

Automated Risk and Utility Management

Andreas Ekelhart and Thomas Neubauer
Secure Business Austria
Vienna, Austria
{ekelhart, neubauer}@securityresearch.ac.at

Stefan Fenz
Vienna University of Technology
Vienna, Austria
fenz@ifs.tuwien.ac.at

Abstract

Information security breaches pose major threats to the reliable execution of corporate strategies and may have negative effects on business value. Information security risk management (ISRM) provides an effective approach for assessing, mitigating, and evaluating information security risks. Existing ISRM approaches are highly accepted but demand very detailed knowledge about the IT security domain and the actual company environment. This paper presents the AURUM prototype that supports decision makers in selecting security measures according to organization-specific technical and economical requirements.

1 Introduction

Existing information security risk management approaches (e.g. CRAMM [7], NIST SP 800-30 [14], CORAS [8], OCTAVE [1], EBIOS [3], and recently ISO 27005 [9]) and their tool implementations are requiring, especially in the risk assessment and risk mitigation phase, very detailed knowledge about the IT security domain and the actual company environment. Up to that point in time, organizations mostly fall back on best-practices, information security standards, or domain experts when conducting the risk assessment and mitigation phases. Several problems arise with these approaches: (1) existing tools often require manual and unguided inventory of the organizations resources and do not support for an automatic or semi-automatic inventory of IT resources, (2) while the underlying knowledge model has to allow the incorporation of general information security domain knowledge and specific knowledge about the considered organization to allow the determination of its current information security compliance, existing tools often lack essential information and connections, (3) existing approaches often do not provide adequate results because they aim at evaluating candidates through a single (aggregated) indicator criterion, and neglect the consideration of

benefits. In order to obtain an investment's adequacy and efficiency multiple objectives should be considered (e.g., to allow to determine the (business) value of IT security investments) (cf. [10, 4]), (4) management decision makers should not be confronted with a single "optimized" solution when having to decide on their security infrastructure but should rather be invited to analyze and explore different scenarios and, thus, to participate in and to control the investment selection process. Therefore, tools should support intuitive and interactive decision support methods that are prepared in a way business people can relate to (cf. [2]).

In order to address these reservations and demands outlined above, this paper presents our AURUM¹ prototype and shows its salient features compared to existing implementations. AURUM is based on previous work (cf. [5, 6, 11] for the concept of the security ontology and [12] for interactive decision support).

2 Tool Description

AURUM was designed to minimize the interaction necessary between user and system and to provide decision makers with an intuitive solution that can be used without extensive knowledge about the information security domain. However, the solution is also capable of providing expert users with detailed information on different levels of granularity. Figure 1 shows the high-level architecture of AURUM. According to the requirements, the security ontology provides each AURUM module with knowledge regarding the general information security domain and the specific security status of the considered organization. The AURUM - Inventory module incorporates interfaces to several third-party inventory and network scanning solutions to support the system characterization phase. The AURUM - Bayes module generates and modifies the Bayesian network for the threat probability determination. The AURUM - Risk module is the central module which uses the Bayes module and connects to the security ontology web service

¹ derived from *AURUM: Automated Risk and Utility Management*




Figure 1. AURUM architecture

to calculate the risk levels for assets. This is done by gathering probability values from the Bayesian network and multiplying those with the defined impact values, stored in the ontology. Furthermore, all changes to the ontological data repository are handled from this point.

Figure 2 demonstrates the schematic layout of the working area. Section 1 summarizes information on (a) the business processes and its dependence on assets, and (b) the assets's physical locations in the organization. Section 2 - the main area - provides the decision maker with (a) detailed information about the selected asset, (b) a graphical representation of the selected business process together with the assets needed for the execution of the selected business process, and (c) the graphical representation of the physical location model together with assets. Information provided in Section 2 depends on the selection the decision maker made in Section 1 (the same holds analogously for the dependence between Section 2 and 3). Section 3 displays (a) the risk level for the selected asset, (b) a list of threats and their calculated probabilities, and (c) implemented and not implemented controls with their calculated effectiveness figures.

