Anzenna

CROWDSTRIKE

CrowdStrike + Anzenna

Turn Endpoint Telemetry into People-Centric Security Operations

CrowdStrike Falcon is the leading EDR platform for protecting endpoints from external threats. But endpoint protection alone leaves gaps when threats come from within your organization. Employees upload sensitive code to ChatGPT, grant OAuth access to risky third-party apps, and exfiltrate data through SaaS channels. Endpoint telemetry alone can't tell that story.

Anzenna is the agentic layer for insider risk management. It extends CrowdStrike by combining endpoint telemetry with identity, SaaS, code, and HR signals from 100+ integrations to build a complete picture of every person's activities. AI agents autonomously investigate every detected threat and trigger remediation instantly. Your security team gets full visibility into user behavior from endpoints to the cloud, and can act on risk the moment it appears, without adding headcount.

The Reality: CrowdStrike Sees Endpoints, Not People

CrowdStrike Falcon delivers strong endpoint telemetry: malware detection, process execution monitoring, USB activity, and network connections. For external threats, this coverage works well. But insider risk operates differently. Malicious or negligent insiders don't trigger endpoint anomalies. They use legitimate credentials and authorized applications to move data where they shouldn't.

CrowdStrike captures a large volume of telemetry, but the signals that matter for insider risk stay buried in noise. The platform wasn't built to highlight identity context, behavioral patterns, or SaaS-layer activity. The data may exist in LogScale, but turning it into action requires dedicated query expertise most teams don't have.

What CrowdStrike Sees

- User authenticated successfully
- Browser launched and connected to cloud services
- Files accessed from the local directory
- Network connection established to external domains
- No malware detected, no suspicious process behavior

What Gets Buried or Missed

- The user received a departure notice from HR yesterday
- That browser session uploaded 47 files to a personal Google Drive account
- Those files contained proprietary source code for your flagship product
- The same user granted OAuth access to an unknown third-party app with full read/write permissions
- Five minutes later, they shared a confidential M365 folder with an external email address

CrowdStrike's endpoint telemetry shows normal behavior. The user has valid credentials, is using authorized applications, and is not infected with malware. But this is a textbook insider threat in progress. The endpoint data exists. Without behavioral analytics that tie identity, SaaS, and HR context together, these signals never get flagged as a coordinated risk.

Why 'We're Already Covered' Is a Dangerous Assumption

Security teams often believe insider risk is handled because they've invested in strong tools. But each tool operates in its own silo. None provides the full picture.

"We Have CrowdStrike. Our Endpoints Are Protected."

Reality: CrowdStrike protects endpoints from malware and unauthorized access. It captures telemetry on user activity but doesn't highlight or prioritize post-authentication behaviors such as uploading code to AI tools, sharing files via SaaS, or granting OAuth permissions. These activities leave a minimal footprint at the endpoint and get lost in the noise.

"We Have MDR. Someone Is Watching."

Reality: MDR services monitor EDR alerts for external threats. They face the same challenge: insider risk signals exist in telemetry, but no one highlights or correlates them with SaaS activity, OAuth grants, AI interactions, or cloud data movement.

"We Have DLP. Data Exfiltration Is Blocked."

Reality: DLP excels at pattern matching but struggles with intent. It can't distinguish legitimate backups from preparation for theft and doesn't detect OAuth grants or browser-based uploads that bypass DLP controls.

"We Have a SIEM. Everything Is Logged."

Reality: IEMs aggregate logs without user context. Analysts manually pivot between systems to understand identity, access, and behavior. Investigations stretch from minutes to hours.

"How Many People Do I Need to Run This?"

Reality: Zero new hires. Anzenna automates the investigation workflow that currently consumes your team's time. Your existing analysts review finished cases instead of manually correlating alerts across six different consoles. Organizations handling 200+ CrowdStrike alerts per month recover 40 to 60 analyst hours by automating investigation, employee notification, escalation, and follow-up. One analyst can handle 50+ cases per day when the investigation arrives already completed.

Anzenna's Four Core Pillars

People and Priority First

Every activity is tied to a user identity, enriched with HR context and historical behavior. Your security team focuses on high-risk users exhibiting concerning patterns rather than individual alerts across fragmented systems. A CrowdStrike alert becomes: "Sarah Chen (Senior Engineer, departing, 92nd percentile risk score) has a critical threat on her MacBook Pro" instead of just "Device abc123 has a detection."

Highlights What Matters and Adds Behavioral Analytics

CrowdStrike captures endpoint telemetry. Anzenna highlights what matters by adding behavioral analytics that CrowdStrike doesn't provide, including the identification of anomalous patterns across AI tool usage, OAuth-granted apps, SaaS platforms, and browser-based activity. Raw LogScale data that would otherwise require dedicated query expertise gets automatically parsed, classified, and attributed to named employees.

