

Studying IT Security Professionals: Research Design and Lessons Learned

David Botta, Rodrigo Werlinger, André Gagné
Konstantin Beznosov, Lee Iverson, Sidney Fels, Brian Fisher
University of British Columbia, Vancouver, Canada

{botta,rodrigo,andreg,beznosov,leei,sfels}@ece.ubc.ca, fisher@cs.ubc.ca

ABSTRACT

The HOT Admin Field Study used qualitative methods to study information technology security professionals. Both the nature of the field and the difficulty of gaining access to participants had implications for the study design. We present the lessons learned, and offer suggestions overcoming the challenge of participant recruitment.

1. INTRODUCTION

Information technology security management (ITSM) in large organizations is exceptionally challenging due to the increasingly high numbers of application instances, resources, and users interacting with business processes that are growing in complexity. Yet little is known about how individuals perform ITSM, their roles and responsibilities within organizations, and how effective their existing tools and practices are at protecting organizations and employees while still allowing productive collaborative work [10].

This field study is the first phase of the project *HOT Admin: Human, Organization, and Technology Centred Improvement of IT Security Administration*.¹ The project investigates methods and techniques for developing better tools for managing IT security from the perspective that human, organizational and technological influence the ability of security practitioners to do their job well.

Our approach in this ongoing field study is to obtain stories of IT security practitioners' daily interaction with tools, communication with other people, and deployment of technologies. A high-level view of the organization would not have revealed the gritty details involved, while human-factors studies of tool interfaces would have considered neither the technical ecology nor the organizational pressures.

This paper describe methods of the study, outlines the challenges that we have encountered so far, and shares the lessons learned. We also discuss alternative approaches.

¹hotadmin.org

2. STUDY DESIGN

The output of the study was intended to be used by the later phases of the HOT Admin project, which are (1) devising a methodology for evaluating the effectiveness of IT security management tools, and (2) developing and evaluating techniques for designing more effective tools for IT security professionals. To focus the study design, we developed a hierarchy of research questions. The questions turned out to require both macro- and micro-level information. The macro-level information concerned (1) demographics such as the level of a participant's education and the size of their organization; (2) the approach to achieving organizational goals; and (3) the types of tools used. The micro-level information comprises detailed examples of the relationships between tool use, tasks, and organizational pressures. Our approach was to obtain stories of IT security practitioners' daily interaction with tools, communication with other people, and deployment of technologies. A high-level view of the organization would not have revealed the gritty details involved, while tool usability studies would not have considered the technical ecology and organizational environment. To gain insight into these issues, we decided to use qualitative research methods.

To see the use of tools in action, we needed a level of data granularity like that demonstrated in Maglio et al.'s [8] observation of a problem-solving episode in administering a web application. They used a distributed cognition approach, in which they paid particular attention to the representation of information as it propagated from one medium to another across a network of people and systems. Like Maglio et al., we were also interested in how people construct common understanding in order to solve problems, as discussed by Clark [3]. This level of granularity requires work shadowing. However, particularly with security, illustrative events are not likely to avail themselves to the convenience of researchers. In order to capture such events, a researcher would have to be present for extended periods of time, which was not feasible in our case. Therefore we adopted the approach called contextual interview [1]. Nevertheless, a close up view does not necessarily reveal the goals that people have in mind. In order to learn about how security practitioners use their tools to achieve their goals [2], we needed to conduct semi-structured interviews.

We employed a pre-interview questionnaire and a semi-structured interview. We also plan to conduct contextual interview with some of the study participants.

The pre-interview questionnaire (10 to 15 minutes in length) provided information about skill and training, what tools were preferred, and enabled us to tailor the semi-structured interviews to individuals. During these interviews (1 to 1.5 hours in length), the participants were encouraged to tell stories and give examples. As security practitioners are busy, we did not want to request too much at one time. Therefore, we took a graduated approach. The first contact letter only asked the participant to complete the short questionnaire. One of the questions in the questionnaire was whether the participant would be willing to be personally interviewed. After the semi-structured interview, some participants were asked whether they would like to participate in a contextual interview.

We sought both complexity and diversity in our target organizations. Our success in achieving diversity was only partial, because participants were difficult to recruit. Kotulic and Clark also point at this difficulty [6]. We approached post-secondary educational institutions, research organizations, financial, insurance, and energy organizations in Greater Vancouver, Canada, for our study. Most of our participants were from academic or research organizations. Especially with commercial organizations, our experience showed that recruiting participants who are both able and willing to talk about IT security involves a long, slow process of relationship building.

