# Breaking the Rules: When Static Detection Logic Reaches Its Limits, What's Next?

Correlation rules haven't kept pace with the evolution of cyberthreats. As a result, security operations centers (SOCs) everywhere are feeling the pressure.

Security professionals are intimately familiar with static rules: the basic programming logic that performs an action when specific conditions are met. They reliably detect the known triggers they were written for, flagging events for SOC analysts and supporting business needs. For example, an engineer can write a rule to fire every time a user is locked out of their account. More sophisticated rules can even identify the context, such as where and how the lockout happened.

However, engineers write these rules manually. This process becomes convoluted and high maintenance as an enterprise's needs and use cases evolve. Many cybersecurity professionals may remember a time where rule count was a standard benchmark for SIEM performance. The assumption was that more rules meant better coverage.

This doesn't square with reality. More rules don't equal more protection; instead, they can introduce duplication, redundancy, and alert fatigue. They also accrue technical debt. As the environment changes and new rules are written, what happens to the old ones? Often, they contribute only complexity. Rule hygiene—the laborious retirement of obsolete rules—becomes another time-consuming process for an already overstretched SOC.

For a simplistic line of defense, rules require considerable effort and resources. Engineers must write correlation rules to detect an event, create query logic for alerting, and then factor in additional if-then-else logic for unexpected activity. And all of this requires testing and upkeep. Detection engineers become stewards of an entire language within the SOC, creating content to support everything from collection to reporting. But this endless labor only leads to complexity.

## An Overreliance on Rules Can Put Security at Risk

Rules still contribute meaningfully to the SOC. For audit and compliance, organizations may need to demonstrate they have rules to perform various functions. The data they produce is valuable for threat detection.

The problem is that rules follow a strict logic. They only trigger for the specific events they're written for and reveal limited context. They may show that a user has been locked out after too many failed login attempts, but is the lockout due to human error or a brute-force attack?

To find out, analysts would need to search through myriad alerts or overwhelming raw logs to determine a coherent sequence of events. Or, they would have to write queries so long and convoluted that the effort would not be worth the benefit. Such practices are neither sustainable nor scalable with an under-resourced team.

All the while, the SOC remains oblivious to the most dangerous risks: adversary persistence, lateral movement, and insider threats. Rules are insufficient for connecting the dots between a user, their IP address, their host, and the series of actions they take. When rules fire, they reveal a single fingerprint of a fleeting moment. When attacks involve multiple actions across multiple systems, correlation rules simply fall apart.

Today, every threat is multifaceted. How relevant are rules in the SOC, and what should replace them? Before answering, it's worth reflecting on the recent trajectory of cyberattacks.

## When Did Correlation Rules Reach Their Tipping Point?

Think of the major cyberattacks from the last five years. The attacks were advanced, but so were the intricate and expensive technologies designed to stop them. How were the breaches successful?

In all cases, the attacks were highly adaptive and used varied tactics and techniques. They could insinuate themselves into their target environments and "live off the land," persisting without detection. And any threat detection, investigation, and response (TDIR) workflow that relies on flagging singular events instead of identifying a continuum of suspicious activity is powerless against such adversaries.

But there wasn't a single tectonic shift that upended established paradigms. The evolution of digital technologies has been both swift and subtle. The world doesn't look much different today than it did yesterday, but compared to five years ago, there have been significant changes. This is perhaps why correlation rules remain prevalent in threat detection, even though they can no longer form the foundation of a modern detection strategy.

Like the proverbial frog in the slowly boiling pot, these gradual yet radical shifts make it easy for SOCs to believe they have the right tools to guard against cyberthreats, even when they don't.

### The 2000s: Large-Scale Attacks and Polymorphic Malware

The diminishing returns of correlation rules in threat detection date back to the origins of polymorphic malware. Such programs, which could adapt at runtime to infiltrate different operating systems, appeared in the 1980s but had achieved scale by the 2000s. Anyone working in cybersecurity at this time remembers how tracking static signatures quickly became futile.

The 2000s redefined the potential scale of malware attacks. At the turn of the millennium, the ILOVEYOU computer worm marred digital commerce and communication. This malware invaded 45 million computers in 24 hours and eventually infected 10% of internet-connected computers worldwide. By this point, cybersecurity researchers were already looking beyond correlation rules. The 2010s would raise the stakes once more.

### The 2010s: The Rise of Ransomware and Nation-State Threats

With the mass adoption of cloud computing came the end of the traditional network perimeter. The SOC was no longer accountable only for defending data behind a firewall. This not only complicated threat detection but also incentivized cybercrime, as personal data and credentials became more valuable.

