# AURUM: A Framework for Information Security Risk Management

Andreas Ekelhart
Secure Business Austria
Vienna, Austria
Email: ekelhart@securityresearch.ac.at

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

Thomas Neubauer
Secure Business Austria
Vienna, Austria
Email: neubauer@securityresearch.ac.at

***Abstract*—As companies are increasingly exposed to a variety of information security threats, they are permanently forced to pay attention to security issues. Risk management provides an effective approach for measuring the security through risk assessment, risk mitigation and evaluation. Existing risk management approaches are highly accepted but demand very detailed knowledge about the IT security domain and the actual company environment. This paper presents AURUM - a new methodology for supporting the NIST SP 800-30 risk management standard - and provides a comparison with the GSTool and CRISAM in order to highlight the benefits decision makers may expect when using AURUM.**

## I. Introduction

Security breaches pose major threats to the reliable execution of corporate strategies and may have negative effects on business value, e.g., on profit, shareholder value, or reputation (cf. [1] [2] [3]). As a consequence, companies are steadily increasing the amount of resources for protecting corporate assets. For example, total global revenue for security products and service vendors amounted to $21.1 billion through 2005. From 1999 to 2000, the number of organizations spending more than $ 1 million annually on security nearly doubled, representing 12% of all organizations in 1999 to 23% in 2000 (cf. [4]). Although companies consider security as one of the most important issues on their agenda, many companies are not aware how much they spend on security and if their investments in security are effective. Risk management is a crucial element for ensuring long-term business success because it provides an effective approach for measuring the security through the identification and valuation of assets, threats, and vulnerabilities and offers methods for the risk assessment, risk mitigation and evaluation.

However, while existing approaches (e.g. CRAMM [5], NIST SP 800-30 [6], CORAS [7], OCTAVE [8], EBIOS [9], and recently ISO 27005 [10]) for implementing an adequate risk management strategy are highly accepted within the community they 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) best practice guidelines such as the German IT Grundschutz Manual [11] or the French EBIOS standard [9] provide excellent knowledge about potential threats, vulnerabilities, and countermeasures, but without a domain expert the organization is usually unable to consider all the complex relationships between relevant IT security concepts, which results in a non-holistic IT security approach endangering the organization in performing its mission [12] [13] [14] [15], (2) to check which concrete infrastructure elements are endangered by certain threats the organization has to manually map the knowledge from best-practice guidelines to their actual infrastructure [16], (3) especially information security standards such as ISO 27001 [17] are stating only very abstract implementation suggestions for risk mitigation; concrete countermeasures or combinations thereof are mostly missing [18], (4) the determination of threat probabilities is mostly based on subjective perceptions, instead of objective evaluation [19] [13] [18], and (5) retaining a domain expert for the entire process is a very expensive but effective way for ensuring business continuity; we have to keep in mind that these experts act as a single point of failure and in most cases the organization is not able to compare their decisions with a reference model.

In order to address these reservations and demands outlined above, this paper presents a methodology for supporting the entire NIST SP 800-30 risk management standard [6]. Our prototype named AURUM[1] is based on previous work (cf. [20]–[23] for the concept of the security ontology and [24]–[26] for interactive decision support). Beside introducing this new risk management approach, we provide a comparison of AURUM with CRISAM[2] and the GSTool[3] and itemize for each step how our approach performs in terms of usability, time exposure for conducting the entire risk management process, completeness of the threat/vulnerability identification and the control recommendations, and the granularity of control implementation suggestions.

The entire NIST SP 800-30 risk management process (cf. [6]) is subdivided into three main processes, namely (1) risk assessment, (2) risk mitigation, and (3) risk evaluation. The

[1]derived from *AUtomated Risk and Utility Management* (according to http://wordnet.princeton.edu we define Utility as a measure that is to be maximized in any situation involving choice)

[2]CRISAM: www.crisam.net/, last access: 1. September 2008

[3]GSTool: www.bsi.bund.de/gstool/index.htm, last access: 1. September 2008

978-0-7695-3450-3/09 $25.00 © 2009 IEEE

risk assessment process identifies potential risks and their impacts, in order to recommend preventive and risk-reducing countermeasures. In the risk mitigation process the identified risks are prioritized and adequate preventive countermeasures are implemented and maintained. After the countermeasure implementation, a continual evaluation process determines whether the implemented risk-reducing countermeasures are decreasing the risk to an acceptable level or whether further controls are required. The following subsections will illuminate the subprocesses in detail and show how AURUM supports the individual steps.

## II. System Characterization

Due to the fact that the NIST SP 800-30 risk management methodology (cf. [6]) is concerned with assessing risks for IT systems, the first step of the methodology requires the definition of the system boundaries. Since (1) questionnaires, (2) on-site interviews, (3) document reviews, and (4) automated scanning tools are irreplaceable for conducting holistic risk assessment, we concentrate our effort on improving the efficiency of automated scanning tools and questionnaires by combining them with our ontological framework.

What resources and information are used and/or required by the defined system? An answer to this question requires a systematic inventory of hardware, software, existing physical countermeasures, system interfaces, data, information, and persons who support or use the IT system. The following techniques are used to gather the required information: (1) questionnaires, (2) on-site interviews, (3) document reviews, and (4) automated scanning tools. Other risk management approaches and information security standards such as [17], [9], or [11] propose similar system characterization approaches and information gathering techniques.