2.1 System Characterization

The basic version of the security ontology provides the user with an extensive set of initial information on the security infrastructure that is usually needed in organizations. An organization that decides on using the security ontology as a basis for information security risk management has to initialize the ontology once with organization specific information. This phase is supported by the AURUM-Inventory module (for a detailed description see [6]) which is able to capture the data of software and IT-related infrastructure elements automatically (operating system, IP address, patch level, etc.). This enables us to enhance the efficiency of the system characterization step significantly, since the inventory of IT-related infrastructure elements is one of the most labor-intensive steps. Collecting such detailed device data enables, in the case of software-related threats (e.g. malware or errors in standard software), the mapping of software vulnerabilities on the current IT infrastructure in order to visualize threatened systems immediately. Besides

the IT component inventory AURUM provides the following two options for updating information on the organization's assets: **(1) Process Model:** AURUM allows to use business process models as a basis for identifying corporate risks. Section 1 provides an overview of the business processes selected for the specific risk assessment. By selecting one of these processes, the tool provides a graphical representation of the business process in Section 2. Additionally, all assets needed for executing this process are displayed. When the user selects one of those assets, the tool provides further information on threats, vulnerabilities, risk levels and potential controls (cf. Section 2.2) in Section 3 (cf. Figure 2). In order to support organizations already using business process management tools, features to import business processes under consideration as well as the mapping between services and processes from such tools (e.g. Adonis or ARIS) are planned, **(2) Physical Model:** Based on the data stored in the security ontology, AURUM allows to generate a building model including the location of all assets. In return this model can be used for adding assets to the ontology and, of course, for simulating different scenarios (e.g. for identifying the optimal location for valuable assets). Section 1 provides an overview of the corporate assets and their location in the building. By selecting one of those assets, the tool displays a graphical representation of the asset location and connected business processes in Section 2. In analogy to the process model, clicking on one of the assets provides the user with further information on threats (and connected vulnerabilities), risk levels and potential controls (cf. Section 2.2) in Section 3 (cf. Figure 2).

2.2 Threat and Vulnerability Assessment

The threat tree (located in Section 3) shows the potential threats to the selected asset, including a priori threat probabilities and organization-specific attacker profiles. By selecting a threat from the tree representation, valuable information such as a threat description in natural language is displayed. Furthermore, affected security attributes (confidentiality, integrity, and availability) are provided. In addition, a threat can be a consequence of other threats (e.g. unauthorized physical access can be the result of missing key management) and can itself potentate other threats (e.g. unauthorized physical access gives rise to data disclosure). Note, that this step only shows those threats to the risk manager, which are relevant for the organization and the considered asset. The same step analyzes potential vulnerabilities which are present in the defined scenario. This includes the consideration of vulnerabilities in the field of (1) management security (e.g. no assignment of responsibilities, no risk assessment, etc.), (2) operational security (e.g. no external data distribution and labeling, no humidity control, etc.), and (3) technical security (e.g. no cryptography solu-

tions in use, no intrusion detection in place). For each threat highly granular vulnerabilities, which a threat could exploit, have been defined and modeled in the ontology. A description of each vulnerability in natural language complements the vulnerability presentation. For each of the vulnerabilities a mitigation control is assigned, thus implementing a control closes a vulnerability. To enhance the understanding, each control is enriched by a natural language description. With these functions in place, a user knows exactly how to protect his organization from specific threats: mitigating vulnerabilities by implementing recommended controls.

Up to the current point the decision maker is aware of the considered system, potential threats and corresponding vulnerabilities, which allow threats to become effective. The control analysis step determines which controls are already in place and which controls exist to mitigate the probability that a threat exploits a certain vulnerability (e.g. the *unauthorized physical access* threat exploits the vulnerability *no access regulation control* which could be mitigated by the installation of an entry checkpoint, security guard, or an access system).

To facilitate the aspect of automatic compliance checks regarding our defined mitigation controls, each control further incorporates formal implementation descriptions. The implementation area in Section 3 shows the actual implementation measures for a control. The underlying formal control descriptions can be executed as rules against the organizations concrete modeled environment to identify which assets are in compliance. The compliant assets can be displayed using the location model (cf. Section 2.1). Most often it is not sufficient just to know if a certain control is implemented or not within a given context. The most important thing to know is, if the implemented control is appropriate or not to achieve the acceptable risk level of the considered asset. Since the risk level determination requires besides the probability component an impact component, each asset is rated regarding to its importance for the organization's mission regarding confidentiality, integrity, and availability.

