Security Practitioners in Context: Their Activities and Interactions with Other Stakeholders within Organizations

Rodrigo Werlinger[a,**], Kirstie Hawkey[a], David Botta[a], Konstantin Beznosov[a,*]

[a] *University of British Columbia, 2332 Main Mall, Vancouver, Canada*

Abstract

This study investigates the context of interactions of IT security practitioners, based on a qualitative analysis of 30 interviews and participatory observation. We identify nine different activities that require interactions between security practitioners and other stakeholders, and describe in detail two of these activities that may serve as useful references for usability scenarios of security tools. We propose a model of the factors contributing to the complexity of interactions between security practitioners and other stakeholders, and discuss how this complexity is a potential source of security issues that increase the risk level within organizations. Our analysis also reveals that the tools used by our participants to perform their security tasks provide insufficient support for the complex, collaborative interactions that they need to perform. We offer several recommendations for addressing this complexity and improving IT security tools.

Key words:
Security Tools; Usable Security; Security Practitioners; Collaboration; Qualitative Analysis

1. Introduction

Security of information technology (IT) has become a critical issue for organizations that 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

*Corresponding author
**Principal author
Email addresses: rodrigow@ece.ubc.ca (Rodrigo Werlinger), hawkey@ece.ubc.ca (Kirstie Hawkey), botta@ece.ubc.ca (David Botta), beznosov@ece.ubc.ca (Konstantin Beznosov)

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 interdependency 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 details on how IT professionals with security responsibilities (hereafter referred to as *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 human-computer interaction (HCI) researchers and tool developers to improve communication and IT security tools. Furthermore, such understanding is needed to develop tests for measuring the usability of security tools in real, complex scenarios (Redish 2007).

We argue elsewhere that human, organizational, and technological factors influence the ability of security practitioners to do their job well (Botta et al. 2007; Werlinger et al. 2008a). To understand how these factors play out in IT security, we conducted 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 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 were supplemented by a 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: (1) the activities of security practitioners that required collaboration, and (2) communication between them and other stakeholders within the organization. A preliminary version of this work appeared in Werlinger et al. (2008b). 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 for evaluating IT security tools. Second, based on these results, we

propose a model that shows which factors make interactions between security practitioners and other stakeholders complex. We also 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 employed 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 need 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 to provide more integration between communication tools and IT security tools. This improvement might be accomplished through IT security tools that allow online 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 IT, collaborative work in the context of IT security, supporting complex tasks, and communication models. In Section 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 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 5, we provide in-depth descriptions of interactions during two activity scenarios: responding to security incidents and developing security policies. In Section 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 7, we discuss the implications of our findings for researchers and practitioners, including opportunities for improved tool support. We discuss the limitations of our research and opportunities for future work in Section 8. We conclude the paper in Section 9.

2. Background and Related Work

Prior research has examined computer supported collaborative work (CSCW). For example, studies by Carroll et al. (2006) and Mohammed and Dumville (2001) propose general frameworks for understanding team effectiveness, and suggest future directions to improve empirical methods and the design of systems that support collaborations. Similarly, Neale et al. (2004) propose a multifaceted framework (i.e., joint awareness, communication, collaboration, work coupling) for evaluating distributed CSCW applications. Neale et al. point out

that the joint awareness for joint activities in the workplace is a macro-level manifestation of Clark's (1996) micro-level common ground. In Clark's theory, people build common ground by establishing, on a moment-by-moment basis, that they mutually understand each other. Although these frameworks integrate different facets of collaboration (e.g., activity awareness, common ground, coupling of work), the complexity and importance of IT administration and IT security have motivated specific empirical studies on collaborative work within these two domains. In this section, we discuss relevant empirical research.

2.1. Empirical Research on IT Collaborative Work

Barrett et al. (2004) and Haber and Bailey (2007) used ethnographic methods to study system administrators in context. Their findings touch upon a broad spectrum of IT administration (e.g., databases, web servers, operating systems), and show that IT administrators collaborate to manage risk, system complexity and system scale. They also find that IT tools do not provide proper support for the collaborative tasks performed by IT professionals. Since a good deal of security management is done by system administrators, their findings are relevant to ours.

At the same time, preliminary analysis of our HOT Admin data (Gagné et al. 2008) and work by Haber and Kandogan (2007) indicate that the practice of IT security has characteristics that differ from other IT practices. For example, security practitioners work in a particularly fast-paced and constantly changing environment; troubleshooting can be more complex, requiring both deep and broad knowledge of IT systems and of the organization (Gagné et al. 2008; Haber and Kandogan 2007). Furthermore, stakeholders within the organization often have a negative perception of security practitioners, which requires the use of persuasion during communication (Gagné et al. 2008). In the study we present here, we examined the interactions and tools used during security-related activities (see Sections 5 and 6), which are not necessarily the same as in the context of IT administration in general. Where possible (e.g., Sections 6.3 and 8), we contrast our results with the ones found by Barrett et al. (2004). Also, in the context of IT security, we build on the recommendations by Haber and Bailey (2007) to (1) integrate IT tools with each other and with monitoring and management tools, and (2) enable shared views on a system for the sake of collaboration.

2.2. Empirical Research on IT Security Collaborative Work

Björck (2005) used grounded theory (Glaser and Strauss 1967) to understand the challenges in establishing a balanced management system for IT security. The data for his study came from 29 semi-structured interviews with eight IT security managers, thirteen consultants, and eight auditors from different Swedish companies. His finding is that sound communication capabilities are one of the success factors for the formation and certification of IT security management systems.

Kandogan and Haber present two different studies related to IT security administrators. Kandogan and Haber (2005) evaluated 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) analyzed ethnographic data from 16 field studies of IT administrators to determine differences between IT system and security administrators. Their findings are that, unlike collaborations of other system administrators, security administrators' collaborations require a greater learning component as new vulnerabilities are discovered daily. Security administrators also have a warfare-like approach; they do not have the incentive of sharing information as widely as other IT administrators do, because this information can be used by attackers against the administrators' systems.

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 nine interviews of a diverse cross-section of intrusion detection experts. Their conclusions are 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) present an investigation on 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. Their conclusion is that security tasks have a high level of interdependency, requiring contributions of other individuals and resources.

Flechais and Sasse (2007) present a study on 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 IT security. Their conceptual framework is populated 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 IT security errors, and that communication breakdowns cause security vulnerabilities.

A preliminary analysis of our data (Botta et al. 2007) identifies 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. The results suggest that

IT security responsibilities are distributed among many individuals, and that security-relevant tools should be flexible enough to (1) be used in combination with other tools, and (2) be tailorable to various unique situations that often involve distributed cooperation. In this current paper we report on activities where security practitioners have to interact with other stakeholders and how security tools can provide better support for these interactions.

