

Security and Reliability Requirements for Advanced Security Event Management

Roland Rieke¹, Luigi Coppolino², Andrew Hutchison³,
Elsa Prieto⁴, and Chrystel Gaber⁵

¹ Fraunhofer Institute SIT, Darmstadt, Germany

² Epsilon S.r.l., Naples, Italy

³ T-Systems, South Africa

⁴ Atos Research & Innovation

⁵ Orange Labs - France Telecom

Abstract. This paper addresses security information management in complex application scenarios. Security Information and Event Management (SIEM) systems collect and examine security related events, with the goal of providing a unified view of the monitored systems' security status. While various SIEMs are in production, there is scope to extend the capability and resilience of these systems. The use of SIEM technology in four disparate scenario areas is used in this paper as a catalyst for the development and articulation of Security and Reliability requirements for advanced security event management. The scenarios relate to infrastructure management for a large real-time sporting event, a mobile money payment system, a managed services environment and a cyber-physical dam control system. The diversity of the scenarios enables elaboration of a comprehensive set of Security and Reliability requirements which can be used in the development of future SIEM systems.

Keywords: security requirements, security information and event management, SIEM, architecting trustworthy systems.

1 Introduction

Security information and event management (SIEM) systems provide important security services. They collect and analyse data from different sources, such as sensors, firewalls, routers or servers, and provide decision support based on anticipated impact analysis. This enables timeous response to (or prevention of) attacks as well as impact mitigation by adaptive configuration of countermeasures. However, there are also a number of constraints for current commercial solutions. These constraints include the inability of systems to consider events from a multiple organisations or the ability to provide high degree of trustworthiness or resilience in the event collection environment.

The project MASSIF [3], a large-scale integrating project co-funded by the European Commission, addresses these challenges with respect to four industrial domains: (i) the management of the Olympic Games information technology (IT) infrastructure [12]; (ii) a mobile phone based money transfer service, facing high-level threats such as money laundering; (iii) managed IT outsource services for large distributed enterprises; and (iv) an IT system supporting a critical infrastructure (dam) [4].

In undertaking the development of next-generation SIEM concepts and constructs, it became clear that the Security and Reliability of the SIEM itself are critical to the successful deployment of SIEM in a particular environment. With this in mind, we set about analysing each of the mentioned scenarios in some detail, to create an explicit list of Security and Reliability requirements. The intention is that these requirements can be used to guide and assess SIEM development, and ensure that these important attributes are incorporated.

2 Large Scale Scenarios in Four Industrial Domains

In this section, four deployment scenarios for SIEM technology are introduced. The elements of the scenario which can benefit from further SIEM development are also outlined in each case. From the introduction of the scenarios and their unique characteristics, a set of consolidated requirements for a next-generation SIEM can be compiled.

2.1 Scenario 1: SIEM Technologies Used in the Olympic Games

The Olympic Games is one of the largest and most high profile sporting events that takes place, and there is a large technical infrastructure to support many aspects of the games both asynchronously and in real-time. SIEM infrastructure is used with the Olympic Games systems, to protect the games IT infrastructure from any undesired and/or uncontrolled phenomena which could impact any part of the result chain and associated services. The nature of this kind of event presents a big challenge to SIEM infrastructures, for example the next London 2012 games cater for 79 days of competition, 26 sports, 94 venues, 17,000 athletes, 20,000 journalists, 70,000 volunteers, 4,000 IT team members, 900 servers, 1,000 network and security devices and more than 10,000 computers deployed. One of the new challenges will be the amount of data generated from the results systems, representing 30% more than in the Beijing Olympics in order to provide real-time information to fans, commentators and broadcasters world wide. The intensity and complexity of this kind of sporting event presents a big challenge to SIEM infrastructure, mainly, due to two very characteristic features: the number of security event types (about 20,000), and the volume of generated events to be handled (around 11,000,000 alerts per day). However, the most critical aspect that a SIEM system faces in the Olympic Games is that those security events must be processed and reacted upon in real-time.

