

Service-oriented Event Assessment – Closing the Gap of IT Security Compliance Management

F. Majer¹, M. Nussbaumer¹, D. Riexinger², V. Simon¹

¹Steinbuch Centre for Computing
Karlsruhe Institute of Technology
D-76128 Karlsruhe, Germany

²IBM Germany
D-68016 Mannheim, Germany

[majer | nussbaumer]@kit.edu
dieter.riexinger@de.ibm.com
volker.simon@student.kit.edu

Abstract: Frequently, Security Monitoring is equated with network intrusion detection. However, Security Monitoring has a much broader scope. It also comprises detection of insider attacks. Since the Enron bankruptcy, monitoring of privileged access to financial data has become a legal requirement stipulated for example in the Sarbanes-Oxley Act (SOX 404). Monitoring of privileged access requires evaluation of its necessity, permission, and correctness. As a result, detection of privileged access is not sufficient and must be reviewed in its business context. Data from various sources combined with business process contexts establish a sound basis for the assessment of a privileged access. Usually, the required data is spread over different data sources within an organization offering heterogeneous interfaces of any kind. Security administrators use multiple applications and data interfaces which result in a time-consuming and error prone process. Security Monitoring is, on the contrary, all about attack detection and prevention in a timely manner. This paper introduces the concept of service-oriented context determination, which efficiently describes relationships between data snippets stored in multiple data sources. Exploiting the architectural paradigm of service-oriented architecture (SOA), the concept establishes an integrated view of complex relationships and supports immediate reactions on suspicious events in the IT infrastructure.

1 Introduction

The collapse of Enron Corporation in late 2001 and WorldCom in 2002, and other large corporations led to a general weakness of the stock markets and the foreign exchange value of the dollar [7]. The ACFE Report of the Nation on Occupational Fraud & Abuse lists 1134 occupational fraud cases, each causing a medium loss of USD 159.000 [1]. In the recent past, Société Générale uncovered “an exceptional fraud” which resulted in a

loss of 4.9bn Euros [19]. The fraud had been committed by a trader in charge of "plain vanilla" hedging on European index futures. In this case the trader was granted privileged access to trading systems allowing him to get unauthorized access to financial resources. The numerous scandals in the US and the European Union and continued weakness of the financial markets have led to a tightening of accounting and auditing regulation. The Sarbanes-Oxley Act of 2002 introduced new rules for the regulation of corporate governance and financial practices [18]. The European response has been a proposal that broadens the scope of the 8th Council Directive on Company Law, which primarily pertains to the approval of statutory auditors in EU member states [4]. The laws aim to ensure the accuracy of data reported in financial statements by improving the internal controls and accuracy of the processes.

Today's most business processes rely on information technology. Therefore, the internal controls must include the IT infrastructure and the IT processes summarized by the term IT Compliance Management [9]. In this context, implementing legal requirements and internal controls including the IT infrastructure is a real challenge [10, 12, 16].

In response to legal requirements, many organizations have already implemented identity management solutions, as well as change and configuration management solutions. However, a significant and important security compliance gap which has not been successfully addressed by compliance management solutions is the detection and prevention of unauthorized access to sensitive data and misuse of privileged access rights.

The ITIL Security Management process [17, 11] describes the structured fitting of security in the management organization based on the Code of Practice for Information Security Management defined by ISO/IEC 27002 [8]. In this context, event monitoring plays an important role with respect to the effective implementation, monitoring, and enforcement of regulatory requirements [6]. Specifically, this holds for event correlation, which is software technology that observes, classifies, and correlates message streams and informs human or technology recipients of system events, enabling the recipients to react in an appropriate way. Despite filtering of security-relevant events and correlation of security incidents there remains a considerable amount of events which must be analyzed by human recipients in a timely manner. Automated provisioning of relevant context information offers a high potential for optimally reducing reaction time, effort and risk.

2 Requirements Analysis

This chapter defines a set of common requirements for an efficient solution for privileged access analysis. The requirements were elicited and recorded during a study conducted in a security operations center responsible for monitoring user activity. The first subchapter introduces a scenario demonstrating the complexity of the problem domain and the weaknesses of current approaches.

