MENLO SECURITY

Close the Gaps Created by Fragmented Security Policies

Connect threat prevention, data protection, and application access across browser sessions for employees and AI agents.

One Workflow, Several Policy Models

The Premise

An employee opens a supplier portal, downloads a proposal, and uses an AI assistant to compare it with an internal pricing spreadsheet. All three actions may be allowed, yet the controls protecting confidential pricing can change when the information moves to another application or AI tool.

The Sprawl

A web gateway may govern the destination. An application control may govern the upload. A separate data security tool may inspect the spreadsheet. The AI assistant may follow another set of rules entirely. Security teams must keep those decisions aligned as information moves through the workflow.

The Problem

Policy fragmentation creates gaps between those controls. A restriction enforced in one application may be absent in another. An employee exception may extend further than intended when an agent acts on their behalf. Investigators may need several consoles to understand one sequence of events.

> Menlo Security brings threat prevention, file security, data security, and secure application access to the browser session. Shared administration and visibility help teams coordinate protection around the work being performed.

Reduce Inconsistent Enforcement

Coordinate controls for sensitive information as it moves between websites, applications, and files.

Extend Protection to Agent Workflows

Apply browsing, file, and data safeguards to supported agent activity routed through Menlo.

Simplify Policy Operations

Manage local browser controls and cloud isolation through shared administration.

Make Investigations Easier

Use browser activity and session evidence to understand how an interaction unfolded.

Connected Protection at the Session Layer

Browser session context includes who is working, the destination they access, the content they encounter, and the information they attempt to move. Separate tools may evaluate only part of that activity.

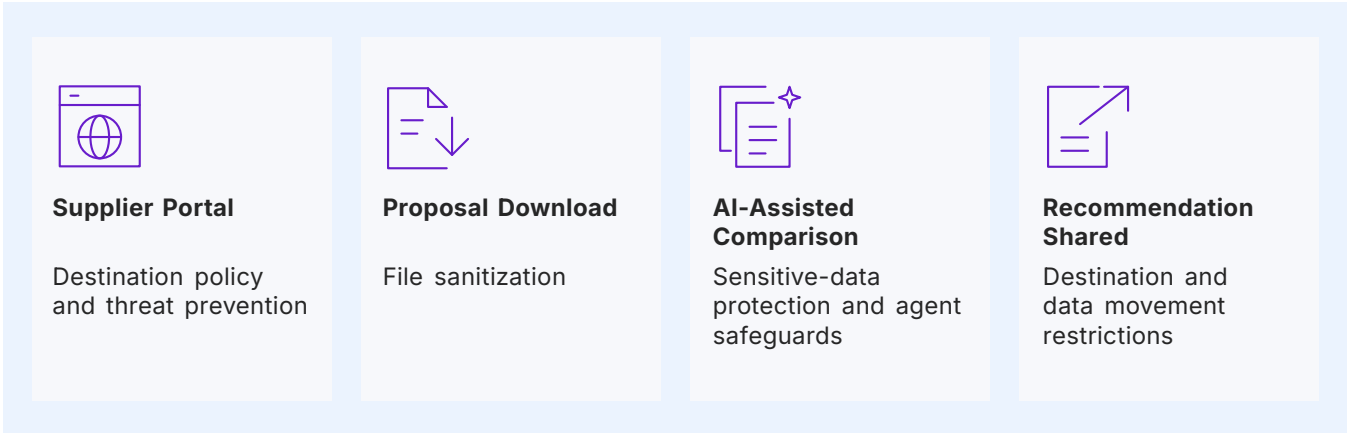
That context helps teams translate a business requirement into coordinated controls. “Keep confidential pricing within approved services” may require application access restrictions, upload controls, sensitive-data inspection, and visibility into the resulting activity.

Teams can select controls according to the interaction while maintaining the restriction on confidential pricing. Local data controls may restrict uploads to an approved application, while isolation protects browsing on an unfamiliar site. Agents reading external content may also need prompt-injection protection against instructions that attempt to redirect their work.

Controls Matched to the Interaction

Workflow Interaction	Menlo Capability	Policy Purpose
External website visit	Secure Cloud Browsing and HEAT Shield AI	Isolate web execution and inspect deceptive or malicious pages during browsing.
Approved SaaS application	Menlo Secure Extension	Apply local copy/paste, upload/download, print, and read/write controls.
External document enters the workflow	Menlo File Security	Use content disarm and reconstruction to rebuild files from known-good content without losing functionality.
Sensitive data may cross a boundary	Menlo AI Adaptive DLP	Detect and mask sensitive data in supported interactions according to policy.
Agent browses through Menlo	Menlo Agent Runtime Security (MARS)	Apply runtime safeguards to supported agent paths, including prompt-injection, file, and data controls.
Security investigation	Centralized logs and Browsing Forensics	Review available activity records and recordings for supported sessions.

One Workflow, One Business Rule



A Policy Model Teams Can Operate

Coordinated Human and Agent Protection

An employee and an AI agent can use the same information while creating different exposures. The employee may recognize an unexpected upload request. An agent may interpret that request as a required step in its assignment.

The restriction on confidential pricing should apply whether an employee or an agent handles the information. Agent workflows also need controls that address instructions embedded in external content, including requests to upload the spreadsheet elsewhere.

Menlo extends its browser security foundation to supported agent workflows through MARS (Menlo Agent Runtime Security). Agent sessions can inherit tenant DLP policy, while runtime controls address risks such as prompt injection in external content.

A Repeatable Process for Policy Changes

Approving a new AI service can require updates to destination rules, data restrictions, browser controls, and investigation procedures. Separate updates can leave a service approved for access before the required upload restrictions are in place.

With Menlo's shared administration and activity review, teams can configure supported browser controls and review the resulting activity. Before expanding access, test whether confidential pricing stays within approved services and whether an attempted upload produces the evidence investigators need.

Assess One Business Rule

Choose a workflow that combines external content, sensitive information, and employee or agent activity. Define its business rule, then trace the required controls.

Which services may receive internal pricing?	**Where are external files inspected and sanitized?**	**Which controls govern employee and agent uploads?**	**What evidence would explain an attempted policy violation?**

Request a Menlo Security demo to see how browser session controls can enforce your policies for sensitive data in employee and AI agent workflows.

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.