MENLO SECURITY

Steps You Can Take Now to Counter GenAI Risks

The use of generative AI in the enterprise is skyrocketing. Learn how to embrace its advantages while mitigating its risks.

In a companion paper, How GenAI is Shaping the Modern Workspace, we explored the use of generative AI in the enterprise based on our own metrics and reporting by industry experts and analysts. In that paper, we offered insights into some trends we're seeing around the use of GenAI and the ways that GenAI is evolving, including its increasing use in phishing and deepfakes. And we discussed the importance of automating protections against the rising threats posed by GenAI use, both those posed by the actions of uninformed users inside the enterprise and the malicious use by outsiders.

In this paper, we are expanding on that last insight about steps you can take to protect your organization as you navigate the rise of GenAI in the enterprise. We'll also discuss solutions you can implement right now to mitigate its risks and challenges.

It is Time to Implement a Strategy for the Safe Use of GenAI

We know that global traffic to GenAI sites is rising. In a single quarter in 2025, Menlo Security metrics revealed a 70 percent increase in such traffic. Furthermore, the use of GenAI is expected to continue rising around the world, putting AI use on a collision course with regionally specific regulations and laws. In fact, in some major firms, including Meta, Box, Shopify, Microsoft, and Amazon, the use of AI is increasingly required, particularly in cases where higher-level human reasoning is nonessential.[1]

[1] https://www.washingtonpost.com/business/2025/06/03/ai-workplace-duolingo-shopify-employees/

SOLUTION BRIEF

In addition to GenAI's rapid growth, every enterprise needs to be concerned with the fact that attackers are using AI themselves, including creating fake AI sites that impersonate legitimate ones. Remember how easy it used to be to detect phishing attempts? Not any more. AI is helping attackers create emails and SMS messages that are completely convincing. They're embedding malicious JavaScript in dynamic webpages, creating deepfake content, and more.

Even if you have not established strong policies around the use of GenAI in your enterprise, you must defend your organization from attackers who are making use of it today.

The good news is that there are guardrails you can implement now, including the Menlo approach to addressing the issues GenAI introduces. While there are many different directions from which to approach the challenge of using GenAI securely, Menlo can be enabled in days, requires no costly rip-and-replace, and is compatible with the tools and technologies you probably have in place today, including your browser.

The Menlo solution is built around our unique Secure Cloud Browser, which functions like a hardened digital twin of users' local browsers, executing all active content in the cloud. This has never been more significant than today, as dynamic webpages contain many elements pulled from different sources throughout the browsing session. Attackers are using AI to rewrite elements like JavaScript to evade detection. The Secure Cloud Browser scans traffic continuously and automatically, and patented Adaptive Clientless Rendering ensures seamless performance.

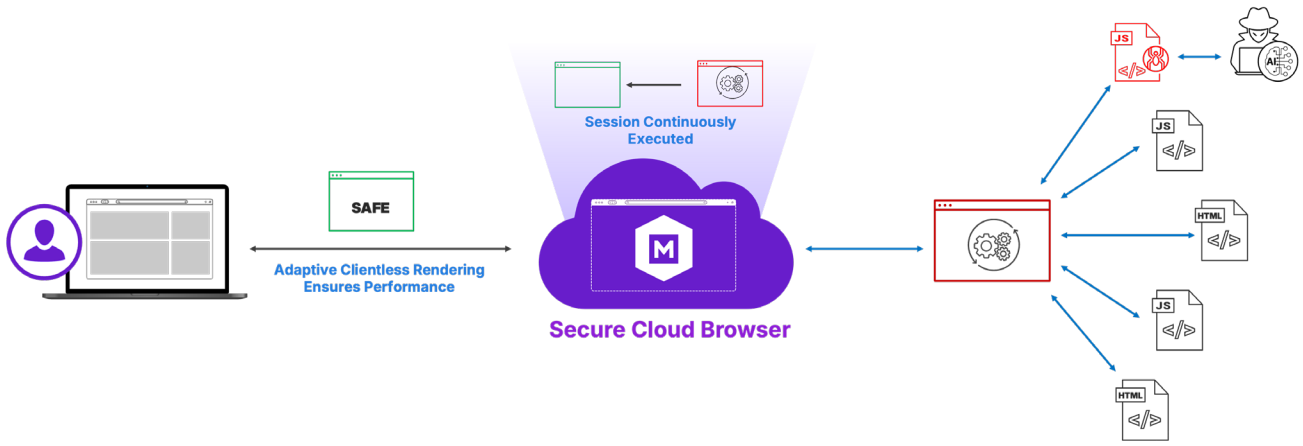

Figure 1: The Menlo Secure Cloud Browser provides continuous protection from malicious web elements throughout the entire browsing session without a performance impact or changes to the endpoint.

