MENLO SECURITY

Right-Sized Browser Security:

A Buyer's Guide to Protection in Every Session



GUIDE

Executive Summary

You already know that enterprise users spend most of the day in the browser. You know that the browser is where phishing and social engineering exploits begin, leading to ransomware and more. You know that data can leak out of the browser in simple copy/paste GenAI activities, typically run by users trying to boost their productivity. Browser security has been an issue for years, and the fact that it persists may lead you to believe the problem is overstated or already handled. More often, the real reason it persists is something different.

In this Buyer's Guide, we will examine issues that may have prevented you from deploying a browser-specific solution as well as some questions to ask prospective vendors. In the end, you will be able to answer why the browser security gap persists and identify concrete next steps you can take to close it.

Everyone's Job Touches the Browser, but No One Function Owns It

Most enterprises run a security stack that includes the firewall, the Secure Web Gateway (SWG), Data Loss Prevention (DLP) tools and systems, and Endpoint Detection and Response (EDR). The network team typically owns the firewall while the security team takes point on the policies—in smaller organizations, this may all be handled by one group. The SWG historically sits next to the proxy and firewall and is owned by network security or bundled into Security Service Edge (SSE) programs. DLP controls can be highly fragmented, often run in monitor-only mode due to false-positive friction, and are owned by a variety of teams. EDR may be run by the SOC and administered by the endpoint team.

Every element of your security stack is doing what it was built to do. The problem is that while each tool may see part of an issue that is happening in the browser session, none of these tools shows what is happening there from the start to the finish.

A good example is the case of GenAI. Devices at the network/proxy layer, including firewalls and SWGs, can see which AI domains users are going to, but cannot see what is happening on the page. CASBs and SSEs can find sanctioned and unsanctioned AI apps, but only if the apps have already been cataloged, and with thousands of new AI domains being registered every day, it is impossible to keep up. Endpoint tools may be able to see clipboard or file activity, but without context and only on managed devices. As summed up by the 2026 Verizon Data Breach Investigations Report, 67% of users are using non-corporate accounts on their corporate devices to access AI services, while shadow AI has become the third most common non-malicious insider action detected in the Verizon Report's DLP dataset.

Much of that shadow AI risk lies in data that is copied and pasted into prompts or in files that are uploaded or downloaded, which is activity that happens inside the browser session.

An Example You're Living with Today

Another place where the browser security gap takes center stage involves server-side vulnerabilities, such as current PHP issues. Please note that while these specific CVEs are center stage today, this pattern of an exploit with a long tail of fall-out will continue to exist, as it has every time a major vulnerability is discovered and exploited. While larger organizations will patch their servers quickly, the primary danger may not come from there. Instead, picture a smaller company that you have done business with for years, such as an office supply distributor, a payroll/benefits provider, or even a small MSP. These companies often feature servers that are built, hosted, or maintained by others. Such servers could effectively be

turned into a supply chain watering hole. The compromised supplier portal doesn't attack the supplier; it attacks the visitor. This approach is particularly useful for attackers because it exploits your trust in this supplier. You trust their site. You communicate with them regularly, and you expect to get content from them. Let's look at a few examples of what can go wrong.

> The site that your team has trusted for years is the one that your existing filters, your security stack, and your users will not question. That leaves you wide open to subsequent attacks that use the trusted site as a vector.

Example One—The Phish

In this example, the attacker has modified the supplier site to enable it to serve up a fake login overlay page or credential prompt. If the user goes to the portal frequently, they likely go through a bookmark; reputation or URL filters in your SWG will wave this traffic through because the site is trusted and the certificate is valid. The attacker didn't have to lure anyone; the user came on their own.

The issues start when you consider that the site has already been compromised. The site looks normal to the supplier, but every visitor is walking into a trap. One way that this could play out is that as the page loads, the user is presented with an overlay that claims that their session has expired or that they need to log in to view new pricing. The user enters their credentials, and, in the case of an adversary-in-the-middle attack, the malicious page relays the login to the real identity provider. The user sees what they expect to see, and MFA is legitimate—only it is legitimate on the attacker's behalf.

The identity provider then delivers a valid session token and returns it, passing through the attacker's hands. The attacker can now go into the site without a login screen or MFA prompt. If the lure only captures the portal login, this is relatively low value. But effective attacks use the portal as a stage to get a bigger prize. If the overlay page asks for a Microsoft, Google, or SSO login, the credentials being captured are the user's corporate identity. If the overlay asks for OAuth consent, the attacker gets a corporate token. This attack works because all of the content assembles in the browser, against a backdrop of trust. Now consider the types of communications that you get from such a supplier. Most users would not hesitate to open a link to a receipt or go to the portal if a sale was announced.

