opentext™

GUIDE

Advanced insider threat detection with user behavior analytics

A guide to using behavioral analytics for smarter, faster threat detection

Contents

Introduction

Insider threats represent one of the most complex challenges facing modern organizations. While external attacks often dominate headlines, insider threats pose a more insidious and challenging risk. These threats arise within your organization—from malicious employees, negligent actions, or compromised accounts. Insider incidents often go undetected for extended periods (almost three months) and at an average annual cost of $16.8 M, they are one of the costliest types of breaches to address.[1]

Traditional detection tools rely heavily on static rules or known attack patterns, which are designed to catch predictable external threats. But insider threats are different. They don't follow predictable paths, and their activity often blends with expected user behavior. The result? Security teams are overwhelmed by false positives, wasting time investigating harmless anomalies while real threats slip through unnoticed.

This is where behavioral analytics transforms the game. Instead of relying solely on predefined rules, it establishes a dynamic understanding of what "normal" behavior looks like across your organization. By continuously learning and adapting, it identifies deviations that could indicate insider threats, such as unusual access to sensitive files, unauthorized changes to critical systems, or suspicious activity patterns. This advanced approach provides the clarity and context your security team needs to act quickly and effectively.

In this guide, we'll break down the fundamentals of insider threats and show you how leveraging behavioral analytics can provide unmatched visibility and control. Whether you're looking to enhance your existing security operations or start building an insider threat program, this guide will equip you with the knowledge and strategies to protect your organization from within.

1 Ponemon Institute, *2023 Cost of Insider Risks: Global Report*

1 Understanding insider threats

Insider threats are security risks that originate within an organization—by employees, contractors, partners, or anyone with access to systems, data, or facilities. Unlike external attackers, insiders already have a level of trust and access, making their actions more challenging to detect and mitigate. These threats often fly under the radar, blending seamlessly into day-to-day operations until it's too late.

Insider threats can generally be categorized into three main types:

1. Malicious insider threats

These involve individuals who intentionally abuse their access to harm the organization. Their motives can vary—from financial gain to revenge or even collusion with external threat actors.

Example: A disgruntled employee deliberately exfiltrates sensitive customer data to sell on the dark web or to damage the company's reputation.

2. Negligent insider threats

More common than malicious insiders, negligent threats occur when insiders unintentionally compromise security due to careless or risky behavior. These incidents are not malicious but can be just as damaging.

Example: An employee uses weak passwords, clicks on phishing links, or shares sensitive files through unsecured personal devices.

A recent Gartner survey emphasized the risk that organizations face from negligence. It found that 93 percent of respondents admitted to knowingly engaging in actions that increased cybersecurity risks, often prioritizing convenience over security. Another 69 percent of employees admitted to intentionally bypassing their enterprise's cybersecurity guidance.[2]

3. Compromised insider threats

Compromised insiders are individuals whose accounts or systems are taken over by external attackers. These threats are especially dangerous because they give attackers legitimate access while masking their activities.

Example: A phishing attack successfully compromises an employee's credentials, allowing the attacker to escalate privileges and move laterally through the network.

While these threats originate externally, once inside, these threat actors' actions look legitimate and can be indistinguishable from other employee traffic, making them a type of insider threat despite being external to the company.

