Risk Reduction Overview

Hellen Janine Havinga, Olivier Theobald Sessink

HAL Id: hal-01403999
https://inria.hal.science/hal-01403999

Submitted on 28 Nov 2016

Risk Reduction Overview

A visualization method for risk management

Hellen Nanda Janine Havinga[1], Olivier Diederik Theobald Sessink[2]

[1]Rijkswaterstaat, Central Information Services, Delft, The Netherlands
`hellen.havinga@rws.nl`
[2]Ministry of Defense, The Hague, The Netherlands
`odt.sessink@mindef.nl`

Abstract. The Risk Reduction Overview (RRO) method presents a comprehensible overview of the coherence of risks, measures and residual risks. The method is designed to support communication between different stakeholders in complex risk management. Seven reasons are addressed why risk management in IT security has many uncertainties and fast changing factors, four for IT security in general and three for large organizations specifically. The RRO visualization has been proven valuable to discuss, optimize, evaluate, and audit a design or a change in a complex environment. The method has been used, evaluated, and improved over the last six years in large government and military organizations. Seven areas in design and decision making are identified in which a RRO is found to be beneficial. Despite the widely accepted need for risk management we believe this is the first practical method that delivers a comprehensive overview that improves communication between different stakeholders.

Keywords: Design · Security · Residual risk · Risk management · Security measure · Visualization

1   Introduction

Risk management in IT security is complex. Large numbers of security measures and many stakeholders make risk management in large organizations even more complex. Risk management involves the balance between the residual risks for the business and the costs of the measures that are taken to reduce the initial risks. In this article we define risk as the product of the chance that a threat causes damage to the business and the damage of that threat to the business. A 5% chance for $100 total damage, for example, will justify the costs for a $5 measure. The costs include not only the cost of the implementation of the measure itself, but also the costs of any loss of functionality.

There have been various approaches to quantify and model the security costs and benefits [6, 7, 11], [13]. However, in large organizations both the chance and the damage are mostly unknown. Furthermore, the estimates of the effect of measures on

chance and impact also have a high degree of uncertainty. Those approaches are therefore not practical for real-world risk management in situations with large numbers of security measures. There also have been several approaches to present risk management visually [16, 17]. However, none of these methods present an intuitive overview of the coherence of all risks and measures required for risk management.

1.1 Challenges of Real World IT Risk Management

There are four main reasons why chance and damage are uncertain in IT security and thus why risk management is complex. First, the known vulnerabilities in IT systems change with a high rate. Every computer runs millions of lines of code, and thus the existence of bugs is almost a certainty [12]. At a certain moment, for example, it may seem impossible to gain unauthorized access to a system, a week later there may be a zero-day exploit and an experienced hacker may gain access, one week later an exploit is released on the internet and access is possible for every "script kiddy", and again one week later the vulnerability is patched and unauthorized access seems impossible again. Second, the IT environment itself changes continuously, which changes both the chance and the potential damage to the business. For example the introduction of new software or a new network connection to an external system. New technology developments such as cloud computing and ubiquitous computing that introduce completely new security challenges extend this challenge even more. Third, the chance that a threat causes damage is influenced by unknown external factors. It is for example difficult to quantify how much effort an external entity is willing to take to gain access to your information, or to quantify the number of backdoors in the software you use, or to quantify the bypass rate of your security measures [2]. Last, the cost of the damage is hard to estimate. The damage incurred when IT systems are unavailable, interrupting business processes or critical infrastructure, or when sensitive information is disclosed, affecting competitiveness or causing reputation damage, is not easy to express in monetary terms [3], [5], especially in the public sector where information if often sensitive for political, sovereignty, or privacy reasons [1]. Furthermore, the total cost of ownership of measures is difficult to express in monetary terms. Especially since most measures cause a decrease in productivity as a side effect.

