Challenges, Collaborative Interactions, and Diagnosis Performed by IT Security Practitioners: An Empirical Study

by

Rodrigo Werlinger

Electrical Engineer, Universidad de Chile, 2000

A THESIS SUBMITTED IN PARTIAL FULFILMENT OF
THE REQUIREMENTS FOR THE DEGREE OF

Master of Applied Science

in

The Faculty of Graduate Studies

(Electrical and Computer Engineering)

The University of British Columbia

(Vancouver)

July, 2008

Abstract

This thesis investigates four different aspects of information security management: challenges faced by security practitioners, interactive collaborations among security practitioners and other stakeholders, diagnostic work performed by security practitioners during the response to incidents, and factors that impact the adoption of an intrusion detection system in one organization. Our approach is based on qualitative analyzes of empirical data from semi-structured interviews and participatory observation. For each theme under study, the contributions of the qualitative analysis are twofold. First, we provide a richer understanding of the main factors that affect the security within organizations. Second, equipped with this richer understanding, we provide recommendations on how to improve security tools, along with opportunities for future research.

Our findings contribute to the understanding of the human, organizational, and technological factors that affect security in organizations and the effectiveness of security tools. Our work also highlights the need for continued refinement of how factors interplay by obtaining more rich data (e.g., contextual inquiry), and the need to generalize and validate these findings through other sources of information to study how these factors interplay (e.g., surveys).

Contents

Abstract ii

Contents iii

List of Tables viii

List of Figures ix

Acknowledgements xi

Dedication xii

Statement of Co-Authorship xiii

1 Introduction 1

1.1 Motivation 1

1.2 Methodological Approach 2

1.3 Related work 4

1.3.1 Challenges to IT security 4

1.3.2 Collaborative interactions 5

1.3.3 Diagnostic tasks 6

1.3.4 Deployment of an IDS 6

1.4 Contributions 6

1.5 Structure 9

Bibliography 10

2 Challenges for security practitioners . . . 14
2.1 Introduction . . . 14
2.2 Background . . . 15
2.2.1 Human factors . . . 15
2.2.2 Organizational factors . . . 16
2.2.3 Technological factors . . . 17
2.3 Methodology . . . 17
2.4 Building an Integrated Framework of Challenges . . . 19
2.4.1 Human factors . . . 19
2.4.2 Organizational factors . . . 20
2.4.3 Technological factors . . . 22
2.5 Discussion . . . 23
2.5.1 Cross analysis . . . 23
2.5.2 A holistic view of challenges and their interrelationships . . . 24
2.5.3 Opportunities for future research . . . 24
2.6 Conclusion . . . 28

Bibliography . . . 29

3 Interactions with other stakeholders . . . 32
3.1 Introduction . . . 32
3.2 Related Work . . . 34
3.2.1 Empirical Research on IS Collaborative Work . . . 35
3.2.2 Summary . . . 36
3.3 Research Methods . . . 37
3.3.1 Participant Recruitment . . . 37
3.3.2 Data collection from multiple sources . . . 38
3.3.3 Data Analysis . . . 39
3.4 Analyzing Interactions in Context . . . 39
3.4.1 Activities Requiring Interactions with Other Stakeholders . . . 40
3.4.2 Communications Channels Used during Interactions . . . 42

3.4.3 Security Tools Used within the Context of Interactions . . . 43
3.5 Interaction Scenarios . . . 45
3.5.1 Interactions in Responding to Security Incidents . . . 45
3.5.2 Development of Policies . . . 50
3.6 Modeling the Complexity of Interactions . . . 52
3.6.1 Organizational Attributes . . . 54
3.6.2 Multiple stakeholders . . . 55
3.6.3 Multiple security-related activities . . . 56
3.6.4 Consequences of the complexity of security interactions . . . 57
3.7 Implications of Findings . . . 58
3.8 Conclusion . . . 61

Bibliography . . . 64

4 Diagnosis of Security Incidents . . . 67
4.1 Introduction . . . 67
4.2 Related Work . . . 69
4.2.1 ITSM: Background . . . 70
4.2.2 Guidelines for Security-Incident Response . . . 71
4.2.3 Diagnostic Work during Security Incidents . . . 72
4.3 Methodology . . . 73
4.3.1 Data Collection . . . 74
4.3.2 Data Analysis . . . 75
4.4 Results . . . 75
4.4.1 Preparation Phase . . . 76
4.4.2 Identification Phase . . . 80
4.5 Discussion . . . 86
4.5.1 How Security Practitioners Diagnose Security Incidents . . . 87
4.5.2 Opportunities for improving IT security technology . . . 88
4.6 Conclusion . . . 92

Bibliography . . . 93

5 Case study on an Intrusion Detection System . . . 97

5.1 Introduction . . . 97

5.2 Related work . . . 99

5.2.1 IDS Phases . . . 99

5.2.2 IDS Usability Challenges . . . 100

5.2.3 Support and Evaluation . . . 101

5.3 Methodology . . . 101

5.3.1 Data collection . . . 102

5.3.2 Data analysis . . . 103

5.4 Anatomy of an IDS . . . 104

5.4.1 The Deployed IDS . . . 105

5.5 Investigating IDS Usability . . . 105

5.5.1 Issues Deploying an IDS . . . 106

5.5.2 Advantages and Disadvantages of IDSs . . . 111

5.6 Discussion . . . 114

5.6.1 Considerations before deploying an IDS . . . 114

5.6.2 Configuring and Validating an IDS . . . 115

5.6.3 Ongoing Usage . . . 117

5.7 Conclusion . . . 119

Bibliography . . . 121

6 Conclusions . . . 125

6.1 Contributions . . . 125

6.1.1 Challenges theme . . . 125

6.1.2 Interactions theme . . . 126

6.1.3 Diagnosis theme . . . 127

6.1.4 Deployment of an IDS theme . . . 127

6.2 Applications . . . 128

6.3 Limitations 128
6.4 Future work 129

Bibliography 130

Appendices

A Coding examples for the Challenges theme 139
A.1 Examples for the list of challenges/factors that affect security practitioners 139
A.2 Interplay among challenges/factors 150

B Detailed notes from participatory observation 156
B.1 Deciding on the Purpose of the IDS 156
B.2 Integrating the IDS in the Network 157
B.3 Initial configuration of the IDS 158
B.4 Effectiveness of the Graphical User Interface 160
B.5 Configuring for Multiple Stakeholders 161

C UBC Research Ethics Board's Certificate 162

List of Tables

2.1 Profile of our participants and their organizations 18

2.2 Challenges participants described for implementing security controls 20

3.1 For each type of organization, we indicate the number of unique organizations and the total number of participants interviewed. These participants held various positions, including Managers (with security tasks), IT Practitioners (with security tasks), Security Managers, and Security Specialists. 38

3.2 Types of activity in which IT-security communication occurs 63

4.1 For each type of organization, we indicate the number of unique organizations and the total number of participants interviewed. These participants held various positions, including IT Managers (with security tasks), Security Managers, Security Specialists and IT Practitioners (with security tasks). 75

4.2 Sequence of tasks to respond to security incidents. 87

5.1 Participant Information (Semi-Structured Interviews) 103

A.1 Table that shows the axial and open codes. The name of the open codes match the challenges/factors that affect security practitioners 139

A.2 Table with quotes that describe the interplay among challenges/factors 150

List of Figures

2.1 A holistic view of challenges and their interrelationships. These interrelationships were extracted from the stories of our participants (see Appendix, Section A.1). . . 25

3.1 Responding to security incidents. Thicker arrows indicate more frequent interactions. For simplicity, only interactions between security practitioners and other stakeholders are shown. 47

3.2 Response to an incident that triggers multiple and complex interactions among stakeholders. Dashed lines indicate two possible actions depending on the cooperation from the client. End-users are behind other agents, clients of the participant's organization. 49

3.3 Communication flow diagram for developing security policies. Thicker arrows indicate more frequent interactions. For simplicity, only interactions with security practitioners are shown. 51

3.4 Factors that make interactions more complex for security practitioners within organizations. 53

4.1 Diagram of the network connections for the IDS. The IDS has one connection to the monitored network and a second to the management network. 77

4.2 Error message when the license of the IDS could not be validated 78

4.3 Adaptation of the flow communication diagram showed in previous Chapter (Chapter 3) with the collaboration among different stakeholders to respond to a security incident. We now highlight the diagnostic aspects of such collaboration, including the monitoring tasks performed by our participants on the organization's systems. Thicker arrows indicate more frequent collaboration. For simplicity, only collaboration between security practitioners and other stakeholders is shown. 82
4.4 Forces influencing tool reliability (false positive rate) 89

5.1 System configuration options of the IDS in the back. On top, Configuration options of the IDS's rules (bottom right) and status of alarms. 104
5.2 Network diagram used during one discussion about the installation of the IDS. The IDS has a connection to the management network and another to the port of the switch that transports internal traffic from the firewall. To compare configuration of the firewalls, it would be necessary to include another connection to the external traffic (dashed line). 108

C.1 UBC Research Ethics Board's Certificate 163
C.2 Amendment UBC Research Ethics Board's Certificate, page 1 164
C.3 Amendment UBC Research Ethics Board's Certificate, page 2 165
C.4 UBC Research Ethics Board's Certificate, first renewal 166
C.5 UBC Research Ethics Board's Certificate, second renewal 167

Acknowledgements

I would like to thank all my colleagues from the Laboratory for Education and Research in Secure Systems Engineering (LERSSE) for their constant support and feedback, and the participants who took part in our study. Special thanks to my supervisor Konstantin Beznosov and my friends and mentors Kirstie and Kasia. They have done their best to teach me not only how to be a good professional, but also a better person.

I do not have words to express my gratitude to my family, Gloria, Mila and Igor, for their love and support. They are the best that could have ever happened to me. I also want to thank my parents–in–law for their support, specially my mother–in–law for devoting one year of her life to help us survive the process.

To My Family: Gloria, Mila and Igor

Statement of Co-Authorship

A version of each chapter of this thesis has been either published, accepted, or submitted for publication. The author of this thesis performed all the qualitative analysis shown in chapters 2, 3 and 4. He also authored the corresponding papers, under the supervision of the co-authors who provided feedback and guidance throughout the research process. In chapter 5, the qualitative analysis was performed with Kasia Muldner. Chapter 6 includes a section that was initially written by Kirstie Hawkey for chapter 3. Below are the details for each chapter.

- Chapter 2: A version of this chapter has been accepted for publication. The author of this thesis wrote all the sections of this chapter.

  R. Werlinger, K. Hawkey, and K. Beznosov. Human, Organizational and Technological Challenges of Implementing IT Security in Organizations. *In HAISA '08: Human Aspects of Information Security and Assurance (13 pages, to appear)*, July 2008.

- Chapter 3: A preliminary version of this chapter has been published. A full version of this chapter has been submitted to a journal for publication. The author of this thesis wrote all the sections of this chapter.

  R. Werlinger, K. Hawkey, and K. Beznosov. Security practitioners in context: their activities and interactions. In CHI '08 extended abstracts on Human factors in computing systems, pages 3789—3794, 2008.

  R. Werlinger, K. Hawkey, and K. Beznosov (2008) Security Practitioners in Context: Their Activities and Interactions with Other Stakeholders Within Organizations.

- Chapter 4: A version of this chapter has been submitted to a journal for publication. The author of this thesis wrote all the sections of this chapter, except for sections 4.1 and 4.2.

R. Werlinger, K. Muldner, K. Hawkey, and K. Beznosov (2008) Diagnostic Work during Security Incident Response: A Qualitative Study.

- Chapter 5: A version of this chapter has been accepted for publication. The author of this thesis wrote sections 5.3, 5.4.1, 5.5.1, and Appendix B, which contains the data from participatory observation, used for most of the qualitative analysis of this chapter.

  R. Werlinger, K. Hawkey, K. Muldner, Pooya Jaferian, and K. Beznosov. The challenges of using an intrusion detection system: Is it worth the effort? In Proc. of ACM Symposium on Usable Privacy and Security (SOUPS) (12 pages, to appear), July 2008.

- Chapter 6: Section 6.3 was adapted by the author of this thesis from an initial version written by Kirstie Hawkey for chapter 3. The new version of this section acknowledges the limitations of the qualitative analysis presented in the whole thesis.

All these publications are within the context of HOT Admin project. This project was proposed and designed by Konstantin Beznosov, Sidney Fels, Lee Iverson, and Brian Fisher to investigate how Human, Organizational and Technological factors affect security in the organizations.

A qualitative approached was undertaken by the HOT Admin project in order to reach a rich understanding of IT security management. One part of the qualitative data consists of semi-structured, *in-situ* interviews with security practitioners. The author of this thesis participated as an interviewer in two of these interviews (I20 and I32 [1]). The other source of qualitative data, used for the analysis shown in Chapters 3, 4, and 5 comes from participatory observation in one Canadian Academic Organization. This activity was performed completely by the author of this thesis.

[1] I32 is I29 in Chapter 3

Chapter 1

Introduction

This thesis studies four aspects related to the work of security practitioners within organizations: their challenges, collaborative interactions with other stakeholders, diagnostic work during response to security incidents, and factors that affect the adoption of an intrusion detection system (IDS).

1.1 Motivation

The study of information security needs to consider not only technical aspects, but also human and organizational aspects, in order to understand how to improve security levels (Beznosov and Beznosova 2007; Botta et al. 2007b; Kotulic and Clark 2004). A central actor within the human, organizational, and technical dimensions is the security practitioner, the IT professional who has security responsibilities in the organizations (e.g., managing firewalls). However, to date there is little empirical evidence about security practitioners' work and what activities they perform to protect the organizations (Botta et al. 2007b; Björck 2005).

Prior research has highlighted several key areas where a better understanding is required to improve the support provided to security practitioners. First, although there are studies that investigate how specific organizational factors such as size and sector impact the effectiveness of information security (IS) (Kankanhalli et al. 2003; Chang and Ho 2006), there does not exist an integrated framework that comprises the elements that make security tasks challenging for security practitioners. Such a framework could help organizations identify their limitations with respect to implementing security standards as well as determine if they are spending their security resources effectively. Second, security tasks are highly collaborative in nature, and security tools do not provide support for that level of collaboration (Botta et al. 2007b; Goodall et al. 2004a; Kandogan and Haber 2005). However, a current lack of a rich understanding about how security practitioners interact and communicate with other stakeholders makes it difficult for HCI researchers and tool

developers to improve communication and IT security tools. Third, despite the fact that security practitioners perform intensive diagnostic work during response to security incidents, few studies have investigated how to improve security tools to provide better support to this type of diagnostic task (for exceptions see Goodall et al. (2004a); Riden (2006)). To improve the support given to security practitioners when performing diagnostic tasks, it is necessary to understand how security practitioners perform these tasks and the factors that impact the diagnoses of security issues.

This thesis aims at filling these gaps highlighted by the prior research, by investigating how these three aspects i.e., challenges, interactions and diagnostic tasks, impact the work of security practitioners, and the development and evaluation of security tools. This thesis also investigates in detail one specific security tool, an intrusion detection system (IDS). Specifically we analyze the human, organizational and technological factors that influence the adoption of an IDS in an organization.

1.2 Methodological Approach

The methodology we use in this thesis is based on a qualitative approach. This decision comes from the lack of previous empirical studies showing how different aspects of security management affect the design of security tools.

To obtain rich empirical data on the work of security practitioners and their tools, we needed a level of granularity like that demonstrated in Maglio et al.'s (2003) 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 (1996). 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 because of the recruitment constraints mentioned in the next paragraph. Therefore we aimed at collecting data via contextual interviews (Beyer and Holtzblatt 1998). 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 (Bodker 1991), we needed to conduct semi-structured interviews.

The qualitative data described above have been collected in the context of HOT Admin project (see Hawkey et al. (2008) for an overview of this project). This project has provided two sources of data: questionnaires and semi-structured in-situ interviews with security practitioners from both the academic and private sectors. The difficulties associated with recruitment of participants for the study of information security (Botta et al. 2007a; Kotulic and Clark 2004) imposed two limitations on the data provided by the project. First, recruitment difficulties meant that we were not able to be selective in the recruitment process (e.g., target specific kinds of security professionals as type of organization, age, etc.) Second, it was not possible to collect data using contextual interviews, due to the recruitment difficulties. These limitations were partially addressed by the opportunity that the thesis author had to perform participatory observation (Fetterman 1998) in one academic organization in Canada. The thesis author recognized the value of this opportunity to collect as much data as possible, in order to supplement the data from the interviews in two specific topics: development of security policies and deployment of an intrusion detection system.

The analysis of the data was performed considering pre-designed analysis themes. These themes are directly related to each area under study, e.g., challenges, interactions, diagnosis during security incidents, and implementation of an IDS. To perform the analysis for each theme, qualitative description (Sandelowski 2000) was used. The alternative analysis technique is Grounded Theory (GT) (Charmaz 2006), which is particularly appropriate when there is a lack of formal theory related to the analysis themes, as is the case with the data here. However, the thesis author chose qualitative description instead of GT due to the above-cited limitations related to the type of data that was available. Specifically, given the recruitment difficulties, to avoid losing participants, the HOT Admin project strategy was to contact potential participants and perform subsequent interviews as soon as individuals expressed interest in participants. This strategy meant that the analysis of the data from the interviews was well behind the collection of the interviews; we started the analysis when 14 interviews had already been performed. This delay did not allow for the application of theoretical sampling interview by interview (interview questions were adjusted only three times, before interviews 15, 22, and 27), limiting the dynamic process of building theories as suggested in GT. Despite this, the qualitative description method used in this thesis

has GT overtones, as both methods rely strongly in constant comparison (Sandelowski 2000). The difference in terms of findings is that the results of our analysis are very close to the data provided by our participants, and do not necessarily correspond to high-level interpretations of the data, as is usually the case in GT.

The qualitative analysis of each chapter was made as follows. For Chapters 2 and 3, the data were coded iteratively, starting with open coding and continuing with axial coding. Posterior analysis was based on further elaboration of "memos" (Charmaz 2006) written during the coding process. For Chapters 4 and 5, the qualitative analysis was based on narrative analysis of the interview transcriptions and notes from participatory observation, to identify excerpts that pertained to diagnostic work and the implementation of intrusion detection systems (IDSs), respectively. The next step consisted of organizing the excerpts into different stories or "memos" (Charmaz 2006) describing diagnostic work during security incident response and challenges to deploy an IDS.

The other themes investigated in the HOT Admin project (e.g., tasks and tools, errors) were analyzed by other researchers. Given that these researchers were working on the same data for the analysis, it was possible to do triangulation of the findings, by discussing the interpretations that each researcher had of the data related to his or her theme with other researchers. This triangulation was made at the level of "memos", as comparison at lower levels of analysis (e.g., open or axial codes) had shown to be time-consuming and did not result in dramatic changes in the interpretations of the data (Botta et al. 2007b).

1.3 Related work

Each chapter provides a detailed survey of related work for the theme under study. Here we summarize these relevant studies to show how the investigation of this thesis fits and contrasts with prior work.

1.3.1 Challenges to IT security

Previous research has studied separately the human, organizational, and technological factors that challenge the adoption of security within organizations. We define human aspects as those related to cognition at the individual level, as well as culture and interaction with other people.

Organizational aspects are those related to the structure of the organization, including size and managerial decisions around IT security. Technological aspects involve technical solutions such as applications and protocols.

Within human factors, prior work has studied how communication of security risks (Koskosas and Paul 2004; Tsohou et al. 2006) and human errors (Kraemer and Carayon 2007) influence security management. Kankanhalli et al. (2003) and Chang and Ho (2006), propose that the organizational factors of top management support, and size and sector of the organizations impact the effectiveness of information security. As for technological factors, Audestad (2005) suggests that technical complexity is one of the elements that limits the achievement of 100% security.

In contrast with other studies that investigate factors affecting information security separately, in Chapter 2 we aim to empirically build a framework with a comprehensive and integrated list of information security challenges. We also study how these challenges interplay and suggest research opportunities to improve security processes and technologies, considering the human and organizational factors.

1.3.2 Collaborative interactions

Prior work has examined computer supported collaborative work (CSCW) (e.g., Carroll et al. (2006); Mohammed and Dumville (2001)); these studies propose general frameworks for understanding team effectiveness, and suggest future directions to improve empirical methods and the design of systems that support collaborations.

The complexity and importance of information security has motivated specific empirical studies on collaborative work in the context of information security. Some of these studies have identified the collaborative nature of security tasks (e.g., Knapp et al. (2005) and Björck (2005)), while others have proposed the need for better collaborative features for security tools (e.g., Botta et al. (2007b) and Goodall et al. (2004a), in the context of general tasks and tools, and an IDS respectively).

Our study for this theme, presented in Chapter 4, was focused on interactions between security practitioners and other stakeholders; we adopted a qualitative approach to collect empirical data on how security practitioners interact with other stakeholders when they perform their security tasks. Based on the analysis of these data, we offer several recommendations on how to provide better support to the collaborative interactions of security practitioners.

1.3.3 Diagnostic tasks

The study of security incidents has focused on the development of guidelines for incident response (Casey 2002; Stephenson 2004; Mitropoulos et al. 2006) and informal case studies that typically involve incidents from only one organization (e.g., Riden 2006; Schultz 2007). A key aspect of incident response is intrusion detection. Studies suggest that intrusion detection work is challenging due to its highly collaborative nature, which drives the need for analysts to coordinate with other stakeholders (Goodall et al. 2004a;b).

Our study on the diagnostic work of security incidents, presented in Chapter 5, includes the qualitative analysis of semi-structured interviews with 13 security practitioners who performed diagnostic tasks from 7 different organizations. This analysis shows the intensive use of diagnostic work by security practitioners not only during the detection, but also during the investigation of security incidents. We also describe the tasks, skills, and tools that were discussed by participants in their stories. Finally, we identify opportunities to improve the support that security tools provide during diagnostic tasks.

1.3.4 Deployment of an IDS

Prior work has centered on describing some of the challenges related to the use of IDSs within organizations. For example, Goodall et al. (2004a;b) show how the highly collaborative nature of the detection of security incidents and the need of specific knowledge of the organization's unique network environment can represent an obstacle to the use of IDSs. Other studies have shown technical challenges related to the high volumes of information generated by IDSs, which can be solved via the use of better visual interfaces (Komlod et al. 2005; Malécot et al. 2006).

Our study shows not only technical, but organizational challenges to the adoption of IDSs; we present some of the aspects that organizations should cover when deciding the adoption of an IDS, from the planning to the implementation and operation stages.

1.4 Contributions

We provide a summary of our findings and corresponding contributions related to each theme below.

- Chapter 2 presents the qualitative analysis based on the theme of human, organizational and technological challenges faced by security practitioners within organizations. From our analysis of 27 semi-structured interviews, we identify human, organizational, and technical challenges that security practitioners face when implementing security controls in their organizations, e.g., lack of security training of other stakeholders (human), distribution of IT responsibilities (organizational), and mobile access to applications (technical). We also show how these challenges interplay and propose research opportunities for the improvement of IT security technologies from a holistic point of view. For example, security technologies need to take into account that security practitioners have to effectively communicate security issues to other stakeholders who have different perceptions of risks and work in a distributed environment (e.g., multiple stakholders involved in the administration of IT systems).

- Chapter 3 presents the qualitative analysis on the activities and collaborative interactions performed by security practitioners. We identify nine different activities that require interactions between security practitioners and other stakeholders. Furthermore, we provide detailed descriptions of two activities that may serve as useful references for usability scenarios of security tools. We also propose a model of the factors contributing to the complexity of the interactions between security practitioners and other stakeholders. The discussion is centered on how this complexity is a potential source of security issues that increase the risk level within organizations. Our qualitative analysis also reveals that the tools our participants use to perform their security tasks provide insufficient support for the complex, collaborative interactions they have to perform.

  We offer several recommendations for addressing this complexity and improving IT security tools. For example, our findings show that security practitioners sometimes have to combine several tools to perform their security tasks and communicate with other stakeholders, which required copy-pasting between tools, making exchanges of information during interactions error prone. In this vein, an opportunity for improvement is better integration between communication and IT security tools. This improvement might be accomplished through IT security tools that allow on-line collaboration between security practitioners and other stakeholders during the detection and analysis of malicious network traffic.

- Chapter 4 shows the theme of diagnostic work during the response to security incidents. Based on empirical data from 13 interviews with security practitioners who responded to security incidents and participatory observation in one academic organization, we identify the tasks, skills, strategies, and tools that security practitioners use to diagnose security incidents. Our analysis shows that the diagnosis of security incidents is a highly collaborative activity, which may involve practitioners developing their own tools to perform diagnostic tasks. Furthermore, our findings suggest that diagnosis during security incident response is complicated by practitioners' need to rely on tacit knowledge, as well as usability issues with security tools.

  We offer recommendations to improve technology that supports the diagnosis of security incidents, including criteria to evaluate security tools in the context of diagnostic work. Some examples of tool improvements we provide relate to: the tradeoff between task complexity and tool reliability, the need for tools to support tailorability and correlation of high volumes of data, as well as the need for multi-faceted simulation support.

- Chapter 5 presents the qualitative analysis on the adoption of an IDS in one organization. Our analysis reveals that SPs have both positive and negative perceptions related to the utility of an IDS. The analysis also revealed several issues encountered during the initial stages of IDS deployment. In particular, practitioners found it difficult to make decisions about where to place the IDS and how to best configure it for use within a distributed environment with multiple stakeholders. We provide recommendations for mitigating these challenges through better tool support.

1.5 Structure

Each chapter of this thesis has a similar structure: (1) Introduction that summarizes and motivates the investigation of each theme; (2) Related work showing relevant studies in the area under analysis; (3) Methodology explaining the recruitment of participants and the use of qualitative description for the analysis of the data; (4) Results, describing the main findings from the stories told by participants; (5) Discussion, elaborating on interpretations of the data and, where possible, recommendations for improving security tools or processes; and (6) Conclusions, summarizing the findings of each chapter and proposing ideas for future work. Chapter 6 summarizes the contributions of this thesis, the limitations of our research, and directions for future research.

Bibliography

J. Audestad. Four reasons why 100% security cannot be achieved. *Telektronikk*, 1:38–47, 2005.

H. Beyer and K. Holtzblatt. *Contextual Design, Defining Customer-Centered Systems*. Morgan Kaufmann Publishers, San Francisco, CA, 1998.

K. Beznosov and O. Beznosova. On the imbalance of the security problem space and its expected consequences. *Information Management & Computer Security*, 15(5):420–431(12), 2007.

F. J. Björck. *Discovering Information Security Management*. Doctoral thesis, Stockholm University, Royal Institute of Technology, 2005.

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.

D. Botta, R. Werlinger, A. Gagné, K. Beznosov, L. Iverson, S. Fels, and B. Fisher. Studying IT security professionals: Research design and lessons learned. position paper for the CHI Workshop on Security User Studies: Methodologies and Best Practices, April 28 2007a.

D. Botta, R. Werlinger, A. Gagné, K. Beznosov, L. Iverson, S. Fels, and B. Fisher. Towards understanding IT security professionals and their tools. In *SOUPS*, pages 100–111, Pittsburgh, Pennsylvania, July 18-20 2007b.

J. M. Carroll, M. B. Rosson, G. Convertino, and C. H. Ganoe. Awareness and teamwork in computer-supported collaborations. *Interact. Comput.*, 18(1):21–46, 2006. ISSN 0953-5438. doi: http://dx.doi.org/10.1016/j.intcom.2005.05.005.

E. Casey. Error uncertainty and loss in digital evidence. *International Journal of Digital Evidence*, 1(2), 2002.

S. E. Chang and C. B. Ho. Organizational factors to the effectiveness of implementing information security management. *Industrial Management & Data Systems*, 106:345–361, 2006.

K. Charmaz. *Constructing Grounded Theory*. SAGE publications, 2006.

H. H. Clark. *Using Language*. Cambridge University Press, Cambridge, England, 1996.

D. M. Fetterman. *Ethnography: Step by Step*. Sage Publications Inc., 1998. ISBN 0761913858.

J. R. Goodall, W. G. Lutters, and A. Komlodi. I know my network: Collaboration and expertise in intrusion detection. In *CSCW*, volume 6390, November 2004a.

J. R. Goodall, W. G. Lutters, and A. Komlodi. The work of intrusion detection: Rethinking the role of security analysts. In *Proceedings of the Americas Conference on Information Systems (AMCIS)*, pages 1421–1427, 2004b.

K. Hawkey, D. Botta, R. Werlinger, K. Muldner, A. Gagne, and K. Beznosov. Human, organizational, and technological factors of it security. In *CHI'08 extended abstract on Human factors in computing systems*, pages 3639–3644, 2008.

E. Kandogan and E. M. Haber. Security administration tools and practices. In L. F. Cranor and S. Garfinkel, editors, *Security and Usability: Designing Secure Systems that People Can Use*, chapter 18, pages 357–378. O'Reilly Media, Inc., 2005.

A. Kankanhalli, H.-H. Teo, B. C. Tan, and K.-K. Wei. An integrative study of information systems security effectiveness. *International Journal of Information Management*, 23, 2003.

K. J. Knapp, T. E. Marshall, R. K. Rainer, and F. N. Ford. Managerial dimensions in information security: A theoretical model of organizational effectiveness. https://www.isc2.org/download/auburn_study2005.pdf, 2005.

A. Komlod, P. Rheingans, U. Ayachit, J. Goodall, and A. Joshi. A user-centered look at glyph-based security visualization. In *Proceedings of the IEEE Workshops on Visualization for Computer Security (VizSEC)*, pages 21–28, 2005.

I. V. Koskosas and R. J. Paul. The interrelationship and effect of culture and risk communication in setting internet banking security goals. In *ICEC '04*, pages 341–350. ACM Press, 2004. ISBN 1-58113-930-6. doi: http://doi.acm.org/10.1145/1052220.1052264.

A. G. Kotulic and J. G. Clark. Why there aren't more information security research studies. *Information & Management*, 41(5):597–607, 2004.

S. Kraemer and P. Carayon. Human errors and violations in computer and information security: The viewpoint of network administrators and security specialists. *Applied Ergonomics*, 38:143–154, 2007.

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.

E. L. Malécot, M. Kohara, Y. Hori, and K. Sakurai. Interactively combining 2d and 3d visualization for network traffic monitoring. In *Proceedings of the 3rd international workshop on Visualization for computer security (VizSEC)*, pages 123–127, 2006.

S. Mitropoulos, D. Patsos, and C. Douligeris. On incident handling and response: A state of the art approach. *Computers and Security*, 25(5):351–370, 2006.

S. Mohammed and B. Dumville. Team mental models in a team knowledge framework: Expanding theory and measurement across disciplinary boundaries. *Journal of Organizational Behavior*, 22 (2):89–106, March 2001. ISSN 0894-3796.

J. Riden. Responding to security incidents on a large academic network, 2006.

M. Sandelowski. Whatever happened to qualitative description? *Research in Nursing & Health*, 23(4):334–340, 2000.

E. E. Schultz. Computer forensics challenges in responding to incidents in real life setting. *Computer Fraud & Security*, 12:12–16, 2007.

P. Stephenson. The application of formal methods to root cause analysis of digital incidents. *International Journal of Digital Evidence*, 3(1), 2004.

A. Tsohou, M. Karyda, and S. Kokolakis. Formulating information systems risk management strategies through cultural theory. *Information Management & Computer Security*, 14(3):198–217, 2006.

Chapter 2

Human, Organizational and Technological Challenges of Implementing IT Security in Organizations[2]

2.1 Introduction

Recent research has recognized that technological factors are not the only key to the effectiveness of information security controls; there is also a need to understand the impact of human and organizational factors (Beznosov and Beznosova 2007; Botta et al. 2007; Rayford B. Vaughn Jr. and Fox 2001). A better understanding of how different *human*, *organizational*, and *technological* elements interplay could explain how different factors lead to sources of security breaches and vulnerabilities within organizations (Kraemer and Carayon 2007).

This paper reports on the challenges that security practitioners face within their organizations. We used qualitative methods to understand factors that affect the adoption of best security practices within organizations. Our data consisted of 34 questionnaires and 27 interviews with security practitioners from different organizations (18 from academia and 9 from private organizations). Our results not only validate and extend other studies that address challenges that security practitioners face, but also provide an integrated framework that classifies these challenges. This framework can help organizations identify their limitations with respect to implementing security

[2] A version of this chapter has been accepted for publication. R. Werlinger, K. Hawkey, and K. Beznosov. Human, Organizational and Technological Challenges of Implementing IT Security in Organizations. *In HAISA '08: Human Aspects of Information Security and Assurance (13 pages, to appear)*, July 2008.

standards as well as determine if they are spending their security resources effectively. It also provides a way to understand how different factors interplay, for example, how the culture of the organization's people and decentralization of IT security trigger security issues that make security management more difficult. We also elaborate on several opportunities for researchers and developers to improve technology and processes that are used to support the adoption of security policies or standards within organizations. To illustrate, we found that security processes should consider that security practitioners have to effectively communicate security issues to other stakeholders who have different perceptions of risks and do not have security as a first priority within the organization.

We first present related work (section 2.2) on IT security challenges. We then describe our methodology (section 2.3), including our research questions and participant profiles. We present results (section 2.4) as an integrated framework of human, organizational, and technological challenges. We perform a cross-analysis of the findings and discuss the interplay between the challenges (section 2.5). We end this section by grounding our findings in prior research and discussing opportunities for future research before providing final conclusions (section 2.6).

2.2 Background

Our results build upon prior work that addresses a subset of the human, organizational, and technological elements that challenge the adoption of security within organizations. We define human aspects as those related to cognition at the individual level, as well as culture and interaction with other people. Organizational aspects are those related to the structure of the organization, including size and managerial decisions around IT security. Technological aspects involve technical solutions such as applications and protocols. These definitions of human, organizational and technical aspects were adapted from Beznosov and Beznosova (2007)[3]

2.2.1 Human factors

From the human point of view, adoption of security practices poses several challenges for security practitioners. For example, effective interactions and communications are required to reach

[3]Beznosov and Beznosova (2007) define technical, human and social aspects.

a mutual understanding about security risks among different stakeholders. Koskosas and Paul (2004) study how security risks are communicated in financial organizations. They conclude that risk communication "plays a significant role at the macro-goal level of security management," and affects the setting of banking security goals. Tsohou et al. (2006) recognize that risk management is basically a human activity and propose the use of cultural theory to classify the different perceptions of security risks that stakeholders might have. Depending on the classification, security professionals should adopt different strategies to communicate and reach common risk perceptions with other stakeholders. Garigue and Stefaniu (2003) elaborate on the importance of reporting in order to communicate security concerns within organizations. They conclude that reporting on security issues is both a science and an art, with much human judgement necessary to interpret the reports from security tools.

Human errors represent another threat for best security practices[4]. Kraemer and Carayon (2007) identify and characterize elements related to human errors in the field of information security. They populated a conceptual framework with qualitative data from 16 interviews with network administrators and security specialists. Their analysis shows that organizational factors such as communication, security culture, and policy are frequent causes of errors in the context of information security and that communication breakdowns cause security vulnerabilities.

2.2.2 Organizational factors

Kankanhalli et al. (2003) propose a model that relates organizational factors such as organization size, top management support, and type of industry with the effectiveness of information security controls within organizations. From 63 surveys, they conclude that management support is positively related to the implementation of preventive security efforts. They found that financial organizations invest more resources in controls to deter bad security practices than other organizations and that larger organizations invest more in deterrent measures than smaller ones. Similarly, Chang and Ho (2006) study the factors that influenced the adoption of the IT security standard BS7799 in various organizations in Taiwan. From 59 surveys, they also conclude that factors such as top management support, size, and organization type are related to the implementation of se-

[4]We use the definition used by Kraemer and Carayon (2007) for human error: human but non-deliberate accidental cause of poor computer and information security (CIS) (e.g., an accidental programming error that causes a computer to crash under certain circumstances, or the unintended cutting of a communications cable during excavation).

curity controls. Additionally, they find that the uncertainty of environmental elements, including high-speed change of technology, competitors' behaviors, customers' security requirements, and changes in legislation affect security management.

Knapp et al. (2006) surveyed 936 security professionals about the importance of top management support in predicting policy enforcement and security culture within organizations. They conclude that this factor is critical for implementing security controls within organizations. Similarly, Straub and Welke (1998) study the impact of management training on the implementation of security plans in two tech services organizations. They conclude that managers are not aware of the full spectrum of actions that can be taken to reduce risks, but they will employ security planning techniques if they receive training about these techniques.

2.2.3 Technological factors

Technological complexity is another challenge for security practitioners. Audestad (2005) suggests that one of the reasons for not reaching 100% security is because of the complexity of technology. This complexity makes it extremely difficult for the decision makers to manage the big picture and design security policies that cover all the possible configurations of the systems. Welch and Lathrop (2003) studies the complexity of wireless networks and the challenges they pose to security practitioners. Jiwnani and Zelkowitz (2002) describe security testing of systems as a lengthy, complex, and costly process. They propose a taxonomy to classify vulnerabilities and assist security practitioners in the prioritization of resources to patch them.