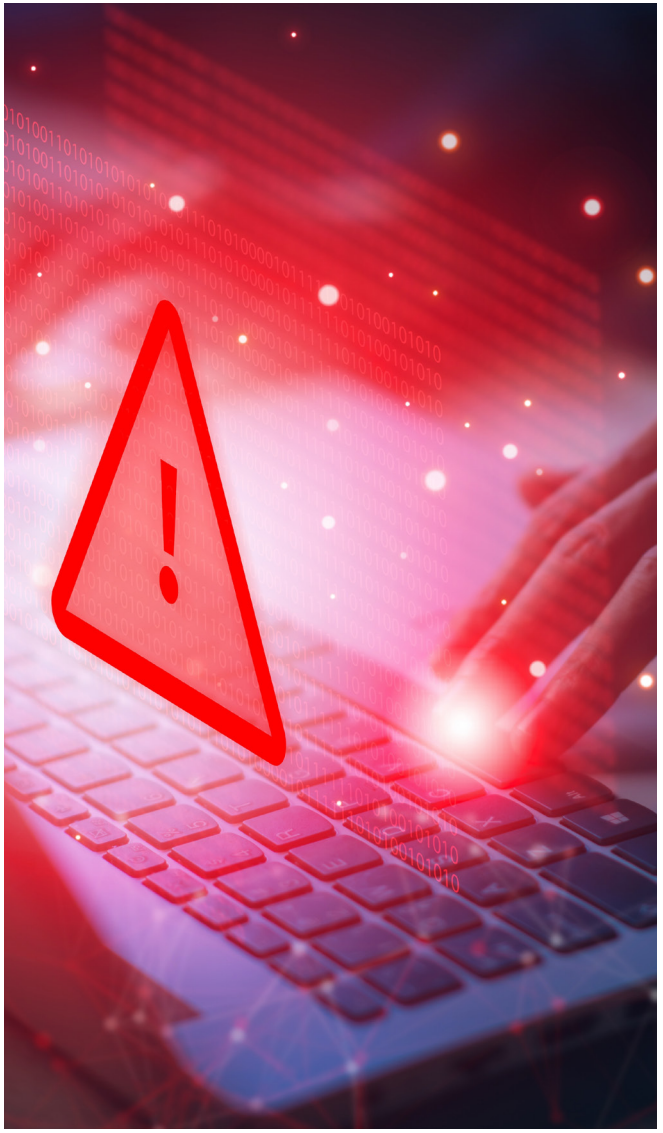

2 Gartner, *Leadership Vision for 2024*

2 Why traditional methods fall short

While effective against many external threats, traditional cybersecurity methods struggle to address the unique challenges posed by insider threats. Static rules, signature-based detection, and manual investigation have inherent limitations that make it difficult to detect the nuanced behaviors of insider activity, especially when the intent isn't overtly malicious, or the attacker takes a low and slow approach.

Limitations of static rules

Static rule-based systems rely on predefined criteria to identify threats, such as unusual login attempts or access to restricted files. While these rules can detect straightforward anomalies, they often miss the subtle, context-driven behaviors associated with insider threats.

Example: An employee working late and accessing sensitive files may trigger alerts, even if their actions are legitimate. Conversely, a malicious insider carefully operating within expected parameters may go undetected.

This rigidity results in an overwhelming number of false positives, burdening security operations center (SOC) teams with noise and diluting their ability to focus on real threats.

Challenges with signature-based detection

Signature-based tools are designed to recognize known patterns of malicious activity, such as malware signatures or IP addresses associated with known attackers. While effective against external attacks, they are virtually useless against insider threats that don't involve malware or known attack vectors.

Example: An employee emailing sensitive data to their personal account or a compromised insider moving laterally within the network typically won't match any existing signature.

The burden of manual investigation

Many organizations rely on manual investigation to compensate for the gaps left by automated tools. This approach is time-consuming and labor-intensive, requiring analysts to sift through massive amounts of data to uncover suspicious activity.

SOC fatigue: The high false-positive rate generated by static rules and signature-based detection overwhelms SOC teams, leading to alert fatigue. Over time, this fatigue diminishes their ability to respond effectively, increasing the risk of missing actual threats.

Difficulty detecting non-malicious insider threats

Traditional methods are particularly ineffective when it comes to identifying negligent or compromised insiders:

Negligence: Actions like using weak passwords, clicking on phishing links, or sending sensitive files over unsecured channels don't always trigger alerts because they aren't inherently malicious. Yet, these behaviors can lead to severe security incidents.

Compromised accounts: Attackers using legitimate employee credentials can bypass many traditional detection mechanisms. Their actions appear normal because they exploit existing privileges and behaviors.

Organizations need solutions that adapt to ever-changing environments, reduce false positives, and empower SOC teams to detect and respond effectively to insider threats—whether malicious, negligent, or compromised. Advanced tools like behavioral analytics along with AI-driven insights can bridge these gaps, helping organizations cut through the noise and focus on genuine risks.

In a landscape where 74 percent of security breaches include a human element (Verizon DBIR),[3] relying on outdated methods leaves organizations vulnerable. Insider threats demand an approach that goes beyond static rules and manual intervention to focus on dynamic, behavior-based insights.

3 Verizon, *DBIR Report 2023 - Summary of Findings | Verizon Business*, 2023

3 Behavioral analytics: The key to insider threat detection

Behavioral analytics is a security methodology that focuses on understanding how users, devices, and systems typically operate and then identifying when something significantly differs from the norm. Unlike traditional approaches that rely on known signatures, static rules, or predefined scenarios, behavioral analytics operates on dynamic patterns. It learns continuously, evolving alongside your organization's workflows, employee roles, and technology stack.

