

Chapter 7

The Building Security in Maturity Model as a Research Tool

CONTENTS

7.1	Introduction	201
7.2	Background	202
7.3	Questionnaires in Software Security	202
7.4	A Case Study	203
7.5	Discussion	205
7.6	Conclusion	206
	Acknowledgment	207
	References	208

7.1 Introduction

Measurement of software security is difficult; it is next to impossible to take two pieces of code and decide which is “more secure” than the other [4]. To tackle this problem, bright minds had the idea to instead try to measure second-order effects, i.e., to study the activities related to software security that are performed by successful software development organizations.

The Building Security In Maturity Model (BSIMM)[6] has been used successfully for years by the software security company Cigital¹ to measure the software security maturity level of their clients. The BSIMM report and framework is released with a Creative Commons Attribution-ShareAlike license², which implies that it is freely available to anyone who wants to use it for whatever purpose, including self-assessment.

In this chapter we try to establish whether BSIMM is also suitable as an academic research tool, and discuss possible adjustments that could make it more tractable. The remainder of the chapter is structured as follows: In Section 7.2 we present relevant background related to BSIMM. In Section 7.4 we present a case study where BSIMM was used by a third party to perform a maturity assessment of a set of software development organizations. We discuss further in Section 7.5, and conclude in Section 7.6.

7.2 Background

The starting point for the first BSIMM survey in 2008 [7] was to study the software security activities performed by nine selected companies. The nine companies were presumably far ahead in software security, and the activities that were observed here formed the basis of the framework in Table 7.1. Representatives from Cigital physically visited each company, and these first surveys were done by Gary McGraw and Sammy Migues personally, using a whole day for each company.

The purpose of BSIMM is to quantify the software security activities performed in real software development projects in real organizations. As these projects and organizations use different methodologies and different terminology, a framework that allows describing all initiatives in a unified manner has been created. The BSIMM framework consists of twelve practices organised into four domains; Governance, Intelligence, Secure Software Development Lifecycle (SSDL) Touchpoints, and Deployment (see Table 7.1). Each practice has a number of activities on three levels, with level 1 being the lowest maturity and level 3 is the highest. For example, for practice Strategy and Metrics, SM1.4 is an activity on level 1, SM 2.5 is an activity on level 2, and SM 3.2 is an activity on level 3. In total, there are currently³ 113 BSIMM activities.

7.3 Questionnaires in Software Security

Questionnaires has been a popular empirical research tool in many scientific disciplines, particularly in Organizational Behavior and patient-centered medical research. The popularity is partly due to the fact that it enables collection of data from

¹<http://www.cigital.com>

²<https://creativecommons.org/licenses/by-sa/3.0/>

³New activities are added as they are observed in the field, and activities are promoted or demoted as their relative importance is determined to change. In the latest update of the BSIMM report, from BSIMM 6 to BSIMM 7, one new activity was added, and 4 existing activities were assigned new levels.

Table 7.1 The BSIMM Software Security Framework

Governance	Intelligence	SSDL Touchpoints	Deployment
Strategy and Metrics	Attack Models	Architecture Analysis	Penetration Testing
Compliance and Policy	Security Features and Design	Code Review	Software Environment
Training	Standards and Requirements	Security Testing	Configuration Management and Vulnerability Management

a large number of subjects in a relatively short time, and with modern online survey tools the data collection can also be made largely automatic. The method is not uncontroversial, though, as many have criticized the use of self-report questionnaires due to single-method induced variance or bias [10]. Boynton and Greenhalgh [2] note that the design of a questionnaire is not a trivial thing, and that a poorly designed questionnaire is unlikely to yield useful results. They also emphasize that if there already exists a validated questionnaire suitable for the task at hand, this also provides the additional advantage of being able to compare the results with previous studies.

Spector discusses self-report questionnaires used in Organizational Behavior studies [10], and finds that whereas Negative Affectivity may influence responses to a questionnaire, it may also influence other data collection methods, and this alone cannot be a reason to eschew questionnaires. Spector also highlights that longitudinal studies can go some way toward mitigating concerns regarding causal relations.

