10 Reasons SIEM Should Remain Dedicated to Security

The Diverging Needs of Security, IT Operations, and Application Performance

Introduction

As organizations strive to streamline costs and unify their monitoring tools, decision makers might entertain the tempting idea of utilizing security information and event management (SIEM) solutions for IT operations management (ITOM) and application performance management (APM) use cases. Alternatively, they might consider deploying an ITOM tool for security purposes. This white paper serves as a cautionary guide, dismantling such notions by highlighting why they prove impractical and potentially detrimental.

In the world of IT research and consulting, a cadre of analysts specialize in the trifecta of people, processes, and technology within the context of security operations and security operations centers (SOCs). The toolset in this sphere includes SIEM; user and entity behavior analytics (UEBA); and security orchestration, automation, and response (SOAR). These security-focused analysts often collaborate with ITOM experts to strategize the most effective methodologies for Fusion Centers, which integrate network operations centers (NOCs) and SOCs. This joint venture involves evaluating tools that can deliver on both security and ITOM outcomes, and possibly even address APM and observability use cases.

However, empirical evidence indicates that attempting to repurpose SIEM solutions for ITOM use cases (or vice versa) is a misguided strategy. Contrary to initial impressions, employing a single, multipurpose tool does not result in substantial cost savings in terms of licenses, storage, or computational resources. Furthermore, the risks and operational complications associated with deploying a single tool are considerable, and the organizational pitfalls overshadow any perceived financial benefits.

This white paper will elaborate on the key reasons why organizations should avoid using SIEM for ITOM or APM/observability tasks.

The myth of the one-size-fits-all tool

Consider a family asked to describe their perfect car. On the surface, it seems simple — after all, a car just needs four wheels, an engine, brakes and a steering wheel, right? But on deeper exploration, it becomes apparent that the "perfect car" differs drastically between the family members, each bringing unique needs and preferences to the table. Collectively, their demands might be something like this:

- Roomy enough for the whole family or equipment, yet compact enough for easy parking in tight spaces
- Robust and powerful for managing speed and hauling large loads, but also energy efficient for great mileage
- Versatile enough for daily errands, shuttling kids to soccer practice, off-roading on weekends, and the occasional race among friends

Does such a universally perfect car exist? The answer is a resounding "No." That's why we have options like Toyota Camrys and Honda Accords for everyday commuting, minivans for family outings, Jeeps for off-road adventures, and sports cars for racing. Just imagine the conflicts that would arise if this all-in-one car had to be shared among parents and teenage children.

Now, transpose this scenario to organizations in the context of big data analytics tools. At first glance, it might seem that all big data analytics solutions are fundamentally the same, featuring collection, ingestion, and storage of logs and events, plus some form of analytics and reporting, as well as case management. Therefore, shouldn't we be able to use the same tool for ITOM, APM/observability, and security? When we consolidate the unique requirements of these distinct functions, we realize that no single tool can efficiently fulfill all of them:

- A tool to collect and manage all types of data, such as logs, traces, and context
- A single instance of data storage
- Efficient operation with a reasonable total cost of ownership (TCO)
- Capability to handle security needs, including threat detection, investigation, and response, case, and incident management
- Suitability for ITOM use cases such as capacity management
- Applicability for APM and observability teams to optimize diverse applications
- Ability for all teams to access the tool simultaneously, while adhering to role-based access control (RBAC) principles and separation of duties
- Different alerting and notification needs, reporting requirements, and levels of business criticality and service level agreements (SLAs)

Is there such a comprehensive tool with a reasonable TCO? No.

As the title of this paper suggests, we will explore 10 primary reasons explaining the distinct nature of ITOM, APM/observability, and security tools. These reasons can be classified into four categories:

A. User needs
   1. Different stakeholders or owners
   2. Varying user profiles, leading to unique workflows
   3. Diversity in use cases resulting in unique outcomes

B. Data characteristics
   4. Varied data sources across use cases
   5. Variances in logs originating from the same data source
   6. Field variations within the same log from the same data source
   7. Different context enrichment sources

C. Analytical outcomes
   8. Differing analytics methods and rules
   9. Complexity in parsing and modeling

D. Governance practices
   10. Unique requirements for compliance and segregation of duties

In the following sections, we'll delve into each of these differences in further detail.