How it works

Behavioral analytics involves creating detailed profiles—or "baselines"—of what normal activity looks like for each entity in your environment. For a user, a baseline might capture factors such as:

- Usual login times and locations
- Typical file access patterns (e.g., which documents they read, create, or modify)
- Preferred devices and network segments used
- Frequency of administrative actions taken, if any

Behavioral analytics builds a statistical model of each entity's expected patterns by aggregating and analyzing large volumes of event data over time. Whenever current activity deviates from these established patterns, the system flags it for further scrutiny. For example, if an employee who normally logs in from the company's headquarters between 8 am and 6 pm suddenly connects from a distant location at 3 am and attempts to access sensitive design documents they've never touched before, that's a meaningful deviation.

A well-designed behavioral analytics solution doesn't just send an alert for any deviation or anomaly. Using AI and advanced machine learning, it understands the risk of individual and groups of anomalies and applies a relevant risk score based on the big picture. This is a critical step in providing SOC teams with high quality leads based on behavioral indicators of compromise (bIOCs).

User behavior baselines and deviations

Establishing baselines is the foundation of behavioral analytics. It's not just about logging actions, but understanding their frequency, context, and relevance. By doing this at scale—across all users, systems, and devices—behavioral analytics identifies anomalies that may appear benign in isolation but are suspicious in aggregate. This context-driven approach is especially effective when distinguishing between harmless variations (like a user working late occasionally) and genuine risk indicators (like a user suddenly downloading proprietary data at odd hours for multiple consecutive days).

The system doesn't just send a generic alert when a deviation is detected. It provides context around what changed and why it's unusual, guiding analysts to the root of the issue faster. These insights reduce time wasted on false positives and arm security teams with a richer understanding of potential risks.

What is a false positive?

A false positive is when a security system wrongly flags legitimate activity as a threat. These errors overwhelm SOC teams and divert focus from real threats. Behavioral analytics reduces false positives by accurately identifying truly anomalous behavior.

What is a false negative?

A false negative occurs when a security system fails to detect a real threat, allowing malicious activity to go unnoticed. This is especially dangerous for insider threats, where subtle, atypical behaviors might evade traditional detection. Behavioral analytics helps minimize false negatives by identifying nuanced patterns that signal potential risks.

Behavioral indicators of compromise (bIOCs)

Cybersecurity has long focused on traditional indicators of compromise (IOCs)—such as known malware signatures or suspicious IP addresses—which work well for detecting external attacks that follow a pattern but struggle with insider threats. That's because insider threats often involve activity that appears legitimate on the surface—such as a user accessing a corporate server or downloading a file. They don't rely on known malicious code or easily identifiable external signals.

Behavioral indicators of compromise (bIOCs) address this gap by focusing on the nature of the behavior itself. Instead of looking for particular malware, bIOCs look for actions that stand out against a user's or group's established norms. For instance:

- An engineer who typically only accesses code repositories is now exploring confidential financial databases.
- An HR staff member suddenly exporting large volumes of employee data to an unfamiliar device.
- An administrative account logging in from geographically distant locations within a short time frame.

Because these anomalies are identified by comparing activity against tailored baselines, bIOCs adapt to unique organizational contexts. They are not static lists of "bad" behavior. Instead, bIOCs are dynamic, growing more accurate and insightful as they process more data and refine their understanding of what "normal" looks like in your environment.