2.3. Complex Tasks

Many IT security tasks are performed in response to real-time issues. The practice of IT security is also complex, sociotechnical, and involves balancing dynamic perceptions of risk. Because *ecological interface design* (EID) (Vicente and Rasmussen 1992; Vicente 1999; 2002) addresses such work environments; it is important to consider the relationship between EID and our work. Although our research does not go so far as to propose user interfaces, it does provide information that is relevant for the EID approach. EID models the work environment by reflecting its constraints and rules hierarchically at different levels of abstraction, in which a lower level comprises a means for achieving the next higher level. In this way, all of the constraints of a system are laid bare, so that practitioners can better diagnose unexpected violations of constraints. This is similar to how a road map, as opposed to a list of directions, enables a driver to adapt to an unexpected detour. Also, the *kind* of information that is to be processed in order to achieve particular goals is identified. This allows the information to be represented appropriately as signals, signs, or symbols, so that the practitioner can bring to bear skills to signals, rules to signs, or knowledge to symbols. Our work identifies high-level IT security goals (Section 4.1) for which multiple stakeholders (e.g., security practitioners, managers, clients, end-users) must interact in order to achieve them. Thus our work contributes to an abstraction hierarchy, and also a flow hierarchy for IT security. Our work also takes a step toward determining the kind of information that is exchanged with respect to different goals during interaction (Sections 5.1, 5.2, and 6.3), and thereby contributes to the *skill, rules, knowledge* aspect of the EID framework.

2.4. Communication Models

Another area of research relevant to ours is investigation of communications among stakeholders in particular and individuals in general. Lloyd et al. (1990) recognize at least four common key parameters in communication models (e.g., Shannon and Weaver (1949); Berlo (1960); Sanders (1976)):a sender, a receiver, a message, and feedback from the receiver to the sender. These models have been used in proposed communication frameworks and theories in specific fields, such as development of computer-base information systems (Guinan and Bostrom 1986), non-spoken verbal communications by disabled individuals (Lloyd et al. 1990), and group work in the context of negotiation systems (Benbasat et al. 1995). Benbasat et al. use Berlo's model structure (Berlo 1960) to outline the various dimensions in group interfaces for computer supported work, with emphasis on negotiation tasks. Berlo's model provides a

concrete list of elements that describe each component of the communication, namely communication skills, attitudes, knowledge, social system, and culture for the sender and receiver; elements, content, treatment, structure, and code for the message; and the five human senses (seeing, hearing, touching, smelling, and tasting) for the channel. Similarly, we use Berlo's model to relate the elements that make security interactions more complex in the context of communication models (Section 6.5).

Other studies have investigated how people select tools to communicate depending on their activities (e.g., Pinelle and Gutwin (2003)) or attitudes and behavior (e.g., Trevino et al. (2000)). Pinelle and Gutwin (2003) worked with clinicians and administrators for three years to develop group support technologies for home care clinicians. They propose several recommendations for designing communication tools that are intended to support the work of multidisciplinary groups that act independently, such as the functional units for managing community-based patient care. These recommendations include flexibility and consolidation of information. Trevino et al. (2000) propose an integrated set of hypotheses that relate the attitudes or behavior of an individual with his or her preference for some communication medium, such as fax, phone, or e-mail. One of their findings was that "individuals are more likely to choose rich media (e.g., face-to-face meetings) and that they are less likely to choose lean media (e.g., letters, fax) when message equivocality is high." We use these previous findings on communication tools to elaborate on some of the communication features that we propose for security tools in Section 7.

2.5. Summary

As discussed above, prior studies have used empirical data to discover 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 IT administrators and specifically 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; we next describe our research methods in detail.

3. Research Methods

Our three primary research questions for the study reported in this paper 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 miscommunication? 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.

3.1. Participant Recruitment

This study of interactions among security practitioners and other stakeholders was 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. 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, making initial contacts difficult. We used two strategies in the HOT Admin project to address these challenges. First, professional connections 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 were interested to meet for a one-hour interview. For a discussion on the effectiveness of these strategies, see Botta et al. (2007). In the next section, we describe the questionnaires, interviews, and participatory observations that composed our study data.

Participants included IT and security managers, and IT and security specialists. Table 1 shows the positions held by our participants across these different sectors. In total, we conducted 30 interviews – numbered I1, I2, ..., etc. It should be understood that, occasionally (I1, 17, I6, I22), interviews contained two participants – a primary interviewee, plus another who added details or confirmed recollections of events. A couple of times (I1 and I17) the secondary participants were recruited for individual interviews about their own experiences (I3 and I18). Altogether, we interviewed 32 security practitioners from 14 organizations (3 academic, 11 non-academic).

Sixteen of our participants provided an estimate of the time they spent on IT security. Two of the IT managers that did so estimated that they spent less than 10% of their time on security-related activities, while the other estimated 20%. Five security specialists gave estimates ranging between 20 to 100 percent of their time (on average 67%). The IT specialists with security duties who gave estimates averaged less time (10%–60%, average of 30%) than the security specialists. No security managers provided this estimate. We discuss the impact on our findings of position and time spent on security in Section 8.

3.2. Data Collection

The data were collected through demographic questionnaires, semi-structured interviews, and participatory observations.

3.2.1. Questionnaires and Semi-Structured Interviews

The questionnaires completed by participants provided demographic information (e.g., job title, type of organization). This questionnaire was the first means of contact for the majority of the interviewees. We obtained 32 questionnaires, which led to 21 interviews. Other means, such as personal contact, led to further interviews for which we do not have questionnaire data.

Table 1: For each organization type, we indicate the number of unique organizations and the total participants interviewed. We indicate the interview in which each participant took part; for those interviews with two participants we indicate whether the participant was the primary interviewee (a) or the secondary interviewee (b). Participants held various positions, including managers with security tasks, regular IT practitioners with security tasks, security managers, and security specialists.

Organization Type	**Position Type**				**Total Org**
	IT Manager	**Security Manager**	**Security Specialist**	**IT specialist**	
Academic (3)	4 (I1a, I15, I17a, I17b/I18)	1 (I2)	4 (I1b/I3, I9, I11, I21)	9 (I6a, I6b, I7, I8, I10, I14, I20, I22a, I22b)	18
Financial Services (2)	-	-	2 (I4, I25)	-	2
Insurance (1)	-	-	2 (I5, I28)	-	2
Scientific Services (1)	-	-	-	2 (I12, I13)	2
Manufacturing (1)	1 (I16)	-	1 (I21)	-	2
Telecommunications (2)	-	1 (I30)	1 (I29)	-	2
Health Services (non-profit) (1)	-	-	-	1 (I19)	1
IT Consulting Firm (3)	-	-	1 (I27)	2 (I23, I26)	3
Total Role	5	2	11	14	32

The semi-structured interviews were audio recorded, transcribed, and anonymized. Our participants answered questions about various aspects of IT security, including their tasks, the tools they use, and the communications they perform to do their job (the interview guide can be found in (Werlinger et al. 2009)). 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 not all topics were discussed at the same level of detail with all participants, due to the nature of semi-structured interviews.

3.2.2. Participatory Observation

While the interviews gave the general context of the tasks performed by security practitioners, we used the ethnographic technique of participatory observation (Fetterman 1998) to provide fine-grained data on the two activities where the observer was involved. The participant observer, a security specialist with four years of experience as a security consultant in a large telecommunications organization, spent 78 hours working under the supervision of a senior IT security professional at an academic organization in Canada. One of the observer's tasks was the development of policies: he 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 was the deployment of an intrusion detection system (IDS): he worked with two security specialists on the installation and configuration of an IDS in the internal network (see Werlinger et al. (2008c) for details). The participatory observation gave insight on issues that were not evident from the interviews (e.g., intensive use of threat analysis when writing policies, lack of customizable access control for the IDS). Consequently, the results of the 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.

