MENLO SECURITY

Secure by Design

Zero Trust Access for the Human and Agentic Enterprise

WHITE PAPER

The access model most enterprises rely on today has a gap that VPNs and VDI were never designed to close. Network-level controls secure the path to an application. They have almost no visibility into what happens inside it.

That distinction matters more than it used to. Your workforce now includes employees on personal devices, contractors on unmanaged endpoints, and AI agents operating at machine speed—all of them working primarily through a browser. The inside of that browser session is where your most significant exposure lives, and it's exactly where your current access architecture goes dark.

This paper makes the case for a different model: Zero Trust access that starts in the browser, covers every user type, and gives both your security and compliance teams the visibility to prove it's working.

Access Policy Isn't One Problem. It's Three.

Your access model must account for fundamentally different user types, each with a distinct risk profile. Treating them as a single population is where most access architectures introduce their first gap.

Managed Users

Managed users need access to internal applications, but not uniform access. A sales rep and a finance analyst may both use the same web-based project management application, but what they can see, edit, and export should be entirely different. Least-privileged access within managed user populations matters as much as the boundary between managed and unmanaged endpoints.

A role-based model that enforces differentiated permissions inside an application, and not just at the door, closes. the lateral exposure that flat access policies leave open.

Third Parties

Contractors, business partners, and M&A counterparties need access to specific applications for specific tasks. The risk is clear, and includes access that's too broad, and controls too coarse to adjust without disrupting everyone else.

What makes third-party access uniquely dangerous isn't any individual user. It's the supply chain behind them. When a business partner has access to your internal applications, so does their security posture, their devices, and their own vendor relationships. A single partner compromise can propagate across hundreds of organizations simultaneously. Access for third parties has to be narrow, specific, and easy to revoke the moment the relationship changes.

Agents

AI agents don't behave like humans, and your access model has to account for that. Three characteristics make them a distinct risk category:

- They scale faster than humans. In many enterprises, the number of agents accessing internal resources will quickly exceed the number of human users.
- They operate at machine speed. Controls designed for human-paced interaction face a different stress test when agents can iterate thousands of times per minute.
- They're obedient by design. Agents follow instructions, including malicious ones injected through a compromised session. The same quality that makes them useful makes them exploitable.

An access model that wasn't designed for agents is already inadequate for a significant portion of the sessions hitting your applications.

Example: Role-Based Access in a Project Management Application

A single project management app serves multiple user types with entirely different access needs:

- Finance needs read access to project costs and timelines
- Marketing needs access to content and integration with creative tools
- Program managers need a full view into progress and timelines.
- Agents may need narrow access to identify and propose quick tasks for which they can help.

Each role should see only what it needs to complete its task. An access model that can't enforce that distinction at the application level is only approximating Zero Trust rather than enforcing it.

Device Type Changes the Risk Profile

User role is one dimension of a least-privileged access policy. The device being used is another. Even when a user's role is consistent, the endpoint they're working from can change the risk calculation significantly.

Managed Devices

Company-issued, IT-maintained devices give you a meaningful baseline. You can enforce patch discipline, apply threat protections, and maintain consistent permissions. For high-sensitivity applications and roles, managed devices remain the right foundation.

BYOD

Personally owned devices introduce a different category of risk. Even with device enrollment, IT may not be able to control patching cadence or password policy, and may have limited visibility into what else is running on the device. When a device is shared—or simply used for personal activities—the exposure compounds: family members visiting risky sites, personal apps that introduce malware, credentials that end up on the wrong machine.

> As analysis on bring your own device (BYOD) and infostealers in the 2025 DBIR shows, even on personal devices, it was relatively common for employees to have corporate account credentials or information that is ripe to be compromised.[1]

The core problem is a gap between access and control. Employees have full corporate access. IT has limited ability to enforce, monitor, or remediate on a device it doesn't own.

Unmanaged Devices

The risk profile for contractor and partner devices mirrors BYOD, with one critical difference in scale. A single vendor compromise doesn't affect one device—it can cascade across hundreds of organizations simultaneously. Browser traffic from unmanaged devices has historically been invisible to most security tools, making it nearly impossible to detect suspicious session behavior before it becomes a problem.

