Architecting a Security Strategy Measurement and Management System

Roland Rieke
Fraunhofer Institute SIT
Rheinstrasse 75
Darmstadt, Germany
roland.rieke@sit.fraunhofer.de

Julian Schütte
Fraunhofer Institution AISEC
Parkring 4
Garching, Germany
julian.schuette@aisec.fraunhofer.de

Andrew Hutchison
T-Systems South Africa
4 Churchill Close, Bellville
Cape Town, South Africa
andrew.hutchison@t-systems.co.za

ABSTRACT

The use of formal models to guide security design is appealing. This paper presents a model driven approach whereby security systems in operation can be assessed and measured against various requirements that are defined when the system is created. By aligning with organisational policy, and business requirements of a specific system, design and operation can proceed in a way that allows measurement of how successfully security objectives are being achieved. This paper describes a model driven approach which overcomes the contextual restrictions of existing solutions. In particular, where models have been used previously these have tended to be predefined and closed models, whereas the approach described here is an extensible model that comprises all parts of the security monitoring and decision support process. By means of interlinked semantic concepts, the proposed security strategy meta model provides a way to model security directives at an abstract level, which can be automatically compiled into specific rules for an underlying framework of monitoring, decision support, and enforcement engines.

Categories and Subject Descriptors

K.6.5 [**Management of Computing and Information Systems**]: Security and Protection; D.2.1 [**Software Engineering**]: Requirements/Specifications; C.2.3 [**Computer-Communication Networks**]: Network Operations—*Network management,Network monitoring*

General Terms

Security,Management,Measurement

Keywords

Security strategy, security monitoring, decision support, security information and event management, information security measurement model, governance and compliance

1. INTRODUCTION

Cyber Security is an area of great global focus, yet it is both hard to manage and – arguably – even harder to measure. But the two concepts go together: if some sort of *measurement* approach can be implemented, it should at least be possible to *manage* systems better and assess whether they are meeting the security objectives that they were designed with. In spite of the fact that technical security solutions are deployed, there are numerous instances of processes or transactions being compromised. Part of the challenge with security implementations is that they are made in isolation from any formal specification or model of what the security profile should look like. In the absence of a holistic view that extends from business process to logical and technical security realisation, there is high potential for gaps or mismatches to occur. Fueling this situation is the fact that life-cycle approaches to security are not easily applied – or measured.

In this paper we argue for a *meta model* approach to drive security from design to implementation, through an *analysis and refinement* approach, and also through a *security measurement* approach which would enable assessment of the system's performance against the security requirements it was designed for. To achieve a meta model approach for security, several phases are required and in this paper we present a *Security Information Meta Model* (SIMM) consisting of: (a) high level goal setting, (b) security requirements, (c) measurement requirements, and (d) objects of measurement. Through applying this model, high level goals for security can be established and defined. Importantly, measurement objectives can also be developed and stated at this point. By proceeding in this way, security can be designed in such a way that it can be measured (and managed). Activities of analysis and refinement are required to move from security requirements to measurement requirements. In this process, objects of measurement also need to be identified.

In order to cover the operational aspects of this concept, we make use of a *Security Strategy Meta Model* (SSMM) [23] that describes the control flow at runtime, independent from the underlying event description language. A specific rule from a *Security Strategy Model* (SSM) that adheres to the SSMM is called *Security Directive* (SD). A distinguished *Security Strategy Component* controls the execution of the SDs. It can execute a SD or parts of it directly or delegate workload to a specialised *Security Strategy Processing Component* (SSPC). The overall aim is to overcome the contextual restrictions of existing solutions, with their predefined and closed models, and rather to provide an extensible model that spans all parts of the security monitoring and decision support process, namely: (i) detecting threatening events; (ii) putting them in context of the current system state; (iii) explaining their potential impact with respect to some security- or compliance model; and (iv) taking appropriate actions. Depending on the outcome of the analysis of these components, other components that implement decision support and enforcement will be triggered. The proposed SSMM together with the framework of SSPCs could be used as a core of a technology platform for an integrated concept for governance, risk and compliance [19]. Furthermore, we consider the proposed approach to be applicable within the design of a cyber attack information system [24], which uses collaborative detection and response mechanisms for high-level situational awareness and coordination of local incident response.

