PixieBrix

SECURE BROWSING

THE BROWSER IS THE NEW FRONT DOOR

Your security stack wasn't built for this.

95%

of organizations reported browser-based attacks in 2025

— Palo Alto Networks

8,000%

growth in AI-driven browser sessions over 2025

— HUMAN Security

5 / 6

top egress channels worrying security teams are browser-based

— Fortinet

MONITOR · WARN · BLOCK — IN THE BROWSER

pixiebrix.com

Executive summary

The browser is the dominant runtime for enterprise work — and the fastest-moving attack surface. The access controls built into SaaS apps, internal tools, and websites were designed for the *app*, not your business processes, and not for AI-driven workflows. Traditional SASE, SSE, and DLP see network flows, files, and endpoint events — not what a user does *inside* a tab: which records were viewed, what was pasted into ChatGPT, whether a refund was self-dealt, or whether a junior agent just performed a senior-only action. That is where the highest-velocity threats land first.

The browser is the new front door for AI productivity — and risk

95%

of organizations reported browser-based attacks in 2025.

Palo Alto Networks

1 in 10

browser sessions were AI-driven by end of 2025 — up 8,000%.

HUMAN Security

5 of 6

top egress channels that worry security teams are browser-based.

Fortinet

What makes PixieBrix different

Real-time monitor, warn, and block in the flow of work. Most browser-layer risk is caught post-hoc — in finance reconciliation, in a QA sample, or from a customer complaint. PixieBrix sits inside the browser and intervenes at the moment of the risky action: pre-send context checks, justification prompts, hard gates on actions that bypass KYC. **An extension, not a browser.** No rip-and-replace, no parallel browser to support, no admin access to the host app — deployed through Chrome Enterprise, Intune, or JAMF policies you already run.

Industry leaders are converging on the same architecture: in January 2026, CrowdStrike announced its **$420M acquisition of Seraphic** to *“decouple security from the browser itself”* and deliver Zero Standing Privilege “for the modern agentic workforce.” PixieBrix has been built on that thesis since day one — with a broader process and compliance control surface than pure security tools.

Six categories of browser-layer risk

The rest of this guide walks through six categories of risk that originate in the browser tab — the layer your network and endpoint controls don't see. For each, we document representative TTPs, what the typical control stack catches (and what it misses), and the PixieBrix capabilities that close the gap.

1. ATO Defense	Social engineering, vishing-coached fraud, and compromised agent sessions targeting customer-facing roles.
2. Data Loss Prevention (DLP)	PII over-exposure on screen, cross-customer leaks, and pastes into shadow IT including public AI tools.
3. Fraud Prevention	Self-dealing refunds and credits, skipped KYC verification, and bypassed approval gates.
4. Compliance, Audit & Governance (GRC)	Out-of-policy responses, process drift across sites and BPOs, ignored critical announcements, audit-trail gaps.
5. Identity, Access & Privilege (IAM/PAM)	Privilege creep where host-app RBAC is too coarse to enforce least privilege at the action level.
6. Insider Risk Management (IRM)	Snooping on VIP and personal records, unauthorized edits, and bulk scraping by departing employees.

How to read the TTP tables. Each row pairs a representative risk with what happens with and without browser-layer controls. The **Capability Used** column maps to the PixieBrix building blocks glossary at the end of this guide.

Adversary catalog

The TTPs in this guide are driven by a small set of adversaries. Mapping browser-layer controls to *who* is doing the action — not just *what* the action is — helps prioritize where to deploy guardrails first. Five of these six adversaries are authenticated insiders whose actions look like normal work to network and endpoint tooling. The browser layer is where intent shows up.

Adversary	Description	Typical TTPs
Malicious employee or contractor	Authenticated insider committing fraud: refund-to-self, gift card theft, comping friends, data exfiltration for resale.	Self-dealing transactions; bulk scraping; skipped KYC.
Negligent employee or contractor	Well-meaning but skips steps, mishandles PII, sends to the wrong customer.	Cross-customer leaks; PII over-exposure; out-of-policy responses.
Coached or colluding employee or contractor	Social-engineered by an external fraudster, or part of an organized scheme. The fraudster talks the agent through an exception flow in real time.	Social-engineered ATO; coached refund or gift-card fraud.
Curious employee or contractor	Snooping on celebrity or VIP records, ex-partners, coworkers — outside any active ticket.	Unauthorized record access; unauthorized record editing.
Departing employee or contractor	Bulk data scraping before resignation or termination; clipboard, print-to-PDF, and rapid record sweeps.	Bulk exfiltration; clipboard misuse; rate-limit violations.
Compromised account	External attacker using stolen credentials or a hijacked session of a legitimate employee.	Anomalous action mix; step-up auth bypass; session takeover.