Agentic AI That Investigates and Remediates Autonomously

When a CrowdStrike detection fires (or any other alert across your connected systems), Anzenna's AI agents activate autonomously. They examine weeks of cross-platform activity, pull the employee's full context, connect individual events into narratives, and produce a structured investigation brief with evidence, confidence scoring, and a recommended response. AI agents investigate every threat, not a subset. For cases that require human judgment, agents deliver the investigation already completed: timelines built, evidence linked, and actions recommended. When threats are confirmed, agents trigger coordinated remediation across the entire stack: isolate endpoints via CrowdStrike's API, revoke OAuth grants, deactivate accounts, and notify the employee via Slack, Teams, or email, all at once. If the employee doesn't respond, Anzenna auto-escalates to their manager and creates a Jira or ServiceNow ticket.

Unifying the Operational Layer

Anzenna adds the missing operational layer to your existing security investments. 100% API-based integration with 100+ platforms, including CrowdStrike, Microsoft, Google, and Okta. No agents, deployment in minutes. Your team uses the same CrowdStrike investment they already have. Anzenna simply makes it work harder.

What Anzenna Adds on Top of CrowdStrike

Capability	CrowdStrike Alone	CrowdStrike + Anzenna
Threat detection	Alerts in Falcon console	Alerts correlated to employee identity, risk score, department, and manager
Who is affected?	Device AID/hostname	Named employee with job title, department, manager chain, and full risk history
Agentic AI investigation	No built-in AI investigation. SOC analysts manually triage in Falcon	Every threat is investigated autonomously by AI agents. Analysts receive structured briefs with evidence and recommended actions
Remediation	SOC analyst triages in Falcon	Coordinated response across endpoints, SaaS, and identity platforms, all triggered automatically
Escalation	Manual follow-up	Auto-escalate to the manager if the employee doesn't respond within a configurable window
Data exfiltration	Raw telemetry in LogScale	Parsed, classified, and risk-scored exfiltration across 6 channels (USB, SCP, Git, AirDrop, network shares, DLP)
Privilege management	Manual admin group changes	Self-service just-in-time admin access with approval workflows and auto-revoke
Device context	Hardware and OS details	Device + owner + risk score + all behaviors across every connected system
Cross-platform correlation	Endpoint signals only	Endpoint + identity + collaboration + code + data platform signals in one risk score

Data Exfiltration Monitoring

Anzenna queries CrowdStrike's LogScale in real time and turns raw sensor events into classified, risk-scored data movement events attributed to named employees. Six categories of exfiltration are tracked:

- **USB file writes:** Files copied to removable drives, with file name, size, and SHA256 hash
- **SCP and Rsync transfers:** Command-line file transfers to remote servers, classified by direction
- **Git operations:** Code pushed, pulled, or cloned to/from remote repositories via SSH and HTTPS
- **AirDrop:** macOS file transfers detected via AirDrop send directory monitoring
- **DLP policy violations:** Sensitive data sent via browser, email, chat, cloud apps, or printing

Most organizations have CrowdStrike LogScale data sitting idle because it requires dedicated query expertise to put into action. Anzenna ships with pre-built queries for all six exfiltration categories, so your team doesn't need LogScale expertise.

Endpoint Privilege Management (EPM)

Anzenna uses CrowdStrike's Real-Time Response (RTR) infrastructure to deliver just-in-time admin access, one of the most operationally impactful features of the integration.

Employee requests temporary admin access through Anzenna

Approval workflow evaluates the request. Low-risk employees can be auto-approved; others require manager's sign-off

Anzenna executes the privilege grant remotely via CrowdStrike RTR

Time-limited access: privileges automatically expire after the configured window

Auto-revoke: Anzenna sends a second RTR command to remove admin rights when the window closes

Full audit trail: every request, approval, grant, and revocation is logged

Replaces standalone PAM tools ($15 to $50/user/month). For a 500-person company, that's $90K to $300K in avoided PAM licensing costs per year.

Remote Auditing via RTR

Anzenna uses CrowdStrike's Real-Time Response to run lightweight audit scripts directly on devices, with no additional agents required:

- **Browser extensions:** Every extension in Chrome, Edge, and other browsers
- **Installed applications:** Full software inventory (Windows registry scan, macOS Applications)
- **IDE extensions:** Extensions in VS Code, JetBrains IDEs that may access source code
- **AI tool audit:** Model Context Protocol (MCP) clients and servers on the device
- **App uninstall:** Remotely removes unauthorized applications (Windows)

Real-World Scenario: The Departing Engineer

DAY 1

Resignation Submitted
Sarah, a senior software engineer, submits her resignation. HR updates her status in Workday. CrowdStrike sees her endpoint activity as normal. She's accessing the codebase she's worked on for three years: no malware, no unauthorized access, no suspicious processes.

DAY 3

Behavioral Changes Begin
Sarah begins downloading repositories to her local machine at an unusually high rate. She uploads multiple code files to ChatGPT for 'documentation review.' She grants OAuth access to a third-party code collaboration tool her team doesn't use. She shares several M365 folders containing architectural diagrams with her personal Gmail account.

