cyber futurists

January 2025 Rev. 2

Emerging Cybersecurity Technology Report

Security Telemetry Pipelines: Market Overview and Dynamics

Chief Futurist: Oliver Rochford

Sponsored by AUGURIA™

cyber futurists
Future-Proof Cybersecurity Advisory

Emerging Cybersecurity Technology Report

Security Telemetry Pipelines

Market Overview and Dynamics

Publication Date: January 2025 Rev. 2

Lead Analyst: Oliver Rochford

Cyberfuturists

https://cyberfuturists.com

oliver.rochford@cyberfuturists.com

Special Thanks to Contributors and Reviewers:

Anton Chuvakin, Prateek Bhajanka, Josh Cowling

Special Thanks to our Sponsor Auguria.

Executive Overview

The vast and growing volume of security data from endpoints, networks, cloud environments, and third-party services has outpaced the capabilities of traditional security tools like SIEM. Enter **security telemetry pipelines**—an emerging technology for managing, curating, and analyzing security data in large-scale, decentralized, or federated security deployments. As organizations adopt hybrid work models and expand their digital ecosystems, these pipelines are emerging as a critical enabler of efficient and effective security operations.

Summary

WHY IT MATTERS:

Traditional security operations struggle with distributing large volumes of disparate data, resulting in alert fatigue, missed threats, and high operational costs. SIEM were designed as data consumers, but due to shifting operational demands, what security organizations need are data refineries and distributors. **Security telemetry pipelines** address these challenges, enabling organizations to efficiently process, integrate, and act on security information across environments.

Security Telemetry Pipelines or Data Pipelines, what's the difference?

Existing data pipeline capabilities do a sufficient job at transforming data and ETL for general-purpose data operations. However, using existing solutions for security use cases incurs a "repurposing tax". Most of the detail and context required by the security Analysts, like typical security data formats, or transformations required for security use-cases must be custom developed and encoded in the form of parsers and rules. Building and managing rules takes a lot of human effort and resources, leading to mission creep and distracting security analysts from their core job: security operations. Security Telemetry Pipeline integrate security context out of the box and ensure that data comes ready for security operations use-cases without onerous manual customization.

PROMISED BENEFITS:

- **Cost Reduction:** Removing redundant and repetitive data automatically reduces data processing and storage costs, especially those related to SIEM solutions charging based on ingest.
- **Efficiency**: Automated and guided transformation, enrichment, and classification significantly reduce manual workloads for security teams, enabling them to focus on more strategic security operation activities like threat hunting.

- **Scalability**: Security telemetry pipelines can handle increasingly complex environments where the data lifecycle means that it is reshaped and redistributed across a growing variety of security analytics and automation solutions repeatedly
- **Better threat detection**: To manage cost and operational deficiencies, security-relevant data is today rarely analyzed fully in aggregate. Analyzing and enriching data viewed as a whole before distributing to individual point solutions improves detection accuracy and helps security teams prioritize threats more effectively across a growing attack surface.
- **Vendor-agnostic**: Organizations can integrate and distribute data across multiple platforms and technologies, avoiding vendor lock-in and leveraging the best available tools to strike the right balance among cost, performance, and scale.

WHAT'S NEXT:

As the market evolves, we anticipate greater convergence between SIEM, XDR, and security telemetry solutions. **Security data lakes**, **Cybersecurity Mesh Architecture (CSMA)**, and **SOCless, or rather, Headless SIEM models** are likely to gain traction as organizations seek more flexible and scalable approaches to managing their security posture. Embracing these emerging solutions will be key to staying ahead in the rapidly evolving cybersecurity landscape. By 2030, the Security Telemetry Pipeline Market will have moved to a mature phase, with most players acquired and a few matured into more fully featured security operations solution providers.

Key Takeaways from the Security Telemetry Pipeline Market Overview

The report discusses the evolving landscape of security operations, highlighting the limitations of traditional approaches and the emergence of new technologies and architectural concepts.

Key Takeaways

Security telemetry pipelines are transforming security operations. They address the limitations of SIEM solutions, which struggle to handle the volume and complexity of data generated, especially at the velocity it is generated in modern IT environments. Security telemetry pipelines automate key processes such as data ingestion, normalization, enrichment, and prioritization for security operations use cases

- **Addressing SIEM Limitations:**

  Security telemetry pipelines resolve critical challenges with traditional SIEM solutions, including their inability to handle the growing volume, complexity, and velocity of security data. By automating processes like ingestion, enrichment, and prioritisation, telemetry pipelines provide ready-to-use, actionable data for security teams.

- **Operational Efficiency and Cost Savings:**

  Organisations can significantly reduce costs by minimising redundant data, which lowers SIEM ingestion and storage expenses. Automation of data transformation and enrichment further enhances operational efficiency, allowing teams to focus on proactive threat hunting and strategic initiatives.

- **Enhanced Threat Detection and Response:**

  By enriching data streams with contextual information and applying advanced analytics, security telemetry pipelines improve the accuracy of threat detection. Analysts can prioritise threats more effectively, reducing response times and mitigating potential damage.

- **Emergence of Flexible Ecosystems:**

  The shift toward decentralised and flexible security architectures is gaining momentum. Solutions like security data lakes, Cybersecurity Mesh Architecture (CSMA), and SOCless or Decoupled SIEM models enable organisations to move beyond traditional, siloed approaches, fostering interoperability and scalability.

- **Complementary Role of SIEM and XDR:**

While SIEM provides broad monitoring and compliance capabilities, XDR excels in real-time threat detection. These technologies can converge and complement each other, but the rise of hybrid architectures blending SIEM, XDR, and data lakes using telemetry pipelines is reshaping the landscape.

- Security Data Lakes for Deep Analysis:

  Security data lakes offer unparalleled flexibility for historical analysis, compliance, and threat hunting. Although they require significant technical investment, their ability to retain and reanalyse raw telemetry data makes them invaluable for long-term security strategy.