2.3 Risk Determination

This phase comprises the probability determination of threats exploiting certain vulnerabilities in the given system. The subsequent impact analysis determines the impact on the organization's ability to perform its mission, if a threat should successfully exploit a certain vulnerability. By combining the threat probability with the magnitude of the impact the organization is able to determine the risk level and thus to plan the necessary actions. In contrast to other approaches (cf. [14, 13]), AURUM focuses on an automated support utilizing the developed knowledge

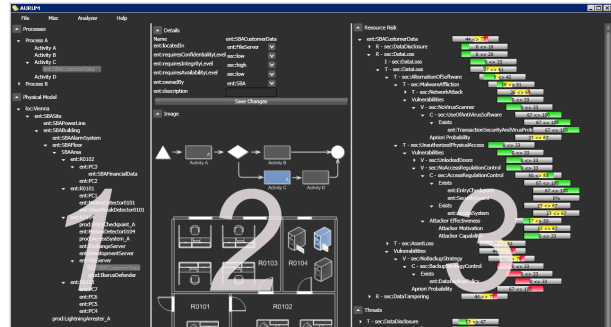


Figure 2. AURUM - risk determination

base and the defined relationships. Figure 2 depicts the AURUM interface which supports the risk manager at the risk determination phase. The exemplary threat probability determination is conducted for the *ent:SBACustomerData* element, representing the customer data of SBA. A Bayesian network for the threat probability determination is generated based on the security ontology. Tree views have been used to visualize this complex network in a comprehensible form (see the upper right section of Figure 2). Since *ent:SBACustomerData* represents a data instance it faces three different risks → data disclosure, data loss, and data tampering. Each risk is the product of the asset's impact rating and the probability of the corresponding threat. The probability of a threat is determined by the probabilities of its predecessor threats, the exploitation probability of corresponding vulnerabilities, the effectiveness of existing control implementations, the a priori threat probability, and the attacker's effectiveness. Since the input values for existing control implementations, a priori probabilities, and the attacker's effectiveness depend on the considered asset this section examines how AURUM determines the input values for *ent:SBACustomerData* in the context of the data loss threat.

The resource risk tree in Figure 2 shows three different types of vulnerabilities: (1) technical - no virus scanner, (2) physical - no access regulation control, and (3) organizational - no backup strategy.

To determine the exploitation probability of the no virus scanner vulnerability, reasoning algorithms check if an instance of the transaction security and virus protection software concept is installed on the file server on which the SBA customer data is located. The physical model in the left section of the user interface shows that the virus scanner Ikarus Defender is installed on the file server. Thus, the virus scanner protects the file server and the SBA customer data from malware affliction and closes the no virus scanner vulnerability. The green bar of the transaction security and virus protection node indicates that the virus scanner Ikarus Defender is highly effective at mitigating the corresponding

vulnerability.

The no access regulation control vulnerability exploitation probability is determined by the effectiveness of an access system, an entry checkpoint or security guard, and the attacker's effectiveness. Because the no access regulation control vulnerability is on the section level (*sec:vulnerabilityOn*) reasoning algorithms check if the required controls are implemented in section *ent:R0104* (the physical location of the file server). As the physical model in the left area of Figure 2 shows, the highly effective *EntryCheckpoint_A* and the medium effective *AccessSystem_A* are located in server room *ent:R0104*. Since the attacker is rated with a 17 - 50% effectiveness and the control combinations are rated with a 33 - 67% and 67 - 100% effectiveness, the no access regulation vulnerability is mitigated to a 0 - 33% exploitation probability.

The no backup strategy vulnerability represents an organizational vulnerability and therefore the existence of an appropriate policy determines the exploitation probability. Since policies are implemented at the organization level, reasoning algorithms check if a data backup policy is implemented for the organization which owns the file server. SBA implemented a low effective data backup policy which covers amongst others the file server's data. Due to the high a priori probability of the data loss threat, the low effective SBA data backup policy mitigates the exploitation probability of the no backup strategy vulnerability to 44 - 78%.

