Why Every Business Needs a Secure AI Workspace

A practical guide to using generative AI without putting your data at risk.

For business owners, CIOs, CISOs, and leadership teams

v2.0

Executive Summary

Generative AI is already embedded in day-to-day work. People use it to draft emails, summarize documents, write proposals, analyze spreadsheets, and accelerate research. Adoption is happening faster than most organizations can govern it. Surveys show most knowledge workers are already using GenAI at work, and many will use whatever tools help them even if the organization does not explicitly authorize it. [1] [2]

That creates a new, widely misunderstood risk. Many organizations assume that "turning off training" in an enterprise AI product makes it safe for business data. Turning off training is important, but it is not sufficient.

The core issue is not whether a public AI provider is "secure." The core issue is control.

The moment a user pastes, uploads, or types sensitive information into a public AI model, that information has left the organization's controlled environment and has been disclosed to a third party. At that point:

- Your data handling and retention are governed by the provider's policies, tooling, and legal obligations, not yours. [3] [4] [5]
- The provider may retain prompts and outputs for operational and safety monitoring, and may retain data longer when legally required. [4] [6] [7]
- Your organization loses the ability to prove, with evidence, what was and was not exposed through AI interactions, unless you have centralized governance and logging.

A simple rule of thumb helps leaders and end users understand the problem immediately:

If you wouldn't post the data publicly, don't paste it into public AI models. [3]

Most businesses feel stuck between two bad options:

- Allow any AI, which leads to uncontrolled usage, personal accounts, and unpredictable exposure of client data, regulated data, and intellectual property.
- Block AI, which often pushes usage underground. People find workarounds, including using personal devices, personal accounts, and screenshots. [2]

A Secure AI Workspace is often the best option for businesses.

A Secure AI Workspace provides one approved front door for AI at work, with guardrails that protect sensitive data before it reaches model providers, and with audit-ready evidence of

how AI is being used. It enables employees with the AI tools they want, while giving leadership, IT, security, and compliance the control and governance they need.

For many organizations, the best way to deploy and run a Secure AI Workspace is through a managed service provider or security partner, using a purpose-built platform such as Liminal. [8]

1. AI is already in your business, whether you planned for it or not

Generative AI tools have moved from experimentation to habit.

Employees use AI to:

- Draft emails, proposals, and client communications
- Summarize long documents and meeting notes
- Brainstorm, outline, and create first drafts
- Translate text or adjust tone
- Analyze and clean up data and spreadsheets

This started informally. Someone tried a tool, saved time, and kept using it. Coworkers saw the output and followed. Often, this happens long before leadership, IT, or security has a chance to set guardrails.

That is why the AI conversation is now in scope for every business. The question is no longer “Will we use AI?” The question is “Will AI usage be governed and defensible, or unmanaged and risky?”

2. The misunderstood risk: “training off” is not the control you actually need

Many vendors offer settings such as:

- “Do not train on my data”
- “Data is not used to improve models”
- “Enterprise plans have stronger privacy controls”

Those settings are important. They reduce one category of risk: using your inputs as training data.

But they do not solve the bigger problem: data exposure and loss of governance when sensitive information is submitted to a third party.

Even if training is disabled:

- Prompts and outputs can be retained in operational logs for abuse monitoring and service integrity, with defined retention windows and legal exceptions. [4]
- Providers may retain chats or sessions when required by law or to enforce usage policies. [6]
- Regulators and government guidance explicitly warn against entering sensitive data into publicly available generative AI tools due to privacy and data protection risks. [3] [7]

A Secure AI Workspace is not a “training off” setting. It is a control plane that governs what can be sent to AI, how it is protected, and what evidence exists afterward.

Important clarification: a Secure AI Workspace should still ensure model providers do not train on your organization’s prompts and outputs. Liminal is designed for enterprise use and avoids training on customer data by default when using its model providers.

3. Why using public AI tools directly inside a business is a bad idea for sensitive data, even on paid enterprise plans

This section is not an argument that businesses should never use ChatGPT, Claude, Gemini, or other public AI products. Many tasks can be done safely when the inputs are not sensitive.

This section is about what happens in real organizations.

In real organizations, sensitive data shows up constantly in normal work. It is embedded in:

- Client emails and client attachments
- Contracts, proposals, pricing, and deal terms
- HR and employee information
- Internal security information and credentials
- Financials, forecasts, and board materials
- PHI, PII, PCI, and regulated records
- Source code, internal architecture, and proprietary processes

Even if leadership says “only use enterprise accounts” or “only use AI for non-sensitive tasks,” users will still end up putting sensitive data into public AI tools, because:

- Many users already have personal accounts and personal habits.
- It’s faster to paste than to think through data classification every time.
- Users often don’t recognize that what they’re pasting is sensitive, especially in context.
- AI is frequently used to rewrite, summarize, or improve content that is inherently sensitive.
- People assume “the risk is minimal” or “everyone’s doing it.”