The participant observer used a diary to make written notes during the participatory observation. These notes captured several aspects related to the activities performed by the observer, including the channels used by the senior security professional in all manner of communications (which frequently included the observer), the information exchanged during meetings, the issues that complicated the policy development process, and the deployment and configuration of the IDS. The observer was required to have not only the ability to record what was happening around him, but also technical skills, good communication skills, and a security background. The technical skills were essential to perform the role of participant observer. When the IDS was to be deployed in the network, the participant observer's responsibility was to configure the IDS and identify the technical issues. Also, the security background was necessary in order to understand the security issues that arose during policy meetings. Without this background, it would have been very difficult to follow the discussions and contribute to the process of writing the policies. The security background was also necessary to convince managers that the participant would benefit the organization, which enabled them to agree to the observation. Although skill and experience were necessary in to perform the participatory observation, in some cases where the observer could influence the decision making process, the exercise of that experience posed the danger of interfering with the observed activities. In particular, the observer was experienced in developing policy, and had to be careful not to give recommendations that might obscure or damage valuable insights into the group's normal process.

3.3. Data Analysis

The audio-recorded interviews were transcribed, sanitized, then analyzed using qualitative description (Sandelowski 2000) with constant comparison and inductive analysis of the data. First of all, we 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 selective coding. The results were then organized according to 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. Analysis of the textual data (interviews and notes from participatory observation) was performed using Qualrus, a qualitative research tool. Further analysis was based on elaboration of "memos" (Charmaz 2006) that were written during the initial coding process. Interview questions were adjusted three times (before interviews 15, 22, and 27), in order to validate emerging theories.

For the overall HOT Admin project, the interview analysis consisted of two main stages. In the first stage, two researchers performed a general analysis

of interviews 1 to 14 to identify general information, including the workplace characteristics and tools used by security practitioners (Botta et al. 2007). In the second stage, five researchers analyzed the full interview corpus, focusing on six specific themes that had emerged during the first stage of analysis. Throughout the analysis, there were periodic meetings during which researchers shared their analysis of a specific theme and received feedback from the other team members. This approach helped mitigate the risk of overlooking connections in the data between the different themes, because the researchers were familiar with the data and each other's findings. Relationships among the themes under study were discussed, which allowed for triangulation of findings, as some themes had a considerable degree of overlap (e.g., interactions with other stakeholders and sources of errors in security management). Additionally, the first author of this paper was involved in the analysis of data for three of the six themes (responses to security incidents, challenges faced by security practitioners, and interactions among security practitioners and other stakeholders), which gave him a broad view of the data.

4. Analyzing Interactions in Context

From the initial analysis described in Section 3.3, a list of fifteen activities performed by security practitioners emerged (Botta et al. 2007). Some of these activities were performed individually – such as monitoring systems, or using documentation – and others required interactions with other stakeholders (e.g., responding to events). During the analysis presented in this paper, we focused on those activities that require interactions between stakeholders; we identified several descriptions in participant interviews of activities in which IT security-related communications occur. We next describe these activities these activities and communications, and then present the communication channels and security tools used by our participants for interacting with other stakeholders. To further illustrate our findings, we provide in Section 5 richer descriptions of the interactions, tools, and miscommunications involved in two of the activities: *security incident response* and *development of policies*.

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: that is, the participants required different strategies for communicating security issues to stakeholders with varying backgrounds and interests. That is, they needed to be sensitive to both the self-image of a stakeholder, and to the stakeholder's perception of the context of the message. By context, we mean the elements beyond the constituent parts of a communication, such as sender, receiver, message, and channel, with their respective attributes (Berlo 1960). (See Fouquier (1988) for a model of *meaning in the message*.) They needed to be proactive in their continuous establishment of mutual understanding in order to build the common ground for joint activities.

To perform security tasks, our participants had to cooperate, coordinate, and collaborate with other stakeholders. These interactions are distinguished by the level of commitment and intensity of the relationship (Winer and Ray 1994). We adapted definitions of *cooperation*, *coordination*, and *collaboration* from Matessich and Monsey (1992) and Winer and Ray (1994), and illustrate them here with examples from our data. Cooperation refers to shorter-term informal relations that exist without any clearly defined mission or structure (e.g., get help from an IT administrator in another country to shut down a phishing site). Coordination requires more formal relationships and understanding of compatible missions (e.g., send a log file to another IT administrator to get back information about the servers involved in a security incident), whereas collaboration is a mutually beneficial and well-defined relationship to achieve common goals (e.g., as a team, develop security policies for the organization). Our participants needed to share information, regardless of the nature of the higher level interactions that they performed.

The distribution of information in a group has been found to influence group judgment (Stasser and Titus 1985), but indirectly through the sharedness of member preferences (Gigone and Hastie 1993). In the same vein of "social sharedness" (Tindale and Kameda 2000), shared representations (i.e., "any task/situation relevant concept, norm, perspective, or cognitive process that is shared by most or all of the group members" (Tindale and Kameda 2000; p. 129)) can enable a minority to persuade the majority. Despite the importance of information sharing in IT security, it was not well supported by the security tools that they used. A similar issue was noted by Denning and Yaholkovsky (2008) for IT activities in general: IT tools are not designed to support formal cooperation, coordination, or collaboration. Whether or not IT and security tools should be redesigned as groupware is an open question; however, our findings indicate that these tools are often used within a group process that will require information sharing.

The three types of interactions (cooperation, coordination and collaboration) were often combined in our participants' duties, although 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 2 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 some cases, they provided more direct support and interacted actively with IT specialists to respond to recommendations provided by the auditor.

To *design services incorporating security requirements*, our participants needed

Table 2: Types of activity in which IT security communication occurs

Activity	**Interviews**		**Stakeholders involved**
	Academia	Industry	
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**

to specify security requirements for new IT services or projects. Design entails ongoing development of shared language (Walz et al. 1993). 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. That is, as consultants to committees, they would check how security requirements were incorporated in the changes that were handled by the committee. 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 needed to *solve end-user IT security issues* when they received notifications about users experiencing security issues with 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 controls 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. In this instance, there was a lack of a shared representation that would function as a "boundary object" (Star and Griesemer 1989).

Our participants also had to *train and educate* other stakeholders on security issues in a variety of circumstances, such as training new employees on the organization's privacy procedures. Participants also had to educate themselves. Participatory observation revealed that security practitioners learn through discussion with each other and establish common ground about how to design and implement security policies. Sometimes during meetings they had to look online for information about new types of attacks or features from security tools. We describe this activity in more detail in Section 5.2, in the context of developing security policies.

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 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 for this distribution of responsibilities. First, "network people" did not manage the firewall policies for controlling 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 reconfigure and administer them.

The remaining two activities are described briefly here, but will be presented in full in Section 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 organizational units.

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 and security tools) used by security practitioners to interact with other stakeholders.

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.

All participants 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 off-site 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 arose 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 a 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 projects' files (I30).

Besides e-mail, other tools like text or video chat were used by at least four participants. Again, perceptions of the usefulness of these tools varied. Three participants (I9, I10, I11) found that text chat was a good tool for getting an immediate response and for asking about specific information (e.g., a system's command syntax), while one participant (I8) felt that the language used in chat was awkward. Although chat is often used, one participant (I11) felt that it was the most error prone form of communication as it was difficult to convey information in a single line. Video chat was preferred because it complemented the advantages of text chat with images. One participant who used video chat (I9) commented that some colleagues did not use it as 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 in-

teract with other stakeholders. Face-to-face communications was felt to allow them to to interact quickly and avoid misunderstandings (I2, I14). 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 for keeping a record of what was discussed. When the electronic recording option was not available, 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.

