A CISO's Guide to Adversary Alignment

"Are We Secure?"

This may be a CISO's least favorite question. Given the complexity of distributed environments and the constant evolution of threats, it's not a question that lends itself to a simple answer. Yet, executives and boards ask it often, expecting reassurance.

Even when the answer happens to be "yes" in a narrow moment, offering that response as a blanket statement isn't responsible. Adversaries, techniques, and conditions change continuously. A single control gap, misconfiguration, or human error can quickly shift the answer in the other direction.

A more useful way to frame the discussion is this:

Have we put the right controls in place to protect our most critical assets from our most relevant adversaries?

This reframing reflects three realities:

- Adversaries pursue objectives that vary by organization.
- They rely on specific techniques to achieve those objectives.
- Security teams need controls aligned to business risk, not generic threat models.

So, instead of answering "are we secure," CISOs are better served by answering a different question: Are we aligned to our adversaries?

Defining Adversary Alignment

Security operations teams are already familiar with adversary behavior. Frameworks such as MITRE ATT&CK® provide a structured way to understand attacker tactics and techniques across the attack lifecycle. That knowledge has long informed threat intelligence, detection engineering, and adversary emulation.

As attackers increasingly automate and adapt their techniques using AI, focusing on observable behavior rather than static indicators becomes

> "What often gets asked is, 'are we secure?' I've been asked that in my career at Fortune 15 companies that I was responsible for protecting. And it's typically after some large problem is discovered, or a large investment is made, or both — 'are we secure?' **And I think that before we can answer that, we must know our alignment to the adversary. In order to do that, you need to consider a pretty wide scope of what or who the adversary can be."**
>
> – Stephen Moore, Chief Security Strategist, Exabeam

Adversary alignment builds on this foundation. It reflects a security operations team's ability to prepare for, detect, and respond to adversary behavior wherever it appears, across identities, endpoints, applications, and automated agents. It also requires understanding how legitimate users and systems behave so deviations stand out quickly and credibly.

While related frameworks such as MITRE D3FEND™ help map defensive countermeasures, adversary alignment focuses first on understanding attacker intent and behavior, then measuring how effectively controls disrupt those behaviors in practice.

Importantly, adversary alignment extends beyond known attack paths. It also accounts for the human, technical, and organizational conditions that allow threats to succeed in the first place.

Who Is the Adversary?

Adversaries are often categorized as external attackers, malicious insiders, or nation-state actors. That view is incomplete.

From an adversary alignment perspective, an adversary can be any force that introduces risk, whether external or internal, intentional or accidental. These adversaries generally fall into three categories.

External Adversaries

External adversaries include criminal groups and other attackers operating outside the organization. They rely on tactics and techniques such as phishing, credential theft, exploitation of public-facing applications, session hijacking, and data exfiltration.

Their methods are well documented, but their execution changes constantly. Detecting them requires continuous visibility into identity, endpoint, network, and cloud activity, as well as behavioral analytics that surface subtle deviations early.

Internal Adversaries

When thinking about internal threats, many organizations focus only on malicious insiders. Those risks are real, but they represent only part of the picture.

Internal adversaries also include legitimate users whose actions unintentionally weaken security, as well as users or service accounts that have been compromised without their knowledge. In modern environments, this category increasingly includes non-human identities and agents that operate at machine speed and scale.

Behavioral analytics play a central role here. By establishing baselines for normal user and agent behavior, security teams can identify anomalies that indicate misuse, compromise, or risky activity before

Understanding Behavioral Analytics

Behavioral Analytics

Behavioral analytics establishes a baseline of what "normal" looks like across users, systems, and automated agents. By analyzing patterns over time, it helps security operations teams identify meaningful deviations that may signal risk, misuse, or compromise, even when no known threat or rule is triggered.

User and Entity Behavior Analytics (UEBA)

UEBA applies behavioral analytics to people and assets such as users, endpoints, servers, and applications. It highlights unusual activity, such as abnormal access patterns or unexpected data movement, and helps teams prioritize investigations based on behavioral risk rather than isolated events.

Agent Behavior Analytics (ABA)

ABA extends behavioral analytics to non-human identities and automated agents, including service accounts, scripts, and AI-driven processes. By monitoring how these agents typically behave, ABA helps surface anomalies that indicate misconfiguration, abuse, or compromise in environments where automation

Endemic Adversaries

Endemic adversaries are embedded in organizational decisions, processes, and priorities. They are not attackers in the traditional sense, but they can be just as damaging.

Examples include:

- Persistent underinvestment in security capabilities.
- Accumulated technical debt and unsupported legacy systems.
- Delayed modernization of identity, logging, or detection tooling.
- Poor visibility into third-party access and integrations.
- Incomplete security integration after mergers and acquisitions.
- Organizational friction that prevents timely decision-making.

