

Technical White Paper

Fluency vs. Traditional Database-Driven Anomaly Detection



INTRODUCTION

As cybersecurity threats become increasingly sophisticated, the ability to detect and respond to anomalies in real-time is critical. Traditional Security Information and Event Management (SIEM) systems, primarily reliant on database-driven approaches, face significant challenges in meeting the demands of modern security environments. This white paper examines the advantages of **Fluency**, a SIEM platform that leverages streaming analytics, over traditional database-driven SIEM products, with a focus on its advanced handling of anomaly detection using archetypes and scope.

THE EVOLUTION OF ANOMALY DETECTION

TRADITIONAL DATABASE-DRIVEN SIEMS

Traditional SIEMs have historically relied on database-driven architectures for anomaly detection. In these systems, data is collected and stored in databases, with queries run at regular intervals to detect deviations from established baselines. While this method has served well in the past, it presents several limitations:

Latency: Anomalies are detected after querying, leading to delays in response.

Resource Intensiveness: Frequent database queries and the need for manual updates to detection models increase operational overhead.

Scalability Challenges: As data volumes grow, scaling traditional SIEMs can become complex and costly.

Lack of Contextualization: Traditional SIEMs often treat anomalies without sufficient differentiation between different types of entities or the broader context of the anomaly within the organization.

FLUENCY'S STREAMING ANALYTICS APPROACH

Fluency simplifies anomaly detection by utilizing streaming analytics, an approach that processes data in real-time as it flows through the system. This method enables the immediate detection of anomalies without the need for periodic queries and incorporates a more sophisticated understanding of anomalies through archetypes and scope:

Real-Time Detection: Anomalies are identified as soon as they occur, allowing for prompt responses.

Continuous Learning: Fluency continuously updates its understanding of what constitutes normal behavior, reducing the risk of undetected threats, while reducing false-positives.

Scalable Efficiency: Fluency's architecture effortlessly scales with data volume, ensuring consistent performance and low operational costs.

Broader Context with Scope: Fluency evaluates whether an anomaly is unique to a specific individual or if it represents a larger trend across the organization, providing a more comprehensive understanding of potential threats.

TECHNICAL COMPARISON: FLUENCY VS. TRADITIONAL SIEMS

	TRADITIONAL SIEMS	FLUENCY
Real-Time Anomaly Detection	In traditional SIEMs, data is collected and stored in a database, with queries executed at defined intervals to detect anomalies. This batch processing approach introduces latency between the occurrence of an anomaly and its detection, potentially delaying the response to threats.	Fluency's streaming analytics processes each data point as it enters the system, enabling immediate anomaly detection. This real-time processing is critical in environments where rapid response is essential to mitigate risks and prevent breaches.
Continuous Learning and Adaptation	Traditional SIEMs typically rely on static models that require periodic updates. These updates are often manual, and the models can become outdated between refresh cycles, leading to potential gaps in detection capabilities.	Fluency employs dynamic algorithms that continuously learn from incoming data, automatically adjusting to new patterns and behaviors. This continuous adaptation ensures that the system remains effective in detecting anomalies, even as the environment changes.
Contextual Awareness: Scope	Traditional SIEMs often apply a one-size-fits-all approach to anomaly detection, treating all entities and anomalies similarly. This can lead to inaccuracies, as the context of the entity (e.g., a system administrator vs. a regular user) and the broader organizational impact of the anomaly are not adequately considered.	Fluency assesses anomalies on both local and global levels, determining whether an anomaly is unique to an individual or part of a broader trend. For example, if a user logs in from a new location, Fluency checks if this is isolated or part of a larger pattern, prioritizing alerts based on the potential impact.
Scalability and Efficiency	As data volumes increase, traditional SIEMs require additional resources to maintain performance. Scaling these systems often involves complex reconfigurations, increased hardware requirements, and higher operational costs.	Fluency's streaming architecture is designed for scalability. By processing data in real-time and maintaining a streamlined state, Fluency can handle large and fluctuating data volumes without compromising performance or requiring extensive resources.
Reduced Operational Overhead	The reliance on database queries and manual model updates in traditional SIEMs leads to significant operational overhead. This includes increased processing power, memory usage, and administrative effort.	Fluency minimizes operational overhead by maintaining an up-to-date state of the system as data flows in. This approach reduces the need for frequent querying and manual intervention, freeing up IT resources for more strategic tasks.

CONCLUSION

Fluency's adoption of streaming analytics provides a robust, scalable, and efficient solution for anomaly detection, surpassing the capabilities of traditional database-driven SIEMs. Its ability to detect anomalies in real-time, continuously learn and adapt, and operate with reduced overhead makes Fluency an ideal choice for modern security environments. As businesses face increasingly complex threats, the need for a SIEM solution that can keep pace with these challenges has never been greater. Fluency delivers the advanced capabilities required to meet these demands, providing a future-ready platform for proactive security management.