As a result, ransomware gained global traction in the late 2010s. If organizations weren't prepared, even the most sophisticated controls would take hours, days, or weeks to generate alerts.

With the vanishing perimeter and the increasing value of online data, nation-state actors become a more significant threat in the cybersecurity landscape. Attacks such as NotPetya in 2017 are still vivid. Russian state actors targeted a popular accounting software in Ukraine to unleash a malicious program that operated like ransomware. There was no way to decrypt files once corrupted, so it functioned more like a data wiper. The attack ultimately impacted more than 2,300 organizations in over 100 countries, with losses exceeding $10 billion.

These factors combined create a far more complex environment for SOCs. As data and systems move to the cloud, they're subject to new risks. As a result, organizations onboard more security solutions to maximize coverage. But these solutions create so much alert noise that more ends up being less.

### The 2020s: Generative and Agentic AI Is Adopted Worldwide

At the end of 2022, the public launch of generative AI programs that could craft convincing messages, images, and video changed the calculus again. In the short time since, it has accelerated in sophistication while attracting trillions of dollars of investment.

It has also placed more powerful tools in the hands of threat actors, who are more numerous and better funded than ever before.

AI doesn't just make social engineering attacks simpler; adversaries use it for reconnaissance to look for zero-day vulnerabilities. Now, agentic AI is emerging as an attack vector that can function like a highly adaptable insider threat. Whether it originates inside the organization or intrudes from outside, it can problem-solve as it moves through a system, finding new paths and creating new code. SOC analysts understand the challenges of attacks that dwell in the environment. For instance, living-off-the-land attacks can progress too slowly for conventional detection methods to catch, which is how an AI-based threat might behave.

The pace of AI adoption means the attack surface is expanding to include autonomous and intelligent elements that the SOC has limited experience in tracking. Monitoring them requires investment in the right capabilities, which is a challenge for many SOCs.

## Rules Are Static, But User Behavior Is Dynamic

With polymorphic malware and AI automations, adversaries don't stop, sleep, or blink. There's no longer time to research, write, and execute rules to combat them.

The world doesn't adhere to the rigid if-this-then-that logic of correlation rules. Rules and signatures are still relevant, but in modern cybersecurity, they require a mechanism to connect them to user activity that unfolds across a multitude of systems and stages.

Machine learning has been successful in this regard. Its capacity to recognize anomalies in large data patterns has proven indispensable. It is the structure upon which user entity and behavior analytics (UEBA) was built. Such technologies originally emerged in response to the variable threat of polymorphic malware. By baselining normal activity at the entity level, it helped the SOC focus on potential risks and follow their development over time.

Yet, different vendors approached machine learning in different ways—some more successfully than others—leading to mixed perceptions about its use. Deploying carefully weighted algorithms to catch nuanced behaviors is extremely challenging, and models must be tuned as the environment changes; it can't be assumed the algorithms will figure it out for themselves.

Not all abnormal behavior is a threat, and machine learning algorithms have been known to produce false positives when they can't piece together the full picture of user activity. As a result, some defenders have reservations about whether behavioral analytics can deliver the insights they need to stop today's threats.

### Powerful Behavioral Analytics Are a Modern Defense

At a time when businesses are championing artificial intelligence, it's worth remembering that machine learning *is* AI. Long before generative AI served as chatbots, cybersecurity developed behavioral analytics to answer critical questions that rule-based practices could not.

The need arose from real-world use cases where analysts repeatedly failed to identify gaps in their defenses. Machine-learning analytics filled a critical void, and as AI improves, so does its applicability in the modern SOC.

With behavioral analytics integrated into a security information and event management (SIEM) system, the SOC can see subtle changes that traditional controls miss, and take on advanced use cases like fraud and insider threat detection.

When the SOC is overly reliant on rules and other alert-generating mechanisms, more is less. There's too much noise for analysts to do their jobs. But with powerful behavioral analytics, less is more. Modern behavioral detection engines now have self-learning, self-tuning capabilities that automate threat detection far better than managing long lists of rules.

For example, Exabeam New-Scale Analytics features a new UEBA detection engine with far fewer rules than in previous iterations, yet it maintains the same level of MITRE ATT&CK® coverage. The engine models behavior across all cloud environments and encapsulates them with a single rule rather than having distinct rules for each cloud provider. This provides more adaptive and expansive monitoring with much less noise.