The structure of the paper is as follows: first the use of *Security Information and Event Management* (SIEM) techniques for information security management in general is discussed; next integration into a system architecture is presented and this is then also contrasted with existing work, positioning how this approach differs from other similar work.

2. USING SIEM FOR INFORMATION SECURITY MANAGEMENT

Information security management is needed to protect information and information infrastructure assets of an organisation against the risks of loss, misuse, disclosure or damage. It specifically describes controls that an organisation needs to implement to ensure that it is sensibly managing the risks. The ISO/IEC 27000-series comprises information security standards in the context of an *Information Security Management System* (ISMS). Specifically, the ISO/IEC 27004 standard [11] provides guidance on the development and operation of measures and measurement, and reporting of the results, with the aim to help organisations to systematically improve the effectiveness of their ISMSs.

2.1 Information Security Measurement

The ISO/IEC 27004 standard defines an *Information Security Measurement Model* (ISMM) that provides a structure linking an *information need* to the relevant *object of measurement*. Furthermore, it describes how the attributes of an object of measurement are quantified and converted to indicators, thus providing a basis for decision making. Figure 1 depicts an abstract view of the ISMM. Specific objects of measurement relevant for the work presented here include the status of information assets protected by the controls and the measurement of process behaviour. This standard further provides a template for an *information security measurement construct* and several examples of concrete measurement constructs. Further examples are given in [15]. The frequency of reporting measurement results is in most of the given examples “monthly”, “quarterly” or “yearly”. The design of measurement constructs as described in the ISO/IEC 27004 standard covers in detail the steps needed to derive the measurement results from a given object of measurement [15]. However, the method for identification of the objects of measurement from the information needs is not specified in detail (cf. the dashed arrow in Fig. 1).

In the following, we show how we map relevant parts of the information security measurement process to a SIEM information flow. While the former usually takes months or years to be updated, through manual inspection and creation of checklists, the latter allows a near-realtime observation of incidents and the mapping of them back to the information security management requirements. Thus, by the application of our model, we expect a semi-automatic and significantly faster update cycle of compliance checks.

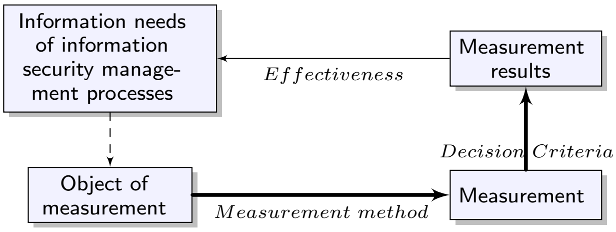

Figure 1: Abstract view of the ISMM

2.2 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. The frequency of reporting measurement results is in most cases very high. In [18], e.g., it is reported that for the Beijing 2008 Olympic Games, more than 12 million IT security events were collected and filtered 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.

The rules for measurements and correlation are usually a mixture of predefined rules from the SIEM system provider and specific rules from the SIEM system user. Compared to the ISMM, a SIEM based approach presents several key advantages:

1. the measurement frequency is much higher,
2. the system is tool based and most actions are executed automatically, and
3. a decision support system or intrusion response system [25] may offer automatic countermeasures.

There are also disadvantages though:

1. the rules are written in vendor specific notation,
2. back-traceability across layers of derived measures from base measures is not always possible,
3. rules don’t necessarily use contextual information,
4. the effects of countermeasures are not clear,
5. there is no clear derivation of the measurement rules from the information needs,
6. therefore, there is no traceback possibility from measurement results to information needs, and

7. because there are also – best practice – rules used from external sources, there is no clear way to express how these rules are related to the company goals.

In order to combine the advantages of both ISMM and SIEM concepts, we propose a Security Strategy Meta Model (SSMM) that addresses the above mentioned disadvantages.

2.3 Interlinked Semantic Concepts

The SIEM information flow is based on rules specifying which system behaviours to observe. However, simply reacting to individual rules is of little help for users who need to understand the actual incident that has been detected and its implications in terms of the security model. For this purpose, we propose firstly, to refine the left side of the ISMM by an *Security Information Meta Model* (SIMM), which should be derived in an measurement requirements elicitation process [20]. The most important objectives of this process are:

Coverage of Security Goals. The requirements elicitation method should ensure that uncovered aspects of high-level security goals are revealed.

Information Needs. A lack of SIEM monitoring capabilities would prevent the derivation of assumptions necessary for the reasoning process. This should be detected in the requirements elicitation process.

