MASSIF: A Promising Solution to Enhance Olympic Games IT Security

Elsa Prieto[1], Rodrigo Diaz[1], Luigi Romano[2], Roland Rieke[3], and Mohammed Achemlal[4]

[1] Atos Research and Innovation (ARI), Atos Origin
{elsa.prieto,rodrigo.diaz}@atosresearch.eu
[2] Consorzio Interuniversitario Nazionale per l'Informatica (CINI)
{lrom}@uniparthenope.it
[3] Fraunhofer - Institut für Sichere (SIT)
{roland.rieke}@sit.fraunhofer.de
[4] Orange - France Telecom SA
{mohammed.achemlal}@orange-ftgroup.com

Abstract. Nowadays, Olympic Games have become one of the most profitable global media events, becoming at the same way more and more attractive target from the terrorist perspective due to their media diffusion and international dimension. Critical for the success of such a highly visible event is protecting and securing the business and the supporting cyber-infrastructure enabling it. In this context, the MASSIF project aims to provide a new generation SIEM framework for service infrastructures supporting intelligent, scalable, and multi-level/multi-domain security event processing and predictive security monitoring.

Keywords: Information Security Management, Security Event Management, Systems Safety, Data Security, Software Protection, Secure Architecture Design.

1 Introduction

Recent terrorist attacks across the world indicated that terrorists continue to target crowded places and show how vulnerable high profile venues and events can be used to perpetrate such incidents for maximum impact across the globe.

Terrorism attacks can adopt many forms, not just a physical attack on life and limb. It can include interference with vital information or communication systems, causing disruption and economic damage. For this reason, in addition to the physical security of the event venues, the cyber-security of the IT event infrastructure should be protected in the same way.

Nowadays, Olympic Games have become one of the most profitable global media events. From the terrorist perspective, Olympics can be seen as one of the most attractive events to commit actions due to their media diffusion, international dimension and symbolic representation. As a consequence of this, security has become a top focus and budget priority. Surpassing all before it in size and scope,

H. Jahankhani et al. (Eds.): ICGS3/e-Democracy 2011, LNICST 99, pp. 139–147, 2012.

security at the Vancouver Olympic Games of 2010 cost an estimated USD $1 billion and included a 15,000-person force of Canadian military, Vancouver police, U.S. security forces, and private contractors to guard the city by air, land, and sea[1]. Vancouver marked a transition into an unparalleled era of Olympic security in terms of cross-national cooperation, planning, and spending. The scope, however, was limited—the majority of funds and efforts aimed to maintain calm during the two-week event and did not address longer-term security concerns.

As the Worldwide IT Partner for the Olympic Games and Top sponsor, Atos Origin integrates, manages and secures the vast IT system that relays results, events and athlete information to spectators and media around the world. The Atos Origin contract with the International Olympic Committee (IOC) is the world's largest sports related IT contract and was recently extended to cover the Sochi Olympic Winter Games in 2014 in Russia and the Rio Olympic Summer Games in 2016 in Brazil. The major challenge is to create an IT solution for each Olympic Games that allows the capture and reporting of every moment of the action and brings it to the world via television and the Internet. Critical for the success of such a highly visible event is protecting and securing the business and the supporting cyber-infrastructure enabling it and naturally, security is a top priority.

With innovation as the cornerstone of its business and strategy, Atos Origin is coordinating the European research project MASSIF (MAnagement of Security information and events in Service Infrastructures, http://www.massif-project.eu/). The MASSIF Consortium consists of 12 project partners from 6 different European countries (France, Germany, Italy, Portugal, Russia, Spain) and South Africa including three different groups of business roles: scenario providers (Atos Origin, Epsilon, France Telecom and T-Systems), scientific partners (Fraunhofer SIT, Institut Telecom, SPIIRAS, C.I.N.I., Universidad Politécnica de Madrid and Universidade de Lisboa) and SIEM developers (Alienvault and 6cure). This paper addresses the challenges that arise in the cyber-security of Olympic Games and how the results of the MASSIF project can help to improve it.

The paper is structured as follows: Section 2 provides an overview of the existing IT infrastructure while Section 3 includes the challenges addressed by the MASSIF project. Section 4 gives an overview of the related work. Finally, Section 5 concludes this paper and provides pointers for future work.

2 Olympic Games IT Infrastructure

Olympic Games are getting more and more huge events, numbers in this context are gigantic. For instance, in the Beijing games 10.708 athletes were competing, 5.600 written press & photographers were accredited, 12.000 rights holding broadcaster staff, 70.000 volunteers, more than 60 competition and non-competition venues (http://en.beijing2008.cn/media/usefulinfo/).