Since the security ontology already provides a highly granular physical infrastructure model, we just have to build a prototypical graphical user interface (GUI) on top of that model to make it usable for the actual user. Besides providing the GUI, the prototype just interprets the underlying OWL version [27] of the security ontology to ensure a high degree of flexibility. With the security ontology on hand, a typical physical infrastructure modeling process is conducted as follows: (1) definition of the organization, (2) definition of site concepts and relating them to appropriate location concepts, (3) definition of building concepts, (4) definition of level concepts (including information about their vertical position within the building), (5) if necessary, definition of areas, (6) definition of sections and section connectors (e.g., doors or windows), (7) modeling tangible assets (including ratings for acceptable risk levels and importance for the organization's mission) and persons, and relating them with their typical physical location (e.g., sections), (8) modeling data and software concepts and relating them with those IT and telecommunication instances on which they are stored on, (9) relating one or more roles to each modeled person, and (10) modeling organizational controls. The entire physical infrastructure modeling process is supported by the security ontology, which ensures by its concept definitions, relations, and formal axioms a consistent and machine-readable infrastructure model. As already introduced in [22] we developed a novel inventory solution for the software and IT-related infrastructure elements which is able to capture the device data automatically (operating system, IP address, patch level, etc.) independent of the used operating system. 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.

- **Usability:** Both, GSTool and CRISAM require a manual and unguided inventory of the organization's resources and do not support for an automatic or semi-automatic inventory of IT resources. AURUM addresses these shortcomings by an automatic IT resource inventory solution which utilizes several third-party network scanning and inventory products to embed the gathered data in our ontological framework. Although the inventory of non-IT resources is also conducted manually, AURUM guides the user at the inventory phase by a typical inventory process and uses the ontological infrastructure model to ensure a formally correct and consistent infrastructure model.

- **Resource Catalog:** The GSTool provides a comprehensive resource catalog and allows the definition of new resource concepts in the given resource classification. CRISAM provides a limited resource catalog which can not be extended by the end user and, thus, significantly decreases the flexibility in the system characterization step. In contrast to CRISAM, AURUM provides a comprehensive and consistent resource classification and allows the user to define new resource concepts in the given classification.

- **Consistency:** The consistency of the infrastructure model created with the GSTool is endangered because the end user is able to define new resource concepts without any restrictions. CRISAM requires the definition of a logical infrastructure model (on which resources depends the considered resource). Because these dependencies are not constrained at all the consistency of the model is not guaranteed. AURUM provides comprehensive resource definitions to ensure consistency in the ontological infrastructure model (e.g. sections (rooms) have to be connected to a certain level or area concept, movable assets have to be connected to sections in which they are located, etc.).

- **Control Inventory:** Both, GSTool and CRISAM do not consider any existing controls in the system characteriza-

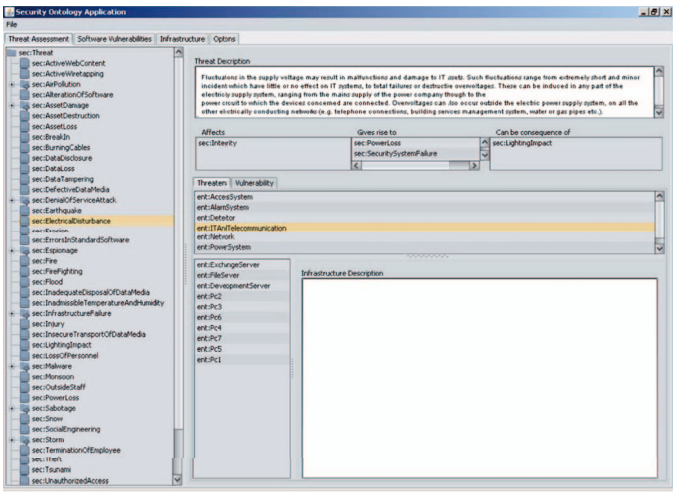


Fig. 1. AURUM - threat view

tion step and circumvent therefore automatic compliance checks in the subsequent steps of the information security risk management process. AURUM incorporates the control modeling in the system characterization step (e.g. modeling fire extinguishers as resources) and recognizes at the same time the risk mitigation potential of each resource. For instance, modeling a fire extinguishing system automatically decreases the fire threat probability in the organization.

## III. Threat Identification

The goal of the threat identification step is to determine potential threats and their corresponding threat sources. Common threat sources are natural threats (e.g., earthquakes, floods, wild fire, etc.), human threats (e.g., active network attacks, theft, unintentional data alternation, etc.), and environmental threats (e.g., power failure, water leakage, ...). At this step it is important to compile a comprehensive list of potential threats as recommended in [6], [9], [11], because the subsequently risk assessment and mitigation steps are taking the results of this step as input for the risk mitigation strategy.

While these standards and best practices often provide an exemplary threat list, the risk manager is not always aware about the nature of each threat. Which threats threaten critical resources? Which threat is a multiplier (i.e. which threat gives rise to other threats)? Which vulnerabilities have to be exploited by a threat to become effective? All these questions are hardly addressed in most of the current risk management standards or best practices. Figure 1 shows our solution concept for that problem, which utilizes the security ontology to present threats and their relationships clearly arranged. The threat tree, located at the left hand side, is the starting point for identifying potential threats to the considered organization. The underlying threat information, including a priori threat likelihoods based on the physical location of the organization, is gained from our OWL-based knowledge base. Each of the items under the sec:Threat root element represents a possible threat to the organization and has to be taken into account in a holistic risk analysis. By selecting a threat from the tree representation, the right area is populated with valuable information. On top a threat description is provided in natural language. Below, affected security attributes (confidentiality, integrity, and availability) are displayed. Furthermore, a threat can be a consequence of other threats (e.g., unauthorized access can be the result of a break in or missing key management) and can itself potentate other threats (e.g., break-in gives rise to unauthorized access or asset damage). Additionally, our proof of concept retrieves the threatened resource concepts for each threat from the security ontology. Due to the fact that we have modeled the organization's resources within the System Characterization Step in an ontological form, we can show only those threats to the risk manager, which are relevant for the organization.