Why Legacy Access Tools No Longer Fit

VPNs and VDI were designed for a world where applications lived on the enterprise network and users were expected to be on-premises or close to it. That world is gone. Both technologies got a boost when the pandemic forced rapid remote access deployment, but the underlying architectural mismatch only widened as the many trends collectively known as digital transformation continued.

VDI: the Right Idea for the Wrong Era

The original premise of VDI was sound: manage and secure virtual centrally, and keep data on centralized servers. The execution has proved far more complex. Servers are no longer centralized. The VDI hardware and software stack is burdensome to maintain. Most SOC and security tools don't perform well in VDI environments, requiring specialized monitoring solutions that add cost and deepen vendor lock-in.

For BYOD and third-party users, which constitute the largest enterprise populations that need remote access, VDI is a particularly poor fit. Its centralized architecture introduces performance bottlenecks over variable networks and devices. For contractors and partners specifically, VDI either forces personal devices to effectively become managed corporate endpoints—which users resist—or delivers a degraded, laggy experience that undermines the productivity it was meant to enable.

[1] 2026 Verizon Data Breach Incident Report

VPN: Broader Access Than Needed, Less Visibility Than Required

VPNs give users the experience of being on the network. The problem is that most applications no longer live there. Even with split tunneling enabled, for many organization enterprise applications, traffic routes into the network, up to the cloud, and back—a path that adds latency while creating the security exposure that has made VPNs a preferred initial-access vector for attackers.

The central architectural problem is built into how VPNs work: they were designed to connect users to the entire network, not to specific applications. According to the 2026 Verizon DBIR, 29% of breaches involved unpatched vulnerability exploitation in edge devices including VPNS, building on findings the previous year that edge-device and VPN exploitation rose roughly 8x, from 3% to 22% of initial-access cases.[2]

The agent problem compounds this further. Agents operate at machine speed and can overwhelm network-based infrastructure. The VPN controls that felt adequate for human-paced access become a liability when attackers or compromised agents can iterate at machine speed. 79% of defenders report AI exploitation speed as their top concern, and it's already showing up in the field.[3]

VPN visibility problems, while not as specialized as those found in VDI deployments, are still daunting. Investigations facing VPN encryption may require time-consuming techniques like key acquisition, deep packet inspection, network flow analysis, or network log reviews to gain visibility into sessions.

Zero Trust: the Strategy, and the Gap in How It's Been Implemented

Zero Trust as a framework has been an enterprise goal since it was introduced in 2010. Its core principles—continuous authentication, continuous authorization, no implicit trust—were first applied at the network layer, which made sense when most applications lived there.

That assumption no longer holds. Enterprise applications are cloud-based. Users are everywhere. The boundary that network-level Zero Trust was designed to protect no longer maps to where work happens.

[2] *Verizon 2026 Data Breach Investigations Report*
[3] *https://www.cio.com/article/4149542/ai-machine-speed-is-breaking-vpn-security.html*
[4] *https://blog.gigamon.com/2025/12/02/how-to-implement-a-zero-trust-architecture-and-why-visibility-comes-first/*

It's also worth being precise about what Zero Trust actually is. It's a strategy, not a product, and the market is saturated with solutions that claim the label while delivering only fragments of the framework.[4] Two specific solution categories are worth distinguishing:

- Zero Trust Network Access (ZTNA): secures the path to applications. VPN replacements and network access control tools fall here.
- Zero Trust Application Access (ZTAA): secures what happens inside the application. VDI and browser security tools fall here.

Most Zero Trust deployments address the first category and stop. The session, where your users spend most of their day and where most attacks arrive, goes ungoverned. your users spend most of their day and where most attacks arrive—goes ungoverned.

Adoption Is Widespread. Maturity Is Uneven.

81% of organizations report implementing Zero Trust, with 52% using a mix of tools and 48% citing cost and resource constraints as barriers.[5] The adoption numbers look strong. The underlying reality is more complicated.

Most deployments, including those that are identity-first and those that are network segmentations only, are piecemeal, because the cost and friction of replacing legacy infrastructure pushes organizations toward partial implementation. Uneven coverage creates a false sense of security. Leadership assumes a deployment is global. In practice, it frequently isn't.