4.3. Security Tools Used within the Context of Interactions

To generate security reports, our participants used tools like Nessus (I9, I12, I23, I25), a tool used to scan an IT infrastructure for vulnerabilities, and McAfee 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's 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 details that would help them 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 (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. Usability issues with encryption systems have been pointed out by the research community before; see for example Whitten and Tygar (1999) and Garfinkel and Miller (2005).

5. Interaction Scenarios

We used communication flow diagrams (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 two scenarios provide a reference for the environment in which security tools should be tested (Redish 2007). For example, a security tool intended to support these scenarios, as shown in Figures 1 and 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. The next sections describe in detail the interplay of interactions, use of resources, and the role of misunderstandings in these two scenarios.

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 descriptions given 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 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 experience 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.

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, which, in this case, were 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 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 that was 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 for an end-user was slow. 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 stakeholders. 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 1 shows the general case, Figure 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. Because 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 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 because the organizations that had detected attacks from the clients' IP addresses were blocking not only those addresses,

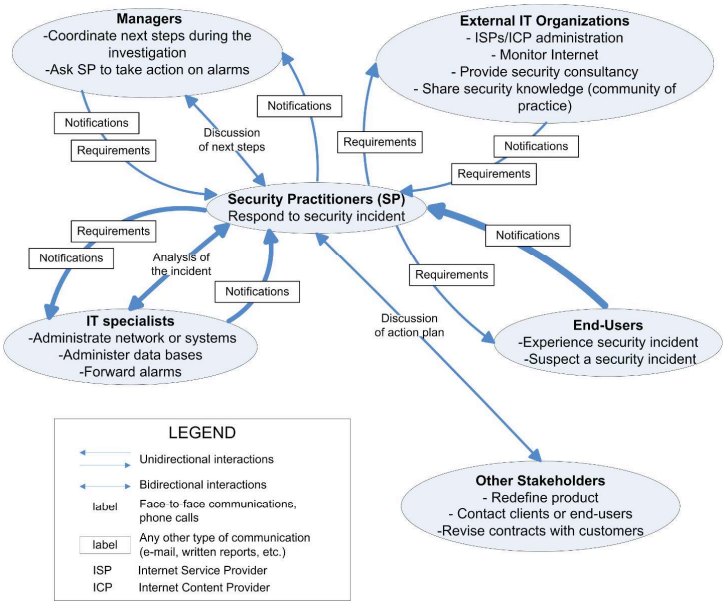

Figure 1: Responding to security incidents. Thicker arrows indicate more frequent interactions. For simplicity, only interactions between security practitioners and other stakeholders are shown.

but also the neighboring addresses within the same network 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 Microsoft 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 layer had to analyze the malicious software and generate a plan to clean and patch the infected machines.

The above examples show how the need to coordinate and respond to requirements from multiple stakeholders might make it necessary to define new procedures that establish formal responsibilities for the various stakeholders involved. For example, one participant (I29) mentioned how the incident illus-

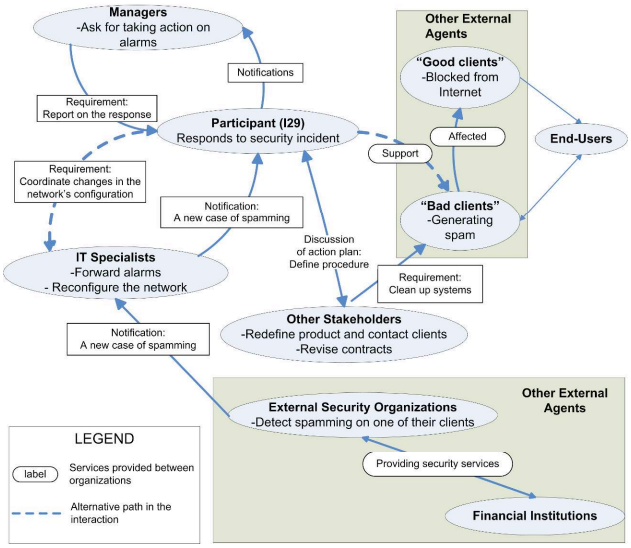

Figure 2: Response to an incident that triggers multiple and complex interactions among stakeholders. Dashed lines indicate two possible actions depending on the cooperation from the client. End-users, who are clients of the participant's organization, are behind other, external agents.

trated in Figure 2 triggered a revision of not only the interactions among 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 the participant's team were able to disconnect clients who 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 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 con-

stantly sending clarification questions through e-mail to avoid misunderstandings.

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. Such a shared representation has to satisfy more than one set of concerns. In the words of Star and Griesemer (1989), "This resolution does not mean consensus. Rather, representations, or inscriptions, contain at every stage the traces of multiple viewpoints, translations and incomplete battles." (p. 413). 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 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. Additionally, data obtained from participatory observation showed that *threat analysis* and *tacit knowledge* (further defined and discussed in Section 6.3.1) about the organization were also important in interactions regarding security policy development. The following results are based 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. 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 in Figure 3.

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. The involved IT specialists iteratively developed the policy text

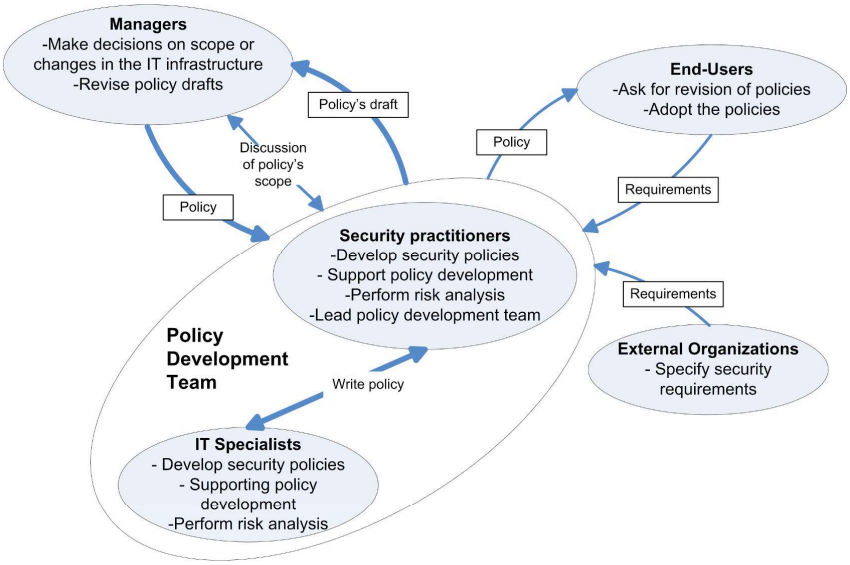

Figure 3: Communication flow diagram for developing security policies. Thicker arrows indicate more frequent interactions. For simplicity, only interactions with security practitioners are shown.

while considering how the tools were able to support the intended controls. During meetings, or individually, they looked for information about some security tools. 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 providing accessible and clear documentation about what security tools can and cannot do.

The group members that we worked with and observed tried to avoid misunderstandings with managers – that is, maintained mutual understanding toward building the common ground of policy – by continually asking for their feedback on, for example, the topics covered by the policies. This practice was necessary because a previous attempt at writing policies had failed because the policies proposed did not meet the expectations of managers.