Why this matters for browser-layer controls. The first four adversaries don't trip network or endpoint signals — they're using their own credentials, on their own managed device, on a sanctioned app. The only place to detect intent is in the browser tab where the action happens. The sixth adversary (compromised account) is what most session-anomaly tools target, but they only see app-level signals; PixieBrix adds action-level behavioral baselines.

1. Social Engineering & Account Takeover (ATO) Defense

Customer-facing agents and front-line staff are the soft target. Browser-layer guardrails catch the social-engineering pattern before the high-impact action lands.

Risk	Example	Without PixieBrix	With PixieBrix	Capability Used
Social-engineered agent (vishing / coached fraud)	External fraudster on a call talks the agent through a password reset, account merge, or refund.	Relies on agent recognizing red flags in real time.	Pattern detection on customer signals (new device + reset + high-value action); agent prompted with risk indicators and required additional verification.	**Contextual guidance** **Process intervention** **Conditional access**
Compromised agent session	Stolen credentials or a hijacked session used to perform high-impact actions.	Detection depends on the host app's session anomaly signals.	Behavioral baselines per agent; deviations in action mix, timing, and velocity flagged or trigger step-up auth.	**Anomaly detection** **Conditional access** **Audit logging**

2. Data Loss Prevention (DLP)

Network and endpoint DLP catches files crossing the perimeter. PixieBrix catches the risky action *inside* the tab — before it becomes a DLP event.

Risk	Example	Without PixieBrix	With PixieBrix	Capability Used
PII over-exposure on screen	Full SSN, card PAN, or DOB visible during routine handle time, screen share, or shoulder-surf.	Sensitive fields visible by default in host app; depends on the host app's redaction support.	DOM-level redaction by default; reveal requires explicit click and is logged with reason.	**DOM redaction** **Audit logging**
Cross-customer data leak	Agent has the wrong tab focused and sends Customer A's data to Customer B.	Caught only by customer complaint or recipient report.	Pre-send context check: ticket ID, recipient, and content cross-validated; mismatch blocks the send.	**Form validation** **Process intervention**
Shadow IT / unsanctioned tools	Agents paste customer data into ChatGPT, translation sites, or personal Gmail to get the job done faster.	Often invisible until a leak occurs.	Detect paste events to non-sanctioned domains; in-app sanctioned alternatives surfaced; high-risk pastes blocked.	**Process intervention** **Contextual guidance** **Conditional access**

3. Fraud Prevention

First-party and collusion fraud lives in the workflow gaps your host apps don't enforce. Browser-layer gates move detection from weeks-later reconciliation to the moment of the action.

Risk	Example	Without PixieBrix	With PixieBrix	Capability Used
Self-dealing financial transactions	Agent issues a refund, gift card, or credit to an account they control or a colluding party.	Caught (if at all) post-hoc in finance reconciliation; days-to-weeks detection lag.	Real-time check at point of action: payee, shipping, and email matched against agent identity and known-bad lists; high-risk actions blocked or require supervisor approval.	**Form validation** **Conditional access** **Supervisor approval**
Skipped verification / KYC steps	High-value action processed without ID check, callback verification, or required disclosure.	Policy violation surfaced in a QA sample weeks later.	Hard gate: action button disabled until the checklist is complete; checklist state persisted to the audit log.	**Mandatory workflow** **Form validation**

4. Compliance, Audit & Governance (GRC)

GRC failures are almost always process failures that scaled. The browser layer is where you enforce policy consistently across geographies, vendors, and shift changes — and where you capture the audit trail your host apps don't.