2.3 Methodology

A better understanding of real world conditions and constraints during the adoption of security practices would help developers and designers make secure systems more usable (Flechais and Sasse in press). None of the studies described in the related work provide a comprehensive, integrated overview of the challenges faced by security practitioners. The goal of our study is to help fill that gap. Our analysis of security challenges is part of an ongoing project whose long term goal is to construct a set of guidelines for evaluating and developing tools used for managing IT security (Hawkey et al. 2008). For the analysis reported here, our primary research questions were: (1) What are the main challenges that security practitioners face in their organizations?; (2)

Table 2.1: Profile of our participants and their organizations

Type of organization	**Interviews**	**Job description**
Financial services 1	I4	IT Security Specialist
Financial services 2	I25	IT Security Specialist
Insurance services	I5	IT Security Specialist
Security consulting services 1	I23	IT Security Specialist
Security consulting services 2	I27	IT Security Specialist
Non-profit medical services	I19	IT Systems Specialist
Manufacturing	I16	IT Manager
	I21	IT Security Specialist
Research institution	I12	IT Systems Specialists
Academic 1	I1	IT Manager
	I3	IT Security Specialist
	I14	IT Systems Specialists
Academic 2	I2, I15, I17, I18	IT Managers
	I9, I11, I24	IT Security Specialists
	I7, I10, I20	IT Systems Specialists
Academic 3	I22	IT Systems Specialists

How do these challenges interplay?; and (3) What are the implications of the challenges on future research?

To answer these questions, we collected empirical data from interviews with security practitioners working in real environments. The strategies we used to address the difficulties of collecting data on how organizations manage IT security are described elsewhere (Botta et al. 2007). For this study, we obtained 34 completed questionnaires that led to 27 interviews with IT professionals with security responsibilities. The questionnaire provided demographic information, while the semi-structured interviews covered various aspects of IT security. Participants answered questions about their tasks, the tools they use, and the challenges of implementing security controls. To reduce interviewer bias, two researchers conducted each interview. This approach ensured coverage of interview questions and allowed the interviewers to probe for details from different perspectives. It is important to note that, due to the nature of semi-structured interviews, not all topics were discussed at the same level of detail with all participants; 23 of our participants explicitly discussed challenges (see Table 2.1 for their profiles).

The interviews were analyzed using qualitative description (Sandelowski 2000) with constant comparison and inductive analysis of the data. We first identified instances in the interviews when participants described the challenges they faced when implementing security controls within

their organizations. These situations were coded iteratively, starting with open coding and continuing with axial coding. Results were then organized by the types of challenges (e.g., lack of resources to implement security controls). Posterior analysis was based on further elaboration of "memos" (Charmaz 2006) written during the coding process. Following a theoretical sampling approach, interview questions were adjusted three times (before interviews 15, 22, and 27), in order to validate emerging theories. For the overall project, four researchers performed the analysis, each focusing their analysis on different themes. The challenges theme had a considerable degree of overlap with other themes (e.g., sources of errors for security practitioners); this made triangulation of analysis possible at the researcher level.

2.4 Building an Integrated Framework of Challenges

Our participants described a variety of factors that made it difficult for them to implement security controls in their organizations. We classified these challenges in human, organizational, and technological categories using the definitions from Beznosov and Beznosova (2007) as explained in section 2.2 (examples of the codes used to analyze the data can be found in the appendix A, Section A.1). Given that some challenges are multi-dimensional and can be classified in different ways depending on the particular interpretation of the researcher, our emphasis is on understanding the different types of challenges that affect the work of security practitioners. Table 2.2 provides a summary of the challenges. We next describe in more detail the human, organizational, and technological challenges identified by our participants.

2.4.1 Human factors

We classified three challenges as human factors: (1) *culture*; (2) *lack of security training*; and (3) *communication of security issues*. These were particularly challenging for participants who had to actively interact with other people across the organization to implement security controls. *Lack of a security culture* within organizations made it difficult to change practices, such as several employees using the same account to access one system (I16). In other cases, employees considered their privileges to access data as a status symbol and resisted the loss of privileges as a result of organizational changes (I5). *Lack of security training* was another issue. It is difficult to

Table 2.2: Challenges participants described for implementing security controls

Type	**Challenge**	**Participants**	
		Academia	Private
Human	Lack of training or experience	I14, I18	I19, I27
	Culture within the organization	I22	I5, I16, I19
	Communicate security issues	I7, I9, I12	I25
Organizational	Risk estimation	I20	I4, I25
	Open environments and academic freedom	I1, I3, I9, I11, I15, I20	NA
	Lack of budget	I2, I3, I18	I16
	Security as low priority	I24	I18, I23, I25, I27
	Tight schedules	I7	I25
	Business relationships with other organizations	I17	I4, I5, I25
	Distribution of IT responsibilities	I2, I11, I17	I16, I21
	Access control to sensitive data	I9, I17, I20	I4, I5, I25
Technological	Complexity of systems	I11	I23
	Vulnerabilities (systems/applications)	I11, I20, I22	I25
	Mobility and distributed access	I14	None

implement security controls when people do not have enough orientation or education about best IT security practices (I19). Both lack of security culture and training influenced the perception of risks that stakeholders have within the organization. When there was not a common view of risks between stakeholders, *communication of security issues* was particularly difficult. For example, two participants (I5 and I14) describe how they tried to avoid communication breakdowns with other stakeholders (e.g., business people) who did not share the same perception of security risks. In these circumstances, the participants assumed the role of "risk evaluators" to explain the risks associated with different business decisions.

2.4.2 Organizational factors

Our participants discussed several challenges linked to the characteristics of their organizations. These included: (1) *risk estimation*; (2) *open environments and academic freedom*; (3) *lack of budget*; (4) *security as a secondary priority*; (5) *tight schedules*; (6) *business relationships with other organizations*; (7) *distribution of IT responsibilities*; and (8) *access control to sensitive data.*

Risk estimation, the consequences if the risks were not mitigated, and the success of mitigation

controls, were all elements our participants found difficult to assess (I20, I25). Stakeholders need security training and experience before they can estimate risks (I14), which made it necessary for security practitioners to try to effectively communicate potential losses for the organization (I25).

An *open academic environment* proved challenging for some participants (I1, I3, I9, I11, I15, I20) who had to adapt their solutions to expectations of academic freedom by faculty members and students: "...that's an interesting trade off all the time. You're constantly trading access versus risk" (I1). This made it difficult to enforce security and implement technical solutions to mitigate risks that could compromise security. For example, one participant (I3) mentioned how difficult it was to monitor and control attacks that could be initiated using the organization's IT systems.

Budget restrictions for security programs was also a challenge discussed by participants. The implementation of security technologies can be costly (I19). It is also difficult to obtain resources for security controls when people do not understand the importance of security (I18).

Security may be a relatively low priority for some businesses: "I come from an outsourcing background where security had very tight processes... What I've learned through this company is we can't always go there... This is not an IT company, it's a manufacturing company" (I16). Participants from the private sector discussed the trade-off between security and the business processes. This trade-off was reflected in specific situations where our participants had to either relax security policies or justify the application of security controls. One participant described how the application of security patches that decreased the performance of certain applications triggered a conflict between IT security people and internal users (I5). A lack of priority for security may also make organizations overlook the need for enforcing security controls when they hire services externally. If security is not part of the big picture, external workers might not be made aware or trained about the security controls in the organization (I17).

Tight schedules as a result of business priorities are a related challenge and may result in human errors that might make the organization more vulnerable (I7). Tight schedules may also result in security controls not being implemented in the systems unless the implementation of security controls is integrated with the development process (I25).

Business relationships with other organizations posed a challenge when the organizations involved did not have similar standards in their security levels. This may also occur when organi-

zations merge or acquire other organizations, resulting in internal silos with different needs and practices in terms of IT security. This problem can be more difficult to solve when IT security is not a main priority of the business (I16). For example, one participant (I4) explained how they had to sacrifice the application of security policies when her organization started to interact with other organizations with different security requirements.

Distribution of IT responsibilities across organizational units was an issue for our participants, particularly for those from academic settings. In the academic organizations we studied, various administrative departments shared the IT networks and systems; within each academic department, at least one employee was responsible for the local IT infrastructure. Some participants believed this distribution diminished the capability of the organization to apply IT security controls: "the decentralized nature does not help." (I2). This challenge of decentralization is similar to interactions with other organizations, as in both cases the decisions on IT security involves distributed entities.

Controlling access to data was an important challenge for our participants (I4, I5, I9, I17, I20, I25). They were concerned about sensitive data distributed in different areas of the organization; this data needed to be accessed by stakeholders from different networks and systems. The problem arose as they did not have a system to control access to data in a centralized fashion.

2.4.3 Technological factors

Our participants were also concerned with technological factors as they tried to implement security policies. The factors we found in our analysis were: (1) *complexity of systems*; (2) *mobile and distributed access*; and (3) *vulnerabilities in systems and applications*. We focus our findings on the first two factors, as they were more related with other organizational factors.

The complexity of systems and the need for having open and secure networks had an influence on the interactions between participants from academia and security vendors. One participant (I15) mentioned how difficult it was for vendors to understand the architecture of the network and offer products that suit his organization's needs. Another participant from the private sector (I23) also mentioned the complexity of the networks and systems as a challenge to implement security controls in organizations. For example, a typical network could have firewalls, DMZs, proxies, switches behind the firewall, routers in front of the firewalls, mail servers and not enough people to

look after the overall security of these interconnected devices. Other organizational factors such as decentralization of IT management, interaction with other organizations, and distributed sensitive data increased the complexity of technical solutions. These technical solutions needed to restrict access from different users with different needs and security requirements.

Mobility and distribution of user access made it difficult to control access to internal resources. Mobility of laptops that can be taken to different places and accessed by people who do not have enough technical expertise was a big problem for one participant (I14). He mentioned how Mondays were particularly bad days as users often came back to work with their laptops infected with malicious software from home usage.

2.5 Discussion

We discuss our results from three different perspectives. First, we perform a cross analysis of the challenges described by participants, considering their organizations and positions. Second, we describe how different challenges interplay. Third, we ground our results in prior research and discuss research opportunities to improve security tools and processes. Where possible, we propose characteristics that these tools and processes should have to support security practitioners in real contexts.

2.5.1 Cross analysis

Our analysis showed no contradictions between the challenges described by managers and other participants; managers discussed factors that either confirmed or complemented the challenges mentioned by other security practitioners. Patterns did emerge from the cross-analysis of participants from different sectors. First, academic institutions face challenges related to academic freedom and the need for an open environment. Second, challenges related to the distribution of IT management were similar for academic and private organizations; in academic organizations there were several independent departments with their own IT infrastructure, whereas in private organizations there was a need for interacting with IT departments from other organizations or from different branches within the same organization. We also found that the need for controlling access to sensitive data was a common concern.

These findings validate and extend prior research as our sample of participants contrasts in quantity and type with those ones used in similar studies (e.g., Koskosas and Paul (2004) performed 15 interviews in three organizations; Kraemer and Carayon (2007) performed 16 interviews in two academic laboratories). However, more data are necessary in order to empirically test these emerging theories. Continued research in this area is important as these factors might be used to predict how effectively security policies are adopted within a given organization.

2.5.2 A holistic view of challenges and their interrelationships

Kankanhalli et al. (2003), Knapp et al. (2006) and Chang and Ho (2006) relate organizational variables such as size, type of business, environmental elements (e.g., customers security requirements), and top management support with security effectiveness, security culture, and enforcement of security policies within organizations. Our framework identifies other organizational variables that make it more complex to perform IT security within organizations. Furthermore, we found human, organizational and technological factors that interplay with each other and directly impact the work of security practitioners (Figure 2.1 illustrates this interplay). This interplay has been partially described in Section 2.4, when the different challenges were described. Other interactions among the different challenges were extracted from the relationships among our axial codes (examples of this relationships are in the appendix A, Section A.2). For example, the challenge communication of security issues was affected negatively by both the human challenge of different perception of risks and the organizational challenge distribution of IT management. The challenge lack of security training negatively impacted the priority given to security. Organizational factors such as an open academic environment, distribution of IT management, interaction with other organizations, and controlled access to data distributed in different departments increased technical complexity.

2.5.3 Opportunities for future research

The challenges we described not only illustrate the complexity of the environment where security practitioners work, but also show the limitations that organizations face when implementing security policies. These challenges also represent opportunities for future research. For example, our analysis showed that effective communication was a challenge for our participants, who needed

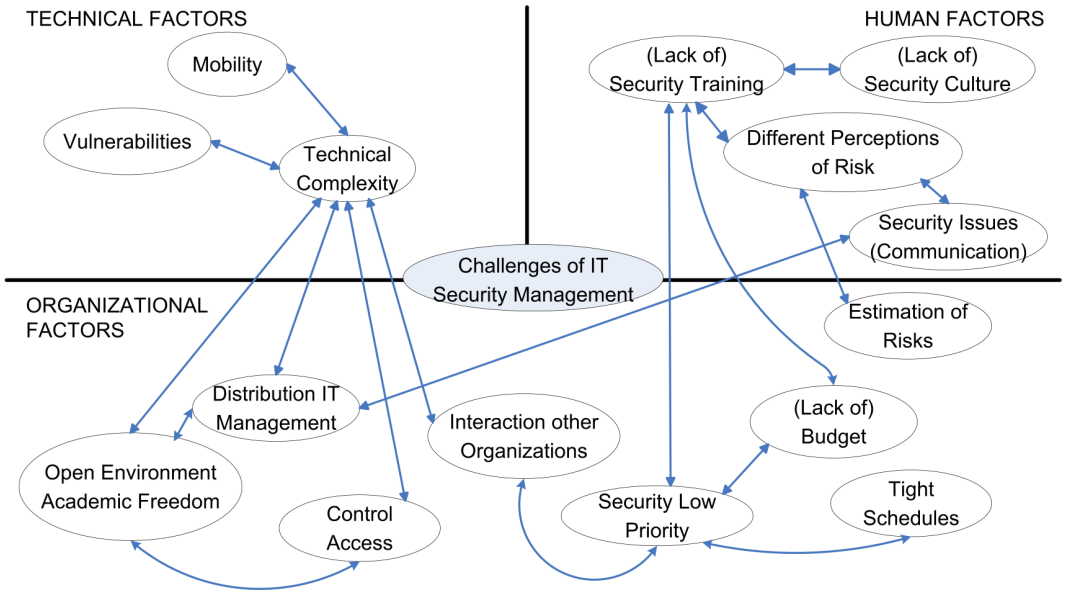

Figure 2.1: A holistic view of challenges and their interrelationships. These interrelationships were extracted from the stories of our participants (see Appendix, Section A.1).

to explain to other stakeholders security risks and the need for security controls. Pattinson and Anderson (2007) highlight the importance of risk perceptions for end-users and how important it is to communicate these risks to them. Koskosas and Paul (2004) study how risks are communicated in financial organizations. They concluded that risk communication "plays a significant role at the macro-goal level of security management." Our study extends this result by showing that the implementation of security processes should consider the organizational culture and the view that different stakeholders (not only end-users) have about security risks. A good starting point for addressing communication issues may be to apply Tsohou et al.'s (2006) proposal of using culture theory to communicate security risks, but focusing only on a subgroup of stakeholders (e.g., managers).

We found that distribution of IT management and the lack of security training of other stakeholders are also factors that negatively impact the effectiveness of communications performed by security practitioners. To address these challenges, security tools might consider the use of flexible reporting (Botta et al. 2007) to communicate security issues (i.e., reports customizable depending on the knowledge or level of the recipient). Our analysis, that included 13 more interviews than the one performed by Botta et al. (2007), also showed that a better integration between security

and communication tools is necessary (e.g., integration of firewall administration tools with e-mail or chat).

Tight schedules for delivering services that include security requirements was another challenge for some participants. Kraemer and Carayon (2007) relate the lack of time, resources, and inconsistent communication among the staff with errors that are introduced into the systems. This implies a direct relationship between tight schedules and the security level of the organization. We propose that security processes and technologies should provide more support on how security practitioners should prioritize their tasks. For example, in the context of security incident reporting, Sveen et al. (2007) propose that organizations should save resources and time by reporting only high priority security incidents. Another potential avenue for improvement is the development of tools that not only show security vulnerabilities, but also give better support to determine how security practitioners should prioritize their tasks considering the level of security risks of the different systems.

Distribution in the context of controlled access to data had two facets: first, to control access from users that are distributed and use different access technologies; and second, to control access to data distributed across the organization and managed by different stakeholders. It seems difficult for those organizations that are highly distributed in nature (e.g., academic ones) to implement centralized, strong security controls able to restrict every access and action. We propose that security processes and technologies must be developed assuming distributed environments. They should be flexible enough to both provide controlled access to highly distributed data and improve communication channels among the different stakeholders that access those data.

Training and education may improve security awareness in organizations (Sveen et al. 2007; Kankanhalli et al. 2003). We argue that the process of designing security policies can be used to train and educate other stakeholders within organizations. When designing security policies, security practitioners have to share their experiences about security incidents, vulnerabilities and culture with other stakeholders. For example, Gonzalez et al. (2005) developed mental models that integrated the fragmented knowledge from different experts. These models identified risks in the transition to integrated operations in the Norwegian oil and gas industry. In the same vein, security policies should not be seen only as artefacts to enforce best IT practices (Thomson and von Solms 2005), but also as a way to share the tacit knowledge that security practitioners have

by explaining the "why" of the controls to other stakeholders. At this point, techniques such as the use of scenarios and anecdotes (Flechais and Sasse in press) look appropriate to spread the tacit knowledge used to build the policies.

We found that, within organizational factors, security as a low priority and lack of resources to implement security controls are related to what Kankanhalli et al. (2003) and Chang and Ho (2006) call organization security effectiveness. They find that the greater the top management support, the more effective security is in organizations, as organizations spend more resources in preventive measures to avoid security incidents. Kankanhalli et al. (2003) propose that penetration testing, security vulnerability, and risk analysis reports can be used to convince top management about the importance of security. They also propose making explicit the tangible business benefits of implementing security controls (e.g., raising customer confidence). However, this is not always possible when the organization does not have security experts with the knowledge to convince other stakeholders. Karyda et al. (2006) propose outsourcing IT security services as a solution for those organizations that do not have resources or the required knowledge to implement security controls or develop security projects. However, outsourcing security seems infeasible when organizations do not perceive security as a priority from the beginning. We argue that more research is needed to both determine the rationale behind the decisions that organizations make in the context of IT security, and the trade-offs between the priority given to resources devoted to IT security and the core business of the organization.

2.6 Conclusion

We used empirical data and prior work to provide an integrated framework of the different *human*, *organizational*, and *technological* challenges that security experts have to face within their organizations. As far as we know, this is the first empirical study that provides a comprehensive list of these challenges in the context of information security. This framework is intended to provide guidance for those organizations and security practitioners that need to identify their limitations to implementing security policies, and determine what is relevant in their decisions in the context of IT security. We discussed how the different challenges interplay and suggested various research opportunities to improve security processes and technologies, considering human and organizational factors in the development of security processes and technologies.

More research is needed to understand how security challenges interplay, as this interaction affects the improvements that organizations can make in terms of their security levels.

Bibliography

J. Audestad. Four reasons why 100% security cannot be achieved. *Telektronikk*, 1:38–47, 2005.

K. Beznosov and O. Beznosova. On the imbalance of the security problem space and its expected consequences. *Information Management & Computer Security*, 15(5):420–431(12), 2007.

D. Botta, R. Werlinger, A. Gagné, K. Beznosov, L. Iverson, S. Fels, and B. Fisher. Towards understanding IT security professionals and their tools. In *SOUPS*, pages 100–111, Pittsburgh, Pennsylvania, July 18-20 2007.

S. E. Chang and C. B. Ho. Organizational factors to the effectiveness of implementing information security management. *Industrial Management & Data Systems*, 106:345–361, 2006.

K. Charmaz. *Constructing Grounded Theory*. SAGE publications, 2006.

I. Flechais and M. A. Sasse. Stakeholder involvement, motivation, responsibility, communication: How to design usable security in e-science. *Int. Journal of Human-Computer Studies*, in press.

R. Garigue and M. Stefaniu. Information security governance reporting. *EDPACS*, 31(6):11–17, 2003.

J. J. Gonzalez, Y. Qian, F. O. Sveen, and E. Rich. Helping prevent information security risks in the transition to integrated operations. *Telektronikk*, 1:29–37, 2005.

K. Hawkey, D. Botta, R. Werlinger, K. Muldner, A. Gagne, and K. Beznosov. Human, organizational, and technological factors of it security. In *CHI'08 extended abstract on Human factors in computing systems*, pages 3639–3644, 2008.

K. Jiwnani and M. Zelkowitz. Maintaining software with a security perspective. *Software Maintenance, 2002. Proceedings. International Conference on*, pages 194–203, 2002.

A. Kankanhalli, H.-H. Teo, B. C. Tan, and K.-K. Wei. An integrative study of information systems security effectiveness. *International Journal of Information Management*, 23, 2003.

M. Karyda, E. Mitrou, and G. Quirchmayr. A framework for outsourcing is/it security services. *Information Management & Computer Security*, 14:403–416, 2006.

K. J. Knapp, T. E. Marshall, R. K. Rainer, and F. N. Ford. Information security: management's effect on culture and policy. *Information Management & Computer Security*, 14(1):24–36, 2006.

I. V. Koskosas and R. J. Paul. The interrelationship and effect of culture and risk communication in setting internet banking security goals. In *ICEC '04*, pages 341–350. ACM Press, 2004. ISBN 1-58113-930-6. doi: http://doi.acm.org/10.1145/1052220.1052264.

S. Kraemer and P. Carayon. Human errors and violations in computer and information security: The viewpoint of network administrators and security specialists. *Applied Ergonomics*, 38:143–154, 2007.

M. R. Pattinson and G. Anderson. How well are information risks being communicated to your computer end-users? *Information Management & Computer Security*, 15, 2007.

R. H. Rayford B. Vaughn Jr. and K. Fox. An empirical study of industrial security-engineering practices. *The Journal of Systems and Software*, 61:225–232, 2001.

M. Sandelowski. Whatever happened to qualitative description? *Research in Nursing & Health*, 23(4):334–340, 2000.

D. W. Straub and R. J. Welke. Coping with systems risk: security planning models for management decision making. *MIS Q.*, 22(4):441–469, 1998. ISSN 0276-7783. doi: http://dx.doi.org/10.2307/249551.

F. O. Sveen, J. Sarriegi, E. Rich, and J. Gonzalez. Toward viable information security reporting systems. In *HAISA 2007*, pages 114–127, July 2007.

K. Thomson and R. von Solms. Information security obedience: a definition. *Computers & Security*, 24(1):69–75, 2005.

A. Tsohou, M. Karyda, and S. Kokolakis. Formulating information systems risk management strategies through cultural theory. *Information Management & Computer Security*, 14(3):198–217, 2006.

D. Welch and S. Lathrop. Wireless security threat taxonomy. *Information Assurance Workshop, 2003. IEEE Systems, Man and Cybernetics Society*, pages 76–83, 2003.

Chapter 3

Security Practitioners in Context: Their Activities and Interactions with Other Stakeholders Within Organizations[5]

3.1 Introduction

Information security has become a critical issue for organizations, which need to protect their information assets from unauthorized access and continue business activities after security breaches. Recent studies have shown the need for more empirical evidence on how human and organizational factors impact security effectiveness in organizations (Beznosov and Beznosova 2007; Botta et al. 2007; Kotulic and Clark 2004). Studies also suggest that security practitioners could benefit from better tools to perform their tasks (Botta et al. 2007; Goodall et al. 2004; Kandogan and Haber 2005).

Prior research has found that IT security responsibilities are distributed in nature (Botta et al. 2007; Knapp et al. 2005). Security activities are performed by groups that usually have a "coordinator", not necessarily a manager, who coordinates other IT specialists to perform IT security activities. Security administration has been found to require collaboration among stakeholders at many levels in the organization (Kandogan and Haber 2005). As such, there is a high level of in-

[5]A preliminary version of this chapter has been published. R. Werlinger, K. Hawkey, and K. Beznosov. Security practitioners in context: their activities and interactions. In CHI '08 extended abstracts on Human factors in computing systems, pages 3789—3794, 2008. A full version of this chapter has been submitted to a journal for publication. R. Werlinger, K. Muldner, K. Hawkey, and K. Beznosov (2008) Security Practitioners in Context: Their Activities and Interactions with Other Stakeholders Within Organizations.

terdependency of security tasks, as they depend strongly on the contributions of other individuals and resources (Knapp et al. 2005). However, these previous studies do not provide detail on how security practitioners interact and communicate with other stakeholders within the organization, or how these interactions vary depending on the security activity being performed. What these studies do identify is the need for a better understanding of how the tools that are used by security practitioners (e.g., intrusion detection systems, vulnerability scanners) support collaboration and information-sharing (Botta et al. 2007; Goodall et al. 2004; Kandogan and Haber 2005). The current lack of a rich understanding in these areas makes it difficult for HCI researchers and tool developers to improve communication and IT security tools. Furthermore, such understanding is needed to develop tests that measure the usability of security tools in real, complex scenarios (Redish 2007).

We argue that human, organizational, and technological factors influence the ability of security practitioners to do their job well (Botta et al. 2007; Werlinger et al. 2008b). To understand how these factors play out in IT security, we have been conducting a field study as part of the HOT Admin research project (see Hawkey et al. (2008a) for an overview of the themes under analysis). The field study has provided us with two sources of data: questionnaires and semi-structured in-situ interviews with security practitioners from both the academic and private sectors. The data are supplemented by ongoing participatory observation in one academic organization in Canada.

In this paper, we present an analysis of our empirical data using qualitative description (Sandelowski 2000) focused on pre-designed themes of analysis (e.g., tasks, interactions). A preliminary version of this work appeared in (Werlinger et al. 2008a). The contributions of our study are threefold. First, we analyze the interdependency of IT security tasks by showing the different roles, types of communications, and resources used by IT security professionals in real contexts. Our results include a list of nine different activities that require interactions among security practitioners and other stakeholders. We also describe in detail two of these activities, *responding to incidents* and *developing policies*, which may be used as scenarios to evaluate IT security tools. Second, based on these results, we propose a model that shows the factors that make interactions between security practitioners and other stakeholders complex, and relate this complexity to potential security issues that erode the level of security in organizations. Third, we highlight the implications of our findings for other researchers working on improving practices and tools used by

security professionals. Our findings suggest that the IT security tools used by security practitioners provide insufficient support to address the complexity of their interactions as they collaborate, cooperate, and coordinate with other stakeholders. We offer several recommendations to improve these tools, and give specific examples of how developers could implement our recommendations. For example, security practitioners had to combine several tools to perform their security tasks and communicate with other stakeholders; copy-pasting outputs of tools as inputs for other tools can make interactions error prone. In this vein, an opportunity for improvement is more integration between communication and IT security tools. This improvement might be accomplished through IT security tools that allow on-line collaboration between security practitioners and other stakeholders during the detection and analysis of malicious network traffic.

The remainder of this paper is organized as follows. We first discuss related work, focusing on empirical studies of collaborative work in the context of information security. In Section 3.3, we describe the research methods used to investigate the interactions among security practitioners and other stakeholders, including recruitment of participants, our data collection from multiple sources, and our data analysis. In Section 3.4, we analyze these interactions in context, identifying those security activities that require interactions with other stakeholders, the communication channels used during interactions, and the security tools used within the context of these interactions. In Section 3.5, we provide in-depth descriptions of interactions during two activity scenarios: responding to security incidents and developing security policies. In Section 3.6, we develop a model of the complexity of interactions, which includes factors arising from organizational attributes, multiple stakeholders, and multiple security-related activities. This model includes security issues that may arise as a consequence of such complexity. In Section 3.7, we conclude with a discussion of the implications of our findings for researchers and practitioners, including opportunities for improved tool support.

3.2 Related Work

Prior research has examined computer supported collaborative work (CSCW) (e.g., Carroll et al. (2006); Mohammed and Dumville (2001)); these studies propose general frameworks for understanding team effectiveness, and suggest future directions to improve empirical methods and the

design of systems that support collaborations. Although these frameworks integrate different facets of collaboration (e.g., activity awareness, common ground), the complexity and importance of information security has motivated specific empirical studies on collaborative work in the context of information security. The next section provides relevant empirical research in this area, followed by a summary that contrasts this previous research with our study.

3.2.1 Empirical Research on IS Collaborative Work

Björck (2005) used grounded theory to understand the challenges in establishing a balanced management system for information security. The data for his study came from 29 semi-structured interviews with 8 IT security managers, 13 consultants, and 8 auditors from different Swedish companies. He finds that sound communication capabilities are one of the success factors for the formation and certification of information security management systems.

Kandogan and Haber present two different studies related to IT security administrators. Kandogan and Haber (2005) evaluate security administration tools through 40 days of naturalistic observations of security administrators at a US university. Based on real situations faced by their participants, they give recommendations about future development of IT tools, including improvement of support for collaboration and information-sharing tasks performed by security administrators. In the second study, Haber and Kandogan (2007) analyze ethnographic data from 16 field studies of IT administrators to determine differences between IT system and security administrators. They find that security administrators, unlike other system administrators, have to collaborate intensively to manage the risk and complexity of their tasks.

Botta et al. (2007) used a qualitative approach to identify the goals, responsibilities, tasks, and tools used by security practitioners within organizations. This initial analysis emerged from 14 interviews with security practitioners, with 10 of them from academic institutions. It suggests that information security responsibilities are distributed among many individuals, and that novel tools are needed to support collaboration among these individuals.

Goodall et al. (2004) report on the expertise and collaboration necessary to administer intrusion detection systems (IDSs). The data used for their analysis was derived from 9 interviews of a diverse cross-section of intrusion detection experts. They conclude that security work is collaborative both within organizations and distributed across the Internet, and that IDSs do not properly support distributed collaborative work.

Knapp et al. (2005) investigate how to model the managerial constructs that most influence the effectiveness of IT security. As part of their study, they surveyed 936 security professionals about the interdependency of IT security tasks. They conclude that security tasks have a high level of interdependency, requiring contributions of other individuals and resources.

Flechais and Sasse (in press) studied how security is applied in the development of e-Science projects. In this type of software development project, the goal is to have systems that are secure enough to guarantee to the researchers (a highly distributed community of users) that their information is safe. At the same time, the systems must be usable enough that other researchers will be encouraged to share their information. From their analysis, they propose a model of socio-technical secure system design. Their model recognizes three different factors that affect security design: the responsibility, knowledge, and motivation of different stakeholders. The model also proposes that effective communication between stakeholders is necessary so that relevant security design information is considered.

Kraemer and Carayon (2007) identify and characterize elements related to human errors in the field of information security. They populate a conceptual framework with qualitative data from 16 interviews with network administrators and security specialists. Their analysis suggests that organizational factors such as communication, security culture, and policy are the most frequently cited causes of information security errors, and that communication breakdowns cause security vulnerabilities.

3.2.2 Summary

As discussed above, prior studies have used empirical data to demonstrate that security practitioners work in a distributed, interdependent, and collaborative environment, where communication breakdowns may create security vulnerabilities.

Previous studies also point to the need for a better understanding of how security and communication tools support interactions among security practitioners and other stakeholders. We designed our study to satisfy this need; we adopted a qualitative approach to collect empirical data on how security practitioners communicate and interact when they perform their security tasks. The next section explains our research methods in detail.

3.3 Research Methods

This study of interactions among security practitioners and other stakeholders is part of the HOT Admin research project, which has the long-term goal of developing a set of guidelines for evaluating and designing tools used for managing IT security.

Our three primary research questions for the study described here were: (1) When and how do security practitioners interact with other stakeholders?; (2) What tools do they need to interact effectively?; and (3) What factors are responsible for miscommunications? In order to answer these questions, we needed empirical data about security practitioners working in real environments. We used qualitative methods to obtain and analyze these data.

3.3.1 Participant Recruitment

Collecting data on how organizations manage IT security poses several challenges (Botta et al. 2007; Kotulic and Clark 2004). Practitioners do not have time to participate, they are not willing to disclose security information, and their contact information is not publicly available. We used two strategies to address these challenges. First, professional contacts of the research team served as initial contacts, who recommended other security practitioners who might be interested in taking part in the study. Second, a graduated recruitment approach was taken; potential participants were asked only to answer a short questionnaire that had a final question asking whether they are interested to meet for a one-hour interview. For a discussion on the effectiveness of this approach, see Botta et al. (2007). In the next section, we describe the questionnaires, interviews, and participatory observations that comprise our study data.

The total number of participants was 32 in 30 interviews. This difference is due to the fact that in two of the interviews (I6 and I22), two participants answered the questions. Eighteen of our participants worked at academic organizations, while the other fourteen came from ten different organizational sectors. Participants included IT and security managers, and IT and security specialists. Table 3.1 shows the positions held by our participants across these different sectors.

Table 3.1: For each type of organization, we indicate the number of unique organizations and the total number of participants interviewed. These participants held various positions, including Managers (with security tasks), IT Practitioners (with security tasks), Security Managers, and Security Specialists.

Organization Type	**Position Type**				**Total**
	IT Manager	**Security Manager**	**Security Specialist**	**IT specialist**	
Academic (3)	4	1	4	9	18
Financial Services (2)	-	-	2	-	2
Insurance (1)	-	-	2	-	2
Scientific Services (1)	-	-	-	2	2
Manufacturing (1)	1	-	1	-	2
Telecommunications (2)	-	-	2	-	2
Non-Profit Organization (1)	-	-	-	1	1
IT Consulting Firm (3)	-	-	1	2	3
Total	5	1	12	14	32

3.3.2 Data collection from multiple sources

Questionnaires and Semi-Structured Interviews

The questionnaires completed by participants provided demographic information. The semi-structured interviews covered various aspects of IT security. Our participants answered questions about their tasks, the tools they use, and the communications they perform to do their job. To reduce interviewer bias and obtain data from different perspectives during the interviews, each interview was conducted by two researchers. This team approach also ensured coverage of interview questions. It is important to note that, due to the nature of semi-structured interviews, not all topics were discussed at the same level of detail with all participants.

Participatory observation

We have also been using an ethnographic approach (Fetterman 1998) to collect more data about the different roles and the nature of communications performed by security practitioners in real settings. This approach consists of participatory observation at one academic organization in Canada. The observer spent over 75 hours working under the supervision of a senior IT security professional. One of the tasks of the observer has been the development of policies; he has participated in eight meetings with IT specialists to write and update a set of internal policies

with respect to data classification, secure browsing, and remote connections. Another task of the observer has been the deployment of an intrusion detection system (IDS); he has worked with two security specialists on the installation and configuration of an IDS in the internal network. The results of this participatory observation were used to cross-validate and complement findings from the interviews about the interactions performed during the development of security policies, and the features that security tools should provide to support better collaboration among security practitioners and other stakeholders.

3.3.3 Data Analysis

The interviews were analyzed using qualitative description (Sandelowski 2000) with constant comparison and inductive analysis of the data. We first identified instances in the interviews when participants described interaction with other stakeholders in performing a task. These situations were coded iteratively, starting with open coding and continuing with axial and theoretical coding. The results were then organized by the different activities that provided context for the interaction, as well as communication channels, tools, general resources (skills and knowledge) mentioned as being necessary for interaction, and the sources of errors identified by participants during communications. Posterior analysis was based on further elaboration of "memos" (Charmaz 2006) written during the coding process. Following a theoretical sampling approach, interview questions were adjusted three times (before interviews 15, 22, and 27), in order to validate emerging theories. For the overall HOT Admin project, five researchers are performing analysis, each focusing on different themes. The interaction theme presented in this study had a considerable degree of overlap with other themes (e.g., sources of errors in security management), which made triangulation of analysis possible at the researcher level.

3.4 Analyzing Interactions in Context

We identified several stories in our participant interviews of activities in which IT security-related communications occur. After describing these activities and communications, we present the communication channels and security tools used by our participants for interacting with other stakeholders. To further illustrate our findings, our results conclude with descriptions of the

interactions, tools, and miscommunications involved in two of the activities: *security incident response* and *development of policies*.

3.4.1 Activities Requiring Interactions with Other Stakeholders

We identified nine security activities where participants had to interact with other stakeholders. These interactions represented a challenge for our participants; they required different strategies for communicating security issues to stakeholders with varying backgrounds and interests. To perform security tasks, our participants had to coordinate, collaborate, and cooperate with other stakeholders. We used Matessich and Monsey's (1992) definitions to frame our analysis: *collaboration* is a mutually beneficial and well-defined relationship to achieve common goals; *cooperation* is characterized by informal relationships that exist without any commonly defined mission, structure or planning effort; *coordination* is characterized by more formal relationships and understanding of compatible missions. Although these three types of interactions were often combined in our participants' duties, some tasks were characterized by a bigger influence of one or two of them. For example, participants mainly coordinated time and resources with other stakeholders to perform security audits. Table 3.2 on page 60 shows the nine activities described by our participants, as well as a summary of stakeholder interactions for each activity. Next, we give a brief description of each activity.

