Using GenAI Tools

How you can get safely started

May 2025

REVISION HISTORY

DATE	AUTHOR	NOTES
2025-05-07	R. Zaso	Initial Publication

Disclaimer: AI technologies are evolving rapidly. While care has been taken in content creation, AI-generated responses may contain inaccuracies. You are encouraged to use this content as a foundation for independent research and verification. As with any publicly available AI tool, care must be taken not to share personal, proprietary, or sensitive content. Never input confidential data (e.g., client names, strategies, proprietary code) into third-party AI platforms unless there is an approved enterprise license with data controls. Assume all inputs and outputs may be logged by the provider unless stated otherwise. For secure AI deployment, use enterprise-grade or self-hosted models with proper governance.

This document was produced through human-AI collaboration, with human review and final approval.

Table of Contents

Introduction

As generative AI tools gain traction across industries, many organizations are eager to explore their capabilities—**even without enterprise licenses**. For businesses or functional teams experimenting with consumer-facing generative AI tools—such as ChatGPT Free, Claude.ai, Gemini, Grok, Microsoft Copilot, Perplexity.ai; the path forward must strike a balance between innovation and information protection. These publicly accessible platforms provide a powerful entry point, but they operate outside enterprise-level data controls. This document outlines how organizations can begin using these tools responsibly while exploring options that align with privacy-first principles.

Getting Started

1. **Identify the AI Tool:** First decide which tool best suits your needs. Each platform offers different strengths and interfaces. It's common to use more than one tool depending on your needs.

Tool	How to Access Public Version	Best for	Starting Tips
ChatGPT (OpenAI)	chat.openai.com (Free and Plus versions)	Conversational AI, brainstorming, drafting, coding help	Use ChatGPT-4 (via Plus plan); disable chat history for privacy
Claude (Anthropic)	claude.ai	Long-context tasks, summarization, writing	Use Claude 3 Opus for free with email signup; avoid PII in inputs
Gemini (Google)	gemini.google.com	Google ecosystem tasks, quick info retrieval, document analysis	Use with personal or Gmail-based Workspace account; review terms
Grok (XAI)	grok.com	Tracking breaking news, trending topics, or current events	Sign in with an xAI account (or create one), and select Grok 3 or Grok 3 mini from the model dropdown menu

Tool	How to Access Public Version	Best for	Starting Tips
Microsoft Copilot	copilot.microsoft.com	Microsoft 365 integration, Excel/Word support, summarizing emails	Available through free Microsoft account; limit sensitive inputs
Perplexity (Perplexity.ai)	www.perplexity.ai	Research-focused queries, uses retrieval-augmented generation (RAG) to enhance factual accuracy, real-time web search for fast summarization and citations	Use for factual queries; avoids hallucinations; no login required for basic use

2. **Access the Platform:** Once you've selected a tool, you'll need to access its public interface. This generally requires:
    a. **Visiting the Website:** Most of these tools have a dedicated website (e.g., chat.openai.com for ChatGPT, gemini.google.com for Gemini, claude.ai for Claude, copilot.microsoft.com for Microsoft Copilot).
    b. **Creating an Account (if required):** On some platforms, creating a free account is necessary, and you'll typically have the option to sign up with your email address or a social media account.
3. **Get Familiar with the User Interface:** The user interface is typically a chat-like window where prompts or questions can be entered. Take time to understand the layout and any available features or settings. For example:

OpenAI ChatGPT

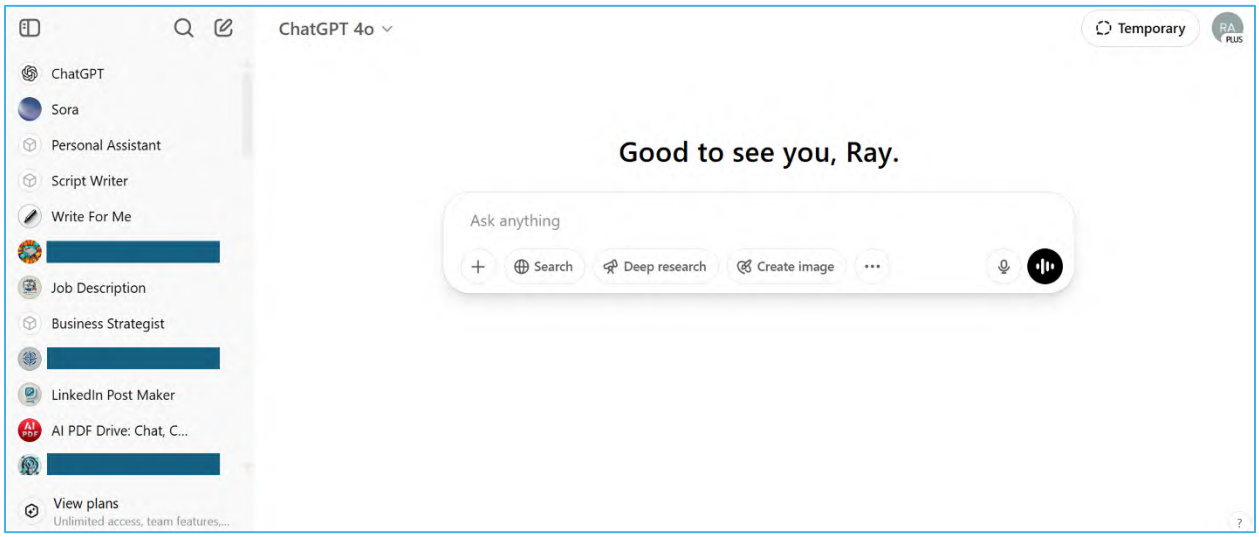

Anthropic Claude

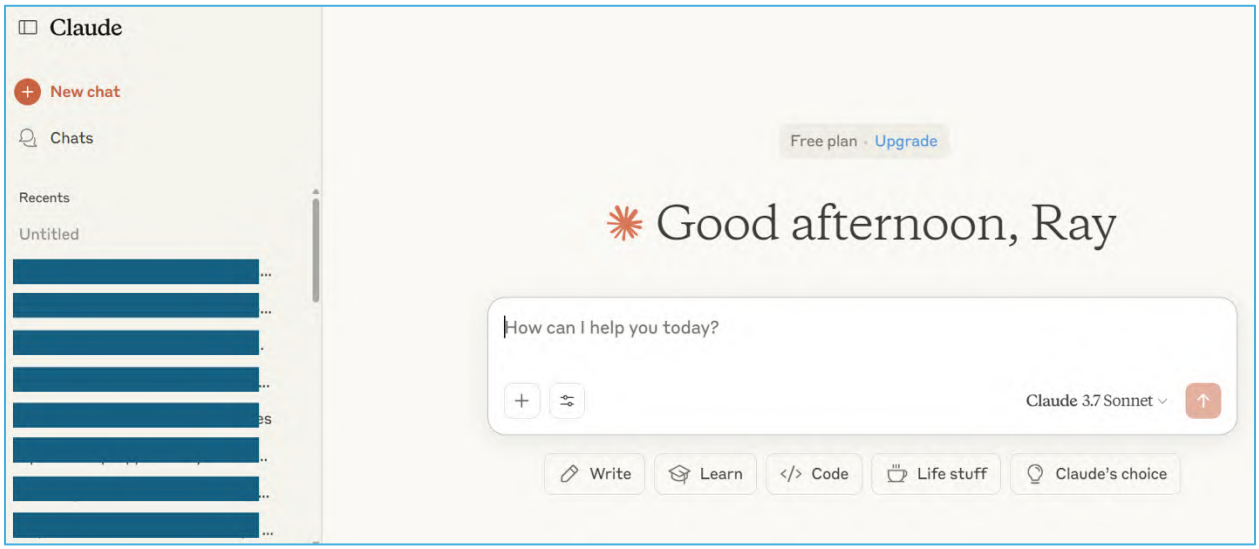

Google Gemini

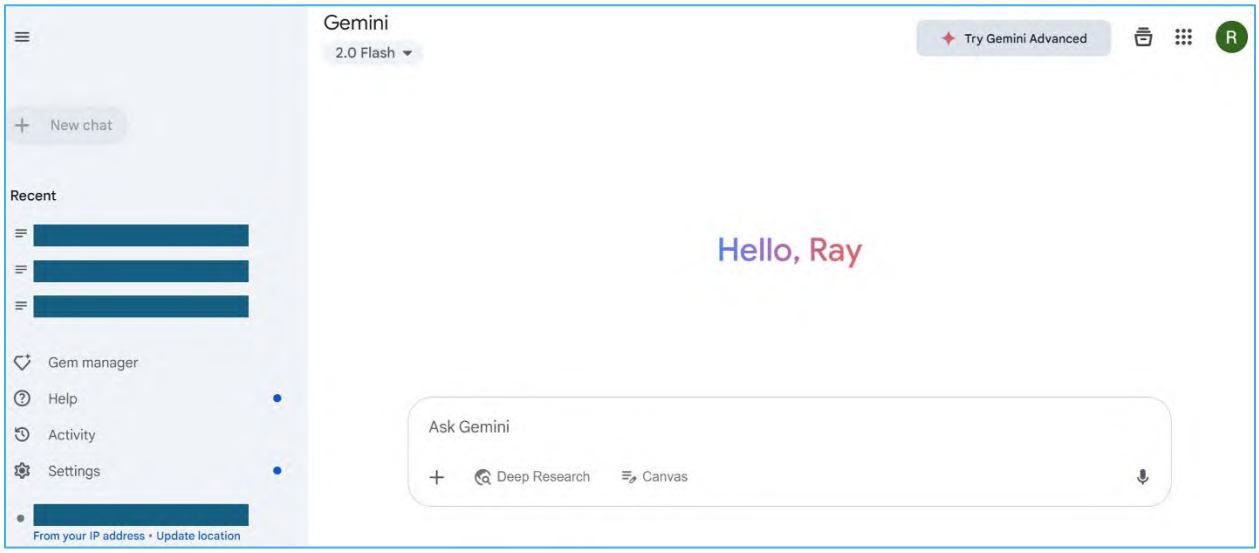

Grok

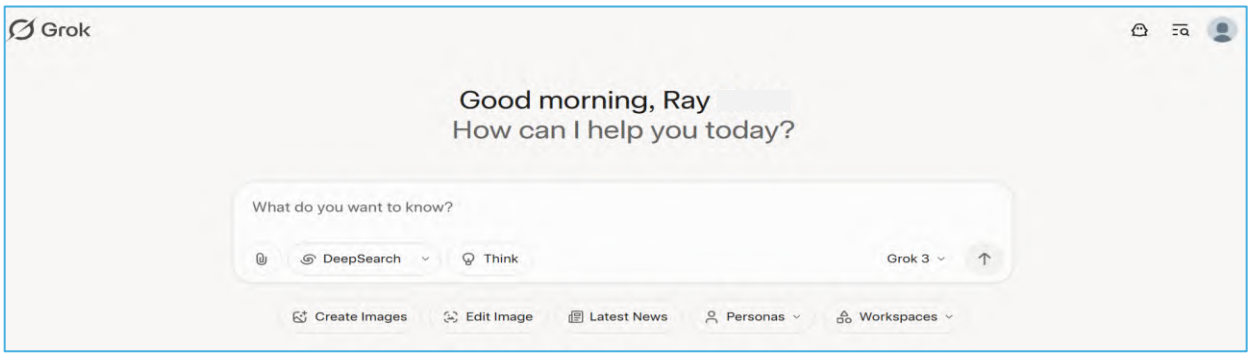

Microsoft Copilot

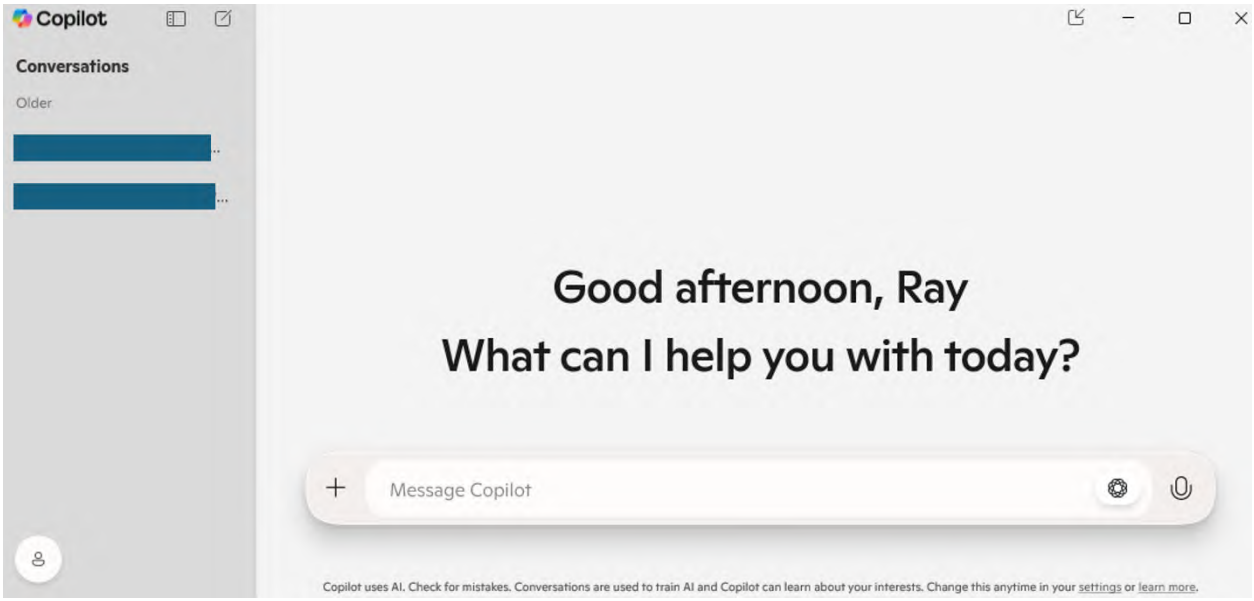

Perplexity

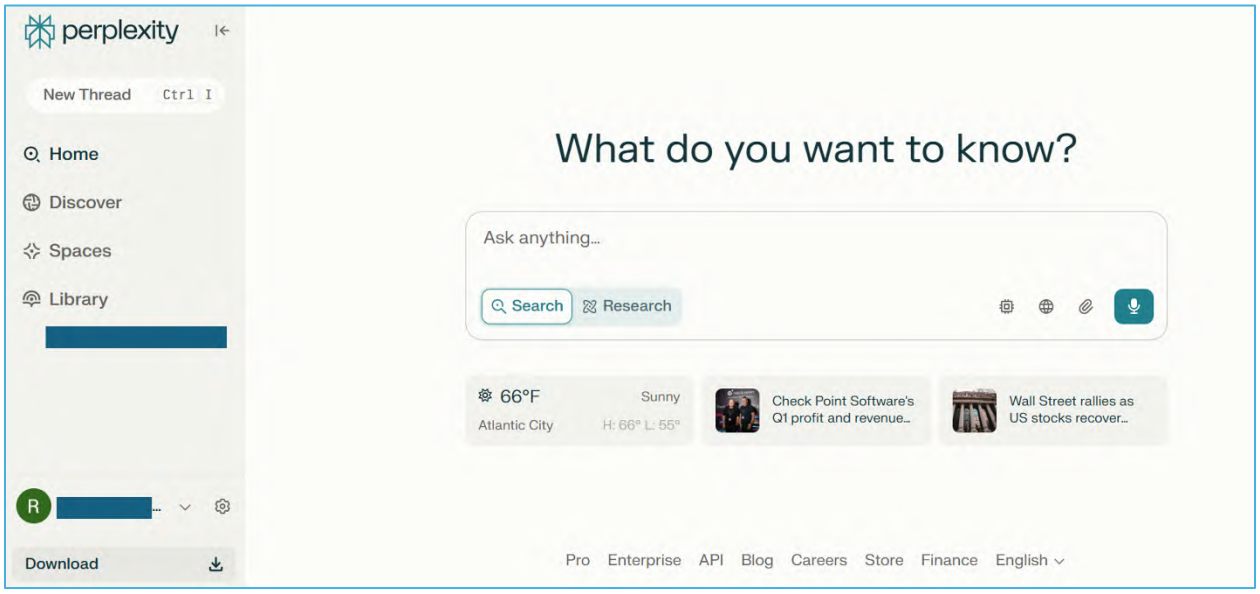

4. **Now Start Interacting:** Start typing in queries or prompts. It's helpful to start with simple requests to understand how the model responds and then gradually increase complexity. For instance, you could ask for summaries, creative writing samples, explanations of concepts, or help with brainstorming.

5. **Explore Features and Limitations:** Public versions might have certain limitations compared to enterprise versions, such as usage caps, slower processing times, or fewer advanced features.

a. **Check the official pricing pages:** Most platforms will have a dedicated pricing page that outlines the different tiers (free, pro, enterprise, etc.) and explicitly lists the features available in each.
b. **Usage quotas or limits:** This might be the number of messages per day/hour/month, the number of API calls, or the size of files you can upload. For example, free tiers often have lower message caps.
