INSIDER THREAT

Insider Threat Detection and Investigation

Behavioral analytics, automated investigation and threat hunting for insider risk programs — on the data you already have.

The problem insider risk teams are solving

Insider threat programs sit between security, HR, legal and physical security. The signals that matter are spread across all four, and most are individually unremarkable: a badge-in, a file deletion, a permission change, an outbound message. Risk shows up in the combination, not the event.

Most programs run on a UEBA platform or a UEBA module inside their SIEM. Both identify behavioral anomalies and leave the investigation to the team — producing hundreds of anomalies a day, a queue nobody can clear, and a program measured on volume rather than outcomes.

> **An anomaly is not an insider threat. It is the beginning of an investigation.**

Behavior analytics should be the beginning, not the end

AiStrike does not replace behavioral analytics. It makes behavioral analytics the first step in a larger operating model.

LEGACY	Telemetry → behavioral anomaly → risk score → human investigation
AISTRIKE	**Telemetry → behavioral anomaly → investigation → evidence → risk determination → governed response**

Once suspicious behavior is identified, AiStrike gathers and correlates evidence across identity, endpoint, cloud, physical access, DLP and applications — asking the questions an experienced insider-threat analyst would ask. Was this normal for this user? For their peers? What happened immediately before and after? Did physical and digital behavior align? What sensitive systems or data were involved?

How AiStrike approaches insider threat

Behavior analytics	Baselining, peer-group comparison, rare-event and first-occurrence scoring, and identity correlation — the full set of behavioral requirements an insider program needs.
Investigation that adapts	Every anomaly is investigated end to end. The path is driven by the evidence, not a workflow written in advance — so two identical starting signals do not produce two identical investigations.
Federated data access	Analytics run on raw telemetry in your existing SIEM, security data lake or enterprise data store, without requiring another centralized copy of all the data.
Search and threat hunting	Built-in search with prepackaged insider-threat hunt packs, so the team can sweep for multi-stage behavior rather than waiting for an event to breach a threshold.
Program separation	Insider investigations carry HR, legal and privacy obligations a shared SOC tenant was not designed for. AiStrike supports the access segregation and data-handling controls insider programs require.

A worked example

A legal department employee badges into a restricted R&D lab after hours, accesses an R&D system her peer group has never touched, deletes 342 files from three project folders in six minutes, then clears the Windows security log on the same host.

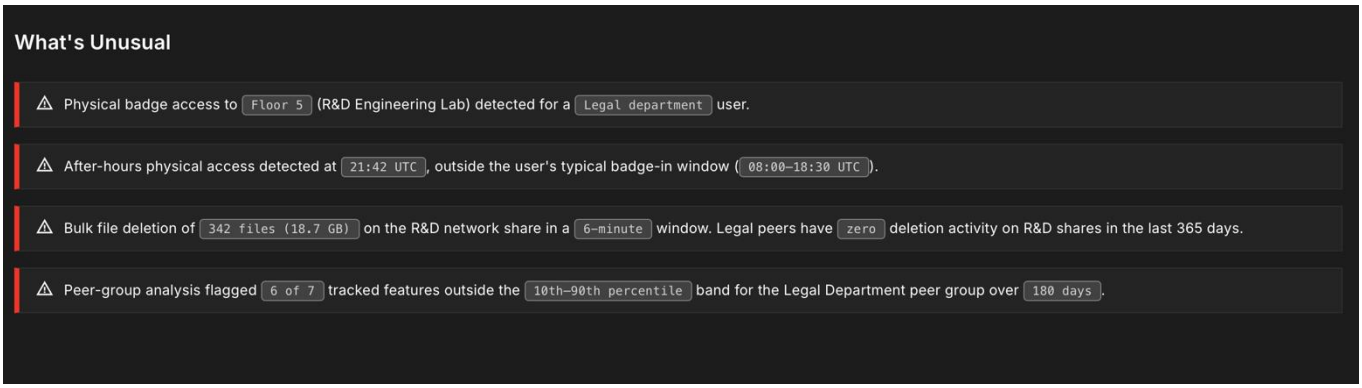

Four independent behavioral deviations, each with its own baseline and peer comparison.