Get Rid of "Shadow AI"

Individual users were well ahead of enterprises when it came to the adoption of GenAI tools and services. Many of those folks have been using their own credentials to sign up for GenAI services, typically opting for the free tier. This phenomenon is known as shadow AI, and its use creates thorny issues for the enterprises where these users work.

First, the enterprise loses visibility into and control over essential data and resources in such a situation. Second, most of the free-tier services use data submitted in the form of queries 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.[2]

> "The trouble is that hidden [AI] tools compromise security, break compliance rules, and foster inaccurate or biased outcomes. 3"
>
> – Cloud Security Alliance

In addition to visibility issues, enterprises must contend with attackers' use of AI tools to create convincing attacks and deepfakes, and their ability to mimic well-known AI tool sites. In its research, Menlo observed malicious sites that featured "ChatGPT," "Copilot," or "Gemini" in the domain name—and these fake sites often use the top-level domain ".ai." Expecting end-users to parse a site's validity is far too risky in the age of AI.

Once you and your enterprise security team have identified areas where AI is appropriate, you should select the best tool(s) to fit the need. But that is only part of the solution, because it is equally important to prevent access to other AI tools. To accomplish that goal, you need visibility into enterprise browser traffic.

Once the appropriate tools have been selected and visibility established, it is time to educate users. One of the easiest ways to do that is to redirect users to the sanctioned tool when they navigate to any site classified as GenAI. A notification can alert the user of the organization's safety precautions and the reason for the redirect.

During this period, it is helpful to have a way to look more deeply into web traffic to ensure that all users are compliant.

[2] TELUS Digital Digital Experience Survey, February 2025

[3] AI Gone Wild: Why Shadow AI Is Your IT Team's Worst Nightmare, Cloud Security Alliance, 3/4/2025

How Menlo Can Help

The easy answer to this problem is to use a secure web gateway (SWG) to block access to sites classified as GenAI. Menlo goes a step further, automatically redirecting users to sanctioned GenAI tools as they browse, and educating them on use of the new tool with customizable banner pages. Menlo Browsing Forensics can be set to record all traffic categorized as GenAI, and provides a video-like record of the full user session.

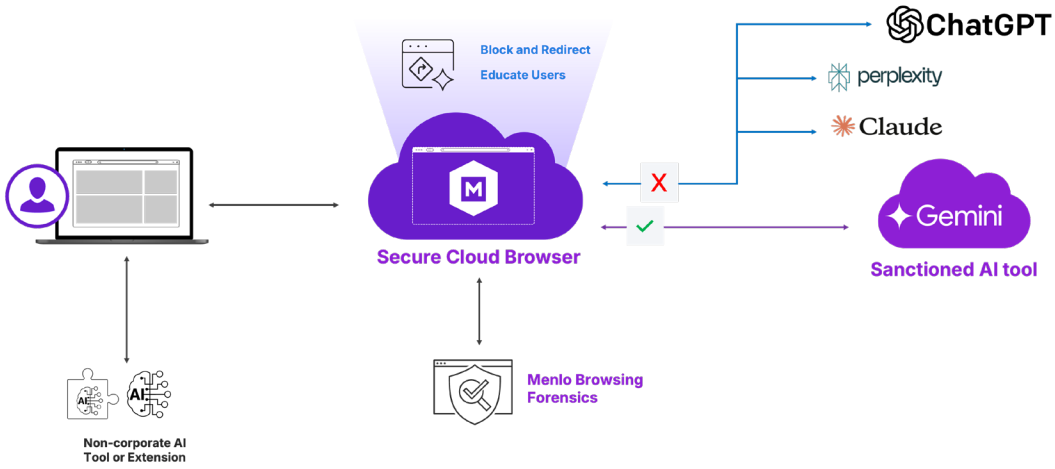

Figure 2: Automatically block access to non-sanctioned sites in the Menlo Secure Cloud Browser as you educate users. Enforce compliance with visibility from Menlo Browsing Forensics.

The reality, however, is that AI as a category is iterating so quickly that no classification system can really keep up. That's why it is vital to go beyond the tool itself to consider what user behavior should be allowed and to ensure that data is automatically protected.