c. **ChatGPT:** The free version has limitations on GPT-4o usage, with a certain number of messages allowed every 3 hours, which can vary based on demand. Paid plans like Plus and Team offer higher limits.
d. **Gemini:** The free version has daily limits on chat requests and code completions. Paid tiers offer significantly higher limits and access to more advanced models.
e. **Claude:** The free version has a daily message limit, while the Pro version offers a significantly higher limit with a reset cycle every 5 hours.
f. **Microsoft Copilot:** The free version recently removed usage limits on its "Think Deeper" and voice features. However, enterprise versions within Microsoft 365 might have different billing structures based on message consumption.
g. **Perplexity.ai:** Check their specific subscription tiers for any listed usage limits on queries or advanced features.
h. **Feature comparison tables:** These tables often clearly show which features are exclusive to the enterprise or paid versions.

For businesses, it's common to establish a usage and governance policy before allowing employees to access AI tools. As a fundamental best practice for ensuring responsible and secure use, every organization leveraging public GenAI tools should establish a comprehensive internal usage and governance policy. This policy needs to be clearly defined and should include:

- **Prohibited Data:** Explicitly list the types of sensitive information that should never be entered into public AI tools.
- **Acceptable Use Cases:** Define appropriate and safe ways to leverage these tools for work-related tasks.
- **Prompt Engineering Guidelines:** Provide guidance on how to formulate prompts to avoid revealing sensitive details.
- **Account Management (if applicable):** Stipulate guidelines for account creation and management.
- **Regular Training and Awareness:** Emphasize the importance of ongoing education on data privacy risks associated with AI tools.
- **Consequences of Policy Violations:** Clearly state the repercussions for not adhering to the policy.

Addressing Data Privacy Concerns

If you are using these tools in the public domain, meaning you do not have an enterprise license, **data privacy should be your primary concern**. Using publicly available GenAI tools raises legitimate data privacy concerns, since the data you enter is processed by the provider's or other third-parties and may be used to improve the models.