User needs

1. Different stakeholders or owners

The notion of shared tools can initially seem enticing, with the promise of cost sharing among different teams including ITOM, APM/observability, and security. However, such arrangements often get more complicated when divergent stakeholder priorities arise. Balancing a range of complex use cases in a single tool often leads to performance degradation across the board. Which use cases are more complex, require more compute or are data-intensive? What about security and compliance? Which use cases will be decommissioned to make room for a newly prioritized use case? How are local and international privacy standards being applied?

Ultimately, the organizational and operational overheads associated with shared tools can often eclipse the potential savings of jointly used tools. In fact, these overheads could not only nullify the cost benefits, they could lead to inefficiencies and organizational conflict, making such an arrangement more disruptive and less viable.

2. Varying user profiles, leading to unique workflows

The distinct nature of ITOM, APM/observability, and security results in unique user profiles and, consequently, different workflows.

For example, SIEM and security log management (SLM) users, focused on threat detection, investigation, and response (TDIR), operate in a very different workflow from those working in APM/observability. The latter are primarily concerned with measuring and improving the performance of applications.

Similarly, ITOM users follow another distinct workflow aligned to tracking the operational aspects of an IT organization, its assets, and its resources. So, when adopting the people, process and technology approach, it becomes evident that different workflows and processes require different tools. The attempt to blend these distinct workflows into a single tool might not only compromise efficiency, but also lead to conflicts and operational challenges.

3. Diversity in use cases, resulting in unique outcomes

Trying to serve a wide range of use cases and outcomes with a single, general-purpose tool can be analogous to trying to "converge a toaster and a refrigerator," as Apple CEO Tim Cook once quipped. The resulting product may be neither satisfactory nor efficient.

ITOM, APM/observability, and security operate using distinct use cases and desired outcomes. Attempting to unify these within a single tool is asking a lot and increases the degree of difficulty in achieving the intended outcomes. Instead of attaining efficiency, such a strategy might lead to suboptimal performance across all disciplines. SIEM, on the other hand, is a solution that was designed by security professionals specifically for use by other security professionals.

Data characteristics

4. Varied data sources across use cases

Different use cases require a diverse range of data sources. For example, ITOM needs data that reflect the state of the IT infrastructure, providing detailed insights into the performance, usage, and availability of assets and resources; APM and observability, on the other hand, require application-centric data that supply intricate details about each application's performance. In contrast, security use cases require many additional data sources, streaming feeds, and user and application context.

While it might seem convenient gathering as much data as possible, it's important to avoid irrelevant data sources. The costs associated with collecting, storing, and processing irrelevant data can balloon quickly, often with diminishing returns. And data volumes are only growing.

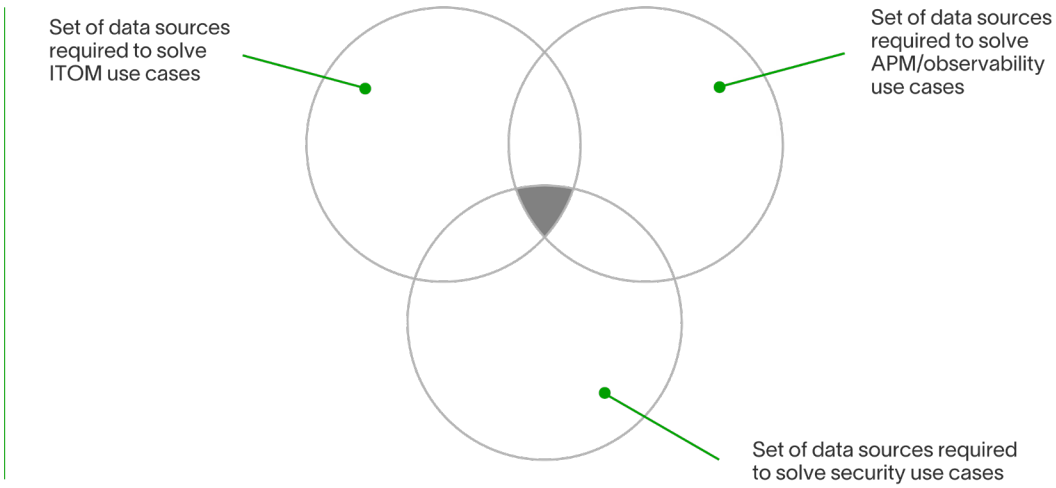

Figure 01 A Venn diagram representing the specific data sources required for each individual use case. Each use case demands one or more unique data sources.