Sufficiency of Monitoring Capabilities. Assumptions can be derived from the monitoring capabilities for reasoning whether the given requirements are fulfilled under these assumptions. This reasoning process can't be successful if monitoring capabilities are insufficient or can't be assigned to entities used in the current abstraction level of the system model.

Traceability. The derived relations between security event measurements, the associated security requirements and corresponding assumptions, and the security goals can be used to identify the concrete high-level goals affected by the measurements.

Secondly, we propose to use an SSMM that comprises all aspects of SIEM functionality as well as countermeasure configuration support in order to cover the operational aspects of the overall security management goal. By means of interlinked semantic concepts, the SSMM provides a way for users to model incident detection rules at an abstract level, which are automatically compiled into specific rules for an underlying *Complex Event Processing* (CEP) engine. Thus, the SSMM serves as a generic and extensible model on top of a specific rule language used for actual event evaluation. At the model level, these rules are linked to environmental conditions, countermeasures, and explanations based on an external security model. The SSMM is constructed of four parts : *on*, : *if*, : *do*, and : *why*, which are derived from the measurement requirements on the one hand, and which refer back to the SIMM on the other hand (cf. Fig. 2).

The : *on* part specifies the incident detection by means of an *event stream property*. This addresses the first of the "disadvantages" mentioned previously, by abstracting from specific event processing languages. Once specified, event stream properties can be reused across different CEP engines and do not require users to be expert in such systems. This supports a separation of duty, where security engineers can concentrate on modelling security information measurement and do not necessarily need to be knowledgeable regarding all the technical details of a CEP engine.

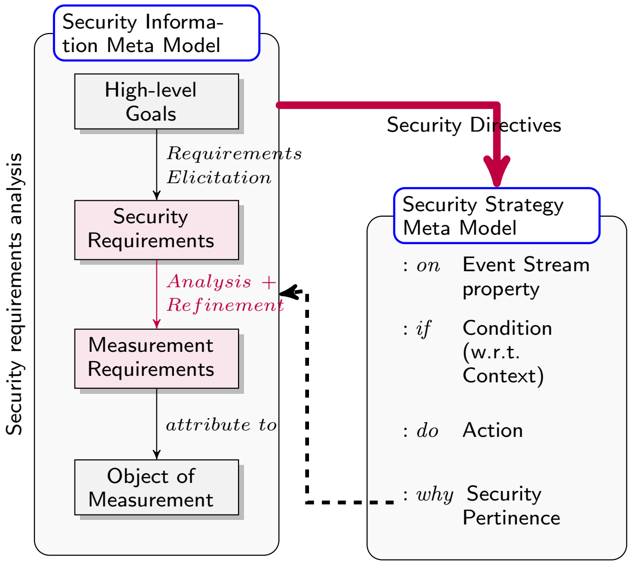
Security Information Meta Model
High-level Goals
Requirements Elicitation
Security Requirements
Analysis + Refinement
Measurement Requirements
attribute to
Object of Measurement
Security requirements analysis
Security Directives
Security Strategy Meta Model
: on Event Stream property
: if Condition (w.r.t. Context)
: do Action
: why Security Pertinence

Figure 2: Security Information Meta Model and Security Strategy Meta Model

Addressing the second stated "disadvantage", the model is able to express correlations of incoming events "horizontally" (i.e., as steps in a workflow), as well as "vertically". While most SIEMs focus on *horizontal* correlation, *vertical* correlation is an interesting feature, because it allows the linking of information across different levels of abstraction, such as events from an intrusion detection system with the currently threatened protection goal. To address the third "disadvantage", the : *if* part of the model allows for inclusion of context information. This is of special importance for stateful incident detection, as encountered in the monitoring of ongoing processes, and also to increase the likelihood of discovery of a targeted attack. Addressing the fourth "disadvantage", our model allows combination of SIEM functionality – for detecting incidents – with an actual handling of these. This is modelled by the : *do* part of the model that refers to countermeasures to be taken, ranging from a simple reporting, to autonomous re-configurations of the system. In [22], we have shown how such an autonomous and goal-driven re-configuration can be realised. In order to close the traditional *plan-do-check-act* [8] cycle of Information Security Measurement, addressing the fifth, sixth and seventh "disadvantages", we finally need to link security incident detection to the high-level security requirement. This is achieved by the : *why* part of the SSMM, represented by the dotted arrow in Fig. 2. The : *why* part should contain information similar to the *information security measurement construct* defined in the ISO/IEC 27004 standard.