While the resource risk tree comprehensively represents the threat probability calculation, it may confuse the user with the complex representation. Therefore, the threat view in the middle right section of the user interface represents relevant threats and their probabilities. If the user wants to decrease a certain threat probability, controls depicted in the lower right section have to be implemented. In the case of the SBA customer data, AURUM suggests amongst others to implement an anti theft device for the file server, a safety door for the server room, and a fire extinguishing system in the server room.

Due to the hierarchy of the Bayesian threat probability network, control implementations on different levels affect the final threat probability with different intensities. Therefore, AURUM supports the user with weights for each control implementation in a specific threat context. In the case of the data loss threat the following weights have been determined by AURUM: data backup policy (0.1666), locked doors policy (0.0763), access system (0.0509), entry checkpoint or security guard (0.0509), transaction security and virus protection software (0.0416), anti theft device (0.0277), lightning arrester (0.0138), safety door (0.0046), and fire extinguisher (0.0046). With this results on hand the user will rather implement a sound data backup policy instead of investing a lot of money in expensive safety doors to protect the customer data of SBA.

As an example the following scenario is assumed: amongst other control implementations the SBA customer data is protected by a low effective safety door and a low effective data backup policy which results in a data loss probability of 31 - 64%. To lower the data loss probability for the valuable customer data the risk manager is forced to implement more effective control implementations. The aforementioned weights of relevant control implementations revealed that the risk manager should prefer the implementation of a sound data backup policy instead of investing in a safety door. The Bayesian threat probability determination confirms these recommendations. If a high effective safety door and a low effective data backup policy is entered into the Bayesian network the data loss threat probability decreases by 1% to 30 - 63%. If a low effective safety door but a high effective data backup policy is in place the probability of a data loss decreases by 15% to 16 - 49%. The following combination confirms the minor influence of the safety door on the data loss probability: setting besides the data backup policy also the safety door to a high effectiveness results also in a threat probability of 16 - 49%.

2.4 Control Evaluation and Implementation

This step involves the identification and evaluation of controls or combinations thereof regarding their cost/benefit ratio. As a result, controls which are suitable to mitigate the risk to an acceptable level at the lowest possible costs can be incorporated in the control implementation plan. At this point management knows which risks are not acceptable for the organization and therefore, measures have to be taken (in terms of controls which could mitigate or eliminate the identified risks). For each vulnerability, appropriate controls are identified, taken from best practice standards such as the German Baseline Protection Manual. Offering these controls equips the decision makers with effective countermeasures to lower the risk level and thereby protect their business. As controls only provide information on the class of safeguards that should be used (e.g., Fire Extinguisher), instances must be identified that are finally implemented into the organization. Therefore, potential control implementations are evaluated according to a defined resource- and benefit categories (e.g., costs, effectiveness, reliability) in order to precisely target the company's specific business needs in line with economic demands. This analysis does not only consider cost and benefit in monetary terms but includes non-financial objectives. All potential controls identified are rated against the chosen criteria using data from the security ontology. Using the potential controls and their ratings in each category as input, all Pareto-efficient combinations of safeguards are determined (i.e., there is no other solution with equally good or better values in all objectives

and a strictly better value in at least one objective). All solutions taken into consideration have to be feasible with respect to two sets of constraints: The first set relates to limited resources (e.g., development costs or maintenance costs). The second set ensures that at most a maximum – or at least a minimum – number of safeguards from given subsets (e.g., from a certain type of safeguards such as fire-walls) is included in the feasible solutions.