Risk	Example	Without PixieBrix	With PixieBrix	Capability Used
Out-of-policy responses	Agent promises a refund, discount, or commitment outside guidelines, or uses non-approved language.	QA catches a sample; the rest reaches the customer.	Real-time guidance, approved-template insertion, and blocked phrasings flagged before send.	**Contextual guidance** **Snippet manager** **Form validation**
Process drift across sites / vendors	Manila, Austin, and Bogotá teams handle the same workflow differently; one site has higher fraud loss.	Addressed via training and QA; drift returns under attrition.	Same enforced workflow regardless of location, language, or BPO vendor; updates pushed centrally.	**Mandatory workflow** **Centralized deployment**
Critical announcement ignored	Policy change, fraud alert, or new scam pattern not read; agents continue old behavior.	Email blast or LMS module; acknowledgment is self-reported and partial.	In-app blocking announcement on next session; acknowledgment captured per agent.	**Announcements** **Acknowledgment tracking**
Audit-trail gaps	Host app logs the action but not the context — which records were viewed, what guidance was shown, what was redacted.	Investigations rely on screenshots and agent recall.	Browser-side telemetry of what the agent saw, when, and what they did; tamper-evident logs shipped server-side.	**Audit logging**

5. Identity, Access & Privilege Governance (IAM/PAM)

Host-app RBAC is usually too coarse for least-privilege enforcement at the action level. PixieBrix lets you layer fine-grained access logic over the host app — without waiting on the vendor.

Risk	Example	Without PixieBrix	With PixieBrix	Capability Used
Privilege creep / under-RBAC'd host app	Junior agent or contractor performs senior-only actions because host app RBAC is too coarse.	Tickets often raised to the vendor; workaround is separate accounts or shadow processes.	Role-based UI hiding/disabling layered on top of host app RBAC; finer-grained gating without host app changes.	**Conditional access** **Element hiding**

6. Insider Risk Management (IRM)

Insider risk is the hardest to detect with network or endpoint tooling — the actions look like normal work. Browser-layer guardrails add the *intent* signal: justification at access time, rate limits, and telemetry on what was viewed.

Risk	Example	Without PixieBrix	With PixieBrix	Capability Used
Unauthorized record access (snooping)	Agent opens VIP, celebrity, ex-partner, or coworker record without an active ticket.	Logged in CRM but rarely reviewed; relies on retrospective audits.	Justification prompt at access time; high-sensitivity records gated behind manager approval; real-time alert on out-of-pattern access.	**Process intervention** **Justification capture** **Audit logging**
Unauthorized record editing	Agent overwrites customer contact email so the customer doesn't receive a CX quality survey.	Caught (if at all) from a customer complaint.	Justification prompt at access time and write-time validation.	**Process intervention** **Justification capture** **Audit logging**
Bulk data scraping / exfiltration	Departing agent opens hundreds of records, copies to clipboard or local file, or uses print-to-PDF.	May be caught by network DLP if tooling is in place; often missed entirely.	Per-agent rate limits on record opens and exports; clipboard activity telemetry; anomaly alerts to security ops.	**Audit logging** **Anomaly detection** **Conditional access**

Scope and limitations

PixieBrix is a defense-in-depth control. It is not a replacement for the rest of your security stack, and it does not address every browser-layer risk. Because we'd rather you discover the limits during evaluation than after deployment, here is what PixieBrix explicitly does *not* do:

Does not replace host-application authentication, RBAC, or server-side authorization.

PixieBrix layers fine-grained gating *on top of* the host app's identity and authorization model. If the host app's server accepts a privileged action, PixieBrix cannot retroactively un-do it. Use it to enforce least-privilege at the action level, not as the primary auth boundary.

Does not prevent out-of-band exfiltration — screen photos, dictation to a phone, memorization.

Anything that leaves the screen via a path that doesn't touch the browser is out of scope. Pair PixieBrix with screen-recording detection, watermarking, or workforce policies if these vectors matter to your threat model.

Does not catch fraud via direct API access, a different browser, or an unmanaged device.

PixieBrix is browser-side. Adversaries who bypass the managed browser entirely — by calling the host app's API directly, using a personal device, or installing a fresh browser without policy — are out of scope. Enforce managed-browser-only access at the IdP or SASE layer.

Effectiveness depends on the extension being installed and enabled.

Managed deployment via MDM, group policy, or browser-management policy (Chrome Enterprise, Edge for Business, Intune, JAMF) is required. Force-install and disable-removal policies are recommended in production.

Does not provide native EDR, NDR, or SIEM functionality.

PixieBrix emits browser-side telemetry — what the user saw, when, and what they did — and ships it to your existing observability and SIEM stack. It complements EDR/NDR/SIEM; it does not replace them.

