MENLO SECURITY

How AI Is Shaping the Modern Workspace

Insights Into the Use of Generative AI Based on Menlo Telemetry

Table of Contents

How AI is Shaping the Modern Workspace

Insights Into the Use of Generative AI Based on Menlo Telemetry

Artificial Intelligence:
From Science Fiction to Everyday Reality

Few technologies have promised greater productivity and decreased drudgery—and lived up to the hype—in the way that generative AI (GenAI) has. The actual cost savings resulting from the use of GenAI are impossible to calculate, as it is now integrated into browsers and applications, along with a plethora of tools and services. But as time has passed between ChatGPT's release in late 2022 and today, most agree that the overall impact is likely in the hundreds of billions of dollars (U.S.).

As a rapidly growing field, there are many questions about AI, but one thing is certain: there is no going back. As Stan Lee once told us, "With great power comes great responsibility." And, though the adage was applied specifically to the powers gained from radioactive spiders at the time, it couldn't be truer than it is now in our new world of AI. AI is already demonstrating its power in the modern workspace by streamlining business operations, automating tasks, informing better business decisions, and improving customer experiences. Enterprises now have the responsibility to ensure that its use is secure and appropriate.

In this paper, we consider both the benefits and hazards of AI, based on telemetry captured by Menlo Security, to give you an inside look into the ways you can optimize the many advantages of AI, while sidestepping the pitfalls.

A Unique Vantage Point for Viewing the AI Landscape

As a result of digital transformation, many essential enterprise apps and services have moved from a client/server model to cloud delivered and SaaS enabled. In other words, they're accessible via the web browser.

Menlo Security has focused on browser security for over a decade, and we have seen the browser move from a simple app that rendered HTML to what is essentially an operating system in its own right. It has become the launching point for most applications, tools, and services, including AI.

> The heavy reliance on web access has given Menlo Security an unobstructed view of the ongoing changes in the AI space.

The adoption of AI has been sufficiently large that it has changed the overall pattern of web traffic. In fact, web traffic to AI tool sites grew from seven billion visits in February 2024 to 10.53 billion visits in January 2025, a 50 percent increase.[1] While some AI deployments are moving to desktop or private applications, it is estimated that up to 80 percent of GenAI is still accessed via the browser, whether that access is through web-based AI sites or integration into existing web services. That stands to reason, because:

- The browser is the most popular and widely used application across desktops and mobile devices, ensuring broad accessibility across different operating systems and devices.
- The browser already interacts with other online tools and services, making it easier to integrate into the various AI services.
- The task of developing and deploying web-based apps is well understood, so a web-based approach speeds time to market for AI services.

The First Thing You Need to Know About AI Is That it's Not All the Same

The AI most of us are familiar with is reactive by nature—it looks for patterns and rules and requires significant management from humans. Responses to queries are generated from the original large language model (LLM) in the AI platform, and it can incorporate new information submitted in the form of queries if allowed to do so.

Many enterprises are starting to build their own AI, taking advantage of commercially available LLMs based on open source models, which may then be customized. Such an implementation requires specialized skill within the organization but will deliver full transparency—in other words, it will be clear where the AI tool got its information. Another approach is to purchase a closed source LLM from a third-party provider. This method is easier and faster to implement and is protected by the provider, but the methodology by which conclusions are reached can be somewhat opaque.

Agentic AI, however, is built on a different premise. Agentic AI is designed to be proactive, understanding user goals in context and taking steps autonomously to meet those objectives. It can learn from interactions, interpret data, and adapt to new information, making it ideal for automating complex processes.

While agentic AI is beyond the scope of this paper, it serves to illustrate how rapidly the industry is changing, as well as the importance of a good foundation.