Advanced SIEM System Contribution. The Olympic Games scenario is valuable to demonstrate the enhancement on scalability, processing enormous amounts of generated data events in real-time. Furthermore the scenario can contribute to validate the cross-layer correlation of events (service, application, infrastructure) from multiple sources.

2.2 Scenario 2: Mobile Money Transfer Service

Use of Mobile Phones to effect payment is a widely used service, particularly in developing markets where banking systems may not be as dense or available as in developed

countries. Characteristics and challenges of authentication, confidentiality, integrity and mobility all have to be considered in this scenario.

From a SIEM perspective, mobile money transfer is an interesting and challenging scenario for the unique attributes that the scenario presents. Indeed, this scenario is quite complex because it requires to analyze past and present data and to extract information from raw events. It is also very sensitive to the performance of detection as the rate of false positives and true negatives should be optimised. Finally, all this should be done while keeping the service scalable and secure. The service allows end users to convert cash to “electronic money” (and vice versa) at merchants, who act as distributors and act as channel users. The electronic money can be used to pay purchases at the merchants’ or for bills such as electricity. Furthermore the electronic money can also be transferred between the end users. End users access the service with their mobile phones and distributors can access the service either via mobile phone or directly on the Internet. Both means of access are handled by front-end servers that then access the back-end servers containing the transactions etc.

Advanced SIEM System Contribution. Like any other money transfer service, the service is exposed to the risk of money laundering and other types of fraud. The money laundering risk implies misuse through disguising illegally obtained funds to make them seem legal, and more generally the fraud risk implies any intentional deception made for financial gain. In addition, any money transfer service that has part of its infrastructure exposed via the Internet and/or the end user can access the service using electronic means (a mobile device such as a phone or a pad in this use case), has an increased exposure to fraud, via both attacks against the service infrastructure itself and the abuse of normal service functionality. The objective of including this scenario is to achieve greater protection and transactional integration of SIEM protection through next generation SIEM services. The ultimate intention is to protect the money transfer service against fraud both by detection and application of relevant counter-measures.

2.3 Scenario 3: Managed Enterprise Service Infrastructures

The use of *managed services* by businesses is an increasingly used model, whereby elements of IT and infrastructure are “outsourced” to specialist service providers. In some instances, services are provided by an outsourcer via shared platforms, giving customers economies of scale. In other instances, managed services are performed by a provider on the infrastructure belonging to a customer. Mixed approaches are also possible, and an extrapolation of this can be viewed as occurring when such services are provided in a “cloud based” mode. Provision of Security Information and Event Management services for customers is a valuable complement to the management which an outsourcer or service provider can deliver. The purpose of including a managed enterprise service infrastructure scenario was to consider just such cases: where the services of large enterprises are managed, and a SIEM service is used to collect, inspect and react to large scale security events from member systems and devices.

Advanced SIEM System Contribution. There are a number of limitations of SIEM systems, encountered by managed security service providers, that are not adequately addressed by current SIEM solutions. For this reason, such a SIEM deployment is

interesting to consider when looking at next-generation SIEM requirements. Some of the issues that can be identified in particular are: (i) insufficient resilience of the SIEM infrastructure itself to withstand large scale attacks; (ii) inadequate trustworthiness of source data within the SIEM; and (iii) inadequate disaster recovery capabilities of SIEM systems. Solutions to the limitations that current SIEM systems present will improve the resilience and business continuity capabilities of large companies, through enabling managed service providers to detect and address security events more proactively. It is considered that work on next-generation SIEM systems could address some of the identified problems through the following focus areas:

1. Providing guidelines on the minimum requirements for event data to enable successful event correlation.
2. Providing guidelines on the impact of the unavailability of certain event data on successful event correlation and management.
3. Guaranteeing the trustworthiness of event sources.
4. Improving correlation modelling for better analysis of complex environments (and for better automated correlation processing in complex environments).
5. Improving the resilience and business continuity capabilities for large enterprises.

