European Journal of Operational Research 248 (2016) 318–327

Contents lists available at ScienceDirect

European Journal of Operational Research

journal homepage: www.elsevier.com/locate/ejor

Innovative Applications of O.R.

Optimal selection of IT security safeguards from an existing knowledge base

Andreas Schilling[*], Brigitte Werners

Ruhr University Bochum, Faculty of Management and Economics, Universitätsstraße 150, Bochum 44780, Germany

ARTICLE INFO

Article history:
Received 19 December 2014
Accepted 18 June 2015
Available online 24 June 2015

Keywords:
Knowledge-based systems
Information security
Decision support model
Combinatorial optimization

ABSTRACT

In this paper, a combinatorial optimization model is proposed to efficiently select security safeguards in order to protect IT infrastructures and systems. The approach is designed to provide very concrete decision support for an organization as a whole or separately for specific systems. It can be applied in practice without requiring the decision maker himself to collect extensive input data. This is accomplished by using an existing comprehensive and highly accepted knowledge base as a basis for decision making. For our analysis, we use the publicly available IT baseline protection catalogues of the German Federal Office for Information Security (BSI). The catalogues contain more than 500 threats and over 1200 safeguard alternatives to choose from. Applying our model, it is possible to make use of this knowledge and determine optimal selections of safeguards according to given security requirements. The approach supports the decision maker in establishing an effective baseline security strategy.

1. Introduction

Information technology (IT) is constantly spreading into more and more areas of organizations and it is a critical factor to be successful in the global economy. The loss, manipulation, disclosure, or simply the unavailability of information caused by IT security incidents may lead to expenses, missed profits, or even legal consequences. Incidents originate from different actors with different motives. Professional and amateur hackers, malicious employees, industrial spies, or even terrorists try to hack into systems to gain access to information or simply to create damage. They search for vulnerabilities and will use any weak link in the security chain of an organization. In the constant struggle to make systems more secure, organizations are always trying to find new ways to adequately address security issues. If too little is done, security is weak and attackers will probably succeed at some point. On the other hand, if the wrong measures are taken, the organization may be wasting precious resources that could have been used elsewhere.

It has become evident that making systems more secure is a difficult task due to the high complexity of IT infrastructures and the large amount of data required to make informed decisions. To address this problem, a considerable amount of research has been done to find ideal IT security budgets and to determine how to invest based on risk and financial measures. The problem is that existing approaches require the decision maker of an organization to provide a lot of exact input data like exact threat and vulnerability probabilities, asset valuations and other fine-grained parameters. However, these values are very difficult to obtain in practice and, in addition, are very critical to the quality of proposed solutions. Such approaches are intended to accurately address the problem but are often impractical when it comes to real-life applications. On the other hand, approaches that require less information often remain vague in their results and require the decision maker to fill in the gaps himself.

To address these issues, an approach is needed that requires an organization to provide as less data as possible but still produces very concrete investment recommendations. To achieve this, we propose to treat the problem on a different layer of abstraction than existing models and to utilize existing IT security knowledge. In fact, both aspects are closely related: instead of developing a model first and forcing the decision maker to gather required input data, we took a reversed approach: we established an extensive knowledge base first, extracted relevant information, and built a model on top of it. The data required can be obtained from common IT security practices and standards which are available from various sources including, but not limited to, International Organization for Standardization (ISO), National Institute of Standards and Technology (NIST), and German Federal Office for Information Security (BSI). These practices help organizations by guiding them on how to establish an effective basis for security. The information is mostly available in form of catalogues, guidelines, recommendations, or standards which have to be followed more or less strictly.

When analyzing existing data, it becomes evident that exact probabilities, financial impacts, etc. are not available, not even for the past.

[*] Corresponding author. Tel.: +49 234 32 28311.
E-mail addresses: andreas.schilling@rub.de (A. Schilling), or@rub.de (B. Werners).

http://dx.doi.org/10.1016/j.ejor.2015.06.048
0377-2217/© 2015 Elsevier B.V. and Association of European Operational Research Societies (EURO) within the International Federation of Operational Research Societies (IFORS). All rights reserved.

But the data still provide a valuable basis for decision making, although on a different level. The question is: is it really necessary to know a threat's exact probability of occurrence? Is it necessary to know the exact financial impact in case of a security incident? If the goal is, for example, to accurately estimate expected losses, the answer is "yes". But if the objective is to protect an IT system or infrastructure, we propose an alternative to facilitate the establishment of effective security safeguards without requiring this information. By abstracting the problem further and using highly aggregated "criticality coefficients", we eliminate the dependence on exact parameters like probabilities and financial evaluations. The aggregated parameters are generated based on an extensive knowledge base with real-world security data. The idea is that exact threat probabilities or expected financial losses are not required to decide which safeguards should be implemented. Knowing which threats are more critical relative to each other is sufficient to allocate safeguards effectively. Effective, in this context, means that safeguards are selected such that their combined impact on the reduction of security threats is maximized.