By 2028, 33% of enterprise software applications will include agentic AI, up from less than 1% in 2024, enabling 15% of day-to-day work decisions to be made autonomously.[2]

[1] Generative AI's Continued Growth: Traffic to Top 50 AI Tools Up 50% in 11 Months, Similarweb, April 2025
[2] Intelligent Agents in AI Really Can Work Alone. Here's How. Gartner, October 2024

Menlo Observations and Insights

During a 30-day period in May–June 2025, Menlo Security analyzed generative AI interactions from hundreds of global organizations. As observed in the past, the use of GenAI crosses organizations of every size and industry.

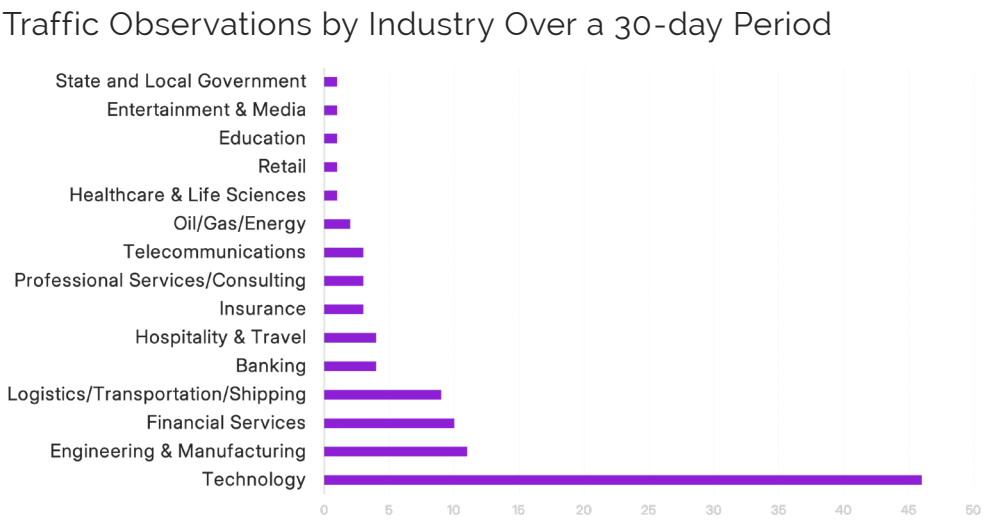

Figure 1: Percentage of traffic to sites classified as GenAI by vertical industry, as seen by Menlo

Browser Traffic to GenAI Sites Continues to Grow

Web traffic to sites classified as GenAI has continued to grow, albeit less quickly than in 2024. This is in keeping with what has been seen across industries. Menlo observed:

> Over 5.6 million site visits
> in a single 30-day period.

Geographical Uptake of GenAI May Foreshadow Upcoming Issues

Regionally, we observed the heaviest traffic to sites classified as GenAI in the Americas region and in the United States in particular. AI use has grown steadily in the U.S., perhaps because some of its underlying technologies, including neural networks and deep learning, began at American universities, and many advances that led to publicly available GenAI tools were developed by U.S. companies. It's worth noting that the rapid adoption of GenAI in this region began with individual users, a fact that has caused issues of its own, as we'll address later in this report.

While the Americas led in traffic volume, the Asia-Pacific region showed the most significant growth in the adoption of GenAI. China and India are at the forefront of this trend, with 75 percent and 73 percent of surveyed organizations reporting GenAI implementation, respectively.[3]

Growth in the Europe, Middle East, and Africa (EMEA) region has been steadily growing, although at a markedly slower rate than other regions. This finding is consistent with industry analysts' observations. The region's slower growth can be partially attributed to the EU's comparatively strict regulations around privacy and transparency.

Insight 1: GenAI Use Is Accelerating, And You Need To Take The Wheel

Any thought that enterprise use of AI would soon plateau have clearly gone by the wayside. Given the growth and diversification in AI tools, it's become increasingly urgent to create an overall enterprise strategy around AI.

