Bottomline™

The Inside Scoop on Insider Threat Management

USING TECHNOLOGY TO STOP FRAUD AND THEFT FROM THE INSIDE OUT

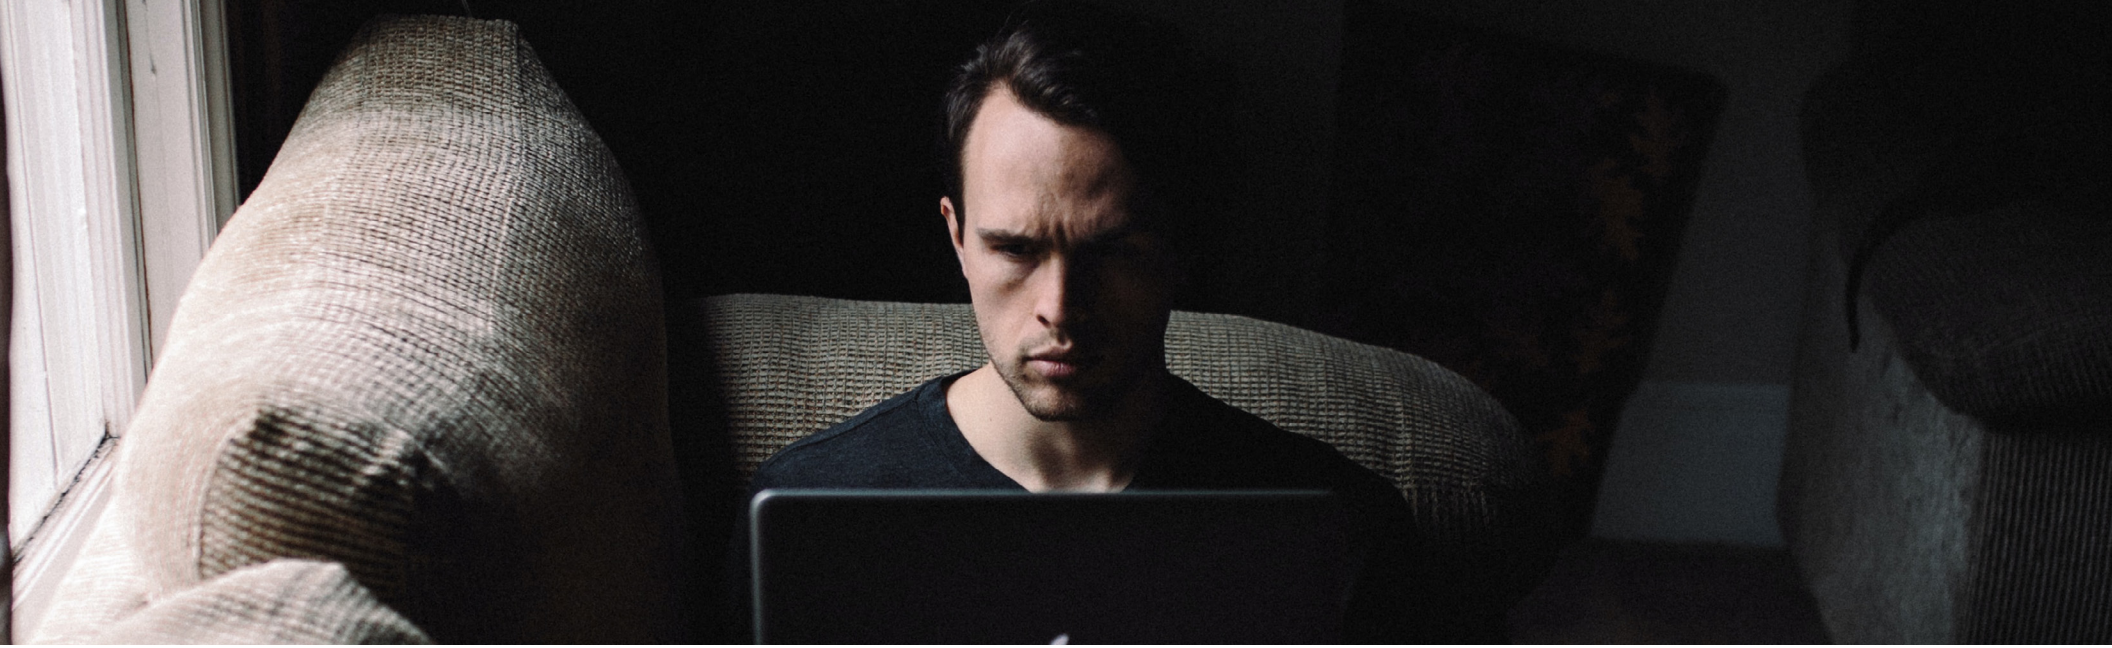

“Detect and Contain” Is the Name of the Game

Most employees are ethical. A certain percentage, however, are not. These malicious actors work from inside your organization to steal, sell, or manipulate data for their own gain. The risk they pose to your company is real and it is growing.