2.4 Scenario 4: Critical Infrastructure Process Control (Dam)

The features of dam infrastructures are strictly related to the aims they are conceived for. Dams are mostly used for water supply, hydroelectric power generation, irrigation, water activities and wildlife habitat granting. Dams represent fundamental assets for the economy and the safety of a country, such as they are counted among critical infrastructures. So, monitoring of a dam is essential since an accident would have dramatic consequences. The amount of parameters to be monitored to assess the safety of a dam and foresee possible failures or anomalies is enormous, and this huge data flow must be analyzed under real time constraints. Each of these parameters is measured using different sensors, such as inclinometers and tiltmeters, crackmeters, jointmeters, earth pressure cells, turbidimeters, and thermometers. In addition to the above mentioned parameters measured by the sensors, other components are necessary for the full control of dam. Some essential elements are: data collectors, human machine interaction interfaces, data storing units, command and data gateways and signal buses. In other cases there is the need also to integrate different subsystems existing.

Advanced SIEM System Contribution. The current SIEM solutions hardly facilitate the introduction of new technologies to improve the efficiency of the security event detection. At the same time, they usually lack in the capability to support heterogeneous systems and technologies. Introducing SIEMs to jointly manage all different aspects related to the security in the monitoring of a dam can be a very powerful mechanism to increase the overall security of such critical infrastructures. However, currently available SIEMs solutions are focused on the management of digital and information security related events and are designed specifically for this type of applications. This may make complex or even impossible the development of applications targeting security of critical infrastructures in a wider sense. For instance, creating an application capable of correlating network and host events that may indicate a cyber-attack with suspicious

activities detected by the dam surveillance system may greatly improve the security of the whole monitoring process but may introduce some implementation difficulties. SIEMs are not designed to deal with this kind of scenarios and so, encompassing security events coming from different application domains within the same application may be troublesome. In particular, the current technologies usually neglect the possibility to correlate physical and logical events, which can improve the effectiveness of the detection process.

In order to secure the dam control system, today recognized as a critical infrastructure and hence of public interest, regulations must be considered. Indeed, any activity of the dam operators strictly follows well-known rigid procedures. For example, the opening of a gate without alerting the control center is not admitted. Unfortunately, the current SIEM technologies insufficiently exploit regularities characterizing dam systems. In particular, procedures could be encoded in patterns and they could be exploited for detecting anomalies in control system. All this information could contribute to make the security system aware of the context in order to correctly interpret the meaning of some evidences. Introducing such features in a SIEM solution moves the focus of the analysis from a system level view to the business process model of the system.

3 Consolidated Guidelines for Next Generation SIEM

Based on the four scenarios described, and the diverse set of circumstances that they cover between them, a set of consolidated recommendations, to guide the design and development of next generation SIEM platforms, is identified and grouped in five topics.

3.1 Guidelines Concerning Advanced Security Services

Besides issues like dependability, redundancy and fault tolerance, analysis of the four scenarios considered reveals the need for enhanced *security-related* features of future SIEM platforms. These features go beyond what is currently supported by existing solutions. Overall a lack of capability to model incidents at an abstract level is perceived. From the scenarios investigated, and the current SIEM limitations observed, the following guidelines have been identified for next-generation SIEMs with respect to security:

Correlation Across Layers of Security Events. Advanced SIEM systems needs to support enhanced correlation across layers, from network and security devices as well as from the service infrastructure such as correlation of physical and logical event sources. This is due to the variety of systems issuing inputs that can give insights to security only when combined. An example is the off-site monitoring and the on-site management of the dam's configuration.

Multi-level Security Event Modelling. Multi-level security event modelling will enable provision of more holistic solutions to protect the respective infrastructures. The Olympic Games Scenario stipulates that it would be of interest to understand the effects of technical events on the user or process level of the system.

Analysis of Malicious Behaviour Using Attack Graphs. Many of the security issues mentioned in this document originate from complex malicious actions or patterns of actions (e.g., the laundering of money in the mobile money transfer scenario or the misuse case of *Low and Slow* attacks in the Olympic Games infrastructure).

Predictive Security Monitoring. Predictive security monitoring allows to counter negative future actions, proactively. There is a crucial demand for early warning capabilities. Moreover, the limitations with regards to the Managed Enterprise Service point to the fact that dealing with unknown or unpredictable behaviour patterns is not sufficient in current SIEM solutions.

