Addressing Misalignment Between Information Security Metrics and Business-Driven Security Objectives

Christian Fruehwirth
Aalto University
(Helsinki University of Technology)
Otaniementie 17, Espoo, Finland
christian.fruehwirth@tkk.fi

Stefan Biffl, Mohamed Tabatabai, Edgar Weippl
Vienna University of Technology
Institute of Software Technology
Vienna, A-1040 Austria
stefan.biffl@tuwien.ac.at

ABSTRACT
Companies, which approach information security management from a business perspective, invest in using security metrics to measure the degree to which their security objectives are being met.

The decision however, on which particular security metrics to use, is surprisingly often based on an uninformed process and disregards the company's security goals and capabilities. Like a factory owner, who bought a new tool, without considering which business goals it should support and whether the staff is actually equipped to operate it, introducing metrics without considering security goals and security capabilities can lead to ineffective operation. Practitioners complain in this context about their security metrics being too complex to use, requiring data that is expensive to gather, or simply measuring the wrong thing. Existing frameworks such as the SSE-CMM or ISO 27000 series provide generic guidance on choosing security objectives and metrics, but lack a method to guide companies in choosing the security metrics that best fit their unique security objectives and capabilities.

In response to this problem we present a method with a tool that supports matching security metrics with the objectives and capabilities of a company. Our method helps companies in deciding which metric best suits their particular context, by determining which metric is 1.) efficient to apply using a companies given capabilities and 2.) provides the maximum contribution to the company's security objectives. The method is supported by existing research in the field of value-based software engineering and has been developed based on the established "Quality Function Deployment" (QFD) approach.

Initial experiences from applying the method suggest that the method improves the selection process off security metrics.

Categories and Subject Descriptors
K.6.M [**Security**]

General Terms
Management, Measurement, Security.

MetriSec2010, September 15, 2010, Bolzano-Bozen, Italy.

Keywords
Security Management, Security Metrics, Business-driven ITSM.

1. INTRODUCTION
While security metrics are a foundation of managing IT security, Jaquith described the challenge of choosing a metric: "*But what kind of numbers? Ay, there's the rub.*" [8]. The driving force behind most efforts in quantifying security lies in the fact that information security management has evolved from a technical into a business issue for many organizations [5]. This means that information security now serves clear business objectives. Along with this new perspective comes the need to measure security aspects in an effective and preferably quantitative [19] way. Scholars and security professionals in the industry have responded to this need for measurement by developing increasingly sophisticated security measurements and metrics. A metric in the context of security is described by Saydjari as "*a system of related measures enabling quantification of some characteristic*" [17]. A measure as "*a dimension compared against a standard.*" [17]. There is a large variety of such security metrics available today that can fulfill the need for measurement [1][8][20]; however it is hard for organizations to decide which particular set of metrics is the "right" one for them to use in a given context. There are several reports [17][18][19] that mention this problem, but to our knowledge no comprehensive and context-aware solution has been proposed. In this work, we refer to *context* as the set of the business-driven security goals that an organization wants to support by measurement and the security capabilities that are available within the company to do so.

Through our work with the Finnish software industry we came across many practitioners who acknowledged the importance of security metrics but were experiencing significant problems when they introduced metrics in their organizations. The most common complaints about the security metrics used were: too complex application, required data is not available or too expensive to gather. These complaints are in line with the literature where Jaquith describes a 'good' metric as having the exact opposite characteristics [8] (more in section 2).

We refer to this situation, where a company uses a subjectively 'bad' metric, as a *misalignment* of the used security metric with the company's security capabilities and security objectives. We argue that a major reason for the occurrence of such misalignments is the decision process used to select the security metrics in the first place. In many real-world cases, the metric selection is performed in an uninformed process that disregarded the company's individual security capabilities and security objectives, thus leading to the described misalignment.

The remainder of this work addresses this problem and connects it to previous research in section 2. Section 3 proposes a method as potential solution. Section 4 summarizes experiences from applying the method and section 5 concludes the results.

2. RELATED WORK AND RESEARCH ISSUES

This work specifically relates to three topics in exiting research: 1.) measuring security with security metrics, 2.) defining security objectives and requirements and 3.) capabilities in organizations.

