was software procurement happening outside procurement. Someone expensed a file-sharing tool, and the data went somewhere the business did not choose. Shadow AI has the same shape and a sharper edge, because the thing being handed over is not storage. It is content.

When a bookkeeper pastes a client trial balance into a chat window to get a plain-English summary, no file is uploaded and no software is installed. A browser-only AI session may add nothing to a software asset inventory. That is why an inventory-only audit can miss it even when browser, endpoint or network telemetry still exists.

Book a 30-minute shadow AI scoping call

> **The working definition Fusion Computing uses.** Shadow AI is any use of an AI system to process business information where the business has not chosen the vendor, the plan tier, or the data-handling terms. The employee's intent is irrelevant to that definition.

That last point matters more than it sounds. Intent does not change the control problem. A well-meaning employee can still put business data into an account the company did not approve. Do not assume misconduct; assess the evidence first. Treating discovery as an investigation into misconduct produces a workforce that hides the tools better, which is a worse position than the one you started in.

Why the Canadian picture sharpened on May 6, 2026

Most writing on this subject is American vendor-survey material. The Canadian position rests on something firmer, including the joint federal and provincial principles for generative AI.

On May 6, 2026, the Office of the Privacy Commissioner of Canada, the Commission d'accès à l'information du Québec, and the privacy commissioners of British Columbia and Alberta published the findings of a joint investigation into OpenAI, PIPEDA Findings #2026-002.

For GPT-3.5 and GPT-4, the regulators found the original collection of personal information from public and licensed sources could not rest on implied consent; under PIPEDA that issue was conditionally resolved after OpenAI committed to further mitigation. Training on user interactions was assessed separately under both the appropriate-purposes and consent issues, and the conclusions there turn on OpenAI's mitigations. Cite the specific issue rather than a single overall outcome.

On user data specifically, the Offices weighed the benefits of training on user interactions against the residual privacy risk in light of OpenAI's mitigations. OpenAI committed to filtering tools that mask personal information, formal retention policies, improved transparency, and quarterly compliance reports.

> *Read that finding as an operator, not a lawyer. It does not say your staff broke the law. It says that for GPT-3.5 and GPT-4, your own regulators found the consent basis for training on user interactions was invalid. For a shadow-AI review, check the vendor terms, model and settings that applied when the data was entered rather than assuming every consumer account or current model handled data the same way.*

The practical consequence is narrow and useful. If client information was pasted into a consumer account, you have a dated official finding about those models to inform the assessment. For any other tool or model, check the terms, settings and data practices that applied at the time. That is a better starting point for a privacy assessment than a vendor's marketing page.

Sector context helps size the exposure. Statistics Canada reported that 19.2% of Canadian businesses used AI to produce goods or deliver services in the second quarter of 2026. Rates vary widely by sector: information and cultural industries at 42.3%, finance and insurance at 40.4%, and professional, scientific and technical services at 32.4%.