- Challenges in Adoption:

  Implementing advanced architectures like SOCless or Decoupled SIEM models requires technical maturity, skilled personnel, and a security-driven organisational culture. Overcoming these hurdles is critical to maximising the value of these solutions, and a strong use-case for telemetry pipelines.

- Real-Time Capabilities for Modern Threats:

  Modern telemetry pipelines support real-time enrichment and detection, enabling organisations to identify and respond to emerging threats without delays caused by batch processing.

- Vendor-Agnostic Solutions for Greater Flexibility:

  Telemetry pipelines allow organisations to integrate data across diverse platforms, avoiding vendor lock-in and enabling them to leverage the best tools for specific needs.

- Market Evolution and Future Trends:

  The convergence of SIEM, XDR, and telemetry pipelines is driving innovation in the security operations market. By 2030, the market will likely mature, with key players offering integrated solutions. Early adopters can gain a competitive edge by embracing these technologies today.

Introduction

As organizations struggle to make sense of the vast amount of security data generated from distributed networks, endpoints, cloud environments, and third-party services, traditional solutions such as Security Information and Event Management (SIEM) platforms are reaching their limits. Modern security telemetry pipelines provide a flexible, scalable approach to ingesting, processing, and analyzing security data, transforming how security teams respond to threats.

THE NEED FOR SECURITY TELEMETRY PIPELINES

Security telemetry pipelines are becoming essential for modernizing security operations. The hybrid work models, distributed environments, and rapid cloud adoption have left organizations struggling to manage a massive influx of security data from diverse sources. SIEM tools, though useful for aggregating and correlating data, were not designed for the scale and complexity of current environments. Often metered by the volume of ingested data, SIEM also often suffers from high operational costs. Today's security operations are far more diverse and distributed than last decade, requiring cheap and fast data proceeding at scale. SIEM was originally designed as a single centralized destination to normalize, aggregate, and correlate security data across an organization. Modern security operations increasingly require data refinement and redistribution.

In contrast, security telemetry pipelines can automate and generalize data ingestion, normalization, transformation, refinement, and enrichment, typically integrating a data lake for raw data storage, replay, and compliance, and can route data optimally prepared to different security data destinations, including SIEMs, data lakes, SOAR, AISecOps or cloud object storage. significantly reducing the burden on security teams. They improve the efficiency of SIEM, XDR (Extended Detection and Response), and security data lakes, allowing organizations to gain more actionable insights, prioritize alerts effectively, and scale their security operations.

Market Drivers

THE EVOLUTION OF SIEM AND XDR: COMPLEMENTARY OR COMPETITIVE?

SIEM has long been the cornerstone of security operations, providing organizations with centralized log management and real-time analysis. Since its inception, SIEM has gone through multiple generations, evolving from basic log aggregation to advanced platforms that incorporate machine learning, cloud integration, and data lakes. Its role has expanded beyond compliance monitoring to encompass more sophisticated threat detection and incident response capabilities.

In recent years, XDR has emerged as a potential complement, or in some cases, competitor to SIEM. XDR builds on the capabilities of Endpoint Detection and Response (EDR) by integrating telemetry from multiple sources, such as endpoints, networks, and cloud services, into a unified threat detection platform. Unlike SIEM, which is often seen as broad but complex to configure, XDR offers more streamlined, out-of-the-box capabilities focused on threat detection and response.

While some argue that XDR could eventually replace SIEM in certain environments, the reality is more nuanced. The two technologies serve different, if partially overlapping, use cases, and many organizations benefit from deploying them together. SIEM's flexibility allows it to handle a wider range of use cases, including compliance and broader security monitoring, whereas XDR excels in real-time detection and response with less manual customization required. In many ways, they can be viewed as complementary, sitting on opposite ends of a security spectrum that also includes security data lakes.

The key differences between SIEM, XDR, and Data Lakes can be summarized as follows:

	SIEM	**XDR**	**Security Data Lake**
Use-case Coverage	SIEM platforms provide broad coverage for security monitoring, incident detection, compliance reporting, and log management. They are versatile, integrating with various tools and technologies, which makes them the go-to for security operations centers (SOCs) managing heterogeneous environments. But they are designed as data destinations, not as data distributors.	XDR is more specialized, with a deep focus on detecting, investigating, and responding to threats across multiple security layers—endpoint, network, identity, and cloud. It's designed to simplify and centralize detection and response, making it an ideal fit for threat-centric security operations.	SDLs focus on long-term data storage and analysis. Unlike SIEMs or XDR, data lakes are not built primarily for real-time threat detection. Instead, they shine when it comes to historical data analysis, threat hunting, and forensic investigations. They offer an unmatched ability to retain high volumes of raw telemetry, making them indispensable for organizations with compliance, auditing, and in-depth threat research needs.
Detection & Analytics Engineering Effort	SIEMs require significant customization and engineering effort to create detection use cases and tailor the	XDR platforms, especially those aligned with a specific vendor ecosystem, come with pre-built,	Typically, data lakes necessitate significant engineering effort because they are not pre-designed for real-

	platform to specific organizational needs. While many SIEMs offer an overwhelming collection of out-of-the-box content, this requires configuration and much of the power lies in the ability to build custom queries and rules.	curated detection capabilities. These platforms are designed to require minimal engineering effort, making it easier to implement threat detection without heavy customization.	time detection. Advanced analytics and machine learning models must be layered on top to extract meaningful insights. However, the payoff is the ability to run deep analytics across years of data, which neither SIEM nor XDR can typically achieve without performance bottlenecks.
Customization	Offers the most flexibility when it comes to out-of-box customization. Organizations can build bespoke use cases, alerts, and workflows tailored to their unique environment. SIEMs are designed for custom rule creation, making them a powerful option for enterprises that need fine-grained control over their security monitoring.	Customization in XDR solutions is often more limited because they are built around a specific detection framework and vendor ecosystem. While this reduces engineering overhead, it also limits flexibility, especially when integrating with non-native tools.	Data lakes offer unparalleled flexibility for custom analytics, as they allow you to ingest, store, and analyze any type of data from any source. However, they don't come with pre-built security rules, so customization depends heavily on building analytics layers that can sift through the raw data and uncover meaningful patterns.
Data Integration	SIEMs can aggregate data from diverse sources, firewalls,	XDR platforms typically integrate tightly with specific	Data lakes are vendor-neutral and can integrate with