2.1 Security metrics

The first question to ask when introducing a security metric, is "What makes a *good* security metric?". The issue is widely discussed in the literature [8] [19] [18] [17] and authors agree that desirable characteristics are: 1.) consistency of measurement, 2.) inexpensive to gather, 3.) expression as a cardinal number or percentage and 4.) using at least one unit of measure. Stefani et al uses these criteria to establish meta-metrics [18] for the evaluation of metric candidates in the e-commerce domain. Stefani's work is useful because it shows that not all metrics are equally useful in a given context. Jaquith identified this problem as well by stating, that a good metric needs to be "*contextually specific*", and "*relevant enough to decision makers so they can take action*" [8]. Hence, we argue that 1.) in order to be relevant to decision makers, metrics need to be linked to the organization's security objectives and 2.) in order to be able to take action, the metrics need to be in line with (i.e. executable by) the company's security capabilities.

2.2 Research Focus

Based on Jaquith's argument, the focus of this work is to develop a method that can help companies choose the 'right' security metric from a set of candidates. It is not the goal of this method to evaluate the metric per se, as to which is the "best" but rather to provide a guideline on which would be most suitable given the company's security objectives and available capabilities.

2.3 Security objectives and security capabilities

Many existing information security frameworks, like the ISO 27000 series, provide companies with excellent support in defining their security objectives [7]. Traditional security objectives typically include the assurance of confidentiality, integrity and availability of an organization's information assets in a value-neutral form [12]. We focus on business-driven security objectives because they enhance traditional objectives with the concept of value. The following example illustrates the difference between value neutral and business driven objectives:

A value neutral security objective: "*Protect the integrity of the data warehouse*."

A business-driven security objective: "*Use 60% of security resources for protecting the integrity of the data-warehouse (and 20% for confidentiality and 20% for availability")*.

If the introduction of a security metric should not be a waste of money, the metric needs to contribute in some form to the company's security objectives. A metric can contribute to a security objective by strengthening capabilities that are employed by the company to achieve the objective. For example: A metric that measures the amount, type or source of malicious content detected by the company's web-proxy over time will improve the company's capability to conduct security awareness trainings (by better targeting the training lessons towards real world threats). In return, improved security awareness training is likely to reduce the number of security incidents that need to be fixed by the resident IT department, thus contributing to the objective of lowering incident management costs.

More detailed descriptions of the various contributions of security metrics are available in the literature [8][20][19] and not in the scope of this work. Instead, we focus on measuring the 'extent' of the contribution. The extent to which a metric can contribute to an objective is determined by the 'fit' of the metric with the objective and by how 'well' it can be applied.

The ability to apply a security metric is determined by the company's individual capabilities. The 'how well' is measured by the range of results that can be expected from the application. This is in line with a similar description of "*process capabilities*" by Paulk in the Capability Maturity Model (CMM) [14].

Consequently, we argue that a company's security capabilities represent the connection, or *alignment mediator* between security metrics and objectives. Thus, in order to determine the most suitable metric from a set of candidates they need to be evaluated in the light of these given capabilities.

Figure 1 summarizes this concept and illustrates the differentiation between well-aligned and misaligned security metrics with the example of three different metric candidates.

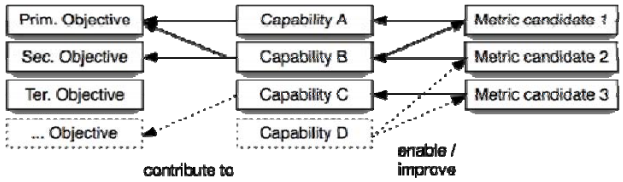

Figure 1 - Relation between security objectives, capabilities and metrics; Differentiating well-aligned from misaligned metrics

Metric candidate 1 in Figure 1 is enabled by capability B and contributes to the primary security objectives through improving capability A. Thus we consider metric candidate 1 to be well aligned.

Metric candidate 2 on the other hand is less well misaligned: Even though it is expected to strengthen capability B, which contributes to two objectives, the metric candidate's application requires the, currently unavailable, capability D.