Both with transcription of interviews and analysis of interviews, the team had to develop a rigorous procedure of handling the files in such a way that their confidentiality was always observed. For the analysis of the interview data, only sanitized transcripts were used. Names and other references that could be used for identification were replaced with random 4-digit numbers. The research team shared interview records and transcripts in encrypted form. The master copies of the audio files were kept on encrypted discs in a locked office, with the understanding that the copied would eventually be deleted.

2.1 Recruitment

We observed three key challenges in recruiting participants: (1) participation in the study was seen by the chronically overworked IT professionals, and especially by their supervisors, as an uncompensated burden, (2) the potential disclosure of IT security procedures, practices, and even tools in use went against common organizational culture of carefully restricting outside parties access to such details, and (3) since our participants were the backstage people whose contact information was not published on the company web sites or other publicly accessible sources, just finding ways to make first contact with them would be next to impossible without buy-in from the gatekeepers, i.e., management personnel.

To address the first challenge, we developed a graduated recruitment strategy so that the work burden was minimal to begin with. We initially asked potential participants only to answer a short questionnaire, the final question of which asked if the participant is willing to give a one-hour interview. At the end of this contextual interview, we asked some participants if they would be willing to allow us to shadow them in their workplace.

Graduated recruitment also helped in building trust between participants and the researchers in order to overcome the second challenge. We actively educated potential participants about the purely academic (i.e., noncommercial) and worthwhile goals of the HOT Admin project and the study itself. In addition, prior background of the principle investigator as a security professional himself seemed to aid with both (a) building trust through speaking the language (and jargon) of IT security, (b) developing professional contacts.

To address the third challenge, we used two approaches. Some participants were recruited directly, through professional contacts of the research team. Project team members developed and maintained such contacts by participating in the meetings of a regional security special interest group and presenting at a regional forum for IT security professionals. Although professional contacts ended up being most effective in the recruitment, they were too few.

To recruit other participants, we contacted managers of IT departments and met or interviewed them to solicit their cooperation. With their cooperation, we asked for recommendations of employees they felt would be knowledgeable and/or were involved with security management in their organization. In all cases, we obtained—directly or through the participants—management permission before involving our participants in the study.

Once identified, we contacted participants by e-mail. Our letter of first contact contained a brief description of the project and its goals, its policy about the privacy of the participants and the confidentiality of the collected data, and an invitation to complete the online questionnaire.

Our gaining of approval for the interviews was necessarily top down. Personal introductions were often mediated by IT managers. The request for participation could therefore easily be perceived as coercive. Also, we had to be particularly careful about securely handling the data, and to tell our participants how we went about it.

2.2 Questionnaire

The questionnaire was included in the tail of the e-mailed first contact letter. Interested participants responded by replying with the completed answers within the body of the e-mail, or clicking on a link with the web version of the questionnaire. We wanted to provide this simple and convenient interface for responding to the questionnaire, because we expected the potential participants to be unable to devote much time or attention to a questionnaire [5].

The pilot-tested questionnaire had 21 questions ranging from general background and responsibilities to questions about the IT system and security management, in addition to requesting participation in the follow-up interview. The questionnaire was not intended to gather quantitative data; rather, it was used to gather information that would help us better focus the semi-structured interview. For example, if in the questionnaire the participant mentioned interacting significantly with other individuals in the organization, we would be alerted to ask about the nature of these interactions. The questionnaire was not helpful for predicting which participants would give the most valuable interviews.

2.3 Semi-Structured Interview

The semi-structured interview allowed participants to tell stories that provided information beyond the current situation or time-frame. The interviewer had the opportunity to inquire about a wide range of aspects of security management, from minute routine details to long-term goals.

The following is a small sample of the questions comprising our semi-structured interview:

- What did you do yesterday?
- How do you interact with different types of people during the course of your work?
- Is there anything special about your organization that makes IT security management more difficult; for example, a rapid turnover of users, or special relationships with other organizations, or something else?
- What do you wish for in your tools?

Depending on the job roles that the interviewees played, we found it useful to quickly move to the topic of tools because stories about tool use (1) tended to be detailed and concrete, and (2) led easily into detours concerning communication with other people, prioritization of tasks, and organizational idiosyncrasies.