1.2 Challenges of IT Risk Management in Large Organizations

Large organizations add three additional challenges to the complexity of IT risk management. First, where in a small organization a single administrator might be solely responsible for the overall security, in large organizations there are many different roles involved, such as business owner, information security officer, authorizing official, functional application manager, solution architect, and system administrator. This separation of roles causes few people to have the required overview and knowledge to link security measures to chance and impact on the business required for a good cost/benefit analysis. Second, large organizations have large numbers of information systems which are interconnected in many ways: sharing hardware, net-

work infrastructure, storage, and data. Defining a strict boundary for a single information system is therefore almost impossible. A security breach on any level might affect many information systems. NIST 800-37 [9] recommends segmentation of systems with guards in between, but this is not considered feasible for most networks except when dealing with very high classification levels. The chance that a security breach occurs is difficult to estimate because there are thousands of IT components and thousands of unique security measures that affect the chance. The damage is difficult to quantify because a security breach might affect many information systems. Last, the cost/benefit analysis becomes rapidly more complex due to the large number of business processes with different security requirements in large organizations. A new security measure may, for example, reduce the risk for one business unit, but may decrease productivity for another business unit.

1.3 Risk Management Methods and Standards Used in Large Organizations

To overcome the challenges of complexity, large organizations use generic risk management methods and security baseline standards for the generic infrastructure, such as NIST 800-37 [8], CRAMM, and ISO/IEC 27005 [10], to guarantee a minimum security level for all systems. In order to be sustainable over time these methods and standards use generic threats and describe only generic measures that allow for different implementations; only some common cases are described. The generic measures are not clearly linked to threats to the business [4]. When the threats change over time, the implementation might need to be re-assessed and updated. Furthermore, whether or not a specific implementation complies with the standard has to be justified by the IT security architect.

Every implementation thus needs an argumentation why it meets the requirements of the generic measures and how it tailors or supplements baseline security controls. The IT security architect needs to provide insight into three aspects to justify his decisions. First, the different successive complementary and independent technical and procedural measures. Second, the risks which are reduced by each security measure. Last, the residual risks that have to be accepted.

So although these standards help larger organizations to improve the overall security, they do not eliminate the need for communication between those that design the measures and those that have to accept the residual risks.

2 Objectives

The challenges in IT risk management as described have led to the objectives for the RRO method. First, present the relation between risks, measures and residual risks in an intuitive way. Second, present the security design in a way that gives people with different skills and roles the opportunity to either discuss, evaluate or audit if a design meets the required security level. Third, present the security design in a way that people can evaluate the residual risks that a design imposes on the business. Last, the presentation should be applicable independent whether or not the design is already

implemented or in design phase. Summarizing, the presentation should help the business to improve risk management: to improve communication and clarify the link between threats, security measures, and residual risk and make IT security designs more comprehensible and auditable.

3   The Risk Reduction Overview

The Risk Reduction Overview (RRO) consists of two parts: a flowchart representation and an appendix. The flowchart provides an intuitive overview of the coherence of all risks, measures, and residual risks. The position and relation of measures in the flowchart show if a measure is successive to another measure (which provides defense in depth), complementary to another measure (the measure reduces a different aspect of the risk), or independent from other measures (the measure acts on a different risk). The appendix provides the details on each risk, measure, and residual risk.

**Fig. 1.** Basic elements of a Risk Reduction Overview

The flowchart is based on two basic elements: risks and measures. There are three types of risks: the initial risks, the residual risks, and the final residual risks (Fig. 1). All paths in the RRO flowchart follow the same format: initial risks are identified and measures are applied to reduce these risks, which lead to residual risks. The flowchart starts with all the initial risks that are identified for the particular design. All initial risks are followed by one or more measures, which are followed by residual risks and more measures, and finally end with a final residual risk. Arrows depict the flow. The flow is not necessarily linear; multiple measures from different flows may lead to the same residual risk, and multiple residual risks may follow a single measure. Arrows are drawn from risk to the resulting residual risk. When the measure itself introduces a new risk, an arrow can be drawn from that measure to a new risk.