In [23], we have defined a language from which to form a model that satisfies these requirements.

3. INTEGRATION INTO SYSTEM ARCHITECTURE

We now describe a mapping of the SSMM to the components of a proposed monitoring infrastructure. The goal is to enable the inclusion of existing engines, which need not know about the overall security strategy but only receive specific tasks in their respective language. We first describe the concept and then continue with a prototypical implementation.

3.1 Security Strategy Processing Components

Conceptually, the implementation of the processing of the SSMM is composed of SSPCs. The main components and some optional components of the proposed system architecture are illustrated in Fig. 3. A distinguished *Security Strategy Component* controls the execution of the SDs. It can execute a SD, or parts of it, directly or delegate the workload to specialised components. The Security Strategy Component initially gets the SSM from the *Security Information Modeller*. It parses the SDs of the SSM, identifies the responsible SSPCs for each subtask, and distributes a respective configuration to the relevant SSPC. The CEP engine normally processes the : *on* part of the SD. The security monitoring probes, which are described at an abstract level in the SSM, have to be compiled to the configuration language of the actual CEP engine, if an engine specific specification is not given in the : *on* part of the SD. Optionally, the events could be processed directly. Furthermore, other event processing components such as intrusion visualisation could be triggered. The : *if* part of the SD can be processed by several different components, responsible for different aspects of the domain or several domains. One component, which will be needed in most implementations, is that responsible for the provisioning of the network state information. Other components could, e.g., provide cyber-physical models, workflow specifications, business process information or process visualisation. For example, in the project MASSIF [3] we are currently implementing an advanced SIEM architecture that comprises an *Attack Modeling and Security Evaluation Component* (AMSEC) [13] and a *Predictive Security Analyser* (PSA) [9]. The AMSEC component provides attack scenario recognition by real-time event analysis and prognosis of future attack steps by recognition of the attacker model. The PSA component provides advanced, application aware security monitoring capabilities. Specifically, it supports close-future process behaviour simulation and prediction of possible security violations. Prior to the start of the engine, the process description and security goals/events are transformed into “PSA understandable” models, which are then used for the continuous real-time analysis and close-future simulation. Thus, : *if* components such as AMSEC and PSA provide situational awareness with regard to network state, attack state, and application state.

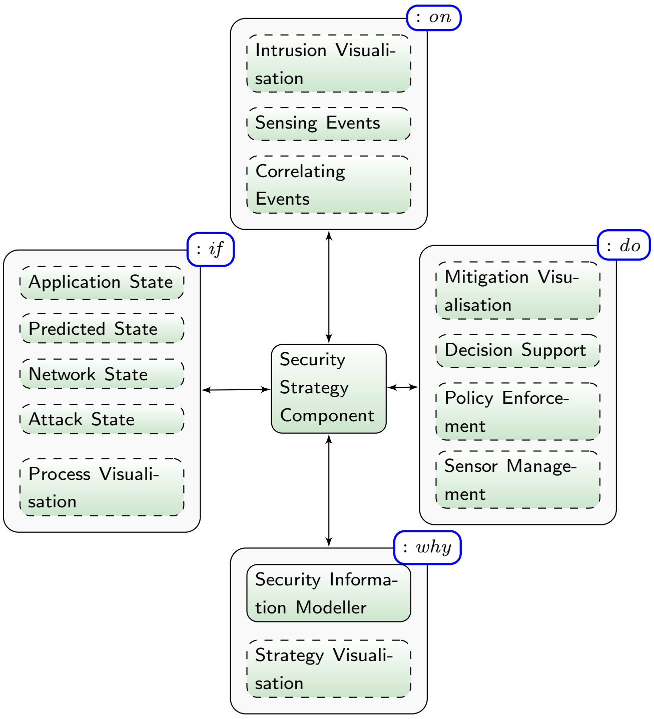

Figure 3: Conceptual components of the framework

Depending on how the : *if* condition evaluates, the respective : *do* components will be triggered. These components can implement, e.g., *simulative mitigation visualisation*, *decision support*, *policy enforcement* or *sensor management*. A sensor management component can control the configuration of sensors in a monitored system, e.g., the (de-)activation and the adaptation of the sampling rate to an optimal level [5].