6. Modeling the Complexity of Interactions

The two scenarios described in Section 5 illustrate the richness and complexity of interactions performed by security practitioners. This section proposes a model that integrates our findings and presents the factors that determine the

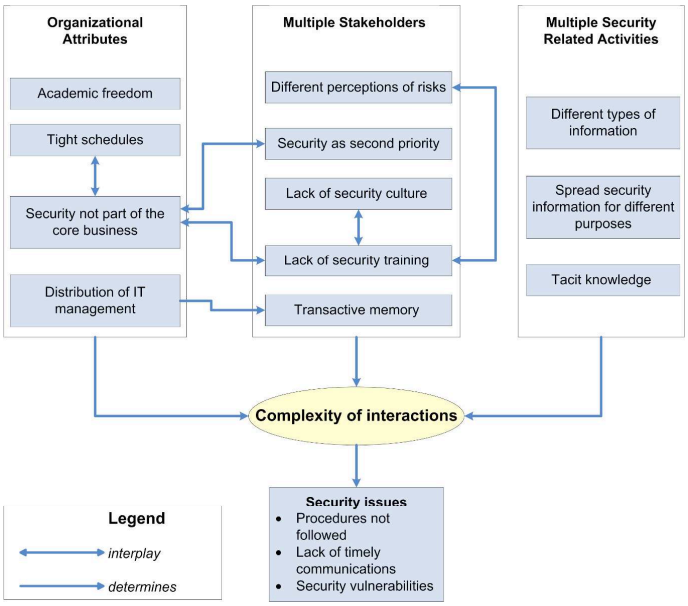

Figure 4: Factors that make interactions more complex for security practitioners within organizations.

complexity of interactions required to perform security tasks. Shown in Figure 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 grouping 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. Future work is needed to validate these categories and connections.

This model can be used to explain the complex interactions that security practitioners face when performing their activities (see Section 4). For example, when designing services with security requirements, 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 shows that the complexity of interactions for security practitioners causes security issues that altogether increase an organization's vulnerability and risk.

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.

6.1. Organizational Attributes

Tight schedules made interactions more complex for the security practitioners we interviewed. 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 (e.g., 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 (2007) 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 could 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 (2007) 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.

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 among 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 a second priority*, and did not have a *security culture or training*. These characteristics combine to determine what Flechais and Sasse (2007) identify as motivation of the stakeholder. Our participants constantly had to persuade other stakeholders who each 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 cooperation. 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.

The need for interaction with multiple, internal stakeholders to perform tasks is a characteristic that was also mentioned for some activities performed by other types of IT practitioners (see Barrett et al. (2004)). These IT practitioners, however, did not report any interactions with external stakeholders. Our analysis suggests that security practitioners not only need to interact frequently with internal stakeholders, but also with external ones (e.g., IT specialists who work in other organizations). For example, to respond to incidents, security practitioners received notifications from external organizations about malicious traffic generated locally, had to ask for support to shut down phishing sites, and had to provide support to clean up systems in other organizations (see Section 5.1). It appears that this kind of interaction is less common for other IT practitioners,

who usually have to respond to events that involve stakeholders and systems within the organization. Further study will be required to validate and refine this finding.

Interactions with external agents during security incidents pose another challenge in the work of security practitioners: they are in a scenario that is usually unregulated and has no standards to share information. Further, organizations involved in security incidents sometimes provide IT services to each other, without contractual specifications on how to respond jointly to malicious events (e.g., shut down connections from one organization in case of a major compromise). This characteristic means that security practitioners not only need to know about a wide variety of systems in the organization (Haber and Kandogan 2007), but they also need to understand the technical and legal relationships that the organizations establish with each other. Other legal aspects in the work of security practitioners are mentioned by Gagné et al. (2008). This seems to be another trait of IT security practice, distinguishable from other IT practices.

6.3. Multiple Security-Related Activities

Multitasking is a characteristic that was noted by Barrett et al. (2004) in the work performed by some IT practitioners (web and database administrators). Our results also 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 a 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.

6.3.1. Tacit Knowledge

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, and how these stakeholders would adapt security practices in their tasks.

A great deal of knowledge in the practice of IT security has to do with idiosyncratic situations, in which the knowledge can *in principle* be articulated, but to do so during practice would be impractical. Furthermore, much of this knowledge is internalized to the extent that practitioners can quickly trouble shoot in very complex environments. To refer to this knowledge as being merely undocumented falls short of conveying this important aspect of IT security practice. The following lays out how tacit knowledge can include knowledge that is articulable.

Concerning tacit knowledge, Polanyi (1966) endeavors to explain what the phrase "we can know more than we can tell" means. For example, we can recognize a familiar face in a crowd, but cannot exhaustively tell how we recognize it. We achieve "an integration of particulars to a coherent entity to which we are attending" (Polanyi 1966; p. 18). Polanyi regards the integration of particulars as an *interiorization*, so that "instead of observing them [the particulars] in themselves, we may be aware of them in their bearing on the comprehensive entity which they constitute" (p. 18). Further, as a blind man with a cane attends to the world of objects that are revealed by the touch of the cane (as opposed to tending to the cane itself), we extend our bodies to include our instruments. Polanyi points out that intense scrutiny of the particulars will efface their meaning within the comprehensive entity. But, "the destruction can be made good by interiorizing the particulars once more. The word uttered again in its proper context [...] and the details of a pattern glanced at once more from a distance: they all come to life and recover their meaning and their comprehensive relationship" (pp. 18–19). Altogether, "tacit knowing is shown to account (1) for a valid knowledge of a problem, (2) for the scientist's capacity to pursue it, guided by his sense of approaching the solution, and (3) for a valid anticipation of the yet indeterminate implications of the discovery arrived at in the end" (p. 24).

Notably, the practice of IT security includes an intense scrutiny of the particulars, which seems to belie the aspect of *tacit knowing* of not being able to articulate *how* one knows. But with ongoing daily experience with an application of technology within an idiosyncratic organization, this "destruction . . . [is] made good by interiorizing the particulars once more." Thus *tacit knowing* accounts for how security professionals recognize valid problems, have a sense of how to approach solutions, and have an anticipation of the implications. In any case, Nelson and Winter (1982) emphasize that cost matters. They say, "Whether a particular bit of knowledge is *in principle* articulable or necessarily tacit is not the relevant question in most behavioral situations. Rather, the question is whether the costs associated with the obstacles to articulation are sufficiently high so that the knowledge *in fact* remains tacit" (p. 82).

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 who did not have a security background. Not following security procedures generated communication overhead. For example, one participant (I2) highlighted the consequences of not following a change-management procedure aimed 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 that had already been 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) cited 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. 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.

