

PROCIRCULAR / KEYNOTE

Shadow AI

The risk hiding in plain sight.

Aaron Warner

Founder & CEO, ProCircular

Translational AI Center / Iowa State University
May 2026



A quick word about where I'm coming from.

AARON WARNER

PROCIRC

Aaron R. Warner is the Founder and CEO of ProCircular, a cybersecurity consulting firm headquartered in North Liberty, Iowa, serving midmarket organizations across healthcare, higher education, manufacturing, and financial services. Before founding ProCircular in 2016, Aaron spent 22 years as CIO and CTO of Integrated DNA Technologies, a global biotech leader later acquired by Danaher for \$2.2 billion. He holds an MBA from the University of Iowa's Tippie College of Business, is a Certified Information Systems Security Professional (CISSP) and Security+ engineer, a Novell 3.12 CNE and a member of the FBI/DHS InfraGard partnership.

Aaron is a regular cybersecurity contributor for KCRG-TV9, and was named the 2023 SBA Small Business Person of the Year. ProCircular has earned Inc. 5000 recognition and delivers offensive security, managed detection and response (MXDR), governance and compliance (vCISO), AI security, and incident response services to organizations navigating an increasingly complex threat landscape. Aaron is a frequent speaker on cybersecurity strategy, nation-state threats, and the intersection of technology risk and business leadership.



How We Protect Organizations



OFFENSIVE SECURITY

Penetration Testing & Red Team

External, internal, wireless, application, and API pen testing. Adversary simulation, social engineering, vishing, and phishing campaigns. We find your weaknesses before attackers do.



MANAGED XDR

24/7 Detection & Response

Continuous monitoring with SentinelOne, high-fidelity alerting, threat hunting, and automated containment. MXDR covers endpoints, cloud, identity, and network in a single platform.



GRC & vCISO

Governance, Risk & Compliance

Gap assessments, framework mapping (NIST CSF, ISO 27001, CMMC, HIPAA), vendor risk management, policy development, and fractional CISO leadership for boards and executives.



AI & CYBERSECURITY R&D

AI Governance & Innovation

AI risk assessments, shadow AI discovery, Colorado/EU AI Act compliance readiness, and security architecture for AI-enabled organizations. Led by our VP of AI & Cybersecurity R&D.



INCIDENT RESPONSE

When Bad Things Happen

Rapid triage, forensic preservation, root cause analysis, and remediation for ransomware, BEC, data breaches, and supply chain compromises. Available on retainer or emergency engagement.



ATTACK SURFACE MONITORING

See What Attackers See

Continuous external reconnaissance using the same tools threat actors use. Monitors for exposed services, typosquatting, certificate issues, and cloud misconfigurations.

The headline.

The fastest-growing security risk inside your organization isn't an attacker.

It's something your own people are doing to be helpful.

That, and a software vendor you already trust quietly turning on an AI feature without telling you. That one is harder to fix because nobody on your team did anything wrong.

Two faces of the same problem.

WHAT YOUR PEOPLE DO

- Paste the customer contract into ChatGPT for a quick summary.
- Install an AI browser extension nobody approved.
- Build an automation in a no-code tool against production data.
- Forward sales decks to a personal account so the AI assistant can clean them up.

WHAT YOUR VENDORS DO

- Turn on AI features by default in HRIS, ERP, CRM, or your code repo.
- Quietly add language to the terms allowing your data to train models.
- Ship cross-tenant integrations the security team never reviewed.
- Embed agents that act on tokens issued months or years ago.

Both bypass IT. Neither was approved. Both create exposure your DLP can't see.

What the IBM report actually proves.

97%

of organizations with an AI-related breach lacked proper AI access controls

\$670K

added to average breach cost when shadow AI was a factor

65%

of shadow AI breaches involved customer PII (vs. 53% global)

247 days

average lifecycle of a shadow AI breach, longer than any other category

63%

of breached organizations had no AI governance policy at all, or were still in the middle of writing one.

Of the ones that had a policy, only 34% audited for unsanctioned AI.

Source: IBM 2025 Cost of a Data Breach Report (Ponemon Institute, 600 organizations)

Why this is different now.

For twenty years, small and mid-sized businesses had an unspoken advantage. **They weren't worth the effort to attack carefully.**

AI eliminated the cost of attention. The math no longer works in your favor.

+1,265%

rise in phishing volume,
late 2024

+442%

rise in voice phishing
over the same period

7 sec

of audio is enough to clone
your CEO's voice

Every business with a domain name and an email server is now worth the effort.

Three vectors that show up in every assessment.

01

EMPLOYEE PASTE

An HR coordinator pastes termination details into ChatGPT to polish the wording. They don't think they sent employee data outside the firm. They did.