The objective of *security audits* for our participants was to find vulnerabilities in the IT infrastructure and generate reports with recommendations for other IT specialists. These reviews could be in the context of formal audits performed either by internal departments or by external audit companies, or as part of less formal internal checks within the IT department. When our participants performed the audits, they had to interact with other IT specialists to communicate and explain the vulnerabilities found in the systems. In other cases, they provided support and interacted actively with IT specialists to respond to recommendations provided by the auditor.

To *design services incorporating security requirements*, our participants had to specify security requirements for new IT services or projects. They had to plan the deployment of new services with other specialists, such as remote access, integrated solutions for collaborative environments, and internal customized services. They also had to participate in committees to approve new projects or changes in the infrastructure, checking how security requirements were incorporated in

the changes. Typical issues that our participants needed to address as consultants were: where to place access controls, what antivirus protection to use, and which security vendors to choose. For this last issue, our participants needed to interact with potential vendors involved in the project, in order to request specifications or evaluate security features of the products offered.

Our participants had to *solve end-user IT security issues* when they received notifications about users experiencing security issues in their computers (e.g., malicious software). Depending on the type of request, they had to either get more information from the users (either by phone or e-mail), or visit them *in situ* to check their computers.

To *implement security controls* such as access control policies for the internal resources, interaction was necessary with other departments within the organization. Usually these interactions were motivated by a lack of consolidated databases of employees and active users of the systems. For example, one of our participants had to coordinate with Human Resources to verify the list of active users in their database systems.

Our participants also had to *train and educate other specialists* on security issues in a variety of circumstances, such as training new employees in the organization's privacy procedures.

Mitigation of vulnerabilities started with notifications from IT providers or security entities identifying new vulnerabilities in the systems. These notifications triggered interactions among our participants. In these cases, participants forwarded the information to other specialists, both to notify them and to confirm the vulnerability with them.

Administration of security devices was another activity described by participants. For example, one participant (I24) had to administer the network's firewalls, even though there were IT specialists who were devoted to operating and maintaining the devices in the network. There were two main reasons to have this distribution of responsibilities. First, "network people" did not manage the access control policies configured in the firewall to control traffic transmitted from one part of the network to the other. Second, there was a historical reason: our participant had started the installation of the firewalls in the network, and had the expertise necessary to re-configure and administer them.

The remaining two activities are described briefly here, but will be presented in full later in section 3.5 as illustrative scenarios of interactions, tools, and sources of errors. To *respond to security incidents*, our participants needed to actively interact with other stakeholders. For

example, to verify the reasons for spikes in e-mail or traffic in a highly distributed IT environment, our participants needed to correlate their information with that of other IT specialists to find out the physical location of the affected devices. *Development of policies* generally involved committees comprising different IT specialists, managers and executives from the affected areas.

These nine activities described by our participants show the diversity of IT security-related tasks and the importance of interactions in performing them. The scenarios themselves also speak to the need for intimate knowledge of the organization in order to involve stakeholders from pertinent areas. The next sections elaborate on the main tools (communication channels, security tools) used by security practitioners to interact with other stakeholders.

3.4.2 Communications Channels Used during Interactions

Participants used multiple communication channels to interact, such as e-mail, text and video chat, phone calls, and face-to-face meetings. These channels were used to broadcast information, receive notifications, share documents, gather information, send requirements, and report security issues.

Our participants all relied heavily on e-mail. They reported using e-mail to broadcast information to other IT specialists and to share documentation. E-mail was also reported to be easier to track and read from remote locations, such as home, than other solutions like ticketing systems (I3 and I15). Nevertheless, participants' perceptions about the effectiveness of e-mail varied. For example, one participant (I4) claimed that misunderstandings arise easily through the casual language common in many e-mails and expressed the need for care about how things were written. The same participant (I4) also compared e-mail unfavorably with verbal communication in situations that required clarification. In contrast, three participants (I3, I5 and I30) thought e-mail was useful to formalize and clarify what they had discussed during meetings.

The large quantity of e-mails from systems and people was reported to be an issue. However, one participant (I9) was able to diagnose at a glance by noting the number of new e-mails in certain folders (the more e-mails from specific systems, the more likely a problem existed).

Keeping a record of communications was important for participants. One participant (I21) was careful to keep two CD-ROM copies of all e-mail. For access control administration, an e-mail reply from an authorized person might be taken as proof of authorization for access when only

logged-in users can use the e-mail system. Another participant included copies of the e-mails in project's files (I30).

Besides e-mail, at least four participants used other tools like text or video chat to communicate. Again, perceptions of the usefulness of these tools varied. Two participants (I9 and I11) found text chat a good tool for getting an immediate response and asking about specific information (e.g., a system's command syntax), while two other participants (I8 and I11) found it distracting, with no guarantee of response. Video chat was preferred because it complemented the advantages of text chat with images. However, one participant (I9) commented that some colleagues did not use video chat because they found it unnatural, with shifts between what is seen and what is said, and with each party unable to see the eyes of the other.

Seven participants (I1, I4, I8, I11, I14, I15, I30) stated that they preferred to use verbal communication (e.g., face-to-face or phone) when they had to interact with other stakeholders. Face-to-face communications allowed them to quickly interact and avoid misunderstandings. Two participants (I14 and I30) mentioned the use of whiteboards to support face-to-face communications. One of them (I30) had access to electronic whiteboards, which were very useful to keep a record of what was discussed. When the electronic option was not possible, the participant took pictures of the whiteboard.

Internal web sites were used to keep track of meetings (I2, I30). These sites were also used to show information to end-users about their IT security services. For example, in order to reduce the overhead of questions from end-users, one participant (I10) employed an internal web site to show users how their spam filters were configured.

Communication systems mentioned by our participants also included an incident-tracking system used by the helpdesk of the participants' organizations (I1, I3, I21). This type of system automatically kept a record of incidents and their resolution, generating tickets to be sent to IT specialists when users reported a problem about the IT infrastructure.

3.4.3 Security Tools Used within the Context of Interactions

To generate security reports, our participants mentioned tools like Nessus (I9, I12, I23, I25), a tool used to show the vulnerabilities of the IT infrastructure; and McAffee ePolicy Orchestrator (I3, I4, I14), a tool used to summarize the virus activity of the systems. One participant (I9), who

coordinated the mitigation of vulnerabilities with other IT specialists, explained the flexibility of Nessus' reports in terms of how easy it was to browse through their links and check the vulnerabilities at appropriate levels of detail. This flexibility allowed him to have a general overview of the vulnerabilities, whereas other specialists could have a detailed view of the information to mitigate the vulnerabilities.

Our participants also mentioned other reporting features that security tools should include. For example, security tools should generate reports that can demonstrate to other stakeholders the economic benefits of applying security controls (I3, I24). Reports should specify what is "normal" traffic in the network and what is not, based on correlation features (I3); and reports should help security practitioners to prioritize their activities, showing security risk levels according to systems' vulnerabilities and compliance of the IT infrastructure with patches, antivirus tools, and countermeasures for new vulnerabilities (I4).

Reports and notifications also came from the different systems that our participants monitored. Three participants (I3, I12, and I25) described how they wrote scripts to monitor the systems, correlate data, and send alarms by e-mail to themselves when an anomaly was detected. Other participants (I2, I9, I22) mentioned how they received notifications generated by scripts created by other IT specialists.

Another important requirement mentioned for communicating security information was the use of an encrypted communication channel (e.g., virtual private networks or VPNs). Two participants (I26 and I29) reported the need to transmit sensitive information (e.g, a report about a security incident or a list of passwords) and protect it from attackers who could be sniffing the network. However, both participants mentioned that they were unable to send encrypted information by e-mail. One participant (I29) said that the organization did not provide the tools necessary to encrypt e-mail, and another participant (I26) said that her clients found the process of encrypting and decrypting e-mails too complex.

3.5 Interaction Scenarios

We used communication flow diagrams[6] (Beyer and Holtzblatt 1998) to show the interactions between security practitioners and other stakeholders during two activities performed by our participants: *responding to security incidents* and *developing policies*. These scenarios provide a reference for the environment in which security tools should be tested (Redish 2007). The next sections describe in detail the interplay of interactions, use of resources, and the role of misunderstandings in these two scenarios.

3.5.1 Interactions in Responding to Security Incidents

Responding to security incidents was the activity most commonly mentioned by our interviewed participants. Interactions during security incidents were complex, involving collaboration, coordination, and cooperation. These interactions were also characterized by the use of multiple communication channels for sharing knowledge among different specialists during the investigation.

From the stories told by our participants, we built a communication flow diagram showing the exchange of information among the main stakeholders involved in responding to a security incident (details in Figure 3.1). These stakeholders include the *security practitioner* who responds to an incident and interacts with: (1) *IT specialists* who administer other systems (e.g., networks, databases); (2) *other stakeholders* from different areas (e.g., business, legal), who intervene depending on the incident (e.g., contacting the end-user, revising contracts with customers); (3) *end-users* who usually experienced the consequences of the security incident; (4) *external IT organizations* that administer systems interconnected in some way with the organization experiencing the incident (e.g, Internet service providers); and (5) *managers* from the organization, who need to be notified about the incident and coordinate the next steps. The notification information typically included: (1) *notifications* about new incidents, malicious traffic, or status of the investigation; (2) *requirements*, which usually consisted of messages for retrieving network or system configuration, or for starting the investigation of an incident; and (3) face-to-face or phone communications to *discuss* or *analyze* a security incident.

[6] A communication flow diagram shows how work is divided across people and how these people coordinate and interact to perform the required tasks to finish the job (Beyer and Holtzblatt 1998).

Security practitioners received notifications of security incidents from different stakeholders, especially from end users and other IT specialists. For example, one participant (I22) worked in an organization that controlled the access to library contents. This participant constantly referred to the need to interact with different stakeholders in order to receive notifications of anomalies, i.e., alarms, that could be related to malicious activity; an alarm might be triggered internally, by: (1) an IT specialist who detected peaks of traffic on the gateway servers; (2) a user who reported that the service was slow; or (3) directly by the systems that generated alarms upon the detection of traffic patterns in the network or servers (these systems are omitted from Figure 3.1 for simplicity). Alarms may also be triggered externally, by external stakeholders such as a content provider who detected unusual use of some of the resources in his databases. The information exchanged also varied with the type of notification: an e-mail including log files when the incident was detected by a vendor or another IT specialist, or just a phone call reporting that a service was slow in the case of a user. In the same vein, depending on the incident, a combination of communication channels may be necessary during the investigation. One participant (I15) described how, during an incident that compromised the performance of the whole network, communications included e-mails to notify people about the incident and share general information, as well as phone and face-to-face communications to make sure the practitioners had the same understanding of the situation.

Security incidents usually triggered multiple and complex interactions among the actors. For example, notifications from end-users saying that their Internet connections were slow might imply the participation of: (1) IT specialists, who were experts in specific operating systems; (2) the security practitioner who intervened when there was a compromise of data and; (3) end-users, who had to give more details about what was happening with their computers.

While Figure 3.1 shows the general case, Figure 3.2 describes a particular, complex case of interactions during a specific security incident, where more external agents are involved. This case was described by one participant (I29), whose organization received notifications from external organizations that had detected spam attacks coming from IP addresses administered by the organization where the participant worked. As these IP addresses were used by clients from that organization, this participant had to interact with other internal stakeholders (commercial and legal departments) to contact the clients. Most of the clients were not aware of any problem in

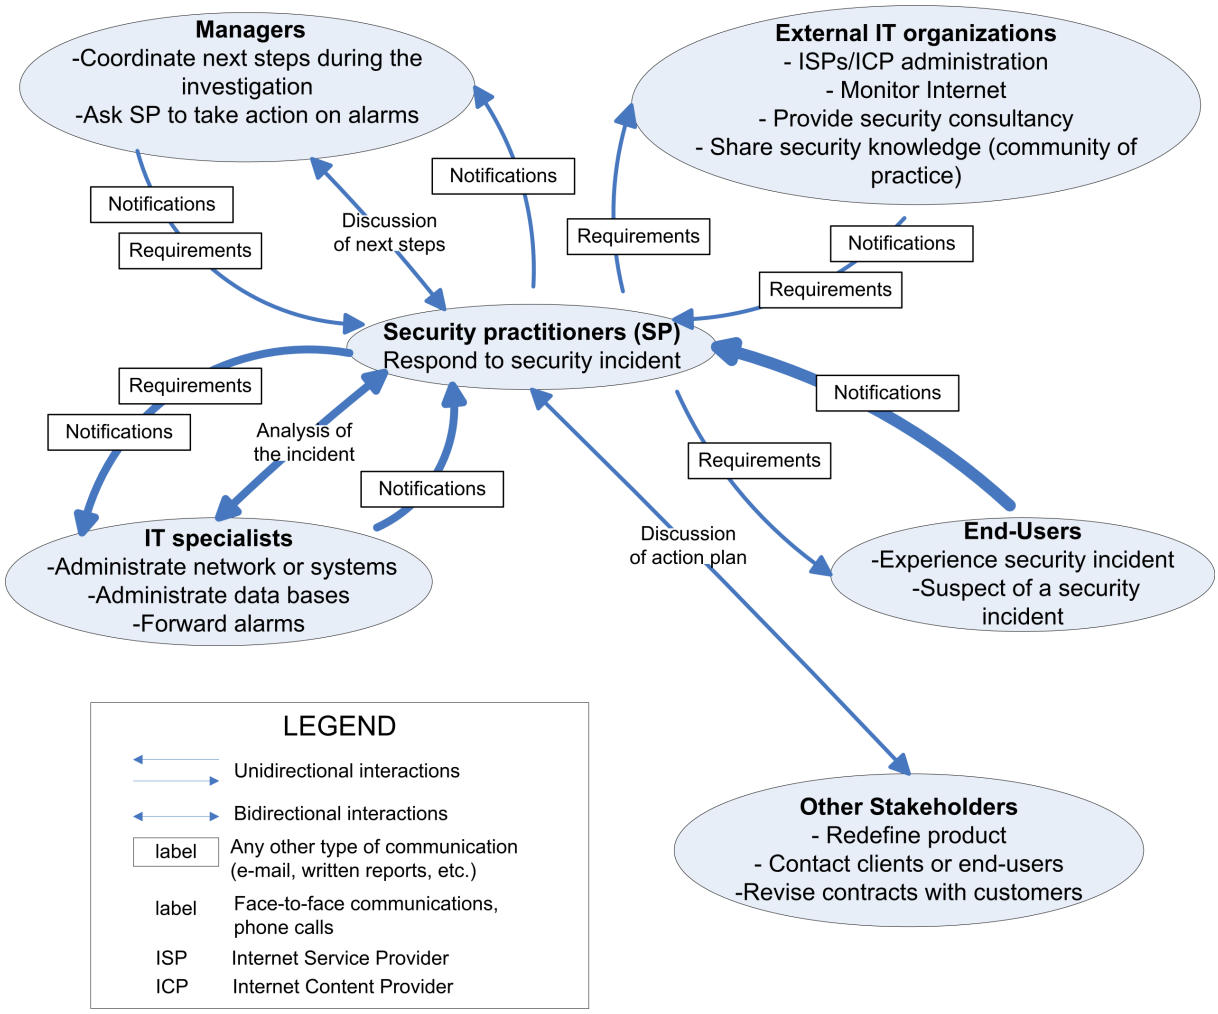

Figure 3.1: Responding to security incidents. Thicker arrows indicate more frequent interactions. For simplicity, only interactions between security practitioners and other stakeholders are shown.

their systems when they were notified about the situation. Some clients were very cooperative and promised to solve the problem; others claimed that they were victims of an external agent, and needed support from the participant's organization to clean their systems. In some situations, clients did not want to cooperate; the participant (I29) had to coordinate with other specialists to block Internet access from these clients' IP addresses. This step was necessary as the organizations that had detected attacks from the clients' IP addresses, were blocking not only those addresses, but also the neighboring addresses within the same segment. This blocking caused "good clients", who were not involved in the incident, to be unable to access their services due to the malicious traffic generated by the systems of the "bad clients". During the investigation of the incident, the participant (I29) also received requests from internal managers asking about the status of the investigation of the incident.

Another large-scale incident, in terms of the number of devices compromised by malicious software, represented an interesting challenge in terms of interactions. One participant (I4) described how, as the "owner" of an incident, he had to coordinate the activities of internal ad-hoc groups that were in charge of responding to the incident. Their main objective was to clean those organizational MS Windows machines that had been infected by a virus. The ad-hoc group consisted of approximately 20 people, most of them network and MS Windows specialists. They were organized in two layers: the first layer was in charge of evaluating the damage in terms of services affected. The other group had to analyze the malicious software and generate a plan to clean and patch the infected machines.

The above examples show the need to coordinate and respond to requirements from multiple stakeholders might make it necessary to define new procedures to establish formal responsibilities for the various stakeholders involved. For example, one participant (I29) mentioned how the incident illustrated in Figure 3.2 triggered a revision of not only the interactions between the internal specialists working on the investigation, but also of the contracts that this organization had with its clients. This revision included secure and responsible use of the Internet services. Similar conditions were also mentioned by another participant (I15), who described how they were able to disconnect from a network those clients that were saturating the network with malicious traffic and affecting other clients sharing the same resources.

Our participants had to interact with external stakeholders to receive support during the

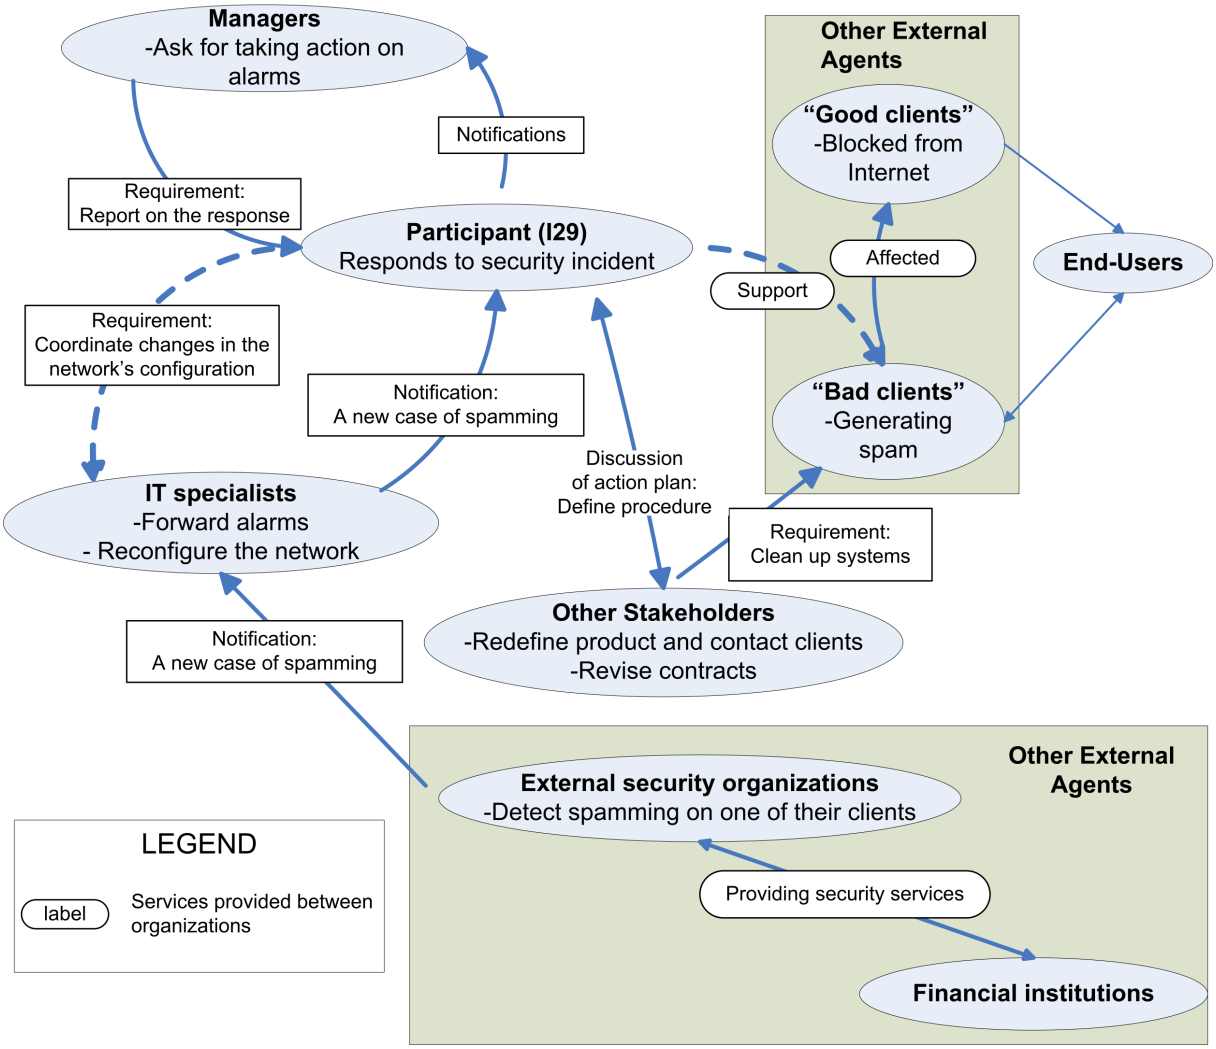

Figure 3.2: Response to an incident that triggers multiple and complex interactions among stakeholders. Dashed lines indicate two possible actions depending on the cooperation from the client. End-users are behind other agents, clients of the participant's organization.

investigation of security incidents. For example, one participant (I13) was trying to find the cause of a suspected security incident: *"So we are at that stage where we are trying to track down, looking through archives of a mailing list to see if anyone else has had similar problems."* Another example of external interactions occurred during a phishing attack. One participant (I4) had to coordinate with an administrator in Germany to take down a phishing web site.

Misunderstandings stemming from a lack of communication can make investigation of security incidents more difficult. For example, changes on the database servers that were not communicated promptly to network administrators made it more difficult to determine the cause of an availability incident (I7). Avoiding miscommunication was described as being important during the response to security incidents. For example, one participant (I3) reported constantly sending clarification questions through e-mail to avoid misunderstandings.

3.5.2 Development of Policies

In addition to incident response, our interview analysis also showed that interactions were extensive during development of a security policy. *Security practitioners* had to interact with: (1) *IT specialists* affected by the policy, who actively participated in developing the policies; (2) *external organizations* that might specify security requirements to be formalized in the security policy; (3) *end-users*, who might ask for revisions to a security policy and were affected by security policies; and (4) *managers*, who defined the scope of the policy and revised the policies. The exchanges of information during the development of policies included: (1) *drafts of the policy*, (2) the *policy* itself, (3) *requirements* about what the policy should include, and (4) meetings to *discuss* and *write* the policies. Figure 3.3 shows in detail the stakeholders involved during the policy development process and the corresponding flows of information.

As in security incidents, participants had to use multiple communication channels to interact with other stakeholders and get feedback from managers (see Figure 3.3). Additionally, data obtained from participatory observation showed that *threat analysis* and *tacit knowledge* about the organization were also important in interactions regarding security policy development. The following results are based primarily on the richer data that our participatory observation provided. We observed a policy development group of security and IT specialists led by a security practitioner, in an organization that did not have a centralized department devoted to IT security

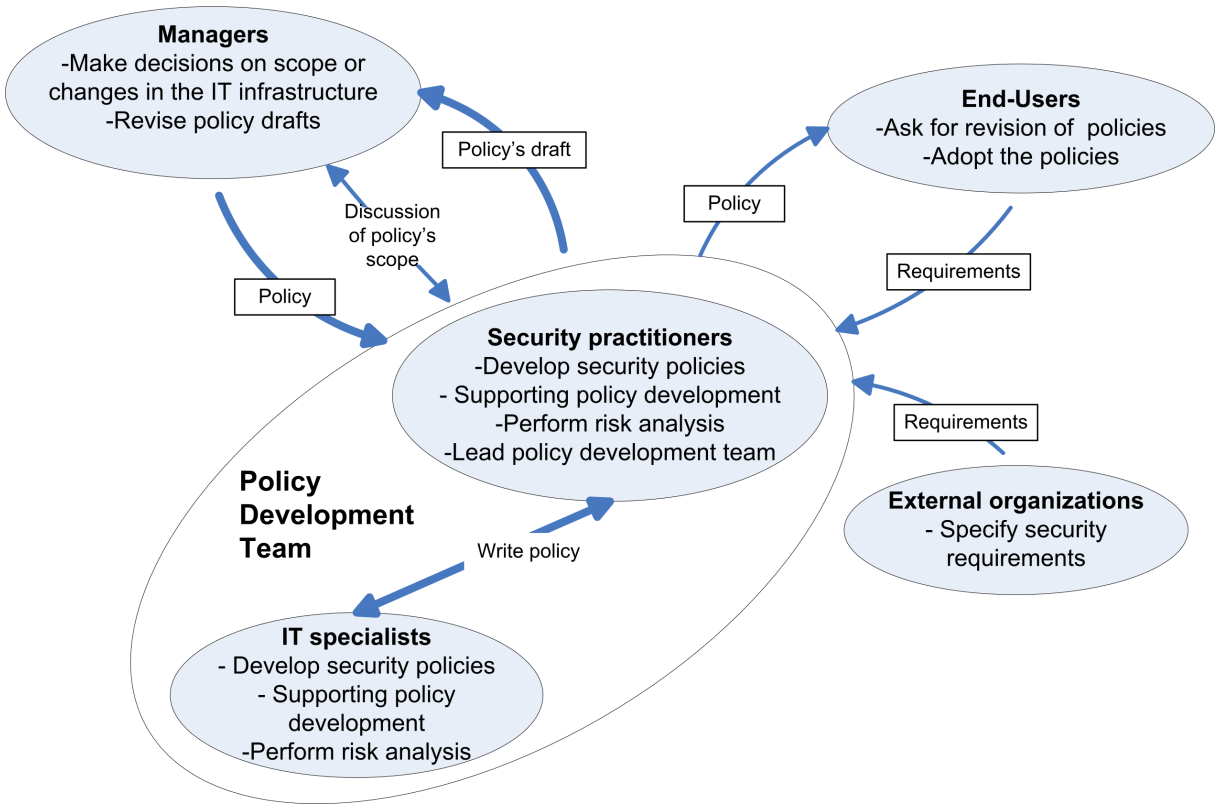

Figure 3.3: Communication flow diagram for developing security policies. Thicker arrows indicate more frequent interactions. For simplicity, only interactions with security practitioners are shown.

(for a discussion of centralized versus distributed security within organizations, see Hawkey et al. (2008b)). An internal web site accessible by all members of the group was the main repository for the drafts and related documents used during the policy writing process. E-mail was also used to share documents with the whole group (for simplicity, these systems are omitted from the diagram).

Threat analysis was necessary in order to cover all possible circumstances in which the policy should apply. Threat analysis allowed our participants to map different risks with the text in the policy. Tacit knowledge was required to devise "implementable" policies, in terms of matching security principles (e.g., confidentiality of sensitive information) with the tasks of different stakeholders. For example, our participants had to know how different specialists made use of the information on the servers, before imposing restrictions on the use of that information.

Another issue uncovered during the participatory observation was related to the knowledge of IT security tools. Our participants needed to know how general IT and security tools could

be used to implement the principles stated in the policies. IT specialists involved in the process had to iteratively complement the policy text, considering how tools were able to support the implementation of the controls stated. For example, for a policy related to data protection, the requirements concerning encryption of critical data made it necessary to study how different encryption tools could be adapted to the organization's needs. This process of understanding how different encryption tools could be used in real settings not only made the process of writing the policy longer, but also confirmed the general finding of Botta et al. (2007) of the importance of accessible and clear documentation about what security tools can and cannot do.

Group members we observed and worked with tried to avoid misunderstandings with managers by continually asking for their feedback on, for example, the topics covered by the policies. This practice was necessary since a previous attempt at writing policies had failed because the policies proposed did not meet the expectations of managers.

3.6 Modeling the Complexity of Interactions

The two scenarios described in Section 3.5 illustrate the richness and complexity of interactions performed by security practitioners, and can also be used as a reference for the complex environment where security tools should be tested (Redish 2007). For example, a security tool intended to support the scenario described in Figures 3.1 and 3.2 must support not only correlation of information from unrelated sources, but also the integration of communication features so that security practitioners can interact with the different stakeholders involved.

We now present a model that integrates our findings and presents the factors that determine the complexity of interactions required to perform security tasks (see Figure 3.4). This model is also used to discuss how such complexity might affect the security of the organization. In building the model, continued posterior analysis allowed us to group our findings into a hierarchical construction of categories.

The central, most general category of our model is *complexity of interactions*. This complexity is determined by three different high-level categories: organizational attributes, multiple stakeholders, and multiple security-related activities. Each high-level category has detailed subcategories, which include relationships with other subcategories that arose from our analysis in Chapter 2. Future work is needed to validate these categories and connections.

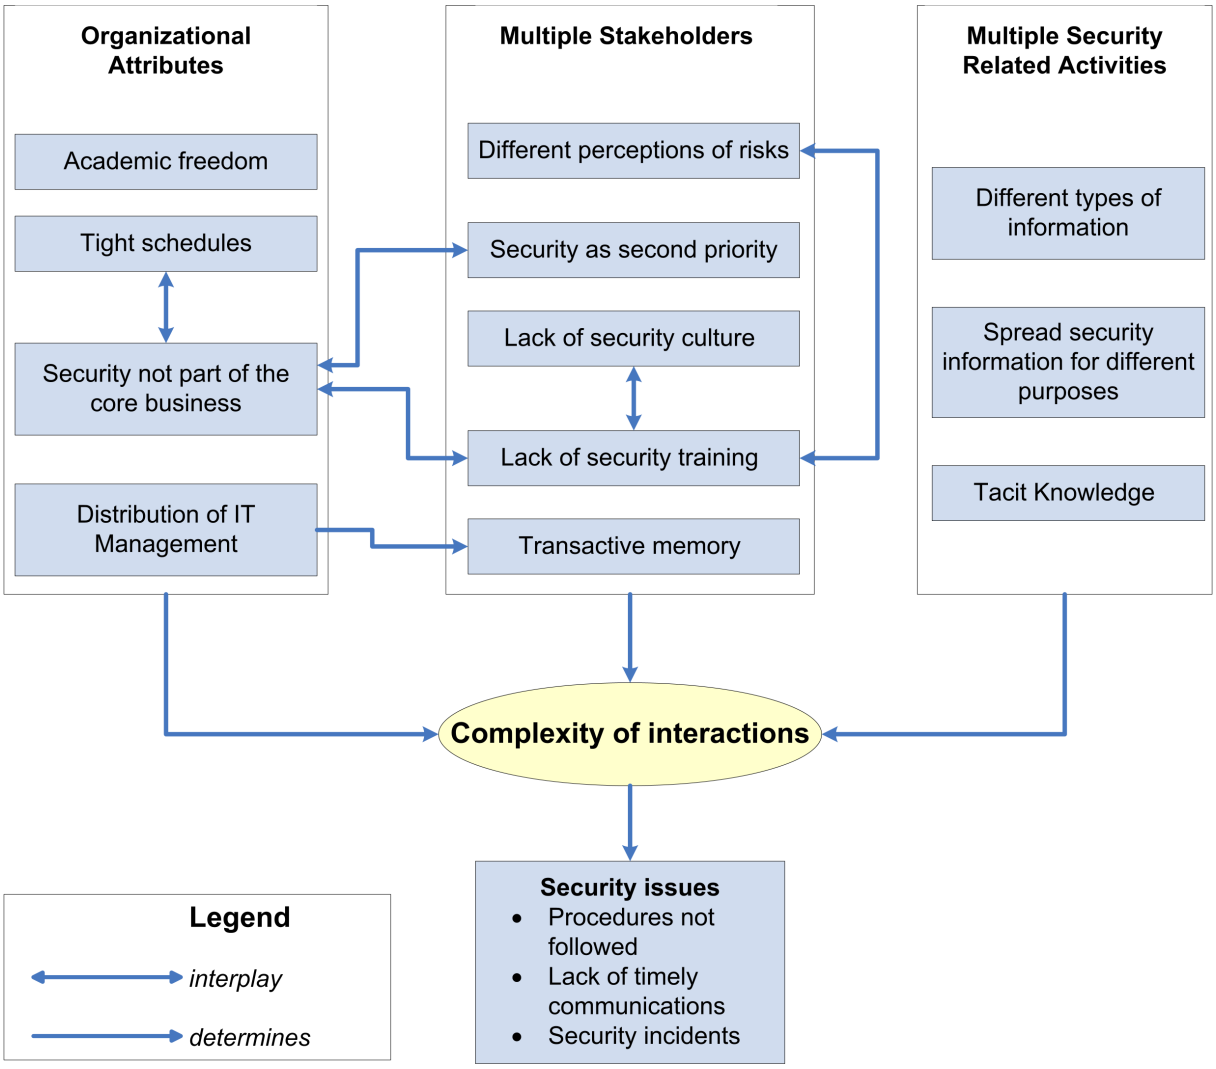

Figure 3.4: Factors that make interactions more complex for security practitioners within organizations.

This model can be used to explain the complex interactions that security practitioners face when performing their activities (see Section 3.4). For example, when designing services with security requirements (second activity in Table 3.2), security practitioners who work in a company with the *organizational attribute* of not having *security as part of its core business*, would have to convince other stakeholders of the need to consider security controls from the beginning of the project. Because of this organizational attribute, the *multiple stakeholders* involved in the project (e.g., different IT specialists), would: (1) *not have security training or work in a security culture*; (2)*not have IT security within their priorities*; and (3) *have different perception of security risks*. These factors make it difficult for security practitioners to explain the importance of security controls to the other stakeholders involved in the project. Another dimension of this complexity is given by the *multiple security-related activities* performed by security practitioners; they have to manage their priorities and the *types of information* involved when they have competing priorities with other security tasks (e.g., responding to a security incident). In this example, the consequences of the complexity of interactions might be the *lack of timely communications* about the new project and, in the end, the lack of security controls in the service developed. In general, our analysis has shown that the complexity of interactions for security practitioners causes security issues that make organizations more vulnerable and increase their security risks.

We next describe each factor that contributes to the complexity of interactions for security practitioners, and also elaborate on the security issues that this complexity raises for organizations.

3.6.1 Organizational Attributes

Tight schedules made interactions for the security practitioners we interviewed more complex. Participants had to effectively communicate what was important in terms of security, without oversimplifying the importance of security controls. When security was not a priority within the organization (i.e., a manufacturing organization that does not have security as part of its core business), it was more difficult for security practitioners to devote the time required to analyze and apply security in the organization's projects.

When *security was not integrated in the core business*, it made it difficult for our participants to communicate security principles that should have been considered from the beginning of the different projects within the organization. In this vein, Flechais and Sasse (in press) point out that

the application of security requirements during project implementation increases costs, although they do not specify which types of costs are in play. Our analysis shows that when security was not integrated in the projects, there was more communication and interaction overhead for our participants, who had to interact more actively with other specialists to try to understand the design of the project and propose security controls.

We found that *distribution of IT management* made our participants rely on other specialists to integrate different sources of information (e.g., to match IP addresses with contact information from end users). In these cases, where communication was usually in the form of requests made by e-mail, a lack of a prompt response can cause delays in the investigation of the detected anomaly.

Academic freedom was a factor that made interactions more complex in academic institutions. The main issue was the lack of standardization within the organization in terms of priorities and stakeholder knowledge of IT security. The results of Flechais and Sasse (in press) also show the complexity of academic environments in terms of the high variation in security knowledge of the stakeholders involved. This factor was directly related to the different perceptions of security risks that various stakeholders had within academic institutions. Failure to arbitrate conflicting perceptions of risk can compromise the organization.

3.6.2 Multiple stakeholders

The involvement of multiple stakeholders was another factor that made interactions more complex. In most of the participants' organizations, IT security-related activities required interaction between a variety of different stakeholders. Knapp et al. (2005) also identify this characteristic of interdependency of IT security tasks. Our analysis expands their results and highlights how this interdependency makes interactions more complex for security practitioners.