Configuration Guidance for Public Domain GenAI Tools

When leveraging public-facing generative AI tools, it's critical to adjust certain settings prior to use—particularly when working with sensitive, strategic, or enterprise-adjacent content. The following breakdown outlines key settings to review across the most prominent tools on the market: **ChatGPT (OpenAI), Gemini (Google), Grok (xAI), Claude (Anthropic), and Perplexity.ai.**

Platform	Settings to Adjust
ChatGPT (OpenAI) Free (GPT-3.5), Paid (GPT-4, GPT-4o)	**Chat History & Training**: • **Setting**: *Turn off "Chat history & training"* • **Why**: Prevents your interactions from being used to train future models. • **How**: Settings > Data Controls > Toggle off. **Custom Instructions**: • **Setting**: *Set role preferences and communication style."* • **Why**: Helps tailor the assistant's responses to your business tone and priorities. **Default Workspace** *(Team/Enterprise Accounts)*: • **Setting**: *Ensure secure workspace settings are active.* • **Why**: Enterprise accounts offer API keys and SSO for better control over access and data.
Claude (Anthropic) Claude 3, available via Anthropic or integrations like Slack / Notion	**Data Retention**: • **Setting**: *Check platform policies—Anthropic doesn't train on conversations unless explicitly permitted.* • **Why**: One of the more privacy-friendly platforms by default. **System Prompts (for API users)**: • **Setting**: *Establish clear role behavior at the start of sessions.* • **Why**: Claude responds better when given a defined persona or constraints. **Integration Permissions** (e.g., Slack bots): • **Setting**: *Restrict which channels Claude can read or post in.* • **Why**: Prevents accidental exposure of sensitive messages.
Gemini (Google)	**Data Usage**: • **Setting**: *Disable "Help improve Gemini by saving activity."*

Platform	Settings to Adjust
Gemini in Google Workspace, Gemini Pro (API & Bard successor)	• **Why**: Avoid sharing prompts or data with Google for model improvement. • **How**: Google Account > Data & Privacy > Gemini Activity. **Workspace Context Sharing** (if integrated): • **Setting**: *Limit app access to documents or email by default*. • **Why**: Gemini can read across Workspace apps—ensure scope is necessary. **Personalization Settings**: • **Setting**: *Adjust model tone and preferred domains*. • **Why**: Enhances relevance and reduces ambiguity in complex queries.
Grok (x.AI) Grox.com, X.com	**Data Usage:** • **Setting:** Opt out of allowing your posts and interactions to be used for training Grok • **Why:** Prevents your public posts, interactions, inputs, and results with Grok from being used to train xAI's models, reducing the risk of sensitive data being incorporated into AI training sets. • **How:** On X, go to Settings > Privacy and Safety > Grok > Uncheck "Allow your posts as well as your interactions, inputs, and results with Grok to be used for training and fine-tuning." On the Grok mobile app or Grok.com, go to Settings > Data Controls > Deselect "Improve the Model." **Conversation History Deletion:** • **Setting:** Delete your Grok conversation history. • **Why:** Removes your past interactions from xAI's systems within 30 days (unless retained for security or legal reasons), minimizing data retention and potential exposure. • **How:** On X, go to Settings > Privacy and Safety > Grok > Click "Delete Conversation History" and confirm. On the Grok mobile app or Grok.com, go to Settings > Data Controls > Select option to delete data. **Private Chat Mode:** • **Setting:** Use Private Chat (ghost icon) for sensitive interactions. • **Why:** Ensures conversations are not saved or used for model training, providing a higher level of privacy for sensitive queries. • **How:** In the Grok interface, select the ghost icon to initiate a Private Chat session.
Perplexity.ai	**Pro Mode / Copilot**: • **Setting**: *Enable "Copilot" for guided, context-aware querying*.

Platform	Settings to Adjust
Web + Pro version; integrates real-time search	• **Why**: Adds structure and provides more useful results for research-heavy tasks. **History Settings**: • **Setting**: *Clear session history after sensitive queries.* • **Why**: Prevents others (shared devices or accounts) from seeing past prompts. **Citations & Search Sources**: • **Setting**: *Toggle inclusion of certain domains or request peer-reviewed sources.* • **Why**: Enhances trustworthiness for enterprise-grade research.
Microsoft Copilot (365 Copilot + Copilot Pro) Integrated into Microsoft 365 apps (Word, Excel, Outlook, etc.), available via enterprise and individual Pro subscriptions	**Data Access and Document Permissions**: • **Setting**: *Restrict access to specific documents or folders within SharePoint or OneDrive.* • **Why**: Copilot pulls data contextually from Microsoft 365 files—tight permission management is critical to avoid accidental data exposure. • **How**: Via Microsoft 365 Admin Center or per-file permissions. **Copilot Activity and Logging** *(Admin Control)*: • **Setting**: *Enable or review logging and audit trails for Copilot-generated content.* • **Why**: For traceability, compliance, and user accountability in regulated industries. • **How**: Microsoft Purview and Audit Logs can capture Copilot use. **Privacy Configuration** *(Enterprise & Pro users)*: • **Setting**: *Ensure organizational policies are enforced via Microsoft Purview.* • **Why**: Enterprise versions are built with data residency, encryption, and tenant-level data separation—this should be fully activated. **User Education and Prompt Best Practices**: • **Setting**: *Train users not to enter confidential information into natural language prompts.* • **Why**: While Copilot is secure within M365, human error (e.g., over-disclosure) remains a risk. **Plugin and Add-In Management**: • **Setting**: *Review and disable third-party plugins if unnecessary.* • **Why**: Reduces risk of data leakage through non-Microsoft extensions accessing Copilot-generated data.

