

EBOOK · INSIDER RISK


What DLP and UEBA Can't See

Two mature categories, billions in spend, and insider risk is still handled after the fact. This is an examination of why, and what the missing capability actually is.

DLP and UEBA are good technologies aimed at the wrong question. Both assume insider risk can be specified in advance. It cannot be, because the thing being detected is a person's intent, and intent does not appear in a policy or a baseline.

Inside: the reactive default · what DLP was built for and the five things a policy cannot express · what UEBA was built for and why deviation is not risk · the gap mapped across seven scenarios · insider risk as an investigation problem.

Above Security



THE PROBLEM

Most organizations are stuck being reactive

Ask a security leader whether insider risk matters and the answer is immediate. Ask them to describe their insider risk program and the conversation slows down. That gap between recognition and capability is the subject of this book.

62%

of breaches involved the human element, up from 60% the prior year

VERIZON DBIR, 2026
\$4.92M

average cost of a breach that began with a malicious insider, the costliest vector measured

IBM COST OF A DATA BREACH, 2025
260 days

mean time to identify and contain a malicious insider breach


IBM COST OF A DATA BREACH, 2025
45%

of employees are now regular AI users on corporate devices, up from 15% in a year

VERIZON DBIR, 2026

Functioning insider risk programs have historically existed in a narrow band of organizations: large financial institutions, intelligence agencies, heavily regulated enterprises. Not because those organizations understood the problem better, but because they could afford what it actually demands. Standing investigative teams. Permanent coordination between security, HR, and legal. Case building by hand, one incident at a time.

Everyone else inherited a default, and the default is a chain that only starts once the damage is done:



THE SHAPE OF THE PROBLEM

That default is forensics, not risk management. By the time security is engaged, the decision that mattered was made weeks earlier. The work that follows is documentation, and documentation is not prevention.

HOW LEGACY TOOLS TRIED AND FAILED TO SOLVE THE PROBLEM

DATA LOSS PREVENTION

- Requires predefined static rules
- Holds up on the known cases and fails on the edge cases, which is precisely where insider risk lives

UEBA


- Flags anomalies without context
- Tells you something changed, never whether the change matters

SECTION ONE · DLP

DLP is not broken. It is answering a different question.

The DLP model is coherent and, on its own terms, effective. It remains the backbone of most compliance programs for good reason, and the sequence it runs has not fundamentally changed in twenty years.

THE DLP OPERATING MODEL



That sequence was designed for a world that made sense at the time. Data lived in files, files lived on endpoints, and they left through a small number of legible exits: email, removable media, print. Watching the exits was a fair approximation of watching the risk.

DLP ACQUIRED THE INSIDER RISK MANDATE LARGELY BY DEFAULT

It was the only control positioned to observe what employees did with company data, so when insider risk arrived on the roadmap it landed on the tool already standing there. That was an accident of adjacency rather than a design decision, and much of the frustration security teams feel toward DLP traces back to it.

Where DLP earns its keep

PLAY TO THE STRENGTH

Regulated data with stable definitions. Card numbers, health records, and national identifiers look the same on Tuesday as they did on Monday. Known channels with known rules, and demonstrable control for an auditor.

Where the model runs out of road

THE STRUCTURAL LIMIT

Insider incidents are not fundamentally about files moving. They are about a person deciding something. DLP observes the artifact of that decision, sometimes, if the artifact happens to cross a monitored boundary.

DLP THINKS IN DOCUMENTS



The unit of analysis is the object and its movement: this file, this channel, this rule. Precise, auditable, defensible. Also the reason a DLP console can be entirely healthy while an insider incident is underway three feet away.

SECTION ONE · DLP

Where DLP falls short: 5 things a policy cannot express

These are not tuning problems. A rule is a conditional statement evaluated against an observable event, and the following five things are either not observable, not events, or not reducible to a condition.