The final missing element in most Zero Trust architectures is visibility. Without the ability to see and prove how controls are working, adjustments become reactive and only happen after something goes wrong.

How Menlo Closes the Gap

The Menlo Browser Security Platform operates at the point where most modern work actually happens: the browser session itself. Where traditional ZTNA secures the path to an application but remains blind to what happens inside it, Menlo isolates, inspects, and governs every session in the cloud, regardless of whether the user is an employee, a contractor, or an AI agent.

The difference is similar to knowing knowing someone walked through the door and knowing what they did once inside.

[4] *https://blog.gigamon.com/2025/12/02/how-to-implement-a-zero-trust-architecture-and-why-visibility-comes-first/*

[5] *Rangrej, S. (2025). Current State of Zero Trust Assessments, Future Directions for Improvement. Clareus Scientific Science and Engineering, 2(9), 07-16. DOI: 10.70012/CSSE.02.056*

Menlo Secure Application Access

Menlo Secure Application Access enforces least-privileged access for every user through the Menlo Cloud. Every user accesses only what their role requires and nothing more.. Policies apply to users, groups, source IPs, and geolocations, without requiring complex network infrastructure or broad network exposure.

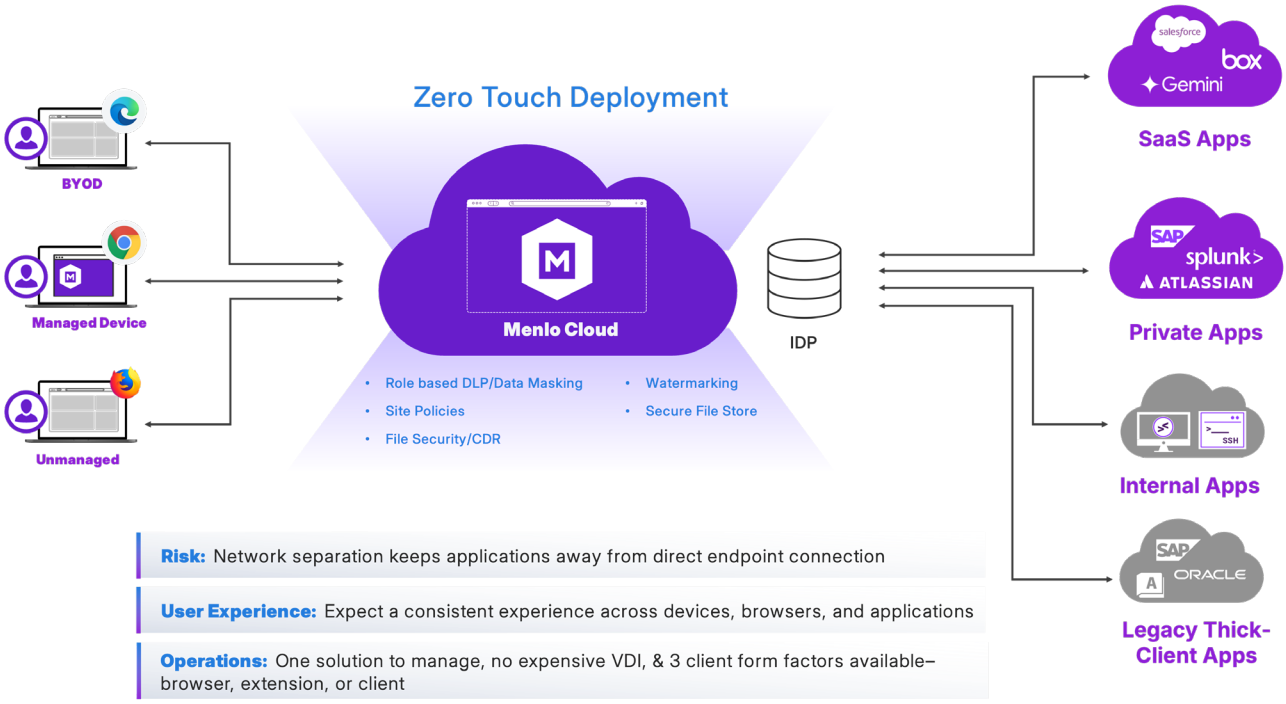

As users' browser traffic passes through the Menlo Cloud, the user's browser is replicated in a hardened, disposable container. That container renders the application and delivers content to the endpoint within the user's local browser. The result is a clean separation between the application and the endpoint: malicious requests involving parameter tampering, web scraping, or API abuse are blocked at the cloud layer. Even a compromised endpoint can't issue direct requests to the application server—all requests run from the Menlo Cloud, not from the device.