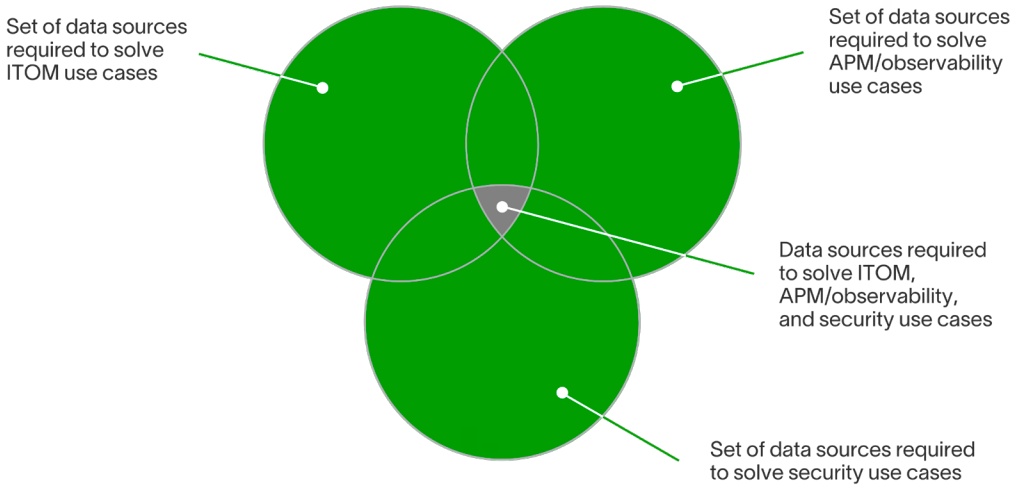

Figure 02 A Venn diagram presenting all the data sources that would need to be considered to address all three use cases, regardless of the number of solutions employed

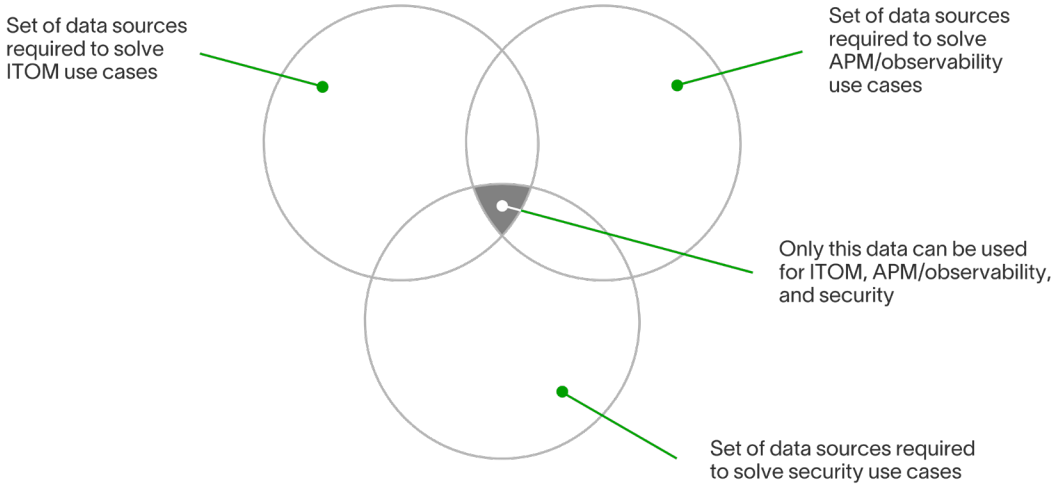

Figure 03 A Venn diagram showing the potential savings realized by using a joint tool for all three use cases

As illustrated in Figure 3, the utilization of a joint tool for all use cases yields only minimal savings in terms of data collection and storage. This implies that the potential drawbacks and complexities associated with a joint tool far outweigh any marginal benefits.