Metric candidate 3 also requires the same unavailable capability D and only improves capability C, which does not contribute to the top security objectives. Thus we consider candidate 3 to be misaligned.

If a company had to decide which of the above 3 security metric candidates to introduce or spend money on, they would probably be quick to rule out candidate 3, but remain less clear on the choice between metric candidate 1 and 2.

How can we support this type of differentiation in a structured process and make it available in a form that supports a company's metric decision?

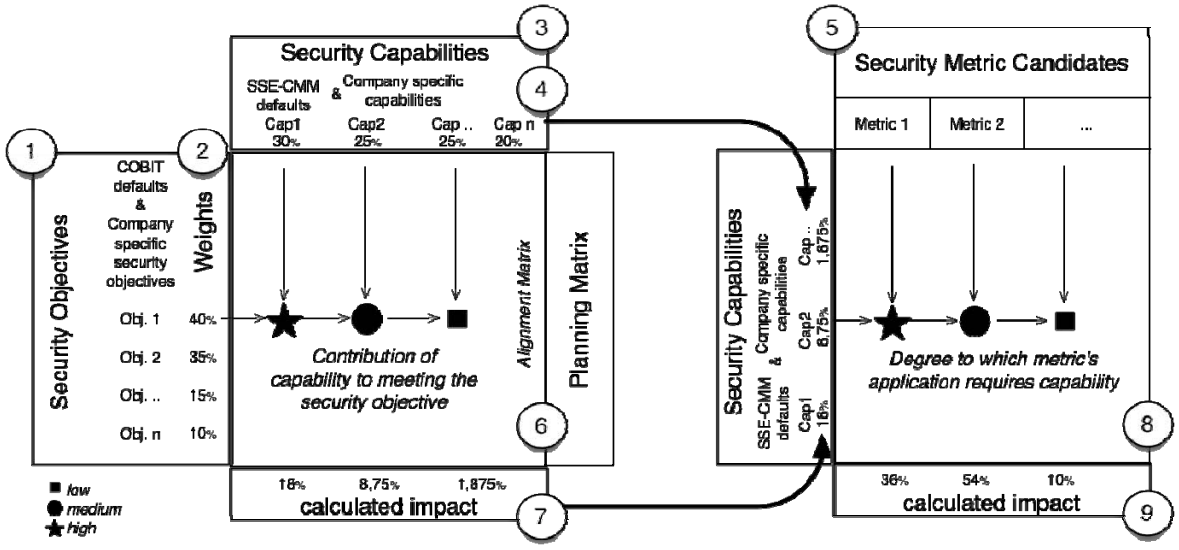

Figure 2 - Alignment Method overview.
From left to right: Matrix1: Security objectives against capabilities; Matrix2: Security capabilities against metric candidates

2.4 Structured alignment process and Quality function deployment (QFD)

Alignment processes for the purpose of connecting capabilities with business objectives have been discussed in the literature on software process improvement and value based software engineering. Oza et al [13] developed a suitable solution using the Quality Function Deployment (QFD) approach[1]. QFD structures links between information by creating alignment matrices, called 'Houses of quality' (HoQ) [13]. QFD has been used successfully t in production planning by the Toyota auto company [6] and authors like Richardson et al[16], Liu et al [10] and Oza demonstrated that it can be adapted for new domains like software engineering. In the security context Kongsuwan[9] and Mead[11] recently analyzed its application in optimizing the quality level of information security and security requirements elicitation.

We chose to use QFD as basis for the alignment process because of these authors' experiences, its ease of use and proven track record in the industry.

3. METHOD DESCRIPTION

In this section we introduce the proposed alignment method and illustrate its usage. The goal of the method is to construct two alignment matrices, where the first matrix aligns security objectives against security capabilities, and the second matrix security capabilities against security metrics. Figure 2 shows a schema of the first matrix on the left side and the second matrix on the right.

The method is executed in 10 consecutive steps, which are divided into 3 phases: 1.) Preparation, 2.) Performing the Alignment and 3.) Analysis. The circled numbers in Figure 2 mark the elements of the alignment matrices that are created during the 10 steps.

The method is ideally executed in a workshop-type setting, where the main stakeholders of the company's security projects are present. The participation of these stakeholders ensures maximum leverage of domain specific knowledge during the alignment process. The following describes the process of constructing and filling the matrices.