What this scoping means for procurement. If your evaluation criteria treat PixieBrix as a replacement for IdP, EDR, SIEM, DLP, or the host app's RBAC, the fit will be poor. If you're scoping it as a complementary control that closes the browser-action gap your existing stack doesn't see, the fit is strong.

Capability building blocks

Each PixieBrix guardrail is composed from a small set of low-code building blocks. IT and operations teams combine these capabilities to ship custom controls in minutes — embedded in the SaaS, internal apps, and websites your team already uses. The same building blocks appear across every category in this guide.

Audit logging	Log to PixieBrix telemetry or your o11y/SIEM. Captures field changes and custom events.
Conditional access	Screen by role, attribute, or live activity (e.g., WFM shift data).
Element hiding	Hide irrelevant UI by user or line of business — panes, menus, and actions.
DOM redaction	Mask sensitive data with optional audit logging or justification capture.
Form validation	Custom validation rules with API checks — catch invalid combos and flag fraud.
Contextual guidance	Overlays, tooltips, and guided tours for completing processes.
Mandatory workflow	Enforce required steps or a specific order of completion.
Justification capture	Collect justifications for overrides — access, redaction, validation, and more.
Supervisor approval workflow	Gate actions on supervisor or manager approval via custom forms.
Announcements	Show announcements in a side panel or as a screen overlay.
Acknowledgment tracking	Require and track acknowledgments as auditable events.
API-based automation	Automate privileged steps via API to enforce least-privilege access.

How PixieBrix fits your stack

vs. Secure Enterprise Browsers (Island, Talon)

Secure enterprise browsers replace the browser itself, which usually requires user retraining, migration, and a parallel toolchain to support. PixieBrix runs as an extension inside the Chromium-based browsers (Chrome, Edge, and others) your team already uses, layering guardrails into existing workflows without forcing a switch. Deploy through Chrome Enterprise policies, Group Policy, Intune, JAMF, and other standard MDM tools.

vs. DLP tools

DLP catches files crossing network or endpoint boundaries. PixieBrix sits *inside* the browser, catching risky actions, paste events, and policy violations in the flow of work — before they ever become DLP events. The two are complementary: PixieBrix is the early-warning and intervention layer; DLP remains the perimeter backstop.

Secure by design

PixieBrix processes data on-device — nothing leaves the browser unless you configure telemetry to flow to your SIEM or observability stack. The control plane runs behind a WAF with continuous app security monitoring, and integrates natively with the IdP, SSO, SIEM, and DLP your team already uses. PixieBrix is SOC 2 Type II certified, with full security documentation — penetration tests, architecture reviews, and policies — available on request.

> **The market is converging on this architecture.** On January 13, 2026, CrowdStrike announced its **$420M acquisition of Seraphic**, a browser-runtime security extension. In the announcement, CEO George Kurtz framed the architecture explicitly: *"By decoupling security from the browser itself, we can turn any browser into a secure enterprise browser, without forcing change or slowing productivity… defining the future of Zero Standing Privilege for the modern agentic workforce."* That is the same architectural thesis PixieBrix has shipped from day one — applied not only to security but to the full surface of process, compliance, and access controls that live in the browser tab.

In production

Coached-fraud attack at a major local delivery platform. Junior contact-center agents were targeted by an external fraud ring that talked them through an exception flow to issue gift cards. Host-app RBAC could not prevent the workflow — it was a legitimate, approved-by-policy exception path. The fraud was caught only after reconciliation, with material loss already booked.

With PixieBrix, the platform conditionally hid the exception workflow for the agent tier being targeted, gated remaining access behind supervisor approval, and added browser-side telemetry on access attempts. The attack pattern stopped. Total time to deploy the intervention: hours, not a vendor release cycle.

PixieBrix is also deployed in production at TaskUs and other global BPO and enterprise customers (case studies at pixiebrix.com/customers).

"PixieBrix is our secret weapon." — SVP Digital, publicly traded BPO

Glossary

Quick reference for the security and access-control terminology used in this guide.

