

Extracting Useful Information from Security Assessment Interviews

Jeffrey M. Stanton
 Syracuse University
jmstanto@syr.edu

Isabelle J. Fagnot
 Syracuse University
ifagnot@syr.edu

Abstract

We conducted N=68 interviews with managers, employees, and information technologists in the course of conducting security assessments of 15 small- and medium-sized organizations. Assessment interviews provide a rich source of information about the security culture and norms of an organization; this information can complement and contextualize the traditional sources of security assessment data, which generally focus on the technical infrastructure of the organization. In this paper we began the process of systematizing audit interview data through the development of a closed vocabulary pertaining to security beliefs. We used a ground-up approach to develop a list of subjects, verbs, objects, and relationships among them that emerged from the audit interviews. We discuss implications for improving the processes and outcomes of security auditing.

1. Introduction

The Information Systems Audit and Control Association (ISACA) defines an audit as, “the process of generating, recording and reviewing a chronological record of system events to ascertain their accuracy” (see www.isaca.org). The apparent simplicity of this definition hides the messy reality of conducting comprehensive security assessments in the context of intact organizations. Despite their messiness and complexity, however, security assessments provide an important mainstay of positive security practice for information security professionals. Periodic, comprehensive analysis of an organization’s methods and measures for information protection can help to ensure that the organization’s investments in technology and human expertise are having their desired affect on the organization’s security status. An old management saw states that, “you can’t control what you can’t measure.” In the world of security, one of the ways that security managers and professionals measure how their programs are working is by conducting periodic security assessments. In the most effective

organizations, such assessments generally seem to comprise a comprehensive review of risks, threats, vulnerabilities, technical controls, governance, and a host of closely related issues. Assessments often take a period of months to complete and may frequently involve some of the organization’s most expert and most powerful personnel. The end result of such assessments is usually a set of reports that make recommendations on how an organization should proceed with future investments and security program modifications.

The term “state of the art” applies aptly to security assessments, because despite the wide range of possible techniques, security assessment remains much more of an art than a science. Within this art, two main schools exist. Quantitative assessment focuses on capturing an empirical, reproducible, and objective stream of data on the performance of a set of information systems (or even of larger, sociotechnical systems). To take a simple example, a vulnerability scanning tool could be used to scan a range of network addresses and count up the number of existing vulnerabilities in a set of hosts. The average number of vulnerabilities per host, the total number of critical vulnerabilities, the maximum number of vulnerabilities and a variety of other statistics numerically represent the readiness of (some of) the organization’s information systems to fend off certain types of attacks. Such quantifications are particularly valuable when assessing performance over time, because they provide a method of comparing apples to apples. The use of common, generally accepted metrics in auditing processes helps to ensure that the data will be comparable to those collected in the past and any that might be collected in the future.

One shortcoming of such quantitative methods is their narrowness. Knowing the average number of vulnerabilities per host tells us something important about where we are, but tells us little about how we got here or where we might be going next. To obtain a richer picture of what is happening in the organization with respect to security, a different school of assessors

relies on a more descriptive or qualitative strategy for analysis. Through in depth questioning of the organization's information security professionals, the assessors learn about the organization's information systems architecture, formal security policies, typical security practices, weaknesses in policies and practices, as well as the preparedness of the staff, the supportiveness of management, and the competence of outsource providers. To an increasing extent, these qualitative assessments reach out beyond the information technology department and include analysis of regulatory constraints, human resource practices, end user attitudes, and a host of other sources of information. In this paper we begin an exploration of strategies for making sense out of one particular element in this wide-ranging data. Using a "ground-up" approach, we have analyzed a large body of security assessment interviews with an eye towards systematic analysis and use of the data contained therein. We imagine a future in which recordings of a set of interviews with technical support staff, end users, human resource managers, or executive managers might be subjected to automated processing that would reveal the prevalent security culture, commonly held beliefs about the importance of controls, and behavioral norms for the protection of information. For the present, that goal is a long way off, but we believe the first steps along the way involve obtaining a comprehensive overview of what people in the organization actually say when they are questioned about information security. We need to know whether what they say is relevant, and we need to learn how to connect what they say to the overall security outcomes of the organization. To begin the process of addressing these questions, we transcribed a complete set of security assessment interviews from assessments conducted in 15 organizations and we processed them at a detailed level to extract the essential security-related vocabulary and relationships. We believe that the results of this analysis begin to give some key insights into the next steps needed to make better use of stakeholder interviews in the overall security assessment process.