This is particularly true when considering how the global growth of AI could affect compliance with data security and privacy regulations. The EU AI Act, for example, marks the first comprehensive AI law in the world, and other countries are due to follow. Organizations, particularly those with a global footprint, must seriously consider how the use of AI will affect compliance with existing regulations around data security and privacy.

The Number of AI Sites and Applications is Growing

Menlo Has Observed Over Twice as Many Domains as Apps

Over 6,500 Domains	Over 3,000 Apps

The distinction between domains and apps is important, as a single app can have multiple domains associated with it. Attackers have taken advantage of this fact, building "AI" sites of their own using combinations of well-known app names to add legitimacy.

This scam often shows up with extensions. Such threats were recently reported by MalwareBytes, which cited analysis from researchers at Extension Total, who noted that cybercriminals had managed to take over the accounts of at least 36 Google Chrome extensions that provide AI and VPN services.[4] The compromised extensions include "Bard AI Chat," "ChatGPT for Google Meet," "ChatGPT App," "ChatGPT Quick Access," and others.

[3] https://fintechnews.sg/110219/ai/apac-emerges-as-a-leading-genai-adopter/
[4] https://www.malwarebytes.com/blog/news/2025/01/google-chrome-ai-extensions-deliver-info-stealing-malware-in-broad-atta

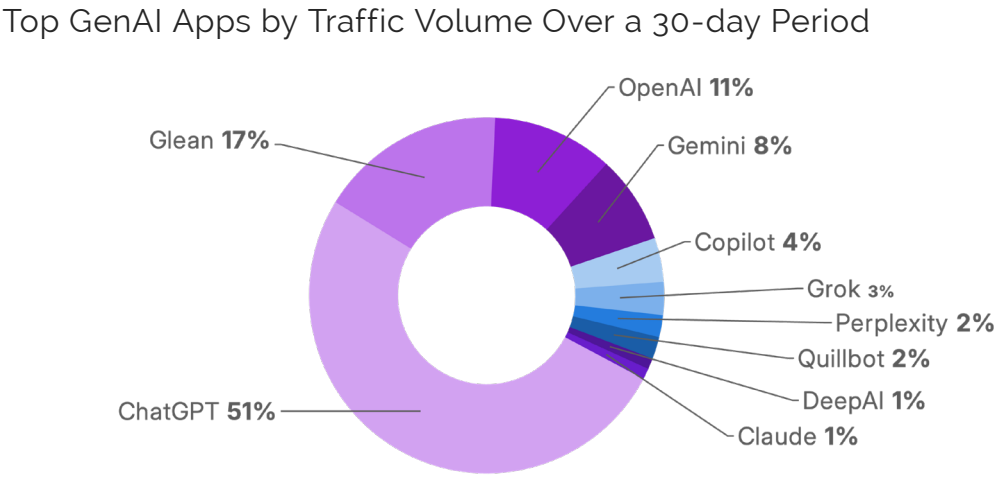

Figure 2: Percentage of traffic classified as GenAI by app as observed by Menlo

When considering apps, ChatGPT saw the most traffic by far. OpenAI shows up as the third most popular, a reminder that it also offers DALL-E, a text-to-image generation model, and Sora, a text-to-video model. Both are available as part of ChatGPT for paid subscribers.

ChatGPT is well established as one of the first AI tools to become widely available, so its dominance here is unsurprising. It is noteworthy, however, because while ChatGPT offers a variety of tiers to businesses, including their Plus subscribers and enterprise customers, the vast majority of users (estimated at over 95 percent) are working on the free offering, which shares queries and responses with the foundational model.

> As of May 2025, ChatGPT has 400 million active weekly users, including 15.5 million Plus subscribers and 1.5 million enterprise customers.[5]

The high use of Glean, which is designed to aggregate company data, is also notable. In addition to Glean's more general Assistant tool, the company offers Agents, which sync with enterprise data and can plan and execute workflows autonomously. The Glean system also features security measures and a detailed permission structure, which are vital.

