ClearSkies™

CORE FUNCTION BRIEF

UEBA

User and Entity Behavior Analytics, a core function of the ClearSkies iISOC platform

Detection that learns what normal looks like for each person, then notices when it breaks.

Contents

SECTION 1

Executive Summary

IN ONE SENTENCE.

UEBA is the core function that learns what normal activity looks like for each individual user and entity, then raises a detection when live behavior breaks that pattern, when an observed value matches known-malicious threat intelligence, or when correlation reveals what a single log source could not show on its own.

Signature-based detection recognizes what it has already been told to look for. That is its strength and its limit. A tool an attacker wrote yesterday has no signature. A set of actions that is entirely ordinary for an administrator can be the clearest possible warning sign when performed by someone in accounts payable. And an endpoint alert that names a device but not a person costs an analyst time at exactly the moment speed matters most.

UEBA answers this by learning rather than matching. Over a training period, the function builds a personal profile of normal activity for each user: which programs they run, which destinations they contact, and at what volume. Once that baseline is established, activity outside it raises a detection, whether or not the behavior has ever been seen anywhere before. Alongside the learned baselines, the function checks observed values against curated threat intelligence and correlates alerts to the user actually responsible for them.

The catalog currently comprises eleven detections across four categories, built on three analytic methods. Each detection publishes into the same shared model every other core function reads, so a behavioral finding is scored, routed and acted on through the same pipeline as any other detection rather than in a separate console.

THE FUNCTION IS PRECISELY SCOPED.

UEBA produces behavioral detections. It does not compute how urgent one is, which is the Risk Scoring Formula. It does not decide who works it, which is Intelligent Alert Routing. It does not decide or execute a response, which is SOAR and Playbooks. And it does not curate the threat intelligence its reputation-matching detections consume, which is Threat Intelligence. Each of those is a separate core function with its own brief.

11 detections

Across four behavioral and reputation categories

3 analytic methods

Learned baseline, reputation match, correlation

Per-user normal

Not one global threshold applied to everyone

SECTION 2

Why Signature Matching Falls Short

Detection content written against behavior somebody has already seen cannot, by construction, recognize behavior nobody has seen yet. Four failure modes follow, and each is answered by a named stage of the mechanism in section 3.

2.1 Threats with no signature to match

Attackers routinely bring in custom tooling, scripts, or malware built for a single campaign. Nothing about it appears on any list, so a detection approach that works by comparison against known-bad has nothing to compare against and stays silent.

2.2 Activity that is only unusual for this specific person

Running a remote administration tool is unremarkable for an IT administrator and a serious warning sign for someone in a finance role. A single global threshold applied across an organization cannot express that difference, so it is set either loose enough to miss the finance case or tight enough to bury the administrator in noise.

2.3 Alerts that do not name a person

An endpoint tool raises an alert about a device. Working out who was actually logged in and active at that moment is manual work an analyst performs under time pressure during an active incident, and the answer determines who to contact and what to contain.

2.4 Checks tied to a single log source

A reputation check implemented inside one product only inspects that product's logs. The same malicious destination contacted through a different system, and recorded in a different log, passes unexamined simply because nothing was looking at that log.

2.5 The cost of the workaround

The usual response is more rules and tighter thresholds, which increases alert volume without addressing any of the four gaps, and a manual enrichment step where an analyst reconstructs identity and context by hand. Cost scales with alert volume and with the number of log sources in scope, while the underlying problem, detection that has no notion of what is normal for a specific person, is untouched.

The scale of that volume is measurable. Microsoft and Omdia, in the State of the SOC, 2026, put the false-positive share of alerts at 46%, and found that highly fragmented security operations spend around 40% more on operational labor than consolidated ones. Neither figure is improved by adding rules, which is the response the four gaps above usually attract.

The share of intrusions that turn on valid credentials or previously unseen tooling, the two things a behavioral baseline is built to catch, is not sourced for this brief.

SECTION 3

How UEBA Works

UEBA is one continuously operating mechanism native to the TDIR core, not a separate analytics product. Normalized activity enters at one end, and a behavioral detection carrying its method, its category and the responsible user leaves at the other.

> *Figure 1. The five-stage behavioral analytics mechanism.*

3.1 The five stages

Each stage is named once here and referenced by name throughout the rest of this brief, paired with the failure mode from section 2 that it answers.