Go Beyond Sanctioned Tools to Sanctioned Use

Much of the AI use that results in data leakage is unlikely to be malicious, but is instead a case of users trying to save time and improve productivity. This situation mirrors the "insider threat" scenario that Menlo has reported on before.

You can safeguard the use of GenAI tools by applying the same methods you would use to prevent insider threats to internal applications. By viewing GenAI sites as sensitive applications themselves, you can protect data and assets by blocking content by file type and limiting or prohibiting copy-and-paste functions. Even watermarking is helpful to raise user awareness about how to treat the sensitive information users may be handling.

The key here is that these rules must be enforced automatically. Once again, since most user interaction occurs in the browser, this is the logical place to build in such enforcements. By treating GenAI like any other application that requires limitations for safe use, many of these requirements can be met simply.

How Menlo Can Help

Menlo content form rules, including copy/paste, upload/download, watermarking, and character limits can be enabled in just a few clicks. This functionality is built into the Menlo Secure Cloud Browser, and serves to remind users to pay attention to prevent data leakage. Visibility can be provided via Menlo Browsing Forensics to help with enforcement.

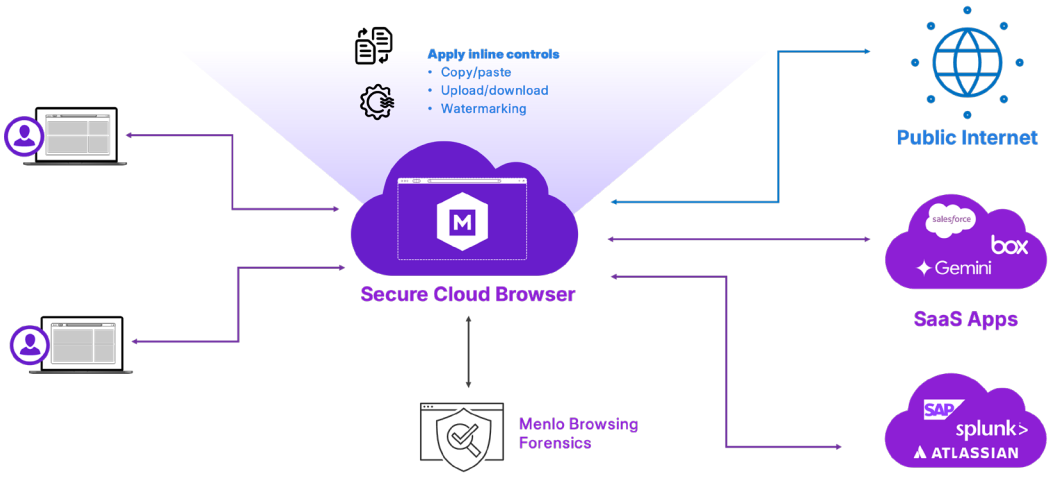

Figure 3: Enforce policies automatically and inline with the Menlo Secure Cloud Browser, with enforcement and visibility from Menlo Browsing Forensics.

Safeguard Your Data

If you have not yet instituted and begun to enforce easy-to-use DLP tools, you need to start. If users are uploading content to GenAI tools, particularly in forms like PDF, which often contain internal or otherwise sensitive data, the results can be disastrous, even if the exposure is unintentional. One way to provide protection for these potentially sensitive documents is to forbid uploads altogether. Another method is to customize DLP rules to flag sensitive information.

How Menlo Can Help

Menlo comes with over 300 DLP dictionaries that you can apply to content with a few clicks. You can also build your own dictionary to provide still more customized DLP controls. These features are built into the Secure Cloud Browser, making it easy to provide protection where it matters most. If content that is being uploaded to a GenAI site triggers a DLP event, you can automatically block the upload or simply log it.