New-Scale Analytics also delivers more with less through an improved Session Data Model. Unlike conventional correlation logic that captures a snapshot in time, Exabeam UEBA uses an open-ended correlation window to track behavior over longer periods. This helps identify low-and-slow threats, such as living off the land (LOTL) attacks, while saving analysts from having to connect hundreds of disparate events.

Having this adaptive tool in the SIEM helps SOC teams make the most of their time. They can quickly understand risks without being distracted by hundreds of simultaneous detections.

Alert fatigue is a real and dangerous problem. It's crucial to automate effectively. While correlation rules still matter, their triggers should feed into a continuous timeline of activity rather than produce disjointed alerts. To streamline TDIR workflows, sophisticated, adaptive analytics are a must-have.

Ultimately, a leader's priorities are reflected in where they spend their time. CISOs and SOC leaders need to consider what they're dedicating their teams' time to, and if they're truly achieving results.

## Agentic AI in Modern Threat Detection

Agentic AI can be a terrible weapon for adversaries, but for the same reasons, it can be a powerful defense for SOCs. In addition to with enabling analysts to detect threats using plain language instead of complex queries, AI agents also allow for dynamic rule creation and monitoring.

By enabling adaptive risk scoring, agentic AI has already proven its capacity to reduce the false positives once associated with UEBA. It handles the grunt work of flagging abnormal behavior, scanning for related events, and grouping detections to automatically create cases so analysts can perform high-value work.

Does this revitalize the role of rules? Not entirely. Agentic AI may help identify areas where additional rules could serve a purpose and even help create them, but the fact remains that rules don't scale. As the list gets longer, they yield diminishing returns.

Rather, AI agents are the next layer for a modern SOC. Just as behavioral analytics builds upon correlation rules to provide context, agentic AI builds upon behavioral analytics to help analysts explain risks and automate their response.

This is how cybersecurity evolves. New technologies don't render the old ones obsolete; they complement each other to create a stronger security structure. Exabeam New-Scale Analytics is built with a newer, faster analytics engine, but the data science that drives detection hasn't changed. What's new is the data science of automation and response, thanks to AI.

AI is also changing what it means to be a detection engineer. Soon, handwriting-targeted signatures won't have as much relevance, but reviewing and managing rules will be necessary. AI assistants will write, adapt, and test rules, while engineers prompt and direct them.

AI is also changing analyst workflows. In the past, analysts had to investigate every alarm, which could take hours. When well-designed AI agents are paired with a powerful detection engine, those hours can be reduced to minutes. Instead of reading countless log file lines, analysts read an articulate, automated report and verify its findings. They can validate their understanding and accelerate their response times, reducing the organization's exposure.

Additionally, AI agents provide continuity in investigations, notifying analysts of developments and directing their attention to the most urgent needs. Information lost during shift turnover can compromise an investigation. One of the biggest benefits of agentic AI is its capacity to help SOC teams stay aware and engaged.

Automated summaries help analysts streamline investigations and rapidly catch up on critical information. Better yet, AI agents can retain awareness of specific cases, responding to requests in the context of the situation. This human-AI interaction will determine the efficacy of the next generation of SOC teams.

There's an important caveat: like UEBA, AI agents need to be conscientiously trained and tuned. An AI agent that pursues its goal without appropriate guardrails acts like an insider threat and exposes dire vulnerabilities.

This is why with Exabeam Nova, the multi-agent AI solution from Exabeam, extensive care was taken to constrain and direct the capabilities of each agent. They're purpose driven to execute specific tasks extremely well.

## A More Effective Way to Play by the Rules

Correlation rules still have a place in the SOC. Some organizations have well-crafted rules designed for their environment that continue to make a difference. Some handmade rules are simply business detection logic, which will always have value.

And machine analytics platforms still use rules to determine if anomalous activity is occurring. Rules will always be there, but how they are written and used will change, and their numbers will decrease.

In a modern SOC, correlation rules need to supplement more sophisticated methods of threat detection, rather than serving as the main means. SOC teams need a platform that supports the transition from rule-based to behavior-based detection so they can continue to use the workflows they excel at while leveling up their deployment of AI.

New-Scale Analytics powers the New-Scale Operations Platform and combines with Exabeam Nova. It serves as the default engine in Exabeam New-Scale SIEM and can be added to an organization's existing SIEM to drive smarter defense for any environment.

## About Exabeam

Exabeam is a leader in intelligence and automation that powers security operations for the world's smartest companies. As a global cybersecurity innovator, Exabeam provides industry-proven, security-focused, and flexible solutions for faster, more accurate threat detection, investigation, and response (TDIR).

exabeam

**Learn more at www.exabeam.com →**