- **Threat Catalog Size:** The GSTool uses the comprehensive threat catalog of the German IT Grundschutz Manual, which provides very detailed natural language descriptions about common information security threats. In contrast to the GSTool, CRISAM does not provide any information on threats. The threat catalog of AURUM is based on the German IT Grundschutz Manual and the French EBIOS [9] standard to provide a widely accepted and formal threat knowledge base including also threat dependencies. Due to the generic structure of the knowledge model we are able to incorporate further information security knowledge sources.

## IV. Vulnerability Identification

Starting from the threat report produced in the previous step, the vulnerability identification step analyzes potential vulnerabilities which are present in the defined system. 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 solutions 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.

- **Vulnerability Catalog Size:** Both, GSTool and CRISAM do not provide any vulnerability catalog to support the vulnerability identification step. AURUM provides a comprehensive vulnerability catalog derived from several best-practice guidelines and information security standards. Incorporating vulnerabilities in the ontological in-

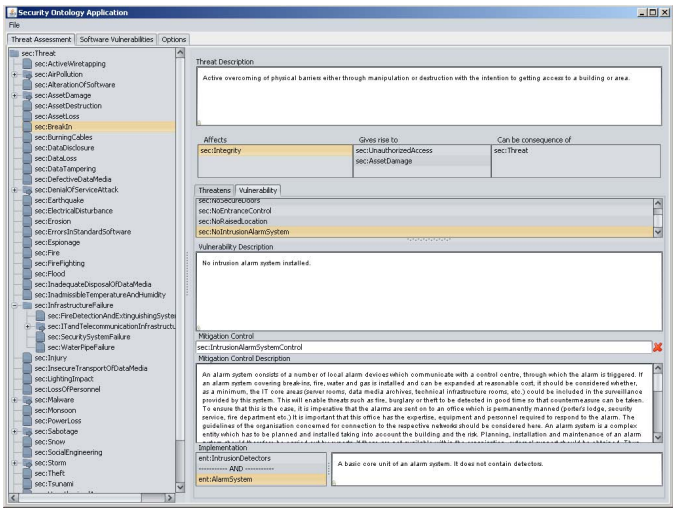


Fig. 2. AURUM - vulnerability and control view

formation security model, enables highly-granular formal descriptions of threats and their relation to risk-mitigating controls.

## V. Control Analysis

Up to the current point the organization is aware of the considered system, potential threats and corresponding vulnerabilities, which allow threats to become effective. The control analysis step determines which controls (either technical such as encryption mechanisms, or nontechnical controls such as security policies) are already in place and which controls are planned to mitigate the likelihood that a threat exploits a certain vulnerability.

To facilitate the aspect of automatic compliance checks regarding our defined mitigation controls, each control further incorporates formal implementation descriptions. The implementation area in Figure 2 shows the actual implementation measures for a control. Referring to our example, the threat break-in exploits the vulnerability No Intrusion Alarm System which could be mitigated by the installation of an intrusion alarm system and an intrusion detector (motion detector, glass break sensor, or heat detector) in every section. The fact that the control has to be implemented in each section, is defined by the *vulnerabilityOn* relation, which states that the vulnerability is on the section level and the vulnerability is mitigated if the control is appropriately implemented for a given section. The underlying formal control descriptions can be executed as rules against the organizations concrete modeled environment to identify which parts of the building are in compliance. Most often it is not sufficient just to know if a certain control is implemented or not within a given area. The most important thing to know is, if the implemented control is appropriate or not to achieve the acceptable risk level. Therefore, we rate each resource regarding to its importance for the organization's mission in terms of confidentiality, integrity, and availability. Since the risk is defined as the product of potential impact and threat likelihood, we aim at incorporating control influences into the threat likelihood.

- **Control Inventory:** Neither GSTool nor CRISAM explicitly support the inventory of existing controls. Instead, both solutions indirectly conduct the control inventory by control questions which have to be answered by the user. AURUM includes a guided physical control inventory in the system characterization step. The user just inventories the organization's resources and AURUM decides if the modeled resources can be used as a control. Organizational controls are also assessed by control questions to map policy and guideline contents to the ontological information security model.

- **Granularity:** GSTool conducts the control analysis based on control questions manually for each resource. In contrast to the GSTool, CRISAM conducts the control analysis based on control questions manually for resources or entire resource groups to decrease the effort for the user. AURUM utilizes the data gathered during the system characterization step to conduct the control analysis automatically for each resource.