According to Ponemon Institute’s 2022 Cost of Insider Threats report, insider threat incidents increased by almost 50% over the prior two years and the mean containment time for an insider incident is 85 days. This is concerning since every day of delay represents additional time for bad actors to continue their unscrupulous activities.

Since it is not possible to prevent dishonest employees from making the attempt to defraud you, there is only one recourse to minimize the risk to your business: detect and contain the greatest number of insider threats in the shortest amount of time.

A wide range of technology solutions exists to aid in this endeavor, such as endpoint detection and response (EDR) systems, agent-based technologies, managed services, data loss prevention (DLP) solutions, and more. A layered approach that takes advantage of complementary capabilities typically delivers the strongest outcomes. Providing an effective solution typically requires dealing with three main factors:

1. Getting the Data

How does the technology **capture data** related to employee behavior?

2. Analyzing The Information

How does the technology **provide insights** into employee behavior?

3. Investigating the alerts

How does the technology **support investigators** in documenting and assessing employee behavior?

In the following pages, we answer these three questions by looking at the different ways technologies approach them.

Getting The Data: How Does the technology capture data related to employee behavior?

The first consideration is understanding the various methods insider threat technology solutions use to collect data on how employees are interacting with personal, confidential, or proprietary customer or company information. There are three primary ways data is captured: through log files, via agents, and from the corporate network.

LOG-BASED DATA

Log files are generated by the application or system itself. Logs typically record user actions, including login and logout times, access to files and directories, changes to files and data, and system commands executed. They also capture a wide range of system and network events, such as firewall activities, network connections, system configurations, and software installations. Log files include timestamps for each event, which are essential for establishing timelines and identifying patterns of behavior.

All this data helps understand who is doing what on the network and, to some degree, within applications. However, even the most detailed log files have their limitations. The greatest of these with regard to insider threats is that log files focus on events that change something, e.g., an update to customer details or a payment transaction. They do not usually capture activities that do not change anything, such as an employee making inquiries, browsing through customer information, or mining data: actions that may be precursors to theft or fraud.

Insider Threat Scenario: Log Limitations

Michael is a customer service representative at a well-established bank. He is planning a future theft and so begins a covert data-gathering operation. He crafts inquiries that extract small, inconspicuous portions of sensitive financial information, such as customer account balances or transaction histories. To avoid suspicion, he schedules these queries to run during non-peak hours, simulating routine account maintenance tasks.

Michael's queries are carefully structured to retrieve limited and seemingly random financial data, ensuring they do not trigger significant deviations in query logs. His access and query patterns closely align with his legitimate job responsibilities, making it difficult for the bank's security personnel to differentiate his malicious actions from his regular duties based solely on log files.