02

VENDOR-EMBEDDED AI

Your CRM, HR system, and code repository all shipped AI features in the last six months. Most were on by default. None went through your security review.

03

AUTONOMOUS AGENTS

AI agents now act on their own credentials, with broad permissions, against systems your IT team never granted them. They're identities. You aren't governing them like one. And your employees are building them every day.

Each one needs a different defense. None of the three are caught by traditional DLP.

The places it hurts most.



Healthcare

Protected health information walks out through a chatbot. HIPAA doesn't make exceptions for accidental.



Higher Ed

Research data, FERPA records, grant work. Faculty and grad students moving fast with whatever AI tool helps them publish.



Financial Services

PII and customer financial data flowing into AI summaries. Regulators are already asking for the audit trail.



Manufacturing

IP, process knowledge, supplier contracts. Your competitive moat being summarized by a model you don't own. It's also a dependency you don't control and can't fix if it breaks.

AND THE LAW HAS CAUGHT UP: Colorado AI Act in effect Feb 1, 2026. EU AI Act high-risk provisions land August 2026.

PART TWO

Three(ish) stories from the last eighteen months.

Each one is a different failure mode. All three happened to companies you've heard of.

Case 00 / Halloween Sucked/ October 2025

WHAT HAPPENED

Aaron's House

More than a year of development using a variety of tools as AI evolved, from cut and paste to AI through IDE integration and into Claude Code 2.0.

The codebase looked great and most functionality worked. The development tool, piloted by Claude Code, decided to die on the afternoon before Halloween. The last changes were not stable, and the rest of the night was wasted trying to get the entire system online because of a broken toolset.

THE PATTERN

9

When you write code in AI, you build dependencies in the code AND the development platform.

You wrote an app. Now you can't fix it.



Read the AI clauses in your renewal terms. Your vendor already shipped a feature you didn't approve.

Case 01 / Salesloft Drift / August 2025

WHAT HAPPENED

salesloft

drift

Drift sells an AI-powered chatbot that integrates into Salesforce. To work, it holds long-lived OAuth tokens for every customer who connects it.

Attackers compromised Salesloft's GitHub, moved into their AWS environment, and walked out with active customer tokens. Then they used those tokens to log into Salesforce instances directly.

It looked legitimate. The connection was supposed to be there. The token was valid. Nobody's user account was compromised.

THE BLAST RADIUS

700+

customer organizations affected by one vendor's compromise

INCLUDING

Cloudflare

Palo Alto

Zscaler

CyberArk

Your vendor's AI integration is your attack surface. Even the security companies got hit.

Case 02 / Asana MCP / June 2025

WHAT HAPPENED

Asana

Asana shipped Model Context Protocol so customers could connect their Asana data to ChatGPT, Claude, and other large language models.

A logic flaw in the MCP server allowed cross-tenant access. Customer A's queries could surface tasks, project metadata, and uploaded files that belonged to Customer B.

It went undetected for over a month. Customers found out when Asana told them to go look at their MCP access logs.

THE NEW FAILURE MODE

Cross-tenant AI data leaks

AI integrations sit on top of multi-tenant data. When the integration has a bug, the data of every customer using it is on the table at once.

AI features ship faster than they get tested. Cross-tenant flaws are now a category, not an outlier.

Case 03 / Slack AI training / May 2024

WHAT HAPPENED

slack

Slack added a paragraph to its privacy principles saying it analyzed customer messages, content, and files to develop AI and ML models.

Every workspace was opted in by default. The opt-out was an email to a specific address with a specific subject line.

Nobody noticed for nine months. Then one Hacker News post went viral and Slack walked it back over a weekend.

THE PATTERN

9

months that the policy was live before anyone noticed.

Your terms changed. Nobody told you. The default was on.

Read the AI clauses in your renewal terms. Your vendor already shipped a feature you didn't approve.

What the three stories have in common.

01

AI integrations are the new attack surface

None of these were endpoint compromises. Two were SaaS-to-SaaS through embedded AI. One was a default-on policy change. Your EDR can't see any of it.

02

Cross-tenant leaks are now a category

When AI features ship fast and span multiple customers' data, a single bug puts everyone using the feature at risk simultaneously. This is new. Plan for it.

03

Default-on AI changes your data posture without notice

Your CRM, your collaboration tool, and your code repo all added AI features in the last year. Most were opt-out. Some you still haven't reviewed.

You're managing what your software vendors do without telling you.



PART THREE

What the market is doing about it.

Three of the four largest security vendors shipped a Shadow AI product in the last 30 days.

Four approaches the market has converged on.