## VI. Likelihood Determination

The likelihood determination is concerned with the probability that a threat exploits a certain vulnerability within the given system. Therefore, the organization has to deal with the following factors: (1) motivation and capability of the threat source, (2) nature of the vulnerability, and (3) existence and effectiveness of current controls. We utilized Bayesian networks to determine the organization-specific threat likelihood in a more objective way. The likelihood $TL_T$ of each threat $T$ is influenced by three components: (1) the a priori likelihood $ATL_T$ of threat $T$, (2) the likelihood of those threats which act as an enabler for the considered threat $\{TL_{T_{p1}}...TL_{T_{pn}}\}$ (e.g., a high break-in likelihood rises the likelihood for the unauthorized physical access threat), and (3) the vulnerability exploitation likelihood $VEL_T$ of those vulnerabilities which could be exploited by the considered threat $T$. Since the security ontology already provides information on threat dependencies, threat a priori likelihoods, and corresponding vulnerabilities including a severity rating for each vulnerability we are able to use that knowledge to set up the Bayesian network for a more objective threat likelihood determination: First, the threat net is generated by querying available threats and their interrelations (*givesRiseTo* relation) from the ontology. In a second step, all relevant vulnerabilities for a given threat can be revealed by interpreting the *exploits* relation within the ontological framework. Within the third step, we identify for each vulnerability by relation *mitigatedBy* those controls which are able to mitigate the considered vulnerability. Since we conducted a comprehensive inventory within the system characterization step and incorporated it into the ontological framework, we can determine which controls are already implemented within the organization. Therefore, the exploitation likelihood of a given vulnerability, can be

determined by combining the effectiveness of existing control implementations and the effectiveness of the attacker in case of a deliberate threat source.

Note that each threat likelihood is calculated for each resource, since the determination of already implemented controls is always bound to the considered resource. To get a specific threat likelihood over the entire organization, the individual threat likelihoods per resource have to be aggregated. Thus, AURUM enables the risk manager to deal with an overall as well as resource-specific risk if necessary.

- **Structured Approach:** Both, GSTool and CRISAM do not incorporate any threat probabilities to determine the risk. AURUM provides a mathematically and formally sound threat probability determination based on location-dependent a priori probabilities stored in the security ontology.

- **Control Consideration:** In contrast to GSTool and CRISAM, AURUM incorporates existing controls in the threat probability determination. Since the developed ontological information security model, provides information on threat dependencies, AURUM is able to determine control influences over the entire threat net. For example: since the probability of the smoke threat is highly dependent on the probability of the fire threat, an implementation of an automatic fire extinguishing system would decrease besides the fire threat probability also the smoke threat probability.

- **Attacker Consideration:** Compared to GSTool and CRISAM, AURUM incorporates the nature of a potential attacker in the threat probability determination. The vulnerability exploitation probability and subsequently the threat probability and the actual risk depend besides the effectiveness of relevant controls on the effectiveness of a potential attacker.

- **A Priori Threat Probabilities:** In contrast to GSTool and CRISAM, AURUM provides the possibility to store location-dependent a priori threat probabilities to enable the structured and formally sound determination of organization-specific posterior threat probabilities for each resource.

## VII. Impact Analysis

Understanding the adverse impact of a successful threat exercise of a vulnerability is necessary to determine the risk level and thereby the basis for the subsequent control recommendations. While in most risk assessment approaches, such as proposed in [6], [28], [29], [30], the impact of threats is determined through interviews and workshops involving the system and information owners, AURUM focuses on an automated support utilizing the developed knowledge base and the defined relationships. Instead of rating the impact of specific threat occurrences, we emphasize on rating the importance of inventoried assets. Information and technology owners rate the business impact regarding loss of the defined security attributes. Thus, for each asset the impact in case of loss of availability, loss of confidentiality and loss of integrity are rated separately in terms of High, Medium and Low. The reader should note that the scale can be adjusted with respect to the organization's requirements. As can be seen, instead of rating the impact of each threat directly, which includes understanding the threat in detail, knowing the threatened assets and all implications on business processes and then deciding on the aggregated impact, we reduced the problem on rating the impact for individual assets independent of specific threats.

In a next step we want to assess the adverse impact of a specific threat. Due to the semantic relations between a threat and threatened asset classes, we automatically obtain a collection of concrete threatened assets in an organization (taken from the inventoried resources, cf. Section II). In addition, for each threat the security attributes put at risk by the threat are added to a threat description. Hence, we compare the security attributes at risk, gained from the threat, with the impact categories defined for each threatened asset. Note, that we always calculate the impact for threat/asset pairs, as each asset might cause a different impact. In case a security attribute affected by the threat has been defined as impact relevant (impact on loss of the security attribute has been rated) for a threatened asset, impact on the organization owning the asset, must be expected. To determine the magnitude of impact we apply the impact level assigned for the asset. The following example is given to clarify the impact determination in case of a threat occurrence: The threat of a computer virus puts the security attribute availability at risk. A file server, located in the organization, on the other hand has been identified as business critical and thus the impact in case of unavailability has been set to High. Due to the relationship between threats and assets we know that the server is threatened by the computer virus. Comparing the information reveals that the server's availability is threatened and thereby the organization. Because the impact if the server is unavailable has been set to High we can expect a high impact on the organization in case of a computer virus attack. This result is exactly the impact level of the threat exposure. Of course also more than one security attributes can be affected by a threat, in this case the highest impact level would constitute the overall threat's impact level on the organization. Another important nuance to mention is the case that more than one threatened assets are identified in the organization, e.g., if a threat is defined on the concept level and more than one instances of threatened asset concepts exist (e.g., most likely there are more servers in an organization). In this case, the overall threat impact level is defined by the highest impact level over all threatened assets.

- **Traceability and Granularity:** In the GSTool, the user manually assigns the protection requirements (normal, high, very high) for each target resource regarding the three basic categories of confidentiality, integrity and availability. The requirement rating is directly bound

to impact levels, provided by the standard in natural language. Furthermore, the protection requirements of subordinated resources show up as recommendation for superior resources. In the risk analysis view the user manually states which threats are mitigated, given from the threat catalogue. If all threats are mitigated the resource is regarded as protected, visualized by green color. No connection between the threats and the rated security attributes is taken into consideration.