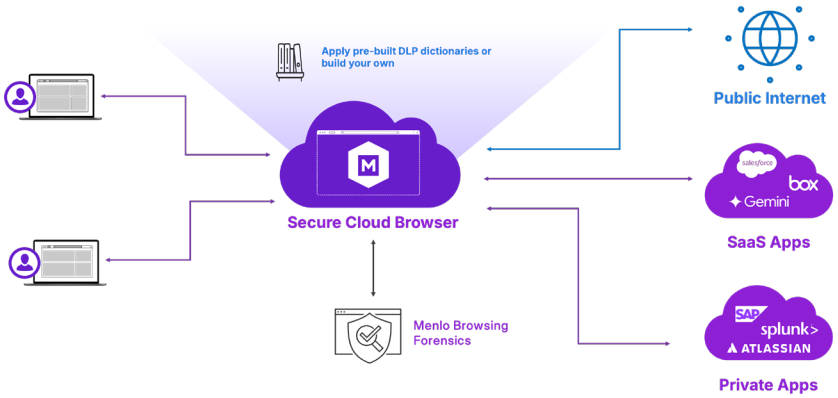

Figure 4: Enforce DLP policies automatically and inline with DLP dictionaries in the Secure Cloud Browser, with enforcement and visibility from Menlo Browsing Forensics.

Enable Content Inspection

If you decide to allow downloads from AI, it is essential to ensure that users are not also bringing in malware. The ability to process active scanning on files to determine that no malware or malicious active content is being included can help to keep the enterprise safe.

How Menlo Can Help

The Menlo Secure Cloud Browser automatically inspects all content unless you choose to disable it. Protections include file hash checks, antivirus, and sandboxing, and are applied to all content, including password-protected files and archives. If content does not pass inspection, you can choose to block it or render it safe to use by disabling active content.

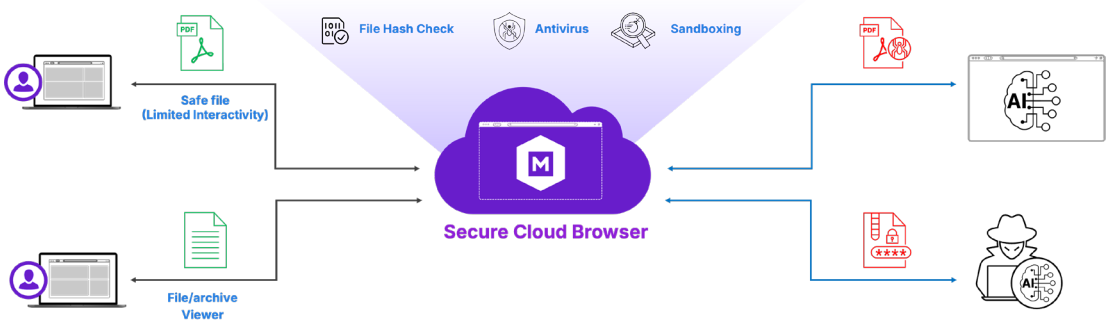

Figure 5: Content inspection is automatic and inline, even for encrypted files and archives. Content that fails inspection can be blocked or rendered "safe."

If retaining full functionality is important, Votiro by Menlo leads in the content disarm and reconstruction category. Once connected, all content is automatically inspected and sanitized unless policy dictates otherwise.

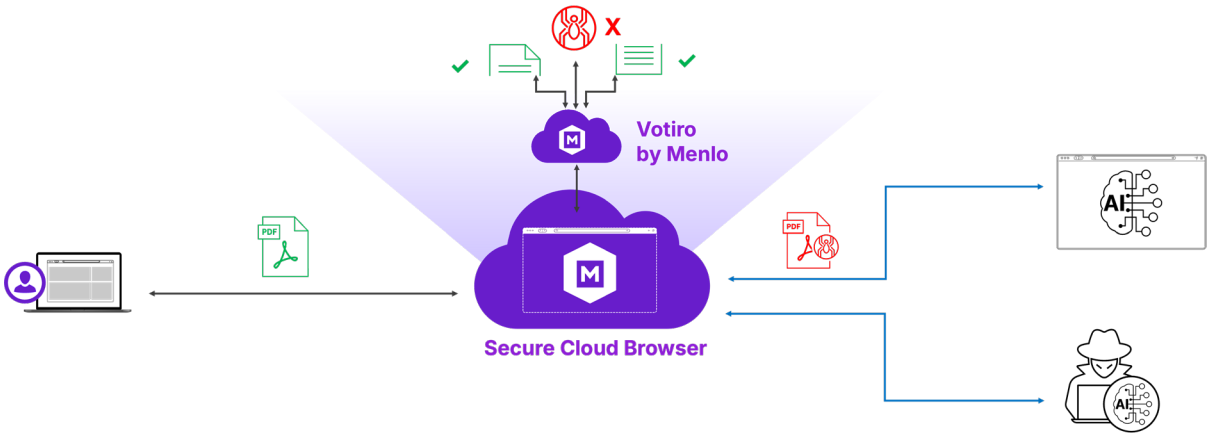

Figure 6: Votiro fully disassembles content, then "defangs" and reassembles it, rendering it both fully functional and safe.