Example Two—The Embedded Malware

Now imagine the same scenario, where the attacker has compromised the supplier site already. In this case, the attack involves a file, like an invoice or even the sales announcement referenced above. The user downloads the file, which looks exactly like what they expected, not realizing that there is malware present beyond the visible content. When the user opens the file, it looks normal, but malware, such as a loader whose job is to fetch the real payload, runs in the background. The loader could install an infostealer or remote access tool.

The issues around PHP vulnerabilities are just one example. In the wake of every major vulnerability, Menlo Security typically observes:

- **Long-Term Phishing Hubs:** Attackers set up credential harvesters buried deep within subdirectories of high-reputation servers, letting them run undetected for long stretches.
- **Sneaky Malicious Redirects:** Injected scripts continuously hijack legitimate website visitors and send them off to scareware, drive-by downloads, or secondary exploit kits.
- **Search & Advert Poisoning:** By riding on the domain authority of compromised institutional sites, malicious links can easily maintain top search engine rankings and advertisements.

Which Browser Are We Talking About?

One of the most daunting elements of browser security is the sheer number of possibilities that a real solution would have to cover. There are many popular browsers that are available as a managed enterprise option and a consumer option, and there are many different versions of each. Browser vendors are constantly patching their offerings as well as bringing new features to market. Complicating matters still further, users tend to be loyal to the browsers that they have been using, making a change difficult to manage.

What Are Your Options?

Browser security is rapidly moving from a nice-to-have to a requirement, which can be extremely difficult to implement given the huge variety of browsers and devices that you need to consider. If you are still on the fence, or are not satisfied with your current solutions, it can be useful to take a step back and consider some options.

Rely on Native Controls

Chrome Enterprise Core/Premium and/or Microsoft Edge for Business deliver important features, but the limiting factor is that they do so only for their own browsers, and most enterprises run more than one browser. Another factor is that managed browsers do not translate well to BYOD or unmanaged devices, which account for a significant share of users in today's enterprise.

Depend on Existing Security Tools

As we've pointed out earlier, the browser security gap is real. All of your security tools are doing what they were built to do, but the gaps between the tools are the ones that are being actively exploited. It is understandable to want to maximize the benefits that can be derived from existing tools, but not at the cost of enterprise security. The fact is that attackers have moved into the browser because that's where users, corporate credentials, and data meet. They are well aware of the browser blind spot and have been making the most of it for years. Adversarial use of AI has only accelerated this trend, helping attackers scale phishing and ransomware campaigns.

Avoid Browser Security for Fear of Disrupting User Workflows

This is another understandable concern, particularly when you consider slow proxies, broken sites, and latency. Any friction becomes a helpdesk ticket, which is a major reason that enterprises have resisted wholesale browser replacement. As advancements emerge, however, it is worthwhile to take another look.

Adopt a Replacement Browser or Extension

Replacement browsers claim to deliver a suite of protections, but can face an uphill battle due to complex deployment and user pushback. Browser extensions that promised to deliver security combined with convenience offered compelling benefits, but many standalone products have been acquired by larger companies and absorbed into their overall security products. This makes considering them a more involved process than it was even a year ago.

Why Act Now?

Given these issues, including the reasons you might be tempted to do nothing, let's look at the situations that make now the time to act.

A Trusted Supplier Site That May Already Be Compromised

The PHP example earlier in this guide is not a one-off. Every major server-side vulnerability leaves behind a long tail of sites that stay unpatched long after larger organizations have moved on. Your users visit these sites because they trust them, and your existing filters let the traffic through for the same reason. URL reputation, certificate checks, and category filtering all confirm that the site is the one you expect. None of them can tell you that the page has been changed.

Contending with these attacks means judging what a page does rather than where it comes from, and making sure that files from trusted sources are safe to open. Some typical issues in such situations include:

- **Fake login overlays and adversary-in-the-middle exploits:** These types of exploits take advantage of user trust to prompt for a login that might not be needed. The attacks work differently but both share the goal of tricking the user into entering their credentials. These exploits can be combined with an overlay page that sends the user to an adversary-in-the-middle, (AiTM) proxy.
- **Long-term phishing hubs, malicious redirects, and poisoned search results:** These exploits, which can linger on long after the initial attack, can essentially turn a trusted site with valid credentials into a supply chain watering hole. These exploits take advantage of the fact that users are accustomed to getting regular communications from such sites.
- **Weaponized invoices and announcements:** This attack, like the others above, uses the victim's trust against them. In these cases, the malware is buried inside a file that corresponds to what the user would expect to get from a supplier. Such malware is usually invisible to the human eye, but can have devastating results.