[5] https://nerdynav.com/chatgpt-statistics

Attackers Are Jumping on the AI Bandwagon, Too

Later in this report we will discuss our observations around the use of fake AI sites to entice users. While fakes are significant, they are far from the only way that attackers are using AI themselves. In fact, according to the Verizon 2025 Data Breach Investigations Report, the percentage of AI-assisted malicious emails doubled (from roughly five percent to 10 percent) over the past two years.[6]

This rise is consistent with observations from the Menlo Threat Intelligence Team. Earlier this year, the team found and analyzed a phishing campaign that utilized a phishing kit called "Greatness." This kit is part of a phishing-as-a-service (PhaaS) infrastructure that is available for purchase via Telegram. The Greatness PhaaS includes email templates, command-and-control (C2) panel access, the kit, Telegram support, and other Telegram options. The ability to build campaigns from a kit makes the process even more simple.

Not only can AI help attackers build "better" phishing content by clarifying spelling and appearances, it can be used to create hyper-targeted "spear phishing" attacks. Menlo Threat Labs found that threat actors are increasingly using AI-powered attacks to exploit browser vulnerabilities, resulting in a 130 percent YoY increase in zero-hour phishing attacks and revealing a surge in GenAI imposter sites,[7] which are discussed specifically later in this report. With AI, hackers can automate phishing attacks, personalize messages, mimic human behavior, and evade security filters, making them harder to detect.[8]

One tactic being used to great effect is the combination of AI with Open Source Intelligence (OSINT). Traditionally used by defenders, OSINT has been used by security teams to gather publicly available information on their organization and get in front of attackers. The combination of OSINT with AI, however, can be used to build effective spear phishing attacks, automate the previously time-consuming process of social engineering, build convincing deepfakes, and more. When combined with AI web-scraping bots to gather data, the results can be devastating.

Insight 2: As GenAI Expands, It May Be Learning More About Your Business Than It Should

As GenAI sites proliferate and become more specialized, we may see an increase in the sensitivity of the data and resources that are required to complete their tasks. Dedicated AI tools that access internal company applications and data, along with those that correct and write code, have significant implications. While there is every indication that vendors take data security and privacy seriously, their use still introduces more risk for enterprises to manage. The use of free tools, which share queries and responses with their foundational models for training purposes, must be carefully considered by any enterprise concerned with data leakage and loss.

[6] Verizon 2025 Data Breach Investigations Report
[7] https://www.menlosecurity.com/resources/state-of-browser-security-report
[8] https://www.webasha.com/blog/how-hackers-use-ai-for-creating-spear-phishing-attacks-the-next-gen-cyber-threat

In order to ensure that all AI tools work smoothly together, it is essential for enterprises to have an over-

arching strategy about data security and privacy. It is also important to begin setting overall DLP rules, particularly for tools that have access to sensitive content. Logging is certainly a good first step to provide clarity on what content is going where, but logging cannot be the final action.

In addition to the rapid rise in apps and services, there's been a significant increase in domains, which is noteworthy. It is common for a domain to have many different URLs associated with it, but the rise in the GenAI category demands caution, because URLs that are "close" to real URLs can indicate the use of an attack technique called "typosquatting." The similarity makes it easier for bad actors to draft on the success of legitimate sites, as you will see later in this report.

Another important observation is the use of AI to build "bigger, better" phishing and malware delivery. GenAI makes it relatively simple for attackers to collect details about their targets, resulting in highly targeted spear phishing attacks arising from an AI analysis of the target's online presence, data harvesting using bots to scrape and gather information on the target, deepfake audio and video communications, and even fake chatbots. Less glamorous but equally effective, attackers are now using GenAI to create or re-skin JavaScript elements used to build dynamic webpages, often evading detection.

How Convincing Have Deepfakes Become?

Last year, a finance worker at an international firm was duped into paying out $25 million to fraudsters who had used deepfake technology to pose as company employees, including the chief financial officer.