2.1 Scenario “Unauthorized Change”

Stephan is security administrator working in a team of security specialists monitoring a world-wide security infrastructure 24 hours a day, 7 days a week. At 6 p.m., Stephan’s security console displays the following security event:

„17:54:03.00: User admin12 logged in with privileged access on host obaserver1.company.com“

Knowing that administrator activity during business hours on an online banking server is unusual, Stephan searches for planned administrative changes in the change management system. Since there is no entry for a scheduled change on the affected server, Stephan must follow a pre-defined process to investigate the suspicious event. Following the process Stephan checks multiple conditions to assess the event:

1. Identify user: Stephan resolves the user name admin12 using the local user identity management system (UID System) and determines the owner of the user name. Using the hostname and the login name, he finds the user “John”.
2. Determine user status and role: Stephan navigates through the human resources directory (HR System) and brokering system. He finds out that the person belongs to the company since 2001, currently working as investment broker with an upper limit of 500.000 EUR.
3. Classify host system: Stephan extracts the complete hostname from the asset management system (Asset Database) and looks up the name in the list of classified systems (CSV-Export). Since the system is listed as a critical asset, Stephan must perform further investigations with high priority.
4. Assess violation: Analyzing the log files of the system (Event Data), Stephan detects a modification of transactions on a brokering application. The application controls the investment activities of the company. The stored transactions will be executed at the end of the business day (6 p.m.) by a batch job which leads Stephan to classify the security event as critical.
5. Gather contact information: Using the human resources directory, Stephan gathers contact details of John and John’s manager but cannot reach any of them. Stephan contacts the security department and gets informed that both persons have left the building.
6. Analyse historical data: The examination of the event records associated with the user and the host reveals the existence of a similar case five months ago. In that case John has logged in to the system without doing any modifications to the system.

The next day reveals that John successfully modified a colleague’s transactions and that the manipulations caused a significant loss for the company. Figure 1 presents an overview of the set of data and their relationships used in our scenario.

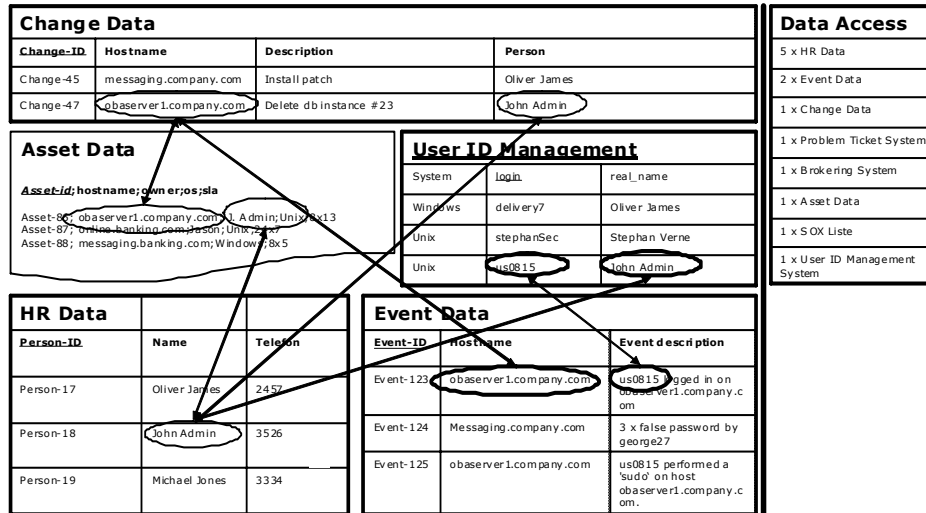


Figure 1: Interrelated data sets of the example scenario

2.2 Requirements for an Efficient Event Assessment

The adequate assessment of suspicious events (security relevant events) requires sophisticated information about the business context which can be derived using data stored in different data sources. Due to the evolutionary character of an organisation's information systems and structures, the necessary data is stored in heterogeneous systems, mostly offering solely proprietary interfaces. Furthermore, data protection policies and laws require limited access to this data. Heterogeneous applications and data silos constrain an efficient linking of data required for the compliance management.