AURUM provides an interactive interface that offers the decision maker information on the specific selection problem while the system ensures that the final solution will be an efficient one. The decision makers learn about the consequences of their decisions and get information on the gap (in each category) between the existing solution and the potential solutions. We are using a search based procedure, which starts from an efficient portfolio and allows the decision maker to iteratively “move” in solution space towards more attractive alternatives until no “better” portfolio can be found. Our approach is based on interactive modifications of lower and upper bounds for one or more objectives. The tool starts with displaying bars representing resource and benefit categories that are assigned with units (cf. the *Analyzer* Section in Figure 3). Two movable horizontal lines with small arrows at one side represent lower and upper bounds and are intended to restrict the set of remaining solutions in a step-by-step manner (e.g., by raising the minimum bound in one of the objectives) or for expanding it (e.g., by once again relaxing some bounds) according to the decision makers’ preferences. In all of these cases, the system provides immediate feedback about the consequences of such choices in terms of the remaining alternatives. Currently we have added the following categories in the AURUM tool: effectiveness, weight (influence on the corresponding threat probability), initial costs and running costs. Continuing our example with the customer data of SBA, we identify the threat with the highest threat probability, namely *T - sec:DataTampering* with a probability of 44 - 77%. Management should address this threat first to reduce the overall risk, thus, we start the calculation of existing security portfolios (cf. Figure 3). For each control class (e.g., *Anti-Virus Software*) concrete products can be added to the ontological knowledge base. An initial set of products has been incorporated already, but can be easily adapted and extended by organizations using the tool. The reader should note, that for demonstration purposes example instances such as *EntryCheckpoint A* and *LockedDoorPolicy A* have been added. The AURUM tool initially finds 150 possible solutions by combining concrete control implementations. As a next step, the decision maker starts to apply his given restrictions/preferences: In our example we have restricted funds and thus set the maximum initial costs to 9.800 and maximum running costs to 1.000 (compare the orange InitialCosts and RunningCosts bars and the upper red

lines indicating our financial preferences). In addition, we demand at least medium effectiveness of the portfolios to be considered (compare the green Effectiveness bars and the lower red line). Each portfolio is represented by a vertical bar; as can be seen, only 6 portfolios are remaining satisfying our preferences. In Section 3, we are provided with a solution list, displaying detailed information about the remaining solutions, comprising the exact figures of the category values and the candidates the portfolio implements. For example our selected portfolio with the ID 8332, requires an intrusion and detection system software of type C, the anti-virus software IkarusDefender, a locked doors policy of type A and an entry checkpoint of type A. This portfolio offers the highest effectiveness following our restrictions, but has higher initial- and running costs than other solutions.

In further iterations, the decision maker continues playing with minimum and maximum bounds and by doing so can learn about the consequences of his/her decisions. After several cycles of restricting and once again expanding the opportunity set, the decision maker will finally end up with a solution alternative that offers an individually satisfying compromise between the relevant objectives. Note that he does not need to explicitly specify weights for objectives nor to specify the form of his/her preference function or to state how much one solution is better than another during any stage of the whole procedure. Instead, ample information on the specific selection problem is provided to him and the system ensures that the final solution will be an optimal (i.e., Pareto-efficient) one, with no other feasible solution available that is “better” from an objective point of view.

3 Conclusion

This paper presented the AURUM tool which provides compared to existing information security risk management support approaches the following benefits: (1) the ontological information security knowledge base ensures that the information security knowledge is provided in a consistent and comprehensive way to the risk manager, (2) modeling the organization’s resources within our ontological framework ensures that resources are modeled in a consistent way, (3) the incorporation of existing best-practice guidelines and information security standards ensures that only widely accepted information security knowledge is used for threat/vulnerability identification and control recommendations, (4) the proposed Bayesian threat probability determination ensures that the threat probability determination is based on a more objective level, (5) threat impacts can be automatically calculated after resources have been rated initially, (6) controls to reduce risks to an acceptable level are offered automatically, (7) the use of interactive decision support allows decision makers (e.g., the risk manager) to