3.1 Preparation Phase

The purpose of the preparation phase is to elicit and weight the three input-components of the method: security objectives, security capabilities and metric candidates.

3.1.1 Step 1: Elicit Security objectives

The participating stakeholders (or more generally, the user of the method) are presented with a default list of security objectives and asked to expand the list by brainstorming objectives that are specific to their company's context. Most companies will have several security objectives in common (e.g. "protect the customer data from unauthorized access"), thus using a pre-selected list of default objectives reduces the workload on the workshop participants and frees resources to focus on security objectives that are specific to the company (e.g. "minimize the amount of unencrypted communication between on- and off-site staff"). The "Control Objectives for Information and related Technology" (COBIT) framework [3] is used as basis to create the list of default objectives.

3.1.2 Step 2: Weight Security objectives

The combined list of objectives (defaults + context specific security objectives) then undergoes a two-staged reduction process, which consists of a voting and weighting procedure based in parts on Boehm's collaborative requirements negotiation approach [2].

First, each participant is given a fixed number of votes and casts them on the security objectives he or she perceives as most

[1] More information on Quality Function Deployment is available from the QFD Institute at http://www.qfdi.org/

important. The list of objectives is reduced to the top 10 objectives and the rest is discarded.

In the second stage, each participant is asked to distribute a total of 100 points among the remaining objectives, with the more important ones receiving more points. The standard deviation of points assigned to a single objective by different participants is used as indicator for the degree of consensus among the participants about this objective's relative importance. The workshop moderator uses this information to identify and resolve disagreement about the importance of individual security objectives.

The resulting list of weighted security objectives is placed on the Y-axis of the first matrix, as depicted on the left part of Figure 1.

Optional: Benchmark security objectives:

Stakeholders may rate their company's planned and current performance in fulfilling each of the selected security objectives (target/actual comparison). The rating uses a scale from 0 to 5, similar to CMMI levels, where 0 refers to not meeting the objective at all and 5 to meeting it to the fullest extent. The benchmark is based on QFD's planning matrix and not part of the alignment process or matrix output in our method, thus considered optional. It can be useful however to visualize which objectives require the most improvement efforts and which may remain unchanged.

3.1.3 Step 3: Elicit security capabilities

The elicitation of security capabilities follows the same process as the security objectives in step 1 (see 3.1.1). The users are presented with a default list of capabilities generated from the Systems Security Engineering Capability Maturity Model (SSE-CMM) framework[2] and asked to expand it with company specific capabilities. After expansion, the list undergoes the same voting and weighting procedure as in step2 (see 3.1.2).

3.1.4 Step 4: Weight security capabilities

The remaining capabilities are weighted using the same 100-point system described in step 2 (see 3.1.2), with stronger capabilities (i.e. the expected results from using these capabilities are above average expectations) receiving more points than weaker ones. The resulting list of weighted security capabilities is then placed on the X-axis of matrix 1, and the Y-axis of matrix 2.

3.1.5 Step 5: Elicit Security metrics candidates

The set of security metric candidates that are considered by the company are placed on the X-axis of matrix 2. The composition of a candidate set can either be determined on a case-by-case basis or drawn from examples in the literature, like [1][8].

3.2 Perform the alignment

Steps 1 to 5 completed the frames of the two matrices (see . Figure 2): Matrix 1 with security objectives as rows and capabilities as columns; Matrix 2 with the same security capabilities as rows and metric candidates as columns. The cells in the matrices will now be filled with *Low*, *Medium*, *High* or *None* alignment symbols (depicted in the legend of Figure 2).

[2] More information on the Systens Security Engineering – Capability Maturity Model (SSE-CMM) is available online at http://www.sse-cmm.org

3.2.1 Step 6: Align security objectives with capabilities

Each cell in matrix 1 is filled with one alignment symbol '`impact`' that refers to the impact of a security capability on the company's ability (columns) to meet a particular security objective (rows). For example: how much does the capability to "Log data-access policy violations on firewalls" contribute to fulfilling the objective of "ensuring customer data base integrity". The process of filling the matrix with impact symbols follows the previous weighting of the security objectives and capabilities, with the more important objectives and capabilities tackled first. The filling process is performed by the method users (e.g. the workshop participants).