CRISAM uses the following approach: For each resource the user manually answers control questions on a scale from A (highly implemented) to F (insufficient). The endangered security attributes, which can be selected from a comprehensive set of attributes, are defined for each question a priori by the CRISAM team. In addition a weight is assigned for each control question. By aggregating the answers for each resource an overall rating per resource is derived. This rating is compared to the policy defined target rating and results in a relative rating for the current resource by means of a bond rating system. If the final rating for a resource is too low, visualized by the color red in the resource tree, controls have to be implemented. After the control questions are re-answered and the target rating is achieved, the icons turn green. The worst ratings are always passed to connected resources. In addition, the dependency of resources is rated for each security attribute (4 step scale - Low to Very High). To each value in this scale a monetary value can be assigned by the user.

In contrast to the GSTool and CRISAM our model is based on formal threat descriptions instead of control questions. This allows us to reason automatically which resources are threatened including the expected impact level. By this solution the result is less dependent on human answers on control questions and hence reproducible, comprehensive and consistent.

## VIII. Risk Determination

As a final step in the risk determination, the mission risk is calculated by multiplying the ratings assigned to risk likelihood and the potential impact. The possible values for Likelihood and Impact and how to collect them in an organization are shown in Section VI and VII. Both of these values are calculated by utilizing the introduced security relationship model. As the impact depends on concrete assets, we also calculate the risk for each asset individually. A risk level matrix is a valuable tool in calculating the risk. The aim of a risk level matrix and the resulting risk score, which quantifies the risk, is to provide a consistent and objective methodology to prioritize threats and the next steps. We follow the NIST SP 800-30 guideline and construct a 3x3 risk matrix based on inputs from the threat likelihood (High, Medium and Low) and threat impact (High, Medium and Low).

The possible risk levels in our matrix comprise High, Medium and Low. To determine these levels, the probability for each threat likelihood level is expressed as follows: 1.0 for High, 0.5 for Medium, 0.1 for Low. Regarding the threat impact the following values are assigned: 100 for High, 50 for Medium, and 10 for Low. Now it is possible to multiply the threat probability with the impact values. The risk scale to interpret the results is given below:

- High ( >50 to 100)
- Medium ( >10 to 50)
- Low (1 to 10)

By now we know the individual risk for every threatened asset in case a threat occurs. These asset bound risk levels are required in the subsequent control recommendation step as countermeasures could be required individually for each asset. To gain a single risk level for a threat, the highest risk level over all assets is assigned. In the traditional process, it is important to define the meaning of a specific risk level and the actions senior management must take. While this is a valuable approach, it is not clear which assets caused the risk level and where to apply countermeasures. AURUM on the contrary offers a more detailed and fine grained approach as introduced in the following section.

- **Traceability:** The GSTool does not incorporate probabilities for threat occurrences and thus no risk level can be calculated. Instead, the IT Grundschutz Manual approach defines lists of relevant threats and required countermeasures according to a typical office environment. By answering control questions a variance analysis between the recommended countermeasures in the IT Grundschutz catalogues and those already implemented can be conducted.

  CRISAM offers neither threat catalogues nor related probabilities. The risk management view in CRISAM displays the variance of the achieved finance ratings, which are derived by answering the control questions, for each resource with the policy defined target rating.

  In AURUM the risk is calculated by multiplying the determined impact level with the derived probability for each resource. In contrast to the other tools we present a final risk level to the user which takes the impact and probability into account.

- **Standard Compliant Controls:** Regarding control recommendations, the GSTool offers a comprehensive set of threats and corresponding controls, taken from the IT Grundschutz Manual, to mitigate those. For each resource a set of recommended controls is provided. To support the implementation process a responsible person for implementation, a schedule, priority and costs can be stored. For each control question in the CRISAM tool, the user can create one or more measures to implement the control. These measures include information on the responsible person for implementation, a schedule, priority and costs. No explicit information on the origin of the control questions is given. Contrasting the other candidates, our solution recommends controls to mitigate vulnerabilities. All controls are derived and explicitly

linked to established best-practice guidelines and information security standards. This makes it, e.g., possible to specific information security standard as target.

## IX. Control Recommendations

Up to this point the overall risk of a threat, or a risk level for each threatened asset, can be calculated. Initially, we sort threats by their risk level, which provides the organization's decision makers with a thorough overview of current risks. In addition, for each threat the organization's assets at risk and their properties can be queried, which gives the rationals for the calculated risk levels. As mentioned in Section IV, also vulnerabilities which render threats possible can be inspected. At this point management knows which risks are not acceptable for the organization and therefore, measures have to be taken. In this step of the process, controls which could mitigate or eliminate the identified risks, as appropriate to the organization's operations, should be provided [6]. To support this recommendation step, we consolidate the security model. For each vulnerability, appropriate controls are modeled, taken from best practice standards such as the IT Grundschutz Manual. Offering these controls equips the decision makers with effective countermeasures to lower the risk level and thereby protect their business. In contrast to the traditional process, this solution provides a thorough knowledge base about countermeasures and thus 1) saves time, 2) avoids that effective solutions are simply forgotten, and 3) provides effective controls in compliance with best-practice standards. While this is already a valuable approach, further improvement can be achieved by a model based test for existing countermeasures and excluding those from the recommendation set. Implementation instructions are modeled for each control as axioms, which can be used to test if implementations already exist.