The request was not in an email with a bad link or via a credible voice mail. The employee sat through a video conference with several people who looked and sounded just like colleagues he recognized. They were all faked.

The fraudsters had stolen these employees' ID cards and used AI deepfakes to trick facial recognition programs by imitating the people pictured on the identity cards.[9]

DLP in an AI Context

According to a recent study, 75 percent of customers believe that generative AI introduces new data security risks,[10] and the prospect of data loss or data leakage dominates the concerns. The term "data loss" refers to sensitive data being transmitted beyond a controlled environment, and data loss prevention (DLP) can be achieved with DLP software or a security framework that controls the flow of sensitive data between end users and internal resources.

While data loss is a legitimate concern in the enterprise, data leakage, in which sensitive information is inadvertently exposed, can be a bigger issue, particularly when it comes to GenAI. That is because many users in the enterprise engage with AI models via direct text input.

[9] https://www.cnn.com/2024/02/04/asia/deepfake-cfo-scam-hong-kong-intl-hnk
[10] https://www.amplifai.com/blog/generative-ai-statistics#generative-ai-statistic-56

Menlo customers can apply a variety of DLP rules, such as restrictions on copy and paste functions, including character counts, as well as controls that may trigger event logging, and the ability to block certain types of content from being uploaded. As a result, Menlo can observe each time these customer policies are seen.

IN A SINGLE MONTH, MENLO OBSERVED:

- **155,005** copy attempts
- **313,120** paste attempts
- **327 visits** to watermarked GenAI sites

Copy and paste restrictions can be indicated whenever enterprises are concerned about data leakage, loss, or transfer; the need becomes even more urgent in the context of GenAI. Users may well not intend to transfer sensitive information in their attempt to summarize or reword content, but it happens. Another potential issue is around data protection policies, which may require such restrictions.

Watermarking is a way to subtly remind users that they could be dealing with sensitive content. Still another method is to limit the number of characters that can be entered into an AI tool.

Of course, copy and paste is far from the only way that users interact with GenAI; some Menlo customers choose to tag content using specific DLP rules, as well. As Menlo examined traffic going to sites classified as GenAI, we found content uploads that triggered DLP rules.

CUSTOMER-DEFINED RULE	PERCENTAGE
SENSITIVE	70%
RESTRICTED	25%
PII	5%

The amount of data specifically called out as personally identifiable information (PII) may seem low, but it is important to realize that PII data also shows up as a subset of the much more comprehensive "sensitive" classification.

Content classified as "sensitive" can contain a plethora of data, including:

- PII, via direct identifiers such as full names, Social Security and National Insurance numbers, passport numbers, and more, as well as indirect identifiers including biometric data like fingerprints, facial recognition information, and retinal scans
- Protected health information (PHI), such as medical histories and lab results
- Financial information, including credit card, bank account, and routing numbers
- Intellectual property (IP), including trade secrets, proprietary code, customer lists and databases, and unpublished financial, product, and marketing plans
- Access credentials and system information, which can include usernames and passwords, PINs, private keys and certificates, network specifics, and admin details
- Legal and regulatory compliance data that can tie back to regulations like GDPR, HIPAA, PCIDSS, CCPA, and more

"Restricted" information differs from "sensitive" information and can be broadly defined as information that should not be shared for a variety of reasons.