3.2.2 Step 7: Calculate capability rating

After the 1st matrix is completed, the cumulative impact on security objectives is calculated for each capability:

```
For each Capability(ii){
 Capability(ii).impactSum +=
 (Objective(i).weight *
  Capability(ii).impact(i) *
  Capability(ii).strength)
}
```

The value of

`Objective(i).weight` and

`Capability(ii).strength`

are percentages and were determined in the weighting processes of step 2 and 4. The value of

`Capability(ii).impact(i)`

is determined by factors assigned to each of the three alignment symbols (e.g. *4* for "*High*", *2* for "*Medium*", *1* for "*Low*" and *0* for *"None"*). The value of the individual factors depends on the company and their specific environment, thus experts of the local domain should be used to set them.

3.2.3 Step 8: Align security capabilities with metrics

The previously identified security capabilities are arranged against the metric candidates. The calculated cumulative capability impact '`impactSum`' from matrix 1 is used as capability weight on the Y-axis of matrix 2. The alignment is performed similar to matrix 1, by filling the matrix cells with *Low*, *Medium*, *High* or *None* alignment symbols. In matrix 2 however, each cell is filled with two, alignment symbols to reflect the two-way relationship between capabilities and metrics described in Figure 1. The first symbol '`capRequirement`' refers to the degree to which the application of a security metric candidate (rows) requires, or benefits from, a particular capability (columns). (E.g. how much does the metric "Logging coverage in % of hosts"[3] require or benefit from the capability to "Log data-access policy violations on firewalls".) The second symbol '`capContribution`' refers to the potential contribution of a metric candidate on strengthening a particular capability.

3.2.4 Step 9: Calculate metric candidates rating

After the 2nd matrix is completed, the cumulative alignment score of the metric candidates are calculated:

```
For each MetricCandidate(iii){
```

[3] For more information see [8] pp.56, table 3-3

```
 MetricCandidate(iii).alignmentScore +=
 ( Capability(ii).strength *
  MetricCandidate(iii).capRequirement(ii)) *
 ( Capability(ii).impactSum *
  MetricCandidate(iii).capContribution(ii))
}
```

The factors of for the symbols in

`MetricCandidate(iii).capRequirement(ii)` and

`MetricCandidate(iii).capContribution(ii)`

are determined as in step 7.

3.3 Step 10: Analysis

The first matrix helps the company to determine which capabilities have the strongest impact on reaching the set security objectives (see step 7 in Figure 2). The second matrix enables an evaluation of the metric candidates' capability requirements and their expected contribution to the security objectives (see step 9 in Figure 2). Higher alignment scores indicate that a metric candidate is more likely to benefit from existing, strong capabilities and contribute to the improvement of capabilities that support the company in achieving their security objectives. A sorted list of metric candidates based on their scores represents an important decision support tool on which metric(s) are the most suitable alternative(s) for introduction in the company. In addition to the prioritized list of metric candidates, different patterns of symbols in both matrices allow for more complementary analysis. (e.g. determining the number of capabilities with a high contribution to a particular security objective vs. objectives that are not supported by at least 1 capability).

3.4 Tool support

The tool support for the method provides an opportunity to apply it in practical contexts. The proposed tool was developed using a spreadsheet metaphor and is designed for use in collaborative settings like group workshops. In such a workshop scenario, a moderator takes control of the tool and operates it according to given instructions. The moderator feeds the needed data to the tool based on the input from workshop participants. The tool support allows the participants to focus more on the substance of the discussion as all data arrangement and matrix development is taken care of by the tool. For example, the tool visualizes the level of agreement between the participants in the objectives and capabilities weighting process (see step 2 in section 3.1.2 and step 4 in section 3.1.4) and enables the moderator to direct the argument to those items where consensus has not been reached.

A screenshot of the tool prototype is presented in the appendix.