STAGE	WHAT HAPPENS	THE FAILURE MODE IT ANSWERS
Baseline	Over a defined training period, a personal profile of normal activity is built for each user and entity: the programs they run, the destinations they contact, and the volume at which they contact them.	Activity that is only unusual for this specific person
Resolve	Observed activity is attached to the user or entity actually responsible for it, drawing on the normalized entity model the Centric-AI Fabric maintains.	Alerts that do not name a person
Evaluate	Live activity is tested by the method applicable to it: deviation from the learned baseline, a volume anomaly against that baseline, a match against curated threat intelligence, or a similarity comparison against the organization's own trusted domains.	Threats with no signature to match
Detect	A detection is raised carrying its analytic method, its category and the resolved user, so a reviewer can see not only what fired but on what analytic basis.	Checks tied to a single log source

STAGE	WHAT HAPPENS	THE FAILURE MODE IT ANSWERS
Publish	The detection is written into the shared model, where the Risk Scoring Formula scores it and Intelligent Alert Routing assigns it. Baselines are refreshed as legitimate behavior changes.	Threats with no signature to match

Reputation-matching detections skip the learned baseline entirely and enter at Evaluate, because a known-malicious value is malicious regardless of whose behavior it appears in. Every detection still passes through Resolve and Detect, so a reputation match is attributed to a person exactly as a baseline deviation is.

3.2 The three analytic methods

Every detection in the catalog uses one of three methods, and this brief names which applies to each, because the methods carry different evidentiary weight.

METHOD	HOW IT DECIDES	WHAT IT IS GOOD AT, AND WHAT IT IS NOT
Learned per-user baseline	Compares live activity against a profile of that specific user's normal, built over a training period.	Catches previously unseen tooling and behavior. Requires a completed training period before it produces useful output.
Reputation and similarity matching	Compares an observed value, an IP address, a domain or a full web address, against curated threat intelligence or against the organization's own trusted-domain list.	Immediate coverage with no training period, and independent of which log recorded the value. Limited to what intelligence already knows, except for the similarity check, which catches an unreported look-alike domain on first use.
Correlation	Attaches identity and context to an alert another tool has already raised.	Removes manual enrichment work. Adds no new detection of its own; it makes an existing one actionable.

3.3 How a baseline is judged

Quality is measured rather than assumed, on definitions a provider can report and a customer can audit. The measures below are reported natively and per tenant, defined here because the same metric name is measured differently by different vendors. Target values are agreed per engagement rather than asserted.

MEASURE	THE DEFINITION USED	WHAT A POOR VALUE TRIGGERS
Baseline maturity	In-scope users and entities with a completed training period, as a share of the total.	A review of coverage gaps, and of whether the training period suits that population.
First-seen rate stability	The trend in first-seen detections per user after training completes, rather than during it.	Investigation of whether the baseline is converging or the training window is too short.
Identity resolution rate	Endpoint alerts attached to a responsible user, as a share of all endpoint alerts received.	A review of the identity and session data available for the affected systems.
Detection mix by method	The distribution of raised detections across learned baseline, reputation matching and correlation.	A check on whether baselines are contributing at all, or reputation matching is carrying the function alone.
Source coverage	Log sources in scope for the reputation-matching detections, as a share of ingested sources carrying a checkable value.	A review of which ingested logs are being left unexamined.

A first-seen detection is not by itself a finding. ***Every user legitimately runs a new program or contacts a new destination sometimes. A first-seen detection states that something has not been observed before for this person, which is a reason to look rather than a conclusion. The Risk Scoring Formula weighs it alongside everything else known about the entity, which is why UEBA publishes rather than escalates.***

SECTION 4

UEBA in the Platform

4.1 What it draws on, what it writes back, and what it does not do

What it draws on. Normalized telemetry from the sources the platform collects: Windows event and process logs delivered through Endpoint Threat Monitoring and Response, firewall logs, web gateway logs, VPN authentication logs, and alerts raised by endpoint detection tools. It draws the normalized

entity model the Centric-AI Fabric maintains, used at Resolve, and the curated indicators Threat Intelligence publishes, used by the reputation-matching detections at Evaluate. Each is an input to a specific stage rather than a general capability UEBA generates itself.

What it writes back. Behavioral detections, each carrying its analytic method, its category and its resolved user, published into the shared model. The Risk Scoring Formula scores them, Intelligent Alert Routing assigns them, and where a detection is confirmed, SOAR decides the response. The relationship is one way: nothing returned to UEBA changes a detection already published, and a baseline changes only because observed behavior changed, never because a downstream function asked for it to.

What it does not do. It does not compute a risk score, decide assignment or execute a response. It does not curate the threat intelligence its reputation-matching detections consume, which is Threat Intelligence, and it does not author the platform's rule-based detection content or attach ATT&CK technique identifiers, which is Detection Factory. Where a hunter tests a hypothesis against the same telemetry, that is Threat Hunting: UEBA runs continuously and unprompted rather than in response to a hypothesis.