Top File Types Tied to DLP Events Related to GenAI Traffic

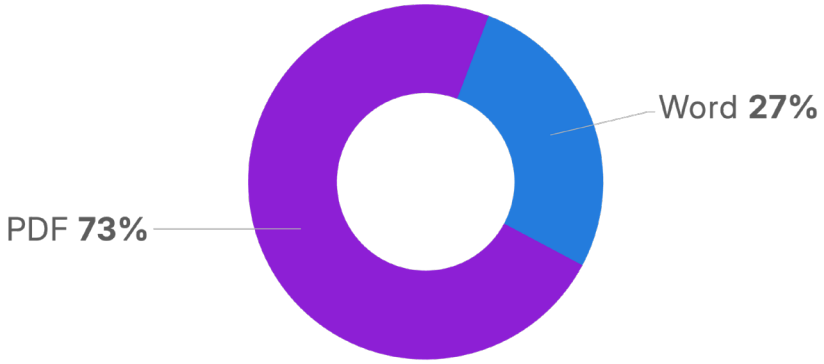

Figure 3: DLP events (allow/block/log) by file type observed by Menlo over a 30-day period

Downloads From AI Sources Should Not Automatically be Trusted

Content uploaded to GenAI sites are not the only area of concern for the enterprise. It is also important to consider content that may be downloaded from such sites.

Top File Types Downloaded From AI Tools

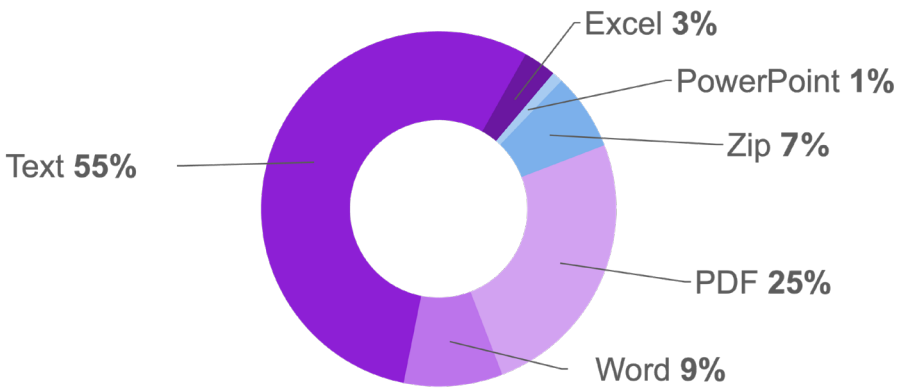

Figure 4: File types downloaded from sites classified as GenAI observed by Menlo over a 30-day period

Files coming into the enterprise from AI tools cannot be overlooked as a threat vector. PDFs have long been used to deliver a variety of malware, including JavaScript exploits, embedded executables, or phishing links.[11] Menlo telemetry showed that the majority of downloads from sites classified as GenAI were either PDF or text files.

[11] https://nordlayer.com/blog/can-pdf-have-virus

This is troubling, because PDFs can provide the perfect environment for embedded malware. According to the IBM X-Force 2025 Threat Intelligence Index, "...PDF malware disguises malicious URLs by encrypting them, hiding them in compressed streams or using hexadecimal representations which can also hinder automated analysis of email security solutions."[12] The incidence of text and Word files is also high on the list, and the combination makes content inspection a must have.

Shadow AI Takes DLP to Another Level

Making DLP issues even worse is the fact that many employees are indulging in "BYOAI," or Bring Your Own AI—also known as shadow AI—and are working with the free tier of their chosen tools.

The issues posed by BYOAI are two-fold. First, the enterprise loses control in such a situation, and second, most of the free-tier services state that they use data submitted to train their models. In fact, a recent survey found that nearly seven out of 10 (68 percent) enterprise employees who use GenAI at work say they access publicly available GenAI assistants such as ChatGPT, Microsoft Copilot, or Google Gemini through personal accounts, and more than half (57 percent) have admitted to entering sensitive information into them.[13]

As stated in the 2025 Verizon Data Breach Investigations Report, "Some of the most common use cases of GenAI tools—such as summarization or coding assistance—often invite the user to upload confidential documents and codebases to achieve them."[14] The Verizon Report goes on to say that a large numbers of users who were utilizing GenAI, "were either using non-corporate emails as the identifiers of their accounts (72%) or were using their corporate emails without integrated authentication systems in place (17%), most likely suggesting use outside of corporate policy."