	endpoints, cloud services, and more, giving them a vendor-agnostic edge. They are designed to integrate with a broad spectrum of tools and technologies, making them a strong choice for organizations with heterogeneous environments.	vendor ecosystems, which can be a double-edged sword. While this tight integration simplifies deployment and enhances native detection capabilities, it also limits flexibility for organizations that use products from multiple vendors.	any system capable of generating data. They are designed to be a central repository for all security telemetry, and they work well in multi-cloud or hybrid environments, where diverse data types from different vendors need to be stored and analyzed.
Focus	SIEMs are designed for broad security use cases, covering everything from compliance and log management to threat detection and incident response. Their primary strength is their versatility. They can be tailored to fit almost any enterprise security need.	XDR is purpose-built for threat detection and response. Its entire focus is on identifying and mitigating threats quickly across multiple security layers, making it ideal for security teams focused on active threat management.	The focus of data lakes is long-term data retention and advanced analysis. They are less about real-time detection and more about enabling deep dives into historical data, threat hunting, and big-picture security analytics. For example, data lakes are perfect for reconstructing attacks that happened months or even years ago, something that's beyond the reach of many SIEMs and XDRs.

Table 1: Key differences between SIEM, XDR, and Security Data Lakes.

Rather than viewing SIEM, XDR, and Data Lakes as entirely separate categories, it can be helpful to think of them as points on a spectrum of security information management solutions. Some vendors are blending SIEM and XDR capabilities, creating hybrid solutions that aim to offer the best of both worlds. As the market evolves, we may see further convergence between these technologies.

The Rise of Security Data Lakes and SOCless Models

As organizations grapple with ever-growing volumes of security data, security data lakes are becoming an attractive alternative to traditional SIEM and XDR architectures. Unlike SIEM, which requires predefined schemas for data ingestion, security data lakes offer basic data storage and replay capabilities, allowing organizations to store massive amounts of raw data and analyze and reanalyze it as needed.

While SIEM and XDR dominate the market, some forward-thinking organizations are exploring alternative approaches. Netflix's "SOCless Detection Team" concept, introduced in 2018, showcased a stack of different data and micro-solutions built with a security data lake architecture. Google's Autonomic Security Operations Model is also gaining traction and advocates for a highly automated approach to security operations.

This approach is particularly appealing to lean-forward organizations, typically found among Fortune 1000 companies, FAANG (Facebook, Amazon, Apple, Netflix, Google), and other tech giants. These companies often have the resources and expertise to build custom solutions tailored to their specific needs.

Security data lakes offer outstanding flexibility and power, but they require significant investment in both technology and skilled personnel to implement and maintain effectively. While they currently represent less than an estimated 5% of the security operations platform spend, this is expected to grow to around 10% by 2028, driven by increasing interest in this approach among larger, more technically sophisticated organizations and service providers.

Netflix's "SOCless Detection Team" concept is a prime example of how forward-thinking organizations are leveraging security data lakes to move beyond traditional SOC models. By utilizing a stack of micro-solutions built on a security data lake architecture, Netflix has created a highly flexible and scalable security environment, reducing the need for a large, centralized SOC.

However, security data lakes require significant investment in both technology and skilled personnel, making them more suitable for large or well-resourced organizations with the resources to implement and maintain such systems effectively.

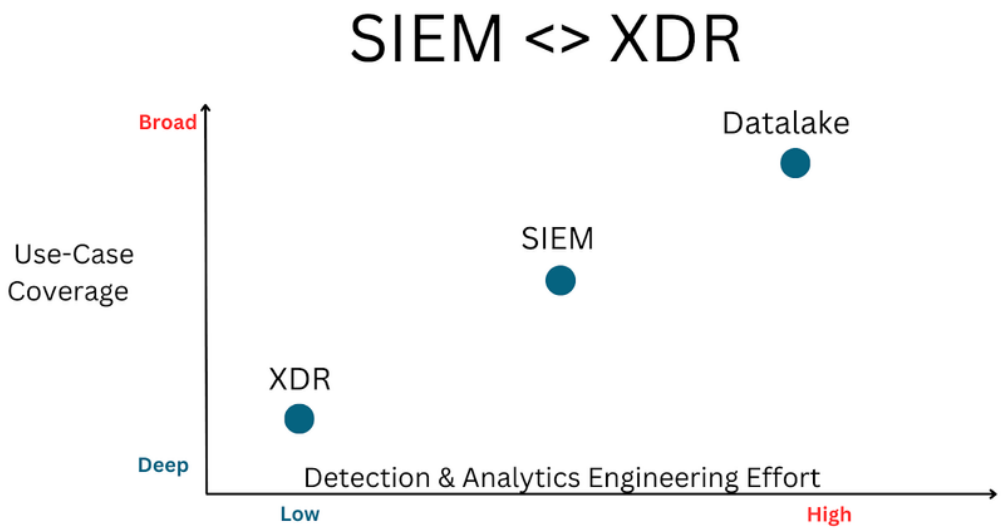

FIG 1: XDR, SIEM, and Security Data Lakes exist on a spectrum

SOCless Architectures: Distributed Security Operations

The concept of SOCless detection, as pioneered by companies like Netflix, represents a radical shift in how security operations are conducted. Rather than relying on a traditional SOC with centralized monitoring and incident response, SOCless models distribute security responsibilities across the organization, leveraging automation and developer-centric tools to detect and respond to threats.