Protect the Enterprise From the Malicious Use of AI

GenAI has made it easier than ever for attackers to get a foothold into the enterprise. Protection from fake AI tools is a good reason to establish and enforce an acceptable use policy, but protections need to go deeper. It is futile to ask users to take on the task of parsing today's phishing scams themselves, as this content has been designed specifically to convey legitimacy, and is often backed by painstaking research of the target, deepfake content, and more. Users must be safeguarded from such threats, and protection must be automatic.

How Menlo Can Help

Today's phishing attacks are more dangerous than ever before. With the help of AI, common "telltales," such as bad spelling or grammar, can be eliminated. It is now possible to rapidly iterate on these threats, increasing the volume of attacks until something gets by defenses. Today's phishing attacks can originate from SMS, social, or even voice platforms, rendering traditional email-only defenses inadequate.

AI has also enabled the creation of highly personalized spear phishing attacks, combining deepfake audio and video personalized with a complete view of the target's life.

Menlo HEAT Shield AI works automatically to catch these threats in real time, inline, before they can get to the user's endpoint, regardless of the source. Because it does not rely on signatures (which can only catch threats that have been seen before), HEAT Shield AI can stop all phishing attacks, even those that are completely new.

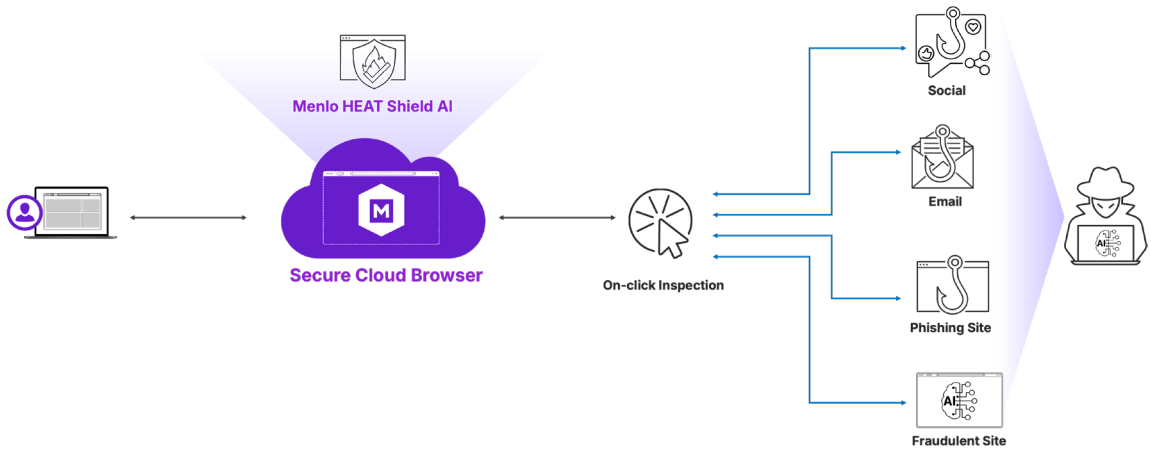

Figure 7: Menlo stops AI-driven threats with HEAT Shield AI, built to capture highly evasive adaptive threat (HEAT) attacks like these and stop them the moment the user clicks.

Protect Internal Applications While Enabling Zero Trust Access

It is vital to consider the flipside of data leaving the enterprise and look at possible threats embedded in content coming in, particularly in the case of unmanaged devices. If a user uploads personal information to a résumé site, for example, there may be more to the returned PDF than a nice layout.

While you can enable and enforce protections for your employees on managed devices, it is impossible to fully control what happens on a BYOD endpoint or a third-party's device. At the same time, these users need to get to your sensitive internal applications to do their jobs. It is essential to provide a zero trust access method that combines security and productivity.

How Menlo Can Help

Menlo Secure Application Access simplifies the process of enabling least-privileged access for all users, even if the device is not owned or managed by the enterprise. With Menlo, any malware that may be present is automatically blocked by the Secure Cloud Browser when the user accesses internal or SaaS apps. As GenAI grows in popularity, so too does the chance that malware could show up on devices that are not managed; with Menlo, you can enable access without worrying that malware will come along.