Additional Considerations

1. **Avoid Sharing Sensitive Personal Information (SPI):** This is the most crucial step. Your client should absolutely refrain from inputting any information that could be considered personally identifiable, confidential, or proprietary. This includes:

    - **Direct Identifiers:** Names, addresses, phone numbers, email addresses, social security numbers, financial account details, medical records.
    - **Indirect Identifiers (in combination):** Job titles combined with company names, specific project details that could reveal sensitive information.
    - **Proprietary Business Data:** Trade secrets, unpublished financial information, strategic plans.

    Don't assume your data won't be used for model training—even if the provider claims not to use it unless opted in. Review the Terms of Service agreements.

    Don't assume access is private if using a shared or managed browser/computer.

2. **Formulate Generic Prompts:** When seeking assistance try to frame requests in a general way without revealing specific sensitive details. For example, instead of asking, "How can I improve the marketing strategy for our new product X which has these specific features and targets this niche market?", you could ask, "What are some general marketing strategies for a new product launch?"

3. **Paraphrase and Anonymize Data:** If you need to work with data that contains sensitive elements, first paraphrase and anonymize it before inputting it into the GenAI tool. This involves removing or replacing any identifying information and rephrasing the content to be more general. For example, instead of "The sales figures for John Doe in Q3 were...", you could say, "Sales figures for a representative in the last quarter were...".

4. **Review the Platform's Privacy Policy:** Carefully read the privacy policy of the specific GenAI tool they are using. This document outlines how the provider collects, uses, and stores user data. Understanding the policy can help in making informed decisions about what information to share.

5. **Be Mindful of Conversation History:** Many public GenAI tools retain conversation history. If the platform allows, consider deleting conversations that might have contained any inadvertently shared sensitive information. Also check if there are settings to disable or limit conversation history retention.

6. **Use Tools that Offer Data Usage Controls (if available):** Some platforms might offer settings that allow users to opt-out of having their data used for model training. Explore these settings if available and configure them according to their privacy preferences.

7. **Consider Using Local or Privacy-Focused Alternatives (if applicable):** While the mainstream tools like Gemini, Claude, ChatGPT, and Copilot are primarily cloud-based, there might be emerging privacy-focused AI tools or even open-source models that could be run locally in the

future, offering more control over data. However, these may be less mature and might have limitations in functionality compared to the leading platforms.

8. **Educate Users on Best Practices:** If multiple individuals within your organization will be using these tools, it's crucial to educate them about data privacy risks and the best practices outlined above. Regular training and awareness campaigns can help minimize the chances of accidental data leaks.

By adhering to these guidelines, you can leverage the power of publicly available GenAI tools while significantly reducing the risk of sensitive information being exposed. It's a trade-off between convenience and control, and careful usage is key to maintaining data privacy in this context.

Recommended Approach for Non-Enterprise Use

For organizations exploring the adoption of powerful GenAI tools without the immediate need for an enterprise license, a measured and strategic approach is recommended. Below is a step-by-step onboarding strategy designed to facilitate initial exploration, manage potential risks, and lay the groundwork for a secure and scalable future integration of AI capabilities within your organization. By focusing on low-risk applications and establishing clear guidelines from the outset, this approach allows for valuable learning and adaptation while prioritizing data privacy and responsible innovation.

1. **Define Use Case:** Choose low-risk applications like marketing copy, training outlines, or content summarization.
2. **Select Public Tool:** Start with a publicly accessible model (e.g., ChatGPT, Claude, Gemini, Microsoft Copilot, or Perplexity.ai) that aligns with your use case.
3. **Set Ground Rules:** Publish internal guidelines on acceptable use and input restrictions.
4. **Run Pilot Projects:** Test output quality, team adoption, and identify friction points.
5. **Evaluate Need for Control:** If usage expands, consider hybrid deployment (public + private tools)
6. **Plan a Secure Future State:** Budget for privacy-centric solutions (open-source, API-based, or enterprise-tier tools).

Using public GenAI tools can be a **safe and strategic starting point**—as long as organizations treat these platforms as **public spaces, not private assistants**. With the right safeguards, education, and a pathway to futureproofing with private or local models, your client can unlock the value of generative AI **without compromising data privacy or compliance.**