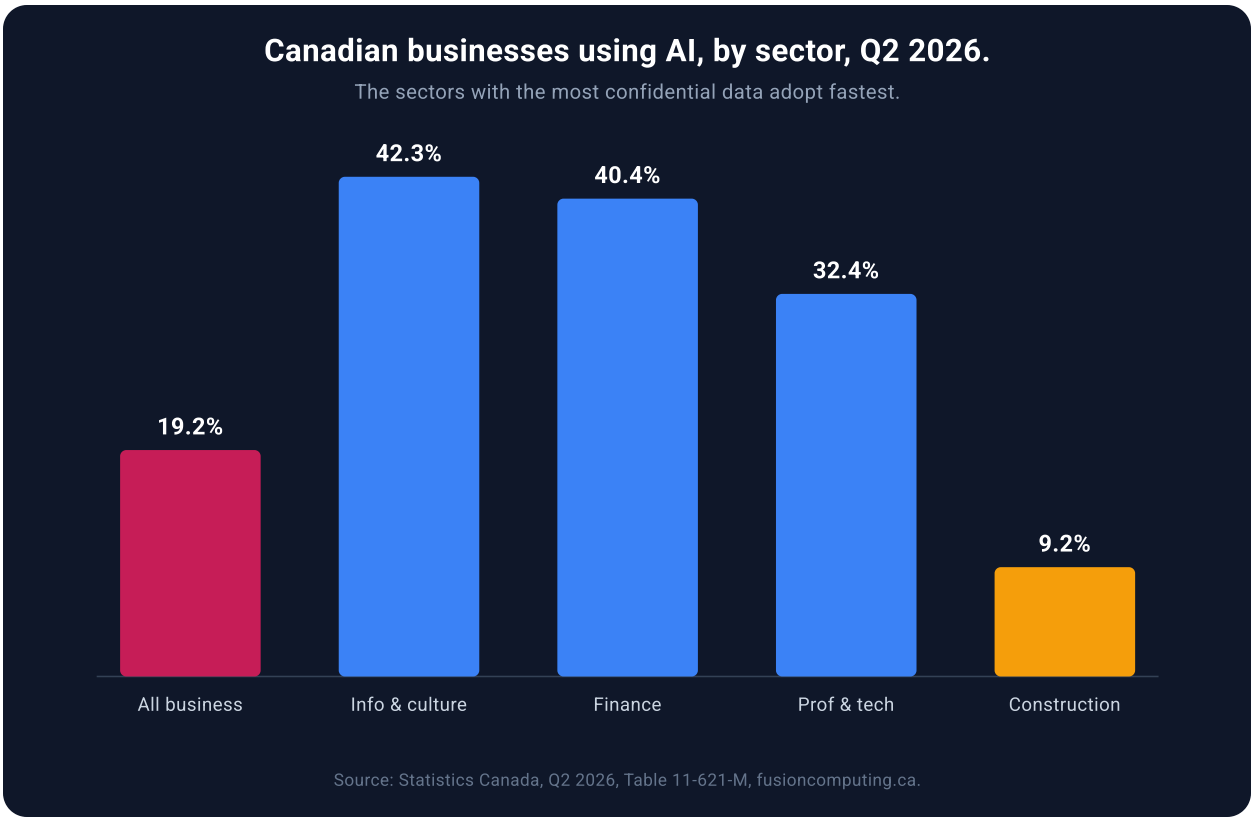

Finance and insurance (40.4%) and professional, scientific and technical services (32.4%) sit well above the 19.2% all-business rate. Source: Statistics Canada, Table 11-621-M.

Those are sector-level AI adoption figures, not shadow AI figures. Statistics Canada's Q2 2026 survey measures business AI use, not unsanctioned use. We did not find a comparable national Canadian dataset for shadow AI, so we are not using vendor-survey percentages here.

Where to look: the four detection surfaces

A shadow AI review is a discovery exercise with a known shape. Each of these 4 surfaces catches a different subset of activity, and coverage depends on identity, device, browser, network and logging configuration.

Surface	What it reveals	What it misses	Effort
Microsoft 365 and Entra	Work-identity sign-ins to registered AI resources and app-consent grants. This does not show prompt content or personal-account activity.	Anything done on a personal account in a private browser profile	Low
Network and DNS egress	Traffic to known AI domains from observed devices. It does not identify the account, prompt or data sent.	Off-network use that bypasses the telemetry source, including some home or mobile-data traffic	Medium
Browser extensions and profiles	Installed AI-related extensions and the sites or permissions they request on managed or enrolled browsers. It does not prove actual page access.	Web use with no extension installed	Medium
Expense claims and cards	Paid subscriptions. A charge proves purchase, not use, frequency or data exposure.	Free-tier use	Very low

Run them in that order, with one exception. In a small firm, Fusion Computing often starts with expenses because it is low-friction and exposes paid subscriptions without adding telemetry. Confirm account ownership and business dependency before treating a charge as evidence of use.