4.2 The detection catalog

The catalog currently comprises eleven detections across four categories. The table states the analytic method behind each and the telemetry it reads, because the method determines what a detection can be relied on to catch and the source determines where it can catch it.

CATEGORY	DETECTIONS	METHOD AND TELEMETRY
Endpoint and process behavior	Security-oriented event detection; first seen process; malicious endpoint activity	Rule-based against a published event catalog, learned baseline, and correlation respectively, over Windows event and process logs and endpoint detection alerts
Network behavior	First seen IP address; unusual network activity by IP address; unusual network activity by destination port	Learned baseline, and volume anomaly against that baseline, over firewall logs and endpoint network visibility
Threat intelligence and reputation	Malicious IP address; malicious domain; malicious URL; domain similarity	Reputation matching against any ingested log carrying the relevant field, and similarity comparison against the organization's own trusted domains over web gateway logs
Identity and access	Geolocation logins	Impossible-travel check against consecutive authentication events in VPN logs

The reputation-matching and similarity detections apply to any ingested log containing the relevant value rather than to one product's logs, which is what gives them coverage across the environment. The specific VPN and web gateway products currently in scope are a deployment matter rather than a fixed platform property.

4.3 A worked example

The chain below is illustrative and technically representative rather than drawn from a named engagement. A finance user's workstation runs a remote administration utility for the first time. Baseline has a completed profile for that user and has never observed the program, so nothing about it matches their normal. Resolve attaches the detection to the specific person logged in at that moment rather than to the device alone. Evaluate raises a first seen process detection, and separately, when the same session contacts an external destination, a first seen IP address detection. Detect records both with their method and category. Publish writes them into the shared model, where the Risk Scoring Formula weighs two independent first-seen signals on a privileged-data-adjacent account and scores the pair well above either alone. For an IT administrator, the identical program would have sat inside an established baseline and raised nothing, which is the whole point of a personal profile.

4.4 Delivery across many tenants

A baseline is personal by construction: it describes one user in one tenant and has no meaning outside that tenant. What travels between tenants is the analytic method and the detection catalog; what never travels is one tenant's baselines, telemetry or detections. A provider therefore runs the same eleven detections across every customer while each customer's notion of normal stays entirely its own.

Because the function is switched on with the platform rather than deployed separately, time to value has two distinct phases, which this brief states plainly rather than blurring. In the first week, the reputation-matching and correlation detections produce output immediately, since they require no training period. Through the first training cycle, the learned-baseline detections are still observing and produce a higher first-seen rate that settles as profiles mature; a tenant is told to expect this rather than surprised by it. At steady state, baseline maturity and first-seen rate stability are tracked as ongoing measures.

SECTION 5

ATT&CK Alignment and Supported Frameworks

UEBA does not map detection content to technique identifiers itself. Where a behavioral detection is translated into mapped detection content, Detection Factory attaches the ATT&CK identifier at that point, and the identifier travels with the detection from there. This function's own contribution is the behavioral finding and its resolved entity, not the technique mapping.

FRAMEWORK OR STANDARD	ITS ROLE IN UEBA
MITRE ATT&CK (Enterprise, Cloud and Identity)	Attached by Detection Factory where a behavioral detection becomes mapped detection content. Neither computed nor validated by this function.
NIST Cybersecurity Framework 2.0	Outcome language used where platform reporting discusses anomalous-activity detection to an executive or board audience.
ISO/IEC 27001 and 27002	Control mapping and audit evidence, where behavioral monitoring of user activity is the control being evidenced.

Behavioral monitoring carries obligations of its own. ***A function that builds a personal profile of each user's activity is processing personal data, and in several jurisdictions that engages works-council consultation, employee notification or data protection impact assessment duties before deployment. The platform produces the evidence those duties call for. Whether a given deployment satisfies them is a legal determination the organization and its counsel make.***

Evidence, not compliance. The platform states which obligations behavioral monitoring produces evidence for. It never states that running the function makes an organization compliant with anything, and control-level framework mapping is held by the Regulatory Frameworks core function.

SECTION 6

Questions Raised in Evaluation

The questions below recur in buyer evaluations, each answered directly against the mechanism described in section 3.