This architecture extends to agentic use cases. Least-privileged access narrows agent reach to the specific task assigned. The Menlo Cloud air-gaps agents from application servers, applies bidirectional DLP to block sensitive data consumption, and stops malicious content in files before it can propagate. Agents perform only their designated functions. Compromised agents can't move laterally into sensitive systems.

File Security: Clean Files, Intact Workflows

Menlo File Security and patented Positive Selection® technology treats every file as potentially malicious. It deconstructs files, removes unsafe elements, and delivers clean, fully functional files back to users' workflows in near-real time—without blocking access or interrupting the task.

Data Protection That Works Inside the Session

Traditional DLP tools protect data at the file level. The problem is that data loss in the modern enterprise rarely happens at the file level, but rather inside the the browser session, through uploads, downloads, copy-paste actions, form submissions, and interactions with AI interfaces. Another problem is that upon detection of sensitive data, traditional DLP simply blocks a file in motion, which compromises productivity.[6]

Menlo Browser DLP governs data movement where it actually occurs: inside the session. Controls apply across uploads and downloads, copy-paste, email, collaboration tools, cloud storage, and SaaS applications—through a single platform, based on user, group, domain, and traffic category.

Menlo AI Adaptive DLP addresses the problem traditional DLP tools can't solve: protecting sensitive data without blocking legitimate workflows. AI-powered detection identifies PII, PHI, financial data, and corporate IP within browser sessions, across uploads, downloads, email, collaboration tools, and AI interfaces. But instead of blocking movement of a file, AI Adaptive DLP masks sensitive content and delivers files. The employee or agent continues working. No help desk tickets. No false positives blocking legitimate work.

Zero Trust that Doesn't Hobble Users

Most organizations pursuing Zero Trust face a version of the same problem: VPNs that grant broad network access and introduce unacceptable risk, or VDI environments that are expensive, complex, and built for a world that no longer exists.

Menlo Secure Application Access replaces both. As a VPN replacement, it enforces precise, application-level connectivity without ever exposing the underlying network. As a VDI replacement, it delivers agentless access to applications directly through the browser, without the latency, infrastructure overhead, or management burden of virtual desktops. The platform extends to non-browser-based legacy applications as well, presenting them to end users—human or agent—through the browser.

The result is a single platform that closes the gap most Zero Trust architectures leave open, not just controlling who gets access, but governing what every user can see and do once access is granted.

[6] A single block sends a user into a repetitive cycle of removing sensitive data, getting blocked again, and trying again. Countless hours are lost.

Visibility: Making Zero Trust Provable

Menlo Browsing Forensics completes the picture. It delivers a detailed view of every browser session, including user actions, third party interactions, and session telemetry, that has historically been invisible to security teams. Access policies don't just exist. They're verifiable.

The combination of session isolation, access governance, and forensic visibility gives your security team what most Zero Trust architectures can't deliver: proof that controls are working.

The Takeaway

Your users are everywhere. Your applications are in the cloud. And the browser session—the place where your people and your agents spend most of their working day—is still the gap that most security architectures don't cover.

Menlo Secure Application Access, combined with the Menlo Browser Security Platform gives your organization the access control, data protection, and session visibility to operate securely across every user type, every device category, and every application, without adding friction for the workforce or burden for your team.

With Menlo, you get zero trust that starts in the browser, and follows your users wherever work happens.

About Menlo Security

Menlo Security eliminates evasive threats and protects productivity with the Menlo Cloud. Menlo delivers on the promise of cloud-based security—enabling zero trust access that is simple to deploy. The Menlo Cloud 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.

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