4. EXPERIENCES FROM METHOD APPLICATION AND LIMITATIONS

We have implemented the method at workshops in two organizational settings. The first concerned the introduction of security measurements in a software service bus environment[4], the second an online community for software business practitioners[5]. Initial experiences (based on the moderator's observations and feedback from the workshop participants) were:

- The method particularly helped highlight that not all organizations have the necessary capabilities to use all metrics.
- The visualization of the alignment matrices was perceived useful and easy to understand.
- The visualization revealed in one case that a previously used security metric (the sum of CVSS vulnerability scores per host) was neither properly supported by the organization's capabilities nor contributed greatly to the security objectives.
- The perceived importance of security objectives greatly differs between different stakeholders and the method's weighting steps (steps 2 and 4) triggered intense discussions between participants. Tool support was found to be very useful in highlighting areas of disagreement and achieving consensus.
- Participants noted that the method would take too long to execute without tool support.
- The typical execution time with tool support was 2 hours and deemed acceptable by all participants.
- The generated ranking of security metrics was perceived most useful when analyzed in combination with the visualized alignment matrices.

4.1 Limitations

The success of the presented method arguably depends on the contribution of the method users, which may induce personal bias during the voting or alignment steps. This issue however can be identified and remedied by a competent workshop moderator. Another limitation lies in the analysis capabilities of the tool support, which in its current prototype state, only provides visual and calculation aids but does not help in discovering emergent patterns from the alignment data in the two matrices. These features and the tool client software are currently under development.

Ultimately one could even argue that the "*companies and governments preoccupation with measuring what is measurable*"[15], what Power described in the economics domain as "*targets culture*" has gone too far and that a formal method to decide on security metrics is not worth the effort. In our experience however the method is lightweight and easy enough to yield quality results in a 1-2 hour workshop. We believe this to be time well spent; in particular for companies that engage with security metrics for the first time on a organization-wide scale.

5. CONCLUSIONS AND FUTURE WORK

This work highlighted the issue that different security metrics are not equally suitable for every company, as different metrics support different objectives and require different capabilities

[4] More information on the Open (Software) Engineering Service Bus (OpenEngSB) is available online at http://www.openengsb.org

[5] More information on the Software Business Community (SWBC) is available online at http://www.swbcommunity.org

which may not be available to the required extent in the particular company.

In response to this problem we presented a method that aligns metrics with the security capabilities and security objectives to help companies decide on the 'right' metric from a set of candidates. The proposed solution helps identify the most suitable security metrics by constructing alignment matrices and a ranked list of metrics based on their capability requirements and the capabilities' contribution to the company's security objectives. The method was developed based on the experience of previous works with the Quality Function Deployment (QFD) approach and integrates existing standard frameworks like COBIT and (SSE-)CMMI.

Practitioners have responded positively to the proposed method and used it successfully in real-world settings. The experiences from applying the method with tool support suggest that it is easy and effective to use in a workshop-type setting. It enabled workshop participants to determine the most suitable metrics and additionally helped identify exiting metric misalignments.

In the future, practical usability of the method will benefit from continuing development of the tool support. On the research side, empirical work will be needed to evaluate the impact of the chosen security metrics on companies' security objectives. The combined results of practical and empirical work will then provide an important contribution to the body of knowledge in information security metrics.

6. REFERENCES

[1] Abedin, M., Nessa, S., Al-Shaer, E., and Khan, L. Vulnerability analysis for evaluating quality of protection of security policies. *Proceedings of the 2nd ACM workshop on Quality of protection*, ACM New York, NY, USA (2006), 49-52.

[2] Boehm, B., Grünbacher, P., and Briggs, R.O. EasyWinWin: a groupware-supported methodology for requirements negotiation. *Proceedings of the 23rd International Conference on Software Engineering*, IEEE Computer Society (2001), 720-721.

[3] COBIT. Information Systems Audit and Control Association. available online at www.isaca.org/COBIT.

[4] Cohen, L. and Cohen, L. *Quality function deployment: how to make QFD work for you*. Addison-Wesley Reading, MA, 1995.

[5] Frühwirth, C. On Business-Driven IT Security Management and Mismatches between Security Requirements in Firms, Industry Standards and Research Work. In *Product-Focused Software Process Improvement*. 2009, 375-385.

[6] Hauser, J.R. and Clausing, D. The house of quality. *IEEE Engineering Management Review 24*, 1 (1996), 24–32.