2. Literature review

Before going more deeply into topics surrounding security assessment interviews, we need to set the context surrounding qualitative security assessment. Researchers choose methods according to their philosophical perspectives (e.g., positivist, interpretive, or critical). Positivists regard the material world as the only reality and consider their view as more objective because they rely on patterns that can be discovered accurately through observation.

Interpretive researchers seek "to understand phenomena through the meanings that people assign to them" (Myers & Walsham [12]). Each person interprets differently the reality of what s/he sees, feels, and comprehends. Interpretive research methods in Information Systems are "aimed at producing an understanding of the context of the information system, and the process whereby the information system influences and is influenced by the context" (Walsham [19]). For critical researchers current social reality is to a great extent a product of history. As Myers [11] stated, this is partly due to the fact that people's ability to change their social reality "is constrained by various forms of social, cultural and political domination." Critical researchers concentrate on studying the current diverging opinions and disagreements with the intention of reducing their grounds.

Qualitative research methods are frequently used to examine social and cultural phenomena under either the interpretive or the critical research perspectives. Qualitative research methods can be grouped under four major categories: action research, case study research, ethnography, and grounded theory (Myers [11]). These categories regroup various techniques that researchers can use to collect empirical data for their studies. These data collection techniques include focus groups, observation and participant observation, interviews (structured, semi-structured, or unstructured), questionnaires (generally open ended), various documents (memos, emails, etc), case studies, and the researcher's impressions and reactions. A frequent way of analyzing qualitative data is by coding it. Coding data involves several steps. Once the researcher is done preparing the data set for analysis (collecting, transcribing, organizing), s/he searches for concepts and themes, those emerging from the data and those suggested by the literature. These codes and themes are at the basis of a coding scheme later utilized by the researcher to code the data set and eventually analyze it. Coding processes are by their nature subjective. As Rubin & Rubin [15] stated, "Though the analysis is based on the descriptions presented by the interviewees, the interpretations in the final reports are those of the researcher". In this paper, we present a method that could eventually provide an analysis of security assessments relying less on the researchers' interpretations and more on the data itself.

The proposed method of analysis differs from and could complement existing methods of data analysis such as content analysis and text analysis. Neuendorf [13] defined content analysis as "the systematic, objective, quantitative analysis of message characteristics." Our method aims at transforming

qualitative data into a more quantifiable data set on which to perform statistical analyses. Further, a difference between text analysis and the development and use of controlled vocabulary lies in the fact that text analysis uses a dictionary which “is a set of word, phrases, parts of speech, or other word-based indicators (e.g., word length, number of syllables) that is used as the basis for a search of texts” (Neuendorf [13]). Controlled vocabulary, in contrast, is both built from and customized to the needs and specifics of a particular domain of inquiry – in our case the area of security assessment.

Audit vs. Assessment. To begin our review of the relevant issues surrounding security assessment interviews, we need to mention that the focus of information systems auditing remains sharply upon the verification of information system records. As we pointed out at the beginning of this paper, such verification is a narrow, though absolutely critical, component of the overall information security picture in an organization. *Security assessment*, in contrast, focuses on a holistic, overall evaluation of an organization’s capabilities for ensuring the protection of its information assets. The distinctions between auditing and assessment, then, lie both in the methods and the scope of the respective analyses: Auditing uses technical methods to substantiate verifiable information security metrics while assessment uses a mixture of technical and social methods to assess overall security readiness.

Unlike auditing, the roots of security assessment are in the discipline of program evaluation. Program evaluation is generally defined as the art or science of analyzing the success of an organized human activity, with an eye towards assisting future decision making about that activity. Contemporary program evaluation shares many similarities with other kinds of social science research: an investigator or a team of assessors conducts an evaluation by using a set of appropriate research methods and measuring a meaningful set of criteria. Such methods may include both quantitative measurements and qualitative data collection. One other distinction between audit and assessment emerges here: program evaluators would be likely to examine social phenomena such as attitudes and beliefs whereas auditors would be unlikely to do so.