Upcoming Cyber Insurance Renewal

2026 cyber insurance applications typically run 12–20 pages or around 40 security questions, and carriers increasingly want proof rather than checked boxes. Questions focus on phishing-resistant MFA, EDR, backups, privileged access, email and web filtering, asset and software inventory, and security monitoring. Browser controls rarely appear by name, but they help provide evidence for several categories:

Questionnaire Category	Browser Security Control That Provides Evidence
Phishing and email/web threat protection	Real-time, intent-based phishing detection (HEAT Shield AI)
Web filtering	Domain- and category-based web governance
Hardware and software inventory	Enterprise-wide browser extension inventory with current risk scores
Data loss prevention	In-session controls: copy/paste, upload/download, print, screenshots, AI prompts
Malware protection	File security: hash checks, AV scanning, cloud sandboxing, CDR
Security monitoring and logging	Identity-tied session logs, searchable and streamed to SIEM

Upcoming SOC 2, HIPAA, or GDPR Audit

Regulatory audits are always taxing, but visibility into browser activity can provide evidence that is hard to gather any other way. Note that browser security provides evidence for these controls; it is not an all-in-one compliance solution.

Compliance Requirement	Framework References	Browser Security Control That Provides Evidence
Phishing and web threat protection	SOC 2 CC6.6 HIPAA 164.308(a)(5)(ii)(B) GDPR Art. 32	Real-time, intent-based phishing detection (HEAT Shield AI)
Malware protection	SOC 2 CC6.8 HIPAA 164.308(a)(5)(ii)(B) GDPR Art. 32	File security: hash checks, AV scanning, cloud sandboxing, Menlo File Security with CDR
Data movement and loss prevention	SOC 2 CC6.7 HIPAA 164.312(e) GDPR Art. 5(1)(f), 25	In-session controls: copy/paste, upload/download, print, screenshots, AI prompts
Unapproved vendors and shadow AI	SOC 2 CC9.2 HIPAA 164.308(b) GDPR Art. 28, 30, 44	Domain- and category-based web governance, plus AI prompt controls

Compliance Requirement	Framework References	Browser Security Control That Provides Evidence
Risk analysis and software inventory	SOC 2 CC6.1–6.8 HIPAA 164.308(a)(1) GDPR Art. 32	Enterprise-wide browser extension inventory with risk scores
Audit logging and activity review	SOC 2 CC7.2, CC7.3 HIPAA 164.312(b), 164.308(a)(1)(ii)(D) GDPR Art. 5(2)	Identity-tied session logs, searchable and streamed to SIEM
Incident scoping and breach notification	SOC 2 CC7.4 HIPAA Breach Notification Rule GDPR Art. 33, 34	Searchable session logs to establish what was accessed or exfiltrated

New Acceptable Use Policy for GenAI with No Means of Enforcement

Although many enterprises have developed acceptable use policies for GenAI, such controls are notoriously difficult to enforce. To be effective, these policies need enforcement and detection mechanisms that capture user behavior where it starts: in the browser.

GenAI AUP Requirement	Where Other Security Tools Fall Short	Browser Security Control That Enables It
Discover which GenAI tools employees use	SWG, firewall, and CASB logs show domains, not what users do on the page. They lag on newly launched AI tools.	Identity-tied session logs showing which users reach which AI tools, plus category-based web governance
Allow approved AI tools, restrict unapproved ones	Domain blocking is all-or-nothing and pushes users to workarounds or personal devices	Domain- and category-based web governance for sanctioned and unsanctioned AI tools
Keep sensitive data out of prompts	Network DLP sees encrypted sessions to allowed domains and can't reliably read typed or pasted prompt text. Endpoint DLP lacks in-page context.	In-session controls on copy/paste and AI prompts