Ineffective interactions might lead to vulnerabilities in organizations. In the context of security, vulnerabilities mean that the level of exposure to attacks is higher. Attacks can come in the form of autonomous software (e.g., worms, viruses) or human-directed intrusions that compromise the internal systems (Kandogan and Haber 2005). Human-directed intrusions are carried out by attackers (also known as adversaries), a special type of stakeholder who competes with security practitioners, looking for attack opportunities (Haber and Kandogan 2007). Although our analysis was not focused on this type of interaction, our data show that most of the tasks and activities described by participants were performed to prevent or respond to attacks (e.g., deploy a firewall to avoid attacks (I2), mitigate vulnerabilities to stop password guessing

attacks (I12), develop a new procedure to respond to an incident (I30)). Security professionals indirectly interact with adversaries, so these attackers should be included amongst the relevant stakeholders. The presence of adversaries impacts the activities and interactions performed by security practitioners, who need to quickly mitigate vulnerabilities to stop active and continuous threats from attackers (Gagné et al. 2008). Security practitioners also have to keep information about vulnerabilities in their systems secret, in order to avoid attackers gaining knowledge about them (Haber and Kandogan 2007). Finally, once an attack successfully exploits a vulnerability, it is very difficult to know the ultimate consequences of the compromise. For example, if a secret *might* have been exposed, then it must be treated as having been *actually* exposed (Whitten and Tygar 1999). All these factors mean that it is critical for security practitioners to avoid interaction breakdowns that might lead to security incidents.

6.5. Revisiting Communication Models

Our model describes components that are also present in communication models (e.g., Berlo (1960); Benbasat et al. (1995)). We next relate some of the factors presented in our model with the ones proposed by Berlo's communication model to understand how our findings fit in the communication process. To begin with, concerning Berlo's first main category – *participants' characteristics*– both our model and Berlo's model indicate *knowledge* as a factor. That is, with security related activities, participants associated knowledge with security training and also with the ability to perform risk assessment. *Culture* is also present in both models. Some participants mentioned that a lack of security culture in the organization made it difficult to persuade other stakeholders, who usually had different motivations, to apply security practices in their day-to-day activities. Some other factors that Berlo includes, such as *skills*, *attitudes*, and *social system*, do not explicitly appear in our model, nevertheless they can be related to our findings. For example, our participants mentioned that *good communication skills* are important in order to interact with stakeholders whose *attitude* to security-related matters was not always positive. This attitude came from the perception that security controls were an impediment to perform business (Gagné et al. 2008). The *social system* can be related to the characteristics of the organization and its security culture.

While Berlo's main categories *message* and *channel* are not explicitly represented in our complexity model; nevertheless, they are critical for security related communications. With respect to *message*, our participants indicated that they had to be very careful about the content and the language of messages communicated to different stakeholders. The way in which participants communicated security issues was key; they needed to persuade other stakeholders to adopt security practices. To do so, they had to decide on (1) the specific words or symbols they were using in the message (Berlo's *code*), (2) what parts of the content they wanted to emphasize using those words or symbols (Berlo's *message treatment* and *message structure*), and (3) the format (Berlo's *message structure*). With respect to *channel*, our participants combined channels of *hearing* (e.g., phone calls) and *seeing* (e.g., whiteboard), depending on their

needs. Technology enables combinations of channels (Benbasat et al. 1995). We elaborate on better communication features for security tools in Section 7.

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 intrusion detection systems. 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 that do not always cover all their needs. For example, our participants needed to avoid the possibility of misunderstandings during communications while keeping track of agreements for future audits.

We next offer 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: Integration of IT tools is a general guideline mentioned by Haber and Bailey (2007). They suggest that IT tools should provide standard reporting to work together with other components and application programming interfaces (APIs) to be integrated with monitoring and management tools. Our analysis showed that security practitioners had to send and receive notifications, reports, and requirements (see Figures 1 and 3) to communicate with different stakeholders. In this vein, 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 a response to a security incident, as illustrated in Figure 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 (and reducing the corresponding risk of error) of copy-pasting outputs from communication to security tools and vice versa.

Security tools should also 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. This guideline is related to the need for sharing views mentioned by Haber and Bailey (2007) in the context of IT tools in general. In the case of security tools, this type of feature can lead not only to the reduction of the communication overhead (as we explain in the next guideline), but also to better consolidation of information buffers (Pinelle and Gutwin 2003). Consolidation of buffers refers to bringing together information

that is fragmented across multiple locations, and to making it visible to other team members. This is an important design guideline for communication tools used by workers that are mobile, widely dispersed, autonomous, and who communicate with each other only intermittently (Pinelle and Gutwin 2003). Our participants exhibited a similar type of behavior when they had to interact with different, independent stakeholders depending on the activity they performed (e.g., response to incidents, as described in Section 5.1).

Reduce communication overhead: Similar to the previous point, security practitioners need tool features in order 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. In 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 for communicating and sharing 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 confidentiality protection (e.g., through data encryption), which should be embedded in security tools that produce reports. For example, in the context of integration of IT tools mentioned by Haber and Bailey (2007), a security tool that generates reports about virus activity should be capable of easily integrating with VPN clients. This integration may mitigate the poor practice of sending sensitive information to external stakeholders without the required protection.

Provide customizable accounts for stakeholders with different goals: We have discussed how security tools can provide better support by sharing views with other stakeholders. Additionally, our participants sometimes required a higher level of flexibility to share not only views in the systems, but also different levels of access to other stakeholders that are using the same tool to collaborate in the same security-related activity. Haber and Bailey (2007) point to the need for IT tools to provide access to sysadmins responsible for different components. Our analysis also showed that, 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), but they also need to provide different configuration options for improving collaboration among IT practitioners from different IT areas. For example, an IDS should have the capability to configure various accounts to monitor different networks or systems independently. To separate these networks or systems, multiple differentiation criteria should be supported: IP addresses, type of operating system used (e.g., Windows, Linux, Mac), or type of network protocol.

Provide reporting options that show the level of risk: Another opportunity for improving security reporting is to provide security practitioners

with better features to interpret and communicate the information from the analyses that security tools perform. In this vein, one potential feature is to show the outputs of security tools in the context of the level of security risk of the organization (e.g., how important the vulnerabilities or alarms are in the general context). However, because of its complexity, finding methods of measuring the level of security risk in organizations is an issue that is being actively discussed within the security community (e.g., securitymetrics.org). Another option for security tools to generate reports that indicate conditions of risk in specific domains of the IT infrastructure. These conditions of risk could include information about the status of security patches and antivirus updates in the servers, and could be correlated with other conditions of risk such us vulnerabilities in network devices or current attacks in the public networks. 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, such as 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: online and automatic generation of different reports for different stakeholders, and the use of different layers of information (general vs. specific).

Correlate data with sources external to IT databases: The need to be able to address new security incident scenarios (see Figure 2) makes it necessary to correlate information in novel ways. For example, in determining who caused a security incident, an IP address may be the only handy information that a security specialist has to go on. Then the specialist has to correlate the IP address 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 individually receive the information at his or her appropriate level of detail, language, and channel (e.g., e-mail, text message, web site). Furthermore, security tool developers 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 letting attacks pass through the internal network, but also to notify the corresponding administrator via e-mail to take action and stop the malicious traffic.

Manage tacit knowledge: Our participants managed tacit knowledge when they (1) provided statements of evaluation while 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 knowl-

edge into explicit knowledge. There is some debate as to whether or not such a thing is feasible (Schmidt 1997), or desirable – why should practitioners give away their stock-in-trade?. Flechais and Sasse (2007) propose the use of scenarios as a more effective tool for helping clearly explain abstract security concepts to other stakeholders. Our findings suggest 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 employed 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 better advantage of this process by involving other stakeholders in it. This process could entail the use of scenarios or anecdotes, as suggested by Flechais and Sasse (2007).