Choosing a strategy for program evaluation depends on the type of information that is needed to make appropriate, accurate decisions on changing and improving a program. Although different experts conceive program evaluation differently, one perspective describes three general strategies for evaluation: goals-based, process-based, and outcome-based (see Clarke [5] for more on program

evaluation). In a goals-based framework, evaluation comprises understanding how well the current status of the program – in this case an information security activity with an organization – matches a set of preset goals (presumably those that were articulated at the time when security systems or personnel were deployed). In a process-based evaluation, the focus is on the workflow and activities that currently exist in the area of information security, rather than on its status at one point in time. This is probably the most common type of security assessment, given that many organization’s security operations grow organically over time rather than being invented as a single, unified initiative. Finally, outcome-based evaluation focuses on the value obtained by program stakeholders, in general the responsible managers in charge of information systems or information security. These categories overlap to some degree and a combination of strategies works better than one or another implemented singularly.

In this perspective, the emerging practice of information security risk assessment is one form of process-based security program evaluation. For the purposes of this paper we lump together general security assessments, security risk assessments, and other types of non-audit assessments into the same general grouping. The important characteristic that sets all of the members of this class apart from the strict definition of auditing is the acceptance and use of qualitative research methods within the context of the evaluation process. For example, Vidalis [17] stated that the qualitative approach to assessment has been more frequently utilized when evaluators were concerned about the impact of employees’ behaviors affecting information security.

Unfortunately, as Kotulic & Clark [9] suggested, the research literature contains few studies pertaining to the processes involved in information security assessment – particularly when such assessment is primarily qualitative in nature. These researchers suggested that a primary reason for this is that organizations are afraid to release information about the status of their information security programs. As a result of this reluctance, however, academic researchers and other involved in the practice of information security assessments have relatively few opportunities to work publicly with data that emerge from security assessment projects. We believe that this gap in research prevents the organizations from finding tested assessment tools and improving their own tools for assessing information security within the organization.

Existing books describing strategies for conducting security assessments propose no straightforward way to compare the differing assessment strategies (e.g.,

Krauss [10]; Senft [16]; Hunton, Bryant, & Bagranoff [8]; Champlain [3]; Herrmann [7]; Peltier [14]), understand their relative strengths and weaknesses, or to determine how their methods may best apply to organizations with different sizes and missions. For example, one of the oldest books in this area, by Krauss [10], proposed a generic security evaluation technique called “SAFE” which was one of the techniques that began the current craze for security checklists. Yet Kraus focused primarily on evaluation for large financial data systems, and had little to say about the organizational context in which those systems were deployed.

Likewise, a very popular second edition of a contemporary book edited by Senft [16], which notably positions itself as a textbook on information systems auditing, contains elements of risk assessment, computer-aided security assessment, policy analysis, governance, outsourcing evaluation – a veritable encyclopedia of auditing and assessment techniques – but provides no clear criteria for knowing which methods to use when or why. This book, along with the several other volumes in this area (e.g., Hunton, Bryant, & Bagranoff [8]; Champlain [3]; Herrmann [7]; Peltier [14]) provide a range of attractive and apparently systematic techniques that assessors might beneficially use to conduct a security assessment, but no strategy for choosing among the techniques, ordering them, or understanding their applicability to organizations of different types or assessments with different goals. We believe that this deficit has resulted from the dearth of research on the processes of information systems auditing.

This is a critical issue, as suggested in an article by Winkler [18] who argued that, “An assessment is potentially the most useful of all security tests, but it is also the hardest to define” because its “report must include comprehensive information on how to secure the client's technical and non-technical vulnerabilities.” The data we collected in 15 small and medium sized organizations we hope will begin to shed some light on how to analyze the data from qualitative security assessments in an efficient manner that will be beneficial to organizations and their employees.

Note that the distinctive approach that we are taking to analyze security assessments’ interviews’ transcripts has occurred in analogous ways in health informatics to analyze interview transcripts from patients, nurses, and doctors (Burnard [1]; [2]). In health care, the need for controlled vocabulary arose because communication between patients, staff, administration, and doctors needed improvement. Controlled vocabulary also enabled hospitals to work

more efficiently by standardizing the terminology used to report and communicate the organization’s activities. In the managerial and behavioral areas of information security we witness the same need arising for standard terminology. (To a certain extent the ‘technical’ areas of information security have standard terminology – e.g., for equipment and protocols – but the quality and universality of these vocabularies vary considerably). The development of a controlled vocabulary for security assessment would allow for the development of assessment interview protocols, questionnaires, analyses, and reports. These improvements could in turn generate better feedback that would potentially enhance the communication among employees especially between end-users and IT staff members.