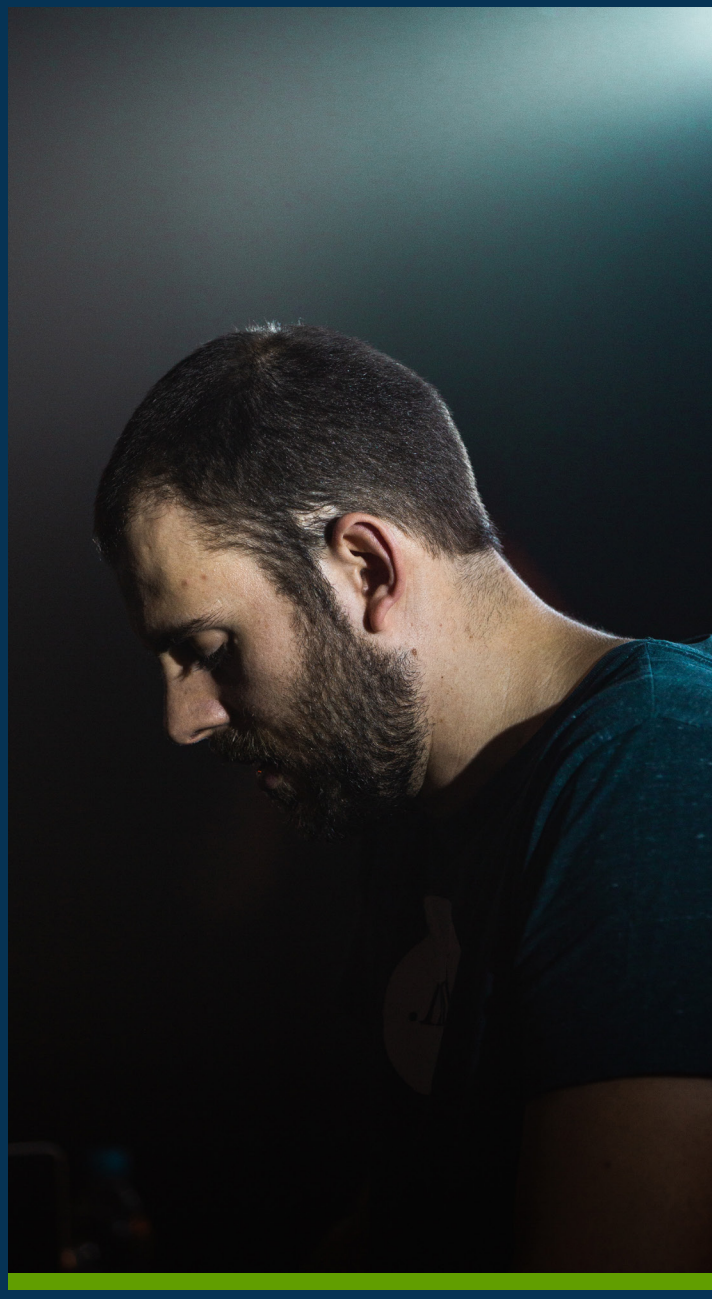

Agent-Based Data

Agent-based solutions rely on agents or software components installed on individual endpoints, servers, or devices to monitor and capture data related to user activities, system configurations, network traffic, and other relevant events. They run in the background, continuously collecting data.

More detail

Agent-based solutions capture more granular and real-time data than traditional log files. They can track individual user actions such as keystrokes, mouse movements, application usage, and file access, providing an expanded view of what is happening on a device or endpoint.

More context

Agent-based solutions often capture context around user activities, such as the sequence of actions leading up to an event. This contextual information can be crucial for understanding the intent behind user actions and for identifying insider threats.

May not recreate user actions

Some Agent-based solutions record user actions as text entries. They may not, however, provide a visual record of the screens a user accessed, nor are they able to replay the interactions that a user took within a system or application.

May not extract field values

Most agents, if not all, cannot obtain the actual field values that appear on the screens an employee looked at or interacted with. These values therefore cannot be incorporated into analysis.

May violate user privacy

Because agents sit within endpoints, servers, or devices, they can capture personal emails, private chat conversations, or even sensitive information entered by employees that is unrelated to work.

If an employee sends a personal email which includes sensitive personal information, an agent that monitors email activity and key strokes would **automatically capture the employee's email content.** This is a major concern given increasingly rigorous privacy regulations such as GDPR.

Network-Based Data

Network-based data capture has been used for many years for analyzing health and performance networks, as well as detecting intruders. Network sniffing can be used also for monitoring user activity within enterprise applications. In this approach, user activity on business-critical systems and applications is captured non-invasively for analysis directly from the corporate network – rather than from endpoints, servers, or devices – without deploying agents. Network-based data capture is readily implemented on mainframes, web-based applications, and SaaS applications, **protecting sensitive data regardless of where it resides.**