All risks and measures in the flowchart have a unique identifier: Initial risks have the identifier I# (in which # is a unique number), residual risks have the identifier R#,

measures have the identifier M#, and final residual risks have the identifier F#. In the flowchart the identifier is followed by a short description of the risk or measure. The appendix of the RRO contains the detailed description for each risk and measure in the flowchart.

3.1 Example

To illustrate the risk reduction flow we use a simplified example: email communication between a network with confidential data and the internet is enabled, and six security measures are proposed to reduce the risks (Fig. 2). This particular example describes a very common application for a risk reduction overview: a change on a secure environment is proposed, and the risk reduction overview is used to show the proposed measures and the residual risks of that change. Figure 2 does not describe a real-life design. Real risk reduction overviews [14] often have over ten initial risks and over twenty measures (Fig. 3).

Fig. 2. Example RRO of email communication between a network with confidential data and the internet

Fig. 3. Example of a real RRO (sensitive information is removed)

3.2 Drawing Method

Before making a RRO, an initial set of risks and measures must have been identified already. The initial risks are the risks if one would take no security measures at all. The initial risks can be derived from the threats to the business combined with all possible vulnerabilities. The measures can be both technical and procedural. Initial risks can be derived from threat analysis and from the standards, but must be completed from expert and domain knowledge. The first step to set up the RRO flowchart is to lay out the set of initial risks horizontally on the top of the flow chart. Then start to lay out the measures below the initial risks and derive the residual risks from the measures. Complementary measures may be placed below to reduce the residual risks further. Place preventive measures, that reduce the chance that a threat occurs, above reactive measures, that reduce the damage when the threat occurs. Continue until all measures are positioned and only final residual risks are left. If two measures provide an identical risk reduction (for defense in depth) they should follow up on each other without residual risk, e.g. measure 3 and 4 in Figure 2. Similar measures and similar risks should be placed near each other, to ease steps two and three.

Fig. 4. Before (left) and after (right) step 2 and step 3

In the second step similar measures are joined together to make the overview more compact and easier to comprehend, as can be seen in Figure 4. This step is greatly facilitated if similar measures have been positioned early in the flow during step one, because re-ordering measures changes all the residual risks. Once the generic measures are applied early in the flow, the residual risks become more specific, so eventually all the detail is still present in the overview.

In the third step similar risks are joined together, as can be seen in Figure 4. It is often possible to rewrite similar risks into a single more generic description, especially early in the flow. Late in the flow residual risks can be very specific and combining them would result in loss of essential detail.

If the number of risks and measures is still too large to give a single comprehensible overview, it is an option to create a summary RRO with less detail, with multiple underlying RRO's that give the required detail. The summary RRO still gives insight in the detailed final residual risks, but summarizes the upstream risks and measures in the risk reduction flow. For example: different measures like fire alarm, fire extinguisher, fire blankets and non smoking areas can be summarized as fire fighting measures. Reviewers can still find the details of fire fighting in the underlying detailed fire fighting RRO, while business owners might only be interested in the summary RRO, that present them the final residual risk of all fire fighting measures. Evaluation of the method has led to the finding that single RRO's with over 30 measures are generally considered not comprehensible anymore

To support quick reading risks and measures should be described in a uniform sentence structure. Start the description of the risks with the vulnerable element. For example: Email attachments may contain malware. Start the description of the measures with the subject that causes the desired effect. For example: Anti virus software scans emails for malware.

The last step is to number and describe each risk and measure of the optimized flowchart in the appendix. The appendix has four sections. The first section describes all initial risks. The second section describes all measures. The level of detail should provide enough information for a reviewer to judge if this measure will indeed reduce the risk to the residual risk. The third section describes all residual risks. The fourth and last section describes all final residual risks. In order to give weight to each of the risks, the description in the appendix should give an indication of the chance that a risk might occur and the damage it may cause to the business. The description should provide a chief security officer or a business owner enough information in a clear and comprehensive way to assess if the final residual risk is acceptable. An indication of the cost of measures and the cost of possible damage could be added to help the decision maker.

Positioning the risks and measures manually can be time consuming, but this process could be automated.