THE BLIND SPOT	WHY THE MODEL CANNOT CLOSE IT	TYPE
Intent	A policy compares an action to a rule. It has no place to store <i>why</i> . Uploading a file to personal cloud storage is one event with two entirely different meanings, and the rule sees one.	UNRESOLVABLE
Sanctioned-channel abuse	Most insider harm travels through approved tools. The CRM export, the Drive share, the direct message. The channel is allowed, the volume is within norms, the person is authorized.	INVISIBLE
Behavior over time	DLP evaluates events one at a time. Insider risk is a sequence: access broadens, patterns shift, then data moves. Each step clears policy in isolation. The arc is the signal, and the arc is not an event.	NOT MODELED
Actions that move no data	Access probing, privilege accumulation, credential sharing, circumventing a control, holding a second full-time job on company time. Nothing leaves, so nothing triggers.	OUT OF SCOPE
The classification dependency	DLP protects what has been labeled correctly. Anything unlabeled, newly created, screenshotted, retyped, or pasted into a browser sits outside the perimeter the policy defines.	STRUCTURAL

THE ATTEMPTED WORKAROUND: POLICIES

Every row above has a workaround, and most teams have built one. The workarounds are policies. Policies require maintenance, maintenance requires people, and the people are the resource the organization did not have in the first place.

A WORKED EXAMPLE: SHADOW AI


Verizon's 2026 DBIR examined more than 850,000 DLP events involving generative AI tools. 45% of employees are now regular AI users on corporate devices, up from 15%, and 67% sign in with non-corporate accounts. Source code is the most submitted data type by a wide margin. DLP saw all of it. What no DLP event can tell you is which paste was an engineer debugging a function under deadline and which was someone assembling a portable copy of the codebase. In the log, those two are the same record.

SECTION TWO · UEBA

UEBA asked a better question

UEBA started from the right instinct. Stop enumerating forbidden actions and start describing normal people. Build a baseline for each user and each peer group, score the deviations, surface what stands out.

THE UEBA OPERATING MODEL



That instinct was correct, and it is why UEBA still belongs in the conversation. It moved the industry from *what is prohibited* to *what is unusual*, a real step in the direction of behavior. For compromised accounts behaving unlike their owners, it works well.

UEBA WIDENED THE VIEW AND THE NUMBER OF ALERTS



UEBA widened the field of view considerably. It also multiplied the number of things demanding attention, and every one of them is a question rather than an answer.

The peer group is where the model shows its seams. Comparing someone to their colleagues assumes the colleagues are a meaningful reference class, which holds for a call center and dissolves in a modern org chart. A platform engineer, a contractor with temporary access, and a sales leader running a territory do not share a distribution. Draw the groups broadly and everyone looks normal. Draw them narrowly and everyone looks anomalous.

THE ARITHMETIC NOBODY RUNS BEFORE PURCHASE

UEBA's output is a score, and a score is statistical distance from a baseline. It is not a finding, it is a prompt for a person to go and produce one. Resolving one alert properly means checking the person's role, their recent changes, whether their manager knows, and the surrounding business context. Twenty minutes on a good day, two days on a bad one. Multiply by the queue.

SECTION TWO · UEBA


Where UEBA falls short: deviation isn't necessarily risk

People depart from their own history constantly, and almost always for legitimate reasons: a promotion, a new project, a deadline, travel, a reorganization, covering for a colleague on leave.

WHAT THE MODEL SCORES	THE BENIGN READING	THE READING THAT MATTERS
400 CRM records exported at 11pm	Building a board deck against a Monday deadline	Packing a book of business before resigning on Friday
First-ever access to the finance share	Covering for a colleague on parental leave	Probing for compensation data ahead of a dispute
Uploads to a personal cloud account	Moving holiday photos off a work laptop	Staging intellectual property somewhere retrievable
Sign-in from an unfamiliar country	A conference nobody logged in the travel system	A second full-time job at a competitor
Print volume triples in a week	Onboarding three new team members	Making an offline copy no control will ever see

Identical telemetry, identical score, two completely different futures. Nothing in the behavioral data separates them, because the separating fact is not behavioral data. It is context: what this person's role changed to last month, whether their manager sanctioned the project, whether they resigned on Friday afternoon.

WHAT TEAMS DO WITH A SCORE THEY CANNOT READ



THE TREADMILL, READ HONESTLY

Every tuning decision in that cycle is defensible. The aggregate is a system tuned toward silence, and a quiet insider risk tool is indistinguishable from a working one. Worse: because the readings in the right-hand column above are rare, a well-tuned engine learns to treat them as exactly the kind of outlier it should suppress.