QUESTION, AS A BUYER ASKS IT	ANSWER
How long before UEBA is useful?	The reputation-matching and correlation detections produce output immediately. The learned-baseline detections require a training period first, and this brief states that as a limit rather than working around it.
Do the reputation-matching detections need a baseline too?	No. They skip Baseline and enter the mechanism at Evaluate, because a known-malicious value is malicious whoever contacts it. They still pass through Resolve, so they name a person like every other detection.
Will first-seen detections bury the team in noise?	First-seen rates are highest during and shortly after training and settle as baselines mature. First-seen rate stability, defined in section 3.3, is tracked precisely so this can be observed rather than assumed, and the Risk Scoring Formula weighs a first-seen signal alongside everything else known about the entity rather than escalating it alone.
Does UEBA decide that a behavioral detection is serious?	No. It publishes the detection with its method and resolved user. The Risk Scoring Formula scores it and Intelligent Alert Routing assigns it.
What happens when a user's role legitimately changes?	Baselines are refreshed as observed behavior changes, so a sustained legitimate change is absorbed rather than alerting indefinitely. A sudden change still raises a detection at the point it happens, which is the intended behavior.
Is this monitoring employees?	It is behavioral detection over activity the platform already collects, and it processes personal data. Deployment obligations, including consultation and notification duties where they apply, are a legal matter for the organization, and the platform produces the evidence those duties call for.
Does UEBA replace signature-based detection?	No. It covers what signature matching structurally cannot, and four of the eleven detections are themselves reputation matches. The two approaches are complementary rather than alternatives.

SECTION 7

The Benefits

Behavioral detection earns its place by catching what nothing else in the stack is built to catch, and by removing work an analyst would otherwise do by hand mid-incident. The outcomes are countable: how much of the population has a mature baseline, how many detections arrive already attached to a person, and what share of detections come from learning rather than from matching a list.

7.1 At a glance

BENEFIT	WHAT PRODUCES IT
Tooling nobody has reported anywhere is still caught	Baseline learns what each user normally runs and contacts, so novelty is detectable without a signature.
The same action reads differently for two people	A profile is personal, so no global threshold has to be set loose for one population and tight for another.
Every detection names a person	Resolve attaches the responsible user at the point of detection, not as manual enrichment during an incident.
A reputation check follows the value, not the log	Evaluate inspects any ingested log carrying the relevant field, rather than one product's own data.
Coverage begins on day one and deepens after training	Reputation matching and correlation need no training period; learned baselines add depth as profiles mature.
Whether behavioral coverage is working is a number	Baseline maturity and detection mix by method, defined in section 3.3, make the answer reportable.

7.2 What changes, and for whom

Per-user behavioral detection changes what each role can see, and what each role no longer has to reconstruct by hand.

ROLE	WHAT CHANGES
SOC analyst	An endpoint alert arrives already attached to the person responsible, so investigation starts from the user rather than from working out who they were.
SOC manager	Baseline maturity and detection mix become visible numbers, so the question of whether behavioral coverage is actually working has an answer.
CISO and risk	Coverage extends to previously unseen tooling and to insider and compromised-account activity that signature matching is structurally poor at catching.
Service provider practice lead	The same eleven detections run across every tenant with no per-customer analytic build, while each tenant's baselines stay entirely separate.

7.3 What sets it apart

▸ **Normal is personal, not global.** A baseline describes one user, so the same action reads differently for an administrator and for a finance user without a rule being written for either.

▸ **Catches what has no signature.** A learned baseline flags a program or destination never seen for that person, whether or not it has ever been reported anywhere.

▸ **Checks travel with the value, not the log.** A reputation match inspects an IP address, domain or web address wherever it appears across ingested logs, rather than inside one product's own data.

▸ **Every detection names a person.** Identity resolution happens at the point of detection rather than as manual enrichment during an incident.

The value in one line. Signature matching answers whether the platform has seen this before. A personal baseline answers whether this person has, which is the question that catches the tool written yesterday and the credential stolen this morning.

SECTION 8

Summary

UEBA is how the platform detects what no signature describes. It is one core function among twelve, scoped to a five-stage mechanism that runs continuously inside the TDIR core rather than as a separate analytics product an analyst queries.

▸ Baseline, resolve, evaluate, detect and publish replace detection that can only recognize what it has already been told to look for.

▸ A profile of normal is built per user and per entity, so the same action reads differently for two different people without a rule being written for either.

▸ Eleven detections across four categories run on three analytic methods, and each detection states which method produced it.

▸ Every detection is attached to the responsible user at the point it is raised, rather than enriched by hand during an incident.

▸ The function detects and publishes, and leaves scoring, assignment, response and technique mapping to the core functions built for those decisions.

Forward-looking statement: any statement about future capability is forward-looking and non-binding, and capability described in the present tense is capability available at the publication date, subject to the deployment configuration in force. This document describes detection capability and does not constitute legal advice on any organization's obligations in respect of employee monitoring or personal data.