[7] Humphreys, T. State-of-the-art information security management systems with ISO/IEC 27001: 2005. *ISO Management Systems 6*, 1 (2006).

[8] Jaquith, A. *Security Metrics: Replacing Fear, Uncertainty, and Doubt*. Addison-Wesley Professional, 2007.

[9] Kongsuwan, P., Shin, S., and Choi, M. Managing Quality Level for Developing Information Security System Adopting QFD. *Software Engineering, Artificial Intelligence, Networking, and Parallel/Distributed Computing, 2008. SNPD'08. Ninth ACIS International Conference on*, (2008), 19–24.

[10] Liu, X.F., Sun, Y., Kane, G., Kyoya, Y., and Noguchi, K. QFD application in software process management and improvement based on CMM. *Proceedings of the third workshop on Software quality*, (2005), 6.

[11] Mead, N.R. and INST, C.U.P.P.S.E. Experiences in Eliciting Security Requirements. (2006).

[12] Neubauer, T., Klemen, M., and Biffl, S. Business process-based valuation of IT-security. *Proceedings of the seventh international workshop on Economics-driven software engineering research*, ACM (2005), 1-5.

[13] Oza, N., Biffl, S., Fr\ühwirth, C., Selioukova, Y., and Sarapisto, R. Reducing the Risk of Misalignment between Software Process Improvement Initiatives and Stakeholder Values. *Industrial Proceedings of EuroSPI*, (2008), 6–9.

[14] Paulk, M.C. and Paulk, M.C. *The capability maturity model: Guidelines for improving the software process*. Addison-Wesley Reading, MA, 1995.

[15] Power, M. *The risk management of everything: rethinking the politics of uncertainty*. Demos, 2004.

[16] Richardson, I. Development of a generic quality function deployment matrix. *Quality management journal 9*, 2 (2002), 25.

[17] Saydjari, O.S. Is risk a good security metric? *Proceedings of the 2nd ACM workshop on Quality of protection*, ACM (2006), 59-60.

[18] Stefani, A. and Xenos, M. Meta-metric Evaluation of E-Commerce-related Metrics. *Electronic Notes in Theoretical Computer Science 233*, (2009), 59-72.

[19] Vaughn, R.B., Henning, R., and Siraj, A. Information assurance measures and metrics–state of practice and proposed taxonomy. *Proc. of Hawaii International Conference on System Sciences*, (2003).

[20] Wang, A.J.A. Information security models and metrics. *Proceedings of the 43rd annual Southeast regional conference - Volume 2*, ACM (2005), 178-184.

7. APPENDIX

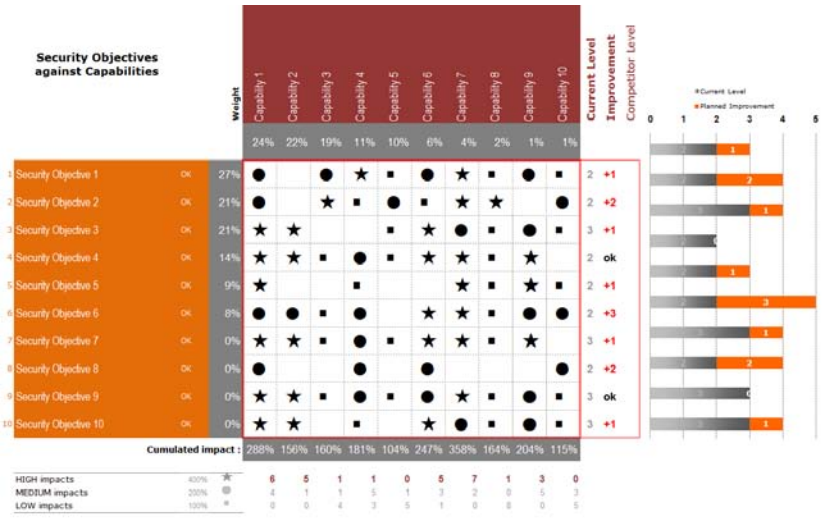

Figure 3 - Screenshot example of an alignment matrix, visualized by the tool support used in a workshop