THE GAP

Why DLP and UEBA can't give the full insider risk picture

DLP and UEBA look like opposites. One defines the forbidden, the other defines the normal. Underneath they share a premise: that insider risk can be specified before it happens. Write down what each actually answers and the gap becomes obvious.

DLP ANSWERS

- Did sensitive data violate a defined policy?
- A yes or no, about an object, against a rule someone wrote in advance

UEBA ANSWERS

- Did a person behave anomalously?
- A distance from a baseline, about an individual, against a model built in advance

INSIDER RISK REQUIRES

- What is actually happening, does it represent risk, why, what else is connected to it, and what should we do?

The third question is not a harder version of the first two. It is a different kind of question. The first two can be answered by a system that compares a thing to a stored definition. The third can only be answered by something that reasons about a situation.

FIVE NORMAL EVENTS, ONE ABNORMAL INVESTIGATION

An account executive over two weeks. Every line is authorized, in-policy, and typical for the role:

**DAY
1**

Signs in to Salesforce and opens a strategic account plan

AUTHORIZED**DAY
3**

Runs the account data through a sanctioned AI assistant

APPROVED TOOL**DAY
6**

Connects a personal application to a work account

PERMITTED**DAY
9**

Communicates with an external destination not seen before

NORMAL FOR SALES**DAY
12**

Working patterns shift the way they do before someone resigns

NOT AN EVENT AT ALL

No policy was violated, so DLP is silent. No single action clears a tuned threshold, so UEBA is quiet. The risk was never in any one row. It was in the sequence, and the sequence is the thing neither category is built to hold.

THIS NOW APPLIES TO ACTORS WHO ARE NOT PEOPLE

Enterprises are handing AI agents legitimate identities, OAuth grants, and SaaS permissions. An agent with standing access to Salesforce or a code repository is not an intrusion to be detected. It is an authorized actor taking technically legitimate actions, which is the shape insider risk has always had.

THE GAP

7 common scenarios: what's still missing

The right-hand column is the part of the job no product in either category performs. Read it as a description of unstaffed work.

SCENARIO	WHAT DLP SEES	WHAT UEBA SEES	WHAT IS STILL MISSING
Data misuse	A policy hit, if the data was labeled and the channel monitored	A volume or timing deviation with a score attached	Whether this person had a legitimate business reason today
Inappropriate access	Nothing. No data left the environment	A first-time-access flag, usually low confidence	Whether the access matches a role they now actually hold
Overemployment	Nothing. This is not a data event	Odd hours and unfamiliar locations, both common in hybrid work	The pattern across weeks that distinguishes a second job from a flexible schedule
Policy circumvention	The bypass itself, if the workaround happens to be monitored	A gap in expected activity, which reads as quiet, not suspicious	That the person is deliberately routing around a control they know exists
Negligent behavior	The violation, after it has already occurred	Little. Careless behavior is often perfectly typical for that user	The moment before the mistake, when a nudge would have prevented it
Coercion	Nothing distinguishable. The actions look authorized because they are	A behavioral shift that resembles a hundred benign life events	The sequence and context suggesting the person is acting under pressure
Malicious insider	The final exfiltration, if it uses a monitored path	Deviations that were individually tuned out months earlier	The buildup: access broadening, then staging, then departure

The right-hand column is not asking for more data. In most of these rows the data already exists somewhere in the environment. What is missing is the reasoning that turns scattered signals into a conclusion about a person.

READ THE TABLE TWICE

The first reading is a product gap analysis, and it is the less interesting one. The second is a staffing analysis. Every cell on the right is a task performed by a human being who understands the organization, or it is not performed at all. That is why functioning insider risk programs have correlated with headcount rather than tooling budget.

ABOVE SECURITY

Insider risk is an investigation problem

Every entry in that right-hand column is the same kind of work, and it is the work a human investigator does: who is this person right now, what is normal for them, what changed, why, and does it indicate risk. Above Security is built on the premise that this is the actual requirement.

Rather than asking teams to define risk in advance, a fleet of AI investigative agents observes activity across identities, SaaS applications, endpoints, data access, and collaboration tools, then reasons about it in real time. Two things make that reasoning possible. The first is first-party telemetry at the browser and application layer, which is where most modern work actually happens and which carries semantics a process or network log cannot: not merely that a file moved, but what was opened, in which application, in what order. The second is an **evidence graph** that ties people, identities, data, and actions together over time, so a sequence spread across three weeks is a single object rather than a query somebody has to think to run.