GenAI Free or Paid Subscription Checklist

GenAI Implementation Planning Checklist (✓)	
1. Clarify Use Case & Intent	
	a. Identify low-risk use cases (e.g., content drafting, summarization, FAQs, internal process documentation)
	b. Confirm business objective(s) and success metrics (e.g., time saved, quality of output, team adoption)
	c. Determine user groups or departments who will participate in pilot
2. Establish Data Privacy Guardrails	
	a. Define what should never be entered into public GenAI tools (PII, financials, contracts, client info)
	b. Create and share internal usage policy or GenAI guidelines
	c. Instruct users to disable chat history or personalization settings (where available)
	d. Mandate use of anonymized or hypothetical data for prompt testing
	e. Track prompts used and responses generated for auditing purposes
3. Choose & Configure Tool(s)	
	a. Public Cloud-based Tools: ChatGPT, Claude, Gemini, Duck.ai, Microsoft Copilot — Signup only; configure settings to enhance privacy
	b. Local / Self-Hosted Tools: LM Studio, Ollama, PrivateGPT — Requires desktop install and/or command line; minimal to moderate IT support
	c. Hybrid Enterprise (optional): Azure OpenAI with private endpoints — IT support required; suitable for mid-sized and regulated environments
	d. Review and document platform privacy policies before use
	e. Select tools based on user technical proficiency
	f. Confirm compatibility with company IT policy or endpoint security standards
	g. Public Cloud-based Tools: ChatGPT, Claude, Gemini, Duck.ai, Microsoft Copilot — Signup only; configure settings to enhance privacy
4. Pilot & Test	
	a. Run controlled 2–4 week pilot using selected use case
	b. Capture user feedback (ease of use, confidence, perceived risks)
	c. Evaluate output quality and assess potential for value scaling
	d. Adjust usage rules or tooling based on pilot results
5. Plan for Expansion or Secure Alternatives	
	a. Determine if more secure options (e.g., local LLMs, private APIs) are needed
	b. Explore long-term options like: - Self-hosted Ollama or PrivateGPT for internal workflows - Using LM Studio for offline experimentation - Adopting Azure OpenAI or private Anthropic/GCP integrations
	c. Work with IT to establish infrastructure and endpoint controls (if moving beyond public tools)
6. Educate & Evolve	
	a. Schedule team briefing or training on GenAI best practices
	b. Provide onboarding docs or cheat sheets for each tool
	c. Encourage knowledge sharing across departments
	d. Monitor AI developments and review usage quarterly
	e. Maintain a library of effective prompts and anonymized examples for team reference

Notes

- Start simple and stay safe: public tools can deliver early wins with the right restrictions.
- As comfort and use grows, migrate toward hybrid or local models for scalability and data control.
- Always treat public GenAI tools as external services, not internal systems.

Using Generative AI Tools Under an Enterprise License

As generative AI tools continue to transform workflows, many organizations are opting to adopt platforms like **ChatGPT (OpenAI), Claude (Anthropic), Gemini (Google), Microsoft Copilot**, or **Perplexity.ai** under **enterprise license agreements**. These subscriptions offer advanced functionality and significantly stronger data protections compared to their public counterparts—making them a strategic option for teams working with sensitive information, regulated data, or proprietary IP.

Enterprise Licensing Key Advantages

Feature	Benefit to the Organization
Data Privacy Guarantees	Enterprise plans typically include **contractual guarantees** that user data is not used to train models, is not retained, and is isolated from public-facing systems.
Dedicated Infrastructure	Enterprise customers may receive **private endpoints**, region-specific hosting, and data segregation, reducing exposure to third-party access.
Identity & Access Controls	Integration with **SSO (Single Sign-On)**, **RBAC (Role-Based Access Control)**, and audit logging supports compliance with internal security standards.
Administrative Controls	Admins can **monitor usage, enforce prompt restrictions**, and **disable features like history saving** or file uploads organization-wide.
Service-Level Agreements (SLAs)	Guaranteed uptime, support availability, and response times ensure business continuity and performance.

How Enterprise Licensing Addresses Data Privacy Concerns

1. **No Training on Customer Data**

   Leading providers like OpenAI, Microsoft, Anthropic, and Google explicitly state that under enterprise agreements, **inputs and outputs are not used to improve the models**. This contrasts with many free/public versions where data may be logged or retained.

2. **Data Residency & Compliance**

   Enterprise customers often have options to select data residency (e.g., U.S., EU), helping meet **GDPR**, **HIPAA**, **SOC 2**, and **ISO 27001** standards.

3. **Zero-Retention Policies**

   With appropriate configurations, enterprises can ensure **ephemeral data handling**, where prompt content is not stored on disk or logs, and no long-term retention occurs.

4. **Vendor DPA (Data Processing Agreement)**

   Executing a DPA with the provider ensures that **data handling obligations are legally binding**, supporting your organization's broader data protection policies.

   Example: Enterprise Use Cases

Use Case	Example Tool
Legal drafting & research with sensitive client data	Claude via Anthropic Enterprise API
Financial modeling in Excel with private spreadsheets	Microsoft Copilot for Microsoft 365
Private knowledge base chatbot for internal operations	ChatGPT Enterprise or OpenAI API with Azure endpoint
Research summaries and citation-supported queries	Perplexity Pro for Teams (coming soon)
Secure internal marketing or HR content generation	Gemini for Google Workspace Enterprise

Enterprise-grade AI tools not only unlock advanced AI features, but also offer the control, security, and scalability needed for responsible adoption in corporate environments. By investing in an enterprise subscription, your organization can innovate confidently while maintaining compliance, privacy, and brand integrity.