Our participants had to communicate with other stakeholders who had *different perceptions of risks*, considered *security as second priority*, and did not have a *security culture or training*. These characteristics combine to determine what Flechais and Sasse (in press) identify as motivation of the stakeholder. Our participants constantly had to persuade other stakeholders who had different motivations, of the importance of security controls. In this process, the participants' communication style was important in approaching stakeholders who did not share the same perception of risks. For example, one participant (I25) expressed the need for diplomacy to achieve cooper-

ation. Koskosas and Paul (2004) studied how risks are communicated in financial organizations. They conclude that risk communication "plays a significant role at the macro-goal level of security management," and affects the setting of banking security goals. Our analysis provides further empirical evidence over a wider range of organizations about the importance and complexity of communicating risks for security practitioners. We show how security practitioners assume the role of "risk evaluators" during interactions with other stakeholders.

Our participants expressed the need to know which stakeholders they must interact with depending on the type of activity. Distribution of IT management heightened this need, as participants had to know who administered what. This requirement suggests that IT security practitioners tend to be centers of *transactive memory*, a kind of mutual understanding about who knows what. "Transactive memory theory is based on the idea that individual members can serve as external memory aids to each other" (Wegner 1986). For example, to respond to security incidents, they needed to know which specialists had to be involved in the investigation, depending on the systems compromised. The need for using transactive memory made interactions more complex, as it required knowing the organization and the roles that each stakeholder had within it.

3.6.3 Multiple security-related activities

Our results show that the IT security practitioners we interviewed had to show significant diversity in the way they communicate, as indicated by the variety of high-level tasks that contextualize their interactions. Eight of our participants (I2, I4, I5, I15, I22, I24, I25, I30) described being involved in at least three different types of activities.

The different activities required that our participants exchange *different types of information*. Examples of the information exchanged were requirements (e.g., write a security policy), reports (e.g., vulnerability scans for audits) and notifications (e.g., security alarms). In order to exchange this information, our participants had to not only use different communication channels, but also needed to manually integrate the outputs of their security tools with the inputs of their communication tools (e.g., attach a report from security scanner to an e-mail, attach log files). Security incident response represented a fairly complex scenario where practitioners needed to use different communication channels to interact with different stakeholders.

The need to *distribute security information for different purposes* made communications more

complex. Our participants needed good communication skills to adapt interactions to the context of the activity; they had to be reactive to solve IT security issues of end users, manage new vulnerabilities, and respond to incidents. They also had to be proactive to perform audits, design new services, implement security controls, educate and train stakeholders, develop policies and communicate risks.

Our participants had to use *tacit knowledge* to perform their activities. For example, in order to write policies, they had to know about other stakeholders' tasks and how security controls would be integrated with those tasks. To integrate security with new IT services, they had to know about the services the organization provided. To implement security access controls, they had to know about the different activities that stakeholders performed depending on their roles.

3.6.4 Consequences of the complexity of security interactions

The complexity of security interactions had implications for the work performed by our participants and for the security of their organizations. For example, several types of miscommunications were mentioned during the interviews, including not following preestablished procedures and not communicating in a timely fashion.

Stakeholders often did *not follow security procedures*, particularly when IT management was highly distributed, security was not considered part of the organization's core business, and there were stakeholders involved without security backgrounds. Not following security procedures generated communications overhead. For example, one participant (I2) highlighted the consequences of not following a change-management procedure aiming at integrating security with other activities, such as the design of new projects and day-to-day operations. When this integration did not exist and security was incorporated as an add-on at the end of the day, security specialists needed much more information and communication with the other stakeholders to understand what had been done and how to apply security requirements to a system already implemented.

Lack of timely communications was another issue mentioned by our participants. High workloads interfered with communication; our participants had no time to notify involved parties of changes during quick responses to incidents. Given the complexity of the IT infrastructure, IT specialists might not anticipate the consequences of local changes in other network domains, and thereby consider it unnecessary to inform other parties about reconfiguration of systems. Lack of timely communications with vendors was also mentioned.

Breakdowns of IT security interaction relate to information errors, according to Hinckley's classification (Hinckley (2001) citepd by Chao and Ishii (2004)). The framework developed by Kraemer and Carayon (2007) suggests that a heavy workload and a lack of formal communications lead to errors that affect the security in organizations. We also found that ineffective interactions can be the source of security incidents, or can increase risk levels. For example, a lack of communication when making changes in firewalls can cause connection problems for other users of the network, or a slow response from a vendor about new patches can expose the IT infrastructure to attacks.

3.7 Implications of Findings

The need for better support for collaboration in security tools has been recognized previously. Goodall et al. (2004) report on this need for one specific type of tool, namely IDSs. IDSs should provide better support for security experts collaborating with other security experts around the world. Our empirical analysis showed that our participants have to use communication channels that are not integrated with their security tools and do not always cover all their needs. For example, they needed to avoid the possibility of misunderstandings during communications while keeping track of agreements for future audits. We next provide guidelines for improving security tools and alleviating the complexity of interactions that security practitioners face when performing security-related activities. We also indicate, where possible, specific opportunities for implementing these guidelines.

Integrate different communication channels: Our participants had to send and receive notifications, reports, and requirements (see Figures 3.1 and 3.3) to communicate with different stakeholders. Security tools would provide better support if they are able to integrate different communication channels, accepting as inputs and producing as outputs data in different formats from different communication or security tools. For example, in the scenario described in Figure 3.1, the security tools used by the security practitioner to obtain reports of malicious traffic in the systems should be able to exchange and process the outputs from the different communication tools used by other stakeholders (e.g., e-mail, Pdf or Html report, text file). In this case, the security tool would integrate and consolidate different sources of information, alleviating the

burden of copy-pasting outputs from communication to security tools and vice-versa.

In the same vein, security tools should provide open interfaces to integrate easily with existing communication tools such as e-mail clients and text chat. This integration would allow not only quick interactions with different stakeholders, but also the option of directly sharing the information generated by security tools according to each stakeholder's access privileges.

Reduce communication overhead: Similar to the previous point, security practitioners need tool features to reduce communication overhead. For example, one of our participants used an embedded feature of a spam filter tool to publish the status of users' e-mails on a web page. This way, he avoided questions from the end users about what happened with their e-mails when a new spam rule was added. This approach represents another opportunity for designing communication support for security tools.

Implement security domains when communicating security issues: Increased flexibility to communicate and share information generated by security tools would still not be enough to support the interaction needs of security practitioners. It is also important to consider the specific constraints of security communications. Our analysis showed that security practitioners need to communicate with external stakeholders frequently. These communications require encryption, which should be embedded in security tools that produce reports. For example, a tool that generates reports about virus activity should provide open interfaces to be integrated easily with VPN clients. This integration would avoid errors of sending sensitive information to external stakeholders without the required encryption.

Provide customizable accounts for stakeholders with different goals: In a distributed IT environment, systems are interconnected but are administered and managed by different IT specialists. In this case, security tools not only need to support different levels of access in a *vertical way* (i.e., regular user vs. administrators), they also need to provide different configuration options for improving collaboration among IT practitioners from different domains. For example, an intrusion detection system that is monitoring different systems and networks should have the option of configuring various accounts to monitor the different networks or systems independently. To separate these networks or systems, multiple differentiation criteria can be provided: IP addresses, type of operating system used (e.g., Windows, Unix), or type of network protocol.

Provide reporting options that show the level of risk: Another opportunity to improve security

reporting is by providing security practitioners with better features to interpret and communicate the information from the analysis that security tools perform. For example, security tools should generate reports that indicate the levels of risk in the IT infrastructure—specifying status of security patches and antivirus updates. This characteristic might help security practitioners to prioritize their tasks.

Provide flexible reporting: Botta et al. (2007) identify the need for flexible reporting to support some security-related tasks, like communication with different stakeholders who have varying levels of expertise. Our current analysis indicates that flexible reporting can be broken down into the following characteristics: on-line and automatic generation of different reports for different stakeholders, and the use of different layers of information (general vs. specific). This last requirement confirms the proposal from Chiasson et al. (2007), of using ecological interfaces to design security systems, showing security information in five levels of abstraction, with different levels of detail depending on the user.

Correlate data that include not only IT databases: The need to be able to address new security incident scenarios (see Figure 3.2) makes it necessary to correlate information in novel ways. For example, it is common that an IP address is the only information that a security specialist has to determine who caused a security incident. Using this technical information, the specialist has to correlate it with internal proprietary databases containing customer information. Security tools should afford the implementation of new types of queries that look for matching information in databases with different formats, implemented with different purposes within the same organization.

Provide notification of configuration changes and alarms in distributed environments: To avoid errors during interaction, our participants used checklists, proactive communications, and training. These strategies may also provide opportunities for tool development. For example, firewall management systems could have a list with contact information from different stakeholders who need to be informed about configuration and other changes. Each stakeholder could respectively receive the information at the appropriate level of detail, language, and channel (e-mail, text message, web site). Furthermore, security tools should consider distributed organizational structures where different IT specialists manage different domains of the networks and systems. For example, a security tool could integrate not only features to monitor and analyze those devices that are let-

ting attacks pass through the internal network, but also notify the corresponding administrator via e-mail to take action and stop the malicious traffic.

Manage tacit knowledge: Our participants managed their tacit knowledge when they: (1) provided statements of evaluation when playing the role of "risk evaluator" and (2) developed training programs for other specialists. Kesh and Ratnasingam (2007) highlight the need for transforming tacit security knowledge into explicit knowledge. There is some debate as to whether or not such a thing is feasible (Schmidt 1997), or desirable (why should they give away their stock-in-trade?). Flechais and Sasse (in press) propose the use of scenarios as an effective tool to help clearly explain abstract security concepts to other stakeholders. Our results show that scenarios might be effective in the process of transforming tacit knowledge into explicit knowledge.

Our analysis also showed that the process of developing security policies could help security practitioners to transform their knowledge between tacit and explicit forms. Using Marwick's (2001) analysis of technologies used for creating organizational knowledge, development of policies can be broken down into the following steps. First, find templates about policies on the Internet, using a browser and a search engine (explicit-to-explicit knowledge). Second, interpret the meaning of other organizations' policies (explicit-to-tacit knowledge). Third, adapt the templates and information found using tacit knowledge of the organization and hold internal meetings to discuss experiences with security issues (tacit-to-tacit knowledge). Fourth, disseminate the policies by presenting them in meetings and on internal web sites (tacit-to-explicit knowledge). We propose that organizations could take more advantage of this process by involving other stakeholders in it. This process could entail the use of scenarios or anecdotes, as proposed by Flechais and Sasse (in press).

3.8 Conclusion

Our qualitative analysis shows the complex environment where security practitioners not only perform security-specific tasks, but also interact with stakeholders with different backgrounds and needs. We have developed a model that shows the factors that make these interactions complex, and the security issues that are a consequence of this complexity.

Security tools used by security practitioners do not provide enough support for the highly

interactive environment they work in. We have offered guidelines to develop more effective security tools. We have also elaborated on two scenarios that illustrate the richness and complexity of the interactions performed by security practitioners, and that can be used as reference environments for evaluating security tools.

We have only begun to answer questions on the complexity of interactions performed by security practitioners. More research is needed to expand and refine our understanding of the interactions with respect to different types of contexts.

Table 3.2: Types of activity in which IT-security communication occurs

Activity	**Interviews**		**Stakeholders involved**
	Academia	Private	
Perform and respond to security audits	I2	I4, I5, I16, I23, I25, I30	1. Coordinate or collaborate with IT specialists 2. Coordinate with auditors
Design services incorporating security requirements	I2, I11, I14, I15, I17	I25, I30	1. Coordinate and collaborate with other IT specialists 2. Coordinate and collaborate with organization's multidisciplinary committees 3. Coordinate with vendors of security technology
Solve IT security issues of end users	I3, I10, I15	I21, I30	1. Cooperate and collaborate with IT specialists 2. Cooperate with external specialists from the organization 3. Coordinate with end users
Implement security controls	I22	I4, I5, I21, I28, I29	1. Cooperate with other IT specialists 2. Coordinate with other areas in the organization (e.g., Human Resources)
Educate and train other employees	I15	I5, I16, I25, I30	1. Cooperate with IT specialists 2. Cooperate with managers/executives 3. Cooperate with end users
Mitigate vulnerabilities	I2, I9, I22, I24		1. Cooperate with other IT specialists 2. Coordinate with vendors of security technology 3. Cooperate with external IT security entities
Administer security devices	I24	I28, I30	1. Coordinate with other IT specialists
Respond to security incidents	I1, I2, I3, I7, I9, I11, I12, I13, I15, I17, I18, I20, I22, I24	I4, I5, I26, I29	1. Coordinate and cooperate with other IT specialists 2. Coordinate and cooperate with specialists from legal department 3. Coordinate and cooperate with external specialists (from the organization) 4. Coordinate with vendors of security technology
Develop security policies	I1, I2, I24	I23, I25, I30	1. Coordinate and collaborate with other IT specialists 2. Coordinate with end users 3. Coordinate and collaborate with managers/executives

Bibliography

H. Beyer and K. Holtzblatt. *Contextual Design, Defining Customer-Centered Systems*. Morgan Kaufmann Publishers, San Francisco, CA, 1998.

K. Beznosov and O. Beznosova. On the imbalance of the security problem space and its expected consequences. *Information Management & Computer Security*, 15(5):420–431(12), 2007.

F. J. Björck. *Discovering Information Security Management*. Doctoral thesis, Stockholm University, Royal Institute of Technology, 2005.

D. Botta, R. Werlinger, A. Gagné, K. Beznosov, L. Iverson, S. Fels, and B. Fisher. Towards understanding IT security professionals and their tools. In *SOUPS*, pages 100–111, Pittsburgh, Pennsylvania, July 18-20 2007.

J. M. Carroll, M. B. Rosson, G. Convertino, and C. H. Ganoe. Awareness and teamwork in computer-supported collaborations. *Interact. Comput.*, 18(1):21–46, 2006. ISSN 0953-5438. doi: http://dx.doi.org/10.1016/j.intcom.2005.05.005.

L. P. Chao and K. Ishii. Design error classification and knowledge management. *Journal of Knowledge Management Practice*, 5, 2004.

K. Charmaz. *Constructing Grounded Theory*. SAGE publications, 2006.

S. Chiasson, P. C. van Oorschot, and R. Biddle. Even experts deserve usable security: Design guidelines for security management systems. SOUPS USM Workshop, July 2007.

D. M. Fetterman. *Ethnography: Step by Step*. Sage Publications Inc., 1998. ISBN 0761913858.

I. Flechais and M. A. Sasse. Stakeholder involvement, motivation, responsibility, communication: How to design usable security in e-science. *Int. Journal of Human-Computer Studies*, in press.

J. R. Goodall, W. G. Lutters, and A. Komlodi. I know my network: Collaboration and expertise in intrusion detection. In *CSCW*, volume 6390, November 2004.

E. Haber and E. Kandogan. Security administrators: A breed apart. In *Proc. of SOUPS Workshop on Usable IT Security Management (USM), 4 pages*, 2007.

K. Hawkey, D. Botta, R. Werlinger, K. Muldner, A. Gagne, and K. Beznosov. Human, organizational, and technological factors of it security. In *CHI'08 extended abstract on Human factors in computing systems*, pages 3639–3644, 2008a.

K. Hawkey, K. Muldner, and K. Beznosov. Searching for the Right Fit: Balancing IT Security Model Trade-offs. *Special Issue on Useful Computer Security, IEEE Internet Computing*, 12(3): 22–30, 2008b.

C. M. Hinckley. *Make No Mistake.* Productivity Press, Portland, OR., 2001.

E. Kandogan and E. M. Haber. Security administration tools and practices. In L. F. Cranor and S. Garfinkel, editors, *Security and Usability: Designing Secure Systems that People Can Use*, chapter 18, pages 357–378. O'Reilly Media, Inc., 2005.

S. Kesh and P. Ratnasingam. A knowledge architecture for IT security. *Commun. ACM*, 50(7): 103–108, 2007. doi: http://doi.acm.org/10.1145/1272516.1272521.

K. J. Knapp, T. E. Marshall, R. K. Rainer, and F. N. Ford. Managerial dimensions in information security: A theoretical model of organizational effectiveness. https://www.isc2.org/download/auburn_study2005.pdf, 2005.

I. V. Koskosas and R. J. Paul. The interrelationship and effect of culture and risk communication in setting internet banking security goals. In *ICEC '04*, pages 341–350. ACM Press, 2004. ISBN 1-58113-930-6. doi: http://doi.acm.org/10.1145/1052220.1052264.

A. G. Kotulic and J. G. Clark. Why there aren't more information security research studies. *Information & Management*, 41(5):597–607, 2004.

S. Kraemer and P. Carayon. Human errors and violations in computer and information security: The viewpoint of network administrators and security specialists. *Applied Ergonomics*, 38:143–154, 2007.

A. D. Marwick. Knowledge management technology. *IBM Systems Journal*, 40(4):814–830, 2001.

P. W. Matessich and B. R. Monsey. *Collaboration: what makes it work. A review of research literature on factors influencing successful collaboration.* Amherst H. Wilder Foundation, St. Paul, MN, 1992.

S. Mohammed and B. Dumville. Team mental models in a team knowledge framework: Expanding theory and measurement across disciplinary boundaries. *Journal of Organizational Behavior*, 22 (2):89–106, March 2001. ISSN 0894-3796.

J. Redish. Expanding usability testing to evaluate complex systems. *Journal of Usability Studies*, 2(3):102–111, 2007.

M. Sandelowski. Whatever happened to qualitative description? *Research in Nursing & Health*, 23(4):334–340, 2000.

K. Schmidt. Of maps and scripts—the status of formal constructs in cooperative work. In *ACM SIGGROUP*, pages 138–147, November 1997. ISBN 0-89791-897-5.

D. M. Wegner. *Transactive memory: A contemporary analysis of the group mind.* In B. Mullen and G. R. Goethals, Editors, Theories of Group Behavior, 1986.

R. Werlinger, K. Hawkey, and K. Beznosov. Security practitioners in context: their activities and interactions. In *CHI '08 extended abstracts on Human factors in computing systems*, pages 3789–3794, 2008a.

R. Werlinger, K. Hawkey, and K. Beznosov. Human, Organizational and Technological Challenges of Implementing IT Security in Organizations. In *to appear in HAISA'08: Human Aspects of Information Security and Assurance (13 pages)*, July 2008b.

Chapter 4

Diagnostic Work During Security Incident Response: A Qualitative Study[7]

4.1 Introduction

Information Technology (IT) has become pervasive in today's organizations, making it especially critical that IT assets are effectively protected against IT-related threats. Traditionally, research has focused on technical aspects of IT security, such as algorithms for intrusion detection systems, firewalls, and virtual private networks (VPNs) (e.g., Chebrolua et al. 2005). Other studies, however, have suggested the need for a richer understanding of how additional dimensions impact IT security management (ITSM), including the human and organizational (Botta et al. 2007b; Beznosov and Beznosova 2007; Rayford B. Vaughn Jr. and Fox 2001). Furthermore, researchers have also proposed investigation is needed of the roles, responsibilities, and tasks performed by *security practitioners*, those professionals responsible for managing IT security within organizations (Botta et al. 2007b).

Diagnostic work, i.e., the practice of noticing and categorizing problems, as well as defining the scope of remediation, is a pervasive feature of ITSM. Diagnosis is particularly prevalent during security incident response, which is one of the primary responsibilities of security practitioners (Botta et al. 2007b; Kandogan and Haber 2005). Despite this fact, however, the field of security incident response is still in its infancy (Killcrece et al. 2005). While a number of organizations provide

[7] A version of this chapter has been submitted to a journal for publication. R. Werlinger, K. Muldner, K. Hawkey, and K. Beznosov (2008) Diagnostic Work during Security Incident Response: A Qualitative Study.

guidelines for the incident response process[8], there are few empirical investigations on how security practitioners respond to incidents (for exceptions, see for instance Goodall et al. 2004a; Riden 2006). The research presented in this paper aims to fill this gap. Such research contributes with real-world data on actual experiences, which can substantially add to the understanding of how to improve support for practitioners during security incident response.

The work reported in this paper is part of a larger project that aims to further the understanding of the human, organizational, and technological factors of ITSM.[9] To date, project researchers have conducted 34 semi-structured *in-situ* interviews with security practitioners from a variety of academic and private sectors. The first author has been also involved in an ongoing participatory observation under the supervision of a senior security practitioner in one academic organization in Canada. We analyzed the interview data using using qualitative description (Sandelowski 2000), focusing on pre-designed themes of analysis (e.g., tools, tasks, interactions). For instance, in the *interactions* theme, our analysis identified nine activities that require security practitioners to interact with other stakeholders, one of which is security incident response (for preliminary results, see Werlinger et al. 2008a). Our work in Werlinger et al. (2008a) and research by others (e.g., Goodall et al. 2004a) have highlighted the highly collaborative nature of ITSM, including security incident response. Furthermore, our preliminary work on security incident response (Werlinger and Botta 2007) has emphasized the need for more investigation of the diagnostic aspects of this critical activity.

In this paper, we present results from our analysis for the theme of security incident response, with the focus on how security practitioners *diagnose and troubleshoot IT systems to detect anomalies and security incidents*. Furthermore, since preparation for security incident response is a key aspect of the response process (Mitropoulos et al. 2006), we also analyze and highlight the diagnostic aspects of the preparation stage. The analysis is based on data from our notes obtained during participatory observation, as well as on a subset of the 34 semi-structured interviews, namely 13 interviews that provided detailed stories on diagnostic aspects related to security incident response.

The contribution of our work presented here is twofold. First, using empirical data from interviews and the participatory observation, we analyze and describe the tasks, skills, strategies,

[8] e.g., Computer Emergency Response Team (CERT), National Institute of Standards and Technology (NIST)
[9] See Hawkey et al. (2008) for a project overview.

and tools that security practitioners use to diagnose security incidents. The results suggest that security practitioners' diagnosis is complicated by human, organizational, and technical factors arising from the multi-faceted nature of ITSM. These factors include, for instance, practitioner's reliance on tacit knowledge during incident response, the need to collaborate with various stakeholders, and the complexity of today's technologies involved in diagnostic process, compounded by usability issues with IT security tools. This enhanced understanding of the diagnostic work during security incident response can support the specification of the complex scenarios in which the tools used during the diagnosis of security incidents should be tested (Redish 2007).

Second, equipped with the understanding of diagnosis during security incident response, we identify opportunities for future research directions related to improving security tools. For instance, our analysis shows that no matter how advanced a security tool is for supporting diagnostic work, practitioners must still customize that tool to fit the specific needs of their organization. This customization process is a challenging task due to the dynamic nature of the IT infrastructure and high expertise demands, including tacit knowledge about the organization and its users. Today's tools, however, provide very little if any support for this customization process. We propose some guidelines for how tools should provide this support, for instance via explicit scaffolding built into security tools to capture practitioners' tacit knowledge. Other aspects of tool improvements that we discuss relate to: the tradeoff between task complexity and tool reliability, the need for tools to support tailorability and correlation of high volumes of data, as well as multi-faceted simulation support.

The rest of the paper is organized as follows. Section 4.2 discusses the related work. Section 4.3 explains the study methodology. Section 4.4 reports the results. Section 4.5 summarizes our findings and provides suggestions for research directions on improving security tools to better support practitioners in diagnostic aspects of security incident response.

4.2 Related Work

In general, diagnosing and responding to incidents tends to be cognitively demanding (Goodall et al. 2004a); consequently, there is work in the Artificial Intelligence community on devising computational support for these processes in a variety of areas, such as medicine, automotive and

security (Rao et al. 1998; Heckerman et al. 1995; Shayman et al. 2000). Although promising, to date computational approaches have at best complemented rather than replaced human experts who perform diagnostic work. In general, to effectively provide any type of computational support, from simple tools to complex adaptive technologies, a solid understanding of how diagnosis is performed in a given field is necessary. As far as IT security is concerned, however, there has been lack of such understanding.

Since a background in ITSM is helpful for understanding diagnostic work during security incident response, we begin with a review of this background, and then discuss security incident response.

4.2.1 ITSM: Background

Role of Communication/Collaboration. There is evidence that both communication and collaboration play a key role during ITSM. Werlinger et al. (2008a) relied on qualitative description to analyze 34 interviews with security practitioners to identify when and how they interact with other stakeholders. The analysis identifies eight activities that involve collaboration and cooperation between security practitioners and end users, managers, and other specialists. Werlinger et al. also analyzed the tools used for interactions and found that existing tools do not provide adequate support. Siegel et al. (2006) performed contextual inquiry of 30 security practitioners at three organizations; one of the key findings show that these practitioner have difficulty effectively communicating with organizational stakeholders, which reduces management buy-in. Kraemer and Carayon (2007) performed 16 semi-structured interviews with network administrators and security practitioners to identify factors contributing to errors in ITSM. Their findings show that factors such as communication, security culture and organizational structure were all responsible for errors and vulnerabilities.

Based on a retrospective of a big security incident in 1988, Spafford (2003) compares the incident to the state of IT security in 2003, and concludes that there are several aspects that have become worse in terms of security since 1988. Although most of these aspects are technical (e.g., increase of security flaws in software), Spafford highlights that the security community has been unable to learn from past experiences the importance of communication during security incident

response. He proposes that the security community should find better ways of not only coordinating during security incidents, but also of distributing the information and reports generated after these incidents.

ITSM Challenges. Research to date suggests that security practitioners operate in a complex and challenging environment and often lack appropriate support. Botta et al. (2007a) analyzed 14 interviews to find that security is distributed across tasks and stakeholders, making coordination challenging. Their other finding is that security practitioners felt tools provided inadequate support. In Haber and Bailey (2007), the authors rely on naturalistic observation to study IT professionals in six organizations. Their major findings are that system administrators need better tool support and compared to end-users, they deal with larger, more complex systems and face a higher risk of failure. Werlinger et al. (2008d) also rely on qualitative analysis of 34 interviews, here to identify the challenges of ITSM. The results correspond to a framework that classifies the challenges according to the human, technological and organizational dimensions, and shows how the challenges interplay with one another.

We now describe guidelines and empirical research that specifically target security incident response.

4.2.2 Guidelines for Security-Incident Response

Given the challenges associated with managing security incidents in ITSM, including preparation, diagnosis and response, a number of guidelines (e.g., Casey 2002; Stephenson 2004) and associations exist that provide support for the incident response process (e.g., Computer Emergency Response Team, CERT, National Institute of Standards and Technology, NIST). Recently, Mitropoulos et al. (2006) have synthesized the information in the various standards as well as existing research to propose a general incident response management framework. In this framework, a variety of stakeholders (e.g., security practitioners, legal advisors, managers) interact to respond to incidents according to the following phases. The *preparation phase* includes activities such as maintaining system archives and resource kits with necessary incident-response tools. During the *identification phase*, security practitioners need to determine if an event actually occurred, which may include audit log collection and system disk imaging. If an incident has occurred, the *containment phase*

disables affected systems and restores them. Compromised systems are rebuild and restored from trusted back-ups during the *recovery phase*. Finally, during the *follow-up phase*, all incident-related information is recorded and documented. In our study, we relied on their model to classify our findings within the *preparation* and *identification* phases.

4.2.3 Diagnostic Work during Security Incidents

Intrusion Detection Systems. One of the tools designed to support practitioners during the detection of security incidents is an intrusion detection system (IDS). Goodall et al. (2004b) and Thompson et al. (2006) relied on data from nine and two semi-structured interviews, respectively, to identify the phases of intrusion detection work, and propose a corresponding framework (IDS setup, monitoring, analysis, response). In this framework, the diagnosis of security incidents occurs in the analysis phase, followed by interventions during the response phase. Goodall et al. (2004a;b) suggest that intrusion detection work is challenging due to its highly collaborative nature that drives the need for analysts to coordinate with other stakeholders. Furthermore, this type of work requires high expertise, both technical and organizational. Unfortunately, attaining this degree of expertise is difficult, as much of the necessary knowledge is tacit and may be organization specific. Werlinger et al. (2008c) analyze data from nine interviews to identify security practitioners' perceptions of the the advantages and disadvantages of IDSs'. Werlinger et al. (2008c) also analyze data from participatory observation to identify challenges related to deploying and maintaining an IDS from a usability perspective. The results show that IDS usability is hindered by lack of technical resources and distributed nature of ITSM. In this paper, we extend these findings via an indepth analysis of the diagnostic aspects related to IDS configuration and deployment.

Case Studies of Diagnosis + Recovery. As well as research on diagnosis of security incidents, there several descriptive case studies of real-life examples related to security incidents. We already described on such case study of a security incident presented in Spafford (2003) above. Casey (2005) presents a case study of an intrusion against one organization and stresses the role for collaboration during incident diagnose and containment. Gibson describes a denial of service attack on his company. The diagnosis of the incident included both technical troubleshooting as well as

interaction with various parties, including the ISP. Riden (2006) describes a series of security incidents on a large academic network ranging from defaced web pages to password guessing to worm-related incidents. Interestingly, key factors contributing to the incidents corresponded to ineffective communication and collaboration between the various security professionals within the organization, which led to inconsistent preventative measures and untimely notification of vulnerabilities. Schultz (2007) describes a variety of sources of information that had to be combined in order to diagnose an incident in one organization. Based on this one organization's experience, the author concludes that diagnosis of security incidents involves a number of challenges, including: (1) validity of information (e.g., is the output from a tool such as an IDS an appropriate form of evidence?); (2) reliability of evidence, as attackers may have compromised it; and (3) completeness of evidence (e.g., are all logs present or where some lost?).

Summary. As we described above, related work shows that ITSM is a challenging endeavor that entails a mix of both technical and other skills, such as communication. As far as security incident response is concerned, the only formal studies that exist investigate a small subset of security incident response, namely a specific tool used to detect security incidents (an intrusion detection system). Although some preliminary case studies do exist, they have only involved a single organization, and have not relied on formal evaluation methodologies to collect and analyze their data. Our work presented in this paper fills this gap, as we now describe.

4.3 Methodology

Prior work has shown that little is known about the diagnostic work security practitioners perform within their organizations. This lack of understanding makes it difficult to develop security tools that support effectively the diagnostic work involved in various facets of ITSM, such as for instance security incident response. To fill this gap, we framed our study with the following research questions:

- How do security practitioners perform diagnostic work when responding to security incidents?
- What tools do security practitioners need to perform this type of diagnostic work?
- How can such tools be improved to better support security practitioners during this diagnostic work?

4.3.1 Data Collection

The two sources of data for our study included: (1) 34 semi-structured *in situ* interviews and (2) participatory observation in one academic organization in Canada. The semi-structured interviews were conducted with a total of 36 security pracitioners, who worked for a variety of organizations (11 different organizations in total from 7 sectors). During the *in situ* interview, participants were asked a variety of security-related questions (e.g., ITSM challenges, ITSM tasks and tools, organizational influences, to name a few). Each interview lasted approximately one hour and was subsequently transcribed and sanitized to preserve the participants' anonymity. As is typically the case with semi-structured interviews, not all participants were asked the same questions, and not all discussed topics relevant to our research questions on diagnostic work. Table 4.1 summarizes information on the 13 participants who did discuss diagnostic work and whose data we considered for the analysis presented here. For presentation purposes, we identify our interview participants according to their original interview number (i.e., I1...I34).

The participatory observation was performed by the first author. The observer took part in two activities: development of security policies and deployment of an intrusion detection system (IDS). In this chapter, we limit our analysis to the observations gathered during the installation and configuration of the IDS. The observer spent 15 hours working with two senior security practitioners who are specialists in their areas (servers, networks) and have worked together in the same organization for several years. These two experts are in charge of the technical security projects in their areas, including the installation of an IDS. It should be noted that the observer is also a security specialist with four years of experience in a large telecommunications organization, although has no prior experience working directly with an IDS.

The participatory observation during the deployment of the IDS has corresponded to two key activities: (1) meetings (a total of three hour-long meetings between the two security specialists and the observer); (2) individual work. The participatory observation started with a meeting, followed by 12 hours of individual work by the observer, followed by two meetings. During the individual work, the observer had brief one-on-one interactions with the specialists to discuss specific issues on the configuration of the IDS. Throughout the process, the observer kept detailed notes on the meetings, the interactions with the security specialists, and the IDS deployment.

Table 4.1: For each type of organization, we indicate the number of unique organizations and the total number of participants interviewed. These participants held various positions, including IT Managers (with security tasks), Security Managers, Security Specialists and IT Practitioners (with security tasks).

Organization Type	**Position Type**				**Total**
	IT Manager	**Security Manager**	**Security Specialist**	**IT Practitioner with Security Tasks**	
Academic (3)	I15	I2	I3, I9	I7, I8, I22, I24	8
Financial Services (1)	-	-	I4	-	1
Scientific Services (1)	-	-	-	I12, I13	2
Manufacturing (1)	-	-	I21	-	1
Telecommunications (1)	-	-	I32	-	1
Total	1	1	5	6	13

4.3.2 Data Analysis

We used qualitative description (Sandelowski 2000) to analyze our data, as follows. First, we analyzed the interview transcriptions and our notes from participatory observation to identify excerpts that pertained to diagnostic work across a variety of security tasks. These tasks included: troubleshooting the installation of an IDS, identifying security vulnerabilities in IT systems and responding to security incidents. To identify diagnostic aspects related to security incidents, we used CERT's definition of a security incident: "any real or suspected adverse event in relation to the security of computer systems or computer networks " (http://www.cert.org/). Second, we organized the excerpts into different stories or "memos" (Charmaz 2006) describing how security practitioners perform diagnostic work and the key challenges they face during the diagnostic process.

4.4 Results

As our results will show, the following three factors were responsible for making the diagnostic work before and during incident response challenging:

1. The highly-specialized knowledge needed during incident diagnosis, including (i) knowledge—predominantly in a tacit form—about the organization (e.g., the tasks performed by different stakeholders), and (ii) highly-specialized technical knowledge of the IT systems (e.g., protocols, operating systems, networking).

2. The need to interact and collaborate with a variety of stakeholders during the diagnosis of security incidents.

3. Usability problems with those IT tools that support the incident-response process.

We now present the results; to do so, we rely on the security incident response model (Mitropoulos et al. 2006) presented in Section 4.2, to classify our findings within two phases of incidence response: *preparation* (i.e., preparing for an incident) and *identification* (i.e., determining if a security breach actually occurred). For each phase, we describe security-related activities that required our participants to perform diagnosis.

4.4.1 Preparation Phase

Security practitioners prepare for security incidents in a variety of ways. For instance, preparation can include maintaining resource kits with necessary incident response tools (Mitropoulos et al. 2006). Preparation can also include vulnerability assessment of IT systems, as understanding the vulnerabilities can guide the incident-response process. We now describe our participants' experiences along both of these dimensions of preparation for a security incident.

Troubleshooting the installation of an intrusion detection system

A security tool that is designed to support practitioners during both the detection and analysis of security incidents is an intrusion detection system (IDS). An IDS monitors and detects abnormal behavior in IT systems, such as, for instance IT networks. Once an incident is detected, the IDS forwards alarms to the appropriate IT professionals. However, the installation and configuration of such software can be extremely challenging and require extensive troubleshooting to correctly configure the system, as we describe below. The majority of the findings stem from the participatory observation work, which as we mentioned in section 4.3 involved the observer and several security practitioners working together to install an IDS.

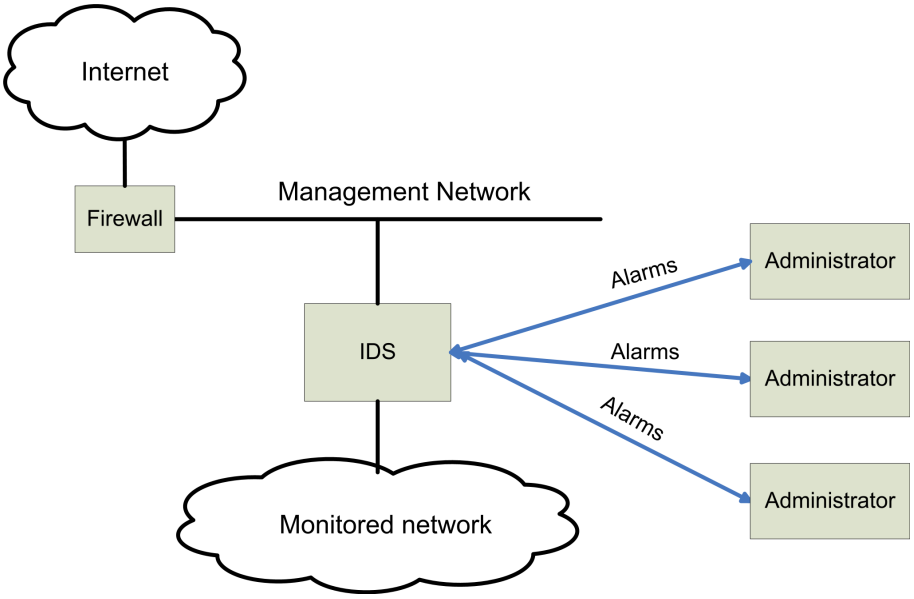

Figure 4.1: Diagram of the network connections for the IDS. The IDS has one connection to the monitored network and a second to the management network.