A *security information modeller component* is responsible for maintaining the security strategy and a *strategy visualisation component* can help to assist in the : *why* determination, thus resolving the traceability requirement.

As a special case, the functionalities of some or all components could also be implemented within one engine. In this case no translation into the specific configuration languages is needed and the SDs could be interpreted directly.

3.2 Prototype Implementation

We have implemented the model in a prototype in order to test whether it can be used, as intended, for detecting security incidents. Additionally, the prototype implementation should confirm that our component based system architecture is applicable and the idea of SSPC extensions is practical. For this purpose, we implemented a simple Security Strategy Component which processes a given SSM and coordinates the different SSPCs for evaluating it. When the SSM is loaded, the Security Strategy Component first parses the : *on* part of a SD and transforms it into a query for the registered CEP engine. Then, the query is registered at the respective engine and the Security Strategy Component receives a callback whenever the query is triggered, i.e., the : *on* part of the SD has been met. In that case, the Security Strategy Component creates an *evaluation context* object, which acts as container during the evaluation process, and writes the attributes received from the event stream to it. The evaluation context is then passed on to subsequent components for evaluating the condition in the : *if* part of the SD. If the condition has evaluated to “true”, the Security Strategy Component loads the action components indicated by the : *do* part and invokes them, passing the evaluation context object as parameter. Because in the prototype, the : *why* part provides merely explanatory reasons, it is not involved in the evaluation process and can rather be explored by users to investigate security implications of the detected incident.

Each of the components, i.e., the Security Strategy Component, as well as the SSPCs, has been created in the form of OSGi bundles and communicates over R-OSGi, a binary RPC protocol. This allows us to dynamically load and unload components, even from a remote repository, so that it is possible to support additional actions by providing respective components in the repository. As an event correlation engine, we have used Esper [2], which comes with its own EPL query language. Listing 1 shows an example of an event description (written in Turtle notation), referring to an anomaly in the traffic to the syslog service, and in Listing 2 its translation into EPL. While the EPL representation is more compact, it is only applicable to the specific CEP engine and does not bear any semantics which could be linked to external models of security requirements and controls.

Listing 1: Modelled Event Condition

```
:historyEvent
      :hasName "HistoryDBServerAnomaly" ;
      :hasCriterion [ :hasParam1 "avgTraffic" ;
                      :hasParam2 42^^float ;
                      :hasBooleanOp :gt ] ;
      :hasExtractor [
            :hasEventChannel [ rdf:type :SyslogChannel;
                  :hasFields "sourceIP" ;
                  :hasName "IPStreamToIPX" ;
                  :hasFields "FIXEDdestIP","traffic" ];
            :hasFunction [
                  :hasParam1 "traffic" ;
                  :providesVariable "avgTraffic" ;
                  :hasScope "30 sec" ;
                  :hasOp :avg ] .
```

Listing 2: Generated EPL Query

```
SELECT   sourceIP?,
         avg(cast(traffic?,float)) AS avgTraffic
FROM     SyslogChannel.win:time(30 sec)
HAVING   cast(avgTraffic?,float)>42
```

The combination of our semantic model with a CEP engine allows us to efficiently evaluate incoming event patterns, while still being able to refer to their semantic description, once the pattern has been found. While an extensive evaluation of our prototype is still outstanding, first results are encouraging and make us confident that it is practically applicable.

4. RELATED WORK

Work related to ours is on the one hand concerned with modelling security-relevant information in a way that creates the possibility to reason about it and link it to the ISMM [11]. On the other hand, we review current SIEM systems to highlight how they could be improved by adding a model-based SIMM. In this paper, we rely on the overall plan-do-check-act cycle [8] and the information flow described by the ISO27004 standard [11].