Developing controlled vocabulary has also been used in other settings to enhance the analysis of interviews. For instance, a human rights group (see http://www.hrdag.org/resources/controlled_vocab.shtml) has built a controlled vocabulary to analyze its data collected from various problem areas around the globe. Indeed, in this group, controlled vocabulary is used to “enable the researchers to decipher the often-complex relationships (...), and ultimately help to answer the question of, “Who did what to whom [because] the controlled vocabulary transforms the collected information into a countable set of data categories, without discarding important information and misrepresenting the collected information.”

Technology solutions to information security problems need to be supplemented by effective management and governance processes. Many sociotechnical researchers would agree with Gordon et al. [6] that conducting an effective security assessment can be a key step for an organization to take in order to facilitate improvements to their security management and governance. Other than Chatterjee et al. [4] – who showed how commitment to a security assessment can impact the success of the assessment process – few articles exist that examine how data collected in security assessments should actually be used in later security governance efforts. We believe that the first step in rectifying this situation is to understand in detail the nature of the data that are gathered from the qualitative components of security assessments. Everyone knows what to do if a quantitative vulnerability assessment shows that 50% of your systems remain unpatched: Patch them. On the other hand, if 50% of your employees are expressing vague reservations about the quality of their communication with the IT security staff, what does this mean? We don’t know, but we aim to find out, and the way we plan to go about it is to first

obtain a systematic understanding of the contents of security assessment interviews.

3. Method

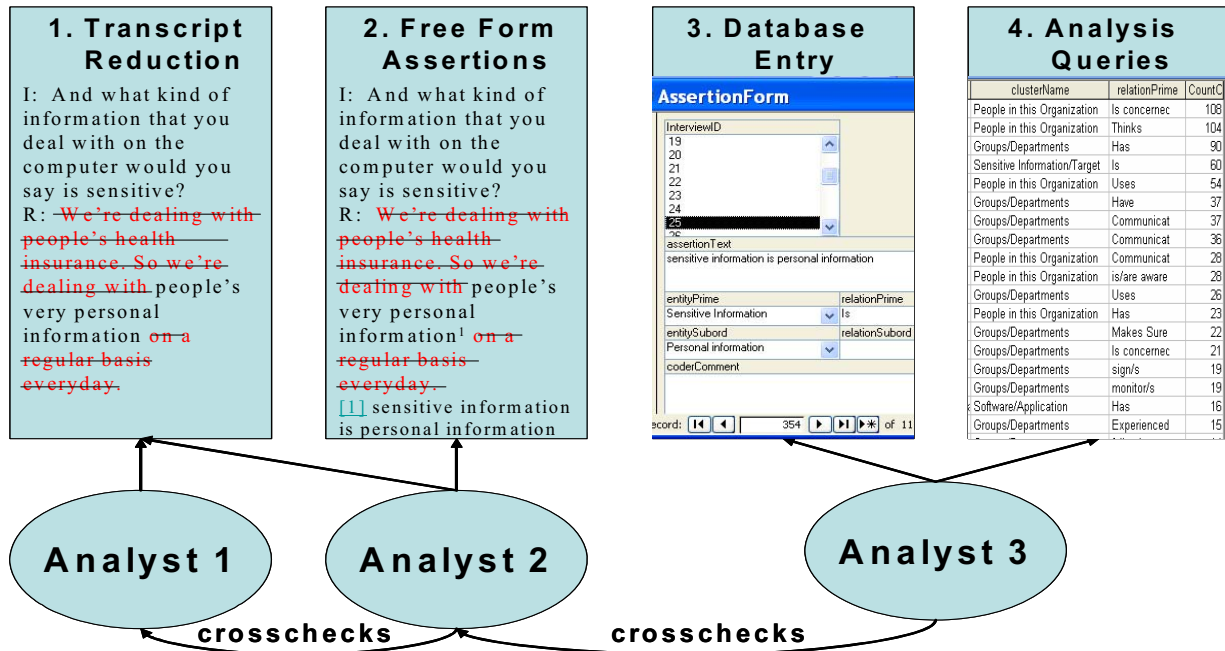
Between October 2003 and August 2004 we conducted comprehensive security assessments in 15 small and medium-sized organizations belonging to various industries as a part of a larger research project on information systems security. The number of employees in these organizations varied from 6 to approximately 2100 employees. Six subject matter experts (graduate students and faculty members in information technology) conducted 68 interviews lasting approximately 20-30 minutes with selected employees, managers, executives, information technology personnel, and information security specialists. The interview protocols were carefully constructed to elicit knowledge, experience, understanding, beliefs, and attitudes related to information security, information privacy, and security policies in the host organization. In addition, within each organization a technical security specialist from our team and two computer science graduate students working as assistants mapped and analyzed the technology infrastructure of the organization. The security assessments typically took four weeks to complete and involved multiple site visits to each organization.