In a SOCless environment, security becomes everyone's responsibility. Automated tools handle much of the day-to-day monitoring and response, while developers integrate security into their continuous integration and delivery (CI/CD) pipelines. This approach reduces the need for a large SOC team and allows for more agile, responsive security operations.

SOCless or SIEMless? Anton Chuvakin and I have been debating what to call this market segment, with neither of us liking the term SOCless. SIEMless, Headless SIEM, and Decoupled and Deconstructed SIEM are all terms we've considered.

While SOCless architectures offer significant benefits, including reduced operational overhead and increased automation, they are not suitable for every organization. A founder recently called it a "Product SOC", because they encounter the model primarily amongst tech and SaaS startups. Implementing a SOCless model requires a high level of technical and engineering maturity and a strong security-centric DevOps culture, as well as the ability to develop and maintain custom detection and response tools.

Emerging Architectural Concepts: Data Fabric and Cybersecurity Mesh Architecture

As security operations evolve, new architectural frameworks are emerging to help organizations manage the growing complexity of their IT environments. Two such concepts are data fabric and Cybersecurity Mesh Architecture (CSMA).

- **Data Fabric**: Defined by Gartner as "an integrated layer of data and connecting processes, data fabric uses continuous analytics to provide seamless data management across hybrid and multi-cloud environments".

https://www.gartner.com/smarterwithgartner/data-fabric-architecture-is-key-to-modernizing-data-management-and-integration

In the context of security, data fabric allows organizations to integrate security data from multiple sources, improving visibility and enabling more efficient threat detection and response.

- **Cybersecurity Mesh Architecture (CSMA)**: Gartner's CSMA is described as a "composable, distributed security architecture designed to secure modern enterprises with complex, distributed IT environments.".

https://www.gartner.com/en/information-technology/glossary/cybersecurity-mesh

Both Data Fabrics and CSMA represent a move toward more flexible, interoperable security ecosystems that can adapt to the unique needs of each organization. Security Telemetry Pipelines are seen as an emerging technology to enable modular and distributed architectures.

Challenges of Integration and the Rise of Cybersecurity Mesh Architecture

As organizations adopt a range of tools—from SIEM to XDR to security data lakes—integrating these systems has become a critical challenge. Many of these technologies lack interoperability, resulting in information silos that complicate obtaining a comprehensive understanding of an organization's security posture.

To address this, some organizations are exploring the concept of a cybersecurity mesh architecture (CSMA). This approach aims to create a flexible, modular security ecosystem by connecting distributed tools and centralizing data control. CSMA allows organizations to deploy different security tools while ensuring they can communicate and share information effectively. This creates a more integrated and interoperable environment, allowing for better collaboration between different technologies.

However, the adoption of CSMA is still in its early stages, and many organizations are hesitant to invest in new architectures without clear evidence of their effectiveness. Similarly, the concept of security data fabrics, which aim to provide a unified layer of data across all environments, is gaining attention but remains largely untested in many environments.

What about security data fabrics?

Some vendors are messaging themselves as security data fabric solutions, but so far no major industry analyst firm has followed suit, possibly because the term has been submitted as a trademark .

Modernizing SecOps with Security Telemetry Pipelines

Security data pipelines are useful for modernizing security operations. Traditional SIEMs are often criticized for alert overload, manual tuning requirements, and their inability to scale as data volumes grow. Security Telemetry Pipelines solve these issues by automating much of the data normalization and processing tasks and ensuring the data is labelled and shaped optimally to enable threat detection.

Security data pipelines don't just do ETL, either. ETL and data transformation are not the real value. There are dozens of solutions to transform streaming event data, from bare-metal Python to observability pipeline solutions like Cribl or Confluent.

If a vendor solely performs deduplication and schema transformation, the value is purely based on cost arbitrage, effectively creating a deploy-and-forget scenario. It is, however, very sticky. Once someone has deployed a security pipeline solution and reduced their license costs, they cannot easily remove it, very similar to a pacemaker.

A modern security telemetry pipeline cleans up, deduplicates, and compacts incoming data before it gets to the SIEM, or a data lake. They can also sort and prioritize events and alerts so that an analyst or automation tool can action them faster and more reliably. This reduces the volume of alerts generated by eliminating duplicates and irrelevant data, allowing security teams to focus on high-priority threats. The result is a more efficient SIEM, one that can keep pace with the scale of modern security environments.

Operating Models

SIEM AUGMENTATION AND OFFLOADING

Replacing SIEM collectors and log shippers with pipelines to reduce SIEM costs and offload compliance data to low-cost storage.

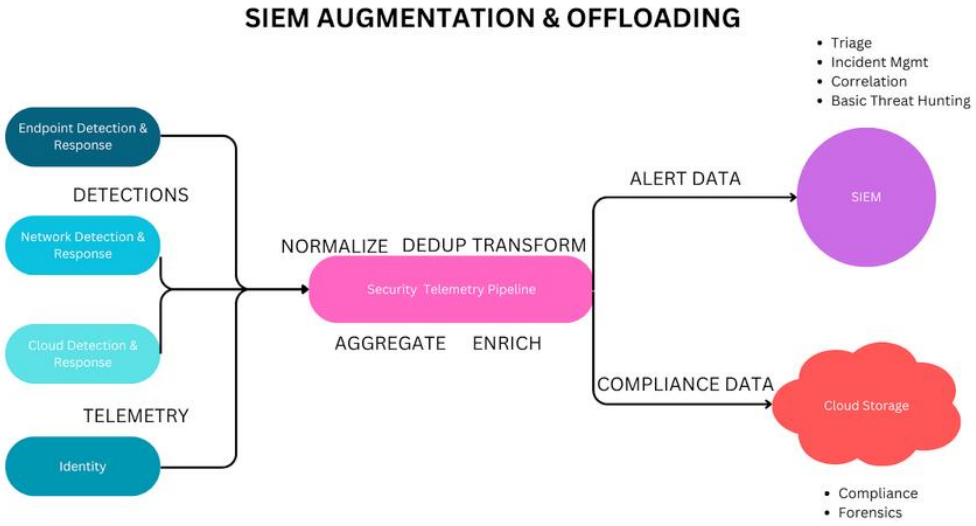

Bifurcated SIEM + Data Lake

Operating a SIEM + data lake, for example, using a SIEM for correlation, reporting, and incident management and a security data lake for broad-scope collection and hunting in long-term historical data.