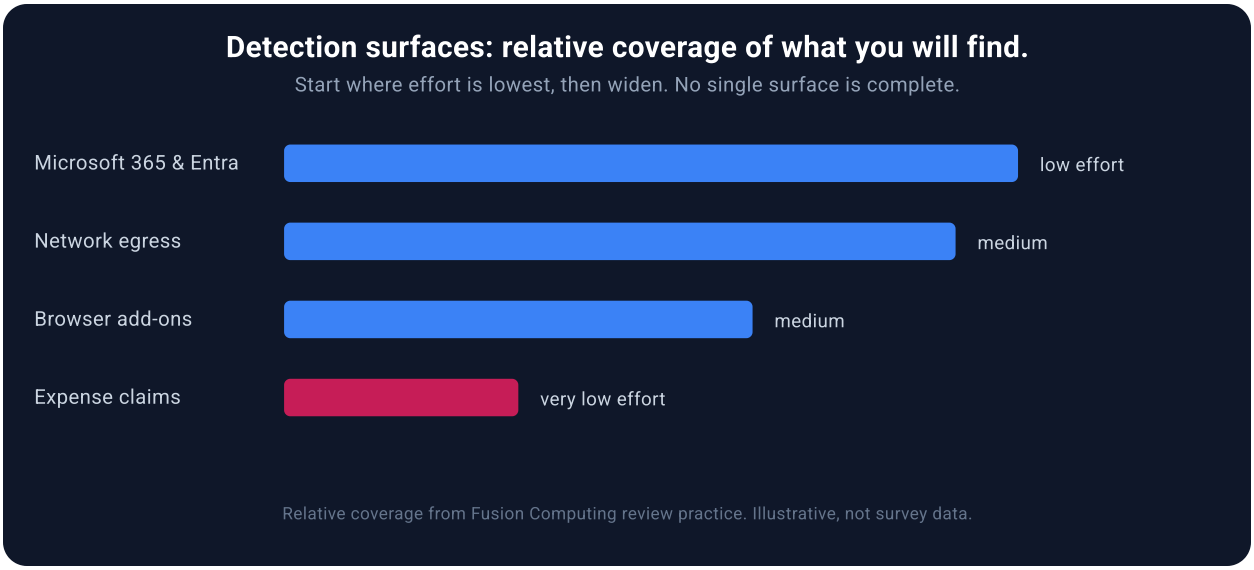

Coverage is illustrative and drawn from review practice. It is deliberately not presented as measured data.

Google Workspace tenants follow the same 4-surface logic with different menus. Admin console security reports replace Entra sign-in logs, and third-party app access control replaces OAuth consent review. The four surfaces do not change.

If you want this run for you rather than run internally, our shadow AI assessment does exactly that: book a 30-minute shadow AI scoping call with a CISSP-led team that has secured Canadian firms since 2012 and Fusion Computing will tell you which surfaces your tenant can actually report on before you commit to anything.

Book a 30-minute shadow AI scoping call

What you will actually find

The expectation going in is often a rogue employee. Assess the evidence before assuming that is what you have.

One pattern worth looking for specifically is an unmanaged personal account signed in on a managed device. The laptop is enrolled, patched and encrypted. The browser profile is personal. Work happens in a tool the business has no relationship with, on a device the business fully controls, and every endpoint control reports green.

> **[FIELD NOTE]** A first-person field observation. The first thing I check is the browser profile list, not the log. A managed device with two Chrome profiles, one signed in with a personal Gmail, tells me more in ten seconds than an afternoon of egress analysis. The device is compliant. The workflow is not.

The second pattern is consent creep. A user grants an AI browser extension or an OAuth app permission to read mail or files, once, 6 or 12 months ago. Nobody revoked it. The tenant is handing out standing access that no one remembers approving, which is why the app-consent review in Entra is worth more than its 5 minutes suggests.

Not sure which of the four surfaces your tenant can report on? Ask a senior engineer before you buy tooling for it.

Five controls to put in place after discovery

Discovery only helps if the findings change configuration. Start with these five controls.

Restrict user consent for applications at the tenant level, so a single click can no longer grant standing access to mail and files. Publish a short approved-tool list, so there is a sanctioned path. Move sanctioned use onto a business tier, where content is not used to train the vendor's models by default.

Then apply conditional access, so work identities reach approved services from managed devices. Add sensitivity labels and data-loss prevention on the categories that actually matter, rather than everywhere at once.

Fusion Computing configures and verifies those controls as part of the managed IT work it already does for Canadian firms. What no provider controls is the vendor's platform behaviour, which is why the plan tier and the contractual terms carry as much weight as the tenant settings.

A sanctioned path matters more than a prohibition. Blocking AI endpoints without offering an approved alternative moves the same activity onto phones and home machines, where you have no visibility at all. If you want the policy half of this, our guide on what belongs in an AI acceptable use policy covers the document itself.

The trap: your monitoring can leak what you are protecting

This is the part most shadow AI advice skips.

Once a firm decides to monitor AI activity, the instinct is to stream everything into the SIEM. Anthropic documents that its Cowork OpenTelemetry export includes the full text of user prompts, along with tool parameters, file paths and user email addresses, and that prompt content is included by default. Anthropic specifically tells organisations to configure filtering or redaction and to plan retention and access before routing those events downstream.

If you want the discovery run for you rather than run internally, book a 30-minute shadow AI scoping call and we will tell you what your tenant can report first.

> **Why Canadian firms bring this work to Fusion Computing**
>
> CISSP-led, a Microsoft Solutions Partner and a CompTIA Managed Services Trustmark holder, securing IT for Canadian SMBs across Toronto, Hamilton, and Metro Vancouver since 2012.