This is one reason why the use of consumer-based BYOAI tools, which are usually free and typically share data, can be so concerning. Users are simply trying to be more productive as they ask the AI tools that they use at home to summarize the contents of a sensitive document, or find the errors in code, but the fact is that they may also be unwittingly sharing information in the process.

Insight 3: You Must Control GenAI As If Your Company's Security Depends On It...Because It Does

The implementation of DLP is a vital first step for using AI in a secure fashion. To be truly effective, it is important for DLP to cover the two most common types of use cases, including those where a user will interact with GenAI directly via text inputs, and instances where a user uploads a document, image, or other content type.

[12] https://www.ibm.com/thought-leadership/institute-business-value/en-us/report/2025-threat-intelligence-index
[13] TELUS Digital Survey
[14] Verizon 2025 Data Breach Investigations Report

Menlo was able to show both the category classification and the file types uploaded because many

Menlo customers have added DLP rules around content, including allow, block, and log. It has become essential, however, for enterprises to go beyond simply logging DLP violations. By understanding the types of content being shared—and where they are being shared—enterprises can set and communicate rules about how to work with AI, and can educate users who may not realize what a data leak or breach could do to the company overall. Beginning this process now, and continuing the discipline, will serve users and enterprises alike.

Controls Must Extend Beyond Upload and Download

In cases where direct text input is allowed, a few steps are necessary. First, the enterprise must have visibility into what AI model is being used. If the model is a free-tier offering accessed via a personal login, it should be banned. Not only do security and IT teams lose visibility in such situations, but these models often train on data submitted by users. Next, enterprises must establish a sanctioned AI model and send users there exclusively.

In cases where the AI tool enables direct input, detailed controls must be built-in, including copy/paste controls, as well as provisions on uploads and downloads. In short, the enterprise must treat AI tools like sensitive applications.

Shadow AI Must Be Controlled or Eliminated

Any DLP problem is concerning, but as the use of AI outpaces enterprise control, anything bad is only going to get worse. Simply notifying users of corporate policy will not provide the sorely needed protections. In order to eliminate shadow AI, enterprises must select sanctioned AI systems or tools that they trust—and mandate their sole use.

While enterprises can (and, arguably, should) control what AI tools are used on their network, there is no way to control what AI tools are used by employees taking advantage of BYOD or third-party contractors and partners. If organizations cannot control which AI tools are used outside the network, they will have to control what is allowed inside. Malware must be stopped at the door.

Not All AI Sites Are What They Appear to Be

Many of today's most serious threats involve malicious sites masquerading as legitimate AI sites. In its 2025 State of Browser Security report, Menlo observed nearly 600 incidents of phishing sites using URLs that imply a connection to legitimate GenAI names—and this trend is not slowing. In March alone, Menlo Protect with HEAT Shield AI blocked more than 40 new sites that purported to be GenAI sites but were actually identified as phishing sites. In addition, Menlo observed malicious sites that featured "ChatGPT" or "Copilot" in the domain name. In 2025, sites impersonating GenAI began to use "Gemini" in the domain name in the majority of tracked instances.

Between 2024 and 2025, researchers tracked thousands of look-alike domain names and impersonation websites (around 2,600 websites surfaced between December 1, 2024, and February 3, 2025) mimicking DeepSeek, as reported in SecurityWeek.[15] New impersonation domains pop up constantly, and while some

are obvious fakes, others are reportedly more difficult to spot. This is only one reason why it is vital that enterprises stop allowing users to make decisions about AI tools on their own.

> In the last year, Menlo identified nearly 600 imposter sites with GenAI-related names that were actually identified as phishing sites.