These conditions create gaps adversaries exploit. Treating them as adversaries forces the organization to confront systemic risk instead of normalizing it.

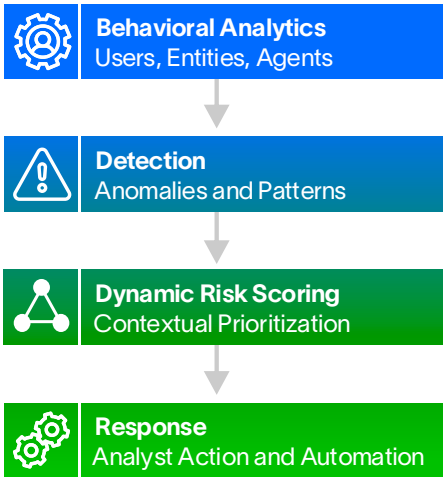

Figure 1.

Behavioral analytics establishes normal behavior for users, entities, and agents. Deviations feed detection logic, inform dynamic risk scoring, and drive prioritized response.

Designing Adversary-Aligned Security Operations

Adversary alignment starts with understanding who the adversaries are, where they operate, and how they behave. Those insights must be shared beyond the security operations team, especially when endemic risks are involved.

For CISOs, this often means documenting control gaps, surfacing risk trends, and engaging executive peers in direct, evidence-based discussions about tradeoffs and exposure.

At the operational level, adversary alignment depends on three foundational capabilities.

1. Detection

Detection answers a basic question: **How early and how reliably can adversary behavior be identified?**

Aligning detections to frameworks such as ATT&CK helps ensure coverage across the attack lifecycle. Behavioral analytics strengthen this coverage by correlating activity across users, entities, and agents to identify anomalies that static rules miss.

Capabilities such as user and entity behavior analytics (UEBA) and Agent Behavior Analytics (ABA) are essential for detecting both malicious activity and risky behavior tied to misconfigurations, compromised identities, AI agents, or automated processes.

2. Speed

Once adversary behavior is identified, response speed becomes critical. Delays increase the likelihood of data loss, operational disruption, and reputational damage.

Security teams need the ability to validate risk quickly, prioritize actions, and contain threats consistently. Exercises such as adversary emulation and threat simulations help teams measure and improve detection-to-response timelines while revealing coverage gaps.

3. Consistency

Detection accuracy and response speed only matter if they are repeatable.

Consistency comes from programmatic processes, automation, and analytics that apply the same logic every time, regardless of analyst workload or experience level. Automation does not just accelerate response; it ensures reliability across investigations and outcomes.

In practice, adversary alignment depends on:

- Clear definitions of adversary behaviors.
- Visibility into how often those behaviors occur.
- Analytical logic that continuously refines prioritization and response.

Three Lenses for Measuring Success

CISOs need clear ways to explain how adversary alignment improves outcomes. Three complementary lenses help frame that impact.

The Risk Lens

The risk lens focuses on identifying and prioritizing deviations from normal behavior. Behavioral analytics assign relative risk based on context and patterns, allowing teams to act before issues escalate.

The Event Lens

The event lens evaluates alert quality. It asks whether detections produce actionable signals and whether analysts can triage events efficiently without noise or duplication.

The Hunt Lens

The hunt lens reflects an active approach to threat discovery. By analyzing attacker behavior and learning from incidents inside and outside the organization, teams uncover blind spots and improve coverage over time.

Together, these lenses create a feedback loop. Risk informs events. Events guide hunts. Hunts refine detection and prioritization.

> **"If you can understand your visibility or your capability gaps to be able to detect a threat across the entire cyberattack life cycle, that's how you align adversary behavior."**
>
> – Stephen Moore, Chief Security Strategist, Exabeam

Adversary Alignment in Summary

Adversary alignment is not limited to the security operations team. It reshapes how organizations understand risk, accountability, and resilience.

By focusing on adversary behavior, including human and agent-driven activity, CISOs gain a more accurate view of exposure. They can demonstrate where controls are effective, where gaps remain, and what outcomes remediation will deliver.

When asked "are we secure," an adversary-aligned CISO can point to evidence: detected behaviors, reduced dwell time, prioritized risks, and measurable improvements. That answer may not always be comfortable, but it is credible,

About Exabeam

Exabeam is the leader in behavior intelligence for the agentic enterprise. As organizations deploy digital workers and confront machine-speed adversaries, Exabeam delivers flexible, industry-proven solutions for insider threat coverage of humans and agents and faster, more accurate threat detection, investigation, and response (TDIR). Learn more at www.exabeam.com.

exabeam™

Learn more at
www.exabeam.com →