Numerical Code	Keyword	Facility name
0	kern	Kernel messages
1	user	User-level messages
2	mail	Mail system
3	daemon	System Daemons
4	auth	Security messages
5	syslog	Syslogd messages
6	lpr	Line printer subsystem
7	news	Network news subsystem
8	uucp	UUCP subsystem
9	cron	Clock daemon
10	authpriv	Security messages
11	ftp	FTP daemon
12	ntp	NTP subsystem
13	security	Security log audit
14	console	Console log alerts
15	solaris-cron	Scheduling logs
16-23	local0 to local7	Locally used facilities

Table 01 23 families of Linux logs[1]

5. Variances in logs originating from the same data source

Logs used across different use cases can vary significantly, even when they are derived from the same data source. For instance, let's consider the various categories of Linux logs. Some of these logs are mainly relevant to ITOM use cases, some align more to APM/observability, while others are more suitable for security use cases.

As previously mentioned, addressing all three use cases requires the union of all relevant logs. Here, again, we observe only marginal storage savings when employing a single, joint solution.

The illustration in Figure 4 underscores that while there is some overlap, most log data are unique to each use case, limiting the efficiency of a single, unified tool. Therefore, it becomes crucial to recognize and appreciate these distinctions while choosing the appropriate tools for different organizational needs.

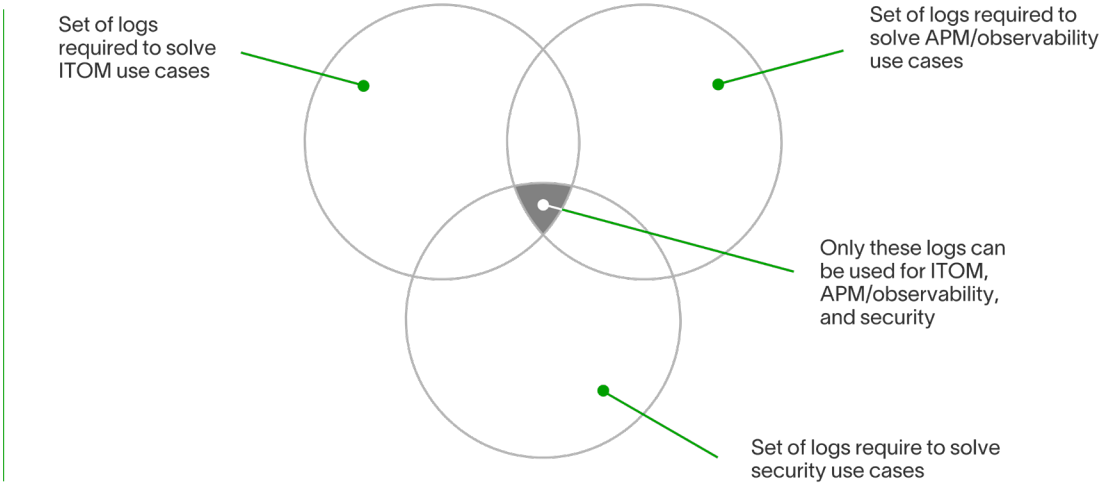

Figure 04 A Venn diagram depicting the limited set of Linux logs that are applicable across all three use cases.

[1] https://devconnected.com/linux-logging-complete-guide/

6. Field variations within the same log from the same data source

Even when using the same log from the same data source is for different use cases, the required fields can vary drastically.

Value	Severity	Keyword
0	Emergency	`emerg`
1	Alert	`alert`
2	Critical	`crit`
3	Error	`err`
4	Warning	`warning`
5	Notice	`notice`
6	Informational	`info`
7	Debug	`debug`

Table 02 The eight degrees of verbosity in Linux logs. Depending on the specific log and the configured verbosity level, a Linux log can contain anywhere from dozens to hundreds of fields.

Contrary to what many vendors might tell end users, parsing all the fields from all logs is not only inefficient, but also prohibitively expensive. Tools should be configured to parse only the relevant fields from the appropriate logs in the pertinent data sources for each specific use case:

- ITOM tools parse only the ITOM-specific fields from ITOM-relevant logs in ITOM-related data sources
- APM/observability tools parse only the APM/observability-specific fields from APM/observability-relevant logs in APM/observability-related data sources
- Security tools parse only the security-specific fields from security-relevant logs in security-related data sources.