Installation of an IDS in a Production Network. To clarify the subsequent discussion, we need to provide some technical details with respect to the IDS. The IDS was installed on a server that had two network connections (see Figure 4.1): one to enable the IDS to monitor the networks (*monitoring connection*) and one to allow security practitioners to manage the IDS, e.g., to specify the IT networks the IDS should monitor (*managing connection*).

To deploy the IDS, the responsible security practitioners had to first validate its license. To do so, they decided to install the management port within one of the organization's networks and connect it to the Internet. This network was protected from the Internet by some firewalls. Although the firewalls blocked traffic that should not have access to the organization's internal systems, the IDS only needed to access networks external to the organization, and so the firewalls were not expected to cause difficulty. When the security practitioners tried to validate the IDS license, however, an error message appeared (see Figure 4.2). The content of this message suggested that the traffic from the IDS was being blocked, contradicting the mental model of the network security specialist.

This was the start of a long troubleshooting process to diagnose the source of the problem. Throughout this process, the observer and the specialists had to formulate their own hypotheses

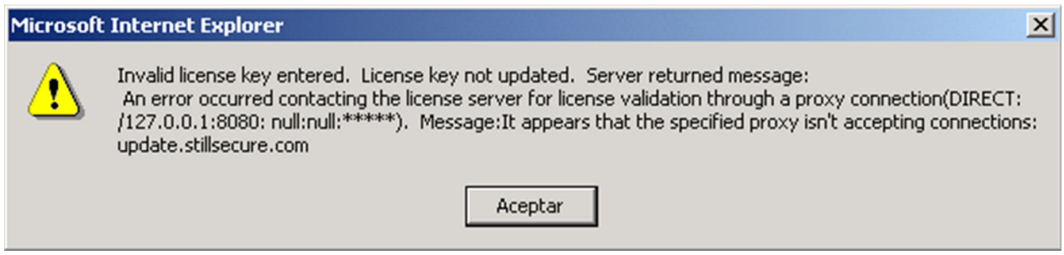

Figure 4.2: Error message when the license of the IDS could not be validated

about the cause of the problem. These hypotheses included: (1) the IDS interface was malfunctioning; this hypothesis was supported by the incomplete and therefore unhelpful nature of the error messages that the IDS showed during its installation and initialization (e.g., "ACPI resource is not an IRQ entry", "smartd failed initialization"); (2) the IDS traffic was successfully leaving the organization, but was blocked by the IDS vendor's server (this hypothesis was deemed unlikely, as it is imperative for vendors to have their servers up and running); (3) despite the practitioners' mental model, the firewall was somehow blocking the traffic from the IDS.

Since the last hypothesis seemed the most probable, the security practitioners investigated it first. To verify that the IDS' traffic was not being blocked by the firewall, the practitioners communicated with the vendor to gain insight into the type of traffic the IDS generated. This communication took place in the form of e-mails, which were highly technical in nature, such as for instance the following: *"The IDS does have internet access, but it is a firewalled network. If the registration traffic conforms to standard TCP/IP connection based traffic, we should have no problem. However, if the registration requires that some UDP traffic be returned from the server, it will undoubtedly fail . . . "*

Given that these communications did not help to resolve the problem, the observer decided to analyze (i.e., *sniff*) the traffic produced by the IDS, to determine if the firewall was configured to block this type of traffic. To do so, he had only one specific tool at his disposal, namely TCPDump. Unfortunately, although he did manage to get this tool running, he could not interpret its output, again highlighting the usability issues with ITSM tools. Specifically, the tool's output corresponded to large amounts of data in plain text, including all the messages from the network to which the IDS was connected. This high volume of data, coupled with a lack of a better interface to visualize all the information, made it difficult for the observer to identify which of the tool's output was

relevant for the task at hand (i.e., determining if the firewall was filtering traffic from the IDS). Complicating the diagnosis was that it had to performed in a *production* network that needed to remain operational. This made it challenging to run tests, as many of them required changes to the firewall that could impact organizational access, and so was something that the security specialists wanted to avoid.

Summary of TroubleShooting Challenges. The security practitioners never isolated the cause of the problem: lack of resources meant that the project was scaled down, and the deployment of the IDS was moved to a smaller, less critical network. Insufficient resources is a commonly-cited ITSM challenge (e.g., Siegel et al. 2006). This story also highlights a number of other challenges specifically related to diagnosing access problems arising when connecting a new system to an organization's IT network. First, in order to use tools designed to monitor organizational networks (such as an IDS), security practitioners must have extensive knowledge of the type of network traffic that is allowed within their organization. A second challenge relates to the fact that it is sometimes necessary to involve external stakeholders (in this case the vendor) in the investigation. Third, tests in production networks are restricted by the fact that these networks need to continue to be operational. Finally, a fourth challenge relates to the usability of security tools; our story above illustrated how tool error messages are often uninformative and misleading, complicating diagnosis.

Detecting System Vulnerabilities

In addition to the installation and configuration of security tools (I2, I3, I4, I9, I32), another aspect of the preparation phase is the detection of IT system vulnerabilities, as these can guide the diagnosis of security incidents. As we describe below, vulnerability analysis is a diagnostic process that requires security practitioners to (1) to corroborate tool output with other data sources; (2) rely heavily on their tacit knowledge about the organizational systems; and (3) collaborate with different IT professionals.

Vulnerability analysis involves the use of specific IT tools called *security scanners* (e.g., Nessus, ISS). Security practitioners used the scanners to determine if the systems were susceptible to known vulnerabilities; the vulnerability list was obtained from public servers maintained by the IT security community[10]. As was the case with tools for monitoring networks (e.g., IDSs), to use

[10]http://www.securiteam.com/securitynews/5AP041FCKE.html

the security scanners effectively, our participants needed to be highly familiar with configuration of their organization's networks, in order to specify the systems to be scanned. Failure to provide accurate information could result in, for example, the tools scanning other organization's networks (I9). These organizations might interpret such unexpected scanning activity as preparation for an attack.

Further complicating the usage of the scanners was lack of accuracy. Participant (I32) described how scanner output needed to be corroborated to (1) discard false positives; and (2) adjust the scanners' interpretation of vulnerability severity. To discard false positives, this participant had to directly access the scanned systems and verify each of the vulnerabilities identified by the scanner, by checking the corresponding processes and applications. If the scanner information was accurate, i.e., a vulnerability, then the participant still needed to confirm the scanner's assessment of the vulnerability's severity, relying on his tacit knowledge of the IT infrastructure to do so. To illustrate, the scanner could report a critical vulnerability, with an accompanying recommendation (e.g., the installation of a security patch), but a security practitioner could assess the severity differently. This occurred, for instance, when the scanner labeled an application as highly vulnerable, but that application was running on a network protected by a firewall. Of course, this does not mean that the vulnerability did not exist, but that the priorities suggested by the scanner had to be adjusted, so that resources could be allocated to mitigate more critical vulnerabilities.

IT system vulnerability diagnosis might involve interaction among security practitioners. For example, a security practitioner found a vulnerability announcement on the Internet - although this practitioner was not responsible for administering the services affected by the vulnerability, he knew who was, and forwarded the information to him (I2). Similarly, a security practitioner used the scanner to identify system vulnerabilities, and then forwarded the scanner-generated report to the responsible administrators (I9).

4.4.2 Identification Phase

Although preparation for security incidents is critical to minimize the damage incurred from an incident, this also heavily depends on security practitioners' ability to effectively detect and investigate an incident. We now discuss our findings pertaining to these aspects (detection & investigation). The results presented here extend the findings of Werlinger et al. (2008b), which

identify nine activities that require security practitioners to interact with other stakeholders; one such activity is security incident response. Here, we extend those results by (1) analyzing security incident response from a broader perspective, rather than focusing only on interactions; and (2) identifying the diagnostic aspects during interactions involved in security incident response (see Figure 4.3).

The diagnostic process during security incident response starts with the *detection* of an anomaly in an organization's IT systems, such as for instance users experiencing slow access to Internet. During this process, our participants performed two types of activities: *monitoring* (A.1 in Figure 4.3) and sending and receiving *notifications* (B.1, C.1 and D.1 in Figure 4.3). Monitoring involves intensive use of IT tools (e.g., IDSs, antivirus), as well as requires tacit knowledge to identify patterns of anomalous activity in the networks. Notification involves extensive collaboration with other stakeholders, who are either directly monitoring systems or indirectly receiving notifications from other stakeholders.

After noticing an anomaly in the IT infrastructure, participants moved to *analysis* of the anomaly. This stage included diagnostic tasks such as: *verification* (A.2 in Figure 4.3), *assessment* (A.3 in Figure 4.3), and *tracking the source of the anomaly* (A.4 and B.2 in Figure 4.3). To perform these tasks, participants required effective (i) communication skills to collaborate with other stakeholders and (ii) analytical skills to generate hypothesis about the causes of the anomaly. When the cause of the anomaly was found, participants moved to containing the incident.

We now describe these various activities and corresponding tools in detail, linking our stories with the diagnostic aspects (A, B, C and D) in Figure 4.3.

Detecting an Incident

To detect security incidents within their organizations, our participants actively *monitored* their organizations' IT systems (Figure 4.3, A.1). Monitoring involved a variety of tools, as well as tacit knowledge about the organizations' IT systems and services. For example, one participant (I3) knew that end-users in his organization typically generate less than 50 e-mails in a given day, and so a higher number of e-mails signaled a potential anomaly.

Examples of the tools security practitioners used to monitor IT systems included antivirus software and intrusion detection systems (IDSs). Antivirus software was used to detect viruses and to generate reports about virus activity in the infrastructure (I3, I4, I12, I24). Intrusion

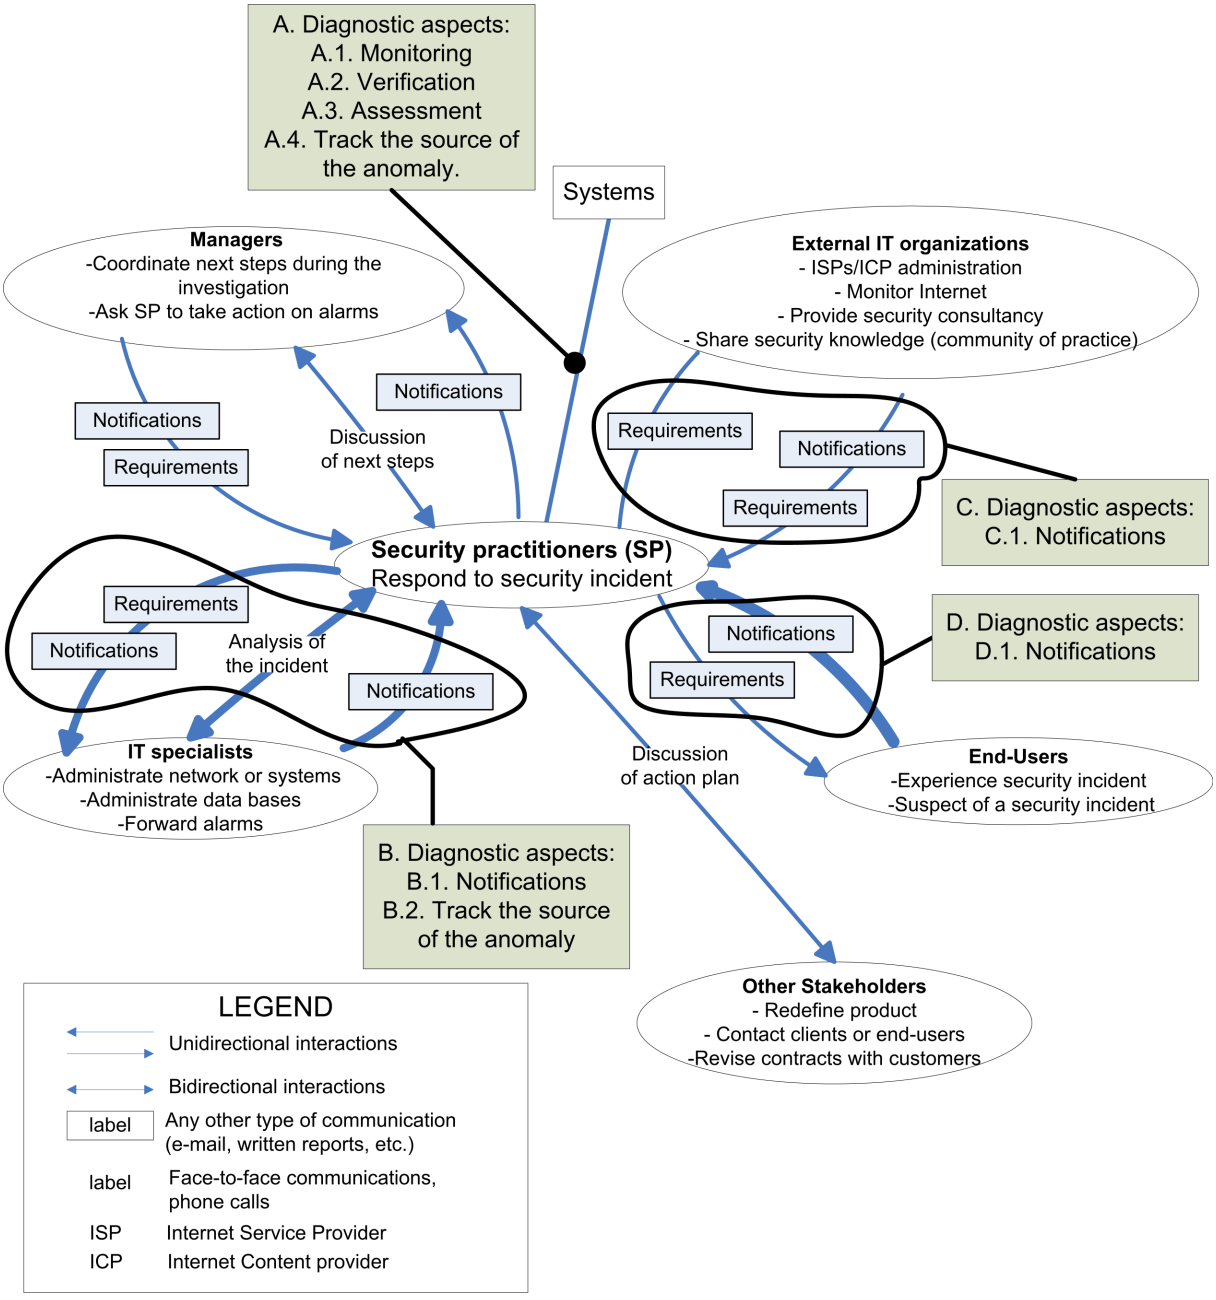

Figure 4.3: Adaptation of the flow communication diagram showed in previous Chapter (Chapter 3) with the collaboration among different stakeholders to respond to a security incident. We now highlight the diagnostic aspects of such collaboration, including the monitoring tasks performed by our participants on the organization's systems. Thicker arrows indicate more frequent collaboration. For simplicity, only collaboration between security practitioners and other stakeholders is shown.

detection systems were used to 'sniff' network traffic and find matches between the traffic and signatures of known attacks.

The usability of monitoring tools in general and IDSs in particular hindered their effective use. For instance, some participants (I4, I9, I12, I24) found it very challenging to use IDSs to generate meaningful reports on monitoring outcomes, largely due to the overwhelming amount of false positives generated by IDSs. To reduce false-positives, an IDS needs to be customized to fit the particular systems of a given organization. However, IDS configuration and customization is a time-consuming and difficult process, and consequently a burden that some of our participants preferred to avoid (I3, I4, I9, I24, I12).

Not all tools our participants mentioned using for detecting anomalies were as complex to use as IDSs, although these tools also sufferred from usability issues. For example, one participant mentioned SmokePing as a tool for identifying when systems were up or down (I13). This tool had two advantages: its output was easy to interpret, and it minimized false positives, because it only showed in red the systems that were unavailable. The tool, however, also had a disadvantage, namely that the alarms it generated did not include any information on the cause of the problem.

As the above examples demonstrate, IT tools typically have cons and pros. In some instances, security practitioners combined tools in unique ways to maximize their utility. For instance, one participant (I12) combined two tools (TCPDump and Ethereal) sequentially to generate and analyze, respectively, the log files he needed. He alternated between the advantages of portability (TCPDump) and good visualization (Ethereal): *"[TCPDump provides] common analysis format . . . it's also a portable format . . . it [Ethereal] shows the SYN and RESET in one colour and then the PUSH commands in another colour. So it is obvious there is content in there."*.

In-house tools. Due to usability issues and budget constraints, our participants sometimes resorted to creating their own tools to detect anomalies in the IT infrastructure (I2, I3, I8, I9, I12, I22, I24). These tools were *scripts*—small programs for the command interpreter of an operating system—that searched for specific patterns, for instance in the networks or log files. According to one participant (I3), scripts relieved the burden of manually analyzing raw log files that were generated by the systems.

In order to create effective scripts, participants needed both technical expertise and knowledge about the IT infrastructure within their organization. For example, a participant (I3), who wrote a script to monitor e-mail traffic, could list at any moment the network addresses of the computers

with suspiciously high number of e-mail communications; this allowed him to selectively monitor some systems more than others. The same participant developed a script to generate only one alarm upon detection of abnormal traffic, to avoid having vast volumes of e-mails associated with the same anomaly. Another participant (I2) explained how in their organization, security practitioners had developed scripts that detected denial of service attacks, and subsequently notified the appropriate administrators, alleviating the burden of a security practitioner having to deal with the notification (see B.1 in Figure 4.3).

Notifications. As was the case with vulnerability assessment, the complexity of IT systems and the lack of resources to monitor *all* systems meant that our participants relied on *notifications* to detect security incidents (see B.1, C.1 and D.1 in Figure 4.3). Our participants received notifications from different stakeholders, including other IT professionals and/or end users. Often, these notifications required interactions among different stakeholders. For example, one participant (I12) described how an external organization (MyNetWatchman) had detected malicious traffic generated from one of the system he administered (see C.1 in Figure 4.3). However, instead of receiving the notification from MyNetWatchman, he received it from another colleague (see B.1 in Figure 4.3) who was notified by MyNetWatchman. This chain of notifications among different security practitioners was also mentioned by another participant who was involved in a response to a phishing attack (I4): *"... we had a person, not even a member of any of our organizations or customers, who emailed our privacy office ... then the privacy office contacted me directly"* (see C.1 in Figure 4.3) Our participants also received notifications about incidents from end-users (see D.1 in Figure 4.3), e.g., in the form of complaints that the Internet access was blocked (I11, I22).

In some instances, monitoring and/or receiving notifications led security practitioners to the detection of anomalies, and their subsequent investigation.

Investigating an Anomaly

Our analysis of the participants' stories showed that the investigation of an anomoly comprised at least three tasks: *verification* (see A.2 in Figure 4.3), *assessment* (see A.3 in Figure 4.3) and *tracking the source of the anomaly* (see A.4 and B.2 in Figure 4.3).

During the *verification* task, security practitioners aimed to confirm, often with alternate data

sources, that a compromise actually occurred. One participant (I3) described this verification purposes: *"I always try and verify by a second or third source. So [I would, again] go back to the Argus [IDS] . . . check the Argus logs and see what's actually happened; . . . then I would . . . go to one of my other logs [say from an MS] Windows box; what have I seen in the logs of the Windows box; was that a real compromise or not."*

If an incident was indeed confirmed, during its *assessment*, security practitioners estimated the incident's magnitude and consequences (I3, I4). One participant (I3) described the assessment process and how it shaped the next steps: *"If it looks like a compromise, I might go through the logs to see what kind of traffic I'm getting from this IP address—everywhere else in the campus; is it scans? is it a successful compromise? So it depends on what I find, depends on what I do."* Another participant, who described a phishing incident (I4), explained how he assessed the attack by checking how many e-mails went out from the organization's e-mail server. His assessment was complicated by the fact that contrary to a typical attack, the server was not inundated with nondeliverable e-mails.

During *tracking the source of the anomaly*, security practitioners aimed to determine the cause of the incident. We next describe stories on how our participants used their tacit knowledge, security training, and collaboration with other practitioners to diagnose the source of an anomaly.

Two participants (I9, I12) used their knowledge about hacking patterns to diagnose the source of an anomaly related to malicious software. One participant (I9) mentioned that diagnosing denial of service attacks was straightforward and could be accomplished by inspecting the volumes of specific network traffic: *"denial of services are easy to spot, cause it's sending millions of the same thing actually over and over and over again, with very little iteration or very little permutation . . . "* Another participant (I12) identified hacking activity by looking for specific type of traffic: *"there is some content here and it looks like IRC [Internet Relay Chat]. So I figure that this is somebody controlling it, the machine . . . [IRC is] very popular with hackers as a control mechanism."*

When the source of an incident was difficult to diagnose, participants found it especially helpful to interact with other specialists, particularly ones who were new to the investigation or had a different background, as they could offer a new perspective. One of the stories we collected exemplified this point: a participant (I13) had to investigate an incident related to the loss of service from the organization's IT systems. He decided to check the systems *in situ*, and asked

for help from another specialist to do so, *"because two eyes are better than one"*. However, the hardware looked normal, and they decided to involve another specialist in the analysis. She thought that the problem was with a small network switch that had not been checked during an earlier inspection; they reset the switch and the network recovered from its failure. In another story we collected, a participant (I11) described needing help from another department's specialist to trace the flow of traffic in a network that was not performing well. As a result of this collaboration, they were able to isolate the device that was slowing outbound traffic: *"We also contacted IT services [to] see if they could see, based on traffic utilization on the network, where it was coming from . . . we finally isolated—hey, it's that new firewall that we just brought up."*

In addition to collaboration, another tactic participants used to identify the cause of an incident involved simulation of the incident. One participant (I13) mentioned how he was collecting information from actual situations where he repeated the conditions of failure: *" So we try to put a proxy in between . . . and then it started crashing . . . [But] as soon as we put in no filtering . . . bad things stop happening . . . "* In another case, a participant (I12) wanted more specific information about the type of malicious traffic that was causing anomalies. He explained how he downloaded the same suspected malicious software to provide such information: *"It's saying . . . downloading a tool from some website. Okay, so I do that, download this tool and run it through the antivirus and it says okay, this is some dial-up . . . "*

Some of the security incidents described above were solved during the analysis process. In other stories told by our participants, an additional step was needed to stop the incident. This step corresponds to the containment phase, according to Mitropoulos et al. (2006). Containment was accomplished in a variety of ways, including: (1) by turning off ports or services in external organizations (the case of a phishing attack, I4) and (2) cleaning up IT systems by reinstalling software (I9).

4.5 Discussion

Our analysis shows that response to security incidents requires intensive diagnostic work. To perform this work effectively, our participants relied on various skills, applied specific strategies, and relied heavily on their tacit knowledge about the IT systems and services within the organization.

Table 4.2: Sequence of tasks to respond to security incidents.

Stage	**Task**	**Skills**	**Security tools**
Preparation	Troubleshooting installation of an IDS	Hypothesis generation Communication	Sniffers
	Mitigating security vulnerabilities	Communication	Scanners
Identification	Monitoring	Pattern recognition	Scripts, IDS
	Receiving notifications	Communication	Incident ticketing system
	Verificating	Hypothesis generation	Scripts
	Assessing	Pattern recognition	Applications to administrate IT systems (e.g., firewall management system)
	Tracking the source of the anomaly	Pattern recognition, hypothesis generation, communication, bricolage	Antivirus
Containment	Shutting down or clean systems	Communication	Applications to administrate IT systems

We summarize these aspects in the next section and then discuss how technology can be improved to better support diagnostic tasks performed by security practitioners.

4.5.1 How Security Practitioners Diagnose Security Incidents

To perform the various tasks during the preparation and identification of security incidents, our participants relied on (1) tacit knowledge about their organizations and IT systems, (2) different security tools, and (3) four key skills: *pattern recognition*, *hypothesis generation*, *communication* and *bricolage* (i.e., dynamic integration of security tools in novel, unanticipated ways (Botta et al. 2007a)).

We should point out that ITSM in general and diagnostic work during ITSM in particular are fairly new fields, and as such, could borrow insights from more mature fields. To illustrate how this could occur, let's focus on the identification phase. During this phase, to isolate the

source of the anomaly, our participants complemented the use of skills with the application of two strategies: *collaboration* and *simulation*. As far as collaboration is concerned, diagnostic work involved dynamic groups of IT specialists to evaluate the different situations and isolate the source of an incident. This strategy of involving various specialists during diagnosis is also employed in so called "High Reliability Organizations" (HROs): organizations that are highly interactive, complex, and tightly coupled[11] (Weick and Sutcliffe 2001). HRO examples are electric and nuclear power plants. When safety incidents occur in these organizations, different teams are dynamically formed depending on the type of incident. Once the safety incident is resolved, these ad-hoc groups are dissolved and do not leave a trace of their existence in the formal structure of their organization. We argue that more research is needed to understand how the diagnosis of security incidents might be improved by adopting strategies used during the investigation of safety incidents in HRO's. For instance, one possible avenue involves investigating how safety incident procedures in HRO's (e.g., in nuclear reactors, as shown in (Park and Jung 2003)) may be applied to the diagnosis of IT security incidents.

4.5.2 Opportunities for improving IT security technology

Security incident response is a multi-faceted activity, where the corresponding diagnosis requires a mix of both strong technical and communication skills. Our participants faced many challenges when diagnosing security-related problems. At least some of these challenges stemmed from insufficient tool support, which in its turn was caused by usability issues, for example, related to tools producing unhelpful or uninformative error messages. Our study identified a number of other aspects of insufficient tool support; we now rely on our analysis to offer suggestions on research directions and guidelines for improving security tools, grounding our discussion in both our participants' experiences and related work.

Task Complexity: A key challenge our participants mentioned pertained to security tools that monitored IT systems and generated alarms upon detection of anomalous events. These monitoring tools generated overwhelming numbers of false positives—i.e., alarms that corresponded to innocuous events—which placed a high burden on security practitioners who had to investigate the alarms. Our analysis suggests that task complexity influences tool reliability (see Figure 4.4, left).

[11] That is, changes in one part of the organization imply changes in other parts.

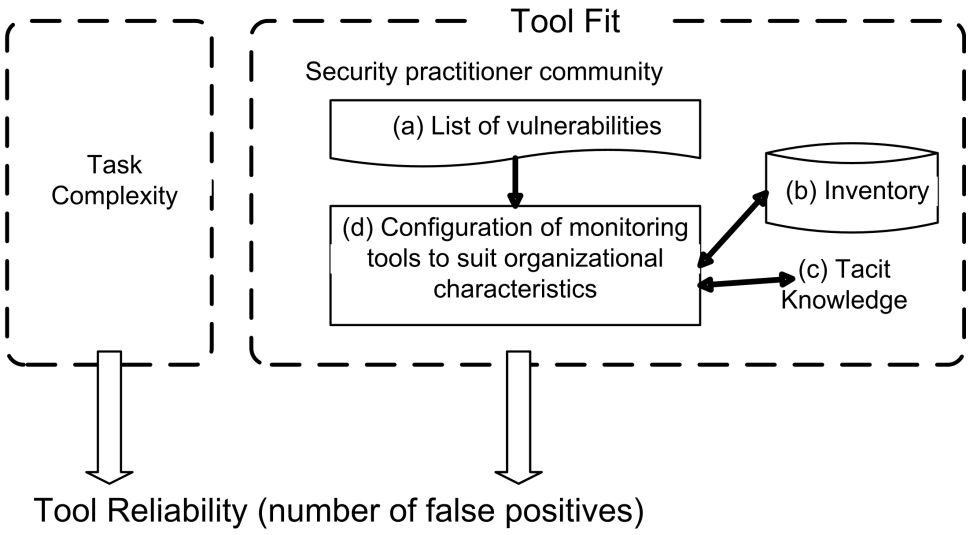

Figure 4.4: Forces influencing tool reliability (false positive rate)

Furthermore, our results suggest there is a tradeoff between the complexity of the task supported by a tool and the tool's reliability: the more complex the task, the less reliable the tool's output for that task. For example, IDS tools perform a variety of complex tasks; these tools generated many more false positives and subsequently required more intervention from practitioners to verify the output than SmokePing, a simple tool that only checked system availability. On the other hand, SmokePing's simplicity was not without disadvantages: its very basic functionality meant that it did not provide information about incidents unrelated to the availability of systems, e.g., attacks to guess the users' passwords.

The above discussion highlights that the tradeoff between task complexity and tool reliability is a dimension that must be taken into account during tool evaluation. In particular, we believe that more research is needed to understand the advantages and disadvantages of security tools designed to perform complex tasks, as compared to tools that are intended for simple tasks. A second dimension that needs to be taken into account when evaluating tools is support for tool integration, as we describe shortly. First, however, we present the second factor influencing monitoring tool reliability.

Customization to Ensure Tool Fit: Security practitioners' ability to configure a monitoring tool to a given organization's characteristics directly impacts the number of false positives produced by that tool (see Figure 4.4, right). Recall that to configure monitoring tools, practitioners relied

on generic lists of attacks and vulnerabilities (Figure 4.4(a)), which are maintained by security practitioners around the world and are available on public servers [12]. Although the lists are a good starting point, highlighting the collaborative nature of ITSM, they do correspond to huge quantities of generic data, making the customization task difficult for security practitioners.

The above point illustrates that no matter how advanced a security tool is, diagnostic work in the context of ITSM still requires customization of the tool to the specific reality of a given organization. The customization requires access to a complete inventory of an organization's IT systems (Figure 4.4(b)). Such an inventory is very costly to create and maintain, given the challenges of ITSM (Gagné et al. 2008). For instance, the dynamic nature of the IT environment means that systems are constantly being upgraded and/or replaced, subsequently requiring practitioners to update the system inventory. Customization also requires intensive use of both tacit knowledge that is usually not shared among practitioners and not explicitly documented (Figure 4.4(c)).

In general, to improve the efficiency of diagnostic ITSM work, more research is needed to investigate how the process of customizing generic list of vulnerabilities could be optimized. One option is to rely on Artificial Intelligence techniques, and so have tools automatically adapt a generic vulnerability list to a given organization's characteristics (e.g., as is done in so-called *anomaly-based* IDSs). Another complimentary option is to design support for transforming security practitioners' tacit knowledge used during tool configuration into explicit knowledge that can be shared with other security practitioners as Gagné et al. (2008) suggest.

Need for tailorable tools: Our participants had to develop their own tools, e.g., scripts, to perform specific tasks related to the diagnosis of security incidents. This fact illustrates how difficult it is to develop standard security tools for the diagnosis of security incidents that fit every organization's needs. Botta et al. (2007b) propose that security tools have to support *tailorability*, so that practitioners can customize tools via their own scripts. Our analysis showed that practitioners require this feature for diagnosing security incidents. In addition to increasing the usability of a tool, support for customization via scripts has a second benefit: they capture practitioners' tacit knowledge.

Depending on the diagnostic work performed, practitioners used scripts either as stand alone tools or in combination with other IT tools via bricolage. While Botta et al. (2007a) show that

[12] https://lists.sourceforge.net/lists/listinfo/snort–sigs

ITSM work in general involves bricolage, our results demonstrate that this skill is also practiced during diagnosis of security incident responses. Note that bricolage is a special instance of *vendor-designed* tool integration[13]. How tools should support the practice of tool integration, however, is an open question. As far as we are aware, there have been no studies of how tool integration in general or bricolage in particular impact tool usability, meaning that novel evaluation methodologies may be needed. Further complicating the issue of tool integration and its impact on usability is that integration must be considered in conjunction with task complexity, since the latter also impacts tool usability. To illustrate, it may turn out that bricolage support is beneficial across the board, from simple to complex tasks; alternatively, it may be the case that bricolage places high cognitive load on security practitioners, making bricolage only beneficial for complex tasks. This rich understanding of the ways in which tools are used during diagnostic work when responding to security incidents can support the specification of the complex scenarios in which these security tools should be evaluated (Redish 2007).

Correlate information from different systems to verify the incidents: To diagnose security incidents, our participants had to correlate different sources of information. To do so, they not only had to understand how various IT systems were related, but also needed security tools that were able to process and relate information from these different sources. To satisfy this need, security tools need to process information from a variety of sources with different formats and structure. For instance, a tool developed by Cisco[14] can integrate with different security devices to correlate information and generate consolidated reports (Cisco Info Center for security monitoring).

In the same vein, security tools that correlate data need to process very large volumes of this data, which in turn must be reified in a meaningful way. Unfortunately, our participants found it difficult to generate on-line reports, needed during diagnosis of security conditions within their IT systems. To deal with this limitation, one option involves abstracting the tasks of data synthesis and visualization, away from the standard security tools towards specialized tools that only focus on these tasks. Abstraction has the advantage of providing a separation of functionality, i.e., tools that collect raw data vs. tools that process that data. This in turn enables practitioners with the flexibility to plug in a variety of devices into the specialized reification tools.

[13]It refers to IT security tools developed to integrate the features of multiple security tools (e.g., a firewall with IDS/IPS features

[14]A major vendor of network devices and monitoring tools.

Multi-faceted simulation support: As we described above, diagnostic work during security incidents involves security practitioners performing *simulations* to verify or investigate an anomaly. Complicating simulation work is that in some instances, it needs to be performed in production systems that needed to remain operational. To address this issue, Fisler et al. (2005) describes an approach for a specific type of simulation involving access control rules. Along a similar vein, Chiasson et al. (2007) propose that any security-system changes should be easily reversible; this guideline ensures that any simulation-introduced problem in a production system is easily reversed. Our results show that diagnostic work during security incident response requires practitioners to perform simulations in distributed systems administered by various practitioners, and so requires collaboration. Since collaboration complicates the simulation process, we propose that tool support for simulation need to address not only the technical factors, but also include functionality that supports collaboration between different IT practitioners as they track the simulations and evaluate their consequences.

4.6 Conclusion

Our qualitative analysis shows the importance of diagnostic work during the preparation and identification stages of the response to security incidents. We have identified the different tasks, skills, strategies, and tools that security practitioners use to detect and classify anomalies during those stages.

The diagnosis of security incidents required high levels of collaboration among our participants and other stakeholders. Participants used different technologies to support their tasks, developing their own tools when they did not have the required security tools for specific tasks.

In our discussion we offer several recommendations to improve the support of security tools to diagnostic work during the response to security incidents. These recommendations include criteria to evaluate usability of security tools in complex scenarios. Further research is needed to expand and refine our understanding on how technology can best provide the required support to security practitioners when they respond to security incidents.

Bibliography

K. Beznosov and O. Beznosova. On the imbalance of the security problem space and its expected consequences. *Information Management & Computer Security*, 15(5):420–431(12), 2007.

D. Botta, R. Werlinger, A. Gagné, K. Beznosov, L. Iverson, S. Fels, and B. Fisher. Studying IT security professionals: Research design and lessons learned. position paper for the CHI Workshop on Security User Studies: Methodologies and Best Practices, April 28 2007a.

D. Botta, R. Werlinger, A. Gagné, K. Beznosov, L. Iverson, S. Fels, and B. Fisher. Towards understanding IT security professionals and their tools. In *SOUPS*, pages 100–111, Pittsburgh, Pennsylvania, July 18-20 2007b.

E. Casey. Error uncertainty and loss in digital evidence. *International Journal of Digital Evidence*, 1(2), 2002.

E. Casey. Case study: Network intrusion investigation – lessons in forensic preparation. *Digital Investigation*, 2(4):254–260, 2005.

K. Charmaz. *Constructing Grounded Theory*. SAGE publications, 2006.

S. Chebrolua, A. Abraham, and J. Thomas. Feature deduction and ensemble design of intrusion detection systems. *Computers and Security*, 24(4):295–307, 2005.

S. Chiasson, P. C. van Oorschot, and R. Biddle. Even experts deserve usable security: Design guidelines for security management systems. SOUPS USM Workshop, July 2007.

Cisco Info Center for security monitoring. Cisco info center web page. http://www.cisco.com/en/US/products/sw/netmgtsw/ps5477/products_data_sheet09186a0080187172.html, April 2008.

K. Fisler, S. Krishnamurthi, L. A. Meyerovich, and M. C. Tschantz. Verification and change-impact analysis of access-control policies. In *ICSE '05: Proceedings of the 27th international conference on Software engineering*, pages 196–205, New York, NY, USA, 2005. ACM. ISBN 1-59593-963-2. doi: http://doi.acm.org/10.1145/1062455.1062502.

A. Gagné, K. Muldner, and K. Beznosov. Identifying differences between security and other IT professionals: a qualitative analysis. In *Proc. of Human Aspects of Information Security and Assurance (HAISA)*, Plymouth, England, July 2008.

S. Gibson. The strange tale of the denial of service attacks on grc.com. URL `http://whitepapers.silicon.com/0,39024759,60026382p-39000404q,00.htm`.