Read that as a general rule rather than a vendor quirk. Detailed AI activity logging tends to capture the content of the interaction, and the content is the sensitive part. A firm that turns on verbose logging to protect client confidentiality can end up with a second copy of that confidential material sitting in a monitoring platform with different retention and different access rules.

> **Decide what may be logged before you turn logging on.** Filtering, redaction, retention and SIEM access are design decisions, not follow-ups. Getting the order wrong creates the exposure the project was meant to close.

What a shadow AI review does not solve

Discovery tells you what is happening now. It does not tell you whether your firm is ready to adopt AI deliberately, which is a separate question our AI readiness assessment answers. It does not write your policy, and it is not 1 of the 5 controls below. It does not decide whether a given workflow is appropriate for your regulator, and in regulated practice that judgement stays with the professional.

It also has a shelf life. A tenant reviewed today looks different in six months, because AI tools and staff both change. Fusion Computing treats shadow AI discovery as a recurring control rather than a project, folded into the same cadence as patching and access reviews.

For firms in regulated verticals, the sector guides go further into the specific duties: AI for law firms, AI for accounting firms, and the ready-to-fill LSO and CPA policy templates.

Is it illegal in Canada for staff to use ChatGPT at work?

No. There is no Canadian law that prohibits an employee from using a general AI assistant, and framing it that way inside your firm will cost you the cooperation you need.

The legal exposure attaches to the information, not the tool. Canadian private-sector privacy law governs how your organisation collects, uses and discloses personal information. Putting personal, confidential, privileged or regulated information into an unapproved AI service can create privacy, contractual or professional-duty issues. Under PIPEDA a transfer to a processor is treated differently from a disclosure to an independent third party, so classify the data flow before deciding the legal consequence.

Which statute applies depends on where you operate and what the engagement is. PIPEDA is the federal default. British Columbia, Alberta and Quebec have their own private-sector regimes, and organisations subject to a substantially similar provincial law are generally exempt from PIPEDA for activity inside that province, while PIPEDA can still reach interprovincial and

international transactions. Fusion Computing confirms which regime applies before recommending controls, because the answer changes the obligations.

Professional obligations sit on top of that, and they are frequently the binding constraint. CPA Ontario published a regulatory standard on the responsible use of AI in professional practice on June 29, 2026. A lawyer's confidentiality duty, a CPA's obligations under a provincial code, and a clinic's health-privacy duties all survive whatever the vendor's terms say. No vendor setting retires a professional duty.

Do we have to report a breach if client data went into an AI tool?

Possibly, and it is a question to answer deliberately rather than assume away.

Under PIPEDA, an organisation must report a breach of security safeguards to the Privacy Commissioner and notify affected individuals where the breach creates a real risk of significant harm. Quebec's regime under the Commission d'accès à l'information requires reporting confidentiality incidents that present a risk of serious injury. Both tests turn on sensitivity and probability of misuse, not on whether a file was technically exfiltrated.

Whether pasting client information into a consumer AI account meets that threshold is fact-specific. It depends on what the information was, whose it was, what the vendor's terms permitted at the time, and what the vendor did with it. The May 6, 2026 joint finding is directly relevant evidence for the middle two of those questions, which is a large part of why it matters to an operator.

> **Document the assessment either way.** PIPEDA requires organisations to keep records of breaches of security safeguards, including ones judged not to meet the reporting threshold. A written determination that no notification was required is itself part of compliance, and it is far easier to write while the facts are fresh.

This is the point where a discovery exercise stops being an IT task. Fusion Computing runs the technical side and produces the evidence, then the firm's privacy lead or counsel makes the notification call. Keeping that line clear protects both sides.

Where to start this week

Ask your bookkeeper for any AI subscription expensed in the last 12 months. Then read the enterprise apps your users consented to in Entra.

If the answer is the second one, ask a senior engineer. Fusion Computing is CISSP-led. We have secured Canadian firms since 2012 and now serve Toronto, Hamilton and Metro Vancouver.

Book a 30-minute shadow AI scoping call

Fusion Computing helps Canadian businesses across Toronto and the GTA, Hamilton, and Metro Vancouver with managed IT, cybersecurity, and Microsoft 365.

Frequently Asked Questions

Q. What is shadow AI? –

Shadow AI is any use of an AI tool to process business information where the business has not chosen the vendor, the plan tier, or the data-handling terms. One example is a consumer ChatGPT, Claude or Gemini account on a personal login, with company or client information pasted in. The employee's intent is irrelevant to that definition.