That structure is then grounded in the specific organization: HR data, reporting lines, prior investigations, and analyst feedback on what turned out to matter. A generic model of suspicious behavior cannot make the distinction the tables in this book keep returning to. A model that knows this person changed teams in June, and that a near-identical pattern was cleared last quarter, can.

An alert

WHAT THE CATEGORIES PRODUCE

A rule name or a score, a user ID, a timestamp, a severity. Everything that makes it actionable is absent, so an analyst opens four consoles and starts reconstructing. The analysis has not begun.

An investigation

WHAT ABOVE PRODUCES

A behavioral timeline, the contextual analysis around it, the reasoning behind the classification, and recommended actions. It lands in case management, where security, HR, and legal work from one record.

BUILT FOR THE STACK YOU ALREADY RUN

Insider risk programs stall on deployment as often as on detection. Of 150 CISOs Above interviewed before launch, 95 said plainly that they did not want another endpoint agent. So Above reads from the EDR already deployed, and enforcement runs through controls the organization already trusts: a block pushed into Microsoft Purview, an action taken through CrowdStrike Falcon, coaching delivered in the browser. Existing DLP does not get ripped out. It gets something upstream that can tell it when to act.

Prevention works the same way. Most insider risk is not a malicious employee stealing data, it is people reaching for the fastest tool in front of them, so real-time nudging intervenes in the moment with a contextual warning or a safer alternative. Careless behavior never arrives as a case at all.



One framework, not seven modules. All seven scenarios on the previous page run through the same behavioral investigation framework, because they are the same question asked about different people. On privacy: the legacy answer was blanket screen recording, which pays a maximum privacy cost for a minimum signal return. Reasoning over focused telemetry needs less collection, not more.

CLOSE

The gaps, closed

THE GAP	HOW A CONTINUOUS INVESTIGATION MODEL ADDRESSES IT
A policy cannot represent intent Section One	Intent is inferred from sequence and context rather than declared in a rule: what preceded the action, what the person's role is, what changed around them.
Sanctioned-channel abuse is invisible Section One	The unit of analysis is the person, not the channel. Authorized actions taken for unauthorized reasons are exactly what behavioral reasoning looks for.
Events are evaluated in isolation Section One	Investigations are built on behavioral timelines, so the arc across weeks is the primary object rather than a byproduct of correlation rules.
Deviation does not equal risk Section Two	Continuous reasoning about role, team, and organizational context separates the benign reading from the one that matters. The system does the separating.
The tuning treadmill Section Two	There is no threshold to tune toward silence. Output volume is governed by how many situations genuinely warrant investigation, not by sensitivity settings.
Risk cannot be predefined The gap	Nothing is defined in advance. The model reasons about behavior as it forms, which is why it does not drift as the organization changes.
Prevention arrives too late The problem	Real-time nudging moves the intervention point ahead of the incident, so a share of the work never becomes an investigation at all.

None of this requires hiring the investigative team the organization could never quite justify. That is the point. The capability that belonged only to institutions with dedicated investigators becomes available to everyone else. And it is worth being clear about the realistic alternative, which for most organizations is not a competing product. It is the status quo: a policy engine nobody has retuned in a year, an anomaly queue nobody reads, and a forensic scramble whenever HR calls.

WHERE TO GO NEXT

If the right-hand column of the scenario table on page 8 describes work your team is doing by hand, or work nobody is doing at all, that is the conversation worth having. Above Security is a proactive insider risk platform that replaces static rules with continuous behavioral investigation. Get in touch at above.security.

SOURCES

IBM Security and Ponemon Institute. Cost of a Data Breach Report 2025. Malicious insider cited as the costliest initial attack vector at an average of \$4.92M, with a mean time to identify and contain of 260 days.

Verizon. 2026 Data Breach Investigations Report. Human element present in 62% of breaches; 45% of employees regular AI users on corporate devices, up from 15%; 67% signing in with non-corporate accounts; more than 850,000 DLP events involving generative AI tools.