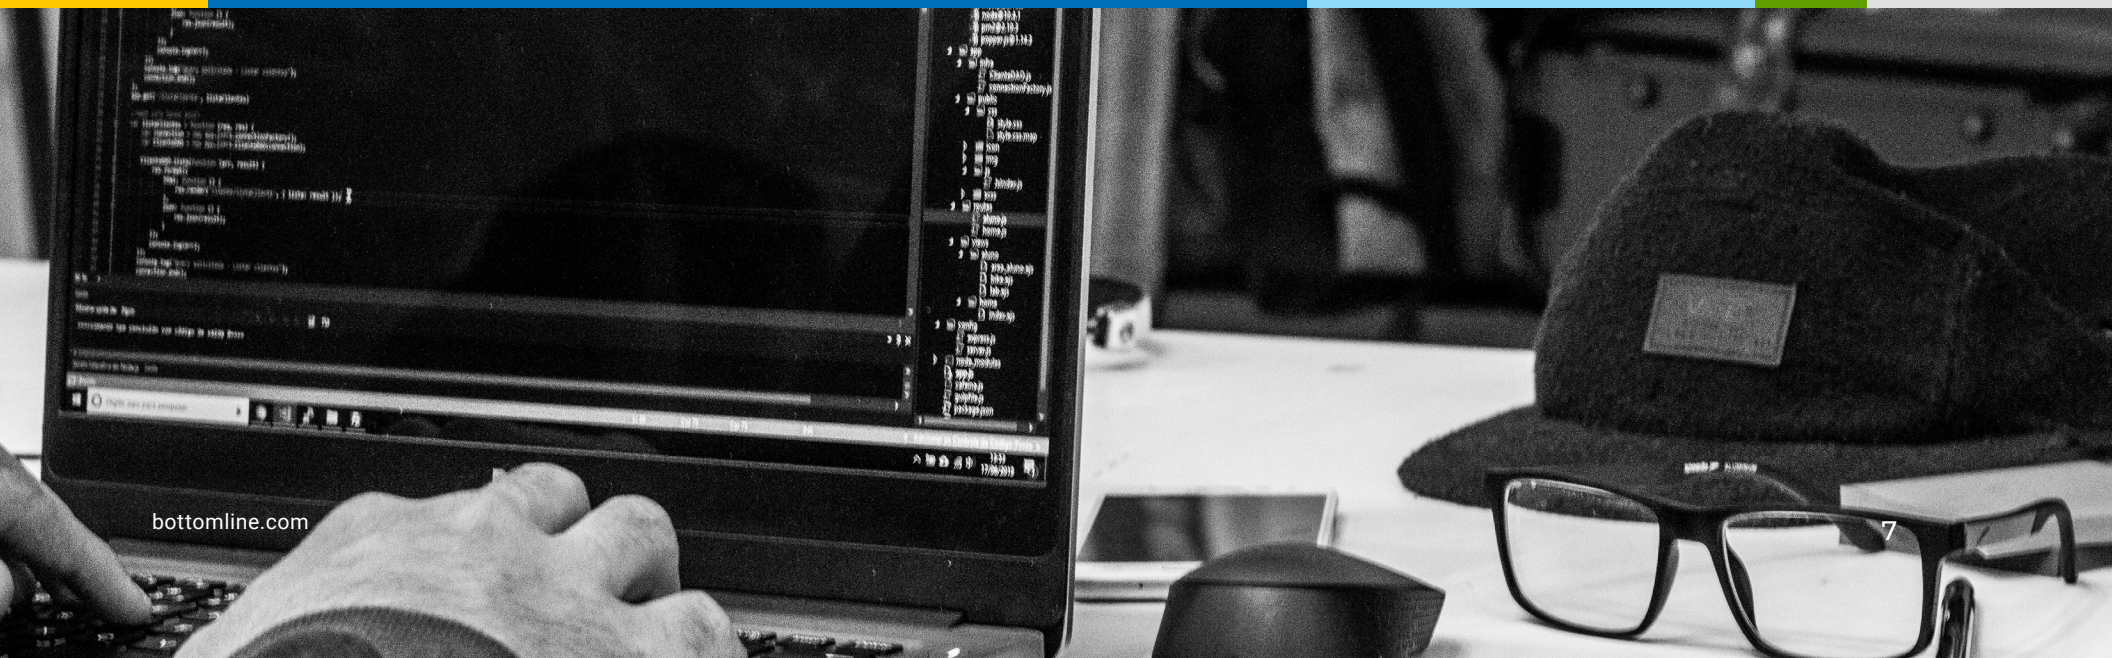

6 Reasons to Say “Yes” to a Non-Invasive Approach

1. ### Monitor access to business data, not employees or endpoints.

   With this approach, it is user access to business data that is being monitored, not the server, device, or endpoint. For example, it does not matter whether the device a user is working on is approved or unapproved, personal or corporate: the business data is constantly monitored at the system or application level.

2. ### Select specific systems and applications for protection.

   Rather than automatically recording all activity on a server, device, or endpoint, network-based data capture can be applied precisely where needed: not only to certain applications, but even to specific screens or particular fields.

3. ### Capture both event and non-event actions.