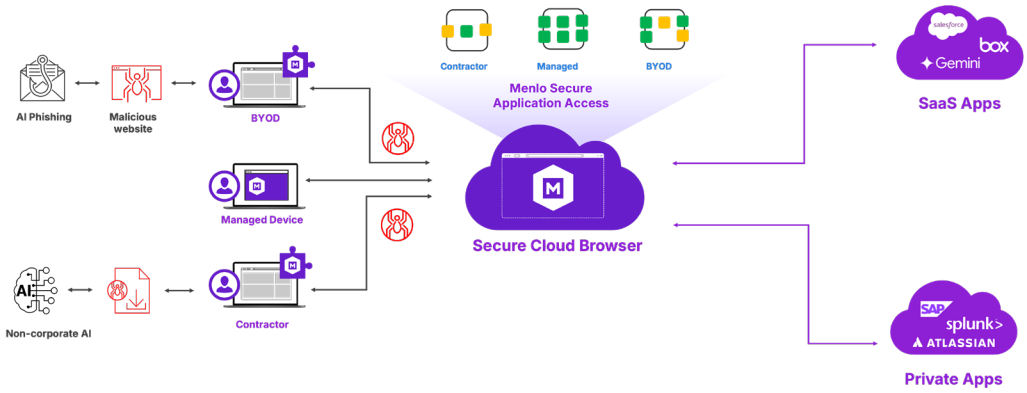

Figure 8: Provide zero trust access for all users and all devices without endangering your internal or SaaS applications.

Don't Forget the Local Browser

While the Secure Cloud Browser provides significant protections, today's leading enterprise browser vendors, Microsoft Edge and Google Chrome, are continuously innovating; a huge reason for their continued popularity. Menlo Browser Posture Manager allows users to choose their own browser type, and use it safely, completing a defense-in-depth strategy.

Current versions of today's enterprise browsers already include AI features, and more will come. Like most new features, these are typically enabled by default. This means you must look more deeply into policies to ensure that new features operate in accordance with your overall enterprise security stance. Your enterprise might choose, for example, to allow GenAI writing assistance, but prohibit this information from being shared back with the public model.

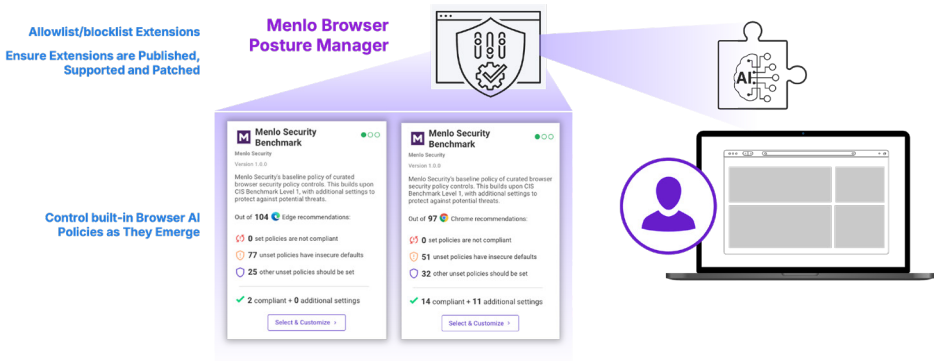

Figure 9: The safe use of GenAI must include the local browser. Menlo Browser Posture Manager makes it easy to manage policies and allow or block extensions, even before security regulations catch up.

In Summary

Menlo has been in the business of securing the world's top enterprise browsers for over a decade. It offers a wide range of security features, allowing you to enable the ones that are a best fit for your enterprise and further tailor them by users or groups. Many of the most important elements of using GenAI safely are built into the solution, applied automatically and inline, and are available at no additional charge. All features are controlled via a single interface, so you don't need to jump between admin screens to get exactly what you want.

With Menlo you can:

- Secure all web traffic for all users, including those on unmanaged devices.
- Automatically provide DLP features, including content form controls and DLP rules, which can be provided via over 300 built-in DLP dictionaries. You can also build your own DLP dictionary to provide even more customized control.
- Fully automatic content inspection, including file hash checks, antivirus, and sandboxing, even for password-protected files and archives.
- Content disarm and reconstruction features from Votiro by Menlo.
- HEAT Shield AI, which protects from the growing threat of evasive, adaptive phishing attacks.
- Secure Application Access, which enables zero trust, least privileged access to internal applications from any authorized user on any device. Content is intermediated by the Secure Cloud Browser, so your enterprise remains protected.

GenAI begins in the browser. Secure it there, with Menlo Security.

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.

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