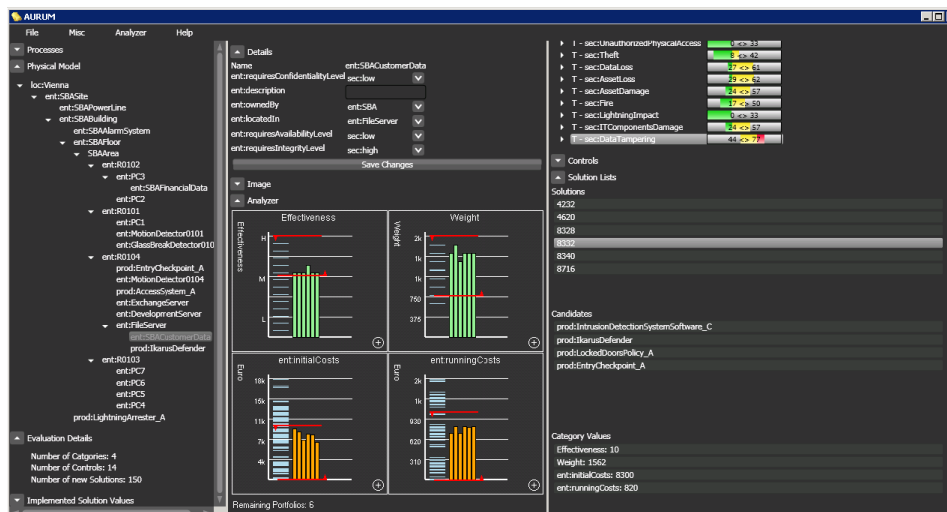


Figure 3. AURUM control evaluation and selection

investigate various scenarios and, thus, to learn about the characteristics of the underlying problem, while the system guarantees that only efficient solution can be selected, and (8) by considering multiple objectives and providing a gap analysis we support decision makers in getting a much better “feeling” for the problem in terms of what can be achieved in some objectives at what “price” in terms of opportunity costs in other objectives.

4 Acknowledgments

This work was supported by grants of the Austrian Government’s FIT-IT Research Initiative on Trust in IT Systems under the contract 813701 and was performed at the Research Center Secure Business Austria funded by the Federal Ministry of Economics and Labor of the Republic of Austria (BMWA) and the City of Vienna.

References

- [1] C. Alberts, A. Dorofee, J. Stevens, and C. Woody. Introduction to the OCTAVE approach. Technical report, Carnegie Mellon - Software Engineering Institute, Pittsburgh, PA 15213-3890, August 2003.
- [2] C. Alves and A. Finkelstein. Investigating conflicts in COTS decisionmaking. *International Journal of Software Engineering and Knowledge Engineering*, 13(5):473–495, 2003.
- [3] DCSSI. EBIOS - Section 2 - Approach. General Secretariat of National Defence Central Information Systems Security Division (DCSSI), February 2004.
- [4] S. Dekleva. Justifying Investments in IT. *Journal of Information Technology Management*, 16(3), 2005.
- [5] A. Ekelhart, S. Fenz, M. Klemen, and E. Weippl. Security ontologies: Improving quantitative risk analysis. In *Proceedings of the 40th Hawaii International Conference on*

System Sciences, HICSS2007, pages 156–162, Los Alamitos, CA, USA, January 2007. IEEE Computer Society. 0-7695-2755-8.

- [6] A. Ekelhart, S. Fenz, T. Neubauer, and E. Weippl. Formal threat descriptions for enhancing governmental risk assessment. In *Proceedings of the First ICEGOV*, volume 232 of *ACM International Conference Proceeding Series*, pages 40–43, New York, NY, USA, January 2007. ACM. 978-1-59593-822-0.
- [7] B. Farquhar. One approach to risk assessment. *Computers and Security*, 10(10):21–23, February 1991.
- [8] R. Fredriksen, M. Kristiansen, B. A. Gran, K. Stolen, T. A. Opperud, and T. Dimitrakos. The coras framework for a model-based risk management process. In *SAFECOMP ’02: Proceedings of the 21st International Conference on Computer Safety, Reliability and Security*, pages 94–105, London, UK, 2002. Springer-Verlag.
- [9] ISO/IEC. ISO/IEC 27005:2007, Information technology - Security techniques - Information security risk management, November 2007.
- [10] V. S. Lai, R. P. Trueblood, and B. K. Wong. Software selection: A case study of the application of the analytical hierarchical process to the selection of a multimedia authoring system. *Information & Management*, 36, 1999.
- [11] T. Neubauer, A. Ekelhart, and S. Fenz. Interactive selection of iso 27001 controls under multiple objectives. In *Proceedings of the IFIPSec2008*, volume 278/2008, pages 477–492, Boston, July 2008. Springer.
- [12] T. Neubauer and C. Stummer. Extending business process management to determine efficient it investments. In *Proceedings of the SAC2007*, pages 1250–1256, 2007.
- [13] M. Stallinger. *IT-Governance im Kontext Risikomanagement*. PhD thesis, Johannes Kepler Universitt Linz, 2007.
- [14] G. Stoneburner, A. Goguen, and A. Feringa. Risk management guide for information technology systems. NIST Special Publication 800-30, National Institute of Standards and Technology (NIST), Gaithersburg, MD 20899-8930, July 2002.