A principled approach to behavioral risk

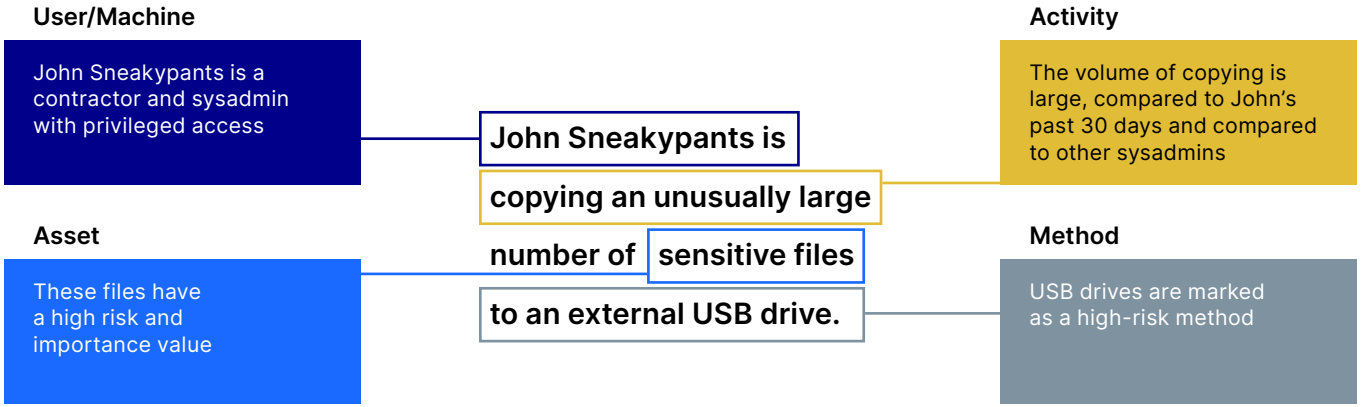

A robust behavioral analytics solution learns directly from your environment, building context from the ground up. Unlike rule-based tools that can be pre-configured with universal data sets, behavioral analytics can't simply plug in someone else's baselines and deliver immediate results. It needs time—often weeks—to observe user activities, gather insights, and mature into a finely tuned system capable of accurately detecting the subtle anomalies unique to your organization.

Why behavioral analytics improves detection accuracy

By centering on user behavior rather than a fixed set of known threats, behavioral analytics dramatically increases the accuracy and relevance of security alerts. It reduces false positives by understanding context (such as time, location, and typical job responsibilities), ensuring that legitimate shifts in work patterns aren't unnecessarily flagged. Simultaneously, it minimizes false negatives by catching threats that don't fit clean signatures or predefined rules—like a trusted insider suddenly abusing their access.

An example alert from a behavioral analytics solution. Where a rule-based threat detection solution would require a large number of predefined cases, a behavioral analytics solution understands the context around anomalies, dynamically creating alerts unbound by rules.

Behavioral analytics equips your security team with a powerful lens to see beyond superficial indicators. It captures what is normal and what is not, enabling organizations to detect subtle, advanced threats before they escalate into full-blown incidents. In an era where insiders can unknowingly or deliberately cause substantial damage, behavioral analytics provides the adaptability, context, and depth of insight necessary to stay ahead of evolving risks.

4 Core use cases for behavioral analytics

1. Detecting data exfiltration

Data exfiltration—removing sensitive data from a secure environment—is a top concern for security leaders and executive teams. Behavioral analytics helps spot subtle changes in a user's activities that might indicate data theft. For example, if an employee who generally accesses small numbers of files suddenly begins downloading large volumes of proprietary documents, behavioral analytics flags the anomaly. Unlike static rules that might only trigger on known "bad" file signatures, this approach focuses on the context and volume of the activity, allowing security teams to quickly intervene before major damage occurs.

2. Identifying credential abuse or misuse

Credential abuse can occur when insiders deliberately misuse their privileges or when external attackers compromise legitimate accounts. Traditional tools may struggle to detect this because, on the surface, the user is performing a "normal" login. Behavioral analytics solves this by comparing current actions to established baselines. A compromised account might access resources at unusual hours, from unexpected locations, or seek information not typically associated with that user's role. By highlighting these deviations, behavioral analytics helps security teams identify compromised credentials early, reducing the window of opportunity for attackers.

3. Monitoring unusual access patterns

Even if an employee is trustworthy and well-intentioned, unusual access patterns can signal risk—through negligence or malicious intent. Consider a user who typically retrieves HR documents during business hours, now accessing them late at night or from a remote location. While any single instance might not be malicious, behavioral analytics detects the pattern and provides context. This contextual awareness allows analysts to quickly determine if the activity warrants deeper investigation, helping prevent threats from slipping through due to a lack of obvious "red flags."

4. Responding to high-risk behaviors, such as privilege escalation

Privilege escalation—gaining unauthorized levels of access within a system—is a key step in many advanced attacks. Traditional controls might only notice that privileges have changed, but not why or whether it's truly suspicious. Behavioral analytics understands the user's typical access patterns and responsibilities. If an account with previously limited access suddenly attempts to modify security settings or view sensitive

financial data, the system recognizes the unusual behavior. Instead of requiring an analyst to manually correlate events, behavioral analytics presents a clear, context-rich signal that a critical incident could be unfolding.