GenAI AUP Requirement	Where Other Security Tools Fall Short	Browser Security Control That Enables It
Restrict file uploads to AI tools	Proxies and CASBs struggle to tell an upload to an AI tool from routine web traffic	In-session upload/download controls, plus file security scanning
Govern AI browser extensions	EDR and MDM have limited insight into browser extensions and the page data they can read	Enterprise-wide browser extension inventory with risk scores
Protect users from fake AI sites and AI-themed phishing	URL reputation misses newly registered lookalike domains	Real-time, intent-based phishing detection (HEAT Shield AI)
Monitor compliance and prove enforcement	Network logs record connections, not user actions, so they can't show the policy was followed	Identity-tied session logs, searchable and streamed to SIEM

The Visibility You Need to Complete a Stalled Investigation

Every enterprise has run into situations like these. Perhaps there has been a breach, but you have no idea where it began. Maybe you suspect that a credential harvester has been at work, but you cannot see it, or your endpoint tools show a mysterious file on user endpoints, but you cannot tell where the file came from. Each reflects an activity that may well have begun in the browser.

SOC Investigation Question	Where Other Security Tools Fall Short	Browser Security Control That Answers It
Was the user phished, and what did they actually see?	Email and SWG logs show a click and a URL, not the page that rendered. A compromised but trusted site has a good reputation and a valid certificate, so the logs look clean.	Real-time, intent-based phishing detection (HEAT Shield AI), plus identity-tied session logs that record the page and the user
Did the user enter credentials or grant OAuth consent?	Identity provider logs show a successful login with MFA passed. In an adversary-in-the-middle attack, that looks legitimate, and nothing shows the page where the user typed.	Identity-tied session logs that place the user on the phishing page, streamed to SIEM to line up with identity provider sign-in events
Where did this file come from?	EDR sees the file on disk and the process it launched, but not the site, page, or session that delivered it.	Session logs that record the download with the user and source page, plus file security results (hash checks, AV scanning, cloud sandboxing, file security with CDR)

SOC Investigation Question	Where Other Security Tools Fall Short	Browser Security Control That Answers It
What data left, and where did it go?	Network DLP can't reliably read encrypted sessions to allowed domains. Endpoint DLP sees clipboard or file activity without page context, and only on managed devices.	In-session controls on copy/paste, upload/download, print, screenshots, and AI prompts, logged by user
Was a browser extension involved?	EDR and MDM have limited insight into browser extensions and the page data they can read.	Enterprise-wide browser extension inventory with risk scores
Who else visited the same site?	Proxy logs show connections, not sessions. Traffic that doesn't route through the proxy leaves no record.	Searchable, identity-tied session logs across all users to scope the blast radius
Can we put the browser evidence next to everything else?	Browser activity either isn't captured or sits in a separate console, so analysts can't build one timeline.	Identity-tied session logs streamed to SIEM alongside EDR, identity, and email telemetry

Introducing the Menlo IRIS Browser Extension

Menlo Security has been providing browser security to some of the world's largest and most security-conscious organizations for over a decade. The Menlo IRIS Browser Extension encapsulates all of the vital capabilities that you need in a form factor that is easy to manage and easy to deploy across the browsers that your users already work with—even those on BYOD or unmanaged devices. If your risk profile changes, or if you determine that you need to add isolation for high-risk data or high-profile users, Menlo has you covered with additional offerings that deliver precisely what you need.

Menlo IRIS also comes with unique features to quickly implement specific security benefits, including HEAT Shield AI for protection from zero-day phishing and social engineering exploits, as well as File Security with Content Disarm and Reconstruction (CDR) to deliver files that are malware-free and fully functional. These capabilities are vital when you are seeking to counter the "long-tail" effects of a widespread vulnerability such as the recent CVEs affecting PHP. Regardless of the high profile and urgency of vulnerabilities like these, remember that there are inevitably companies that lag in patching. If such a company is within your supply chain, particularly if you have done business with the firm for so long that they are implicitly trusted, you can easily become a victim.

Considerations and Questions to Ask Every Vendor

As the browser becomes increasingly powerful and omnipresent, the need to secure it will only grow. Even if you are not ready to consider dedicated browser security today, it is worth examining what your existing security stack cannot see.

Some questions to consider:

- What does your existing security stack see inside the browser session?
- Does your browser security solution depend upon routing traffic through your proxy?
- Which browsers and device types does it support, including BYOD?
- **Insurance:** Can your logs serve as evidence for your cyber insurance application?
- **Insurance:** Which questionnaire controls (phishing protection, web filtering, software inventory, monitoring) does your product help us demonstrate?
- **GenAI:** Can you enforce our GenAI acceptable use policy at the prompt and upload level, including for AI tools that launched last week?
- **Investigation:** Can your logs show what a user saw, where a file came from, and who else visited the same site?
- **Compliance:** Which SOC 2, HIPAA, or GDPR controls can your logs and in-session controls help us evidence?