The Olympic Games, must successfully issue and activate more than 200,000 accreditations for Games that comprise around 300 events representing over 4,500 hours of live competition. Live commentator services are delivered for around 20

sports. More than 15 million information pages are viewed, with peaks of 1 million pages viewed on specific days. Over 3Gb of live results are provided in around 800,000 messages to the Olympic website, broadcasters and sports federations.

The complex, massive IT infrastructure of the Olympic Games is deployed by large teams of people into different environments every other year. Such a major task could potentially pose significant risks, but these can be offset through preserving and sharing the knowledge gained from previous Games.

The Olympic Games have 3 core systems that support the operations of the Games. These systems are summarised below:

Core Games System (CGS). CGS is a set of applications for assisting in the capture and management of data about people who will be attending the Games events and the staff supporting them. Among others, this includes Accreditation and Workforce management (including Volunteers).

Information Diffusion System (INFO). INFO comprises of a set of applications that retrieve and distribute information related to, and supporting, of the Games. The information is provided by different sources e.g. Results system, interfaces with CGS, Weather provider etc. The information is processed and distributed to internal clients e.g. broadcasters, journalists, and other members of the Olympic and Paralympics Families. IT is also sent to external clients e.g. World News Press Agencies (WNPA), sports federations and governing bodies, and Internet Service Providers (ISPs).

Results Systems, are grouped into two sets of systems:

Timing & Scoring Systems (T&S) capture real-time data during the competition. Through electronic feeds to other systems, this data is made available for use on the scoreboard, in TV graphics and other related outputs, by OVR.

On Venue Results Systems (OVR) running at each of the competition venues receives both data from T&S and manually entered data to calculate results of each Olympic event. OVR Systems then distribute the results to INFO.

Concerning the security of the IT infrastructure, for the Beijing 2008 Olympic Games, more than 12 million IT security events were collected and filtered events each day to detect any potential security risk for the Olympic Games IT systems. From these, less than 100 were identified as real issues. All were resolved, with no impact at all on the Olympic Games (http://www.atosorigin.com/olympic_games).

For an event of this magnitude, deadlines are not negotiable, when world-class athletes are ready to compete for gold after years of rigorous training and qualification and viewers are anxious to enjoy such a show, there are no second chances. System disruption or failure is not acceptable. In this context, the main challenge of the SIEM infrastructure in Olympic Games is to protect the games IT infrastructure from any undesired and/or uncontrolled phenomena which can impact any part of the result chain and associated services.

3 Security Management Challenges

Security Information Event Management (SIEM) solutions have become the backbone of the all Service Security systems. They collect data on events from different security elements, such as sensors, firewalls, routers or servers, analyze the data, and provide a suitable response to threats and attacks based on predefined security rules and policies. Despite the existence of highly regarded commercial products, their technical capabilities show a number of constraints in terms of scalability, resilience and interoperability.

The MASSIF project aims at achieving a significant advance in the area of SIEMs by integrating and relating events from different system layers and various domains into one more comprehensive view of security-aware processes and by increasing the scalability of the underlying event processing technology. The main challenge that MASSIF will face is to bring its enhancements and extensions into the business layer with a minimum impact on the end-user operation.

A further goal of the MASSIF project is to integrate these results in two existing Open Source SIEM solutions, namely OSSIM (http://alienvault.com/community) and Prelude (http://www.prelude-technologies.com/) and to apply them to four industrial scenarios, including the Olympic Games IT infrastructure.

Aligned with the security needs of these scenarios, MASSIF challenges can be arranged according to the following dimensions:

3.1 Collection

The data gathering must have the ability to deal with a large number of highly heterogeneous data feeds. The capabilities of the SIEM will be improved by the integration of new types of security tools/probes. This implies that the parsing/processing logic (and code) should be as much as possible decoupled from the specific characteristics of the data format and related technologies. Additionally, the parsing logic and related languages must allow effective processing of virtually any type of security relevant event in cyber-environment, including, in the future, possible extensions to capture and process security events from physical security equipment.

Moreover, 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 to propagate relevant events to the SIEM core platform without loss of information.

These concepts are implemented in MASSIF by the Generic Event Translation (GET) framework. The GET framework relies on grammar-based parsing [2], [3] and compiler-compiler technology to implement effective processing of security-relevant events. The main components of the Generic Event Translation Framework are represented in Fig.1. A brief description of each component is provided in the following:

Generic Event Translation (GET) Manager. This component is responsible for the activation of all the modules which belong to the Generic Event Translation framework. In particular, it is in charge of the generation of new Adaptable Parser modules, as new grammars are added to the system.