8. Limitations and Future Work

While our research approach allowed us to investigate in detail the interactions that security practitioners have with other stakeholders within the context of the security activities, this approach was not without limitations. Both the semi-structured interviews and participatory observation provided us with rich data about the activities and interactions performed by security practitioners. While rich, this data is limited to a relatively 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 interactions performed by participants. Rather, our findings are centered on the commonalities in their descriptions.

The variety of organizations whose employees we interviewed, in terms of industrial sectors (Table 1) and sizes (ranging from fewer than 5 employees to large, multinational companies), has given us a broad perspective about the activities and interactions that security practitioners perform in different settings. All our participants relied intensively on interactions with other stakeholders to perform their activities. Our analysis did not show significant differences among participants concerning the need to involve multiple stakeholders in performing security-related tasks. Nevertheless, this variety of organizations also represents a limitation of our study. Prior research has found that organizational factors such as the size of the organization and the security management model in place may impact the practice of security within an organization (Chang and Ho 2006; Hawkey et al. 2008b; Kankanhalli et al. 2003). We lack data from a sufficient

number of organizations of the same size in each sector to perform analysis considering the effects of organizational factors on our results. For example, a first look at our results suggests that there are differences in communication preferences between academia and private sectors. All of our participants who stated that they preferred verbal communications came from academia. However, the descriptions provided by other participants from the private sector made it possible to infer the importance of verbal communication when performing security tasks. For example, one participant (I5) mentioned the importance of communicating security issues to other stakeholders in person.

Despite these limitations, we believe that our model (Section 6) further improves the understanding of the complexity of security-related interactions. It can be used to explain the interplay between different factors when security practitioners interact with other stakeholders to perform their activities, and it helps to illustrate the consequences of the corresponding complexity for security management. However, more data are needed to expand and refine the model. For example, validation of the model with a large number of participants from organizations of different sizes within each sector may extend the model to include other organizational attributes, such as type and size. Furthermore, variations of the model could be developed that consider 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 model into individual, smaller sub-models specific to each activity performed by the security practitioner.

Another limitation of our study is related to the nature of the data from semi-structured interviews, where not all the participants talk about the same themes with the same level of detail. This made it difficult to perform frequency analysis on the data. For example, using interview data, it is hard to determine how often our participants spent interacting with other stakeholders and how often they worked alone. Data from the participatory observation gives some indication of this. When developing security policies, the observer spent approximately 30% of his time in meetings with the IT specialists. (This number does not include sporadic interactions with the security specialist.) For the implementation of the IDS, the observer spent approximately 20% of his time in meetings. These numbers appear to be similar to the 23% of time that an IT administrator was observed to spend in meetings (Barrett et al. 2004). However, in order to contrast interaction frequencies, more data are necessary to validate this observation and extend it to different activities. Furthermore, longitudinal data would be needed to capture the variations on the activities due to such things as unexpected events, interruptions, and the stage of the activities (initialization vs. closing). Another type of frequency analysis limited by the nature of our data is the relative frequency of each activity performed by security practitioners. Table 2 suggests that response to incidents is perhaps more frequent than the other activities performed. However, such incidents may also be more memorable and more easily recalled during the interview process. Similarly, most of the participants who discussed security incidence response were from academia, which also might suggest that security incidents are more

common in academic environments. Again, more research in this direction is required before drawing conclusions.

As discussed in Section 3.1, we do have estimates from sixteen of the participants about the amount of time spent on security-related activities. Security specialists spent on average 67% on security activities whereas IT specialists spent on average 30%. The three participants who spent the most time on IT security (76%–100%) performed activities such as administering security devices (I24), implementing access control policies (I21), and developing security policies (I24, I25). The three participants who spent the least time on security activities were IT managers (I15, I16, I17). When we examined the interviews, we did not find patterns that suggested differences in the interactions that each group of participants performed during security activities. This result might be explained by the nature of our questions during the interviews, which were focused on security related activities rather than general IT tasks. However, interview data did provide some information on potential differences between the work of security practitioners and the work of other IT practitioners (see Section 6). These results are based on the interactions that are performed in response to incidents and in the development of policies. The activities demonstrate characteristics that seem to typify security practice (e.g., Barrett et al. (2004)). Future work is necessary to better compare security vs. other IT practices, in terms of variables such as *frequency of interactions with external agents*.

Our analysis showed that some activities were more time-sensitive than others. For example, our participants mentioned that the development cycle for policies is very long (e.g., I1, I2), in contrast to the activities with shorter cycles, such as recovering systems when responding to incidents. This difference might impact the use of different tools for interactions depending on the activity; usually participants mentioned the use of communication tools that provided immediate response (e.g., text chat) in the context of troubleshooting and responding to incidents.This observation contrasts with the report of tools used used during participatory observation in the development of policies (e.g., e-mail, Sharepoint to share documents), where immediate response was not necessary. More data could provide greater understanding of how understanding how tools are affected by the type of security-related interaction.

We elaborated in Section 6 on how our complexity model relates to the communication model proposed by Berlo (1960). Recent studies have also proposed communication models for activities that require complex interactions. For example, Keyton et al. (2008) investigate how communication supports collaboration among stakeholders from different organizations. Their model links structural components (e.g., antecedents used for the collaboration, organizations involved) with the collaboration process and its results. In building this model, they collected granular data over nine months, observing and videotaping six multidisciplinary teams who were collaborating to create business plans based on life science research projects. Similarly, our model in Section 6 shows different elements that impact the interactions of security practitioners. Further research is necessary to investigate how these elements can be treated

as variables whose values might determine the effectiveness of security-related interactions (e.g., the bigger the differences in security training, the less effective the outcome of the collaboration process). However, in order to understand the exact relationship between these variables and the final outcome of security-related interactions, it would be necessary to collect fine-grained data on how communications affect the different levels of interactions that security practitioners perform (i.e., coordination, cooperation, and collaboration). These data could be used to build a communication model to evaluate and anticipate the outcome of security-related interactions. Future work in this direction might focus on one of the activities performed by security practitioners. Our study suggests that development of policies and response to security incidents are candidates for further investigation. The former has the advantage of including longitudinal aspects of the organization; its main disadvantage is that policy development can be a lengthy process that would be difficult to observe from the beginning to the end. The latter has the advantage of including a variety of stakeholders within and outside of the organization; its main disadvantage is that it may be difficult to recruit participants or organizations that are willing to reveal how they are impacted by security breaches.

9. 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 of 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 for developing 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 which 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.

Acknowledgments

The authors would like to thank our participants for taking part in our study, members of the Laboratory for Education and Research in Secure Systems Engineering (LERSSE) for feedback on earlier versions of this paper, Craig Wilson and Tara Whalen for improving the readability of this paper, and the anonymous reviewers of our paper for their valuable comments. This work has been supported by the Canadian NSERC Strategic Partnership Program, grant STPGP 322192-05.

References

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

Benbasat, I., Lim, F. J., Rao, V. S., 1995. A framework for communication support in group work with special reference to negotiation systems. Group Decision and Negotiation 4 (2), 113–158.

Berlo, D. K., 1960. The Process of Communication. New York: Holt, Rinehart and Winston.

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

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

Björck, F. J., 2005. Discovering information security management. Doctoral thesis, Stockholm University, Royal Institute of Technology.