- **Concrete Recommendations:** The GSTool provides control recommendations in natural language taken from the IT Grundschutz Manual. These recommendations are mostly highly detailed, but require an expert to draw the appropriate conclusions for the own organization under inspection.
  CRISAM does not incorporate control recommendations, users have to define them on their own in accordance to the given control questions. The control questions sometimes include hints for control implementations on a very high level.
  Our solution provides control recommendations to close vulnerabilities for each resource. The controls are given in natural language to enhance understanding and furthermore, as formal implementation descriptions on a conceptual level. The knowledge base also contains concrete implementation instances, which can be automatically recommended. Another advantage is that those controls (automatically) detected as implemented are not offered as control recommendations.

## X. Control Evaluation and Cost/Benefit Analysis

After identifying all potential controls, they are evaluated and a Cost/Benefit analysis is carried out. Cost/Benefit analysis is an integral part of risk evaluation because investments into security must precisely target a company's specific business needs in line with economic demands. Despite the importance of this step, NIST SP 800-30 gives only a shallow overview. This step involves the definition of the resource- and benefit categories. The careful specification of these categories is of vital importance as these categories should reflect the corporate strategy and security policy of the company. The criteria are company specific and individually customizable, and they can range from monetary quantities (e.g., minimizing the reduction of monetary loss, monetary costs) to intangible values (e.g., user acceptance, implementation hours, loss of reputation). Therefore, this analysis does not only consider cost and benefit in monetary terms but includes non-financial objectives. All potential controls identified in the previous step are rated against the chosen criteria, where the security ontology already provides a selection of objectives (such as confidentiality, integrity, and availability). Using all 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 $K$ objectives and a strictly better value in at least one objective) where the binary variables $x_i \in \{0, 1\}$ indicate whether or not a safeguard $i$ is selected ($x_i = 1$ if so, and $x_i = 0$ otherwise). Of course, 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 firewalls) is included in the feasible solutions.

- **Effectiveness Rating:** The GSTool offers controls taken from the IT Grundschutz Manual. Effectiveness of those controls is not directly given but sometimes information on the effectiveness can be found in natural language in the IT Grundschutz Manual descriptions. Control effectiveness has no influence in this approach, only if all controls are implemented the resource is regarded as protected.
  In CRISAM the user manually defines controls and connects them to control questions. During the creation process he can also define the expected benefit in natural language as well as provide numbers for expected savings. Obviously, only an expert has the necessary knowledge to conduct this step but still there is a risk of forgetting necessary controls, wrong connections, unrealistic figures, input data errors, etc.
  Focusing on effectiveness, our solution provides effectiveness ratings for all control implementations. Those are assigned once by experts in the knowledge base and can be adjusted or extended by users. This approach allows for focusing on highly effective controls

first and provides data for the subsequent cost/benefit analysis.

- **Granularity:** The GSTool does not support cost/benefit analysis. This step has to be conducted externally. In CRISAM it is possible to manually define figures for expected benefits and costs for controls. We consider cost/benefit analysis as an integral part of risk management. As security does not directly generate business value and does not directly improve the net profit, investing in security can only prevent negative events or reduce related adverse effects. Traditional cost/benefit analysis methods do not consider this relation and are ill-suited for the evaluation of security investments, because they fail to properly take into consideration the many important non-financial criteria. AURUM does not aggregate multiple objectives to a single indicator, but leaves them separated and, furthermore, focuses on the selection of whole portfolios of controls. Thus, it can be guaranteed that all portfolio solutions are not only feasible but also are potentially "good" ones from an objective point of view. The ability to consider multiple objectives is of high importance, as an investments adequacy and efficiency for a company is determined by different parameters. Therefore, the evaluation of a number of objectives must be possible and should be accomplished without a priori weighting of objectives and without breaking them down to a common scale.

## XI. Control Selection

Decision makers are often overwhelmed with the high number of solution alternatives and are often not aware if their investments into security are appropriate or effective at all. Therefore, AURUM provides an intuitive 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. The decision maker requires support in making a final determination of the solution that best fits his/her notions out of the possibly hundreds (or even thousands) of Pareto-efficient alternative portfolios identified in the first phase. We are using a search based procedure, which start 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. AURUM is based on interactive modifications of lower and upper bounds for one or more objectives. To this end, the decision support system (DSS) starts with displaying $K$ "flying" bars (cf. Fig. 3) representing resource and benefit categories (such as costs or availability) that are assigned with units (such as "euro" in the case of costs or "points" in the case of availability).

For each objective (cf. Fig. 4) the system provides information on what can be achieved by (i) the efficient solutions (the dark marks on the left side representing the solution space with all efficient portfolios may visually grow together to vertical bars), and (ii) the alternatives that have remained after the decision maker has made decisions in his/her interactive exploration of the solution space (this subset from the solution space is represented by the right bar).