What CrowdStrike Sees

- Normal file access from an authorized user
- Browser connections to legitimate websites (ChatGPT, Gmail, M365)
- No malicious code, no policy violations, no blocked actions

What Anzenna Sees:

- Sarah's HR status changed to 'departed' two days ago
- Download volume is 4x her baseline: 47 files in 6 hours vs. the typical 12/day
- ChatGPT uploads contain code snippets matching proprietary IP patterns
- OAuth app granted full read/write access without IT approval
- External sharing to a personal account outside normal collaboration patterns
- Git clone of 3 repositories she hasn't accessed in 6 months

Anzenna's Response:

Within minutes of Sarah's first anomalous download, Anzenna flags her as high-risk. An AI agent connects her HR departure notice from Workday to the spike in repository clones, the ChatGPT uploads containing proprietary code patterns, the unauthorized OAuth grant, and the external M365 shares. By the time her analyst opens the case, a structured brief is waiting: six correlated signals, a timeline, evidence links, and a recommended response. With one click, the team isolates her endpoint via CrowdStrike's API, revokes the OAuth grant, turns off external sharing, and notifies legal. Sarah receives an automated Slack message explaining the hold. Total response time: 8 minutes.

Without Anzenna, this incident would have required manual correlation across CrowdStrike, Workday, M365 logs, and browser history. It likely would have taken hours, or been missed entirely until Sarah's new employer launched a competing product using your architecture.

Business Impact

Anzenna transforms insider risk from a resource-intensive manual process into an automated, scalable operation. The combination of CrowdStrike endpoint telemetry with 100+ additional integrations gives your team a level of visibility that would otherwise require dedicated headcount and custom tooling to achieve.

Efficiency Gains

Metric	CrowdStrike Alone	CrowdStrike + Anzenna
Investigation time per threat	1 to 4 hours (manual cross-console review)	Minutes. Analysts review finished briefs, not raw alerts
Time from detection to response	Hours (manual investigation, outreach, follow-up)	Minutes. Detection, investigation, notification, and escalation happen automatically
Analyst capacity	1 analyst handles 5 to 10 cases/day	Same analyst handles 50+ cases/day
Time to identify the device owner	5 to 15 min (manual lookup)	Instant (auto-attributed to named employee)
Privilege elevation turnaround	Hours to days (IT ticket)	Seconds to minutes (self-service via RTR)
Data exfiltration detection	Requires LogScale expertise	Automated, pre-built, employee-attributed
Software inventory refresh	Manual or quarterly audits	Continuous via scheduled RTR audits

Cost Savings

Analyst force multiplier
A team handling 200+ CrowdStrike alerts/month recovers 40 to 60+ analyst hours. One analyst covers workloads that previously required three.

Privilege Access Management
Replaces standalone PAM tools ($15-$50/user/month). For 500 users, that's $90K to $300K in avoided licensing costs per year.

Investigation automation
Eliminates 1 to 4 hours of manual cross-console correlation per threat. Every investigation arrives ready for analyst review.

LogScale operationalization
Extracts value from LogScale telemetry that would otherwise require dedicated detection engineers.

Shadow IT detection
Remote RTR audits replace manual device inspections or additional agent deployments.

Tool consolidation
One platform replaces separate tools for employee notification, ticket creation, risk scoring, and compliance reporting.

How Anzenna Works With CrowdStrike

Anzenna connects to CrowdStrike via a secure API connection. You don't deploy agents on endpoints or make architectural changes. Integration takes about 5 minutes and you see value right away.

- 100% Agentless: Connects directly to Falcon's management console. Zero endpoint performance impact.
- Real-Time Data Flow: Endpoint telemetry enriched with identity, SaaS, and HR context.
- Bi-Directional Response: Trigger actions through Falcon's API to isolate hosts and restrict network access while simultaneously acting on SaaS and identity platforms.
- Full Visibility: CrowdStrike secures endpoints. Anzenna highlights the insider risk signals buried in that telemetry and extends coverage to SaaS, AI tools, and OAuth apps.
- Automated Workflows: Employee notification, manager escalation, ticket creation, and remediation, all triggered automatically from CrowdStrike detections.

Appendix: How Anzenna Works, The OODA Framework

Anzenna operationalizes insider risk through a four-stage methodology that turns fragmented data into coordinated action.

Observe

Collect Without Agents
Collect user activities across SaaS, Cloud, endpoint devices, and data movement to build a complete behavioral baseline without agents.

Orient

Enrich with Context
Contextualize raw signals with identity, HR status, and historical behavior patterns to distinguish normal work from actual risk.

Decide

Investigate with Agentic AI
A fleet of specialized AI agents investigates every alert, each handling a different stage of analysis. Your analysts receive finished briefs, not raw data.

Act

Remediate Immediately
Execute automated remediation actions (block, revoke, notify) directly from the platform with or without human approval.

See Anzenna in Action

Ready to see what your CrowdStrike investment looks like with an agentic layer on top? Request a demo at anzenna.ai/demo. In 30 minutes, we'll show you how Anzenna turns your existing endpoint telemetry into a fully automated insider risk operation, from detection through remediation, with zero new agents and zero new hires.

Schedule a Demo

Anzenna

sales@anzenna.ai | www.anzenna.ai
The Agentic Layer for Insider Risk Management