Workspace Security Solutions Overview

The intent-aware Workspace Security platform for human and AI-driven work

Every major shift in how work gets done has reshaped cybersecurity. Networks expanded the perimeter. Endpoints became the front line. Cloud and SaaS changed where applications and data lived. Now AI is changing how people work, how information moves, and how quickly risk can become impact.

Today, work moves across a constantly changing digital workspace. Most of what employees do is legitimate. But under different circumstances, the same behavior can expose sensitive information, violate policy, or become the starting point for an attack.

The challenge is not that organizations lack security tools. It is that traditional tools often see activity without enough context to determine intent, including whether an action aligns with the user's role, the business workflow, the data involved, or the organization's policy.

Ent is the intent-aware Workspace Security platform that helps organizations protect human and AI-driven work as it happens. By understanding behavioral context across the flow of work, Ent helps security teams determine why risky actions are happening and apply the right friction before they become incidents.

Most security tools tell you what happened. Ent helps security teams understand why risky actions occur and intervene before they become incidents.

The New Security Challenge

Productivity itself has become the new attack surface.

Okta's 2025 Businesses at Work report found that the global average number of applications per customer topped 100 for the first time. As work spreads across more environments, security teams face a growing challenge: connecting individual actions and recognizing when routine work begins to create risk.[1]

Human behavior remains central to that risk. Verizon's 2026 Data Breach Investigations Report found that the human element was present in 62% of breaches. That matters because many of the incidents organizations care about most—accidental exposure, non-malicious risk, insider activity, AI prompt leakage, social engineering, and last-mile attacks can look like normal work until the outcome is difficult to reverse.[2]

This leaves security teams in a reactive loop: gathering logs, reconstructing timelines, and judging after the fact whether an action was expected, accidental, or harmful. By then, the data may already be outside the organization or the attack may already be underway.

Why Intent Matters

Intent answers a critical question: why is this action happening right now, and should it continue?

The same activity can mean different things depending on who is doing it and the business context around it. A developer pasting code into an approved AI tool to troubleshoot an issue may be following a normal workflow. In contrast, a marketing employee pasting that same source code into a personal AI account may be creating a data-exposure risk.

Activity tells security teams what happened. Intent helps determine whether the action should be allowed to continue.

This distinction matters because prevention depends on timing. If a tool can only explain risk after the action is complete, security teams are left with investigation and response. If the organization can understand intent while the user is still in the workflow, it can warn, guide, redact, or block where supported before the activity becomes an incident.

The Ent Approach

Ent adds the intent-aware layer. It brings behavioral and workflow context together at the moment of action, helping teams determine whether activity represents legitimate work, a mistake, or meaningful risk.

That context is what traditional tools often lack. EDR sees device events and process activity. DLP monitors governed data movement. Identity tools show logins and access. IRM and UEBA tools flag anomalous user activity but often without the business context to explain it. SIEM and SOC workflows bring alerts together for investigation. Each plays an important role, but none provides the full picture needed to understand why an action is happening in the moment.

Because Ent adds context to the existing security stack, organizations do not need to rip and replace the tools they already use. Ent extends EDR, DLP, SIEM, and SOC workflows with intent-aware context while helping teams rationalize monitoring-heavy insider-risk and UEBA tools over time.

01

Understand intent in real time

Ent reads behavioral context throughout the flow of work to understand why risky actions occur. Instead of treating each event as an isolated signal, Ent evaluates the sequence surrounding the activity: what the user is doing, where the information is going, how the action relates to the user's role, and whether the behavior aligns with organizational policy.

This reduces false positives caused by broad rules that lack business context.

02

Intervene with the right friction

When risk emerges, Ent helps organizations apply a proportional response. An employee making an honest mistake may need a warning or guidance to an approved workflow. An unusual action may require a business justification. Sensitive content may need to be redacted before submission. A higher-risk event may need to be escalated or blocked where supported.

The goal is to apply the right friction at the moment of risk so legitimate work can continue and avoidable incidents can be prevented.

03

Investigate with human-readable evidence

When a case requires review, Ent turns relevant activity into an annotated timeline that explains what happened, why it mattered, and what policy or business context made the action risky. Instead of forcing analysts to stitch together raw logs after the fact, Ent presents evidence in language that Security, HR, Legal, and Compliance stakeholders can understand and act on.

04

Extend the stack you already own

Ent is designed to strengthen existing security investments. Existing tools continue to perform their core functions, while Ent provides the behavioral and intent context needed to make better decisions in real time.

Ent can be deployed as SaaS or self-hosted inside the customer's cloud, giving organizations control over where security data resides.

Customer Outcomes

01

Govern AI without slowing adoption

AI adoption is moving faster than many governance programs can keep up. IBM's Cost of a Data Breach Report 2025 found that 63% of breached organizations studied lacked AI governance policies, while only 37% had approval processes or oversight mechanisms in place.[3]

Broadly blocking AI tools can slow innovation and push employees toward shadow AI, while unrestricted use can expose sensitive information. A more balanced approach is to understand AI usage in context, guide employees to approved workflows, and protect sensitive information before it is submitted to an inappropriate tool or account.

02

Prevent insider risk

Insider risk is not always malicious. By evaluating behavior in context, the platform helps teams distinguish malicious activity, non-malicious risk, honest mistakes, and legitimate work, so they can focus on the cases that matter and respond appropriately.

This helps organizations identify risky behavior earlier, reduce false positives, close investigations faster, and give HR and Legal stakeholders evidence they can understand without interpreting technical logs.