   Every user action within a monitored business system or application is captured, including activities that do not make any changes in the application database such as searches, queries, and browsing.

4. ### Record and replay activity screen by screen.

   Log files and agent-based solutions typically record information in the form of text entries. Network-based data capture supplies richer detail and context by also creating a visual screen-by- screen record that can even be “replayed” like a slide show.

5. ### Extract values from fields for richer analytics.

   In addition to providing a record and replay capability, capturing data from the recorded screens allows for search on the captured screens and analytics on the user behavior.

6. ### Avoid employee privacy issues.

   Because network monitoring focuses purely on user activity that relates to sensitive company data, it is easy to comply with privacy regulations. For instance, employee emails, instant messages, general website activity, and the like all remain completely private because they do not take place within the monitored systems and applications.

Analyzing The Information: How Does The Technology Provide Insights Into Employee Behavior?

Technology solutions differ widely in the analytical capabilities they offer, but they are all the same in one regard: analytics are always limited by the extent of the available data. The less data there is to work with, the more restricted the insights from analytics. Conversely, the more data there is on hand, the richer the insights that analytics can potentially generate. With that in mind, technology solutions typically exhibit **three categories of analytics** depending on the data they draw upon.

Log-Based Analytics

Analytics based on log files can detect:

Anomalous Event-related User Behavior

Excessive updates to customer details, excessive money transfers, unusual login times

Access to Sensitive Resources

Access to files or databases with customer, confidential, or proprietary data

Account Misuse

Users attempting to escalate privileges or access resources not required for their job

Suspicious Patterns

Patterns of behavior that are consistent with known insider threat indicators

Analytics based on log files are limited with regard to:

Anomalous Inquiry-related User Behavior

User behavior which relates to searches, queries, and surfing t hrough customer data cannot be performed based on log files

Complex Insider Threats

It difficult to correlate data and detect threats that span multiple segments in a highly segmented network environment

False Positives

Analytics based on log files often generate large volumes of false alerts which can be time-consuming and resource-intensive to review and assess

Lack of Context

Log files do not provide rich context for interpreting the intent behind a user action

Incomplete Visibility

Bad actors can take measures to cover their tracks, leaving blind spots in the log files

Response Time

Real-time detection may not be possible, allowing malicious actors to continue their activities and cause significant damage

Agent-Based Analytics

Analytics based on data captured by agents include the capabilities of log-based analytics and also provide:

Real-time monitoring for immediate detection and response capabilities.

Rich context which can be helpful in understanding the intent behind user actions.

Endpoint visibility which can detect user activities that do not generate logs

File integrity monitoring to track changes to files and directories at a granular level, helping to detect data exfiltration, unauthorized access, or tampering with critical files.

Behavioral profiling of users and endpoints to learn what constitutes normal behavior and flag deviations, even if those deviations do not generate explicit log entries.

Network traffic analysis to provide insights into communication patterns and potential security risks.

However, as noted earlier, agent-based analytics cannot incorporate application field values from the screens accessed by employees since agents cannot capture this data. This leaves an important gap in detecting potential insider threats which can be committed by accessing corporate applications.

Insider Threat Scenario: Agent Angst

Sarah, a malicious insider working at a financial institution, is colluding with an external partner to defraud clients who have invested in a certain mutual fund. She performs a query to get a list of clients who have shares in this fund and takes a picture of it with her smartphone. Then, Sarah looks up these accounts and takes pictures of the clients' personal information. She sends these photos to her partner, who will conduct the actual fraud. The agent-based analytics the institution uses **do not trigger an alert** because a) Sarah makes no changes to the client accounts, and b) the agent is not able to capture the field value for the fund, therefore there is no way to detect that an employee is reviewing an abnormal number of clients who all have this fund in common.

Network-Based Analytics

Network-based analytics can perform the most critical functions covered in log analytics and then go further by providing capabilities such as:

Field Data Analysis

Because a network-based solution can extract the values of fields on the screens an employee accesses, this data can be used in analytics. This includes screens that an employee only views, including search results or queries. This information can be used to create a highly-detailed profile of user behavior and pinpoint where users are viewing or interacting with business/customer data in potentially inappropriate ways.

Auto-Mapping of Screens

The ability to create a screen-by-screen recording and capture field values allows for auto-mapping of screens utilizing advanced machine learning models. That is, the most important fields on a given screen can be identified so that any activity that involves those fields gets special emphasis, e.g., a heightened sensitivity for generating alerts.