4   Application

The RRO can be used in different stages of the IT security design and decision making process, and can be used to review the design of an already existing implementation. Seven areas are identified in which the RRO is beneficial.

First, by creating the RRO, a security architect has to rethink design decisions of a new design or existing implementation and might find flaws or forgotten details. It forces the security architect to ask himself if the complete set of security measures does indeed cover all the risks it is supposed to. The flowchart provides a clear overview where defense in depth and diversity in defense are applied.

Second, the RRO can help a security architect to optimize a design. Duplicate measures, or measures that do not reduce a risk, are more easily identified. Furthermore, the order of the measures may have an impact on the overall costs of the design. A cheap measure as first layer of defense may, for example, reduce the required capacity for a more expensive second layer of defense.

Third, during review of the RRO, reviewers can check if all risks they expected themselves are present in the overview. A missing risk may be an indication of an insecure design or implementation. A missing risk will make the list of final residual risks incomplete.

Fourth, reviewers of the RRO can check if the suggested risk reduction of a measure is realistic. They are able to see if the assumptions about the effectiveness of technical and procedural measures are correct and if the residual risks are well quantified. Too high residual risk may be an indication of an underestimation of the effectiveness of a measure. Very low residual risks may be an indication of an overestimation. In both cases the corresponding final residual risks will be incorrect too.

Fifth, the chief security officer or business owner of the particular IT system gets an overview of the initial risks, the measures and most importantly the final residual risks. Too high residual risk may be an indication of an insecure design or implementation, very low residual risks may be an indication of too many or too expensive

security measures. If the final residual risks in the RRO are not acceptable for the business, there is a clear gap between the business needs and the security implementation.

Sixth, risks can change over time. When the security infrastructure is in use, new vulnerabilities of systems or measures can occur or the chance that vulnerabilities will be exploited can increase or decrease. The RRO gives for example a computer emergency response team the opportunity to immediately see what effect these changes have on the residual risk. The chief security officer or business owner either accepts the new residual risk or requires new measures to be taken to get the residual risks to an acceptable level.

Last, an existing RRO can be used as a source of inspiration for the design of a new similar environment or to review a new design.

5   Conclusions, Limitations, and Lessons Learned

The RRO method has been used and evaluated in two large organizations over the last six years. The method has been applied to various complex problems in the fields of information security and cyber defense at the Joint IT Command of the Dutch Ministry of Defense, and for cyber defense of critical infrastructure at the Dutch Rijkswaterstaat. In the Joint IT Command the RRO is now a mandatory document for changes that affect residual risk. At Rijkswaterstaat the RRO is now a mandatory document when exceptions from baseline security requirements are requested for critical infrastructure objects and mission critical systems. In both organizations the RRO is used to clarify, discuss and evaluate the security design of innovative products and services, for which no baseline requirements exist. New baseline requirements are extracted from these RRO's.

The RRO method has been found to be beneficial in all seven application areas described in this article. The RRO is found to deliver a comprehensible overview of the coherence of risks, measures and residual risks. Even first time readers with no previous experience with a RRO have little trouble to identify why measures are taken and which residual risks are left. More experienced readers point out that they need less time to review measures and residual risk with a RRO. Especially if there is a large number of measures or risks involved, the RRO gives far more overview than a traditional technical design document does. The concept of risk reduction is understood by people with different skills and roles, and the use of a RRO does improve their understanding of the coherence of the measures and the residual risks. Business owners of information systems point out that the RRO enables them to discuss measures with IT specialists, something they found very difficult in the past.

The RRO method is, however, not a silver bullet for IT risk management. First, having a RRO does not guarantee that the measures are actually correct and that the real residual risks match the described residual risks. Risks change over time; a new vulnerability that was previously unheard of may introduce a completely new risk, and new threats may require new measures. The competence of the security architect is still one of the most important factors. Second, a RRO covering the security of a

large and complex environment will result in a large and complex visualization. The RRO will not make a complex problem simple. Last, the creation of a RRO requires more time than a traditional technical design document that just lists the measures. We do believe, however, that the cost benefit ratio favors the RRO, which is strengthened by the fact that both organizations that have evaluated the RRO have decided to make it a mandatory document in the decision making process.