2.4 Contextual Interview

We also plan to conduct contextual interviews—a.k.a. “work shadowing”—with some of our participants. In a contextual interview, the participant is expected to “teach” the work and correct the interviewer’s misunderstandings. The interviewer, like an inquisitive apprentice, may request the participant to interrupt work in order to explain something.

Contextual interviews are normally video recorded. However, we plan not to employ video for the following reasons: (1) it would physically intrude into the workspace and might cause the participant or the participant’s coworkers to feel uncomfortable; (2) video is both low resolution and likely to resonate in a bad way with the refresh rate of monitors, and is therefore of questionable value in capturing the participant’s microscopic interaction with a tool; (3) such interactions are better captured in a usability study; and (4) a combination of audio, maps, diagrams and still photographs could give us the information we needed.

We suspect that due to the nature of our participants and their jobs observation of their activities might be either out of the question or of dubious worth, since most of our participants spend a lot of time emailing, attending meetings, or doing tasks that have nothing to do with IT security. The sporadic and unpredictable nature of security incidents makes them difficult to document. On many occasions, observation might be of little worth; for example, if a systems administrator checks email to discover 200 automated email messages from servers, and must check the messages against a list of the servers to see if any of the servers are missing, there is not a lot more that observation can reveal about this situation. But there is still hope that in a few cases a contextual interview would be worth while. For example, it

could be enlightening to see a security professional training an intrusion detection system to not send out thousands of false alarms.

At the time of writing, we have not yet performed a contextual interview. Nevertheless, two people have agreed to a contextual interview, out of the 14 we have interviewed. These two are from academic institutions.

2.5 Data Analysis

The discursive nature of the data we collected, combined with a lack of pre-existing theories of ITSM *per se* suggested a bottom-up approach such as Grounded Theory (GT) [4]. Nevertheless, because of practical issues, e.g., those pointed by Locke [7], the presence of a good deal of existing theory, and that observation alone would be unlikely to reveal the social and organizational factors characteristic of the organization, in our methodology we adapted GT to take into account our understanding of the security administration tools and tasks together with a general framework for social cognition, Clark’s [3] theory of psycholinguistic pragmatics.

We filtered the data according to themes, and then practiced GT on the filtered data. For example, we read the transcripts focusing on the themes of tools, tasks, responsibilities, tool pros and cons. Then, with the passages related to each theme, we developed concepts using the GT method of open and then axial coding. The coders worked together on two transcripts, and thereafter worked independently with randomly assigned transcripts. Each analysis was cross checked during meetings, where the codes and categories used for analysis were scrutinized by the project team, and by having selected interviews re-coded by another researcher for comparison, and any differences discussed and rationalized. This triangulation process resulted in refinements, rather than dramatic changes.

3. RESULTS TO DATE IN BRIEF

The job of security management is distributed across multiple employees, often affiliated with different organizational units or groups within a unit and responsible for different aspects of it. Our participants generally would keep a toolkit of system-specific configuration management and monitoring tools as well as tailorable, generic text and information management tools (e.g., grep, shell scripts and e-mail). Their reports of the shortcomings of existing tools focused on need for *tailorability*. They used their toolkits in creative ways to accomplish different tasks in various scenarios. They frequently wrote scripts to complement the functionality of their tools (e.g., Snort), or to perform specific tasks (e.g., analysis of logs, correlation of events). The output of the tools was usually filtered and re-filtered, and compared with output from other sources. While performing security-related tasks, they normally engaged the skills of inferential analysis, pattern recognition, and *bricolage*, as well as design and good communication.

4. SUMMARY OF LESSONS LEARNED

To summarize our lessons learned, gaining entry to organizations required building familiarity, trust, and a sense that our research is worthwhile to the domain of ITSM. A top-down recruitment strategy was necessary; we needed a letter of endorsement from management. Further snowballing

was also helpful. This altogether entailed making assurances that (1) there would be no consequences for refusal; (2) confidentiality would include not revealing anything to management, even whether the prospect actually participated; (3) the participating organizations' and individuals' identities and any other information that could make them vulnerable would be kept confidential; and (4) how the data confidentiality was provided was explained fully.