GenAI Enterprise License Checklist

GenAI Implementation Planning Checklist (✓)	
1. Define Strategic Intent	
	a. Identify high-impact, enterprise-aligned use cases (e.g., legal document generation, finance modeling, HR workflows).
	b. Align GenAI adoption with business goals, digital strategy, and compliance requirements.
	c. Establish success metrics (e.g., cost savings, productivity improvement, cycle time reduction).
	d. Engage cross-functional stakeholders early (Legal, IT, Compliance, Risk, Business Units).
2. Mitigate Data Privacy and Security Risks	
	a. Select an enterprise subscription that guarantees: - Zero retention of inputs/outputs - No model training on customer data - Region-specific data hosting options
	b. Execute a Data Processing Agreement (DPA) with the vendor
	c. Implement administrative controls: - Audit Logs - Role-Based Access Controls (RBAC) - Input/output monitoring
	d. Define rules on allowable input types (no PII, PHI, sensitive client information without masking/anonymization)
	e. Ensure all end users are trained on responsible AI usage guidelines
3. Choose & Configure Tool(s)	
	a. Enterprise Cloud-based Tools: ChatGPT, Claude, Gemini, Microsoft Copilot; configure settings to enhance privacy
	b. Ensure integration with SSO and existing identity platforms
	c. Select tools based on user technical proficiency
	d. Confirm compatibility with company IT policy or endpoint security standards
	e. Vet for data residency compliance (e.g., GDPR, HIPAA, FedRAMP)
4. Pilot & Test	
	e. Run controlled 2–4 week pilot using selected use cases
	f. Capture user feedback (ease of use, confidence, perceived risks)
	g. Evaluate output quality and assess potential for value scaling
	h. Adjust usage rules or tooling based on pilot results
5. Plan for Expansion or Secure Alternatives	
	a. Roll out in tiers—starting with low-risk departments.
	b. Create AI Centers of Excellence (CoE) to guide usage, tooling, and governance.
	c. Maintain a central knowledge base of prompts, workflows, and FAQs.
	d. Regularly refresh training and best practice materials.
	e. Schedule quarterly governance reviews (performance, compliance, ROI).
6. Educate & Evolve	
	a. Develop training tracks for end users, prompt engineers, and team leads.
	b. Establish internal policies on ethical AI use, hallucination mitigation, and user accountability.
	c. Encourage sharing of successful GenAI applications across departments.
	d. Promote transparent reporting and continuous improvement.
	e. Monitor AI developments and review usage quarterly
	f. Maintain a library of effective prompts and anonymized examples for team reference

Adopting GenAI with Confidence and Clarity

Key Precautions for Anybody

Whether using a free, paid, or enterprise license, always take these precautions:

- Thoroughly review inputs: Ensure all names, PII, sensitive company data, and confidential information are redacted or anonymized before entering prompts. A single oversight could lead to unauthorized disclosure.
- Treat public tools as external services: Free and paid subscriptions to public AI platforms lack the governance of enterprise solutions. Assume all interactions are visible to the provider unless otherwise specified.
- Opt for secure alternatives when needed: For confidential or regulatory-sensitive tasks, prioritize enterprise-grade tools or self-hosted GenAI models with robust data controls.
- Stay informed: Regularly review the provider's terms of service and your organization's AI usage policies to stay compliant with evolving standards.

Generative AI is no longer a future-facing concept, it's a present-day catalyst for productivity, innovation, and strategic transformation. Whether you're starting with free public tools or exploring enterprise-grade solutions, the path forward must be intentional, privacy-aware, and aligned with your business goals.

This guide has outlined how to get started safely, implement internal guardrails, explore privacy-preserving alternatives, and plan for enterprise-scale adoption. By taking a measured, governance-first approach, organizations can unlock the benefits of GenAI—faster insights, smarter content, streamlined operations—while minimizing risk.

The key to successful adoption lies not just in selecting the right tool, but in **how you use it**:

- Educate your teams.
- Set clear boundaries.
- Pilot purposefully.
- Scale responsibly.

With the right mix of structure, experimentation, and ongoing education, your organization can treat GenAI as a **strategic co-pilot** not just a productivity shortcut.

For tailored support, implementation assistance, or training sessions, connect with our team at The Zaso Group. We're here to help you navigate GenAI adoption with clarity, confidence, and compliance.

Resources and Contact Information

Curious about what AI can do for your business? Visit us at www.thezasogroup.com for strategies, resources, and expert insights.

Get inspired by transformative conversations on *Technology Reimagined with Ray Zaso*, now streaming on Spotify Podcasts.

Get inspired by transformative conversations on *Technology Reimagined with Ray Zaso*, now streaming on Apple Podcasts.

CONTACT

Ray Zaso, Strategic AI Advisor
The Zaso Group
+1 609-451-7026 office
rzaso@thezasogroup.com
www.thezasogroup.com