Event Dispatcher. This component connects each source of sensor events to the appropriate GET Access Point (GAP), in order to provide it with an Adaptable Parser which is capable of processing the specific event format.

GET Access Point (GAP). It is responsible for orchestrating the translation process of the GET. It is in charge of extracting the content of source messages in the source specific format, using the event parsing capabilities of the Adaptable Parsers and requesting the final conversion to the MASSIF Event Format by the MASSIF Event Manager (MEM).

Format-Specific Grammars. These contain semantic description of the different event formats that are used for the creation of the Adaptable Parsers.

Adaptable Parsers. These components provide the parsing capabilities for the different types of events used in MASSIF. They allow for extraction of the relevant information for the event to be inserted in the MASSIF Event Format.

MASSIF Event Manager (MEM). It translates the event content, extracted by the Adaptable Parser to the MASSIF Event Format, thus allowing the event to be sent to the reliable event bus. It also attaches to each MASSIF Event a timestamp, which is made available by the synchronized time source of the Resilient Architecture.

Sender Agent. It is the component that finally sends MASSIF-formatted events to the reliable event bus.

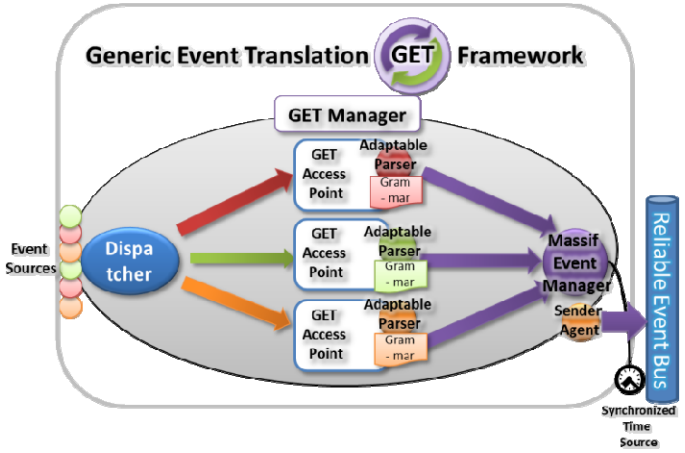

Fig. 1. Generic Event Translation main components

3.2 Processing

The core of MASSIF is an event processing engine capable of handling high input rates and of optimizing the amount of resources based on the actual load [4]. 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. The system must process input data at high rate and provide meaningful results with soft real-time requirements. The engine should be able to aggregate, abstract and correlate heterogeneous events from multiple sources at different levels of the system stack.

3.3 Correlation

MASSIF targets at correlation capabilities across layers of security events, from network and security devices as well as from the service infrastructure such as correlation of physical and logical event sources. 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.

3.4 Resilience

Special emphasis will be placed on providing a highly resilient architecture against attacks, concurrent component failures, and unpredictable network operation conditions. The event flows should be protected, from the collection points through their distribution, processing and archival. 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.

3.5 Timeliness

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.

3.6 Sensitive Information

MASSIF features for forensic support will satisfy the following requirements:

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 be persisted 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 persisted to stable storage (most of the data should be processed in real-time and thrown away).

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

4 Related Work

The research in MASSIF combines aspects of process monitoring, simulation, and analysis as well as trustworthiness and scalability of the complex event processing architecture itself. Relevant contributions from these broad areas are:

Attack Modelling, Simulation and Risk Evaluation. The technology most relevant to the modelling and simulation methods to be developed for MASSIF is commonly called attack-graph analysis, an approach presented by Phillips and Swiler [5] in. Two participants of the MASSIF team namely Fraunhofer SIT and SPIIRAS are actively researching in that area [6], [7].

Predictive Security Analysis. The predictive security analysis in MASSIF will use the method given in [8] to analyse the security requirements. Based on this, the attack models together with the SIEM's information about the current attack state and the process models together with the SIEM's information about the current process state can be used to derive a near future view of possible upcoming security problems [9]. This information can now be used in an ontology-driven approach to select appropriate countermeasures [10].

SIEM Scalability and Trustworthiness. Complex Event Processing (CEP) is a promising technology to improve current SIEM systems. It allows processing of large amounts of streaming data in real time and provides information abstraction and correlation, similarly to SIEM correlation engines. MASSIF will develop new parallel distributed CEP technology that overcomes scalability limitations due to single node bottlenecks or high distribution overhead [4]. The trustworthiness of the SIEM architecture will be improved by utilising secure digital chains of evidence [11].

