KNOWLEDGE GRID
OVERVIEW

Unsupervised Anomaly Detection (UAD)

Finding the threats you don't know to look for

Summary

Unsupervised Anomaly Detection (UAD) finds threats no one has told it to look for — no training data, no predefined rules, no prior threat knowledge required — by analyzing hidden and weak signals in data and reasoning over what is normal and surfacing what is not.

The problem with everything that came before

Almost every detection tool in a modern security stack rests on one of two foundations, and both share the same blind spot.

1) **Signature-based detection** finds only what is already known. It matches activity against a catalog of documented threats. It is effective against known malware and documented techniques — and blind to anything novel, by definition.
2) **Supervised and threshold-based detection** finds only what it has been trained or tuned to find. It requires labeled data or human-defined baselines, produces high false-positive rates when normal behavior varies, and misses slow, low threats that stay within statistical tolerances.

Both approaches answer the same question: "does this match something we have already recognized as bad?" Neither can answer the question that matters most in modern security — "is this normal?" — *without being told in advance what normal looks like*.

The unknown-unknown problem

The most damaging threats are the ones you don't know to look for. If detection depends on prior knowledge — a signature, a rule, IOC, or a labeled example — it structurally cannot recognize what has never been seen before.

How Knowledge Grid does it differently

Knowledge Grid's UAD is fully unsupervised. That means that it builds its own unbiased understanding of normal directly from your telemetry — across users, devices, networks, applications, events, and time — and surfaces meaningful deviations without predefining what to look for. By utilizing the platforms Rough Set mathematics foundation, UAD draws reliable conclusions from incomplete and ambiguous data, without requiring complete information or predefined structure.

Because it reasons from the data itself rather than from a catalog of known threats, UAD requires no training data, no human-defined rules, and no prior threat knowledge. It does not require turning or an extensive baselining period, and can begin reporting anomalies in as little as 48 hours, although 30 days is typical.

Six dimensions of anomaly

Attackers evade detection by staying normal on the axis you happen to be watching. Knowledge Grid watches six — multi-dimensional, single-dimensional, reverse, zero-day, diversity, and distinctive — so evading one still trips another.

Multi-Dimensional	Anomalies that appear when unsupervised features are analyzed in combination with and compared with multiple other attributes.

Reverse	The reduction or absence of expected activity — something that should be happening but isn't. Often the strongest signal, and the one threshold tools miss entirely.
Single Dimensional	Anomalies that appear only when a single attribute is compared to the baseline for that attribute.
Zero-Day	Behavior with no precedent in the observed history — novel patterns that no signature or rule could anticipate.
Diversity	Unusual variety or breadth in activity — an entity suddenly touching far more systems, accounts, or data types than its norm.
Distinctive	Behavior that sets one entity apart from its peers — outliers relative to a comparable population, not just to their own history.

Reverse anomaly detection is worth special note: the ability to detect the meaningful absence of activity is rare in the market and is one of the clearest signals of a sophisticated intrusion — an attacker disabling logging, a process that stops reporting, a heartbeat that goes quiet.

From detection to investigation

Detection is only useful if an analyst can act on it. Knowledge Grid pairs every detection with the context needed to respond.

- **Three generations of threat scoring model evolution** progressively refine how anomalies are ranked, so the most meaningful signals rise to the top.
- **Investigation narratives**, generated by Agentic AI, explain each anomaly in plain language — what happened, why it was flagged, and why it matters — turning a score into a story an analyst can act on.
- **Behavioral and relational context** travels with every finding, preserving the temporal and relational meaning that makes an anomaly interpretable.

Built on the Cognitive Data Platform

UAD isn't a standalone detector bolted onto your stack — it runs on the same Temporal Data Grid and knowledge structures that power the entire platform. That means every anomaly it surfaces arrives already enriched with behavioral history and relational context, so you can recognize not just *that* something is off, but *why*.

Why it matters

Traditional detection	Knowledge Grid UAD
Finds only known threats and defined patterns.	**Finds the unknown — no signature or rule required.**
Requires labeled data or tuned baselines.	**Fully unsupervised — works on day one.**
Detects one kind of anomaly, usually volume.	**Detects six dimensions of anomaly.**
Outputs an alert; the analyst reconstructs context.	**Outputs a scored, narrated, context-rich finding.**

See what Unsupervised Anomaly Detection surfaces in your environment.

Request a UAD briefing: solutions@knowledgegrid.com • knowledgegrid.com