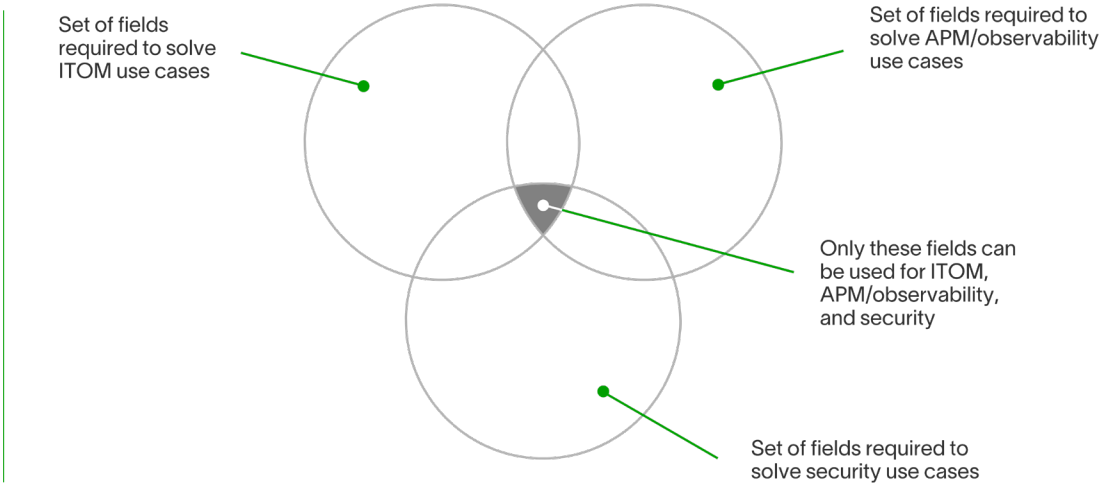

Figure 05 A Venn diagram illustrating the set of fields required for each use case.

This emphasizes the specificity of each use case and reinforces the argument against the one-size-fits-all approach when choosing security/IT tools. Instead, the focus should be on tailoring the tool selection to the requirements of the individual use case.

7. Different context enrichment sources

Raw data used by ITOM, APM/observability, and security tools often require context enrichment from various sources. Consider the disparity in insights provided by these two descriptions of the same event:

1. Error code <123> from IP address 192.168.30.30
2. An authentication failure occurred (indicated by error code <123>) at the workstation with IP address 192.168.30.30. This server operates on Windows 11 version 22H2, is owned by John, and has a MAC address of d0:89:0d:8b:f6:22.

To transition from the first, less informative representation to the second, more detailed, context-rich second one, raw event logs need to be context-enriched using various internal and external sources. In this case, this process would include:

- Correlation with DHCP
- Lookup in CMDB
- Lookup in AD

The need for context enrichment is directly linked to the insights we aim to extract, and these insights are derived from the required use cases and outcomes. Different outcomes and use cases will demand unique context enrichment sources. If too many disjointed enrichment sources are used to serve unrelated use cases, the operational cost of the tool can rise, often offsetting any potential savings associated with a joint tool. User identity and role, for example, are critical to security use cases. Therefore, it's important to carefully consider the cost-benefit balance when deciding on toolsets.

Analytical outcomes

8. Differing analytics methods and rules

The analytics engines and rules vary drastically across different use cases, and are based on the insights that the tools are expected to provide. ITOM and APM/observability mainly concentrate on machine behaviors, such as connectivity, loads, and machine-to-machine communication within IT networks, or the performance of application code within a container. These behaviors can usually be modeled using finite state machines[2], and deterministic approaches[3] provide solutions for most issues. They are often addressed through conditional "if -then-else" rules, or simple mathematical formulas to compute variables like the average round-trip delay between two systems or the mean response time for a specific application.

In contrast, security — especially when dealing with active attackers — is anything but deterministic. The unpredictable nature of human behavior leads to an endless array of attack vectors and creative misuse of an organization's resources. To efficiently address security issues, heuristic approaches[4] are essential, heavily relying on advanced analytics and machine learning models. For example, UEBA is designed to establish norms for behavior, and can flag anomalies typically indicative of an active attack.