Modeling of the Events and Their Relation to Other, Possibly External, Knowledge.

A basic precondition of prediction and simulation as well as of attack analysis is the proper representation of the security requirements and any relevant information about the system as well as any knowledge about the actual and possible behaviour. When reasoning under incomplete information it is not only decisive to properly gather and describe the information available, but it is also required to develop novel methods based on discernibility, probability or plausibility in order to reason about uncertainty.

Securing the Evidence Progressed by the SIEM Components. The misuse case of a sensor compromise, showing that it is vital to be able to trust the information that is received, when using events from sensors like those deployed to monitor the dam or other critical infrastructures.

3.2 Guidelines Concerning Event Processing

Similar to the limitations noted for security, recommendations for *event processing* are also made, based on limitations in current SIEM implementations. The guidelines for a next generation event correlation engine are as follows:

Real-Time. The system must process input data at a high rate and provide meaningful results with soft real-time requirements.

Scalability and Elasticity. The engine should be capable of handling high input rate and should optimize the quantity of resources required based on the actual load. In other words, the system should monitor both input loads and vital parameters, such as CPU utilization, in order to adjust the amount of resources, i.e., provision more resources during peak load times and decommission them during valley load periods.

Handling Streaming and Stored Data. The engine should allow processing and correlation both of streams of events and stored relations (i.e., information stored in a database).

Multiple-Sources. The engine should be able to aggregate, abstract and correlate heterogeneous events from multiple sources at different levels of the system stack.

Pre-defined Correlation Rules and Rule Augmentation Capability. The engine should be shipped with a set of predefined correlation rules to identify well-known attacks. However, it should also support easy and intuitive creation of user-defined rules.

3.3 Guidelines Concerning Advanced SIEM Trustworthiness

Trustworthiness is the ability to provide a service in a way it is expected in terms of safety, security, reliability, availability, and timeliness. The analysis of the input scenarios has resulted in the following guidelines, to improve the general resilience and trustworthiness aspects of a next generation SIEM:

Resilience of the Infrastructure. The infrastructure should be highly resilient under attack, concurrent component failures, and unpredictable network operation conditions.

Security of Event Flows. The event flows should be protected, from the collection points through their distribution, processing and archival.

Protection of the Nodes. The designed mechanisms should offer flexible and incremental solutions for node resilience, providing for seamless deployment of necessary functions and protocols. These mechanisms should take into consideration particular aspects of the infrastructure, such as edge-side and core-side node implementations.

Timeliness of the Infrastructure. The infrastructure should provide for (near) real-time collection, transmission and processing of events, and ensure the corresponding reliable and timeliness generation of alarms and countermeasures when needed. Similarly, features for forensic support should adhere to the following guidelines:

Data Authenticity. Security event data contents, as well as additional/added information related to data origin and destination, must be the reliably stored.

Fault and Intrusion-Tolerant Stable Storage. The stable storage system on which data for forensic use will persist must be tolerant both to faults and to intrusions.

Least Persistence Principle. With respect to sensitive data, only information which is actually needed should be retained to stable storage (much of the data could be processed in real-time and potentially discarded).

Privacy of Forensic Records. Forensic evidence related to security breaches should be made available only to authorized parties.

3.4 Guidelines Concerning Compiler Technologies

In terms of *data acquisition* functionality, it has been noted that next generation SIEM systems should exhibit efficient implementation and/or support for various Features relating to data collection and parsing. Specific guidelines are as follows:

Heterogeneity Support. The data acquisition element must have the ability to deal with a large number of highly heterogeneous data feeds.

High Degree of Adaptability. Seamless integration of new types of security tools/probes should be possible, to improve the capabilities of the SIEM on an ongoing basis.

Peak Handling. The volume of events, to be collected and processed per unit of time, can occasionally increase, resulting in load peaks. The data collection layer should be able to handle such peaks and propagate relevant events to the SIEM core platform without loss of information.