Here is the point to anchor on:

A public AI model is not a safe destination for sensitive data because submitting the data is a disclosure to a third party.

That is why government and privacy authorities recommend not entering personal or sensitive information into publicly available generative AI tools. [3] [7]

If you wouldn't post the data publicly, don't paste it into public AI models. [3]

4. Why "policy only" and "block AI" both fail in the real world

Most organizations try one of these first:

Option A: Write a policy and trust users to comply
Organizations publish acceptable use guidance and tell employees not to paste sensitive information into AI.

Policies are important and can help, but they do not solve the problem on their own. People are busy. They optimize for speed. Many will do what makes their job easier, especially if they believe the risk is small or unclear.

Option B: Block AI tools at the network level
Some organizations try to block access to popular AI sites.

Blocking is understandable, but it often creates two outcomes:

- AI usage becomes harder to see and harder to govern.
- Users find workarounds, including personal devices and personal accounts.

An all-to-common workaround when AI tools are blocked: a user takes a photo or screenshot of sensitive information and submits it to an AI tool on a personal phone, using a personal account.

At that point, the organization has:

- No policy enforcement in the moment the data is shared
- No centralized record of what was exposed
- No evidence to prove what did and did not leave the controlled environment

The goal is not to "ban AI" or "trust AI." The goal is to provide a safer, approved way to use AI that employees will actually use.

5. What a Secure AI Workspace is

A Secure AI Workspace is one front door for AI at work, with guardrails around it.

In practical terms, a Secure AI Workspace provides:

- A single place employees go to use AI for work tasks
- Access to multiple underlying AI models, so users can choose the right tool for the job
- Real-time policies that detect and protect sensitive data before anything is transmitted to model providers
- Centralized logging and reporting that provides audit-ready evidence of AI usage

This changes the posture from "AI is scattered across unknown sites and accounts" to "AI is enabled inside a governed environment."

A Secure AI Workspace does not exist to slow employees down. It exists to keep AI usage productive while making it controlled, defensible, and aligned with existing security and compliance programs.

6. The standard for success: enablement that users accept, plus governance leaders can defend

This is the hard part of AI security. If a solution is too restrictive or cumbersome, adoption falls, and users revert to personal tools. If a solution is too permissive, sensitive data leaks.

A Secure AI Workspace must deliver both:

User enablement

- Familiar workflow and simple UX
- Fast access, minimal friction
- Strong output quality
- The flexibility to use the best model for the task
- Shared capabilities that make it better than personal account

Governance and evidence

- Sensitive data detection and protection before the model call
- Consistent policy enforcement across models
- Role-based access controls and admin visibility
- Audit-ready evidence of what was entered, what was protected, and what was returned
- Reporting that supports security operations, compliance, and client trust

7. What this looks like in practice with Liminal

Liminal is designed to function as a Secure AI Workspace that enables multi-model GenAI while protecting sensitive data and providing governance and observability. [8]

At a high level, the approach is:

Detect and protect sensitive data before anything is transmitted to model providers and preserve evidence of what did and did not leave the environment through AI.

A few capabilities business leaders and security teams typically value:

Multi-model access in one governed workspace

Employees do not need separate subscriptions and separate accounts across different vendors. They can access leading models through one approved environment, with consistent policy and centralized visibility. [8]

No-training by default through the workspace

Turning off training is important, but it's only part of the overall control problem. A Secure AI Workspace should still ensure model providers do not train on your organization's prompts and outputs. Liminal is designed for enterprise use and avoids training on customer data by default when using its model providers.

Sensitive data guardrails before model submission

Liminal is built to detect common sensitive data types, including PHI, PII, PCI, and intellectual property, and apply policies before content is sent to model providers. [8]

Shared "Assistants" for repeatable workflows

A Secure AI Workspace becomes significantly more useful when teams can create and share "Assistants" designed for specific tasks. Instead of every user starting from scratch, approved "Assistants" can be reused across roles and departments. This helps drive adoption in the approved workspace and reduces the temptation to fall back to personal AI accounts for convenience. [9]

Audit-ready evidence and visibility

A major gap with unmanaged AI usage is the lack of evidence. A Secure AI Workspace should create a defensible record of AI interactions, including what was entered, what policy did to the content, and what responses were produced. This supports investigations, governance reporting, and client assurance.

Plain language summary

Organizations don't adopt a Secure AI Workspace because they want "another AI tool." They adopt it because they want AI outcomes without losing control of sensitive data, and they want to be able to prove how AI is being used.

8. Why work with an MSP (Managed Service Provider) or security partner

Most organizations do not want to build AI governance from scratch. They want an outcome:

Enable GenAI for users without exposing sensitive data and be able to prove it.