Figure 2 depicts the situation of a security administrator and his need to manually access multiple applications at the same time.

To this end, the scenario in chapter 2.1 and the current information technology infrastructure lead to the following set of requirements for an approach fostering efficient compliance management:

- **Standardized data access:** To enable the automated gathering of context-relevant information and meet the evolutionary nature of the IT landscape, the different information systems must provide standardized interfaces for efficient and effective access to the data.
- **Relationship description:** For the creation of the relevant business context regarding a security event, mechanisms allowing the description of relationships between information entities of different applications and their usage at runtime must be provided.

- Tool support: To allow for an assessment of the security events and the related business context, the security administrators require a dedicated tool presenting the information adequately.

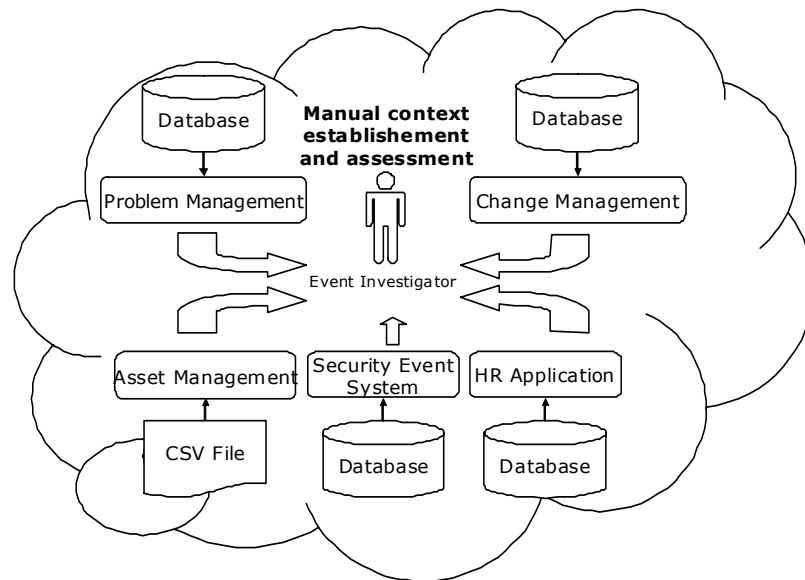


Figure 2: Manual context creation and assessment

3 Concepts for an Efficient Event Assessment

The following sections describe our concepts for the service-oriented context creation in a distributed and heterogeneous system landscape. The first subsection includes a brief architectural overview which helps understand the key aspects of the concept. The following section focuses on a process model for the event assessment as well as creating a homogeneous access layer for data spaces within different heterogeneous systems. The final section introduces mechanisms for descriptions and evaluations of relations and interdependencies between information entities located in different data sources with the aim to create a unified and navigable information space within the problem domain.

3.1 Architectural Overview

Figure 3 depicts an overview of our architectural solution. Within an organization, there are several stakeholders involved in event assessment. Being a complex task and involving multiple information sources, event assessment cannot be done manually. The stakeholders need tool support for efficient and effective event assessment. The tool must integrate data stored in distributed, heterogeneous data sources, combine the data,

and enable the user to recognize the fraudulent actions. As a concept for the integration of heterogeneous data sources, partly within different organizational units, the paradigm of service-oriented architecture (SOA) [3] has been selected, enabling the use of various data for IT Compliance Management processes through a service layer.