High Degree of Expressiveness. The parsing logic, and related Languages, must allow effective processing of virtually any type of security relevant event.

Support for Fast and Reliable Development. Simple Development and configuration techniques and tools must be available. These will make it possible to implement, deploy, and integrate new parsers and collectors in a relatively short time and at a relatively low cost.

Generality and Platform Independence. The parsing/processing logic (and code) should as far as possible be decoupled from the specific characteristics of the data format and related technologies.

Distributed Processing. Whenever possible (and feasible), the data collection and parsing layer should implement parsing, filtering, and correlation functions at the edges and/or at intermediate nodes, i.e. nodes located along the path to the core SIEM correlation engine.

3.5 Guidelines Concerning Legal Aspects

In terms of *legal* considerations, SIEM systems themselves need to be viewed as data processing entities with consideration being given to issues like data retention, data privacy and so on. From the scenarios considered, the following guidelines in terms of legal aspects have been identified:

Data Retention. Data must be retained for a period of time not more than that necessary to the activities for which they were collected. If the data are required for detection and suppression of crime they can be stored for a longer period of time.

Cross-Border Data Transmission. It must be possible to limit the transmission of data outside of certain borders. It should be possible to process data within such a border. If personal data must be transferred to another country, it must be ensured that the level of data protection in the country of destination is adequate.

Minimum and Appropriate Security Measures. Considering state of the art technology, a minimum (but sufficient) set of measures must be taken to preserve integrity, confidentiality, and availability of personal data. More sensitive data require increased security measures.

Data Minimization and Anonymization. Only data strictly needed for security guarantee must be kept, while unnecessary details must be deleted or made anonymous.

4 Related Work

The development of new security relevant systems requires the integration of a security engineering process in the earliest stages of the development life-cycle. This is specifically important in the development of systems, where security is the enabling technology, as in advanced SIEM systems. There are several common approaches to security requirements engineering that may be taken. An overview of such processes is given in [5] and also in [9]. A comprehensive concept for an overall security requirements engineering process is described in detail in [8]. The authors propose a 9 step approach called SQUARE (Security Quality Engineering Methodology). A similar approach based on the integration of Common Criteria (ISO/IEC 15408) called SREP (Security Requirements Engineering Process) is described in [10]. In [6], different kinds of security requirements are identified and informal guidelines are listed that have proven useful when eliciting concrete security requirements. The author emphasises that there has to be a clear distinction between security requirements and security mechanisms. In [7], Hatebur et al. describe a security engineering process based on security problem

frames and concretised security problem frames. The two kinds of frames constitute patterns for analysing security problems and associated solution approaches. [7] specifically addresses accountability by logging.

Though all of the above mentioned approaches may lead to a sufficient level of security for the designed architecture, there is no obvious means by which they can be compared regarding the security requirements that they fulfil. In this paper, we address the first step in every security engineering process, namely the identification of artifacts, such as functional descriptions, dependencies and information flows, the identification of use cases and misuse cases, and stakeholders' information on assets, safety and security requirements. Additionally, we consider state-of-the-art information on existing SIEM systems and challenges identified by other work such as the following:

Security information and event management technology provides log management and compliance reporting as well as real-time monitoring and incident management for security events from networks, systems, and applications. A concise overview of current SIEM systems functionalities is presented in [11]. In [1], current threats are identified and advanced monitoring techniques such as file integrity monitoring, database activity monitoring, application monitoring, identity monitoring, and user activity monitoring are discussed. In [2], some challenges with respect to collecting and analyzing a multi-gigabit network stream are outlined. SIEM systems manage security events but are not primarily concerned with the trustworthiness of the event sources. Compared to traditional IT systems, securing SCADA systems (e.g., in the dam scenario) poses unique challenges. In order to understand these challenges and potential dangers, [13] provides a taxonomy of possible cyber attacks – including cyber-induced cyber-physical attacks on SCADA systems.