SASE	Secure Access Service Edge. Coined by Gartner in 2019. Combines SSE with SD-WAN to deliver network and security as a unified cloud service.
SSE	Security Service Edge. The cloud-delivered security portion of SASE. Components: CASB, ZTNA, SWG, and FWaaS.
SD-WAN	Software-Defined Wide Area Networking. The networking half of SASE.
SWG	Secure Web Gateway. Protective barrier between users and the internet. Filters web traffic, blocks malicious content, prevents access to risky sites.
CASB	Cloud Access Security Broker. Coined by Gartner in 2012. Provides visibility and control over cloud applications via an intermediary deployment model. Four pillars: visibility, compliance, data security, threat protection.
ZTNA	Zero Trust Network Access. Identity-aware, application-level access control that replaces broad network access with per-app verification.
FWaaS	Firewall-as-a-Service. Cloud-delivered next-gen firewall capabilities.
ZSP	Zero Standing Privilege. An identity-security principle that eliminates persistent administrative rights. Instead, just-in-time (JIT) access grants temporary, least-privilege permissions only when needed and revokes them automatically. PixieBrix's API-based automation capability supports ZSP by performing privileged steps on behalf of the user without granting standing privilege.
DLP	Data Loss Prevention. Tooling that inspects files and data flows crossing network or endpoint boundaries to block exfiltration.
EDR	Endpoint Detection and Response. Endpoint-resident agents that detect malware, suspicious processes, and lateral movement.
IRM	Insider Risk Management. Programs and tooling that detect risky or malicious activity by authenticated insiders.
GRC	Governance, Risk, and Compliance. The discipline of aligning controls and processes with regulatory and policy requirements.
ATO	Account Takeover. An attacker gains control of a legitimate user account via credential theft, session hijack, or social engineering.
IAM / PAM	Identity and Access Management / Privileged Access Management. The systems that authenticate users and govern who can do what at the application and infrastructure level.

Frequently asked questions

How is PixieBrix different from a secure enterprise browser like Island or Talon?

Secure enterprise browsers replace the browser itself, which often requires user retraining and migration. PixieBrix runs as an extension inside the browsers your team already uses, layering guardrails into existing workflows without forcing a switch.

How is PixieBrix different from a DLP tool?

DLP catches files crossing network or endpoint boundaries. PixieBrix sits inside the browser, catching risky actions, paste events, and policy violations in the flow of work — before they ever become DLP events.

Does PixieBrix work with internal apps and custom tools, not just SaaS?

Yes. PixieBrix runs as a browser extension on top of the DOM, so it works with any browser-based application — third-party SaaS, internal apps, or custom-built tools.

Where does PixieBrix process my data?

On-device. The PixieBrix extension processes data inside the browser, and nothing leaves the device unless you configure telemetry to flow to your SIEM or observability stack.

How long does it take to deploy a new guardrail?

Most guardrails can be built and deployed in minutes using low-code or AI coding — no sandbox or test environment required.

Can PixieBrix detect when agents paste data into ChatGPT or other AI tools?

Yes. PixieBrix detects paste events to non-sanctioned domains (including ChatGPT, Claude, and other public AI tools), can surface sanctioned alternatives in-app, and can block high-risk pastes entirely.

Can I deploy PixieBrix through my MDM and browser management tools?

Yes. PixieBrix can be deployed and managed through Chrome and Edge enterprise policies, Group Policy, Intune, JAMF, and other standard MDM tools.

Is PixieBrix SOC 2 compliant?

Yes. PixieBrix is SOC 2 Type II certified. Full security documentation — penetration tests, architecture reviews, and policies — is available on request.

YOUR NEXT STEP

Ready to close the browser gap?

See a working PixieBrix deployment built on your highest-risk browser workflow. A solutions architect will surface your top three browser-layer gaps and demonstrate live guardrails that intervene at the moment of risk.

Book a working demo →
pixiebrix.com/schedule-demo

Or read the full product overview → **pixiebrix.com/pixiebrix-secure-enterprise-browsing**

What to expect on the call

✓ A 15-minute working session on your top three browser-layer risks

✓ A live demo of monitor / warn / block on a SaaS or internal app you choose

✓ Walkthrough of MDM deployment, SIEM integration, and SOC 2 documentation

✓ An honest read on where PixieBrix fits — and where it doesn't — alongside your existing SASE, SSE, and DLP

Trust & compliance

✓ **SOC 2 Type II certified** — pen tests, architecture reviews, and policies available on request

✓ **On-device processing** — nothing leaves the browser unless you configure telemetry

✓ **In production** at TaskUs and global BPO and enterprise customers

✓ **Native integrations** with your IdP, SSO, SIEM, and DLP — no rip-and-replace