Each deviation is individually explainable. Together they are premeditated sabotage with anti-forensics. AiStrike assembles the sequence across physical access, endpoint, identity and Windows event telemetry, maps it to the kill chain, and delivers an investigated finding with the supporting evidence — without an analyst having to connect the four.

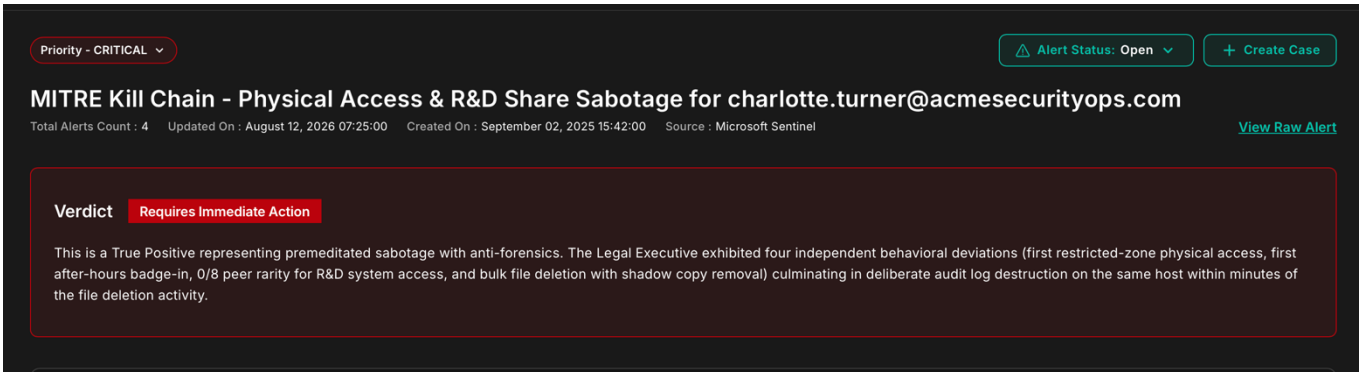

An investigated finding with the evidence attached — not another score.

The threat that never becomes an anomaly

Sophisticated insider activity is low-and-slow. Each action stays inside expected thresholds while the pattern develops over days or weeks. Nothing crosses a line, so nothing is ever scored.

INDIVIDUALLY NORMAL

Access	→	Discovery	→	Staging	→	Transfer
Day 1		Day 4		Day 9		Day 16

COLLECTIVELY SUSPICIOUS

The individual events look normal. The sequence doesn't.

Prepackaged insider-threat hunt packs search available telemetry for multi-stage behaviors and suspicious patterns, so the program can look for what the detection system has not surfaced yet.

Coverage

Mapped against a 48-scenario insider risk framework spanning 14 categories — from data handling and privileged access abuse through fraud, sabotage, anti-forensics and facilities access.

48
scenarios assessed

33
fully covered today

15
partially covered — typically dependent on an HR or third-party feed

14
risk categories

Representative scenarios

SCENARIO	COVERAGE	WHAT AISTRIKE DOES
Unauthorized disclosure of sensitive data	**Full**	Correlates DLP, CASB, email gateway and cloud-storage audit signals into a single investigated finding.
Unauthorized privilege escalation	**Full**	Investigates Entra ID, AD, Okta admin, IGA and PAM activity against the user's own and peer baselines.
Anti-forensics: log deletion, artifact hiding, account deletion	**Full**	Ties cleanup activity back to the originating behavior — the signal that separates sabotage from error.
Circumventing physical access controls	**Partial**	Covered where badge and identity telemetry is available; piggybacking without a badge event requires an additional feed.
Delegated execution via AI agents	**Partial**	Emerging capability — detects the behavioral footprint of agent-driven activity from API gateway, cloud audit and identity telemetry.

What the program gains

Findings, not anomalies

Every anomaly investigated end to end, so the team reviews prioritized insider threats instead of triaging volume.

One platform, one data set

Analytics run on existing telemetry in place — no duplicate ingestion and no parallel alert stream.

A program, not a queue

Search, hunt packs and case handling built for insider work, with the separation the program's obligations require.

Typical data sources

Identity and access (Entra ID, AD, Okta, PAM, IGA) · endpoint and EDR · DLP and CASB · email and collaboration audit · cloud audit (CloudTrail, Azure Activity, M365) · physical access control systems · HR feeds · ERP and finance systems · network, proxy and DNS.