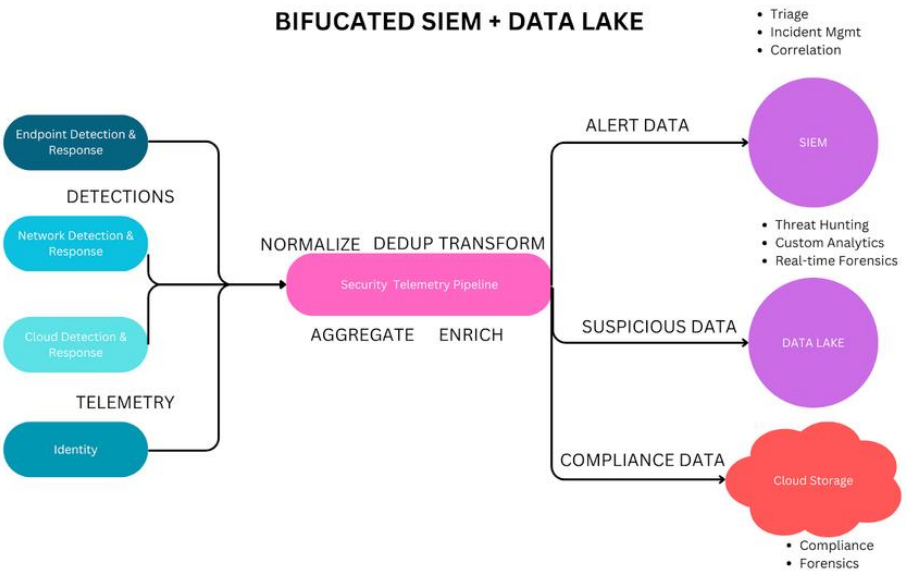

Standalone Security Data lake

Running a standalone security data lake, usually combined with additional SecOps tools, as SIEM replacement

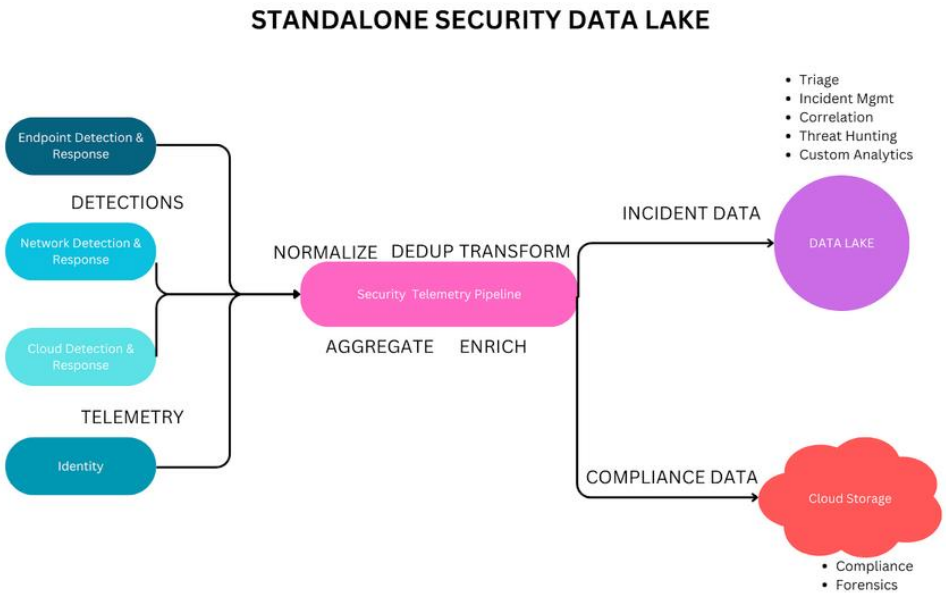

SOCless / Decoupled SIEM

Operating a SOCless/SIEMless Model, eliminating the need for a traditional Security Information and Event Management (SIEM) system or a dedicated Security Operations Center (SOC). This model leverages specialized tools and platforms to provide effective threat detection, visibility, and response capabilities