A validated questionnaire can be a good vehicle for a longitudinal study when tracking the development of certain characteristics, e.g., extent of usage of software security activities. The same questionnaire can be distributed to the same organizations (if not the same respondents), and responses should be comparable. When asking about activities and procedures performed in an organization (rather than by the individual respondent), it may also be less likely that responder defensiveness plays into the answer.

Software developers, and maybe particularly *agile* software developers, are unlikely to have patience with overly long questionnaires that take a long time to fill out. This will certainly also contribute to responder fatigue [9] if the same responder is exposed to multiple iterations of a questionnaire. It is therefore important that the questionnaire is short enough to be perceived as manageable by the respondent. Longer questionnaires may be acceptable to certain respondents that see the intrinsic value of the questions (e.g., managers who are interested in the results), but knowing your audience is vital. For longer questionnaires, a face-to-face interaction may be necessary to ensure completion [9].

7.4 A Case Study

Jaatun et al.[5] performed a study on the software security maturity of 20 public⁴ organizations in a small European country using the BSIMM activities as a basis for a questionnaire. The method used in Jaatun et al.'s study can be characterized as "assisted self-evaluation"; the respondents from the various organizations indicated in a questionnaire which software security activities they do, and then they participated in a follow-up interview with the purpose of clarifying uncertainties and correcting possible errors in the questionnaire. However, the researchers did synchronize their assessment criteria, both before and during the interview phase, in order to ensure that they had an as similar as possible perception of what is required to receive a "yes" for the various activities in the questionnaire. However, it is still possible that researchers may have made different assessments related to what should be approved as an activity.

Since the study was based largely on self-evaluation, there is reason to believe that the resulting "BSIMM-score" is higher than it would be with a review in line with the one made by Cigital in the original BSIMM study, since they were not in a position to verify the claims made by each organization. In concrete terms, this implies that we must assume that it has been easier for the organizations to get an activity "approved" in that study than it would be if Cigital had done the survey in accordance with its usual practice. This means that although these results provide some indications of the maturity level of the evaluated organizations, none of the organizations in this study can claim that they have established their "BSIMM Score." It would also be misleading to compare their results directly with the official BSIMM reports. On the other hand, the validity of the answers in the study was increased because of the follow-up interviews, compared with the results from a pure survey.

One thing that is clear is that the organizations studied vary dramatically, both in maturity level and in what kind of activities they perform. Figure 7.1 illustrates this for the three organizations that received the highest total maturity score among the 20 surveyed. This figure uses the so-called "conservative" BSIMM measure defined by Jaatun et al. [5], where 0.5 points are given if only some activities on level 1 are performed within a practice, 1 point means all activities on level 1 are performed, 1.5 points means all activities on level 1 plus some on level 2 are performed, and so on. We see that the top organization gets a top score in the practice "Code Review", but the next two organizations do only a few activities on the lowest maturity level. None of the three organizations do all of the activities even on the first level in the practice "Strategy and Metrics," whereas the third organization does all the level 1 activities and some level 2 activities in the practice "Standards and Requirements," where the first and second organizations do not even do all the level one activities.

The BSIMM framework is based on the idea that there is a formally defined software security group (SSG), and the activities are centered around this group. Few of the surveyed organizations had such a formally defined group. Several organizations had a manager with more or less explicit responsibility for software security, but then it was usually as part of an overall security responsibility in the organization.

⁴government departments, government-owned or municipality-owned organizations, etc.