Proactive Detection

The ability to capture and analyze both inquiries and transaction initiation performed by employees within the corporate applications enables proactive detection of malicious activity by insiders. For example, a bad actor working in a bank who intends to divert funds from dormant accounts will first need to identify possible victim accounts by using inquiries on customer and account information. By capturing and profiling user access to dormant accounts, such an attempt can be discovered before funds leave the bank.

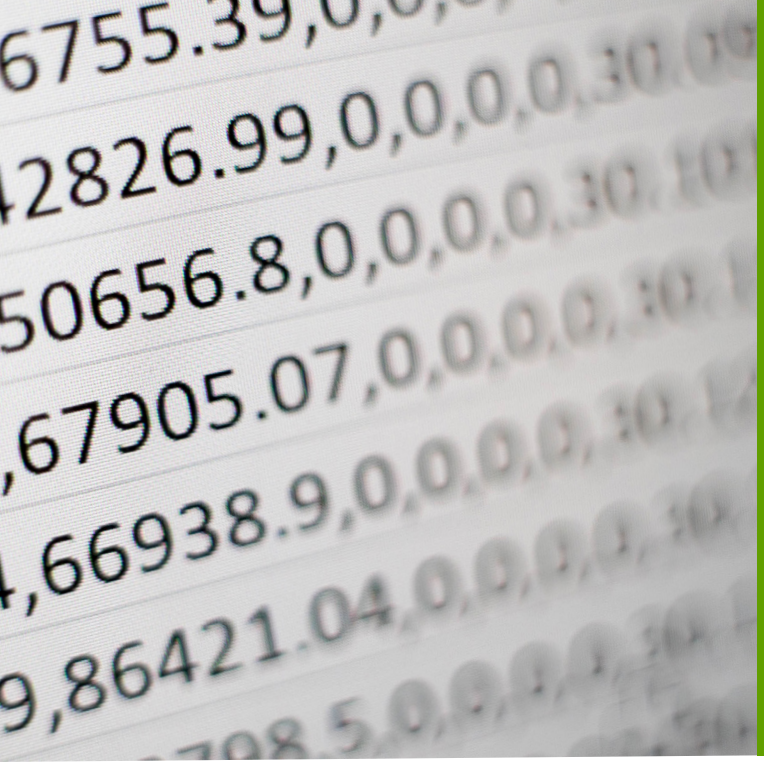

Investigating the alerts: How does the technology support investigators in documenting and assessing employee behavior?

Insider threats are challenging to detect because employees have access to company and customer data legitimately. How, then, can investigators differentiate between employees doing their normal job vs. employees engaging in malicious activity – whether that is preparing for a fraud or theft or committing an actual fraud or theft? The fact of the matter is that two employees could access the same systems, the same applications, and the same data on the same day, yet one of them could be doing their job and the other perpetrating an inside fraud scheme.

Investigators rely upon data and analytics to understand and assess user behavior. They require a high degree of certainty before confronting an employee with their suspicions or allegations, because a mistake – accusing an honest employee who is just doing their job – can cause a great deal of individual harm and even result in reputational damage. Additionally, investigators need hard evidence if the company desires to pursue criminal or civil action.

Once again, log-based, agent-based, and network- based solutions vary significantly in how they help – and sometimes hinder – insider threat investigations.

Log-Based Investigations

Log-files are a staple of insider threat investigations. While the data they capture is more limited than that collected by agents or from the network, it serves as valuable forensic evidence for investigators.

On the plus side, log files enable:

Automation

Some aspects of threat detection and investigation can be automated using log-based data and analytics, allowing investigators to work more efficiently.

Prioritization

Alerts to be prioritized based on risk levels, directing investigators to focus on the most critical threats first.

Reconstruction

Log files can assist investigators in reconstructing events and timelines leading up to an incident so that they understand how it occurred.

Log files also have their negatives:

Wasted time

Searching through and analyzing logs to expose insider threats and compile evidence on complex fraud schemes is often overwhelmingly time consuming.

False positives

Since they have limited visibility into user activity, log-based solutions can generate numerous false positives, diverting investigator attention away from actual threats and causing alert fatigue.

Limited context

While logs provide data, they often lack context Investigators may need to spend additional time gathering context from various sources to understand the full picture.