Putting it all together

By focusing on how users and entities behave, rather than just static rules or known signatures, behavioral analytics provides a more nuanced and accurate form of detection. Whether it's preventing data theft, stopping attackers from exploiting legitimate accounts, or catching unusual access attempts, this adaptable, context-driven approach helps security teams stay one step ahead. The result is a more resilient security posture that not only catches sophisticated threats early but also reduces the noise that often hinders timely response.

5 How to implement advanced behavioral analytics

Implementing advanced behavioral analytics in your security operations center (SOC) can significantly enhance your team's ability to detect and respond to insider threats and other sophisticated risks. Success begins with a well-structured approach that aligns with your current security ecosystem and workflows.

Steps for integration

- **Assess your current environment:** Start by reviewing your existing security tools and data sources—SIEMs, EDR platforms, IAM solutions, and more. Identify coverage gaps and determine where behavioral analytics can provide the greatest boost in visibility and detection. This might mean focusing on identity management first or zeroing in on application servers where sensitive data resides.
- **Select a behavioral analytics solution:** Choose a platform that excels in detecting insider threats, credential abuse, and unusual access patterns. Ensure it supports seamless integration with the tools you rely on, such as Microsoft Defender for Endpoint or Microsoft Entra ID. Look for adaptive capabilities that learn from your unique environment rather than requiring rigid rule updates.
- **Establish data feeds and baselines:** Connect your chosen solution to relevant data sources. Over time, it will establish baselines that define "normal" behavior for each user and system—what files they typically access, when and where they log in, and how they move within the network. Allowing sufficient time to gather these behavioral insights is critical; the quality of your baselines depends on the quantity and diversity of data the system ingests.
- **Integrate with existing workflows:** The real power of behavioral analytics emerges when it's embedded directly into your current operations. Integrate alerts and insights into your SIEM, ticketing systems, or incident response tools so analysts don't need to switch between multiple platforms. This ensures that every suspicious anomaly is contextualized within your existing workflow, making it easier to track incidents, assign responsibilities, and drive faster resolutions.
- **Fine-tune and calibrate:** After a baseline period, review the system's alerts and recommendations. A good system should automatically adjust without human intervention, but if it doesn't, ensure you adjust thresholds as needed and refine what your organization considers "high-risk" behaviors. Because behavioral analytics continuously learns and adapts, ongoing calibration should happen automatically, enhancing precision and relevance over time.
- **Train your team:** Provide dedicated training sessions to help analysts understand how behavioral analytics differs from traditional rule-based detection. Demonstrate how to interpret anomalies, differentiate between harmless irregularities and genuine threats, and leverage the platform's context-rich insights., Your SOC can rapidly move from detection to informed action by building confidence and competency.

"[OpenText Core Threat Detection and Response] has enabled behavioral change in our company. As a CISO I have never had business units reach out to me BEFORE they make changes to the IT or cyber environment."

- CISO, Global internet retailer

6 How OpenText can help

If you are looking for advanced threat detection of insider threats, novel attacks, and advanced persistent threats, look no further than OpenText™ Core Threat Detection and Response. OpenText Core Threat Detection and Response provides a modern, behavioral-driven approach to insider threat management and advanced detection. By focusing on user and entity behavior rather than relying on static rules and known signatures, it empowers security teams to uncover subtle patterns of malicious, negligent, or compromised behavior across their environment—before these actions escalate into damaging incidents.

Threat detection and response extends beyond security operations (SecOps), encompassing DLP, IAM, EDR, NDR, and more.

Your security resilience: now powered by OpenText

Open XDR architecture with OpenText Threat Detection and Response

Data Loss Prevention (DLP)

Typical DLP includes:
- Asset discovery
- Data classification
- Data Encryption
- User management
- Policy management
- Permission management
- Compliance reporting
- Automated user education
- Auto response/playbooks

Now powered by OpenText:
- Find data threats with data permissions management and rules-based logic
- Utilize AI based detection, prevent data exfil, and dynamically evolving ML models
- Cover all advanced threats
- Ingest data access logs with IAM

Endpoint/Network Detection & Response (EDR/NDR)

Typical EDR / NDR includes:
- Endpoint/network level telemetry
- Rich data source
- Rule based response
- Rapid response on machine
- Robust response capabilities

Now powered by OpenText:
- Focus beyond just data threats within the company's perceived perimeters and boundaries
- Utilize behavior analytics to find anomalies of outsiders and insiders that ignore boundaries
- Gain more visibility, see all advanced threats
- Ingest Endpoint/Network data