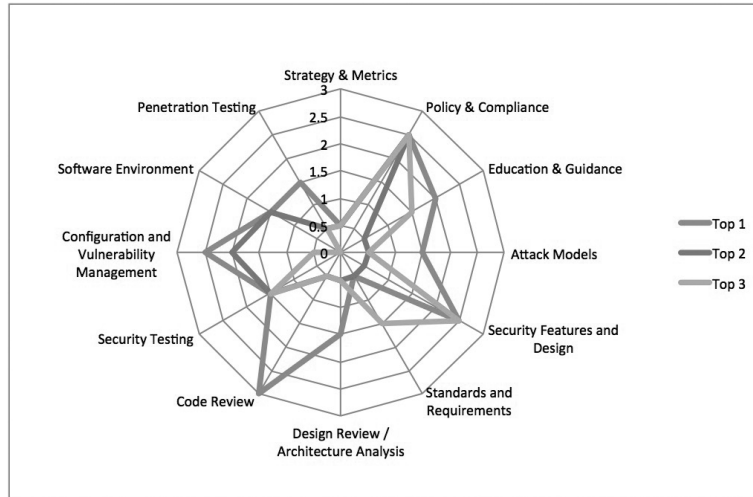


Figure 7.1: Conservative maturity for the three most mature organizations

7.5 Discussion

In personal communication, Konstantin Beznosov stated that he abandoned software security [1] as a research topic because he was unable to get access to the inner workings of the development organizations, and thus was unable to do real empirical research. This may be where the main differentiator lies, since Cigital typically fills the role of a consultant with the target companies, and at the point of performing the BSIMM assessment, they have already convinced the target organization “what’s in it for them.” As BSIMM gains name recognition among mainstream European businesses, it may be that this will also spill over to more academic endeavors; many businesses are interested in knowing more about where they stand when it comes to software security, and many are interested to know how they compare with other, similar organizations.

“The real BSIMM” is not performed using a questionnaire, but using a questionnaire approach significantly lowers the threshold for initiating a software security maturity study. As Jaatun et al. [5] have shown, much of the ambiguity can be resolved by a simple follow-up interview. However, more work is necessary to compare the level of information that can be extracted from an organization using questionnaire and follow-up, vs. embedding one or more researchers in the organization for a day. Although self-assessment is frequently used in other fields such as medicine [3], we cannot overlook that optimistic bias will lead some respondents to overstate their practices [8]. However, it may be equally possible that some respondents may downplay their maturity because they realize that they could be even better; in a larger statistical sample these effects may cancel each other out.

The official BSIMM study [6] has been performed for seven years, and has involved a total of 129 software development firms. The software security activities identified in BSIMM have all been observed “in the wild,” and the list has been fairly stable for the past few years (only 1 new activity in BSIMM7). This is a good argument for using BSIMM as the basis of a survey instrument, since it allows for comparison with a large number of real software security initiatives, and the numbers seem to confirm that these are the right questions to ask.

Another important aspect of BSIMM is that it is the actual performance of activities that is important, not just having the procedures in place. Thus, depending on who is being asked, the answer may be “yes” (because we have the procedures) or “no” (because we never use the procedures). Clearly, selection of respondents must be done carefully, and strategies to mitigate a situation of sub-optimal respondents must be in place. One could be tempted to say that more explanations of each activity would remove ambiguity and doubt, but with 113 activities the questionnaire is already quite long, and takes about an hour to fill out. In the case of an online questionnaire, an alternative might be to first only display the level 1 activities, and only then display the level 2 activities if all the level 1 activities are “fulfilled.” The disadvantage of this approach is that it only covers the conservative metric introduced by Jaatun et al., and not the weighted or high-water-mark metrics, the latter of which is used for comparison purposes in the BSIMM report [6].

BSIMM claims to be descriptive rather than normative, but by ranking activities in maturity levels, there is an implicit statement that some activities are “better” (or more mature) than others. However, a given organization may have good reasons for not doing a certain activity, but this will not be reflected in a study that blindly follows the BSIMM framework. A concrete example of this could be an organization that develops and runs a service that runs in the cloud. In this case, activity SE2.4 “Use code signing” does not make sense, since the source or binaries are never transferred out of the organization’s cloud.

