PIEM

Does Healthcare need a SIEM for Privacy?

by **Daniel Fabbri**, Assistant Professor, Vanderbilt University; Founder & CEO, Maize Analytics.
Ben Flatgard, Founder & Principal, Cycise; former Director for Cybersecurity Policy, National Security Council.

Security Information and Event Management (SIEM) systems are designed for cybersecurity, not patient privacy. Privacy Information and Event Management (PIEM) systems are an emerging class of privacy monitoring system geared for medical record protection.

SIEMs emerged almost two decades ago and have grown into the foundation of a modern security operation. But the SIEM market is evolving to incorporate new technologies such as predictive and behavioral analytics (i.e., User Behavior Analytics) or more specific use cases (i.e., Managed Detection and Response). There's some evidence that SIEM technology is often misused or underused by security professionals, many of whom profess a love-hate relationship with their SIEM.

The healthcare industry should be looking to learn from these lessons to develop a technology platform that serves as the basis of patient privacy and compliance. We propose this platform be called PIEM to reflect its relationship (and integration) with SIEMs.

Security Information and Event Management

The SIEM system was conceived to bring together individual security tools – firewalls, network sensors, access logs. A SIEM ingests disparate data sets—web logs, firewall sensors, remote login, threat intelligence from anti-virus vendors, etc.—consolidates alerting functions across those data sets in real-time, and enables analyses to catch inappropriate activity.

In doing so, SIEMs scale the reach of a security operation, enabling limited staff to conduct sophisticated analyses across huge data sets, monitor activity in real-time, and in some cases, intercede during active cyber-attacks to mitigate damage. SIEMs are designed to, if properly implemented, present the security professional with relevant information at a time when action can be taken.

The Healthcare Privacy and Compliance Challenge

Healthcare privacy officers face an enormous volume and variety of data when protecting patient privacy.

HCA reports over 27 million patient interactions each year. Each interaction results in multiple patient record accesses as different doctors, nurses, and administrators all support the appointment. Some of these accesses, unfortunately, are malicious (e.g., identity theft) or inappropriate (e.g., snooping).

To detect privacy violations and remain in compliance with laws and regulations, privacy officers often conduct manual analyses of a variety of different logs. But at this scale, it is impossible to manually review and analyze accesses with sufficient speed and accuracy.

Some in the privacy community have looked to their security counterparts to adapt SIEM tools to the challenges of protecting patient data. After all, EMR access logs are another set of data to monitor. However, there are stark differences between network monitoring and EMR access auditing. For example, EMR access audits require clinical context to investigate a potential issue, while SIEMs often analyze technical indicators in isolation. Moreover, SIEMs often assume a binary threat model in which an attacker is either malicious or not, while privacy violations may be committed by hospital employees who only occasionally behave badly. SIEMs also have limited functionality for facilitating privacy investigations and reporting. Incorporating a standard SIEM into privacy workflows limits privacy audit effectiveness.

Privacy Information and Event Management (PIEM)

SIEMs may not offer a privacy solution as they exist now, but the concept of SIEMs might be applied in a privacy-centric equivalent. A PIEM (Privacy Information and Event Management) system would serve as a platform to aggregate privacy and compliance data streams. All access logs would be available, accessible, and interoperable on a single platform. Privacy officers could conduct analyses across disparate data sets to correlate behavior, better explain record accesses, and uncover violations. Privacy investigations and reporting would become standardized and streamlined across the enterprise.

To demonstrate the differences between a PIEM and a SIEM, consider the threat environment, data structure, analytics, and user base of the two deployment environments.

In many SIEM deployments, the users monitored are exclusively bad or good (e.g., an outside attacker trying to circumvent a firewall). In contrast, in a PIEM environment, a doctor may legitimately access patient records most of the day, but snoop on a family member or friend in between their appropriate behavior. In healthcare's open access dynamic (i.e., where authenticated users can access any patient's data), the privacy threat environment is full of these inconsistencies.

PIEMs and SIEMs also integrate and analyze distinct data types. SIEMs work mainly at the network or transport layers, analyzing IP addresses or packet contents. PIEMs work at the session and application layers, analyzing the data that are accessed and what the user is doing with the data. Consistent with this scope, SIEMs often analyze unstructured text, with some structured information. PIEMs work primarily on structured logs that are correlated with structured contextual information – understanding the semantics, relationships and types of data are essential for privacy auditing.

Another major difference between PIEMs and SIEMs is the role of context in analyzing accesses.

SIEMs rely on regular expressions, Boolean rules, or statistical anomalies to detect threats. These approaches—that analyze access logs in isolation of clinical context—can detect large scale data scraping (i.e. a user normally accesses 100 records, but today accesses 10,000), but are not well suited to detect small-scale violations.

In contrast, PIEMs attempt to understand the context of an access to ascertain its risk. Much like a privacy officer's current workflow, if the system can understand "why" an access occurred (e.g., because of an appointment, an oncologist treating a cancer patient, etc.), then the access is likely appropriate. To perform these context-based analyses, PIEMs must ingest additional supporting data sets like ICD-10 codes, appointment records, and lab orders to infer the reason for record accesses.

The user base of each tool is also significantly different. Technical security officers operate SIEMs. In contrast, privacy officers, most of whom do not have formal technical backgrounds, will use a PIEM system. The privacy and compliance team, composed of risk managers, lawyers, and human resource staff, require tools that suit their skill set and workflow.

PIEM

Comparing PIEMs to User Behavior Analytics (UBA)

Recently, user behavior analytics systems have gained popularity in the cybersecurity community and are being integrated into SIEMs. UBA systems attempt to model normal access patterns based on a user's historical activity, workflow, and some access context. Access patterns that deviate from the normal activity are escalated for review.

Activity models are a welcome addition to the security professional's toolkit, but need to be refined to detect privacy violations. Because a UBA is looking to model typical user behavior, it is susceptible to being manipulated by users within the organization. For example, when a user accesses a record, he or she controls the specific order of accesses, the fields clicked, and access timing. Because the user has knowledge of the EMR system and its normal workflows, users can circumvent detection by clicking more often or more deliberately.

To prevent such manipulation, PIEMs rely on the context surrounding an access, not the access itself. PIEMs base their analyses on context that users cannot easily manipulate without detection such as appointment information, diagnosis codes or human resource information. Even further, PIEMs enhance the context by filling in missing facts to make sense of the data such as the fact that oncologists treat cancer patients or nephrologists treat dialysis patients (it is important to note that privacy officers must be kept in the loop to approve these facts).

PIEMs are purpose-built for the hospital privacy and compliance use case, while integrating the beneficial SIEM capabilities.

Future of PIEMs

Healthcare organizations will continue to be prime targets of cyberattacks from external and internal actors. To adequately defend patient data, security and privacy monitoring systems must be deployed to detect a wide range of attacks. SIEMs defend against some types of attacks, but are not designed for the specific challenges of protecting patient data. PIEMs allow for additional monitoring at the application layer, cross-referencing clinical context, and allowing non-technical privacy professionals to investigate accesses to protected health information.

SIEMs for Healthcare		PIEMs for Healthcare
Security Professionals	Designed for...	Privacy Officers
Cybersecurity	Use case	Patient privacy
Network Logs, Server Logs, etc.	Type of monitoring	EMR Access Logs
External Attackers, Internal Compromised Systems	Threat	Internal Users' Inappropriate & Malicious Behavior
Regular expressions, Boolean rules, statistical anomalies	Approach	Context-based analysis (using ICD-10 codes, appointment records, labs, HR info, etc.): Was the user involved with treatment, payment or operations?