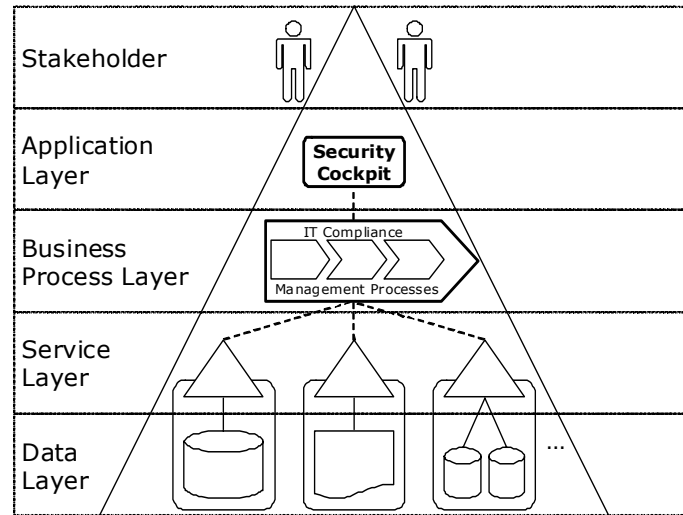


Figure 3: Service-oriented enabling of IT Compliance Management processes

3.2 Service-oriented Context Creation and Assessment

The process for investigating suspicious events in the domain of IT Compliance Management can be divided into four different phases. These phases are data pre-analysis, context creation, context evaluation and resolution. The phase data pre-analysis is initiated by a security event which was classified as suspicious by a dedicated Security Event Management System and needs further investigation. During the context creation phase, information units will be extracted from multiple sources using standardized interfaces. Based on the predefined relationships between data entities within different data sources, the phase context creation addresses the extraction of event-related information from several data sources via standardized interfaces. In this context, common data sources are asset management systems, systems for employee data (HR) as well as change management systems. During the next phase, context evaluation, a support tool (Security Cockpit) provides dedicated views of the context supporting efficient and fast risk estimation. Furthermore an expert system component supports the event investigator by proposing problem solution strategies which are based on former cases and the resolution is conducted in the final phase resolution. Our concept focuses on context creation and context evaluation due to the high potential for reducing manual data collection time and the related error-proneness as depicted in Figure 4.

In order to establish the entire context for security events, information from different heterogeneous and distributed data sources has to be integrated. Our concept leverages a service-oriented architecture making the core enterprise data sources available as standard and reusable web services. Data providers which encapsulate underlying information sources as a service can be used by different consumers. By using the standardized, strongly reusable and evolutionary CRUDS-Interface [14] we achieve a homogeneous service access layer. As a consequence, these services can be embedded and reused in different scenarios [5]. CRUDS is the abbreviation for the methods Create, Read, Update, Delete and Search which are provided by a service in combination with a managed set of specific information or business objects in particular.

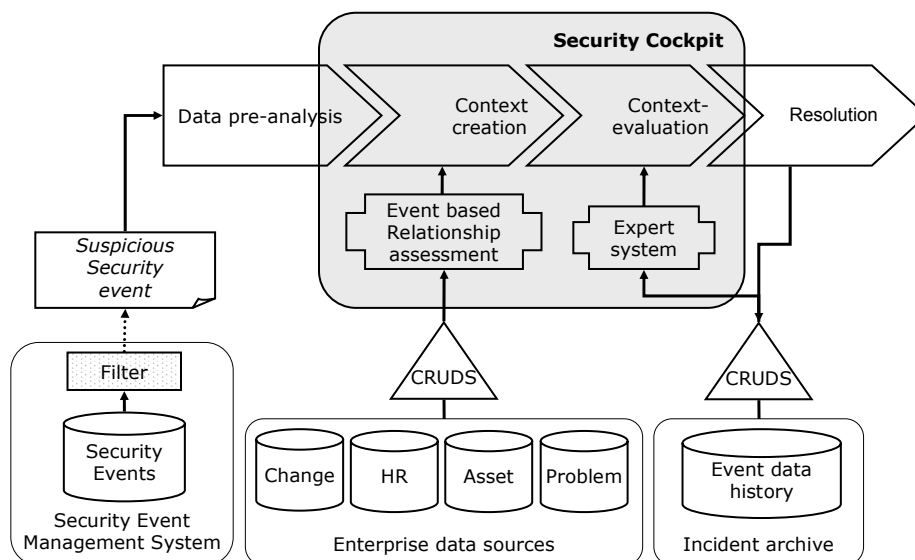


Figure 4: Component model.