These deterministic and heuristic analytics have fundamental differences and require different engines and content (for example, different input data, rules, etc.) to operate effectively. A tool designed to concurrently address ITOM, APM/observability, and security use cases would need to continuously run all these engines simultaneously. This operation would lead to massive computational and storage costs, outweighing any potential savings from a combined tool. Even worse, the added complexity of operating these vastly different engines in parallel could introduce additional costs that would be avoided with dedicated tools.

9. Complexity in parsing and modeling

With the varying fields required to handle ITOM, APM/observability, and security (even when working with the same data sources and logs), specific parsers are necessary for each log. This introduces two major challenges for a single tool trying to cover all three use cases:

1. There would be a significant computational load and temporary storage demand to parse all logs for all use cases.
2. A highly complex data model would be required to accommodate all use cases.

These demands could contribute to performance problems. Only a handful of SIEM solutions can scale beyond a million events per second (EPS). Adding the task of parsing for non-security use cases like ITOM and APM/observability would necessitate a prohibitively high investment to sustain an acceptable performance level. This investment covers not only the license costs (hot/cold storage), but the compute costs of querying across days, weeks, months, or years of data.

[2] https://en.wikipedia.org/wiki/Finite-state_machine
[3] https://en.wikipedia.org/wiki/Deterministic_system
[4] https://en.wikipedia.org/wiki/Heuristic_(computer_science)

10. Unique requirements for compliance and segregation of duties

Segregation of duties is a fundamental principle in demonstrating compliance with external regulations like PCI, HIPAA, and GDPR, as well as internal policies that ensure protection against malicious insiders. For a single tool, it's crucial not to play the roles of both judge and jury, as this could compromise the demonstration of compliance and jeopardize the validity of the process.

Using a joint tool for security and compliance, ITOM, and APM/observability could easily violate key tenets of duty segregation. Potential conflict areas include:

- Ownership governance — Could there be conflicts of interest with a single owner or joint ownership?
- Administrative governance — How would admin privilege sharing be handled?
- Engineering team governance — Who is responsible for developing disjointed use cases?
- User governance — How would role over-provisioning be managed?
- Long term data storage and archiving — How many years of data are you required to maintain?

Maintaining proof of custody, proper accountability, and role-based access control (RBAC) becomes exponentially more complicated when too many individuals have access to the tool. The old adage rings true here: Too many cooks in the kitchen can spoil the broth.

Conclusion

Many organizations would like to have a single joint tool capable of handling security, ITOM, and APM/observability. In theory, this might seem achievable, beneficial, and cost effective. There are even fringe cases where a reduced scope might make this viable for smaller or very specialized organizations. However, a deeper examination of the 10 distinct aspects discussed in this paper reveals that this is generally an unwise strategy.

The drive towards unification should be carefully balanced against the unique requirements, complexities, and costs associated with each field. Attempting to create a "one size fits all" solution might compromise the efficiency and effectiveness of each individual component, diluting the success of them all.

While it's understandable that organizations wish to reduce complexity and cost, the ambitious endeavor of creating a unified solution for security, ITOM, and APM/observability may result in a tool that doesn't perform optimally in any of these areas. It's essential to remember that the more features a tool aims to cover, the more complex it becomes to manage, maintain, and secure, which could lead to elevated costs and decreased performance.

Additionally, compliance requirements and the necessity of maintaining segregation of duties make it critical to ensure tools used for security aren't also used for ITOM and APM/observability.

In summary, while it's an appealing concept, it's often best to stick with specialized tools for specific tasks. The desire for a unified, comprehensive tool should be weighed against the potential tradeoffs in performance, complexity, cost, and compliance. Always remember that the dream of simplicity and unification may turn into a nightmare of complexity and increased overhead.

About Exabeam

Exabeam is a global cybersecurity leader that created New-Scale SIEM™ for advancing security operations. We help organizations detect threats, defend against cyberattacks, and defeat adversaries. The powerful combination of our cloud-scale security log management, behavioral analytics, and automated investigation experience results in an unprecedented advantage over insider threats, nation states, and other cyber criminals. We understand normal behavior, even as normal keeps changing — giving security operations teams a holistic view of incidents for faster, more complete response.