Evaluation Checklist

Use this checklist to compare vendors against the criteria covered in this guide. The Menlo IRIS column shows how IRIS addresses each one. Use the Vendor A and Vendor B columns to score the other solutions you evaluate.

Evaluation Criterion	What to Look For	Menlo IRIS	Vendor A	Vendor B
Protection from compromised trusted sites	Detects malicious pages and files on sites with a good reputation and a valid certificate, such as a compromised supplier portal	Real-time, intent-based phishing detection with the HEAT Shield AI add-on, plus content disarm and reconstruction (CDR) with the Menlo File Security add-on to deliver safe, usable files		
Phishing protection	Real-time detection based on page intent, not only URL reputation. Catches phishing on trusted or compromised sites and newly registered domains.	Real-time, intent-based phishing detection with the HEAT Shield AI add-on		

Evaluation Criterion	What to Look For	Menlo IRIS	Vendor A	Vendor B
Malware and file security	Multiple layers of file inspection before a download reaches the user	File security with hash checks, AV scanning, and cloud sandboxing. Content disarm and reconstruction (CDR) with the Menlo File Security add-on.		
In-session data controls	Policy enforced on user actions inside the page, not just on the domain	In-session controls on copy/paste, upload/download, print, screenshots, and AI prompts		
GenAI governance	Allows sanctioned AI tools, restricts unsanctioned ones, and keeps sensitive data out of prompts	Domain- and category-based web governance for sanctioned and unsanctioned AI tools, plus AI prompt controls		
Browser extension governance	Visibility into every installed extension and the risk each one poses	Enterprise-wide browser extension inventory with risk scores		
Logging and SIEM integration	Logs tied to user identity, searchable and exportable to existing SOC tools	Identity-tied session logs, searchable and streamed to SIEM		
Investigation support	Shows what the user saw, where a file came from, what data left, and who else was affected	Searchable, identity-tied session logs across all users to reconstruct activity and scope the blast radius		
Compliance and insurance evidence	Controls and logs that map to SOC 2, HIPAA, GDPR, and cyber insurance questionnaire categories	Session logs and in-session controls provide evidence for phishing protection, malware protection, DLP, software inventory, and monitoring requirements		
Browser and device coverage	Works across the browsers your users already run, including on BYOD and unmanaged devices	Browser extension that works in the user's existing browser, including on unmanaged devices		
Architecture and user experience	No requirement to route all traffic through a proxy. No added latency or broken sites.	Extension-based delivery that runs locally in the browser, with Menlo Cloud available where needed		
Growth path	A way to add stronger protection without replacing the platform	Upgrade path to full browser isolation for high-risk users and data		

Closing the Browser Security Gap

The browser security gap persists not because the risk is unclear, but because no single team or tool owns the full session, and attackers know it. Whether your trigger is a trusted supplier site that may already be compromised, an upcoming insurance renewal or audit, a GenAI policy you can't yet enforce, or an investigation that stalled at the browser, the first step is the same: find out what your current stack cannot see. Use the questions and checklist in this guide to pressure-test every vendor, including the ones you already have. Right-sized browser security starts with protection in the browsers your users already work in, without rerouting traffic or disrupting workflows, and scales to isolation as your risk profile changes. To see how the Menlo IRIS Browser Extension closes the gaps in your environment, contact Menlo Security to schedule a demo.

About Menlo Security

Menlo Security is the pioneer of the Browser Security Platform (BSP), the industry's first infrastructure designed to govern a hybrid workforce of humans and autonomous AI agents. By centering the browser as the new enterprise operating system, Menlo provides a "Guardian Runtime" that resolves the unique risks created when AI agents operate at machine-speed without human skepticism. The Menlo platform enables the agentic enterprise to scale AI with confidence, providing universal connectivity to legacy data and unified zero-day threat prevention across every session. Trusted by over 1,000 global enterprises—including eight of the ten largest financial institutions and major government agencies—Menlo protects over 8 million users and millions of simultaneous AI agent sessions. Headquartered in Mountain View, California, and backed by investors including JPMorgan Chase, American Express Ventures, and Vista Equity Partners, Menlo is securing the browser for the agentic age. Learn more at www.menlosecurity.com.

© 2026 Menlo Security, All Rights Reserved.

MENLO SECURITY

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