Furthermore, extensibility is an important requirement for our approach as the integration of additional data sources allows the improvement of the event assessment by taking more context information into account and the IT landscape may change due to new IT systems. In both cases, domain experts need to identify and analyze the relevant information contained in the additional systems. The information consists of business objects having unique identifiers (e.g. employee, change ticket, hostname) as well as search patterns or single attributes (e.g. isEmployee, ticketStatus). In a next step, the methods of the CRUDS-interface, which have to be realized according to the needs, are identified. This is done with respect to security and data privacy aspects like data minimality [2]. The Read-method supports the request of business objects by name. If different access rights for multiple fine granulated views are needed, facade methods can be implemented. This means an interface will be extended with dedicated methods which uses the core method of the CRUDS-Interface (in this case READ) but only a small portion of the business object will be returned. Thus, the service consumers' needs

can be adequately fulfilled and specific authorization concepts can be implemented. The Search-method supports retrieval of information by defining search parameters. The methods Create, Update, Delete are not used for the context creation but can give improvement in other scenarios (also outside the domain of IT Compliance Management). A trivial scenario for Update would be the closure of a problem ticket after successful evaluation of an event assessment.

3.3 Description and Evaluation of Data Dependencies

The security event itself is the initial starting point for context creation. By evaluating dependencies between the security event and existing data spaces, the use-case specific requirements will be defined. The specific context realization is carried out by requesting the information through the standardized interfaces defined in chapter 3.2. The description of dependencies between different data units in the area of compliance management is based on ontologies. Therefore, objects of existing domain models - such as persons or IT resources (CIM-Ontology [15]) - are applied and extended. The resulting meta model provides the basis for a standardized data space and allows data spaces of single systems to be mapped to an abstract concept. This way, point-to-point linking between information of a system to other used data spaces is not needed anymore. Thus, the specification and categorization of different security related events as well as their linking to concepts within the ontology, allows the abstract definition of the relevant event context for the event assessment.

To realize a concrete event context, the event must be classified according to the schema of the ontology which in turn defines the abstract information needs of that context. With the initial linking of the data spaces of the existing IT system to the ontology, these needs can be satisfied by accessing the corresponding systems via the interfaces defined in chapter 3.2. These interfaces allow standardized operations on selected sets of business objects and the creation of the event context.

To enrich the event context with further information, the analysis of the attributes of the event and the gathered context is promising. As shown in the example scenario in section 2.1, these attributes often contain identifiers of related objects residing within different data sources. To identify and use these additional data dependencies to extend the event context, a catalogue of search patterns can be created and applied.

4 Implementation

The first prototype to support service-oriented event assessment in a distributed and heterogeneous environment consists of five Web services, implementing the access method to relevant data sources. Based on the business context, we defined the business objects, the data dependencies and implemented the CRUDS-interfaces to access the data sources. In particular, the data of the following systems were considered: A Security Event Management System (IBM Tivoli Risk Manager), a Change Ticket System (Peregrine Service Center), a Problem Ticket System (BMC Remedy), an Asset

Management System based on an IBM DB2 database and a Human Resources directory service.

The Security Cockpit, implemented as a Java Client application, offers a user friendly interface for exploiting the established event context using Web service connectors. One of the connectors loads the IT security events from the Security Event Management System. Based on the data relationships defined in an XML configuration file the requirements for the investigation context are determined (relation based event assessment) and the context is created by calling the Web services.