Some bad actors are taking advantage of AI by creating specialized "AI" tools. One example is a fake AI tool website posing as an app offered by Novaleads, an affiliate online marketing platform. Victims are lured into downloading the tool with the promise of a free, 12-month subscription. Once downloaded, CyberLock ransomware encrypts files across multiple disk partitions.[16]

Conversely, not all AI sites that cause harm are automatically fake. In 2025, millions downloaded DeepSeek. Employees who did so put their company's data at risk—an unsecured DeepSeek database exposed a million lines of log streams containing chat history, secret keys, backend details, and other highly sensitive information, which highlights the need for organizations to scrutinize and secure all AI apps.[17]

Extensions Can Be Equally Problematic

Malicious browser extensions claiming to offer AI capabilities have been popping up almost since ChatGPT was released. One of the first, Chat GPT (note the space) purported to integrate ChatGPT results into Google searches. The extension, which was based on the same open source code as the genuine extension (ChatGPT for Google), was promoted online and offered on the Chrome Web Store before it was discovered. One thing that made the malicious extension so compelling was that it actually functioned as it was supposed to; unfortunately, however, it was also harvesting Facebook and OpenAI tokens in the background, potentially enabling bad actors to gain access to user data.

While this was one of the first malicious extensions masquerading as AI, it was far from the last. A number of threats have been unmasked in 2025, including extensions trading on well-known names, including "ChatGPT for Google Meet," "Bard AI Chat," "AI Assistant – ChatGPT and Gemini for Chrome," "Search Copilot AI Assistant for Chrome," "GPT4 Summary with OpenAI," and others. These extensions injected data-stealing code, collected sensitive user data, compromised user credentials, and more.

[15] https://www.securityweek.com/beware-of-deepseek-hype-its-a-breeding-ground-for-scammers/
[16] https://www.bleepingcomputer.com/news/security/cybercriminals-exploit-ai-hype-to-spread-ransomware-malware
[17] https://info.varonis.com/en/state-of-data-security-report-2025

Insight 4: Fake AI Tools Present A Genuine Risk

The presence of phishing sites claiming to be AI sites is growing, partly fueled by the large number of malicious URLs built around well-known domains. Extensions are another area where users may seek to gain the productivity boost of AI by installing an extension they think will amplify search results, summarize content, or provide writing assistance, among other benefits. It is essential that enterprises provide guidance and enact enforcements to prevent the use of illegitimate or malicious sites and extensions.

In Summary

If AI tools have much to offer today's enterprise, but those benefits are contingent upon the safe use of the technology. Luckily, by securing enterprise browsers and their traffic, you can improve your security posture and lay a solid foundation to handle new and emergent threats. Not only will securing the browser and browser traffic protect you from the potential pitfalls of AI, but it will also enhance the security of your enterprise overall.

There are several key steps you can take right now to establish safe practices around the use of GenAI:

- Eliminate shadow AI
- Set and enforce rules for the secure use of AI tools
- Safeguard your data
- Enable content inspection
- Protect the enterprise from fake AI sites and defend against the growing use of AI in phishing and other attacks
- Protect access to internal and SaaS apps
- Provide defense-in-depth by securing the local browser

For more information on these steps and the Menlo approach to making the most out of GenAI, please contact us.

About Menlo Security

Menlo Security eliminates evasive threats and protects productivity with the Menlo Secure Cloud Browser. Menlo delivers on the promise of cloud-based security—enabling zero trust access that is simple to deploy. The Menlo Secure Cloud Browser prevents attacks and makes cyber defenses invisible to end users while they work online, reducing the operational burden on security teams.

Menlo protects your users and secures access to applications, providing a complete enterprise browser solution. With Menlo, you can deploy browser security policies in a single click, secure SaaS and private application access, and protect enterprise data down to the last mile. Secure your digital transformation with trusted and proven cyber defenses, on any browser.

Work without worry and move business forward with Menlo Security. © 2025 Menlo Security, All Rights Reserved.

MENLO SECURITY

Learn more: https://www.menlosecurity.com
Contact us: ask@menlosecurity.com