Among the 15 organizations, a non-random purposive sample of 71 employees was selected based upon each employee's position within the organization. The discrepancy between N=71 respondents and the N=68 interviews reported elsewhere in this paper arose from the fact that in three of the interviews, two respondents were interviewed together. The purposive sampling strategy attempted to obtain responses from a representative of each key stakeholder group within the respective organization. For example, in a software engineering organization we obtained an interview with a programmer whom we believed represented the larger set of frontline programmers in the organization. Prior to going to the organizations to conduct the interviews, we asked the various organizations to provide us with an organizational chart so that we could identify the employees we would interview. We were careful to choose employees who held key roles regarding information security, information privacy and security policies. We interviewed leaders of the organization, managers of the different departments, end-users and information technology and information security professionals. Generally speaking we interviewed about four or five people from each organization.

Among respondents N=71, 50.7 % were female, 43.7 % were male, and 5.6 % did not wish their gender to be recorded on the transcript for confidentiality reasons. For the purposes of this study we did not record the ethnicity of the respondents.

Figure 1 depicts the process used for the data analysis of the interviews transcripts. This process led to the development of a so-called "controlled vocabulary." Development of this process was overseen and tested by ten subject matter experts: faculty members and undergraduate and graduate students in an information technology Ph.D. program. The aim of the analysis process was first to reduce interview transcripts to essential material by filtering out those utterances that were irrelevant to the topic of information security. This procedure involved deleting superfluous information such as warm up questions-answers, and tangents. At least two reviewers examined each deletion decision to make sure that important information was not elided. Conflicts between a reduction decision and the opinion of a reviewer were resolved by including the originally deleted information. The reduced transcripts formed the basis for the creation of freeform "assertions" statements that would be later used to define a controlled vocabulary. As explained in greater detail, the controlled vocabulary was built up over repeated iterations of examining freeform assertions and trying to represent them with existing elements in the controlled vocabulary. As with other stages of this construction process, reviewers examined both the freeform assertions and the controlled vocabulary analogues of those assertions.

After we regularized the process of analyzing the transcripts and documented it in an analysts' guide, the subsequent processing of remaining transcripts was completed by three analysts who worked jointly. The tasks were divided among three different analysts to ensure maximum objectivity and accuracy of the process. The analysts performed their assigned role in the reduction and analysis; they also crosschecked the work of the other analysts and discussed with them – if necessary, their different choices. The analysis of the interviews' transcripts was done following these four core stages: Analyst 1 reduced the transcript into discrete question-response blocks by deleting the superfluous information and tangents. Analyst 2 received the reduced transcript, checked the work of Analyst 1 and then developed freeform short assertions for each remaining question-response block. Analyst 3 crosschecked the freeform assertions for clarity, accuracy, and completeness. Analyst 3 entered the verified assertions into a standardized database based on the controlled vocabulary.



4. Results and Discussion

The focus of our analysis was on the development of a controlled vocabulary that could subsequently be utilized to analyze security assessment interviews. The successful generation of such a controlled vocabulary might enable effective analysis of security assessment interviews by helping analysts identify key phrases, assertions, and relationships related to the security culture, policies, and practices of an organization.

A controlled vocabulary comprises a set of terms determined by a team of subject matter experts who survey the terminology used in a certain domain and identify the most commonly used key terms. The terms may include entities (subjects and objects) as well as relations (verbs). Other language elements such as modifiers can also be included, but we opted to keep our system as straightforward as possible at this early stage. We did represent negation and tense in relationships, thus providing the opportunity to understand when a respondent asserted that something was not so, and to know whether a respondent was referring to past, present or future. We anticipate that the combinations of entities and relations as offered by interviewees can provide information on security issues, strengths, weaknesses, and risks within the organization.