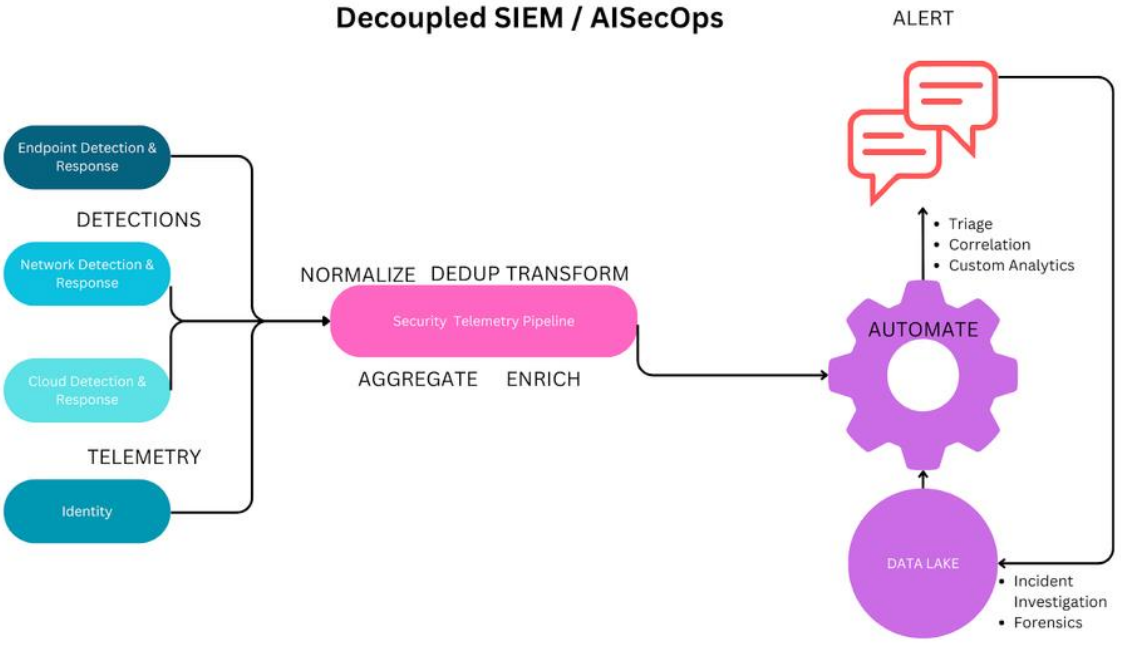

Main Benefits of Implementing Security Telemetry Pipelines

By automating tasks like data ingestion, enrichment, and alert prioritization, security telemetry pipelines reduce the workload for security teams, enabling them to focus on strategic activities such as threat hunting and incident response.

Additionally, these solutions are vendor-agnostic, enabling organizations to select the most effective tools for their specific requirements without being restricted to a single vendor's ecosystem. As security operations continue to evolve, the use of advanced telemetry pipelines will become indispensable for modernizing SIEMs, building federated XDRs, and implementing SOCless detection strategies.

1. **Reducing Costs:** Businesses initially saw telemetry pipelines as a way to cut costs. By removing redundant and deduplication repetitive data, data ingest, egress, and computation costs can be massively reduced. In many cases, this does not lead to lower costs but to a higher ROI as more high-quality data can be ingested.
2. **Increased Efficiency**: By automating data ingestion, normalization, and enrichment, security data pipelines significantly reduce the amount of manual work required by security teams. This allows analysts to focus on more strategic tasks like threat hunting and incident response.
3. **Real-Time Stream Processing**: Security data pipelines allow for real-time data enrichment, enabling faster and more accurate threat detection. Instead of waiting for data to be processed in batches, security teams can act on threats as they emerge.

4. **Scalability**: As security environments become more complex, the ability to scale data processing capabilities is critical. Pipelines ensure that organizations can handle increasing volumes of data without sacrificing performance.
5. **Vendor Agnosticism**: Security data pipelines enable integration across multiple platforms and vendors, ensuring that organizations can use the best tools for their needs without being locked into a single vendor's ecosystem.
6. **Better Threat Detection and Response**: By enriching data and applying advanced analytics, pipelines improve the accuracy and effectiveness of detection and security automation systems. This reduces false positives and ensures that critical threats are identified more quickly.

Cornerstone Capabilities of Security Telemetry Pipelines

STREAMLINED AND OPTIMIZED DATA INGESTION

Data ingestion involves the automated collection of security event data from a wide range of sources, such as servers, networks, cloud environments, applications, and agents. It ensures that security data is collected in real-time or batch processed to detect threats and manage incidents efficiently. Modern security telemetry pipelines support flexible, generalized data ingestion from various data sources, including custom application logs, with the ability to easily adapt to new data sources and formats. On-demand and recommender-guided data transformations enable seamless integration of data into any format required by different consumers, eliminating the need for manually written rules that require constant updating.

FRICTIONLESS DATA INTEGRATION

Data integration is the process of unifying disparate security data streams into a cohesive framework, enabling the correlation of data from different tools and systems for a holistic view of an organization's security posture.

Security telemetry pipelines can create a security data lake by transforming and storing data in highly compressible formats such as Parquet, stored in affordable cloud object storage like AWS S3, Azure Blob, or Google Cloud Storage. Natural language queries can be used to search the data, facilitating on-demand analysis and making data available to downstream analytics tools in the right format.

AUTOMATED DATA HYGIENE AND CURATION

Data hygiene and curation refers to the filtering, normalization, and transformation of security data to reduce noise and ensure only relevant information reaches security teams. With over 80% of security data often being redundant, duplicate, or noisy, data curation plays a crucial role in improving detection accuracy.

Security telemetry pipelines help reduce data overload by filtering out low-value information, such as duplicate data or unimportant firewall logs. Summarization capabilities condense normal, repetitive events into a single event, massively reducing the volume of data that analysts need to process. Additionally, events can be enriched with various contextual information, enhancing the quality of the insights provided.

CONTEXTUAL ENRICHMENT OF SECURITY DATA

Data enrichment involves enhancing raw security event data with additional context from third-party sources, helping security teams prioritize and respond more effectively to threats. Enrichment may include the addition of threat intelligence data, Geo-IP information, or other context that provides deeper insight into the nature of incidents.

Modern security telemetry pipelines use enrichment to provide additional context to security logs, leveraging third-party threat intelligence and detecting anomalies in the data stream. Anomaly detection further helps determine the sentiment of events, allowing security teams to focus on those events that most closely resemble threats. By enabling this enrichment in the data pipeline, SecOps teams can prioritize threats efficiently, reducing the time taken to respond to critical incidents.