J. R. Goodall, W. G. Lutters, and A. Komlodi. I know my network: Collaboration and expertise in intrusion detection. In *CSCW*, volume 6390, November 2004a.

J. R. Goodall, W. G. Lutters, and A. Komlodi. The work of intrusion detection: Rethinking the role of security analysts. In *Proceedings of the Americas Conference on Information Systems (AMCIS)*, pages 1421–1427, 2004b.

E. M. Haber and J. Bailey. Design Guidelines for System Administration: Tools Developed through Ethnographic Field Studies. In *Proc. of 2007 symposium on Computer human interaction for the management of information technology (CHIMIT), 9 pages.* ACM, 2007.

K. Hawkey, D. Botta, R. Werlinger, K. Muldner, A. Gagne, and K. Beznosov. Human, organizational, and technological factors of it security. In *CHI'08 extended abstract on Human factors in computing systems*, pages 3639–3644, 2008.

D. Heckerman, J. S. Breese, and K. Rommelse. Decision-theoretic troubleshooting. *Communications of the ACM*, 38(3):49–57, 1995.

http://www.cert.org/. Cert: Computer emergency response team.

E. Kandogan and E. M. Haber. Security administration tools and practices. In L. F. Cranor and S. Garfinkel, editors, *Security and Usability: Designing Secure Systems that People Can Use*, chapter 18, pages 357–378. O'Reilly Media, Inc., 2005.

G. Killcrece, K.-P. Kossakowski, R. Ruefle, and M. Zajicek. Incident management. Technical report, U.S. Department of Homeland Security, 2005.

S. Kraemer and P. Carayon. Human errors and violations in computer and information security: The viewpoint of network administrators and security specialists. *Applied Ergonomics*, 38:143–154, 2007.

S. Mitropoulos, D. Patsos, and C. Douligeris. On incident handling and response: A state of the art approach. *Computers and Security*, 25(5):351–370, 2006.

M. Park and W. Jung. The requisite characteristics for diagnosis procedures based on the empirical findings of the operators' behavior under emergency situations. *Reliability Engineering & System Safety*, 81(2):197–213, 2003.

M. Rao, H. Yang, and H. Yang. Integrated distributed intelligent system architecture for incidents monitoring and diagnosis. *Computers in Industry*, 37(2):143 – 151, 1998.

R. H. Rayford B. Vaughn Jr. and K. Fox. An empirical study of industrial security-engineering practices. *The Journal of Systems and Software*, 61:225–232, 2001.

J. Redish. Expanding usability testing to evaluate complex systems. *Journal of Usability Studies*, 2(3):102–111, 2007.

J. Riden. Responding to security incidents on a large academic network, 2006.

M. Sandelowski. Whatever happened to qualitative description? *Research in Nursing & Health*, 23(4):334–340, 2000.

E. E. Schultz. Computer forensics challenges in responding to incidents in real life setting. *Computer Fraud & Security*, 12:12–16, 2007.

M. A. Shayman, Emmanuel, and Fernandez-Gaucherand. Fault management in communication networks: test scheduling with a risk-sensitive criterion and precedence constraints. In *he 39th IEEE Conference on Decision and Control*, volume 2, pages 1864 – 1869, 2000.

D. A. Siegel, B. Reid, and S. M. Dray. IT Security: Protecting Organizations In Spite of Themselves. *Interactions*, pages 20–27, 2006.

E. H. Spafford. A failure to learn from the past. In *Annual Computer Security Applications Conference (ACSAC)*, Las Vegas, Nevada, 2003. URL `http://www.acsac.org/2003/papers/classic-spafford.pdf`.

P. Stephenson. The application of formal methods to root cause analysis of digital incidents. *International Journal of Digital Evidence*, 3(1), 2004.

R. S. Thompson, E. Rantanen, and W. Yurcik. Network intrusion detection cognitive task analysis: Textual and visual tool usage and recommendations. In *Proceedings of the Human Factors and Ergonomics Society Annual Meeting (HFES)*, pages 669–673, 2006.

K. Weick and K. Sutcliffe. *Managing the unexpected: assuring high performance in an age of complexity*. Jossey-Bass, 2001.

R. Werlinger and D. Botta. Detecting, analyzing and responding to security incidents: A qualitative analysis. presented at the SOUPS Workshop on Usable IT Security Management (USM), July 18 2007.

R. Werlinger, K. Hawkey, and K. Beznosov. Security practitioners in context: their activities and interactions. In *CHI '08 extended abstracts on Human factors in computing systems*, pages 3789–3794, 2008a.

R. Werlinger, K. Hawkey, and K. Beznosov. Security practitioners in context: Their activities and interactions with other stakeholders within organizations. *submitted to International Journal of Human Computer Studies*, 2008b.

R. Werlinger, K. Hawkey, and K. Beznosov. The challenges of using an intrusion detection system: Is it worth the effort? In *Proc. of ACM Symposium on Usable Privacy and Security (SOUPS) (12 pages, to appear)*, 2008c.

R. Werlinger, K. Hawkey, and K. Beznosov. Human, Organizational and Technological Challenges of Implementing IT Security in Organizations. In *to appear in HAISA'08: Human Aspects of Information Security and Assurance (13 pages)*, July 2008d.

Chapter 5

The Challenges of Using an Intrusion Detection System: Is It Worth the Effort?[15]

5.1 Introduction

Security incident response is one key aspect of maintaining organizational security (Killcrece et al. 2005). A critical task during security incident response is detecting that an incident has occurred. Detection may occur through reports from end-users and other stakeholders in the organization, through detection analysis performed on an ad-hoc basis (e.g., hand-crafted scripts that detect anomalies in server logs), or it may be accomplished by using an intrusion detection system (IDS). In general, an IDS monitors and records events in a computer system, performs analysis to determine if the events are security incidents, alerts security practitioners of potential threats, and produces event reports (Scarfone and Mell 2007). If the IDS also includes mechanisms to block detected intrusions from entering the organizational infrastructure, it is referred to as an intrusion prevention system (IPS). Security practitioners interact with the IDS through a console, which may be used to either perform administrative functions, such as configuration of sensors, and/or to support event monitoring and analysis. Some of the most popular IDSs include Snort, OSSEC HIDS, BASE, Sguil, and Bro.

Intrusion detection (ID) is a challenging endeavor, requiring security practitioners to have a high level of security expertise and knowledge of their systems and organization (Scarfone and Mell

[15] A version of this chapter has been accepted for publication. R. Werlinger, K. Hawkey, K. Muldner, P. Jaferian, and K. Beznosov. The challenges of using an intrusion detection system: Is it worth the effort? *In Proc. of ACM Symposium on Usable Privacy and Security (SOUPS) (12 pages, to appear),* 2008

2007; Goodall et al. 2004b). Traditionally, ID research has focused on technological solutions for improving the accuracy of IDSs (e.g., Chebrolua et al. (2005); Hwang et al. (2007)). Although this is still an active area of research, recent work has also recognized the need to address the human side of ID work (e.g., Goodall et al. (2004b); Komlod et al. (2005); Thompson et al. (2007)). This recognition is driven by the fact that while IDSs automate some aspects of the process, human intervention is very much still required. For instance, although an IDS automatically recognizes potential security threats and generates alerts, the alerts need to be analyzed by a human expert, since many are false positives (as many as 99 percent according to Julisch and Darcier (2002)).

From a usability perspective, much of the research has focused on providing visualizations during the monitoring and analysis phases (e.g., Malécot et al. (2006)), with some claiming these phases to be the most cognitively challenging (Thompson et al. 2006). However, the initial deployment and configuration of the IDS can also be a barrier to its use. The first author has experienced this first-hand while working as a security consultant at a large telecommunications company from 2002 to 2006. This organization's security team wanted to employ an IDS to improve the organization's security, but had two main concerns about incorporating such a system: (1) Were they going to be able to maintain it? (to ease this burden they had the option of outsourcing the network monitoring, but did not want to disclose the log files), and (2) Were they going to learn valuable information from the reports (e.g., were there attacks on their systems that needed to be addressed)? Despite assistance from an external company with the initial configuration of the IDS, the security team was unable to customize it and tune it appropriately for the network they were monitoring within a reasonable time frame.

In this chapter we report on the challenges of using an IDS, with a particular focus on the initial stages of deployment (i.e., decision making, installation, and configuration). Our motivation for this research arose from the first author's prior industry experience as described above. We also noted that other practitioners had similar difficulties with IDSs through our research conducted for the HOT Admin project, which is investigating the human, organizational, and technological factors that influence security management within organizations (see Hawkey et al. (2008a) for an overview and Botta et al. (2007), Gagné et al. (2008), Hawkey et al. (2008b), Werlinger et al. (2008b), Werlinger et al. (2008a) for results to date).

Our findings are based on analysis of nine of the HOT Admin interviews that we conducted

with security practitioners, as well as participatory observation in a large academic organization that is in the process of installing an IDS. This rich set of data has allowed us to identify and describe some of the challenges that impact the ability of security practitioners to successfully deploy and maintain an IDS within an organization. These challenges include deciding on the purpose of the IDS, integrating the IDS in the network, working within a distributed environment, and balancing the tradeoff between limiting the number of false positives to achieve usability of the system, while keeping false negatives at a minimum. While some of these challenges may not have obvious solutions, it is important that security practitioners, researchers, and tool developers are aware of the complexity of the full process of deploying an IDS.

Our work has two key contributions. First, we add to the community's understanding of the factors influencing IDS usability. In particular, while prior work has focused on the challenges associated with the monitoring and analysis phases of IDS work, suggesting that these phases are the most cognitively demanding, our results show that the deployment phase also involves challenges, and that these may be significant enough to hinder the very adoption of an IDS within an organization. Second, we provide recommendations and guidelines for mitigating some of the challenges we identify through better tool support.

The remainder of this chapter is organized as follows. We begin by presenting the related work in Section 5.2 and our methodology in Section 5.3. In Section 5.4, we describe the IDS tool used during participatory observation, and then present our results related to IDS usability in Section 5.5. We discuss our findings in Section 5.6 before presenting conclusions and future work.

5.2 Related work

Before devising support for the human analysts who work with IDSs, it is important to have an understanding of what is involved with ID work, including its phases, challenges, and cognitive demands.

5.2.1 IDS Phases

Based on analysis from nine semi-structured interviews conducted with professionals who were responsible for ID work in their organizations, Goodall et al. (2004b) propose that ID can be

broken into three distinct phases. The *monitoring* phase corresponds to the ongoing surveillance of an IDS, including sifting through the various alerts it generates. When monitoring reveals a potential security event, the *analysis* phase is initiated, which involves in-depth examination to determine if the alert is actually a security event. If a security event is confirmed, the *response* phase involves intervention and reporting of the event. Note that missing from this task analysis is IDS configuration. Thompson et al. (2006) refine the Goodall analysis with data from two semi-structured interviews. They propose that, in addition to the above-mentioned three phases, ID work also involves a *pre-processing* phase. This phase occurs before the monitoring phase and corresponds to the actual IDS setup (e.g., configuring alerts, and/or generating filters for the alerts).

5.2.2 IDS Usability Challenges

Goodall et al. (2004a;b) propose that ID work is challenging due to expertise demands and its highly collaborative nature. ID requires significant expertise, both technical and organizational. Professionals need to have knowledge of their own unique network environment, since what is classified as a security event in one network may not be considered one in another network (Goodall et al. 2004b). Attaining this degree of expertise is difficult, as much of the necessary knowledge is tacit and may be organization specific. Further complicating ID work is its collaborative nature that drives the need for practitioners to coordinate with other organizational stakeholders (Goodall et al. 2004a).

To obtain a fine-grained view of the challenges, Thompson et al. (2006) use data from two interviews to perform a cognitive analysis of the three ID phases (pre-processing, monitoring, analysis, response). In general, they propose that all ID phases are challenging, but that the monitoring and analysis phases are the most cognitively demanding for practitioners. This high cognitive load derives from the need to integrate various sources of information in these two phases, including background knowledge on the network and the user base and information generated by the various tools involved in ID, such as the output of an IDS and network logs.

5.2.3 Support and Evaluation

IDSs generate large volumes of data, which subsequently security practitioners need to inspect. If this information is presented in textual form, this places a high burden on the practitioners to make sense of the data. An alternative is to devise effective visual representation of the data to alleviate some of the cognitive burden and so facilitate the task of identifying security events (e.g., Komlod et al. (2005); Malécot et al. (2006)). For instance, the Intrusion Detection toolkit (IDtk) (Komlod et al. 2005) generates glyph-based visualizations of network data, which may be raw packets or generated by an existing IDS, such as SNORT. IDtk uses color, spatial coordinates and glyph size to create the data visualizations, which aim to support the monitoring, analysis, and response phases of ID work.

To date, although studies have investigated the process of ID, very few usability evaluations of IDSs exist. One exception is Thompson et al. (2007), who compare how different interface types (text vs. visual) support the monitoring and analysis phases through a laboratory experiment with 16 participants (2 professional ID analysts, 14 graduate students). The findings suggest that each interface type has its respective strengths and weaknesses. For instance, a text interface provides access to fine-grained detail, affording flexible interactions and customizations; but it burdens the user with high quantities of data and the need to know the command syntax. A visual interface, on the other hand, can provide an overview of the data, which facilitates the detection of attacks; but it fails to provide fine grained detail and so some attacks may be missed.

5.3 Methodology

Prior work has shown the need for better security tools to detect malicious activity in networks and systems. These studies also propose the need for more usable tools that work in real contexts (Kandogan and Haber 2005; Botta et al. 2007). To date, however, there has been little focus on the pre-processing steps of intrusion detection. We designed our study to fill this gap, as well as to further the understanding of IDS usability and utility, particularly as the IDS is installed and configured in an organization. Consequently, our research questions were:

- What do security practitioners expect from an IDS?

- What are the difficulties that security practitioners face when installing and configuring an IDS?
- How can the usability of an IDS be improved?

We used a qualitative approach to answer these questions, relying on empirical data from security practitioners who have experience with IDSs in real situations. Below we detail our data sources and analysis techniques.

5.3.1 Data collection

We collected data from two different sources. First, we conducted semi-structured interviews with security practitioners. Second, we used participatory observation, an ethnographic method (Fetterman 1998), to both observe and work with two senior security specialists who wanted to implement an IDS in their organization. These two sources of data allowed us to triangulate our findings; the descriptions from interviewees about the usability of IDSs were complemented by the richer data from the participatory observation.

Semi-structured Interviews

For the HOT Admin project, we have conducted to date 34 *in situ* semi-structured interviews with 36 participants from various organizations (16 different organizations from 11 sectors, e.g., post-secondary educational, scientific services, financial services, consulting, manufacturing, insurance, and non-profit). All participants played a role in upholding security in their organizations; their positions ranged from IT manager to general IT staff to security staff. Each interview lasted approximately one hour. The interviews were subsequently transcribed and sanitized to preserve the participants' anonymity. During the interview, subjects were asked a variety of questions pertaining to the nature of security (e.g., challenges, tasks, tools, organizational influences, security culture, etc). Note that due to the diversity of participants' positions as well as the nature of semi-structured interviews, not all participants performed and/or discussed ID work. Information pertaining to the nine participants that did discuss ID is shown in Table 5.1.

Table 5.1: Participant Information (Semi-Structured Interviews)

ID	# Sector	Position
P2	Academic	Security Manager
P3	Financial Services	General Security
P4	Academic	General Security
P9	Academic	General Security
P12	Scientific Services	General IT
P15	Academic	General IT
P20	Academic	IT Manager
P23	Consultant	General Security
P24	Academic	General Security

Participatory Observation

The participatory observation was performed by the first author in one large, distributed post-secondary organization. It should be noted that the observer is a security specialist with four years of experience as a security consultant in a large telecommunications organization, although with no prior experience working directly with an IDS. The observer spent 15 hours working with two senior security practitioners who have worked together in the organization for several years, and are specialists in their areas, namely *servers* and *networks*. These two experts are in charge of the technical security projects in their areas, including the installation of an IDS. This project is currently at the stage where the IDS is connected to a production network, and is ready for tuning.

The participatory observation has consisted of two main activities: meetings and individual work. There have been a total of three, hour-long meetings between the two security specialists and the observer. The work on the IDS started with one meeting, followed by 12 hours of individual work, and continued with two further meetings. During the individual work, the observer had brief one-on-one interactions with the specialists to discuss specific issues related to IDS configuration. Throughout the process, the observer kept detailed notes of the meetings and interactions with the security specialists and of the IDS implementation.

5.3.2 Data analysis

The data from the interviews and participatory observation were analyzed using qualitative description (Sandelowski 2000) with constant comparison and inductive analysis. We first identified

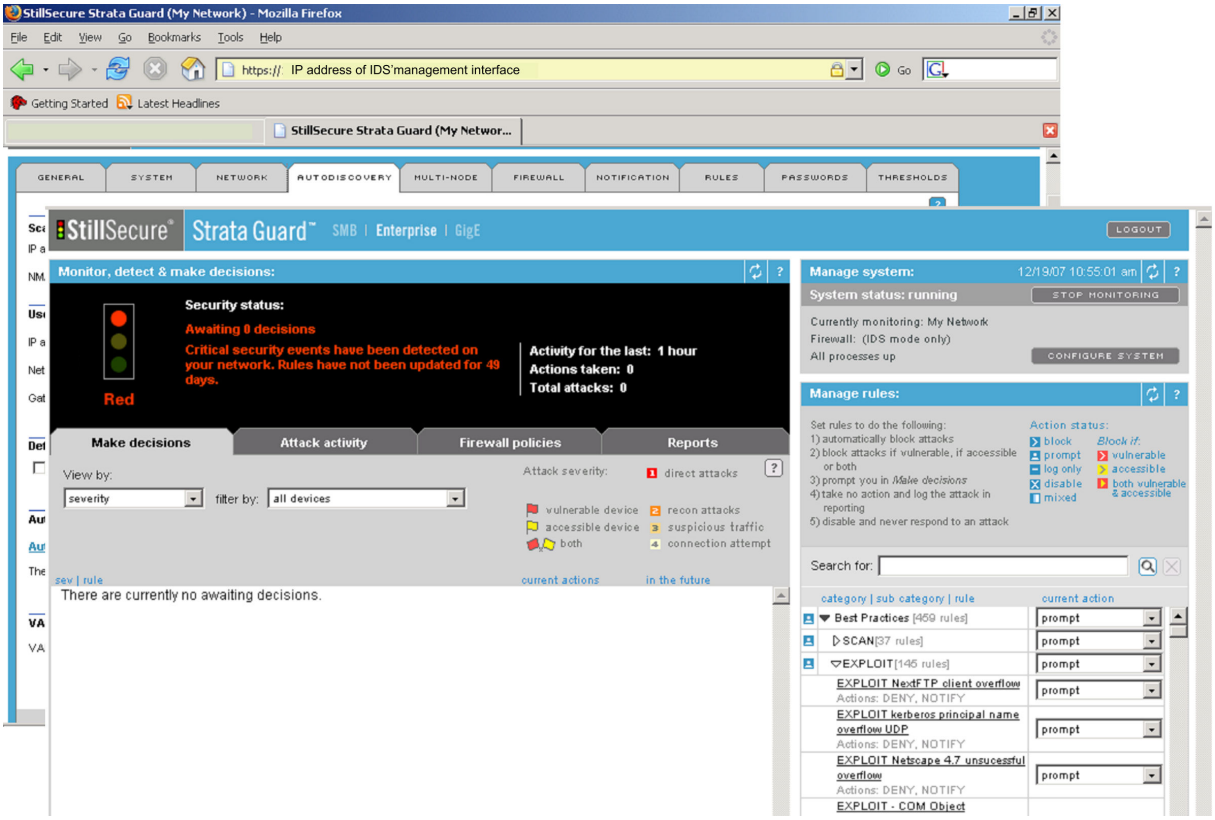

Figure 5.1: System configuration options of the IDS in the back. On top, Configuration options of the IDS's rules (bottom right) and status of alarms.

instances in the interviews when participants described IDSs in the context of the activities they had to perform. We next contrasted these descriptions with our analysis on the participatory observation notes. Results were then organized by the challenges that the participants faced when deploying and maintaining an IDS system.

5.4 Anatomy of an IDS

An IDS is a tool that detects abnormal behavior in systems. For the work reported in this paper, we are interested in those IDSs that monitor and detect attack patterns in network traffic. Such systems are commonly referred to as network IDSs. To monitor the networks, the IDS uses *sensors*, which are probes that are connected in the networks and that passively sniff the network traffic. To detect attacks, the IDS includes an *engine*, which typically performs detection via rules encoding attack patterns or signatures. Finally, the IDS provides mechanisms for administration, such as command line or graphical user interfaces.

5.4.1 The Deployed IDS

The IDS being deployed during the participatory observation was Strata Guard for small to medium businesses, version 4.5 (StillSecure 2008); the choice of system was based on a managerial financial decision. The IDS was acquired approximately five years ago. Since then, the organization has paid a maintenance to StillSecure (the vendor) for updates and general questions about the IDS's operation. Although current Strata Guard IDSs offer the option of being deployed with dedicated hardware (i.e., as an appliance), the version purchased by the organization came as a software package for general purpose servers. Another option, which was not available for the IDS version purchased, is IDS/IPS capability: (i) when operating as an IPS, the tool monitors and potentially intercepts network traffic (i.e., reacts instantaneously to attacks); (ii) when operating as an IDS, the tool monitors traffic and reporting alarms for off-line action.

The Strata Guard software included the following components: Linux operating system, PostgreSQL database, and a graphical user interface (GUI) as shown in figure 5.1, which enables the configuration of some but not all IDS settings (the IDS also includes a command line interface (CLI) that does enable practitioners to configure all aspects of the system). The support service provided by StillSecure gave immediate access to new attack signatures and also the option of opening trouble tickets in case of problems with the system.

During the participatory observation, the Strata Guard system was deployed as an IDS using software installed on an IBM server (Intel Xeon processor, 1 Giga RAM, 30 Giga Hard Drive). The server included two Ethernet ports: one used to monitor traffic, and one to manage the IDS server. To validate the IDS license and download rules to detect new attacks, the IDS needed to have access to the vendor's server (StillSecure) via the Internet, which was realized through its management Ethernet port.

5.5 Investigating IDS Usability

IDS usability evaluations should not be confined to the study of their graphical user interfaces: our data show that security practitioners also emphasize other factors (e.g., organizational) that influence the adoption of an IDS within an organization. We first highlight the main issues that security practitioners had to face during the integration of an IDS in a real network, as uncovered

during the participatory observation. We then present the advantages and disadvantages of IDSs that participants described during the semi-structured interviews.

5.5.1 Issues Deploying an IDS

From discussions with the security specialists during the participatory observation, we learned that the initial objective for the IDS was to monitor traffic on the organization's internal networks. Alarms from the IDS were to be forwarded to the administrators of the appropriate networks. About two years prior to the participatory observation, the IDS had been installed by the security specialists in one particular network domain. However, it soon crashed, possibly due to memory space issues (the IDS GUI did not provide practitioners with functionality to manage the IDS's use of the hard-disk partitions), and/or from additional traffic from a newly-added wireless network. The former hypothesis related to memory issues was based on the fact that the default memory partition size was not large enough to accommodate the logs produced by the IDS; when a partition became full, it seemed the IDS started to overwrite other system partitions not dedicated to the IDS. The security specialists did not have the time to confirm this hypothesis and analyze the exact cause of the system failure, so they decided to start again from scratch and install the IDS in another network. This re-installation was delayed for several months due to high workload and other priorities.

We next describe the main issues the security practitioners addressed and the decisions they made during the current IDS installation, which are distilled from the participatory observer's notes (see Appendix B for details). The issues include not only technical ones, but also human and organizational, providing a rich perspective on the challenges related to installing IDSs. As such, our findings may be useful for researchers and practitioners designing support for IDSs; they may also serve to guide the development of scenarios for evaluating IDSs in real contexts (Redish 2007).

Deciding on the Purpose of an IDS

The target organization's main goal behind the adoption of the IDS was to complement the existing security controls (e.g., firewalls). The security specialists believed that the IDS would make monitoring of the organizational networks more efficient than other alternatives such as having to

manually detect attacks via analysis of the firewall log files, using an IPS, or using an anomoly-based IDS. Manual analysis of firewall logs was deemed too complicated, time consuming, and had no guarantee of obtaining the consolidated attack reports the specialists needed. Automatically blocking traffic through an IPS was ruled out as it would have gone against the open culture fostered by the organization's academic nature. The specialists believed that an anomaly-based IDS would be less effective for their organization, as this organization involves a variety of security protocols and services, with highly irregular network traffic.

Monitoring malicious traffic was not the only purpose that security specialists had in mind for the IDS. They believed that the IDS could provide important statistics about the security of the network, and the security controls they had implemented in the network's boundary. Information about the number of attacks that actually crossed the organization's defenses could give the specialists not only a sense of the security of the internal systems, but would also provide support for proposing new security investments.

The purpose of the IDS was a critical factor influencing details of its deployment and use. For example, to test the security of the network's boundary, it would have been necessary to have a least two probes for monitoring the network, or two different IDSs located before and after the firewalls (see figure 5.2). However, the specialists did not know how to integrate the information from the two points, since it was not clear if the IDS provided functionality for doing so.

Given the limited resources available, the specialists decided to simplify the IDS installation as much as possible, and to install the IDS in the internal network only. We now describe their experience in doing so.

Constraints related to Integrating an IDS in the Network

Despite the fact that the security specialists had tried to simplify the deployment of the IDS by limiting its purpose, the IDS integration proved to be a challenging task, due to a number of organizational constraints. For example, to connect to the IDS, the specialists needed to have available ports (at least two in the case of the IDS used during participatory observation). In addition, they preferred to use the port mirroring feature of the switch connected to the IDS (see figure 5.2) to mirror traffic to the IDS, as this option provided the flexibility to select the traffic that they wanted monitored. These requirements became constraints for our participants, who

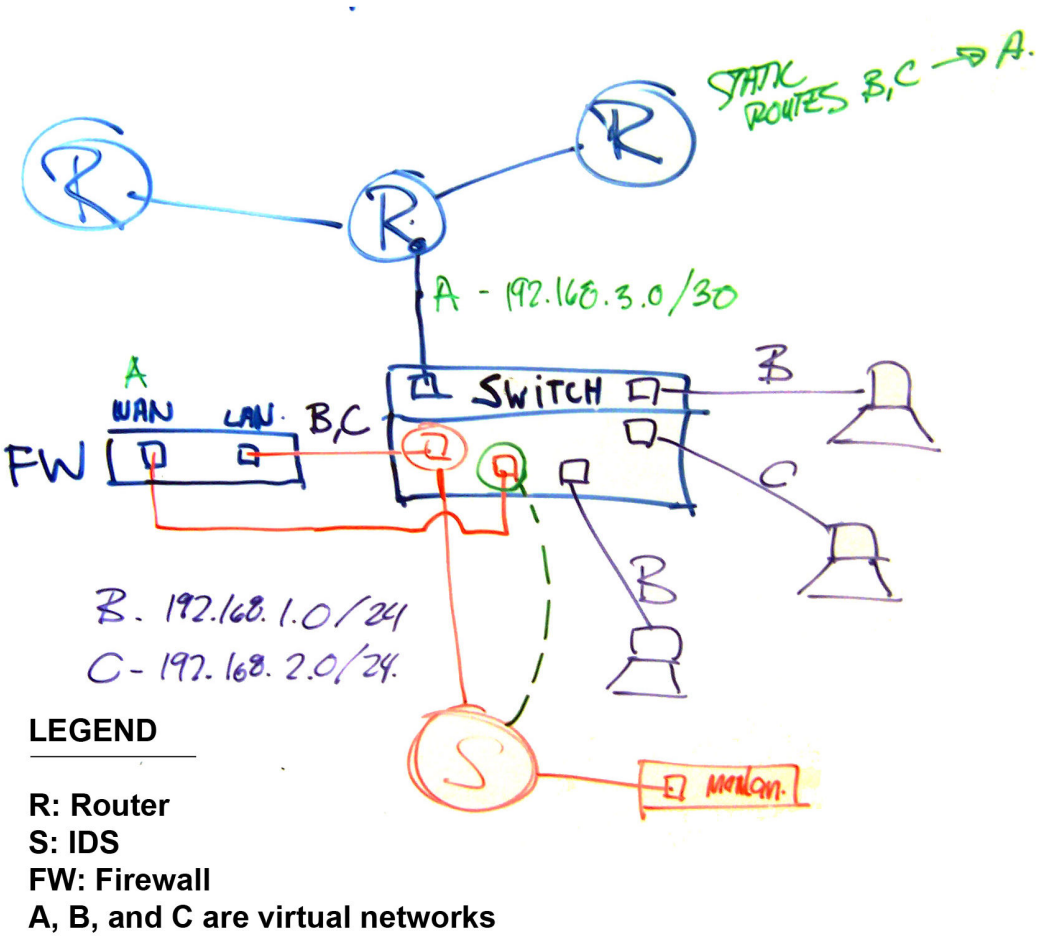

Figure 5.2: Network diagram used during one discussion about the installation of the IDS. The IDS has a connection to the management network and another to the port of the switch that transports internal traffic from the firewall. To compare configuration of the firewalls, it would be necessary to include another connection to the external traffic (dashed line).

could not find the necessary technical resources to connect the IDS in the critical network they wanted to monitor. Consequently, they decided to install the IDS into a less critical network; this decision was also influenced by other factors such as the distribution of IT responsibilities in the organization, as we explain in section 5.5.1.

Using A GUI for the Initial Configuration

Once the practitioners integrated the IDS into the network, the next step involved the installation of the IDS software. This required minimal intervention from the observer, who had to specify only the network settings and two passwords (one for the system and one for the internal IDS database). The GUI integrated with the IDS was intended to alleviate the burden of using a command line interface to administrate the IDS components (e.g., database, security engine) and to provide an easier method of tuning the rules. Specifically, the Strata Guard GUI provides an option (quick tune) to tune the system without the need of going rule by rule and considering the operating systems actually being monitored.

Although the participatory observer has not yet started the IDS tuning process, the initial configuration tasks have revealed some of the shortcomings of the IDS's GUI. For example, the GUI does not allow the user to specify the hard-disk partitions assigned to the filesystem. This configuration option is important to the specialists, given that the pre-defined file space for the logs was too small when the IDS was used in the past. To manage log storage, an additional tool would be necessary. Similarly, the IDS does not provide support for configuring the IDS's security settings. Furthermore, the GUI does not allow users to configure the server's firewall rules, and so this task has to be done via the CLI, a task made difficult by the fact that the rules are non-intuitive and difficult to understand.

In general, although the GUI provided some support for configuring and maintaining the IDS (e.g., disable rules, take action on the alarms), the support was not adequate, given that the IDS was intended to work in a complex environment, influenced by the characteristics of the organization where it was going to be installed. The next section describes some of the key organizational factors influencing the deployment of the IDS.

Working Within a Distributed Environment

The observed organization was highly distributed in terms of IT administration, with various administrators in charge of different interconnected network domains. For these administrators, security usually was not the main priority. These two factors (distribution, security a low priority) triggered specific requirements that had to realized in order to integrate the IDS in the organization. For example, the monitored traffic flowed through various systems that were administrated by different practitioners. Notifications of the alarms the IDS detected in that traffic needed to be sent to the administrators of those systems, who should also be allowed to configure the IDS. Our participants hoped that the IDS would allow different levels of access depending on system characteristics, i.e., operating systems, IP addresses, specific network protocols. However, the deployed IDS did not provide such granularity to define access accounts.

Another issue related to distributed environment is the additional overhead it brings to the IDS project, which the security specialists wanted to minimize. The installation of the IDS in critical networks would have required the intervention of other specialists who administrated different sub-domains of those critical networks. These other specialists were not aware of the project from the beginning and might not have security as a first priority. This factor made our participants decide to discard the installation of the IDS in the critical networks. This decision resulted in a compromise, as the data may have been more interesting from the security point of view if these networks were included. This tradeoff between usability and utility is also discussed in the next section.

Balancing the Tradeoff between Usability and Utility

The security specialists required an IDS that was not only easy to use, but also gave relevant information about the security of the organization's systems. Consequently, the ideal situation would have been to install the IDS in the most critical network domain of the organization to generate meaningful reports about the security level of the networks, with a minimal use of resources. However, this did not occur; as discussed, organizational factors like distribution of IT responsibilities affected the decision to not involve critical networks due to the corresponding overhead of involving multiple administrators.

Another tradeoff between usability and utility was related to how the complexity of IDS configuration varied as a function of the network domains being considered for its installation. Specifically, the specialists could not tell how much more demanding it would be to install the IDS in a large network domain as compared to installing the IDS in a small network domain. This factor also affected the decision of where to install the IDS, as they believed that it would be much easier to install the IDS in a small network domain. However, it seemed that the only way to know how the complexity varied was to complete the full installation process on each of the candidate networks.

Another aspect that security specialists knew required a balance between usability and utility was related with the alarms the IDS generated. They knew that more false positives would require more time from them to investigate the alarms, thereby lowering the usability of the IDS. On the other hand, less false positives would imply less rules running in the IDS and, therefore, potentially more false negatives. Unfortunately, until the tuning process is complete and the IDS is in production, the actual tradeoffs between false positives and negatives will not be known.

5.5.2 Advantages and Disadvantages of IDSs

The results from the participatory observation have highlighted that there are more than just technical factors to consider when installing an IDS in an organization. In this section, we present our analysis of the interviews with various security practitioners, focusing on perceived advantages and disadvantages that IDS afford. As was the case with the results above, our findings span technical, human, and organizational dimensions.

As one of the participants from our field study stated, an IDS is *"one of the most controversial [tools]- some really love it, but some really hate it"* (P24). This controversy is likely rooted in the fact that IDSs have both strengths and weaknesses, and the tradeoff between the two is not always clear, as we discuss below.

Perceived Advantages

Our participants mentioned four key advantages of IDSs, including (1) problem identification, (2) monitoring with privacy, (3) decreased time pressure for maintenance, and 4) reduction of uncertainty.

The first perceived advantage is that an IDS can be a powerful tool to help identify problems

(P4, P24). For instance, P24 stated that the IDS provided *"useful information about what kind of activities are outside a firewall and I want to have something inside the firewall too; to give me some idea whether something managed to go through"*. In identifying problems, an IDS *"makes good business value"* (P4).

Secondly, while security practitioners need to monitor their networks, they also need to maintain privacy of the organizational stakeholders. IDSs can support both of these goals. For example, one participant expressed how Argus Argus Web Page did so: *"Argus is a tremendous tool, it allows us to monitor activity and still respect privacy...because we're not looking at the data portions of the packets, on the header portions"* (P3).

Thirdly, security practitioners are notoriously overworked and juggle a variety of tasks Botta et al. (2007). This sometimes means that they do not have the resources to attend to critical security tasks, such as ensuring that patching of systems happens in a timely manner. As a consequence, the systems become vulnerable and may even be compromised, something that occurred in one participant's organization. According to this participant, an IDS could help with this issue: *"we don't have to run around, for example tomorrow's... patch Tuesday. If we had this intrusion prevention we could patch quarterly. I don't have to run around and neither does anyone else"* (P14).

Finally, one issue that complicates security practitioners' work is related to the inherent uncertainty of their tasks. In particular, our participants mentioned that they are never certain as to the correctness of their activities (P3). An IDS could provide some assurances that everything is in order, e.g., *"...I am going to be considering keeping a closer eye on traffic both in and out, probably with an IDS, so that if there is something weird or not right going in and coming out, what have you, I can at least be alerted to it"* (P20).

Perceived Disadvantages

Despite the fact that an IDS affords advantages, some of our participants were hesitant as to its overall utility, which in turn discouraged them from adopting an IDS in their organization. The disadvantages that the participants mentioned included (1) the expense, (2) the degree of work and time required, (3) the unreliability of the IDS, and (4) the lack of clear utility.

The first disadvantage is that an IDS can be an expensive endeavor: *"so you can easily spend*

a quarter million dollars on an IDS and have 3 people running it" (P4). This is exacerbated by the fact that security is often not a priority, and IDSs fall outside of the mainstream tools, i.e., *"[we do not have a commercial IDS because] we're tight budget-wise and security doesn't get a lot of budget outside of the main stuff, like anti-virus and firewall, and traffic shaper and stuff"* (P3).