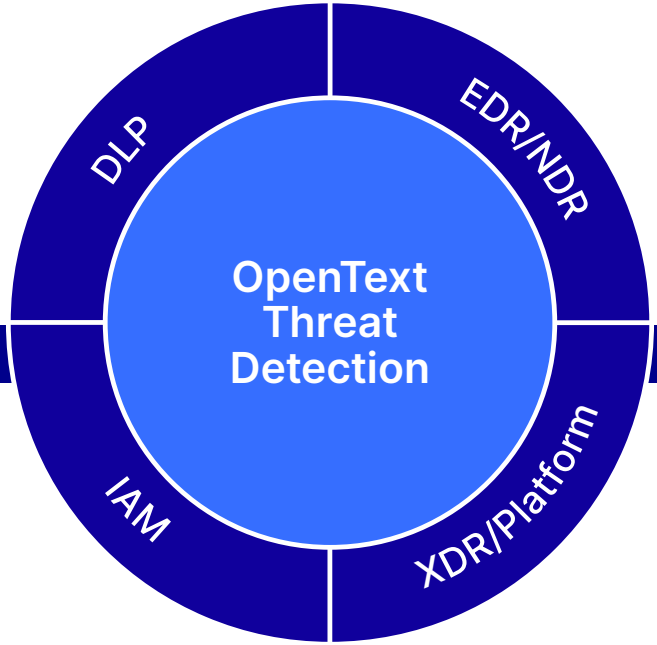

Identity Access Management (IAM)

Typical IAM includes:
- Log access data
- Identity log data
- Authentication
- Compliance reporting
- Zero Trust
- Role management
- Group management

Now powered by OpenText:
- Manages identty access and authentication with threat hunting
- Address insider threat detection linked back to where entry points/ vulnerabilities may be
- Low and slow and APT Detection capabilities added
- Ingest IAM is valuable data source to augment detection

XDR/Platform

Current XDR includes:
- Evolution of EDR/NDR solutions for single pane of glass dashboard
- Integration with most data sources
- Robust response capabilities
- Detection for multiple threat types

Now powered by OpenText:
- Combine behavioral analytics for detection for all types of advance threats with DLP type controls and response: pulls with many data sources including IAM
- Benefit from complex set-up and maintenance: time consuming configurations
- Work with diverse set of response options

Content Management | Service Management | IT Operations Management | Supply Chain Management

Adaptive, behavior-driven analytics

Unlike traditional solutions that require frequent manual rule updates, OpenText Core Threat Detection and Response automatically adjusts its detection baselines, learning from your organization's evolving workflows. As roles, technologies, and business processes change, the solution continuously refines what "normal" looks like, ensuring more accurate detection and fewer false positives.

Seamless integration with existing investments

Whether you rely on Microsoft Defender, Entra ID, or other critical security tools, OpenText Core Threat Detection and Response meshes seamlessly with your established ecosystem. This interoperability means enhanced visibility and fewer blind spots—without disrupting your current operations or forcing you to overhaul your technology stack.

High-context, actionable alerts

When the solution identifies suspicious behavior, it provides context-rich insights in clear, understandable language. SOC analysts no longer waste time deciphering cryptic alerts or investigating benign anomalies. Instead, they can quickly move from detection to decision-making, improving response times and reducing the risk of insider-driven breaches from all levels of expertise.

Reduced SOC workload and alert fatigue:

By filtering out the noise and prioritizing high-risk incidents, OpenText Core Threat Detection and Response empowers analysts to focus on genuine threats. This improves efficiency and enhances analyst morale and effectiveness, contributing to a more proactive and resilient security posture.

Scalable and future-proof

As your organization grows and evolves, OpenText Core Threat Detection and Response scales alongside it. Its behavioral models continue to adapt, positioning your team to stay ahead of emerging risks and evolving insider threat tactics. With the support of OpenText's expert services and ongoing innovation, you gain a solution designed to address today's challenges and tomorrow's unknowns.

Ready to explore how OpenText Core Threat Detection and Response can elevate your insider threat strategy?

Resources

Learn more: Visit our product page to dive deeper into advanced behavioral analytics.

Request a demo: Experience the power of real-time threat detection tailored to your organization's unique environment.

Connect with our experts: Schedule a consultation with our security specialists to discuss your goals and challenges.

Take the next step toward a more proactive, resilient security posture. Your insider threat defense starts here.

opentext™