An MSP or security partner can deliver that outcome by:

- Assessing current AI usage and exposure paths
- Defining practical rules for what data must never be submitted to public AI models
- Deploying a Secure AI Workspace and connecting identity and access controls
- Configuring policies for regulated data and company-specific patterns
- Providing ongoing monitoring, reporting, and tuning as usage evolves
- Helping the business roll out training and adoption that actually sticks

This approach allows your organization to keep moving quickly, while ensuring AI usage stays governed as tools and behavior change.

9. A practical rollout plan

A Secure AI Workspace is easier to deploy when the goal is clear: enable safe usage quickly, without breaking workflows.

A practical approach looks like this:

Step 1: Set the non-negotiables
Define the categories of data that must not be transmitted to public AI models, period. This typically includes PHI, PCI, certain PII, privileged legal content, and core IP.

Step 2: Choose the approved front door
Deploy a Secure AI Workspace users can access easily, that supports the models they actually want to use.

Step 3: Start with guardrails, not complexity
Turn on detection and protection policies that prevent obvious risk patterns. Keep the user experience smooth, then tune from real usage.

Step 4: Enable shared, repeatable workflows
Create and share approved "Assistants" and templates for the common tasks users are already doing. The more useful the approved workspace is, the less users rely on personal tools.

Step 5: Use evidence to improve governance
Review high-level reporting and patterns. Tighten policies where needed. Provide leadership with a clear view of progress and control.

10. Conclusion

Generative AI is already part of business operations. The productivity gains are real, and employees will continue to use these tools.

The central risk is also real:

If sensitive data is pasted into public AI models, it has already left your control.

A Secure AI Workspace is the practical solution. It allows the organization to *say "yes" to GenAI* while keeping sensitive data governed, contained, and protected, and while providing evidence of what was and was not disclosed through AI usage.

Appendix A: Quick guidance for leaders and users

One sentence everyone can remember
If you wouldn't post the data publicly, don't paste it into public AI models. [3]

Examples of what should not go into public AI models

- Patient information, diagnoses, treatments, insurance data
- Payment card data and banking details
- Client confidential information and privileged legal content
- Internal security details, incident reports, credentials, access paths
- Non-public financials, pricing, contracts, and M&A context
- Proprietary code and architecture details

Examples of generally safer AI usage

- Drafting generic emails with no sensitive context
- Brainstorming generic ideas and outlines
- Summarizing documents that are already public
- Rewriting content that does not include sensitive information

When in doubt, use the Secure AI Workspace.

Appendix B: Questions to ask when evaluating AI options

Data control and retention

- What data is retained, for how long, and under what exceptions?
- What operational logs may contain prompts and responses?
- What happens under legal hold or other legal obligations?

Governance and evidence

- Can we prove what was and was not transmitted to model providers?
- Do we have centralized reporting across users and teams?
- Can we investigate and respond to risky usage patterns?

User enablement

- Will users actually adopt it, or will they revert to personal tools?
- Does it provide the models and features users already want?
- Can we create and share approved “Assistants” for common workflows?

References (for footnotes)

[1] BCG, *AI at Work 2025: Momentum Builds, but Gaps Remain*
https://www.bcg.com/publications/2025/ai-at-work-momentum-builds-but-gaps-remain

[2] Microsoft Work Trend Index (2024), *AI at Work Is Here. Now Comes the Hard Part*
https://www.microsoft.com/en-us/worklab/work-trend-index/ai-at-work-is-here-now-comes-the-hard-part

[3] NCDIT (North Carolina Department of Information Technology), *Use of Publicly Available Generative AI*
https://it.nc.gov/resources/artificial-intelligence/use-publicly-available-generative-ai

[4] OpenAI Platform Docs, *Data controls in the OpenAI platform (abuse monitoring log retention and legal exceptions)*
https://developers.openai.com/api/docs/guides/your-data

[5] OpenAI, *Enterprise privacy (retention controls and legal exception language)*
https://openai.com/enterprise-privacy/

[6] Anthropic Privacy Center, *How long do you store my organization's data? (law and policy enforcement exceptions)*
https://privacy.claude.com/en/articles/7996866-how-long-do-you-store-my-organization-s-data

[7] OAIC (Office of the Australian Information Commissioner), *Guidance on privacy and the use of commercially available AI products*
https://www.oaic.gov.au/privacy/privacy-guidance-for-organisations-and-government-agencies/guidance-on-privacy-and-the-use-of-commercially-available-ai-products

[8] PORT1, *Secure AI Workspace resource hub*
https://port1.io/secureai

[9] Liminal blog, *New Feature Roundup: Liminal Helper and Model Agnostic Assistants*
https://www.liminal.ai/blog/new-feature-roundup-liminal-helper-and-model-agnostic-assistants