DATA LINEAGE

Follow the data

Track where data came from, not just what it looks like.
Block sensitive content based on origin.

Cyberhaven, Microsoft Purview



AI-SPM / RUNTIME

Protect the apps you build

Inspect prompts and outputs at runtime. Red-team your own AI before attackers do.

SentinelOne Prompt, Protect AI, HiddenLayer



ENDPOINT DISCOVERY

Find what's already running

Use your existing EDR to inventory unsanctioned AI apps, browser extensions, and local models.

CrowdStrike, SentinelOne, your XDR



IDENTITY & GATEWAYS

Govern AI as identities

Treat agents as users. Route AI traffic through a monitored choke point. Enforce least privilege.

Cloudflare AI Gateway, Microsoft Entra / Agent 365

The new releases

SentinelOne

Prompt Security

- Real-time prompt injection defense
- Pre-production AI red teaming
- Runtime protection for code assistants and home-grown apps

BEST FIT

If you're building AI applications and need active runtime protection.

None of these existed twelve months ago.

Microsoft 365 E7


\$99.00

per user/month, paid yearly

(Annual commitment)

Contact Sales

- ✓ Agent 365—a control plane for agents
- ✓ Advanced identity, network access, and security
- ✓ Desktop, web, and mobile versions of Word, Excel, PowerPoint, Outlook, and other apps
- ✓ 1 to 5+ TB of cloud storage per user
- ✓ Deploy, manage, and update devices at scale with Windows for Enterprise
- ✓ Chat, call, meet, and host webinars with Microsoft Teams
- ✓ Shadow IT discovery capabilities, antivirus, antimalware, phishing, and ransomware protection
- ✓ Unified endpoint management, protection, and security
- ✓ Scalable business analytics with Power BI
- ✓ Extended detection and response (XDR), hybrid identity protection (ITDR), and malicious OAuth app protection
- ✓ Data security and compliance for higher data risk and regulations
- ✓ Secure AI built for work, powered by Work IQ, with access to agents, work-grounded Copilot in apps, and the ability to create agents



CrowdStrike

Shadow AI Visibility Service

Endpoint-based discovery of unapproved AI apps

Threat modeling by CrowdStrike experts

Adversary-informed risk prioritization

BEST FIT

You need to find out what's already running before you can govern it.

Evaluating one of these by year-end.

If you're just starting. Where to spend the first dollar.

If you're heavy in M365...



Start with Purview labels and Microsoft Agent 365. You already pay for most of it.

If you have CrowdStrike or SentinelOne...



Flip on the Shadow AI module before buying anything new. Path of least resistance.

If you build AI applications in-house...



Look at AI-SPM platforms. SentinelOne Prompt, Protect AI, HiddenLayer.

If you're starting from zero on visibility...



Endpoint discovery first. You can't write a policy for tools you don't know about.

And there's a whole ecosystem behind the big three.

DISCOVERY & DATA

Find what's happening, follow the data

Cyberhaven

Nudge Security

Harmonic Security

Lasso Security

RUNTIME & RED TEAM

Protect the AI apps you build

Lakera

HiddenLayer

Protect AI

TrojAI

AGENTIC & AI SOC

Govern agents, accelerate response

Noma Security

Prophet Security

Intezer

Torq

None of these are the safe enterprise pick. All of them are worth knowing about, because at least one will get acquired in the next twelve months.



PART FOUR

Build a culture, not a cage.

If you ban it, your people will work around the ban. And then you'll have less visibility than you started with.

Five principles for a culture of protection.

01

Be open about the risks.

Name what's at stake in language your people understand. Customer data. The contract you're trying to win. The audit you're about to fail.

02

Share the examples.

Tell the Salesloft story. Tell the Asana story. People remember stories. They forget policies.

03

Make the safe path the easy path.

Sanction a tool. Give it a budget. Train people on it. If your sanctioned option is harder than the unsanctioned one, you've already lost.

04

Train your people to protect the firm, not to fear breaking the rules.

Fear gets compliance during the training. It doesn't change behavior on Tuesday morning when someone's trying to finish a deck.

05

Treat AI procurement like security review.

Because it is. Read the AI clauses. Ask vendors what data trains their models. Get the answer in writing.

*Treat AI the way mature companies
(eventually) treated cloud.*

Accept that it's being used, help people to use it right.



awarner@procircular.com

Questions?

Aaron Warner / ProCircular / procircular.com



UP NEXT / LIVE DEMO

SentinelOne

Prompt Security

Real-time prompt inspection. Active jailbreak and injection defense. Visibility into every employee, developer, and agent interaction with an LLM.

Watch how it stops a sensitive paste before the model ever sees it.

[DEMO INSERT]