Secondly, several of our participants stressed that IDSs are also costly as they require a lot of work and time resources (P3, P4, P9, P24, P12). This demand for resources happens both in the pre-processing IDS set-up phase and the monitoring and analysis phases. As far as configuration is concerned, tuning the IDS can be an arduous undertaking that requires both time and expertise: *"tools like Snort, they're great tools, but they require a lot of customization to get it down to something that understands your environment, so you have to turn alarms on and off based on what you're looking for, what's normal, what's not normal. When I first ran Snort in our environment I was getting thousands of flags a day"* (P9). A key issue with fine-tuning an IDS is to reduce the number of false positives (P4, P9, P24, P12), which occur when customization is not done properly. For example, one participant stated *"when I did run Snort in the past, which is looking for pattern matches on incoming traffic, it just had a ridiculous number of false positives"* (P12). Of course, fine tuning also means not blocking legitimate traffic (P3). Unfortunately, it is very difficult to determine how well an IDS is set up (P23). In the monitoring and analysis phases, lack of time was again an issue: *"I don't monitor that as much as I should be because of lack of resources, because it takes too much time... and then investigate the risks on [the IDS]"* (P3).

Thirdly, our participants sometimes found IDS software to be unreliable, which resulted in lost time and potentially important data, e.g., *"it's quite buggy and sometimes it would fill up all the log files so some partitions were filled up because of the humongous amount of logs ...it would just clog it up and you have to reinstall and then you can really kind of clean up the archive logs and stuff like that. It is just a nightmare"* (P24). Another participant mentioned that some IDSs sometimes dropped packets when they became overloaded (P2). This lack of reliability and potential for interfering with regular network traffic was a negative factor in participants' perceptions of the utility of an IDS.

Finally, although IDSs require many resources, their utility is not always clear. It is hard to see improvement in the security processes, *"you don't really notice any improvement"* (P4). Another consequence of the resources required to maintain an IDS is that often, they simply sit idle (*"we*

do have an intrusion prevention system in place but we haven't been using that effectively at all. It just kind of sits there and runs away" P15).

5.6 Discussion

Our findings suggest that the usability of an IDS is not solely determined by the usability of its GUI. We now discuss some of the associated human, organizational and technical challenges practitioners encounter when deploying an IDS, focusing on: (1) considerations before deploying the IDS; (2) the configuration and validation of the IDS; and (3) its on-going usage. Where appropriate, we provide suggestions for addressing the challenges, which are based on three sources: participatory observation, interviews, and guidelines from the literature. While some of these challenges may not have obvious solutions, it is important that security practitioners, tool developers, and researchers are cognizant of the complexity of this process.

5.6.1 Considerations before deploying an IDS

There are number of challenges that impact an organization's decision to use an IDS. First, our interview analysis revealed that IDSs have not gained the same popularity as other *de facto* security tools, such as firewalls. This makes it more challenging for security practitioners to obtain management buy-in. This challenge could be alleviated with concrete data demonstrating an IDS's utility, however, obtaining the data is difficult for two key reasons. First, in order to obtain the data, the IDS needs to installed and configured within an organization, as generic reports may not reflect a given organization's characteristics. Second, once an IDS is installed and configured, the data needs to be transformed to a form readable by various stakeholders, including managers. To alleviate the latter challenge, an IDS should include reporting functionalities that tailor the information according to a user's specific needs. Furthermore, it should provide the ability to compare the outputs of different IDSs or IDS probes. This functionality would allow security practitioners to compare the state of security before and after the implementation of the IDS (a general version of this guideline is suggested in Nohlberg and Backstrom (2007)).

Second, the decision to use an IDS impacts many stakeholders within the organization. These stakeholders need to be involved in the process, to maximize both the stakeholder buy-in as

well as the benefits of installing such a tool. However, doing so comes with a cost due to the overhead needed to manage the involved parties. Consequently, organizations may opt to reduce this overhead, even though this reduces the IDS utility (as was the case for the organization involved in participatory observation). Third, IDS configuration and use requires extensive resources from security practitioners, who typically have other competing priorities. Fourth, our participatory observation revealed that the installation of an IDS requires the participation of security specialists with knowledge and experience not only in network protocols and systems, but also about the organization itself. The observed security specialists had detailed knowledge of the organization, the networks that provided critical services for clients, and even clients' usage patterns.

The last three challenges derive from lack of security budget, tight schedules and security as a low priority in organizations (Werlinger et al. 2008b). To alleviate these challenges, one of our participants proposed that organizations planning to install an IDS should formalize the process via a dedicated project that includes allocation of resources and the responsibilities of the stakeholders involved: *"So we have internally a project approach...- it's going to have some people allocated to it and a certain amount of capital budget. Well then we write it up in a project and it goes through a project approval process through our senior management team."*(P15). Two other participants suggested allocating some dedicated and uninterrupted time for the IDS (P24, P9). To address budget issues, one participant proposed the use of open source tools (P19), an approach suggested by McGann and Sicker (2005). Such tools can afford benefits (Raymond 1998), such as better internal engines (P19, P25); however, our participants believed that these tools suffer from weaker reporting capabilities (P19, P22, P25) and less management buy in (P19), as compared to commercial tools.

5.6.2 Configuring and Validating an IDS

Once an organization makes the decision to use an IDS, the IDS needs to be installed and configured. Our participatory observation revealed a number of challenges related to these steps that we discuss below, along with guidelines to address them.

Collaboratively Evaluating Tradeoffs

One of the main challenges described by our participants during the IDS configuration process was the need for both broad and deep knowledge of services and organizational goals. Without this knowledge, it is difficult for practitioners to weigh the tradeoffs between increased ease of monitoring through a reduction in the number of false positives and the subsequent reduced IDS utility, due to increased false negative.

To obtain this knowledge, the installation of an IDS in the network requires collaboration with different experts in the organization. Our participatory observation showed cooperation between at least two experienced security specialists from the network and server areas respectively.

The Configuration Hurdle

Hill (2006) states that the big hurdle for most users of security tools is not the user interface, but rather acquiring and installing the software. For the security specialists we observed, a factor complicating the IDS installation was uncertainty: they found it very difficult to predict the degree of effort that would be required to configure the IDS in a particular network. In the end, they found it necessary to go through the full installation process to determine the costs and benefits of the different configuration options according to the utility of the events the IDS detected and reported. This characteristic implies that an IDS might be classified as an "all or nothing" security tool, which makes its adoption and use in the organization difficult. This contrasts with other security tools that do not require intensive use of resources in their configuration to assess their benefits. For example, a security scanner can work with its default configurations and still generate useful reports on system vulnerabilities.

Since the configuration of an IDS is the breaking point for many potential users, IDS designers should aim to minimize the resources required to install and configure these tools. The Strata Guard system used during participatory observation provided several features in this direction, such as automatic discovery of the network's devices and a quick tuning option. However, its GUI did not allow the configuration of all the options required to optimize IDS usage (e.g., memory partitions). Furthermore, error messages the IDS generated during the installation were not helpful. Based on these observations and prior work, the following three guidelines aim to improve the usability

of IDSs. First, IDSs should provide facilities for quick configuration, which can be realized, for instance, by grouping related parameter values (Haber and Bailey 2007). Second, IDSs should provide meaningful help during the configuration process or ongoing usage (McGann and Sicker 2005). Third, IDSs should provide documentation on the configuration process (Haber and Bailey 2007).

Determining an Appropriate Test Bed

A challenge our participants encountered during the installation and configuration process was determining an appropriate test bed environment for the IDS. In general, an IDS must be installed in a real environment to have a sense of its benefits; however, inserting the IDS into a production system might be difficult when there are other stakeholders involved who do not see the benefit of altering the networks.

To deal with the complexity of validating IDS configuration, one participant suggested first testing the IDS in a smaller network than the target one, so as to reduce the amount of traffic security practitioners has to contend with when testing: *"we have to redeploy it to a smaller network ... because it used to be on huge networks [and] we had tons and tons of traffic and tons and tons of ... alerts ... [it was] just too much"* (P24). This participant found that testing on a smaller network *"worked quite well"*, as it provided some useful information on network activities. What P24 suggested is a practice called "planning and rehearsal", as advocated in Barrett et al. (2004; 2005).

If an IDS is installed in a rehearsal environment, the tuning will fit that network, but the tuned system may not fit the target environment. This issue highlights the complexity associated with IDS usage. More research is needed to better understand the trade-offs between smaller rehearsal environments to test an IDS, and the configuration impact of moving them to more complex networks that often transport the critical traffic in the organizations.

5.6.3 Ongoing Usage

After an IDS is installed and configured, challenges remain that impact its ongoing usage.

Monitoring an IDS

As discussed above, improving both the back-end of the IDS as well as the visualization of pertinent information for the practitioners monitoring the IDS alerts are active areas of research. In this vein, one of our participants explicitly discussed the need for improved recognition of anomalous network behavior via an IDS that had *"a bit of smarts"*, one that could watch and recognize trends over time (P3). This participant also described how without this ability, an IDS requires more human attention, as it generates alerts for innocuous network traffic that falls outside of the average throughout the year (e.g., in an academic institution before the term starts, there is very heavy traffic coming from web registration). Related work also provides some suggestions to improve monitoring. First, echoing the above-mentioned participant, Thompson et al. (2007) suggest IDSs should provide automatic detection of malicious traffic behavior, realized for instance via pattern recognition techniques. Second, IDSs should provide facilities for practitioners to fine-tune thresholds for generating alarms as well as facilities for suppressing alarms selectively (Haber and Bailey 2007).

A tool that fits the distributed nature of information security management

During our participatory observation, we found that different security practitioners needed to access the output of the IDS, but that doing so was complicated by the fact that these individuals were distributed across the organization. To address this challenge, related work has suggested that an administration tool should provide a shared view of the system state to its users (Haber and Bailey 2007; Barrett et al. 2004; 2005). Furthermore, Barrett et al. (2005) suggest that tools with a shared view should provide proper authentication and authorizations, to ensure access is granted only to appropriate stakeholders. We recommend extending this concept by having the IDS tailor the view according to the needs of a given stakeholder.

Similarly, to facilitate monitoring and alerting, Haber and Bailey (2007) suggest that monitoring tools should provide alarm generation with a configurable destination. This feature enables an IDS to send its alarms through different channels (email, SMS, etc.) to different stakeholders distributed across the organization. In addition, McGann and Sicker (2005) suggests that providing reports in hypertext format would ease the distribution of reports to security practitioners across

the organization. Beyond just providing the option of sending alarms to different stakeholders, we recommend that an IDS also provide features supporting on-line collaboration among these stakeholders. The IDS used during participatory observation could be configured to generate alarms using different communication channels (e.g., e-mail, SNMP), but it did not provide support for real-time collaboration (e.g., to discuss an alarm).

Reporting

We found reporting to be an important feature of an IDS. Reporting can demonstrate the economic value of the tool (not supported by the version of Strata Guard IDS in the participatory observation). It can also ease the burden of monitoring. For example, one participant described first deploying Snort to monitor the network. However, due to weaknesses of its reporting engine, his organization opted to acquire a commercial solution with better reporting features. The IDS should generate reports that help practitioners investigate the alarms. Furthermore, the IDS can help practitioners prioritize their tasks, by assigning priorities to alarms, or assigning each alarm to a practitioner for further investigation (Werlinger et al. 2008b).

More flexible reporting has been recommended for security tools in general (Botta et al. 2007). Flexibility can be afforded along a number of dimensions. As mentioned above, reports should be tailored according to the needs of the specific user reading them (e.g., manager, practitioner). Other options that may increase the utility and usability of reports include supporting a hypertext format (McGann and Sicker 2005) and using dynamic filters to help practitioners analyze large reports easily (Furnell and Bolakis 2004).

5.7 Conclusion

Intrusion detection systems are complex and provide many challenges for security practitioners. Prior IDS research has focused largely on improving the accuracy of these systems and on providing support to practitioners during the ongoing task of monitoring alerts and analyzing potential security incidents. One area that has received little attention is the pre-processing phase of IDS, but the installation and the initial configuration of an IDS can be so challenging that they can serve as a barrier to use. In this paper we have provided an investigation of these challenges through semi-

structured interviews and participatory observation of one such deployment. Our analysis provided insights into the expectations that security practitioners have for an IDS, identified the difficulties they face when installing and configuring an IDS, and provided the following recommendations for improving the usability of ID systems:

- Show economical benefit of the ID system and security controls by comparing the traffic from different points of the network.
- Provide facilities for quick configuration.
- Provide meaningful help during the configuration process or ongoing usage.
- Provide documentation on the configuration process.
- Include pattern recognition techniques.
- Distribution of alarms to different stakeholders using different criteria.
- Include shared views for different users.
- Provide flexible reporting.

One limitation of our work is that only 9 participants from the semi-structured interviews specifically discussed intrusion detection. Furthermore, two thirds of them came from academic organizations, as did those involved in the participatory observations. Although we argue that many of the issues around the deployment of IDS are organization independent, additional data from different organizational types would strengthen our results. Consequently, one aspect of our future work is to confirm and generalize the findings presented here. Additionally, we will begin to apply our findings towards the design of improved user interfaces for intrusion detection systems, focusing our attentions on relieving the burden on security practitioners that is inherent in configuring and maintaining an IDS. Until improvements are made across all phases of ID, it is clear that many security practitioners and organizations will continue to decide that the challenges of using an IDS will not be worth the effort required.

Bibliography

Argus Web Page. Argus intrusion detection and prevention. http://www.qosient.com/argus/, February 2007.

R. Barrett, M. Prabaker, and L. Takayama. Field Studies of Computer System Sdministrators: Analysis of System Management Tools and Practices. In *Proc. of the Conference on Computer Supported Collaborative Work*, pages 388–395, 2004.

R. Barrett, P. P. Maglio, E. Kandogan, and J. Bailey. Usable autonomic computing systems: The system administratorsṕerspective. *Advanced Engineering Informatics*, 19(3):213–221, 2005. URL `http://www.almaden.ibm.com/u/pmaglio/pubs/icac2004.pdf`.

D. Botta, R. Werlinger, A. Gagné, K. Beznosov, L. Iverson, S. Fels, and B. Fisher. Towards understanding IT security professionals and their tools. In *SOUPS*, pages 100–111, Pittsburgh, Pennsylvania, July 18-20 2007.

S. Chebrolua, A. Abraham, and J. Thomas. Feature deduction and ensemble design of intrusion detection systems. *Computers and Security*, 24(4):295–307, 2005.

D. M. Fetterman. *Ethnography: Step by Step*. Sage Publications Inc., 1998. ISBN 0761913858.

S. Furnell and S. Bolakis. Helping us to help ourselves assessing administratorsúse of security analysis tools. *Network Security*, 2:7–12, February 2004.

A. Gagné, K. Muldner, and K. Beznosov. Identifying differences between security and other IT professionals: a qualitative analysis. In *Proc. of Human Aspects of Information Security and Assurance (HAISA)*, Plymouth, England, July 2008.

J. R. Goodall, W. G. Lutters, and A. Komlodi. I know my network: Collaboration and expertise in intrusion detection. In *CSCW*, volume 6390, November 2004a.

J. R. Goodall, W. G. Lutters, and A. Komlodi. The work of intrusion detection: Rethinking the role of security analysts. In *Proceedings of the Americas Conference on Information Systems (AMCIS)*, pages 1421–1427, 2004b.

E. M. Haber and J. Bailey. Design Guidelines for System Administration: Tools Developed through Ethnographic Field Studies. In *Proc. of 2007 symposium on Computer human interaction for the management of information technology (CHIMIT), 9 pages*. ACM, 2007.

K. Hawkey, D. Botta, R. Werlinger, K. Muldner, A. Gagne, and K. Beznosov. Human, organizational, and technological factors of it security. In *CHI'08 extended abstract on Human factors in computing systems*, pages 3639–3644, 2008a.

K. Hawkey, K. Muldner, and K. Beznosov. Searching for the Right Fit: Balancing IT Security Model Trade-offs. *Special Issue on Useful Computer Security, IEEE Internet Computing*, 12(3): 22–30, 2008b.

A. Hill. Shortcuts, Habits, and Sand Castles. In *SOUPS '06: Proceedings of the second symposium on Usable privacy and security*, Pittsburgh, Pennsylvania, 2006. Invited talk.

K. Hwang, M. Cai, Y. Chen, and M. Qin. Hybrid intrusion detection with weighted signature generation over anomalous internet episodes. *IEEE Transactions on Dependable and Secure Computing*, 4(1):41–55, 2007.

K. Julisch and M. Darcier. Mining intrusion detection alarms for actionable knowledge. In *Proceedings of the 8th ACM International Conference on Knowledge Discovery and Data Mining*, pages 366–375, 2002.

E. Kandogan and E. M. Haber. Security administration tools and practices. In L. F. Cranor and S. Garfinkel, editors, *Security and Usability: Designing Secure Systems that People Can Use*, chapter 18, pages 357–378. O'Reilly Media, Inc., 2005.

G. Killcrece, K.-P. Kossakowski, R. Ruefle, and M. Zajicek. Incident management. Technical report, U.S. Department of Homeland Security, 2005.

A. Komlod, P. Rheingans, U. Ayachit, J. Goodall, and A. Joshi. A user-centered look at glyph-based security visualization. In *Proceedings of the IEEE Workshops on Visualization for Computer Security (VizSEC)*, pages 21–28, 2005.

E. L. Malécot, M. Kohara, Y. Hori, and K. Sakurai. Interactively combining 2d and 3d visualization for network traffic monitoring. In *Proceedings of the 3rd international workshop on Visualization for computer security (VizSEC)*, pages 123–127, 2006.

S. McGann and D. C. Sicker. An analysis of security threats and tools in sip-based voip systems. In *In 2nd Workshop on Securing Voice over IP*, June 2005. URL `http://www.colorado.edu/policylab/Papers/Univ_Colorado_VoIP_Vulner.pdf`.

M. Nohlberg and J. Backstrom. User-centred security applied to the development of a management information system. *Information Management & Computer Security*, 15(5):372–381, 2007. ISSN 0968-5227. doi: 10.1108/09685220710831116.

E. S. Raymond. The cathedral and the bazaar. *First Monday*, 3(3), 1998. URL `http://firstmonday.org/issues/issue3_3/raymond/index.html`.

J. Redish. Expanding usability testing to evaluate complex systems. *Journal of Usability Studies*, 2(3):102–111, 2007.

M. Sandelowski. Whatever happened to qualitative description? *Research in Nursing & Health*, 23(4):334–340, 2000.

K. Scarfone and P. Mell. Guide to intrusion detection and prevention systems (idps). Technical report, NIST: National Instutute of Standards and Technology, U.S. Department of Commerce, 2007.

StillSecure. Strataguard ids/ips protection system. http://www.stillsecure.com/strataguard, February 2008.

R. S. Thompson, E. Rantanen, and W. Yurcik. Network intrusion detection cognitive task analysis: Textual and visual tool usage and recommendations. In *Proceedings of the Human Factors and Ergonomics Society Annual Meeting (HFES)*, pages 669–673, 2006.

R. S. Thompson, E. M. Rantanen, W. Yurcik, and B. P. Bailey. Command line or pretty lines?: comparing textual and visual interfaces for intrusion detection. In *CHI '07: Proceedings of the SIGCHI conference on Human factors in computing systems*, pages 1205–1214, New York, NY, USA, 2007. ACM. ISBN 978-1-59593-593-9. doi: http://doi.acm.org/10.1145/1240624.1240807.

R. Werlinger, K. Hawkey, and K. Beznosov. Security practitioners in context: their activities and interactions. In *CHI '08 extended abstracts on Human factors in computing systems*, pages 3789–3794, 2008a.

R. Werlinger, K. Hawkey, and K. Beznosov. Human, Organizational and Technological Challenges of Implementing IT Security in Organizations. In *to appear in HAISA'08: Human Aspects of Information Security and Assurance (13 pages)*, July 2008b.

Chapter 6

Conclusions

This thesis provided a qualitative analysis of four aspects that affect information security in organizations: challenges faced by security practitioners, interactive collaborations among security practitioners and other stakeholders, diagnostic work perform by security practitioners during the response to incidents, and the factors that impact the adoption of an intrusion detection system in one organization. We first summarize the main contributions for each research theme. We continue with the limitations of our study, and finish with directions for future research.

6.1 Contributions

The contributions of this thesis are twofold. First, we provide a richer understanding of the main factors that affect the security within organizations. Second, equipped with this richer understanding, we provide recommendations on how to improve security tools, along with opportunities for future research. We provide a summary of our findings and corresponding contributions related to each theme below.

6.1.1 Challenges theme

We validate and extend prior work by providing a rich description of the challenges to the practice of IT security, as well as integrating these challenges into a framework that organizations can use to identify their limitations with respect to IT security. The research questions related to the challenges theme were: 1) What are the main challenges that security practitioners face in their organizations? (2) How do these challenges interplay? and (3) What are the implications of the challenges on future research?

From our analysis of 27 semi-structured interviews, we identify human, organizational, and technical challenges that security practitioners face when implementing security controls in their

organizations, e.g., lack of security training of other stakeholders (human), distribution of IT responsibilities (organizational), and mobile access to applications (technical). We also show how these challenges interplay and propose research opportunities for the improvement of IT security technologies from a holistic point of view. For example, security technologies need to take into account that security practitioners have to effectively communicate security issues to other stakeholders who have different perceptions of risks and work in a distributed environment (e.g., multiple stakholders involved in the administration of IT systems).

6.1.2 Interactions theme

In the interactions theme, we explore the interactions of IT security practitioners in the context of their organizations, based on a qualitative analysis of 30 interviews and participatory observation. The research questions related to these theme were: (1) When and how do security practitioners interact with other stakeholders? (2) What tools do they need to interact effectively? and (3) What factors are responsible for miscommunications?

We identify nine different activities that require interactions between security practitioners and other stakeholders. Furthermore, we provide detailed descriptions of two activities that may serve as useful references for usability scenarios of security tools. We also propose a model of the factors contributing to the complexity of the interactions between security practitioners and other stakeholders. The discussion is centered on how this complexity is a potential source of security issues that increase the risk level within organizations. Our qualitative analysis also reveals that the tools our participants use to perform their security tasks provide insufficient support for the complex, collaborative interactions they have to perform.

We offer several recommendations for addressing this complexity and improving IT security tools. For example, our findings show that security practitioners sometimes have to combine several tools to perform their security tasks and communicate with other stakeholders, which required copy-pasting between tools, making exchanges of information during interactions error prone. In this vein, an opportunity for improvement is better integration between communication and IT security tools. This improvement might be accomplished through IT security tools that allow on-line collaboration between security practitioners and other stakeholders during the detection and analysis of malicious network traffic.

6.1.3 Diagnosis theme

In the diagnosis theme, we analyze how security practitioners perform diagnostic work during two stages of response to security incidents: preparation and identification. The research questions in the diagnosis theme were: (1) How do security practitioners perform diagnostic work when responding to security incidents?; (2) What tools do security practitioners need to perform this type of diagnostic work?; and (3) How can such tools be improved to better support security practitioners during this diagnostic work?

Based on empirical data from 13 interviews with security practitioners who responded to security incidents and participatory observation in one academic organization, we identify the tasks, skills, strategies, and tools that security practitioners use to diagnose security incidents. Our analysis shows that the diagnosis of security incidents is a highly collaborative activity, which may involve practitioners developing their own tools to perform diagnostic tasks. Furthermore, our findings suggest that diagnosis during security incident response is complicated by practitioners' need to rely on tacit knowledge, as well as usability issues with security tools.

We offer recommendations to improve technology that supports the diagnosis of security incidents, including criteria to evaluate security tools in the context of diagnostic work. Some examples of tool improvements we provide relate to: the tradeoff between task complexity and tool reliability, the need for tools to support tailorability and correlation of high volumes of data, as well as the need for multi-faceted simulation support.

6.1.4 Deployment of an IDS theme

In this theme we investigate the difficulties related to the deployment of an intrusion detection system (IDS) within an organization. The three primary research questions related to this theme were: 1) What do security practitioners expect from an IDS? 2) What are the difficulties they encounter when installing and configuring an IDS? and 3) How can we improve the usability of an IDS?

Our analysis reveals that SPs have both positive and negative perceptions related to the utility of an IDS. The analysis also revealed several issues encountered during the initial stages of IDS deployment. In particular, practitioners found it difficult to make decisions about where to place

the IDS and how to best configure it for use within a distributed environment with multiple stakeholders. We provide recommendations for mitigating these challenges through better tool support.

6.2 Applications

The findings presented in this thesis provide a richer understanding than had been previously available of the different aspects that affect the work of security practitioners. Throughout this thesis, we have proposed three different options with respect to how this understanding can be applied. We summarize these options below.

- First, the recommendations provide a key starting point for tool developers, who can rely on them to improve the design of security tools. For example, as suggested in Chapter 3, security tools can be improved by including features to integrate different communication channels.
- Second, tool evaluators can use the scenarios we describe in Chapters 3 and 4 as *guides* to test security tools in real contexts of interactions and diagnosis of security incidents.
- Third, other researchers can rely on our recommendations for future research to design new studies on information security, as well as to define their own studies based on more specific questions that our research has not addressed.

6.3 Limitations

This thesis investigates in detail various aspects that impact the work of security practitioners within organizations by relying on a qualitative approach. Both the semi-structured interviews and participatory observation provided us with rich data about the each theme considered for the qualitative analysis shown in each chapter. While rich, these data are limited to a small number of security practitioners. Furthermore, during the semi-structured interviews, not all topics were discussed at the same level of detail with all of the participants. Our analysis, therefore, does not include claims about differences in the descriptions made by participants from different levels or organizations. Rather, the findings are centered on the commonalities in their stories.

The variety of organizations whose participants we interviewed, in terms of industrial sectors and sizes (ranging from less than 5 employees to large, multinational companies), has given us a broad perspective about the different aspects under study. Nevertheless, this variety of organizations also represents a limitation of this work; our empirical data lack a sufficient number of organizations to infer the effect of factors such as organization size and sector in our findings. Another limitation related to our sample is the high number of participants from academic organizations, which might bias our findings to the security management perform in academic settings. Despite these limitations, we argue that the models and recommendations shown in each chapter provide enough information to understand the complexity of how human, organizational, and technological factors interplay and impact the development of security tools. However, more data are needed to expand and refine the models and recommendations. For example, validation of the two models presented in chapters 2 and 3 could be performed with a large number of participants from organizations of different sizes within each sector. Furthermore, variations of the models could be developed considering the role of the participant within the organization (e.g., manager versus analyst; security focused versus general IT). Another option is to break down the models into individual, smaller sub-models specific to each activity performed by the security practitioner. As for the recommendations for security tools, the validation and refinement could be performed through the use of surveys or the implementation of prototypes that could be evaluated by groups of security practitioners.

6.4 Future work

Our findings from the qualitative analysis on the four aspects mentioned above (challenges, interactions, diagnostic work, and usability issues with IDSs) contribute to understanding the human, organizational, and technological factors that affect security in organizations and the effectiveness of security tools. Our work also highlights the need for future work by other researchers to continue refinement of how factors interplay by obtaining more rich data (e.g., contextual inquiry), and to generalize and validate these findings through other sources of information (e.g., surveys) in order to study how these factors interplay.

Bibliography

Argus Web Page. Argus intrusion detection and prevention. http://www.qosient.com/argus/, February 2007.

J. Audestad. Four reasons why 100% security cannot be achieved. *Telektronikk*, 1:38–47, 2005.

R. Barrett, M. Prabaker, and L. Takayama. Field Studies of Computer System Sdministrators: Analysis of System Management Tools and Practices. In *Proc. of the Conference on Computer Supported Collaborative Work*, pages 388–395, 2004.

R. Barrett, P. P. Maglio, E. Kandogan, and J. Bailey. Usable autonomic computing systems: The system administratorsṕerspective. *Advanced Engineering Informatics*, 19(3):213–221, 2005. URL `http://www.almaden.ibm.com/u/pmaglio/pubs/icac2004.pdf`.

H. Beyer and K. Holtzblatt. *Contextual Design, Defining Customer-Centered Systems*. Morgan Kaufmann Publishers, San Francisco, CA, 1998.

K. Beznosov and O. Beznosova. On the imbalance of the security problem space and its expected consequences. *Information Management & Computer Security*, 15(5):420–431(12), 2007.

F. J. Björck. *Discovering Information Security Management*. Doctoral thesis, Stockholm University, Royal Institute of Technology, 2005.

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.

D. Botta, R. Werlinger, A. Gagné, K. Beznosov, L. Iverson, S. Fels, and B. Fisher. Studying IT security professionals: Research design and lessons learned. position paper for the CHI Workshop on Security User Studies: Methodologies and Best Practices, April 28 2007a.

D. Botta, R. Werlinger, A. Gagné, K. Beznosov, L. Iverson, S. Fels, and B. Fisher. Towards understanding IT security professionals and their tools. In *SOUPS*, pages 100–111, Pittsburgh, Pennsylvania, July 18-20 2007b.

J. M. Carroll, M. B. Rosson, G. Convertino, and C. H. Ganoe. Awareness and teamwork in computer-supported collaborations. *Interact. Comput.*, 18(1):21–46, 2006. ISSN 0953-5438. doi: http://dx.doi.org/10.1016/j.intcom.2005.05.005.

E. Casey. Error uncertainty and loss in digital evidence. *International Journal of Digital Evidence*, 1(2), 2002.

E. Casey. Case study: Network intrusion investigation – lessons in forensic preparation. *Digital Investigation*, 2(4):254–260, 2005.

S. E. Chang and C. B. Ho. Organizational factors to the effectiveness of implementing information security management. *Industrial Management & Data Systems*, 106:345–361, 2006.

L. P. Chao and K. Ishii. Design error classification and knowledge management. *Journal of Knowledge Management Practice*, 5, 2004.

K. Charmaz. *Constructing Grounded Theory*. SAGE publications, 2006.

S. Chebrolua, A. Abraham, and J. Thomas. Feature deduction and ensemble design of intrusion detection systems. *Computers and Security*, 24(4):295–307, 2005.

S. Chiasson, P. C. van Oorschot, and R. Biddle. Even experts deserve usable security: Design guidelines for security management systems. SOUPS USM Workshop, July 2007.

Cisco Info Center for security monitoring. Cisco info center web page. http://www.cisco.com/en/US/products/sw/netmgtsw/ps5477/products_data_sheet09186a0080187172.html, April 2008.

H. H. Clark. *Using Language*. Cambridge University Press, Cambridge, England, 1996.

D. M. Fetterman. *Ethnography: Step by Step*. Sage Publications Inc., 1998. ISBN 0761913858.

K. Fisler, S. Krishnamurthi, L. A. Meyerovich, and M. C. Tschantz. Verification and change-impact analysis of access-control policies. In *ICSE '05: Proceedings of the 27th international conference on Software engineering*, pages 196–205, New York, NY, USA, 2005. ACM. ISBN 1-59593-963-2. doi: http://doi.acm.org/10.1145/1062455.1062502.

I. Flechais and M. A. Sasse. Stakeholder involvement, motivation, responsibility, communication: How to design usable security in e-science. *Int. Journal of Human-Computer Studies*, in press.

S. Furnell and S. Bolakis. Helping us to help ourselves assessing administratorsúse of security analysis tools. *Network Security*, 2:7–12, February 2004.

A. Gagné, K. Muldner, and K. Beznosov. Identifying differences between security and other IT professionals: a qualitative analysis. In *Proc. of Human Aspects of Information Security and Assurance (HAISA)*, Plymouth, England, July 2008.

R. Garigue and M. Stefaniu. Information security governance reporting. *EDPACS*, 31(6):11–17, 2003.

S. Gibson. The strange tale of the denial of service attacks on grc.com. URL `http://whitepapers.silicon.com/0,39024759,60026382p-39000404q,00.htm`.

J. J. Gonzalez, Y. Qian, F. O. Sveen, and E. Rich. Helping prevent information security risks in the transition to integrated operations. *Telektronikk*, 1:29–37, 2005.

J. R. Goodall, W. G. Lutters, and A. Komlodi. I know my network: Collaboration and expertise in intrusion detection. In *CSCW*, volume 6390, November 2004a.

J. R. Goodall, W. G. Lutters, and A. Komlodi. The work of intrusion detection: Rethinking the role of security analysts. In *Proceedings of the Americas Conference on Information Systems (AMCIS)*, pages 1421–1427, 2004b.

E. Haber and E. Kandogan. Security administrators: A breed apart. In *Proc. of SOUPS Workshop on Usable IT Security Management (USM), 4 pages*, 2007.

E. M. Haber and J. Bailey. Design Guidelines for System Administration: Tools Developed through Ethnographic Field Studies. In *Proc. of 2007 symposium on Computer human interaction for the management of information technology (CHIMIT), 9 pages.* ACM, 2007.

K. Hawkey, D. Botta, R. Werlinger, K. Muldner, A. Gagne, and K. Beznosov. Human, organizational, and technological factors of it security. In *CHI'08 extended abstract on Human factors in computing systems*, pages 3639–3644, 2008a.

K. Hawkey, K. Muldner, and K. Beznosov. Searching for the Right Fit: Balancing IT Security Model Trade-offs. *Special Issue on Useful Computer Security, IEEE Internet Computing*, 12(3): 22–30, 2008b.

D. Heckerman, J. S. Breese, and K. Rommelse. Decision-theoretic troubleshooting. *Communications of the ACM*, 38(3):49–57, 1995.

A. Hill. Shortcuts, Habits, and Sand Castles. In *SOUPS '06: Proceedings of the second symposium on Usable privacy and security*, Pittsburgh, Pennsylvania, 2006. Invited talk.

C. M. Hinckley. *Make No Mistake.* Productivity Press, Portland, OR., 2001.

http://www.cert.org/. Cert: Computer emergency response team.

K. Hwang, M. Cai, Y. Chen, and M. Qin. Hybrid intrusion detection with weighted signature generation over anomalous internet episodes. *IEEE Transactions on Dependable and Secure Computing*, 4(1):41–55, 2007.

K. Jiwnani and M. Zelkowitz. Maintaining software with a security perspective. *Software Maintenance, 2002. Proceedings. International Conference on*, pages 194–203, 2002.

K. Julisch and M. Darcier. Mining intrusion detection alarms for actionable knowledge. In *Proceedings of the 8th ACM International Conference on Knowledge Discovery and Data Mining*, pages 366–375, 2002.

E. Kandogan and E. M. Haber. Security administration tools and practices. In L. F. Cranor and S. Garfinkel, editors, *Security and Usability: Designing Secure Systems that People Can Use*, chapter 18, pages 357–378. O'Reilly Media, Inc., 2005.

A. Kankanhalli, H.-H. Teo, B. C. Tan, and K.-K. Wei. An integrative study of information systems security effectiveness. *International Journal of Information Management*, 23, 2003.

M. Karyda, E. Mitrou, and G. Quirchmayr. A framework for outsourcing is/it security services. *Information Management & Computer Security*, 14:403–416, 2006.

S. Kesh and P. Ratnasingam. A knowledge architecture for IT security. *Commun. ACM*, 50(7): 103–108, 2007. doi: http://doi.acm.org/10.1145/1272516.1272521.

G. Killcrece, K.-P. Kossakowski, R. Ruefle, and M. Zajicek. Incident management. Technical report, U.S. Department of Homeland Security, 2005.

K. J. Knapp, T. E. Marshall, R. K. Rainer, and F. N. Ford. Managerial dimensions in information security: A theoretical model of organizational effectiveness. https://www.isc2.org/download/auburn_study2005.pdf, 2005.

K. J. Knapp, T. E. Marshall, R. K. Rainer, and F. N. Ford. Information security: management's effect on culture and policy. *Information Management & Computer Security*, 14(1):24–36, 2006.

A. Komlod, P. Rheingans, U. Ayachit, J. Goodall, and A. Joshi. A user-centered look at glyph-based security visualization. In *Proceedings of the IEEE Workshops on Visualization for Computer Security (VizSEC)*, pages 21–28, 2005.

I. V. Koskosas and R. J. Paul. The interrelationship and effect of culture and risk communication in setting internet banking security goals. In *ICEC '04*, pages 341–350. ACM Press, 2004. ISBN 1-58113-930-6. doi: http://doi.acm.org/10.1145/1052220.1052264.

A. G. Kotulic and J. G. Clark. Why there aren't more information security research studies. *Information & Management*, 41(5):597–607, 2004.

S. Kraemer and P. Carayon. Human errors and violations in computer and information security: The viewpoint of network administrators and security specialists. *Applied Ergonomics*, 38:143–154, 2007.

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.

E. L. Malécot, M. Kohara, Y. Hori, and K. Sakurai. Interactively combining 2d and 3d visualization for network traffic monitoring. In *Proceedings of the 3rd international workshop on Visualization for computer security (VizSEC)*, pages 123–127, 2006.

A. D. Marwick. Knowledge management technology. *IBM Systems Journal*, 40(4):814–830, 2001.

P. W. Matessich and B. R. Monsey. *Collaboration: what makes it work. A review of research literature on factors influencing successful collaboration.* Amherst H. Wilder Foundation, St. Paul, MN, 1992.

S. McGann and D. C. Sicker. An analysis of security threats and tools in sip-based voip systems. In *In 2nd Workshop on Securing Voice over IP*, June 2005. URL `http://www.colorado.edu/policylab/Papers/Univ_Colorado_VoIP_Vulner.pdf`.