5 Conclusion and Future Work

This paper has described requirements in terms of security and reliability for advanced security information and event management. The approach used to identify requirements is scenario-driven: scenarios relating to a real-time, high profile sporting event infrastructure; a mobile payment system; an enterprise service provider deployment and a cyber-physical environment has been used as catalyst for requirements identification and elaboration.

Based on the key elements and attributes of each scenario, guidelines for security, event processing, trustworthiness, and compiler technologies in next-generation SIEM systems have been elaborated. To consolidate the approach, a conceptual model showing the progression from business process / application / infrastructure to elements of SIEM design and implementation has been introduced. It is considered to be quite unique and beneficial to have such a comprehensive and rigorous set of scenarios to draw upon, and studying and analysing the scenarios presented provides a sound foundation from which to make recommendations for next-generation SIEM systems.

We cannot necessarily claim that the set of recommendations is “complete”, but by developing (and ultimately testing) the proposed items against such a diverse set of scenarios, there is a high probability of addressing a wide range of SIEM requirements. The benefit of multiple scenarios is that associated characteristics which include diverse

requirements including mobility, scalability, real-time processing, potentially hostile device environments and so on. In this light, the security and reliability requirements are considered to be applicable to a wide range of advanced security event management contexts.

Acknowledgements. The authors developed this work in the context of the project MASSIF (ID 257475) being co-funded by the European Commission within FP7.

References

1. Monitoring up the Stack: Adding Value to SIEM. White paper, Securosis L.L.C., Phoenix, AZ (November 2010), <https://securosis.com/research/publication/monitoring-up-the-stack-adding-value-to-siem>
2. Applied Network Security Analysis: Moving from Data to Information. White paper, Securosis L.L.C., Phoenix, AZ (December 2011), <https://securosis.com/research/publication/applied-network-security-analysis-moving-from-data-to-information>
3. Project MASSIF website (2012), <http://www.massif-project.eu/>
4. Coppolino, L., D'Antonio, S., Formicola, V., Romano, L.: Integration of a System for Critical Infrastructure Protection with the OSSIM SIEM Platform: A dam case study. In: Flammini, F., Bologna, S., Vittorini, V. (eds.) SAFECOMP 2011. LNCS, vol. 6894, pp. 199–212. Springer, Heidelberg (2011)
5. Fabian, B., Gürses, S., Heisel, M., Santen, T., Schmidt, H.: A comparison of security requirements engineering methods. *Requirements Engineering* 15(1), 7–40 (2010)
6. Firesmith, D.: Engineering security requirements. *Journal of Object Technology* 2(1), 53–68 (2003)
7. Hatebur, D., Heisel, M., Schmidt, H.: Analysis and component-based realization of security requirements. In: *Proceedings of the International Conference on Availability, Reliability and Security (AREs)*, pp. 195–203. IEEE Computer Society Press (2008), <http://www.ieee.org/>
8. Mead, N.R., Hough, E.D.: Security requirements engineering for software systems: Case studies in support of software engineering education. In: *CSEET 2006: Proceedings of the 19th Conference on Software Engineering Education & Training*, pp. 149–158. IEEE Computer Society Press, Washington (2006)
9. Mellado, D., Blanco, C., Sánchez, L.E., Fernández-Medina, E.: A systematic review of security requirements engineering. *Computer Standards & Interfaces* 32(4), 153–165 (2010)
10. Mellado, D., Fernández-Medina, E., Piattini, M.: A common criteria based security requirements engineering process for the development of secure information systems. *Comput. Stand. Interfaces* 29(2), 244–253 (2007)
11. Nicolett, M., Kavanagh, K.M.: Magic Quadrant for Security Information and Event Management. Gartner Research (May 2010)
12. Prieto, E., Diaz, R., Romano, L., Rieke, R., Achemlal, M.: MASSIF: A promising solution to enhance olympic games IT security. In: *International Conference on Global Security, Safety and Sustainability (ICGS3 2011)* (2011)
13. Zhu, B., Joseph, A., Sastry, S.: Taxonomy of Cyber Attacks on SCADA Systems. In: *Proceedings of CPSCom 2011: The 4th IEEE International Conference on Cyber, Physical and Social Computing*, Dalian, China (2011)