In [10], an approach to create ISO27001-based metrics based on a security ontology is proposed. While it lacks the automatic gathering of measurements, it could serve as a later extension to our SIMM, which is more focused on measurable technical events. Further, linking semantic attacker models to the : *why* part of the SSMM could be promising (cf. the AMSEC model [13]). Another example for a potential information source is the *Engineering Knowledge Base* (EKB) [16], which is an ontology relating to sensor values and combining run-time with development-time models. It is focused on the analysis of industrial automation systems, and is used to define SPARQL or SWRL queries over sensor definitions. As we have a similar goal of finding inconsistencies, we believe that an approach like the EKB could help defining which inconsistencies to look for in event streams, and thus which measurement points might indicate violations of the security requirements. Other approaches of interest to this end are the modelling concepts in [12], where business, application, physical, and technical information is merged and related, as well as concepts to use event-triggered rules for sensing and responding to business situations in [21]. In this paper, we focused on a model to bridge the gap between high-level security measurements and data gathered by SIEM engines, like OSSIM [4], Prelude [17], or Akab [1]. OSSIM detects events at the network layer and stores respective attributes like *IP address* or *port number* in a relational database. Thus, while it is possible to link these attributes to our model, OSSIM itself does not support reasoning over gathered data, nor extending its model. Similarly, Akab [1] is a SIEM appliance for monitoring network events. It uses a proprietary event format and also stores collected events persistently in a database. Prelude [17] is an open source SIEM framework which relies on the open IDMEF [7] event format. Also related to the model-based security information measurement that we envisage are commercial tools RSA Archer, ArcSight ESM, or IBM Tivoli Security Information and Event Manager [6]. Although they also aim at relating incidents to compliance catalogues and corporate policies, they rely on predefined event structures, comprising specific technical attributes [14]. The RSA Archer Threat Monitor manages an assets catalogue and links it to security-relevant information, such as known vulnerabilities and patch levels. It does not however feature an extensible and semantic model which would enable automatic reasoning regarding the implications of a detected incident as it relates to the affected security requirements. It could also make amendments based on information from external sources like our PSA.

5. CONCLUSION

This paper has presented a model driven approach for architecting a security strategy measurement and management system. Concretely, it has described the definition of security objectives for a particular system, and a mechanism for collecting information from operational systems in a manner which enables assessment and measurement of how well the system is fulfilling the security objectives. Existing SIEM solutions are limited, while this approach overcomes these contextual restrictions (typically predefined, closed models) offering an extensible and open model, encompassing all parts of the security monitoring and decision support process, namely: (i) detecting threatening events; (ii) putting them in context of the current system state; (iii) explaining their potential impact with respect to some security- or compliance model; and (iv) taking appropriate actions. The proposed deployment model brings together all parts of security runtime management, namely, detection, reporting, handling, and explanation of security incidents, which are to date covered by different systems, such as intrusion detection systems, CEP engines [2], SIEM systems [4, 17, 1], intrusion response systems, cyber attack information systems [24], and governance, risk management, and compliance systems [19]. So, the model supports an integration of functionalities of these existing systems into one coherent security strategy management framework.

6. ACKNOWLEDGMENTS

The work presented here was developed in the context of the projects MASSIF (ID 257475) being co-funded by the European Commission within the Seventh Framework Programme, the project SealedCloud being funded by the German Federal Ministry of Economics and Technology (BMWi), and the project ACCEPT (ID 01BY1206D) being funded by the German Federal Ministry of Education and Research.

7. REFERENCES

[1] Araknos website. http://www.araknos.it/en.html, 2012. [Online; accessed 16-Sep-2012].

[2] Esper – Complex Event Processing. http://esper.codehaus.org/, 2012. [Online; accessed 16-Sep-2012].

[3] Project MASSIF website. http://www.massif-project.eu/, 2012. [Online; accessed 16-Sep-2012].

[4] AlienValult. AlienVault Unified SIEM. http://alienvault.com/, 2012. [Online; accessed 16-Sep-2012].

[5] L. Baumgärtner, P. Graubner, M. Leinweber, R. Schwarzkopf, M. Schmidt, B. Seeger, and B. Freisleben. Mastering Security Anomalies in Virtualized Computing Environments via Complex Event Processing. In *Proceedings of the The Fourth International Conference on Information, Process, and Knowledge Management (eKNOW 2011)*, pages 76–81. XPS, 2012.

[6] A. Buecker, J. Amado, D. Druker, C. Lorenz, F. Muehlenbrock, and R. Tan. *IT Security Compliance Management Design Guide with IBM Tivoli Security Information and Event Manager.* IBM Redbooks, July 2010. ISBN 0-7384-3446-9.

[7] H. Debar, D. Curry, and B. Feinstein. The Intrusion Detection Message Exchange Format (IDMEF). RFC 4765 (Experimental), March 2007.

[8] W. E. Deming. *The new economics for industry, government, education.* Massachusetts Institute of Technology, Center for Advanced Engineering Study, Cambridge, MA, 1993.