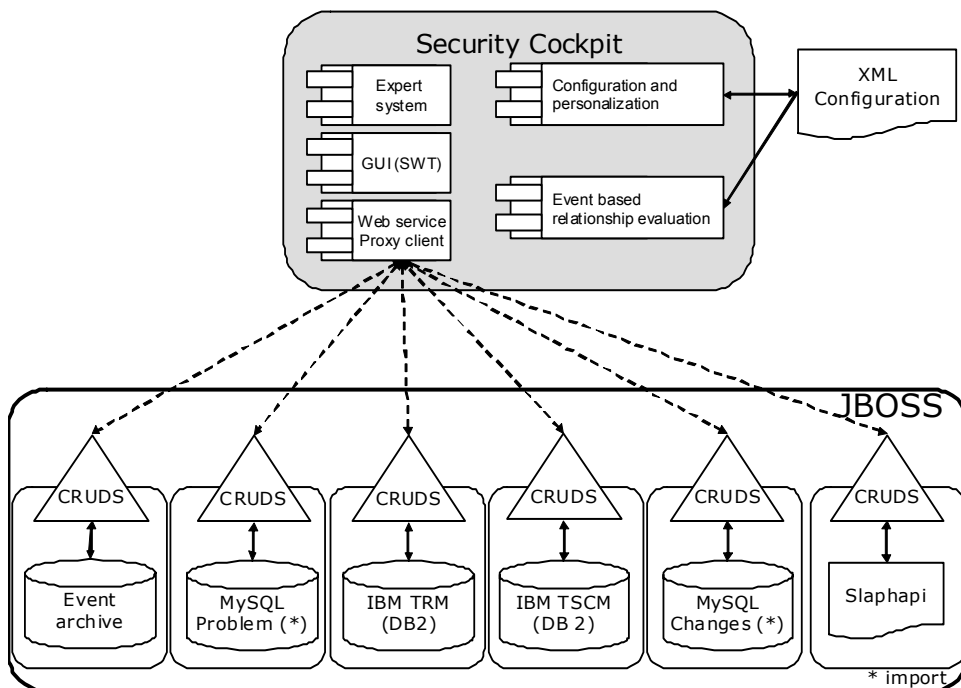


Figure 5: Architectural overview of the prototype.

The Security Cockpit presents the newly created, event- and user role-specific context. The user starts assessing the event in this highly adapted environment in an efficient way. The specific view depends on the user's tool configuration and role [13]. An expert system component supports the user by showing him problem solving strategies and results of similar cases which have been investigated in the past. The expert system exploits a classification system for security related events. Having completed the event assessment, the user documents the solution and which of the historic solutions helped to resolve the current solution. Thus, the expert system improves the knowledge base for future event investigations.

In a pilot project the prototype was tested and has shown big potential for our concept. New data sources were integrated and their information objects linked with each other.

The standardized CRUDS-interface simplified the integration of different data sources and context creation. In summary, the prototype facilitates and improves the investigation of security events with respect to effort and complexity.

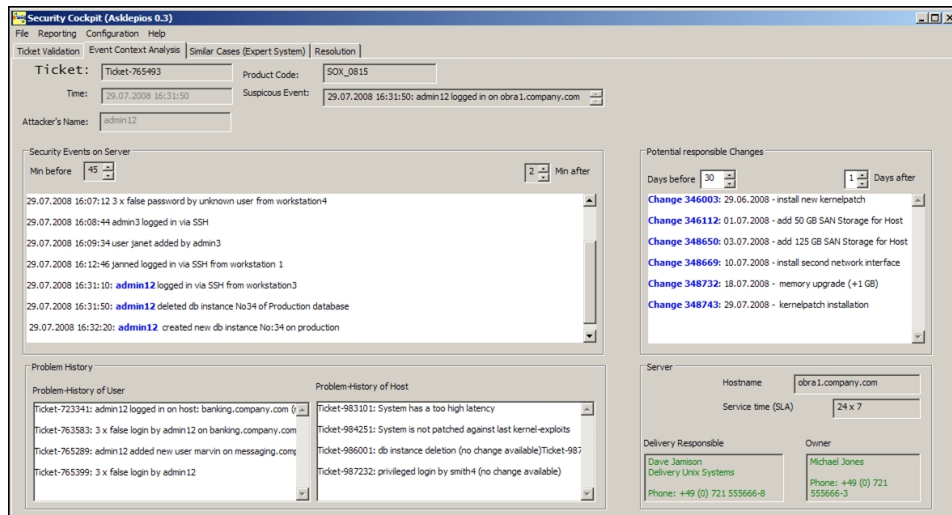


Figure 6: Security Cockpit.

5 Summary and Outline

The assessment of IT security events is a mandatory sub process in the area of IT Compliance Management and helps to reduce the risk of unauthorized access to IT resources. The investigation of security events in an adequate and efficient way requires creating and evaluating the entire context of the business case. Due to the heterogeneity and complexity of today's IT systems and dependencies between information sources this task is challenging.