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

Carroll, J. M., Rosson, M. B., Convertino, G., Ganoe, C. H., 2006. Awareness and teamwork in computer-supported collaborations. Interact. Comput. 18 (1), 21–46.

Chang, S. E., Ho, C. B., 2006. Organizational factors to the effectiveness of implementing information security management. Information Management & Computer Security 106, 345–361.

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

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

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

Denning, P. J., Yaholkovsky, P., 2008. Getting to "we". Commun. ACM 51 (4), 19–24.

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

Flechais, I., Sasse, M. A., 2007. Stakeholder involvement, motivation, responsibility, communication: How to design usable security in e-science. Int. Journal of Human-Computer Studies doi:10.1016/j.ijhcs.2007.10.002.

Fouquier, E., 1988. Figures of reception: Concepts and rules for a semiotic analysis of mass media reception. International Journal of Research in Marketing 4 (4), 331–348.

Gagné, A., Muldner, K., Beznosov, K., July 8-9 2008. Identifying differences between security and other IT professionals: a qualitative analysis. In: HAISA'08: Human Aspects of Information Security and Assurance. Plymouth, England, pp. 69–80.

Garfinkel, S. L., Miller, R. C., July 2005. Johnny 2: A user test of key continuity management with S/MIME and Outlook Express. In: Proceedings of The Symposium on Usable Privacy and Security (SOUPS). ACM Press, Pittsburgh, Penn. USA.

Gigone, D., Hastie, R., 1993. The common knowledge effect: Information sharing and group judgment. Journal of Personality and Social Psychology 65, 959–974.

Glaser, B., Strauss, A. L., 1967. The Discovery of Grounded Theory, Strategies for Qualitative Research. Aldine Publishing Company, Chicago, Illinois.

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

Guinan, P., Bostrom, R. P., 1986. Development of computer-based information systems: A communication framework. SIGMIS Database 17 (3), 3–16.

Haber, E., Kandogan, E., 2007. Security administrators: A breed apart. In: Workshop on Usable IT Security Management, (USM'07) held with the ACM Symposium on Usable Privacy and Security (SOUPS).

Haber, E. M., Bailey, J., 2007. Design Guidelines for System Administration: Tools Developed through Ethnographic Field Studies. In: CHIMIT '07: Proceedings of the 2007 symposium on Computer Human Interaction for the Management of Information Technology. ACM, pp. 1–9.

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

Hawkey, K., Muldner, K., Beznosov, K., 2008b. Searching for the Right Fit: Balancing IT Security Model Trade-offs. Special Issue on Useful Computer Security, IEEE Internet Computing, 30–38.

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

Kandogan, E., Haber, E. M., 2005. Security administration tools and practices. In: Cranor, L. F., Garfinkel, S. (Eds.), Security and Usability: Designing Secure Systems that People Can Use. O'Reilly Media, Inc., Sebastapol, Ch. 18, pp. 357–378.

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

Kesh, S., Ratnasingam, P., 2007. A knowledge architecture for IT security. Communications of the ACM 50 (7), 103–108.

Keyton, J., Ford, D. J., l. Smith, F., 2008. A mesolevel communicative model of collaboration. Communication Theory 18, 376–406.

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

Koskosas, I. V., Paul, R. J., October 2004. The interrelationship and effect of culture and risk communication in setting internet banking security goals. In: 6th international conference on Electronic commerce (ICEC). ACM Press, pp. 341–350.

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

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

Lloyd, L. L., Quist, R. W., Windsor, J., 1990. A proposed augmentative and alternative communication model. Augmentative and Alternative Communication 6 (3), 172–183.

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

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

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

Neale, D. C., Carroll, J. M., Rosson, M. B., 2004. Evaluating computer–supported cooperative work: models and frameworks. In: CSCW '04. ACM Press, pp. 112–121.

Nelson, R., Winter, S., 1982. An Evolutionary Theory of Economic Change. Harvard University Press.

Pinelle, D., Gutwin, C., 2003. Designing for loose coupling in mobile groups. In: GROUP '03: Proceedings of the 2003 international ACM SIGGROUP conference on Supporting group work. ACM, New York, NY, USA, pp. 75–84.

Polanyi, M., 1966. The Tacit Dimension. Doubleday & Company, Inc., Garden City, New York.

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

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

Sanders, D. A., 1976. A model for communication. In L. L. Lloyd (Ed.), Communication assessment and intervention strategies. Baltimore: University Park Press.

Schmidt, K., November 1997. Of maps and scripts—the status of formal constructs in cooperative work. In: ACM SIGGROUP. pp. 138–147.

Shannon, C., Weaver, W., 1949. The Mathematical Theory of Communication. Urbana, IL: University of Illinois Press.

Star, S. L., Griesemer, J. R., 1989. Institutional Ecology,Translations and Boundary Objects: Amateurs and Professionals in Berkeley's Museum of Vertebrate Zoology, 1907-39. Social Studies of Science 19 (3), 387.

Stasser, G., Titus, W., 1985. Pooling of unshared information in group decision making: Biased information sampling during discussion. Journal of Personality and Social Psychology 48 (6), 1467–1478.

Tindale, R. S., Kameda, T., 2000. 'social sharedness' as a unifying theme for information processing in groups. Group Processes and Intergroup Relations 3 (2), 123–140.

Trevino, L. K., Webster, J., Stein, E. W., 2000. Making connections: Complementary influences on communication media choices, attitudes, and use. Organization Science 11 (2), 163–182.

Vicente, K., 2002. Ecological interface design: Progress and challenges. Human factors 44 (1), 62–78.

Vicente, K., Rasmussen, J., Jul/Aug 1992. Ecological interface design: theoretical foundations. Systems, Man and Cybernetics, IEEE Transactions on 22 (4), 589–606.

Vicente, K. J. ., 1999. Cognitive Work Analysis: Toward Safe, Productive, and Healthy Computer-Based Work. Mahwah, NJ: Lawrence Erlbaum Associates, Publishers.

Walz, D., Elam, J., Curtis, B., 1993. Inside a software design team: knowledge acquisition, sharing, and integration. Communications of the ACM 36 (10), 63–77.

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

Werlinger, R., Hawkey, K., Beznosov, K., July 2008a. Human, Organizational and Technological Challenges of Implementing IT Security in Organizations. In: HAISA'08: Human Aspects of Information Security and Assurance (10 pages).

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

Werlinger, R., Hawkey, K., Beznosov, K., January 2009. Auxiliary material for the study of security practitioners in context: Their activities and interactions with other stakeholders within organizations. Tech. Rep. LERSSE-TR-2009-01, Laboratory for Education and Research in Secure Systems Engineering, University of British Columbia.
URL `http://lersse-dl.ece.ubc.ca/search.py?recid=168`

Werlinger, R., Hawkey, K., Muldner, K., Jaferian, P., Beznosov, K., July 23-25 2008c. The challenges of using an intrusion detection system: Is it worth the effort? In: Proceedings of the Symposium On Usable Privacy and Security (SOUPS). Pittsburgh, Pennsylvania, pp. 107–116.

Whitten, A., Tygar, J., 1999. Why Johnny can't encrypt: A usability evaluation of PGP 5.0. In: The 9th USENIX Security Symposium. pp. 169–183.

Winer, M., Ray, K., 1994. Collaboration Handbook: Creating, Sustaining, and Enjoying the Journey, 5th Edition. Amherst H. Wilder Foundation, Saint Paul MN.