[9] J. Eichler and R. Rieke. Model-based Situational Security Analysis. In *Proceedings of the 6th International Workshop on Models@run.time at the ACM/IEEE 14th International Conference on Model Driven Engineering Languages and Systems (MODELS 2011)*, volume 794 of *CEUR Workshop Proceedings*, pages 25–36. 2011.

[10] S. Fenz. Ontology-based generation of it-security metrics. In *Proceedings of the 2010 ACM Symposium on Applied Computing*, SAC '10, pages 1833–1839, New York, NY, USA, 2010. ACM.

[11] I. Iec. ISO/IEC 27004:2009 - Information technology - Security techniques - Information security management - Measurement. *ISOIEC*, 2009.

[12] F. Innerhofer-Oberperfler and R. Breu. Using an enterprise architecture for it risk management. In J. H. P. Eloff, L. Labuschagne, M. M. Eloff, and H. S. Venter, editors, *ISSA*, pages 1–12. ISSA, Pretoria, South Africa, 2006.

[13] I. Kotenko, A. Chechulin, and E. Novikova. Attack Modelling and Security Evaluation for Security Information and Event Management. In P. Samarati, W. Lou, and J. Zhou, editors, *SECRYPT*, pages 391–394. SciTePress, 2012.

[14] Lieberman Software. Common event format configuration guide, Jan. 2010.

[15] K. Lundholm, J. Hallberg, H. Granlund, and T. forskningsinstitut. Informationssystem. *Design and Use of Information Security Metrics: Application of the ISO/IEC 27004 Standard.* FOI-R. Information systems, Swedish Defence Research Agency, 2011.

[16] M. Melik-Merkumians, T. Moser, A. Schatten, A. Zoitl, and S. Biffl. Knowledge-based runtime failure detection for industrial automation systems. In *Workshop Models@run.time*, pages 108–119. CEUR, 2010.

[17] PreludeIDS Technologies. Prelude PRO 1.0. http://www.prelude-technologies.com/, 2010.

[18] E. Prieto, R. Diaz, L. Romano, R. Rieke, and M. Achemlal. Massif: A promising solution to enhance olympic games it security. In C. K. Georgiadis et al., editors, *Global Security, Safety and Sustainability & e-Democracy*, volume 99 of *Lecture Notes of the Institute for Computer Sciences, Social Informatics and Telecommunications Engineering*, pages 139–147. Springer, 2012.

[19] N. Racz, E. Weippl, and A. Seufert. A Frame of Reference for Research of Integrated Governance, Risk and Compliance (GRC). In B. D. Decker and I. Schaumüller-Bichl, editors, *Communications and Multimedia Security*, volume 6109 of *Lecture Notes in Computer Science*, pages 106–117. Springer, 2010.

[20] J. Repp and R. Rieke. Formal Specification of Security Properties. Deliverable D4.2.1, MASSIF Project, September 2011. http://www.massif-project.eu/sites/default/files/deliverables/D4.2.1%20-%20Formal%20Speci%EF%AC%81cation%20of%20Security_v1.0_final.pdf.

[21] J. Schiefer, S. Rozsnyai, C. Rauscher, and G. Saurer. Event-driven rules for sensing and responding to business situations. In H.-A. Jacobsen, G. Mühl, and M. A. Jaeger, editors, *DEBS*, volume 233 of *ACM International Conference Proceeding Series*, pages 198–205. ACM, 2007.

[22] J. Schütte. Goal-based policies for self-protecting systems. In *Proceedings of the International Conference on Advanced Information Networking and Applications (AINA).* IEEE Computer Society, 2012.

[23] J. Schütte, R. Rieke, and T. Winkelvos. Model-based security event management. In *International Conference "Mathematical Methods, Models and Architectures for Computer Networks Security" (MMM-ACNS-12).* Springer, 2012.

[24] F. Skopik, Z. Ma, P. Smith, and T. Bleier. Designing a cyber attack information system for national situational awareness. In N. Aschenbruck, P. Martini, M. Meier, and J. Tölle, editors, *Future Security*, volume 318 of *Communications in Computer and Information Science*, pages 277–288. Springer, 2012.

[25] N. Stakhanova, S. Basu, and J. Wong. A taxonomy of intrusion response systems. *Int. J. Inf. Comput. Secur.*, 1(1/2):169–184, Jan. 2007.