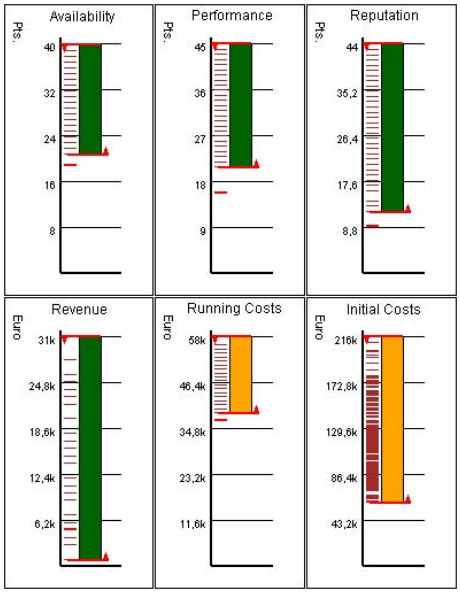


Fig. 3. Status of the DSS at the beginning

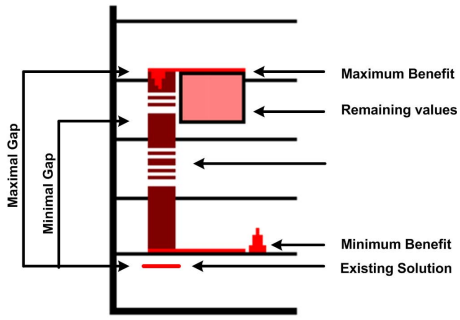


Fig. 4. Subwindow details

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. Let us illustrate this by reducing the maximum allowance for resource A (cf. Fig. 5).

Because this setting has primarily filtered those solutions that come with a relatively high value in "Resource Category A" (and, on average, a somewhat higher need for resource C) but still values in "Benefit Category A", the options in the other objectives have been reduced as well and the position

and size of the flying bars have changed accordingly. Raising the minimum value for Benefit A (e.g., functionality) narrows the set of remaining alternatives even further, since many alternatives with low resource values (e.g., price) drop out (cf. Fig. 6).

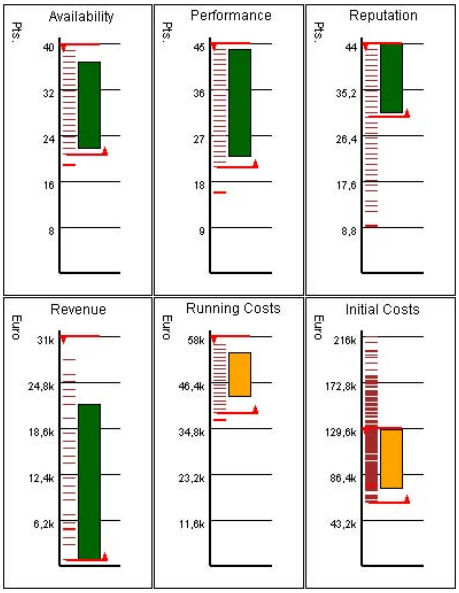


Fig. 5. Status of the DSS after the first setting

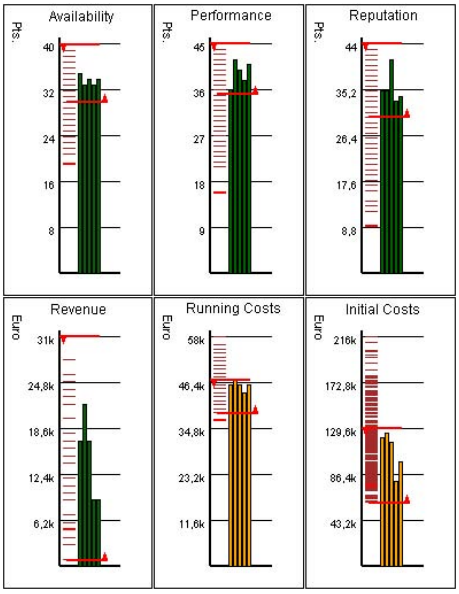


Fig. 6. Status of the DSS after two settings

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 and, thus, gain 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. 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.

- **Interactive Selection:** Whereas the GSTool provides control recommendations in natural language, CRISAM does not incorporate control recommendations. Both tools do not interactively support the decision maker during this phase. We use our knowledge base that contains concrete implementation instances as a basis and provide decision makers with a stepwise and repeatable methodology that gives them plenty of information on the selection problem at hand. This approach provides them with the opportunity to thoroughly explore the set of Pareto efficient solution alternatives until they find the individually most attractive security investment portfolio, while the system at the same time guarantees that only (feasible) efficient solutions are taken into consideration. Finally, the method allows the decision makers to learn more about the characteristics of the specific decision problem and maybe even about their own preferences.

## XII. Conclusions

Companies consider security as one of the most important issues on their agenda, because the increasing number of security breaches poses a major threat to the reliable execution of corporate strategies and may have negative effects on business value. Risk management ensures the consideration of all possible threats and vulnerabilities, as well as the valuable assets. Existing approaches such as best-practice guidelines, information security standards, or domain experts but also risk management approaches that are highly accepted within the community come with shortcomings.

This paper presented a methodology for supporting the entire NIST SP 800-30 risk management process and provides, compared to existing solutions 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 likelihood determination ensures that the threat likelihood determination is based on a more objective level, compared to existing approaches, (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 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. We compared AURUM with common risk management tools, namely CRISAM and the GSTool, in terms of usability, time exposure for conducting the entire risk management process, completeness of the threat/vulnerability identification and the control recommendations, and the granularity of control implementation suggestions.

## XIII. 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] H. Cavusoglu, B. Mishra, and S. Raghunathan, "The effect of internet security breach announcements on market value: Capital market reactions for breached firms and internet security developers," *International Journal of Electronic Commerce*, vol. 9, no. 1, pp. 69–104, Fall 2004.

[2] L. Gordon, M. Loeb, W. Lucyshyn, and R. Richardson, "CSI/FBI Computer Crime and Security Survey," September 2006.