03

Protect data across modern workflows

Sensitive information can leave through AI prompts, browser sessions, collaboration tools, personal accounts, screenshots, and copy-and-paste workflows. Ent extends protection into these supported workflows, where data often moves beyond traditional governed channels.

04

Improve security operations efficiency

Security teams spend significant time reviewing false positives and reconstructing events across disconnected systems. With intent-aware context and human-readable evidence earlier in the process, teams can reduce unnecessary alert review, accelerate investigations, and collaborate more effectively across Security, HR, Legal, and Compliance.

Key Use Cases

Organizations apply intent-aware Workspace Security across several high-priority areas where traditional tools may see activity but struggle to understand meaning in the moment.

Insider risk

Without sufficient context, routine work and serious risk can produce similar signals, leaving investigators to reconstruct the situation after information has already been accessed, shared, or removed.

Ent evaluates activity with role, workflow, organizational policy, and surrounding behavioral context. When meaningful risk emerges, it can intervene while the action is taking place and preserve clear evidence for the teams responsible for review. This helps organizations identify risky behavior earlier, reduce false positives, close investigations faster, and give HR and Legal stakeholders evidence they can understand without interpreting technical logs.

AI and app usage control

Employees often adopt AI tools faster than organizations can approve and govern them. Netskope's Cloud and Threat Report 2026 found that the number of people using SaaS generative AI applications increased threefold in one year, while the number of prompts sent to those applications increased sixfold. During the same period, incidents involving users sending sensitive data to AI applications doubled, with the average organization seeing 223 incidents per month. [4]

Block lists alone may not solve this problem because employees can find new ways around them. Ent helps organizations understand supported AI usage, including situations where sensitive information may be typed, pasted, or uploaded. When risk is identified, it can explain the concern in plain language, request justification, redact sensitive information, or guide the employee toward a safer, approved way to complete the task.

Endpoint and AI data protection

Traditional DLP remains important, but when data moves through AI prompts, personal browser sessions, collaboration tools, screenshots, or clipboard activity, visibility can become fragmented. Content-only rules may also struggle to distinguish legitimate work from inappropriate data movement.

Ent augments existing DLP by evaluating data movement in context, including the user, destination, application, and workflow. This helps organizations determine whether an action is appropriate and apply a response that matches the risk.

This strengthens protection for sensitive information, including intellectual property, customer data, employee records, regulated information, and credentials, without requiring organizations to replace their existing DLP investments.

Last-mile threat prevention

Social-engineering and last-mile attacks increasingly begin with actions that appear to be ordinary user behavior. An employee may be instructed to copy a command from a website, open a system utility, install a fake update, or allow an external party to control their computer. Each step may look harmless when viewed independently.

Ent examines the sequence surrounding the activity and whether it aligns with the user's role and expected workflow. When the behavior indicates meaningful risk, it can guide the employee, interrupt the activity, escalate the event, or block supported actions. This helps organizations stop user-driven attack chains that rely on trusted users and approved tools before they become full security incidents.

CUSTOMER IMPACT

With Ent, a Fortune 2000 organization built and scaled a security program from zero to measurable outcomes in nine months:

9 months
from program launch to measurable outcomes

1,500+ users
including high-risk candidates, remote workers, users of interest

80+ investigations
conducted in partnership with HR and Legal

$5M+ in net impact
through role elimination and cost avoidance

50+ employment actions
supported by investigation findings

2.5 FTE
operating the program, scaling from zero in nine months

Protect Work As It Happens

Effective security should not force employees to choose between productivity and protection. When every situation is treated the same, employees experience unnecessary friction, while security teams inherit more alerts and investigative work.

The response should fit the circumstances. By understanding intent and applying proportionate controls, organizations can reduce risk while allowing legitimate work to continue.

This creates a security program that does more than monitor activity. It helps employees make safer decisions in the moment, gives security teams the context they need to act, and protects the business without slowing it down.

Security has evolved alongside every major shift in how people work. As AI reshapes the enterprise, organizations need a new security layer built for modern work. It must understand not only what users and AI-driven workflows are doing, but why those actions are occurring, and intervene before risk becomes an incident.

This is the role of intent-aware Workspace Security. Ent helps organizations protect work where it happens and as it happens by understanding intent, applying the appropriate level of friction, and extending the security tools teams already use. With Ent, organizations can move from reactive investigation toward prevention while preserving the productivity that drives the business forward.

See how Ent helps security teams detect insider risk, govern AI usage, prevent data loss, investigate incidents with the full behavioral context, and stop risky actions before they become incidents.

Visit ent.ai to learn more

COMPANY AND FOUNDERS

Ent was founded by Elias Manousos and Brandon Dixon, co-founders of RiskIQ (acquired by Microsoft) and members of the team behind Microsoft Security Copilot. Ent is deployed with Global 2000 customers across hospitality, financial services, and defense, and is backed by Decibel, Sequoia, Crosspoint Capital, Craft Ventures, Shield Capital, Felicis, and In-Q-Tel.

SOURCES

[1] Okta, Businesses at Work 2025 — okta.com/reports/businesses-at-work-2025
[2] Verizon, 2026 Data Breach Investigations Report — verizon.com/business/resources/reports/dbir
[3] IBM, Cost of a Data Breach Report 2025 — ibm.com/reports/data-breach
[4] Netskope, Cloud and Threat Report 2026 — netskope.com/resources/cloud-and-threat-reports