Considering this background, this article presents a concept which allows the efficient linking of information between diverse data sources to support IT Compliance Management processes. Core aspects are the definition of standardized interfaces which allow access to data and taking into consideration relations between data units of different data spaces. During a pilot project, a prototype implementing the concept showed its potential in a real world environment and the transferability of the approach to other scenarios.

As a next step we will investigate whether the proposed concept can be exploited to further automate the real-time evaluation of security events. Moreover, an evaluation of the approach regarding the response time to suspicious events as well as the quality of the problem resolution will be conducted. Furthermore, future work will address aspects

of security and privacy related questions as well as reducing cognitive load of users working with the prototype.

Literature

- [1] Association of Certified Fraud Examiners, Report to the Nation on occupational fraud, 2006.
- [2] Bundesministerium der Justiz, Bundesdatenschutz Gesetz 1990 (Neufassung 2003).
- [3] ENDREI M., ANG J., ARSANJANI A. et al, Patterns: Service-Oriented Architecture and Web Services, IBM Redbooks, IBM, 2004.
- [4] Europäischer Rat, 8. EU Richtlinie 84/253/EWG, 1984, URL: <http://eurlex.europa.eu/LexUriServ/LexUriServ.do?uri=CELEX:31984L0253:DE:HTML> (23.05.2008).
- [5] FREUDENSTEIN, P., MAJER F., MAURER, A., RIED D., JULING W., Wiederverwendungsorientierte Dienste für Universitäten 2007, in: Proceedings of INFORMATIK 2007, 37. Jahrestagung der Gesellschaft für Informatik, Band 1, p. 497-501. Bremen, Germany, 2007.
- [6] GIBLIN C. J., MUELLER S., PFITZMANN B., From Regulatory Policies to Event Monitoring Rules: Towards Model-Driven Compliance Automation, IBM Research Paper, 2007.
- [7] GOLTSCHE W., COBIT kompakt und verständlich, Friedr. Vieweg & Sohn Verlag GWV Fachverlage GmbH, Wiesbaden, 2006.
- [8] HARICH Thomas W., ISO/IEC 27001 Eine praxisorientierte Einführung, Mosenstein und Vannerdat, 2008.
- [9] JOHANNSEN, W. und GOEKEN, M., Referenzmodelle für IT-Governance. Strategische Effektivität und Effizienz mit COBIT, ITIL & Co, dpunkt.Verlag, Heidelberg, 2007.
- [10] KAARST-BROWN, M. L., KELLY, S., IT Governance and Sarbanes-Oxley: The latest sales pitch or real challenges for the IT Function? In Proceedings of the 38th Hawaii International Conference on System Sciences, 2005.
- [11] KÖHLER, P. T., ITIL: Das IT-Servicemanagement Framework, Springer, Berlin, 2006.
- [12] LESKELA L., MOGULL R., LOGAN D., Sarbanes-Oxley Readiness Study: Preliminary Results, in: Gartner Study, 2004.
- [13] MAJER F., FREUDENSTEIN P., NUSSBAUMER M., Roadmap towards Lifecycle Support for Highly Distributed Web-based Systems, in: Proceedings of 8th International Conference on Web Engineering (ICWE2008), New York, USA, 2008.
- [14] NUSSBAUMER, M., Entwicklung und Evolution dienstorientierter Anwendungen im Web Engineering, Universitätsverlag Karlsruhe, 2007.
- [15] QUIROLGICO, S., ASSIS, P., WESTERINEN, A. et al., Toward a Formal Common Information Model Ontology, in Web Information Systems - WISE 2004 Workshop, Australia, 2004.
- [16] TAYLOR, H., Why Sarbanes-Oxley and Service-Oriented Architecture may be the best thing that ever happened to you, Wiley Publishing Inc, Indianapolis, Indiana, 2006.
- [17] United Kingdom's Office of Government Commerce, Information Technology Infrastructure Library.
- [18] U.S. Securities and Exchange Commission, Sarbanes-Oxley Act, S. 107–204, 2002.
- [19] WILSON, T., Societe General: How Did It Happen, in: Dark Reading - Risky Business, 28.01.2008. URL: http://www.darkreading.com/document.asp?doc_id=144313 (21.07.2008).