We developed the terms in the closed vocabulary over the course of a 15-week long, iterative process involving a team of 10 subject matter experts (faculty

members, undergraduate and graduate students in an information technology Ph.D. program) who met weekly to develop the vocabulary. In response to the proliferation of specific topics (entities) referred to by respondents, the team also developed a clustering scheme to organize the entities, their definitions and their content. Figure 2 shows the process of iterations to develop both the terms and the clusters of the vocabulary.

We created a relational database to organize the data resulting from the first two phases of the analysis process described above and shown in Figure 1. As a reminder, the first phase consisted of transcript reduction and the second phase involved developing freeform assertions from previously reduced transcripts. An assertion is a shorter, more precise phrase representing the content of a section of text. During this process, we paid careful attention to the data reduction and assertion formation processes to ensure that the information reported by respondents about security issues was represented as accurately as possible to the original verbatims. The 68 interviews (with 71 people) generated a total of 1218 distinct assertions. Table 1 presents the different core elements utilized in the database, their definitions and the total number of elements of that type.

Table 1. Overview of the database

Name	Definition	Total
Coded Freeform Assertions	A short, precise phrase representing the content of portions of verbatim transcript text	1218
Entities	A primary or subordinate subject reported by a respondent	73
Entity Clusters	A grouping of closely related entities	12
Objects	An entity that always appeared as the target of a relation	59
Relations	A group of related verb forms (e.g., variations on the verb "to be"); provides links among entities and objects	27

Given the number of interviews we conducted, we found that on average each interview contained about 18 distinct assertions concerning security. On average the security assessment interviews were approximately 20 minutes long, with a continuous stream of discussion throughout the whole interview. As a result it appears that on average each respondent made one security-related utterance per minute. Given the number of words in a typical interview (about 1600), this result indicated that on average we reduced 80-100 words of discussion down to a single assertion about security. The following two examples illustrate the complete process:

Example 1: *Verbatim:* "Interviewer: And my last question is about your co-workers, do they do anything that you wonder if that would not be the best action, like sharing passwords or leaving passwords on the monitor? Respondent: Everyone in our office knows each others password." *Free-form assertion:* "Employees share passwords." *Closed vocabulary assertion:* <EntityPrime> Employees in this organization <RelationPrime> share(s) <EntitySubord> passwords

Example 2: *Verbatim:* "Respondent: I understand that the network that we have makes it difficult to hack onto our system." *Free-form assertion:* "The respondent feels that the organization's network is difficult to

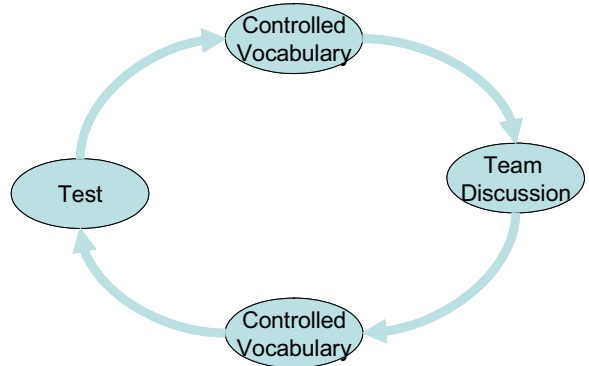


Figure 2: Iterative Vocabulary Development

hack." *Closed vocabulary assertion:*
 <EntityPrime> Respondent/Interviewee
 <RelationPrime> Thinks/Believes
 <entitySubord> Information System(s)
 <relationSubord> Is/Are <Object> Secured Status

As the above examples illustrate, the controlled vocabulary allows for compact and accurate representation of the content and its database representation allows for rapid and accurate analysis. We conducted some preliminary analyses on the assertions so that we could begin to understand what the respondents were discussing. Table 2 shows the topics of discussion by showing the total number of entities mentioned within each of the 12 entity clusters.