The most important observation while the method was evaluated was that adjusting the level of detail and the layout of the flowchart manually requires a significant amount of time. This discouraged some authors to improve a RRO after reviewers had sent their comments, especially since there was no obligation to deliver a RRO at the time. A tool to automate this process should therefore be developed and will be published in the future on the RRO website [15].

6 Application Outside IT Security

The RRO is believed to be applicable in any area in which risk management is an issue, such as public safety, fraud prevention, food safety, physical security, military operation planning and medical hygiene. There are three separate situations in which a RRO is beneficial for generic risk management. First, if there is a need for communication about risks, measures and residual risks. Second, if there are stakeholders in the design or decisions making process that do not have the adequate knowledge to derive the residual risk from the different measures taken. For example if stakeholders are not known with the technology of complementary measures from different knowledge domains. Last, if the number of risks or measures in a certain risk assessment is high and a better overview is required to discuss or evaluate the overall situation. If multiple situations are present, the case for a RRO is even stronger.

7 References

1. Acquisti, A., Friedman, A., Telang, R. Is there a cost to privacy breaches? An event study. In: Fifth Workshop on the Economics of Information Security. Cambridge (2006)
2. Arora, A., Hall, D., Pinto, C., Ramsey, D., Telang, R.: An ounce of prevention vs. a pound of cure: How can we measure the value of IT security solutions? Lawrence Berkeley National Laboratory, University of California (2004)
3. Berinato, S.: Finally, a real return on security spending. In: CIO Magazine, Februari 15, pp.43–52 (2002)
4. Bornman, G., Labuschagne, L.: A comparative framework for evaluating information security risk management methods. In: Proceedings of the Information Security South Africa Conference. ISSA (2004)

5. Garg, A., Curtis, J., Halper, H.: Quantifying the financial impact of IT security breaches. Information Management and Computer Security 11(2), pp.74–83 (2003)

6. Gordon, L. and Loeb, M.: The economics of information security investment. ACM Transaction on Information and System Security 5(4), pp. 438–457 (2002)

7. Hoo, K.J.S.: How much is enough? A risk management approach to computer security. Doctoral Thesis, Stanford University (2000)

8. NIST Special Publication 800-37 Revision 1, http://csrc.nist.gov/publications/nistpubs/800-37-rev1/sp800-37-rev1-final.pdf

9. Joint Task Force Transformation Initiative: Guide for Applying the Risk Management Framework to Federal Information Systems: A Security Life Cycle Approach. NIST Special Publication 800-37, Revision 1. Computer Security Division, Information Technology Laboratory, National Institute of Standards and Technology (2010)

10. Joint Technical Committee ISO/IEC JTC 1/SC 27: ISO/IEC 27005:2011 Information technology — Security techniques — Information security risk management. International Organization for Standardization, Geneva (2011)

11. Longstaff, T., Chittister, C., Pethia, R., Haimes, Y.: Are we forgetting the risk of information technology? Computer 33(12), pp. 43–51 (2000)

12. Martin, R.A.: Managing Vulnerabilities in Networked Systems. Computer 34(11), pp. 32–38 (2001)

13. Neubauer, T., Klemen, M., Biffl, S.: Business process-based valuation of IT-security. In: Sullivan, K. (ed.) Proceedings of the 7th international workshop on Economics-driven software engineering research. ICSE, pp.1–5. ACM, New York, (2005)

14. Risk Reduction Overview example, http://rro.sourceforge.net/examples.html

15. Risk Reduction Overview website, http://rro.sourceforge.net/

16. Roy, A., Kim, D.S., Trivedi, K.S.: Attack countermeasure trees (ACT): Towards unifying the constructs of attack and defense trees. Security and Communication Networks, pp. 929–943 (2012)

17. Schneier, B.: Attack Trees. Dr. Dobb's Journal of Software Tools 24(12), pp. 21–29 (1999)