[3] Computer Economics, Inc., "2005 malware report: Executive summary," January 2006. [Online]. Available: http://www.computereconomics.com/article.cfm?id=1090

[4] M. Bishop, "What is computer security?" *IEEE Sec. Priv. Mag.*, vol. 1, no. 1, pp. 67–69, Jan.-Feb. 2003.

[5] B. Farquhar, "One approach to risk assessment," *Computers and Security*, vol. 10, no. 10, pp. 21–23, February 1991.

[6] G. Stoneburner, A. Goguen, and A. Feringa, "Risk management guide for information technology systems," National Institute of Standards and Technology (NIST), Gaithersburg, MD 20899-8930, NIST Special Publication 800-30, July 2002.

[7] 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*. London, UK: Springer-Verlag, 2002, pp. 94–105.

[8] C. Alberts, A. Dorofee, J. Stevens, and C. Woody, "Introduction to the OCTAVE approach," Carnegie Mellon - Software Engineering Institute, Pittsburgh, PA 15213-3890, Tech. Rep., August 2003.

[9] DCSSI, "Expression des Besoins et Identification des Objectifs de Scurit (EBIOS) - Section 2 - Approach," General Secretariat of National Defence Central Information Systems Security Division (DCSSI), February 2004.

[10] ISO/IEC, "ISO/IEC 27005:2007, Information technology - Security techniques - Information security risk management," November 2007.

[11] BSI, "IT Grundschutz Manual," 2004. [Online]. Available: http://www.bsi.de/english/gshb/manual/download/index.html

[12] M. Vitale, "The growing risks of information systems success," *MIS Quarterly*, vol. 10, no. 4, pp. 327–334, December 1986.

[13] K. Bandyopadhyay and P. Mykytyn, "A framework for integrated risk management in information technology." *Management Decision*, vol. 37, no. 5/6, pp. 437–444, 1999.

[14] C. Jung, I. Han, and B. Suh, "Risk analysis for electronic commerce using case-based reasoning," *International Journal of Intelligent Systems in Accounting, Finance & Management*, vol. 8, pp. 61–73, 1999.

[15] W. Baker and L. Wallace, "Is information security under control?: Investigating quality in information security management," *IEEE Security and Privacy*, vol. 5, no. 1, pp. 36–44, 2007.

[16] R. Baskerville, "Information systems security design methods: Implications for information systems development," *ACM Computing Surveys*, vol. 25, no. 4, pp. 375–414, December 1993.

[17] ISO/IEC, "ISO/IEC 27001:2005, Information technology - Security techniques - Information security management systems - Requirements," 2005.

[18] W. Baker, L. Rees, and P. Tippett, "Necessary measures: metric-driven information security risk assessment and decision making," *Communications of the ACM*, vol. 50, no. 10, pp. 101–106, 2007.

[19] S. Frosdick, "The techniques of risk analysis are insufficient in themselves," *Disaster Prevention and Management*, vol. 6, no. 3, pp. 165–177, 1997.

[20] A. Ekelhart, S. Fenz, M. Klemen, and E. Weippl, "Security Ontologies: Improving Quantitative Risk Analysis," in *40th Hawaii International Conference on System Sciences (HICSS'07)*. Los Alamitos, CA, USA: IEEE Computer Society, Jan 2007, pp. 156–162.

[21] A. Ekelhart, S. Fenz, G. Goluch, and E. Weippl, "Ontological Mapping of Common Criteria's Security Assurance Requirements," in *New Approaches for Security, Privacy and Trust in Complex Environments, Proceedings of the IFIP TC 11 22nd International Information Security Conference, IFIPSEC2007, May 14-16*, ser. IFIP International Federation for Information Processing, H. Venter, M. Eloff, L. Labuschagne, J. Eloff, and R. von Solms, Eds., vol. 232/2007. Sandton, South Africa: International Federation for Information Processing, May 2007, pp. 85–95, 978-0-387-72366-2.

[22] A. Ekelhart, S. Fenz, T. Neubauer, and E. Weippl, "Formal threat descriptions for enhancing governmental risk assessment," in *Proceedings of the First International Conference on Theory and Practice of Electronic Governance*. ACM Press, 2007.

[23] T. Neubauer, A. Ekelhart, and S. Fenz, "Interactive Selection of ISO 27001 Controls under Multiple Objectives," in *Proceedings of the IFIP TC 11 23rd International Information Security Conference, IFIPSec 2008*, vol. 278/2008. Boston: Springer, July 2008, pp. 477–492.

[24] T. Neubauer, C. Stummer, and E. Weippl, "Workshop-based Multiobjective Security Safeguard Selection," in *Proceedings of the First International Conference on Availability, Reliability and Security ARES*. IEEE Computer Society, 2006, pp. 366–373.

[25] T. Neubauer and C. Stummer, "Interactive Decision Support for multiobjective COTS Selection," in *Proceedings of the 40th Annual Hawaii International Conference on System Sciences*, no. 01, 2007.

[26] ——, "Extending Business Process Management to Determine Efficient IT Investments," in *Proceedings of the 2007 ACM Symposium on Applied Computing*, 2007, pp. 1250–1256.

[27] W3C, "OWL - web ontology language," http://www.w3.org/TR/owl-features/, February 2004.

[28] J. Burtles, *Principles and Practice of Business Continuity: Tools and Techniques*. Rothstein Associates Inc., 2007.

[29] T. R. Peltier, *Information Security Risk Analysis*, 2nd ed. Auerbach Publications, 2005.

[30] S. Kairab and L. Kelly, *A Practical Guide to Security Assessments*. Boston, MA, USA: Auerbach Publications, 2004.