The questionnaire was not helpful for predicting which participants would give the most valuable interviews. Although the response rate to our questionnaire was reasonable, recruitment for semi-structured interviews was disappointing. Rather than banking heavily on new recruits, we hope to also be able return to some participants with deeper questions based on emerging theory. During the course of semi-structured interviews, we learned to delve quickly and deeply into questions about the use of tools, because we could draw out *detours* on issues of organizational idiosyncrasies, and user interface issues. Recruitment of participants for work shadowing, in our case in the style of contextual interview, was substantially more poor than recruitment for semi-structured interviews. Nevertheless, both our research questions and arising theory indicate that this is a necessary technique for achieving objectives of our study. Despite interviewing about one participant per week, it was difficult to keep up GT analysis to refocus questions on emerging theory. A mixed thematic and GT approach to analysis appeared to address our research questions.

5. CONCLUSIONS

We invite IT security tools' developers and researchers to consider that the IT security tools must be both tailorable and survive in an arena of *bricolage*. Our participants used tools in different situations that arise out of the complexity of both the technology and the environment. By *testing the use of suites of tools* in real-world task scenarios, one may uncover usability problems that result from differences in assumptions of the various tools, and from difficulty in transferring information between them in order to coordinate a particular task.

We would like to offer the following ways for working around the difficulty of recruiting participants, which so far has been the biggest challenge for us:

- The researcher works with a security tool vendor who recruits two or more of its clients. The researcher employs contextual interviews and other ethnographic techniques to study the clients' use of specific tool(s). Generality would be compromised for depth of information. Intellectual property rights, and what constitutes a valid and original contribution to the body of knowledge, would be issues.
- Perform a longitudinal case study evaluation of security expenditures of an organization. The organization gains the evaluation, and approves what the researcher may publish. (See [11].) Validity of the data would be compromised for breadth of information.
- Security professionals participate in play-acting workshops in which they author, direct, and act in scenarios that illustrate a wide variety of issues that bear on

everyday IT security management, similarly to what Newell et al. did in their usability investigations for senior citizens [9].

Acknowledgments

The HOT Admin project is funded by the Canadian NSERC Strategic Partnership Program, grant STPGP 322192-05. The authors are grateful to all those IT professionals who donated their time and energy to participate in the field study reported in this paper, as well as to Craig Wilson for improving the readability of the paper.

6. REFERENCES

- [1] Hugh Beyer and Karen Holtzblatt. *Contextual Design, Defining Customer-Centered Systems*. Morgan Kaufmann Publishers, San Francisco, CA, 1998.
- [2] S. Bodker. Human activity and human-computer interaction. In S. Bodker, editor, *Through the Interface: A Human Activity Approach to User Interface Design*, pages 18–56. Lawrence Erlbaum Associates, Publishers, Hillsdale, NJ, 1991.
- [3] H. H. Clark. *Using Language*. Cambridge University Press, Cambridge, England, 1996.
- [4] Barney Glaser and Anselm L. Strauss. *The Discovery of Grounded Theory, Strategies for Qualitative Research*. Aldine Publishing Company, Chicago, Illinois, 1967.
- [5] Eser Kandogan and Eben M. Haber. Security administration tools and practices. In Lorrie Faith Cranor and Simson Garfinkel, editors, *Security and Usability: Designing Secure Systems that People Can Use*, chapter 18, pages 357–378. O'Reilly Media, Inc., Sebastapol, 2005.
- [6] Andrew G. Kotulic and Jan Guynes Clark. Why There Aren't More Information Security Research Studies. *Information & Management*, 41(5):597–607, 2004.
- [7] Karen Locke. *Grounded Theory in Management Research*. SAGE Publications, Thousand Oaks, CA, 2001.
- [8] P. P. Maglio, E. Kandogan, and E. Haber. Distributed cognition and joint activity in collaborative problem solving. In *Proceedings of the Twenty-fifth Annual Conference of the Cognitive Science Society*, 2003.
- [9] A. F. Newell, A. Carmichael, M. Morgan, and A. Dickinson. The use of theatre in requirements gathering and usability studies. *Interacting with Computers*, 18:996–1011, 2006.
- [10] E. E. Schultz, R. W. Proctor, M. C. Lien, and G. Salvendy. Usability and Security An Appraisal of Usability Issues in Information Security Methods. *Computers and Security*, 20(7):620–634, 2001.
- [11] Detmar W. Straub and Richard J. Welke. Coping with Systems Risk: Security Planning Models for Management Decision Making. *MIS Quarterly*, 22(4):441–469, 1998.