Q. How do I find out if my staff are using ChatGPT at work? –

Check four places. Microsoft 365 and Entra sign-in logs and app consents show work identities reaching AI services. Network or DNS egress shows devices talking to AI endpoints. Browser extension inventories show assistants that read page contents. Expense claims show paid personal subscriptions. In a small firm, expense records are a low-friction place to start.

Found a personal account signed in on a managed device? Closing tenant-level app consent is the next step.

Q. Is it illegal in Canada for employees to use ChatGPT with company data? –

Not by default. There is no general Canadian ban on employees using ChatGPT. But putting personal, confidential, privileged or regulated information into an unapproved AI service can create privacy, contractual or professional-duty issues. Under PIPEDA a transfer to a processor is treated differently from a disclosure to an independent third party, so classify the data flow before deciding the legal consequence. Professional confidentiality duties sit on top and often bind harder.

Q. What did Canadian regulators decide about ChatGPT in May 2026? –

On May 6, 2026, the federal Privacy Commissioner and the Quebec, British Columbia and Alberta regulators published PIPEDA Findings #2026-002. The investigation covered ChatGPT as powered by GPT-3.5 and GPT-4, and excluded later model releases and other OpenAI services. On collection of personal information from public and licensed sources, they found OpenAI could not rely on implied consent; under PIPEDA that issue was conditionally resolved after OpenAI committed to further privacy-filtering and transparency measures. Training on user interactions was assessed separately under both the appropriate-purposes and consent issues, and the conclusions there turn on OpenAI's mitigations. OpenAI committed to filtering, retention policies and quarterly compliance reports.

Q. Can Microsoft 365 tell me which AI tools my staff use? –

Partly. Entra sign-in logs show work identities authenticating to AI services, and the enterprise applications view shows OAuth consents users have granted. What Microsoft 365 cannot see is a personal account used in a separate browser profile, which the tenant cannot see. Pair it with observed egress telemetry to see whether the device reached an AI service; that still does not identify the personal account or prompt content.

Q. Do we have to report a privacy breach if client data went into an AI tool? –

Possibly. Under PIPEDA an organisation must report a breach of security safeguards where there is a real risk of significant harm, and Quebec's regime uses a risk of serious injury test. Whether pasting client information into a consumer account meets the threshold is fact-specific. Record the assessment either way, including a decision not to notify.

Q. Is an AI policy enough, or do we need technical controls? –

A policy alone governs very little. Restrict tenant-level app consent so one click cannot grant standing access to mail and files, publish an approved-tool list, move sanctioned use onto a business tier, and apply conditional access. Fusion Computing configures and verifies those controls. The policy then describes a reality that already exists.

Q. What does a shadow AI review actually involve? –

A review works the four detection surfaces in order of effort, then reports what was found and which controls close it. It produces evidence a privacy lead can act on, not a verdict. Expect a tenant configuration review alongside it, because discovery only pays off when it changes settings.

Q. Will blocking AI tools just push the behaviour underground? –

It can. Blocking endpoints without offering a sanctioned alternative moves the same activity onto phones and home machines, where there is no visibility at all. An approved path on a business tier, where content is not used to train the vendor's models by default, does more than a prohibition.

Q. How is this different from an AI readiness assessment? –

A readiness assessment answers whether the firm is ready to adopt AI deliberately. A shadow AI review answers what is already happening and what has already left the building. They run in different directions, and the discovery question usually needs answering first, because readiness planning over an unknown estate is guesswork.

Q. Can AI activity monitoring create its own privacy problem? –

Yes, and it is routinely missed. Anthropic documents that its Cowork OpenTelemetry export includes the full text of user prompts, plus tool parameters, file paths and user email addresses, with prompt content included by default. Decide what may be logged, configure filtering or redaction, and set retention and access before enabling any AI activity export.

Q. How often should we repeat a shadow AI review? –

Treat it as a recurring control rather than a project. AI tools and staff both change, so a tenant reviewed today looks different within about 6 months. Fusion Computing folds shadow AI discovery into the same cadence as patching and access reviews for the firms it manages.

Put this guide to work

Fusion Computing is a CISSP-led Microsoft Solutions Partner and CompTIA Managed Services Trustmark holder, securing Canadian businesses since 2012.

- Managed IT support with a 1-hour priority response
- Cybersecurity, managed detection and response, and compliance reviews
- Microsoft 365 management and licensing

Book a 30-minute scoping call

fusioncomputing.ca/contact-us · 416-566-2845
Toronto · Hamilton · Metro Vancouver