Sometimes checklists have an option to specify “Not relevant” to a given question, and it could be worth considering adding this to the BSIMM yardstick as well. Looking at this from a different angle, maybe an organization should first establish the set of software security activities that represent the “holy grail” for them, i.e., the 113 minus any activities deemed to be not relevant. The results should then be compared with this modified yardstick.

From a psychological point of view, it is tempting to ask if there is a threshold where a BSIMM score becomes de-motivating rather than inspiring. If an organization is “flatlining” in almost every practice, management might not even want to tell the employees. This is troublesome on many levels, not least if it leads to the assessment report to be filed and promptly forgotten. If we accept that the BSIMM activities represent “good software security practice,” organizations should most likely strive to implement more activities; simply ignoring the immaturity problem does not make it go away.

7.6 Conclusion

The BSIMM Software Security Framework represents a comprehensive list of good practice software security activities which is a good foundation to build a software security program in a development organization. It may not be possible to replicate the BSIMM study method as it is done by Cigital, but even a questionnaire-based approach can produce useful results when studying software security practices in the real world.

Acknowledgment

This work was supported by the SoS-Agile: Science of Security in Agile Software Development project, funded by the Research Council of Norway, grant number 247678.

References

- [1] Konstantin Beznosov and Philippe Kruchten. “Towards agile security assurance.” In: *Proceedings of the 2004 New security paradigms workshop*. ACM. 2004, pp. 47–54.
- [2] Petra M. Boynton and Trisha Greenhalgh. “Selecting, designing, and developing your questionnaire.” In: *BMJ* 328.7451 (May 2004), pp. 1312–1315.
- [3] J. T. Fitzgerald, C. B. White, and L. D. Gruppen. “A longitudinal study of self-assessment accuracy.” In: *Med Educ* 37.7 (July 2003), pp. 645–649.
- [4] Martin Gilje Jaatun. “Hunting for Aardvarks: Can Software Security Be Measured?” English. In: *Multidisciplinary Research and Practice for Information Systems*. Ed. by Gerald Quirchmayr et al. Vol. 7465. Lecture Notes in Computer Science. Springer Berlin Heidelberg, 2012, pp. 85–92. ISBN: 978-3-642-32497-0. DOI: 10.1007/978-3-642-32498-7_7. URL: http://dx.doi.org/10.1007/978-3-642-32498-7_7.
- [5] Martin Gilje Jaatun et al. “Software Security Maturity in Public Organisations.” English. In: *Information Security*. Ed. by Javier Lopez and Chris J. Mitchell. Vol. 9290. Lecture Notes in Computer Science. Springer International Publishing, 2015, pp. 120–138. ISBN: 978-3-319-23317-8. DOI: 10.1007/978-3-319-23318-5_7. URL: http://dx.doi.org/10.1007/978-3-319-23318-5_7.
- [6] Gary McGraw, Sammy Miguez, and Jacob West. *Building Security In Maturity Model (BSIMM 7)*. <http://bsimm.com>. 2016.
- [7] Gary McGraw, Sammy Miguez, and Jacob West. *Building Security In Maturity Model (BSIMM) Version 6*. 23/11/2015. URL: <https://www.bsimm.com/download/> (visited on 11/23/2015).
- [8] Hyeun-Suk Rhee, Young U. Ryu, and Cheong-Tag Kim. “Unrealistic optimism on information security management.” In: *Computers & Security* 31.2 (2012), pp. 221–232. ISSN: 0167-4048. DOI: <http://dx.doi.org/10.1016/j.cose.2011.12.001>. URL: <http://www.sciencedirect.com/science/article/pii/S0167404811001441>.
- [9] Colin Robson. *Real World Research*. 3rd ed. John Wiley & Sons, 2011.
- [10] Paul E. Spector. “Using Self-Report Questionnaires in OB Research: A Comment on the Use of a Controversial Method.” In: *Journal of Organizational Behavior* 15.5 (1994), pp. 385–392. ISSN: 08943796, 10991379. URL: <http://www.jstor.org/stable/2488210>.