Agent-Based Investigations

Investigators using agent-based solutions have several advantages over those who have access only to log-based capabilities. These include:

Real-Time Alerts

Agents offer real-time monitoring capabilities, allowing investigators to receive alerts as soon as suspicious activities occur. This enables rapid response to insider threats.

Endpoint Data

Agent-based solutions provide detailed endpoint data, including user activities and system events. Investigators can analyze this granular data to understand the actions taken by users and devices.

Pre-configured Rules

Agent-based solutions often include pre-configured rules for detecting certain types of suspicious activity.

Control Capabilities

Some agent-based solutions provide endpoint control capabilities, allowing investigators

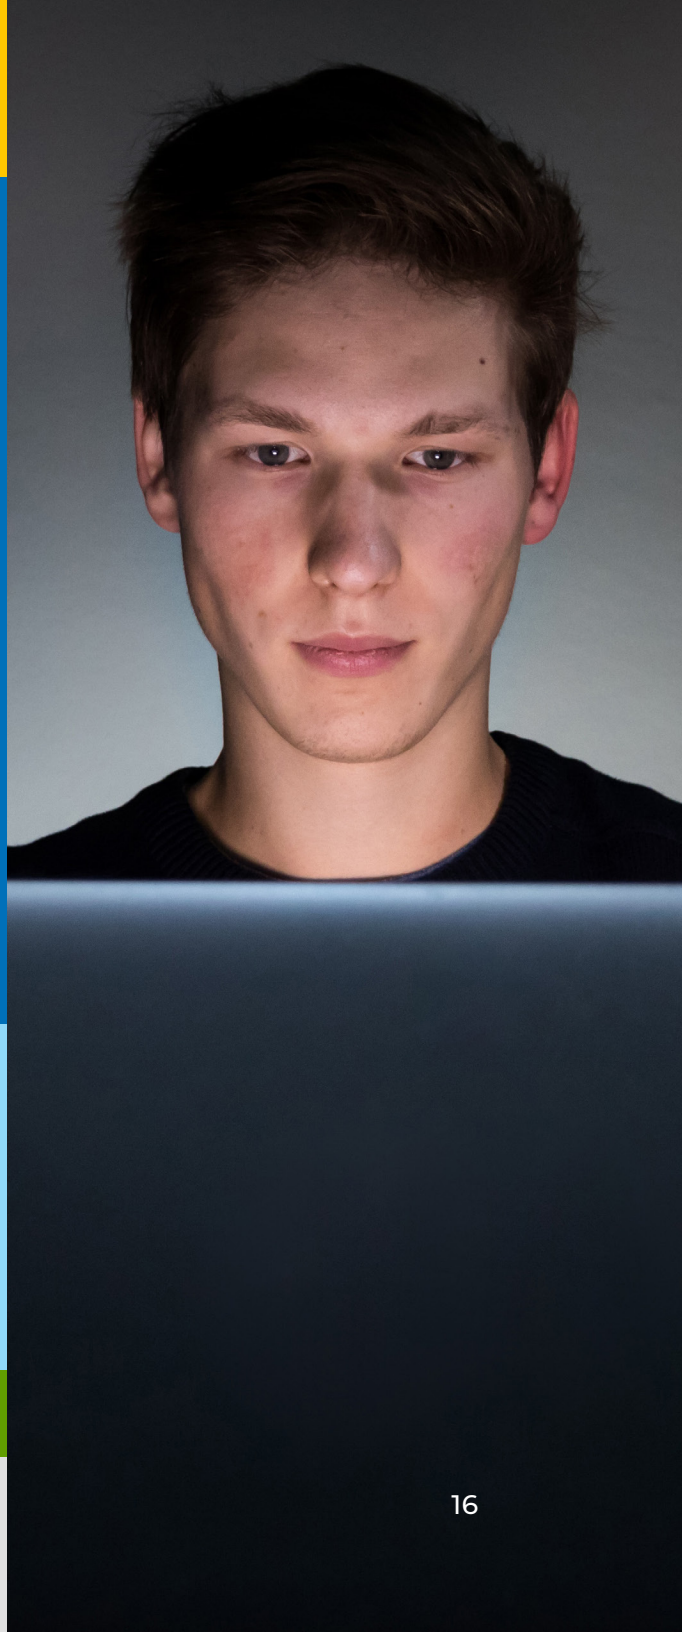

Agent-Based Investigations

Like log-based solutions, agent-based investigators must also deal with a large volume of false positives and be sensitive to privacy concerns. Additionally, maintaining and updating agents on all endpoints requires ongoing effort; outdated or unpatched agents can introduce security vulnerabilities and inhibit investigations.