Alternative Solutions and Technologies

Before the emergence of dedicated security telemetry pipelines, advanced and engineering-led security organizations repurposed a diverse toolset to integrate, transform, and route security data. These include:

- **Scripted Data Pipelines**: Many organizations developed their scripts using programming languages like Python or data query languages like SQL to automate the collection, transformation, and routing of security data. These scripts were often complex and required significant maintenance but provided highly customizable solutions tailored to specific needs.
- **Observability Tools**: Tools like Prometheus, Grafana, and Elastic are sometimes used for collecting, visualizing, and analyzing data. While not specifically designed for security use cases, these observability tools allowed organizations to gain visibility into various aspects of their infrastructure and provided a foundation for detecting anomalies. A parallel observability pipeline market has also sprung up.
- **Streaming Data Solutions**: Technologies like Apache Kafka, Apache Flink, and Google Pub/Sub are used to manage real-time data streams. These solutions enabled organizations to process large volumes of security data in near real-time, which was essential for timely threat detection and response.
- **SOAR (Security Orchestration, Automation, and Response)** platforms, such as Splunk Phantom, Demisto, and Swimlane, have been used to automate response actions and orchestrate workflows across different security tools, including alert triage.

PAYING THE REPURPOSING TAX

For most teams, these sorts of tools incur a “repurposing tax” in the form of needing extensive customization for security use cases or skills and expertise that are non-transferable to other security operations tasks.

Other Emerging Alternatives

- **AISecOps:** These tools leverage GenAI and LLMs to attempt to automate common incident response tasks, including alert triage and incident qualification. The benefit as opposed to SOAR is that they don’t require playbook maintenance. They do not scale well due to the high cost of running LLMs. Examples include Dropzone.ai, Radiant Security, Simbian, and Strikeready.
- **SIEM and Security Data Lake-native capabilities:** SIEM vendors are following suit and are adding basic telemetry pipeline capabilities as well; see, for example, Exabeam or Gurucul. These capabilities focus on optimizing data for the incumbent SIEM platform.

Market Size

Security telemetry pipelines are only now beginning to come out of stealth, but Cribl has reported revenue of $100M+.
An indirect proxy measure may be funding, which has exceeded $50M already in the seed stage.
Cybefuturists expect the market to have a potential TAM of $250M by 2028.
Unless security telemetry pipeline solutions can evolve up or down the security stack to become more SIEM or XDR-like or add adjacent capabilities, the TAM will always be a small percentage of the SIEM market.

Representative Market Players

Auguria

- **Founders**: Keith Palumbo (CEO) and Chris Coulter (CTO).
- **Founding Year**: 2022
- **Funding**: Auguria has raised $6.5M seed funding
- **Investors:** SYN Ventures and S Ventures
- **Headquarters**: Ladera Ranch, California, USA.
- **Website:** https://auguria.io

Abstract Security

- **Founders:** Colby DeRodeff (CEO), Ryan Clough (CPO), Aaron Shelmire (Chief Threat Research Officer), and Chris Camacho (COO).
- **Founding Year**: 2023
- **Funding:** Raised $8.5 million in a seed round, and $15M in Series A.
- **Investors**: Crosslink Capital, Rally Ventures, Liquid 2 Ventures.
- **Headquarters:** Austin, TX, USA
- **Website:** https://www.abstract.security

Databahn.ai

- **Founders**: Nanda Santhana (CEO), Praful Ilamkar (Head of Engineering), and Aditya Tirumalai Sundararam (Head of Product).
- **Founding Year:** 2023
- **Funding:** Undisclosed
- **Investors:** Undisclosed
- **HQ:** Dallas, TX, USA

- **Website:** https://databahn.ai

RunReveal

- **Founders:** Evan Johnson (CEO) and Alan Braithwaite (CTO), both previously holding key engineering and security roles at Cloudflare and Segment.
- **Founding Year:** 2023
- **Funding**: $2.5 million seed round (May 2024)
- **Investor:** Costanoa Ventures
- **Headquarters:** Austin, TX, USA
- **Website:** https://runreveal.com

Tarsal

- **Founders:** Sunny Rekhi (CEO) and Manny Gundampalli
- **Founding Year:** 2022
- **Funding:** $6 million seed funding round in March 2024
- **Investors:** The seed round was led by Harpoon Ventures and Mango Capital, with participation from, Y Combinator, Abstract Ventures, and Backend Capital.
- **Leadership team:** The company recently appointed Barrett Lyon as Chief Technology Officer (CTO)
- **HQ:** New York, NY, USA
- **Website**: https://tarsal.co

Honorable Mentions

Observo.ai

- **Founders:** Gurjeet Arora (CEO) and Ricky Arora (COO).
- **Founding Year**: 2022
- **Funding:** The company has raised funding from at least one institutional investor, GIT1K Ventures.
- **Headquarters**: Newark, CA, USA

- Website: https://www.observo.ai/

Onum

- **Founders:** Pedro Castillo (CEO), Lucas Varela, and Pedro Tortosa
- **Funding:** Raised $28 million in a Series A round in 2024
- **Investors**: Led by Dawn Capital, with participation from Kibo Ventures and Insight Partners
- **Headquarters:** Madrid, Spain
- **Website:** https://onum.com/

Prevalent AI

- **Founders**: Paul Stokes (CEO), Sir Iain Lobban, Andrew France OBE, and Arun Raj (COO).
- **Founding Year**: 2017
- **Funding:** Backed by ISTARI, which became a significant minority shareholder in 2021.
- **Investors:** ISTARI (a cybersecurity platform under Singapore's Temasek).
- **Headquarters**: London, UK, with additional offices in Cochin, India.
- **Website:** https://prevalent.ai

The Future of the Security Telemetry Pipeline Market

By 2030 at the latest, the market will have cycled. A few providers will have developed a broader portfolio of security operations capabilities and solutions and will grow into SIEM, XDR or Security Data Lake companies. Some will be acquired by SIEM and XDR providers seeking to integrate the capabilities themselves. There will always be a much smaller dedicated market for high-end and power-user organizations.