S. Mitropoulos, D. Patsos, and C. Douligeris. On incident handling and response: A state of the art approach. *Computers and Security*, 25(5):351–370, 2006.

S. Mohammed and B. Dumville. Team mental models in a team knowledge framework: Expanding theory and measurement across disciplinary boundaries. *Journal of Organizational Behavior*, 22 (2):89–106, March 2001. ISSN 0894-3796.

M. Nohlberg and J. Backstrom. User-centred security applied to the development of a management information system. *Information Management & Computer Security*, 15(5):372–381, 2007. ISSN 0968-5227. doi: 10.1108/09685220710831116.

M. Park and W. Jung. The requisite characteristics for diagnosis procedures based on the empirical findings of the operators' behavior under emergency situations. *Reliability Engineering & System Safety*, 81(2):197–213, 2003.

M. R. Pattinson and G. Anderson. How well are information risks being communicated to your computer end-users? *Information Management & Computer Security*, 15, 2007.

M. Rao, H. Yang, and H. Yang. Integrated distributed intelligent system architecture for incidents monitoring and diagnosis. *Computers in Industry*, 37(2):143 – 151, 1998.

R. H. Rayford B. Vaughn Jr. and K. Fox. An empirical study of industrial security-engineering practices. *The Journal of Systems and Software*, 61:225–232, 2001.

E. S. Raymond. The cathedral and the bazaar. *First Monday*, 3(3), 1998. URL `http://firstmonday.org/issues/issue3_3/raymond/index.html`.

J. Redish. Expanding usability testing to evaluate complex systems. *Journal of Usability Studies*, 2(3):102–111, 2007.

J. Riden. Responding to security incidents on a large academic network, 2006.

M. Sandelowski. Whatever happened to qualitative description? *Research in Nursing & Health*, 23(4):334–340, 2000.

K. Scarfone and P. Mell. Guide to intrusion detection and prevention systems (idps). Technical report, NIST: National Instutute of Standards and Technology, U.S. Department of Commerce, 2007.

K. Schmidt. Of maps and scripts—the status of formal constructs in cooperative work. In *ACM SIGGROUP*, pages 138–147, November 1997. ISBN 0-89791-897-5.

E. E. Schultz. Computer forensics challenges in responding to incidents in real life setting. *Computer Fraud & Security*, 12:12–16, 2007.

M. A. Shayman, Emmanuel, and Fernandez-Gaucherand. Fault management in communication networks: test scheduling with a risk-sensitive criterion and precedence constraints. In *he 39th IEEE Conference on Decision and Control*, volume 2, pages 1864 – 1869, 2000.

D. A. Siegel, B. Reid, and S. M. Dray. IT Security: Protecting Organizations In Spite of Themselves. *Interactions*, pages 20–27, 2006.

E. H. Spafford. A failure to learn from the past. In *Annual Computer Security Applications Conference (ACSAC)*, Las Vegas, Nevada, 2003. URL `http://www.acsac.org/2003/papers/classic-spafford.pdf`.

P. Stephenson. The application of formal methods to root cause analysis of digital incidents. *International Journal of Digital Evidence*, 3(1), 2004.

StillSecure. Strataguard ids/ips protection system. http://www.stillsecure.com/strataguard, February 2008.

D. W. Straub and R. J. Welke. Coping with systems risk: security planning models for management decision making. *MIS Q.*, 22(4):441–469, 1998. ISSN 0276-7783. doi: http://dx.doi.org/10.2307/249551.

F. O. Sveen, J. Sarriegi, E. Rich, and J. Gonzalez. Toward viable information security reporting systems. In *HAISA 2007*, pages 114–127, July 2007.

R. S. Thompson, E. Rantanen, and W. Yurcik. Network intrusion detection cognitive task analysis: Textual and visual tool usage and recommendations. In *Proceedings of the Human Factors and Ergonomics Society Annual Meeting (HFES)*, pages 669–673, 2006.

R. S. Thompson, E. M. Rantanen, W. Yurcik, and B. P. Bailey. Command line or pretty lines?: comparing textual and visual interfaces for intrusion detection. In *CHI '07: Proceedings of the SIGCHI conference on Human factors in computing systems*, pages 1205–1214, New York, NY, USA, 2007. ACM. ISBN 978-1-59593-593-9. doi: http://doi.acm.org/10.1145/1240624.1240807.

K. Thomson and R. von Solms. Information security obedience: a definition. *Computers & Security*, 24(1):69–75, 2005.

A. Tsohou, M. Karyda, and S. Kokolakis. Formulating information systems risk management strategies through cultural theory. *Information Management & Computer Security*, 14(3):198–217, 2006.

D. M. Wegner. *Transactive memory: A contemporary analysis of the group mind*. In B. Mullen and G. R. Goethals, Editors, Theories of Group Behavior, 1986.

K. Weick and K. Sutcliffe. *Managing the unexpected: assuring high performance in an age of complexity*. Jossey-Bass, 2001.

D. Welch and S. Lathrop. Wireless security threat taxonomy. *Information Assurance Workshop, 2003. IEEE Systems, Man and Cybernetics Society*, pages 76–83, 2003.

R. Werlinger and D. Botta. Detecting, analyzing and responding to security incidents: A qualitative analysis. presented at the SOUPS Workshop on Usable IT Security Management (USM), July 18 2007.

R. Werlinger, K. Hawkey, and K. Beznosov. Security practitioners in context: their activities and interactions. In *CHI '08 extended abstracts on Human factors in computing systems*, pages 3789–3794, 2008a.

R. Werlinger, K. Hawkey, and K. Beznosov. Security practitioners in context: Their activities and interactions with other stakeholders within organizations. *submitted to International Journal of Human Computer Studies*, 2008b.

R. Werlinger, K. Hawkey, and K. Beznosov. The challenges of using an intrusion detection system: Is it worth the effort? In *Proc. of ACM Symposium on Usable Privacy and Security (SOUPS) (12 pages, to appear)*, 2008c.

R. Werlinger, K. Hawkey, and K. Beznosov. Human, Organizational and Technological Challenges of Implementing IT Security in Organizations. In *to appear in HAISA'08: Human Aspects of Information Security and Assurance (13 pages)*, July 2008d.

Appendix A

Coding examples for the Challenges theme

A.1 Examples for the list of challenges/factors that affect security practitioners

Table A.1: Table that shows the axial and open codes. The name of the open codes match the challenges/factors that affect security practitioners

Axial Code	**Open Code**	**Quotes**
(Lack of) security culture	People believing that more access privileges give more seniority	"it doesn't work that way, right. a person could be hired tomorrow and he could be in 129, but that's what he needs to do his job, that more important than [just..] some of the facts that we can have all this really sensitive information and you have this culture of thinking that as you move forward you have to keep getting things added to you." (I5)

Continued on next page

Axial Code	**Open Code**	**Quote**
	People culture as security factor that limits the implementation of more security controls	"The other aspect the particular culture of the company I think plays a big role. What I mean by that is, one of the things that I wanted when I first came here is people had an ID for the LAN that you sign on in the morning and it never expired, the password never expired. What do you mean, never expires? You don't do that, you know, you make it expire in 60 days or 90 days you have to change it. No, no, we've never done that." (I16)
	Trying to change security culture in the organization	"It's been - that's one reason why it taken us five - almost six years to get to the point where we're at now - is it's been a culture change here. And trying to implement a culture change to being at my level in the organization to a higher level is difficult, because usually culture changes come from the top down." (I19)
(Lack of) security training	Not funding for security training	"The training quite honestly is inadequate, not enough funding for training"
	Lack of same background	"Sometimes people don't understand what I am talking about because they don't have the background" (I12)

Continued on next page

Axial Code	**Open Code**	**Quote**
	Difficulties understanding security concepts	"you can't take things away from me. I have seniority. you can add to me, but you cannot take away from me. They don't understand like the security concept of you're doing this job now, you're not doing this job, you don't need that access anymore. so these are some of the major issues that we kind of worked through over the last five years to educate people" (I5)
Different perceptions of risk	Other people do not do risk assessment	"I do my own risk assessment for everything I've responsible for. Unfortunately in my opinion not enough people understand risk management." (I14)
	Explaining security risks	"In my experience these are some of the things that can happen and these are some of the potential situations you'll have to deal with." (I5)
		"The security coordinators take it to the data guardian and explain what the risks are." (I5)
Security issues (communication)	Lack of communication generates incidents	"Well, I mean, yes things happen. Somebody changes a switch or a firewall over here and hasn't adequately communicated that to somebody over here, and part of the research network goes down." (I17)

Continued on next page

Axial Code	**Open Code**	**Quote**
	Lack of communication generates overhead in the investigation of an incident	"for instance I'm being stopped in the hall for a [network] problem, I would come to my cubicle I would not have time to get in synch with the other colleagues right and uh ten minutes later we would find that we pulled the same network device trying to trouble shoot it right." (I8)
Risk estimation	Potential risks in case of a security breach on confidential information of the organization	"Especially once I said to them, look if people download some sort of worm or something like that that mines data and it collects client care information and then sends that out to the internet, you have just compromised basically the security of our information for the agency. And then, like I said the executive directors, who are faced with prison" (I19)

Continued on next page

Axial Code	**Open Code**	**Quote**
	Potential risks in case of a security incident because of malicious software	"In some cases, if I called him I was like, "listen 2389 I really think this is something.... if you're going to go talk to 1691 right now I just thought you should really know this because this is something you gotta explain." if we get a worm comes in here and infects 225 tablet PCs and all the field officers are offline, how are we going to deal with that situation. They have to bring all of those machines back in here so we can re-image them. and we don't have that many people to re-image them so they going to be offline for about a week before we can get these things back out there, you guys gotta keep this in mind as well." (I5)
Open environments and academic freedom	Constraints to apply the same security controls in academic organizations	"there is a lot more sensitivity to academic freedom so some of the things that would be best practices in the security suite of tools we can't necessarily implement quite as easily as a private sector organization" (I1)
	Academic freedom gives more privileges to end-users, increasing levels of risk	"If you don't have firewall to protect it is very vulnerable to get attacked. And the worse I think the university user has more freedom to use that software. Or sometimes they need a research project that they maybe use for some software so they have more powerful user privilege than a user with a typically working company." (I10)

Continued on next page

Axial Code	**Open Code**	**Quote**
	Academic freedom causes more security incidents	"somebody in some lab somewhere sets up a server and maybe they're not a real IT expert, maybe they're a psychologist and they set up a Linux server to get student reports or who knows what they use it for and it'll [just] get hacked all over the place and people would be using it as an open proxy and using it to send spam with and doing all kinds of stuff with it and filling it up with porn and they'll have no idea why, so we've got to come in and figure out what happened and what everybody's responsibility is" (I9)
Lack of budget	Expensive security solutions that are not a priority when work well	"is not cheap in terms of you talk about your firewalls you talk about your anti-virus solutions, your anti-spam solutions, and anti–phishing, there's a lot of money that goes toward that every year, and if everything's running well, they don't see any stuff, they may think its not much of a problem, they don't have to worry about it, right? Let's cut the budget here and cut it there, right?" (I3)
	Lack of budget as a main security issue	"Budget - definitely budget. I have seen a number of times where its really hard to get people to understand that even though you've got no money for budget you still have to try and secure machines, you have to install perimeter firewalls at the very least." (I18)

Continued on next page

Axial Code	**Open Code**	**Quote**
Security as a secondary priority	Security not important unless something goes wrong	"we need to justify because again its low priority because we don't think about it when things are running well, then when something goes wrong you think about it for a couple of weeks, then [it kind of fades] and drops to low priority again right, and so we want something in terms of tools to help keep justify the expenditures" (I3)
	Difficult to make people understand that they need to implement minimal security controls	"So that's the first thing is that units need to have perimeter firewalls, but there's quite often a challenge with getting to understand that you've got 100 workstations installed and you haven't done anything to plan for the security of that." (I18)
Tight schedules	Implement security in projects in reasonable time	"Interviewer: What are other top challenges in your job? I would say that just trying to deploy systems in a reasonable amount of time while also balancing security ...And I would say the biggest challenge is just communicating security in such a way that it is accepted as a business enabler rather than a roadblock." (I25)
Interactions with other organizations	Need to implement specific security controls to connect systems with other organizations	"we have cluster of firewalls that's connected to $SpanBC to do work with 0671 gov. we have more worms and viruses pounding on our firewall through $SpanBC than we have from the raw internet." (I5)

Continued on next page

Axial Code	**Open Code**	**Quote**
	Need to interact with external contractors with no security controls	"the problem is the contactors and temporary people are hired all the time and the information that resides with all the PIs and the upper level staff that says patient data is confidential...And if you are a contractor, you're only here for a few months and are doing something on IT for somebody - you are in, you're out, you don't care." (I17)
Distribution of IT management	Decentralization of security management makes it difficult to implement effective security policies	"the decentralized nature does not help...part of the challenge is that while 4354 does have formal responsible use policies and a policy on access to administrative systems; there aren't formal policies that vest a lot of authority in any group to deal with security." (I2)
	Different access controls depending on the area	"Like you can't get at the Administration systems, and those are very tightly set up and heavily firewalled around those. Then there are other areas of the campus where there is like no other firewall in place because of the way that they need to operate. So it's challenging in that regard." (I15)

Continued on next page

Axial Code	**Open Code**	**Quote**
Access control to sensitive data	Different points to access the data	"there's many different points where this information can be accessed. it can be accessed in databases, it can be accessed through applications, it can be accessed when its in transit, over a network. there's a lot of touch points that are possible, right. you make one error on access control you give somebody access to something that they shouldn't have access to, right" (I4)
	Diverse access control systems	"we have access controls within our banner system and so it's diverse throughout all the different built in.. it's currently mostly in applications, but there's a plan to move it at a.. centrally controlled." (I3)
Complexity of systems	Complexity of firewalls' rules setup	"Interviewer: which particular activities in your job do you think are most error-prone? Participant: I'd say rule-based access setup on the firewall is probably the most and it's one of the things we do basically on a day-to-day basis depending on whats happening on which site. Because understanding what devices are on what networks and what they need to connect to, and if you make an access rule change or a new access rule it can affect their connectivity." (I11)

Continued on next page

Axial Code	**Open Code**	**Quote**
	Networks growing organically	"our networks have sort of grown organically in that we've got sites and we just add in more sites as they get developed. For instance the new building at the 9257 and so we will put in a new device there and set up tunnels, teach the individual the other sites that need to get access to resources there. And now what were doing its getting pretty complex like I say." (I11)
	Too many devices interconnected	"Well, if you look at a complete perimeter defence for any organization there is usually a firewall, there is a demilitarized zone with a couple of boxes; there's a couple of mail servers that are probably using Port Address Translation or something similar. There is usually some switch or other behind the firewall, a couple of routers in front of the firewall, the firewall might be running some sort of proxy server, there may be a proxy in the DMZ; there's a lot of things running there and it's difficult especially for the smaller organizations, even the bigger ones" (I23)
Mobile and distributed access	Mobile laptops get infected every Monday	"Mobile is a big issue. I bring this up many times with management. You know an institute laptop goes home, becomes a toy machine for the family - to bring it back to work Monday morning and it doesn't talk on our network properly because behavior has changed on the laptop." (I14)

Continued on next page

Axial Code	**Open Code**	**Quote**
Vulnerabilities in systems and applications	Need to patch systems quickly in all the systems	"we are also engaged in the software development support if there is a security bug or flaw that becomes known in that software, the onus is on us to get that patch out there so that not just our site, but where it is installed all over the place are able to upgrade it very quickly." (I22)
	Need to monitor new vulnerabilities	"let's say a new vulnerability is announce, the security team is constantly abreast of the news groups and we will notice the vulnerability, assess it...We will then look for tools that will look for vulnerable systems and we will tend to run vulnerability scans, both internally and depending on if it's a particularly serious vulnerability, we'll run campus wide scans." (I22)

A.2 Interplay among challenges/factors

Table A.2: Table with quotes that describe the interplay among challenges/factors

This challenge	**Interplays with**	**Support**
(Lack of) security training	(Lack of) security culture	"you can't take things away from me. I have seniority. you can add to me, but you cannot take away from me. They don't understand like the security concept of you're doing this job now, you're not doing this job, you don't need that access anymore. so these are some of the major issues that we kind of worked through over the last five years to educate people" (I5)
	Lack of budget	"The training quite honestly is inadequate, not enough funding for training" (I3)
	Different perceptions of risk	"I do my own risk assessment for everything I've responsible for. Unfortunately in my opinion not enough people understand risk management." (I14)
Different perceptions of risk	Security issues (communication)	"But he would say to them, you know you are having all of these companies having hours, days of down time and then people say yeah, but that's not happening here. And he would say to them exactly, it's not happening here now. Don't you understand this correlation; that we've implemented all these security features and things like this aren't hitting us any more?" (I19)

Continued on next page

This challenge	**Interplays with**	**Support**
Different perceptions of risk	Estimation of risks	"Especially once I said to them, look if people download some sort of worm or something like that that mines data and it collects client care information and then sends that out to the internet, you have just compromised basically the security of our information for the agency. And then, like I said the executive directors, who are faced with prison" (I19)
Open environment, Academic freedom	Distribution IT Management	"its not a corporation its sort of like a sometimes like a loose concatenation of different bodies which have their different mandates and nobody really wants to come down on everybody with an iron fist or anything like that" (I9)
	Control access	"we can't necessarily implement quite as easily as a private sector organization can simply due to the academic freedoms and expectations and needs faculty and students. I know that's an interesting trade off all the time. You're constantly trading access vs. risk." (I1)
		"its not a corporation its sort of like a sometimes like a loose concatenation of different bodies which have their different mandates and nobody really wants to come down on everybody with an iron fist or anything like that" (I9)

Continued on next page

This challenge	**Interplays with**	**Support**
	Technical complexity	"somebody in some lab somewhere sets up a server and maybe they're not a real IT expert, maybe they're a psychologist and they set up a Linux server to get student reports or who knows what they use it for and it'll [just] get hacked all over the place and people would be using it as an open proxy and using it to send spam with and doing all kinds of stuff with it and filling it up with porn and they'll have no idea why, so we've got to come in and figure out what happened and what everybody's responsibility is" (I9)
Control Access	Technical complexity	"all of the information that we have in our databases has to be controlled, because its either employee information or its customer information. so there's got to be a way to ensure that only the people who need to have access to it have access to it when they need to have access to it, and its locked. that's very difficult" (I4)
Security low priority	Lack of budget	"we need to justify because again its low priority because we don't think about it when things are running well, then when something goes wrong you think about it for a couple of weeks, then [it kind of fades] and drops to low priority again right, and so we want something in terms of tools to help keep justify the expenditures" (I3)

Continued on next page

This challenge	**Interplays with**	**Support**
Security low priority	Lack of budget	"Most organizations are not looking at security as their number one priority. Their main priority is to keep systems running and to keep systems running at a reasonably high speed and high response rate to get as much functionality out of their systems, to pay as little money for their systems as possible. All that conspires against the systems being as secure as they could be and nobody is running Fort Knox; nobody really thinks that their information is important - they think it's important but they don't think it is important enough to devote a hundred thousand dollars a month to protecting it for example." (I23)
Interactions with other organizations	Technical complexity / Control access	"we have cluster of firewalls that's connected to $SpanBC to do work with 0671 gov. we have more worms and viruses pounding on our firewall through $SpanBC than we have from the raw internet." (I5)
Distribution of IT Management	Technical complexity	"Like you can't get at the Administration systems, and those are very tightly set up and heavily firewalled around those. Then there are other areas of the campus where there is like no other firewall in place because of the way that they need to operate. So it's challenging in that regard." (I15)

Continued on next page

This challenge	**Interplays with**	**Support**
Distribution of IT Management	Security Issues (Communication)	"our security guy who is our network guy was changing something in 4831, the 5052, which affected a bunch of researchers in the 2244, which is in other hospital buildings. It was a lack of communication between IT guy at 2244 who is also one of our employees and the security guy" (I17)
Interaction with other organizations	Security low priority	"lately we've been doing a lot more work outside third-parties...you don't necessarily at that point have the luxury of saying this has to be within our goal technology architecture, sometimes you have to make exceptions to that and assume the risks." (I4)
Security low priority	Tight schedules	"Interviewer: What are other top challenges in your job? Interviewee: I would say that just trying to deploy systems in a reasonable amount of time while also balancing security ...And I would say the biggest challenge is just communicating security in such a way that it is accepted as a business enabler rather than a roadblock." (I25)
(Lack of) Security training	Security low priority	"So that's the first thing is that units need to have perimeter firewalls, but there's quite often a challenge with getting to understand that you've got 100 workstations installed and you haven't done anything to plan for the security of that." (I18)

Continued on next page

This challenge	**Interplays with**	**Support**
Technical complexity	Vulnerabilities in systems and applications	"we are also engaged in the software development support if there is a security bug or flaw that becomes known in that software, the onus is on us to get that patch out there so that not just our site, but where it is installed all over the place are able to upgrade it very quickly." (I22)
	Mobility	"Mobile is a big issue. I bring this up many times with management. You know an institute laptop goes home, becomes a toy machine for the family to bring it back to work Monday morning and it doesn't talk on our network properly because behavior has changed on the laptop." (I14)

Appendix B

Detailed notes from participatory observation

B.1 Deciding on the Purpose of the IDS

During the first meeting of the two security specialists involved in the deployment of the IDS, one of the main points of discussion was the type of reports that the IDS needed to provide. The security specialist from the servers area was looking for evidence to show the effectiveness of the rules that were implemented in the firewalls. To obtain this evidence, it was necessary to install two sensors in the IDS system, one before the firewall and the other one after. The differences between the alarms shown by the two probes would give a sense of how well the firewalls were configured. Such a report would shed more light on the investment decisions and business cases that the organization was considering for IT security. For example, a report saying that no attacks were crossing the firewalls and routers would confirm that those devices were saving the organization money, by avoiding security incidents. On the other hand, if the firewalls and routers were not filtering properly, then this would provide our specialist with evidence to support the purchase of firewalls with better functionalities and centralized management.

The security specialist from the network area was concerned about the IDS set up proposed by his colleague, the server security specialist, for two reasons. First, the IDS might be unable to process all the information from the probe set up before the firewall. Second, the information might have little use, as the priority is identification of attacks that could actually penetrate to the internal systems.

The final decision about the purpose of the IDS was determined by practical issues. Given the lack of resources (e.g., time, man hours), the IDS was going to be installed with its basic

configuration, with one probe only. This decision was discussed in parallel with where to position the IDS in the network.

B.2 Integrating the IDS in the Network

The discussion in the first meeting described above was supported by a sketch on the whiteboard of the internal network, including the main routers, switches, firewalls, and servers (see figure 5.1). This diagram had two main objectives. The first objective was to reach a common understanding of the current status of the network. The importance of this shared understanding was evident during the discussion, as each specialist knew unique details about the network. The second objective was to find ports available for connection to the sensor and management IDS ports.

From a technical point of view, the decision about where to position the IDS had several constraints. One of them was the bandwidth of the critical traffic to be monitored, which had to be smaller than 100 Mbps. Another constraint was the routing necessary to reflect in one specific network-device port all the traffic to be monitored. To do so, the traffic had to go through different devices and links that may not have spare capacity. The decision about the location of the IDS was not made in the first meeting, and the discussion continued during the second and third meetings.

During the second meeting the connection of the IDS was discussed in more detail. The initial idea of connecting the IDS to one of the routers was deemed impractical as the network had been reconfigured with new devices. These new devices would require a special module (not installed at that moment) to mirror traffic in one of its ports. The other possibility was to connect the IDS with another device within the same network domain, but the only port available in that device for reflecting traffic was reserved for troubleshooting during the investigation of network anomalies. Within this option, there were also issues with physically carrying the traffic to the room where the IDS was going to be installed. The security specialists discussed if it would be possible to reflect traffic in one device in the middle, and then reflect again this traffic in a second device in the target room. This was deemed infeasible so they had to evaluate a physical extension of the cables to connect the IDS.

Several issues arose during the connection of the management port of the IDS. It was not clear if the IDS' management port should be in the management network or if it was necessary to create

a different VLAN for it. The network specialist proposed to create another VLAN to connect the IDS, but this option was deemed too complex. Another issue was the security of the management port; in the case of a new VLAN, it would be necessary to configure additional firewalls specifically for the IDS.

Given the inconveniences of connecting the IDS's ports, the security specialists began to evaluate other alternative locations for the IDS. This change in location meant that they would be giving up the possibility of monitoring the most important traffic in the network, but did have the benefit of decreased complexity. This situation would have an impact, as the IDS-related reports would include as interesting results as they were originally hoping. The final decision about the location of the IDS was postponed until the third meeting.

During the third meeting, the security specialists continued to discuss the option of installing the IDS system in a less critical network. They finally decided to adopt this last option, connecting the IDS's sensor in the network that carried traffic generated by the organization's internal staff members. These conditions made the project less ambitious, and it was now considered a pilot study. The management port was connected to one production network. In making this decision, the specialists discarded the connection of this port in the management network, which carries all the management traffic from the organization's devices. The main reason for not taking this option was that the security specialists did not want to involve the administrators of the management network, in order to reduce the project overhead.

Another topic discussed in the meetings was related to the configuration of the IDS, described in the next section.

B.3 Initial configuration of the IDS

The security specialists knew from their previous experience that customizing the IDS to the connecting network is a time consuming and iterative process. An IDS that is well tuned should minimize both false positives (i.e.,alarms that correspond to valid traffic) and false negatives (failure to generate alarms for anomalous traffic).

The IDS configuration was done by the observer as part of his individual activities. To be more prepared for the eventual tuning in the real network, the objective for the observer's individual work

was to become familiar with the IDS and its graphical interface. The first task was to reinstall the IDS software on the server. This process, which took 20 - 30 minutes, automatically installed the required components of the IDS system: the Linux operating system, PostrgreSQL database, Snort rules for detecting malicious traffic, and the IDS graphical interface. The information required by the system to finish the initial configuration included: (1) the IDS port IP addresses, which were set for only the management port in one of the organization's internal, secure networks, and (2) two passwords, one for configuring the IDS and another for the database. The strength of these passwords was not checked by the system.

During the installation, the observer noted that the IDS did not allow for customizing the configuration of the system. This was not surprising, as this packaged IDS software is intended to alleviate the burden of having to integrate each of the IDS components manually, performing in the background all the necessary steps to have the system running quickly. However, there were some configuration options that the posterior use of the IDS showed needed customization. These options were related to: (1) the partitions that the system assigned to the filesystem, and (2) the server security settings that prevent unauthorized access of the IDS. These were not shown in the initial setup, and they were not accessible from the IDS's graphical user interface.

The partitions assigned to the filesystem were important because the security specialists knew from their previous experience that the file space for the logs might be too small. They wanted to check that the new version of the IDS had more space for the logs, but again this was a setting that could not be configured within the IDS' graphical user interface and that required the use of additional tools.

The ability to access the IDS security settings was important because the security specialists wanted to know what type of firewall rules, if any, were necessary to protect the IDS's ports. In its IPTable file, the IDS system recommended to not modify the default protection settings. These settings would be hard for a security practitioner to understand, particularly for one who was seeking the usability advantages afforded by the IDS's graphical interface.

Another drawback of the IDS installation and booting processes was that some error messages did not give sufficient information about their cause or consequences. For example, during the installation process, the message: "ACPI resource is not an IRQ entry" was displayed; and during the booting process, the message "smartd failed initialization" appeared. These messages became

very relevant, as the IDS's management port could not initially connect to the central server of the vendor, and it was not clear this issue was due to problems related to those messages or to the configuration of the network's filters.

Troubleshooting of the IDS's Internet connection revealed that it was the network that was blocking the connections. As a consequence, the IDS's management port was moved to another, more open, network where the system started to download the rules from the vendor's server. The next step in the observer's individual work was to develop an understanding of the configuration options that the graphical interface provided, particularly the detection rules.

B.4 Effectiveness of the Graphical User Interface

Through the graphical user interface it was possible to make changes to the IDS rules and to the system's configuration (see figure 5.1). This last option allowed the modification of parameters such as the IP addresses of the ports, the autodiscovery option, and the networks to be monitored.

Without real traffic, it was very difficult to anticipate the types of alarms that the IDS was going to report. The only way to configure the system in such circumstances is by already possessing detailed knowledge about all the valid protocols that the network carried. However, the organization's open, distributed environment included traffic unknown to the security specialists. In such a situation, the organizational security policies may play an important role; for instance, a full set of rules could be disabled if the organizational policies do not exclude certain traffic (e.g., disable rules associated with port scanning).

This inability to anticipate alarms made it clear that the tuning process could not be done off-line; it was necessary to look at real traffic. Unfortunately, predicting the complexity of configuring the IDS in a particular network is very difficult. Consequently, the security specialists did not know if it was worth tuning the IDS for a simple domain of the network (i.e., low traffic, not many different types of devices) versus directly tuning the system for the more important, complex domains.

The IDS interface also provided an option of quick tuning, that looked like a good way of avoiding the specification of all the default rules of the IDS (more than 1,000). However, without real traffic it was impossible to assess the tradeoffs associated with this option.

Another aspect important for the security specialists was the ability to notify other administrators about malicious traffic in their networks. The next section gives details on this issue.

B.5 Configuring for Multiple Stakeholders

The IDS was supposed to detect security events and send alarms to those internal stakeholders who should be notified of security incidents. The security specialists were worried about the benefit of these notifications; they had to be very careful to limit the number of false positive notifications. This meant that the alarms issued by the IDS needed to be preprocessed.

Another functionality that security specialists needed in their collaborative environment was the definition of access accounts to the IDS system with different privileges. For example, some users should be able to look at alarms from specific network domains, without looking at alarms from other domains. However, despite the fact that the IDS was monitoring traffic that was going to different domains, the system did not allow different accounts when it was installed with a single sensor node.

Appendix C

UBC Research Ethics Board's Certificate

This appendix includes the UBC Research Ethics Board's Certificate required to recruit participants in our study. Figure C.1 shows the certificate, Figures C.2 and C.3 show its amendment, and Figures C.4 and C.5 show its renewals.

The University of British Columbia
Office of Research Services and Administration
Behavioural Research Ethics Board

Certificate of Approval

PRINCIPAL INVESTIGATOR	DEPARTMENT	NUMBER
Beznosov, K.	Electrical and Computer Eng	**B06-0413**

INSTITUTION(S) WHERE RESEARCH WILL BE CARRIED OUT

UBC Campus ,

CO-INVESTIGATORS:

Botta, David, ; Fels, Sidney, Electrical and Computer Eng; Fisher, Brian, Computer Science; Gagne, Andre, Computer Science; Iverson, Lee, Electrical and Computer Eng; Werlinger, Rodrigo, Electrical and Computer Eng

SPONSORING AGENCIES

Natural Science Engineering Research Council

TITLE:

HOT Admin: Human, Organization, and Technology Centred Improvement of IT Security Administration

APPROVAL DATE	TERM (YEARS)	DOCUMENTS INCLUDED IN THIS APPROVAL:
JUN 20 2006	1	June 15, 2006, Consent form / Contact letter / May 19, 2006, Questionnaires

CERTIFICATION:

The application for ethical review of the above-named project has been reviewed and the procedures were found to be acceptable on ethical grounds for research involving human subjects.

Approved on behalf of the Behavioural Research Ethics Board
by one of the following:
Dr. Peter Suedfeld, Chair,
Dr. Susan Rowley, Associate Chair
Dr. Jim Rupert, Associate Chair
Dr. Arminee Kazanjian, Associate Chair

This Certificate of Approval is valid for the above term provided there is no change in the experimental procedures

Figure C.1: UBC Research Ethics Board's Certificate

The University of British Columbia
Office of Research Services
Behavioural Research Ethics Board
Suite 102, 6190 Agronomy Road, Vancouver, B.C. V6T 1Z3

CERTIFICATE OF APPROVAL - MINIMAL RISK AMENDMENT

PRINCIPAL INVESTIGATOR:	DEPARTMENT:	UBC BREB NUMBER:
Konstantin Beznosov		H06-80413

INSTITUTION(S) WHERE RESEARCH WILL BE CARRIED OUT:

Institution	Site
UBC	Point Grey Site

Other locations where the research will be conducted:
N/A

CO-INVESTIGATOR(S):

Brian D. Fisher
Andre Gagne
David Botta
Rodrigo Werlinger
Lee Iverson
Sidney S. Fels

SPONSORING AGENCIES:

Natural Science Engineering Research Council - "HOT Admin: Human, Organization, and Technology Centred Improvement of IT Security Administration"

PROJECT TITLE:
HOT Admin: Human, Organization, and Technology Centred Improvement of IT Security Administration

Expiry Date - Approval of an amendment does not change the expiry date on the current UBC BREB approval of this study. An application for renewal is required on or before: June 20, 2007

AMENDMENT(S):

AMENDMENT APPROVAL DATE:
December 2, 2006

Document Name	Version	Date
Advertisements:		
Advertisement Flyer	1	November 25, 2006
Questionnaire, Questionnaire Cover Letter, Tests:		
online questionnaire	1	November 26, 2006
cover letter	1	November 26, 2006

Figure C.2: Amendment UBC Research Ethics Board's Certificate, page 1

The amendment(s) and the document(s) listed above have been reviewed and the procedures were found to be acceptable on ethical grounds for research involving human subjects.

Approval is issued on behalf of the Behavioural Research Ethics Board and signed electronically by one of the following:

Dr. Peter Suedfeld, Chair
Dr. Jim Rupert, Associate Chair
Dr. Arminee Kazanjian, Associate Chair
Dr. M. Judith Lynam, Associate Chair

Figure C.3: Amendment UBC Research Ethics Board's Certificate, page 2

The University of British Columbia
Office of Research Services
Behavioural Research Ethics Board
Suite 102, 6190 Agronomy Road, Vancouver, B.C. V6T 1Z3

CERTIFICATE OF APPROVAL- MINIMAL RISK RENEWAL

PRINCIPAL INVESTIGATOR:	DEPARTMENT:	UBC BREB NUMBER:
Konstantin Beznosov	UBC/Applied Science/Electrical and Computer Engineering	H06-80413

INSTITUTION(S) WHERE RESEARCH WILL BE CARRIED OUT:

Institution	Site
UBC	Point Grey Site

Other locations where the research will be conducted:
N/A

CO-INVESTIGATOR(S):

Brian D. Fisher
Andre Gagne
David Botta
Rodrigo Werlinger
Lee Iverson
Sidney S. Fels

SPONSORING AGENCIES:

Natural Sciences and Engineering Research Council of Canada (NSERC) - "HOT Admin: Human, Organization, and Technology Centred Improvement of IT Security Administration"

PROJECT TITLE:
HOT Admin: Human, Organization, and Technology Centred Improvement of IT Security Administration

EXPIRY DATE OF THIS APPROVAL: June 6, 2008

APPROVAL DATE: June 6, 2007

The Annual Renewal for Study have been reviewed and the procedures were found to be acceptable on ethical grounds for research involving human subjects.

Approval is issued on behalf of the Behavioural Research Ethics Board

Figure C.4: UBC Research Ethics Board's Certificate, first renewal

The University of British Columbia
Office of Research Services
Behavioural Research Ethics Board
Suite 102, 6190 Agronomy Road, Vancouver, B.C. V6T 1Z3

CERTIFICATE OF APPROVAL- MINIMAL RISK RENEWAL

PRINCIPAL INVESTIGATOR:	DEPARTMENT:	UBC BREB NUMBER:
Konstantin Beznosov	UBC/Applied Science/Electrical and Computer Engineering	H06-80413

INSTITUTION(S) WHERE RESEARCH WILL BE CARRIED OUT:

Institution	Site
UBC	Vancouver (excludes UBC Hospital)

Other locations where the research will be conducted:
N/A

CO-INVESTIGATOR(S):

Brian D. Fisher
Andre Gagne
David Botta
Rodrigo Werlinger
Lee Iverson
Sidney S. Fels

SPONSORING AGENCIES:

Natural Sciences and Engineering Research Council of Canada (NSERC) - "HOT Admin: Human, Organization, and Technology Centred Improvement of IT Security Administration"

PROJECT TITLE:
HOT Admin: Human, Organization, and Technology Centred Improvement of IT Security Administration

EXPIRY DATE OF THIS APPROVAL: April 23, 2009

APPROVAL DATE: April 23, 2008

The Annual Renewal for Study have been reviewed and the procedures were found to be acceptable on ethical grounds for research involving human subjects.

Approval is issued on behalf of the Behavioural Research Ethics Board

Dr. M. Judith Lynam, Chair
Dr. Ken Craig, Chair
Dr. Jim Rupert, Associate Chair
Dr. Laurie Ford, Associate Chair
Dr. Daniel Salhani, Associate Chair
Dr. Anita Ho, Associate Chair

Figure C.5: UBC Research Ethics Board's Certificate, second renewal