Network-Based Investigations

Network-based data gathering leverages all the benefits of log solutions and then goes far beyond them to bring insider threat investigations to the next level.

First and foremost, the network-based solution has been proven to reduce investigation time by up to 90%. This dramatic decrease in time is the result of:

Easy-to-understand dashboards, reports, and charts that visualize insights, analytics, risks, and threats.

Screen-by-screen data capture coupled with record-and-replay capabilities that eliminate the need for laborious and time-consuming searches through log files to find relevant data. Instead, investigators can view activity screen by screen or watch employee action unfold as it happened in real time. and devices.

A Google-like index search that enables investigators to look up any piece of data – e.g., a phone number, social security number, last name, etc. – to see what screens the data was shown on, the context of the activity, and who engaged in the screen interaction.

Second, all of this information can be presented as **documented evidence** to make a case against an employee or to initiate legal action. For example, the record-and-replay functionality provides a visual storyboard that shows exactly what the employee was doing in an application. There is no missing data and no alternative perspective that could refute the facts because the action can be viewed as it happened and therefore speaks for itself.

Third, network-based data and analytics open up an entire new world of **proactive risk mitigation** because they can reveal and document suspicious behavior that is not event-related. For instance, if an employee is looking up information in preparation for a fraud attempt, analytics can signal an alert based on the types of accounts being accessed. An investigator can review all the screens the employee was pulling up and verify whether the activity was legitimate or not. By being able to investigate “non-event” activities, investigators

decrease in investigation time

Insider Threat Scenario: Network Know-how

Eric works at a bank and has his sights set on theft. He browses through accounts to find ones that are dormant that could be good targets to divert money from. He believes that he is safe because he is not taking any action that would register as an event in a log file or by an agent.

When Eric finds a likely account, he takes photos of the screens with his smartphone. He then goes home and, using his personal computer, creates an online user ID for the account and changes the phone number to his own. He is careful to use a phone number that is not on his employee record at the bank. Eric is confident that since this activity does not happen on a company device or endpoint, there is no way for the changes to be detected.

What Eric does not realize is that the banking application he is accessing is protected by a network- based solution and every screen he is pulling up and looking at is being captured. In fact, every action he takes on the screen is all being recorded. But that is not all: the system analytics are comparing Eric's current activity to his previous activity and to the activity his role should produce. Whereas he used to access one or two dormant accounts at most per day, he is now accessing at least six. So an alert is triggered.

Anne, an investigator, gets the alert and looks into it. She does not have to search painstakingly through log files or agent-captured data to find what she needs: the alert gives her the ability to replay all the searches and see everything that Eric saw and watch everything he did. Her suspicions are indeed raised when she watches the replay of account after account.

But Anne is not finished. She performs a link analysis and sees that new online user IDs have been created for five accounts – and that, in each case, the phone number has been changed ... all to the same new number! Furthermore, the link analysis shows payment activity: a large money transfer was recently initiated online from one of the dormant accounts.

Anne documents her findings and writes up her report. With this information in hand, the bank is able to protect the other vulnerable customers, restore funds to the defrauded account, and turn Eric over to the authorities for prosecution.

Stop Insider Threats from the Inside Out

The challenges to protecting your business, your data, and your customers from insider threats are expanding constantly. More employees have more remote access than ever before. Business systems are increasingly complex and interrelated, spreading across mainframe, web, and SaaS applications. Traditional methods of detecting and containing insider theft and fraud such as log-based and agent-based solutions have many benefits, yet are also plagued by issues such as false positives, a lack of visibility, and the danger of violating employee privacy.

Bottomline's network-based Insider Threat Management (ITM) system steps into this gap, providing a vital layer of protection for your organization. ITM empowers you to capture data across business-critical systems and applications in real-time, analyze the information with proactive behavioral analytics, and support investigations with screen-by-screen record-and-replay functionalities plus a wide variety of other tools.

It's time to stop insider threats from the inside out. For more information and a demonstration on how network-based Insider Threat Management can protect your organization.

Contact Us

Bottomline™

REV US013024LM

© Copyright 2015 - 2023 Bottomline Technologies, Inc. All rights reserved.

Bottomline®, Bottomline Technologies® and the Bottomline logo are trademarks of Bottomline Technologies, Inc. and may be registered in certain jurisdictions. All other brand/product names are the property of their respective holders.