For our analysis we focused on the IT baseline protection catalogues (or IT-Grundschutz catalogues) (Federal Office for Information Security, 2013b) which are part of a standard security practice provided by the BSI. We chose the IT baseline protection catalogues because they are publicly available free of charge and offer an extensive repository of technical, organizational, personnel, and infrastructural information security knowledge to protect IT systems. The catalogues are in line with the ISO 27000 series and, in addition to custom BSI certificates, it is possible to obtain an ISO 27001 certification based on safeguards listed in the baseline protection catalogues.

Organizations applying these practices to improve the state of their IT security practices can be divided into basically two groups: one group strictly follows the guideline of a standard of their choice and tries to obtain a valid certificate. The other group is not necessarily interested in a certificate, but nonetheless wants to achieve a sufficiently high security level of its IT infrastructure and systems. Certifications fulfill their purpose for the first group of organizations but have only limited value for organizations of the second group. Organizations trying to improve security without a certification face the challenge of selecting appropriate safeguards from the given catalogues. In the following, we mainly address the second group and demonstrate how they can benefit from existing IT security knowledge.

This paper is a first approach to make the large amount of available information security knowledge contained in the IT baseline protection catalogues usable for organizations. For this purpose, we propose a combinatorial optimization model which makes use of the entire baseline protection catalogues. The model is also applicable to an arbitrary subset of components which makes it usable for any use case covered by the knowledge base of the catalogues. It is possible to integrate our approach into an existing risk management process to automate and support the selection of safeguards and thus guide the decision maker in creating an effective risk mitigation plan.

The remainder of the paper is organized as follows: Section 2 outlines literature related to decision making in IT security. Section 3 introduces the source of data which we used as basis for our evaluation and illustrates the data extraction process. Section 4 presents our mathematical model and shows how available data are utilized to support automatic decision making in IT security. In Section 5 we conceive a realistic case study to demonstrate the application of our model and discuss the results and computational performance. Section 6 concludes the paper.

2. Related work

In recent years, the interest in quantitative models for information security investment decisions has increased significantly. This trend is driven by the fact that information security is becoming more important each day and, at the same time, the complexity of IT systems continuously increases. Questions like, "How much security is necessary?", "How much should be spent?", and "How can security be improved?" are becoming more relevant these days. There are several research streams which basically try to solve the security investment problem from different angles. This problem can be broken down into two distinct subproblems where each subproblem is focused on one key issue: (1) what is the optimal amount to invest in security; and (2) what security safeguards should be selected to invest in?

The first question is probably the most-discussed one and there exists a considerable amount of related literature. It is often addressed by traditional risk analysis methods to determine loss expectancies and a return on investment (Berinato, 2002; Bojanc & Jerman-Blažič, 2012; Sonnenreich, Albanese, & Stout, 2006). This is a reasonable starting point, since other investment problems have been treated very successfully with these tools. The security investment problem, however, is problematic because security investments have no return in the classical sense of the word, i.e., there is no incoming cash flow after investing in security. For this reason, risk analysis approaches usually treat prevented losses as a profit: profit = loss reduction × probability of incident. A study by Gordon and Loeb (2002) uses risk analysis to suggest an optimal budget for a risk-neutral decision maker. In their approach, they compare the loss caused by security incidents to the investment required to reduce the related vulnerability. Based on two general classes of security breach functions, they state that the amount to invest is considerably lower than the expected loss caused by an incident. In fact, they find that the amount to invest in security never exceeds 37 percent of the expected loss and in most cases will be substantially less. However, these observations only hold true if the security breach functions meet the condition of decreasing marginal returns in case of security investments. Hausken (2006) examined four additional types of security breach functions with different shapes and found that the amount to invest is no longer capped at 37 percent and different investment strategies should be applied in each case. Wang, Chaudhury, and Rao (2008) proposed a more detailed analysis which makes use of security incident data and statistical methods like the concept of value-at-risk to support a decision. Another general class of approaches falls into the field of microeconomics and uses game theoretic models to treat security as a game between an organization and an attacker to determine an ideal investment level (Baykal-Gürsoy, Duan, Poor, & Garnaev, 2014; Cavusoglu, Raghunathan, & Yue, 2008; Gal-Or & Ghose, 2005; Gao, Zhong, & Mei, 2014; Roy, Ellis, Shiva, Dasgupta, Shandilya, & Wu, 2010).

These approaches address the security investment problem from a business perspective. Financial investments and achieved security have to be balanced. On this level, an organization can obtain insights on a reasonable budget. Which safeguards have to be taken and how the budget should be distributed remains up to the decision maker.

At this point, the second question has to be answered: which safeguards should be selected for implementation within a budget that was determined previously? Most approaches to address this question apply management tools and financial analysis based on measures like annual loss expectancy, return on investment, internal rate of return, net present value, etc. (Bojanc & Jerman-Blažič, 2012; Schilling & Werners, 2013; Sonnenreich et al., 2006; Tsiakis, 2010). Other approaches use real options analysis where dynamic aspects of investments are considered and the flexibility of decision making is utilized Gordon, Loeb, and Lucyshyn (2003); Tatsumi and Goto (2010); Ullrich (2013). An optimization driven approach to select security safeguards is proposed by Sawik (2013) which produces optimal safeguard portfolios. In their study, they use a bi-objective model to minimize expected and worst case losses applying the value-at-risk. Viduto, Maple, Huang, and López-Peréz (2012) also proposed a bi-objective model that considers the trade-off between financial