5 Conclusions and Outlook

The MASSIF project is still at an early stage. However, the challenges that the project aims to achieve will provide a significant advance in the area of Security Information and Event Management (SIEM) by integrating and relating events from different system layers and various domains into one more comprehensive view of security-aware processes and by increasing the scalability of the underlying event processing technology. To address the challenges the MASSIF partners plan to develop a novel SIEM system with the following solutions and implied research and development needs.

In order to enable a highly scalable security situation assessment, the MASSIF event engine will provide a flexible language to express filtering, transformation, abstraction, aggregation, intra-layer and cross-layer correlation as well as storage of security events. The event engine will be able to process with the same language both the real-time event flow as well as stored events for forensic analysis. Additionally, specific collectors to translate from the external languages into the event engine language will be provided.

Ideally, the MASSIF system should be able to analyze upcoming security threats and violations in order to trigger remediation actions even before the occurrence of possible security incidences. Therefore, new process and attack analysis and simulation techniques will be developed in order to be able to relate events dynamically from different execution levels, define specific level abstractions, evaluate them with respect to security issues and during runtime interpret them in context of specific security properties. Novel adaptive response technologies will enable anticipatory impact analysis, decision support and support impact mitigation by adaptive configuration of countermeasures such as policies.

Due to the highly distributed and heterogeneous nature of the various components, and the hostile and unpredictable operational environment, it becomes a challenge to design an integrated solution for the protection of the SIEM framework itself. Therefore, the MASSIF system will be based on a resilient, trust-enabling architecture with trusted collection of security-relevant data from highly heterogeneous trusted networked devices in order to ensure unforgeability of stored security events and to support criminal/civil prosecution of attackers.

Acknowledgements. The work in this paper has been sponsored by the EC Framework Programme as part of the ICT MASSIF project (grant agreement no. 257644).

References

1. McRoskey, S.R.: Security and the Olympic Games Making Rio an Example. Yale Journal of International Affairs (2010)
2. Turmo, J., Ageno, A., Catala, N.: Adaptive Information Extraction. ACM Computing Surveys 38(2) (2006)
3. Campanile, F., Cilardo, A., Coppolino, L., Romano, L.: Adaptable Parsing of Real-Time Data Streams. In: Proc. of The Fifteen Euromicro Conference on Parallel, Distributed and Network-based Processing (PDP 2007), Naples, Italy, February 7-9, pp. 412–418. IEEE Computer Society Press, Los Alamitos (2007)
4. Gulisano, V., Jimenez-Peris, R., Patiño-Martínez, M., Valduriez, P.: A Large Scale Data Streaming System. In: 30th IEEE Int. Conf. on Distributed Systems (ICDCS), Genoa, Italy (2010)
5. Cynthia, A.P., Swiler, L.P.: A graph-based system for network-vulnerability analysis. In: NSPW 1998, Proceedings of the 1998 Workshop on New Security Paradigms, pp. 71–79. ACM Press (1998)
6. Rieke, R.: Abstraction-based analysis of known and unknown vulnerabilities of critical information infrastructures. International Journal of System of Systems Engineering (IJSSE) 1, 59–77 (2008)
7. Kotenko, I., Stepashkin, M., Doynikova, E.: Security Analysis of Computer-aided Systems taking into account Social Engineering Attacks. In: 19th Euromicro International Conference on Parallel, Distributed and network-based Processing (PDP 2011), Ayia Napa, Cyprus (2011)

8. Fuchs, A., Rieke, R.: Identification of Security Requirements in Systems of Systems by Functional Security Analysis. In: Casimiro, A., de Lemos, R., Gacek, C. (eds.) Architecting Dependable Systems VII. LNCS, vol. 6420, pp. 74–96. Springer, Heidelberg (2010)
9. Rieke, R., Stoynova, Z.: Predictive Security Analysis for Event-Driven Processes. In: Kotenko, I., Skormin, V. (eds.) MMM-ACNS 2010. LNCS, vol. 6258, pp. 321–328. Springer, Heidelberg (2010)
10. Cuppens-Boulahia, N., Cuppens, F., Lopez, J., Vasquez, E., Guerra, J., Debar, H.: An ontology-based approach to react to network attacks. International Journal of Information and Computer Security 3, 280–305 (2009)
11. Kuntze, N., Rudolph, C.: Secure digital chains of evidence. In: Sixth International Workshop on Systematic Approaches to Digital Forensic Engineering (SADFE) 2011, Oakland, USA (2011)