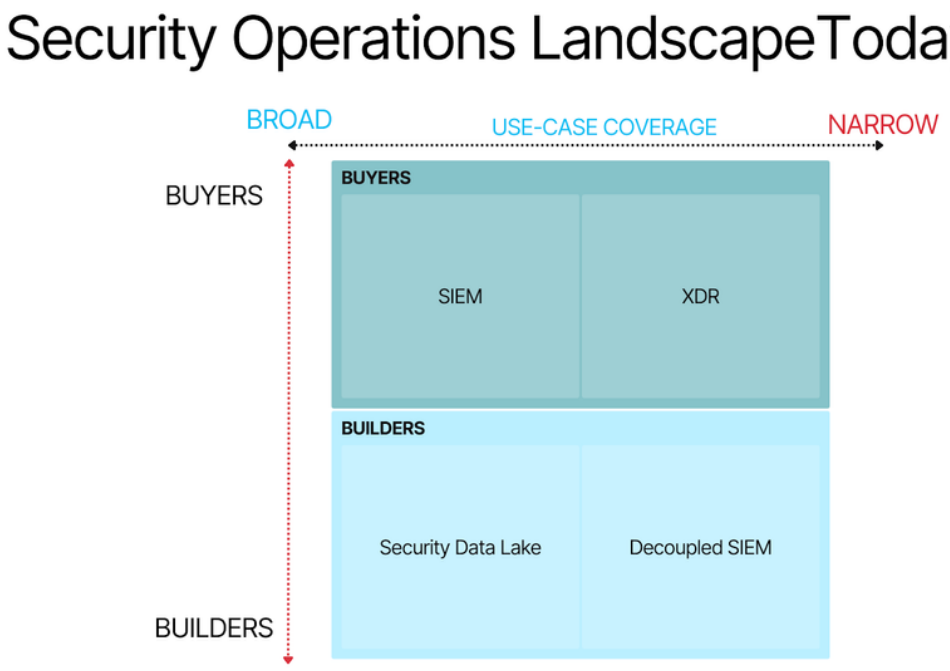

FIG 2: The Misleading SecOps Market View

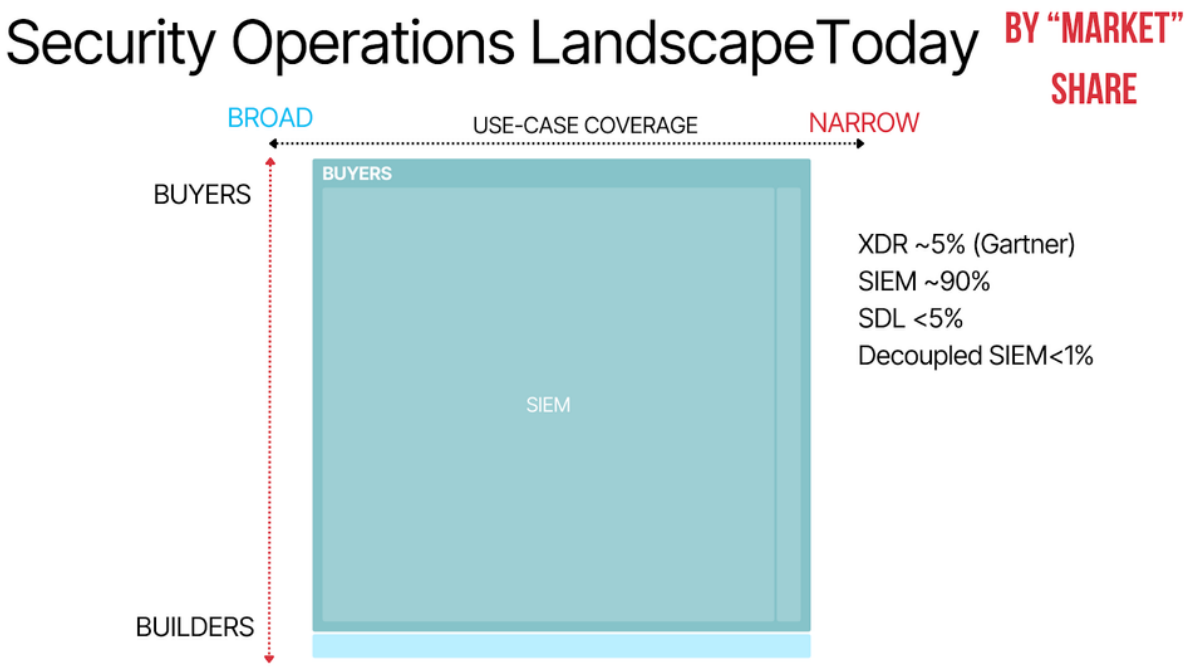

FIG 3: The Reality of the SecOps Market.

Final Words: The Future of Security Operations is Integrated

Security data pipelines are in the vanguard of a new wave of tools transforming how organizations manage security operations and process security data. By automating many of the manual tasks that burden security teams, security telemetry pipelines free up analysts to focus on more strategic activities. Integrating real-time data enrichment and advanced analytics ensures that security operations can scale effectively while staying agile in the face of evolving threats.

As security environments grow in complexity, adopting a security data pipeline model will be essential for organizations looking to modernize their operations, augment their security teams, and build flexible, scalable detection and response architectures. Whether modernizing a SIEM or building a federated XDR, security data pipelines are the key to achieving security mastery in the digital age.

ABOUT OUR SPONSOR

Auguria is a pioneering platform redefining Security Operations (SecOps) by integrating AI and human expertise to modernize SIEM and Security Data Lakes. Its Security Knowledge Layer™ eliminates 99% of irrelevant security data, focusing on the critical 1% that matters. This reduces costs, enhances efficiency, and empowers rapid threat detection and response, transforming modern cybersecurity.

ABOUT CYBERFUTURISTS

Cyberfuturists is a research and advisory firm led by Oliver Rochford specialising in cybersecurity innovation. Focused on emerging technologies like Security Telemetry Pipelines and AISecOps, the firm provides actionable insights to help organisations optimise security operations, reduce costs, and combat modern threats. By combining strategic foresight with deep technical expertise, the Cyberfuturists equipw businesses with the tools to build scalable, resilient, and future-ready security postures.