Table 2. Overview of Clusters

Clustered Entity	# as Prime	# as Subord.	Total
Groups/Departments in this Organization	475	117	592
People in this Organization	436	17	453
Security Countermeasure	38	307	345
Sensitive Information	82	122	204
Security Related Policies	32	152	184
Security Related Attitude	0	125	125
Communication Modes	23	102	125
Security Threat	10	110	120
Software Application	32	41	73
Communication	28	13	41

Clustered Entity	# as Prime	# as Subord.	Total
Quality			
Organizational Behavior	4	31	35
Common Sense	2	18	20

Note that the first numeric column (# as prime) refers to the use of entities in the respective cluster as the subject of a sentence, whereas the second numeric column (# as subord.) refers to the use of the entities as the subject of a subordinate clause or the object of a sentence. Even at this early stage of analysis, the results suggest a number of important points. First, the most salient topic to these respondents was people: themselves and their departments. The entities clustered under “Groups/Departments in this organization,” included the respondents own department, the IT department and other departments in which information security was an important issue (note that these clusters represent an intentional summary of the data for analytical purposes; we retained all of the underlying detail about what departments the respondent was discussing). When the respondents were discussing people, they primarily referred to themselves, their coworkers and managers.

Next, note that the primary security topics on the minds of these respondents comprised security countermeasures and the protection of sensitive information. This intuitively appealing result hides a significant detail that we learned from additional analyses: Respondents seemed more likely to make utterances pertaining to countermeasures and protection of sensitive information in the organizations that the security audits showed had the *most favorable* overall security outlook. Prior to conducting these analyses, we obtained ratings of the overall security status of each of the 15 organizations that allowed us to form a consensus ranking of the organizations from most secure to least secure. In the least secure organizations, people talked mainly about themselves and others. In the most secure organizations people talked more about countermeasures and less about themselves. If we are able to confirm this result with additional analyses, it suggests that the most loquacious discussions of people and security may occur in the organizations that are least well prepared. On a related note, the organizations in which the cluster of entities known as “Common Sense” appeared frequently also seemed to have unfavorable security conditions.

We are continuing to conduct analyses on the entities contained in each cluster and are working on development of a visualization technique to be able to

display the entities and relations together, weighted by their importance overall, their importance within each respective organization, and their importance relative to a single respondent’s utterances. We are also continuing to adjust the content of clusters in order to ensure that the results of any summary analyses accurately summarize the underlying entities and relations. For example, the definition of one cluster is shown in Table 3.

Table 3. Example Cluster

Cluster name	Security Related Policies
Cluster definition:	Governance statements regarding security in the organization and the workplace
Cluster content:	Acceptable use policies; behavioral rules; confidentiality agreement; security policies

Sometimes respondents refer to behavioral norms that people in the organization seem to agree exist, but that have not been codified into a written document. The question of whether one would include assertions about such behavioral norms in this cluster or in a different cluster depends upon whether one accepts or believes that such norms have a similar influence to written policies. If they have substantially more influence or substantially less, or if their influence affects security in a different way, than the choice of clustering such norms with written policy documents might obscure important distinctions that respondents were trying to make about the status of their organizations.

Overall we have just begun to scratch the surface in analyzing the data extracted from these interviews, but based on our preliminary analysis we expect to find that the frequency of occurrence of various assertions made by managers, employees, and technology people in security assessment interviews will vary systematically across organizations and may relate to certain security-relevant outcomes in those organizations.

5. Conclusion

In this paper we suggested that a key distinction between information security auditing and information security assessment arises from the differences in method and scope between these different activities. Whereas security audits can focus attentively on quantifiable records of information systems performance, the broader realm of security assessment typically involves at least some capturing of qualitative data. While qualitative data relevant to

information security apparently can provide a rich source of insight about the security status of an organization, the general lack of sociotechnical research related to conducting security assessments has interfered with our collective abilities to understand what security assessment interviews are telling us. In short, we can interview managers, employees, and technology people and apparently gather lots of interesting information about security, but we are not quite sure what to do with it. We posed the question of whether systematic analysis of security assessment interviews could yield useful information about the security status of an organization.

Based on our preliminary analyses of more than 1000 assertions made in the course of 68 security assessment interviews, we suggest that the answer is yes. Our interviewees discussed themselves, their departments, the sensitive information they were charged with protecting and the security measures used to protect that information. These topics were the most common areas of discussion. Although it is too soon in our analysis process to conclude that these utterances and the relative frequency of their appearance relate systematically to the security status of an organization (or what the mechanism of that relationship might be), we believe that we have established a useful *method* for extracting from security interviews the content that has the highest likelihood of establishing such relationships.

Of perhaps more value, we have gotten a running start on the development of a system of terms – a closed vocabulary as we have called it – that we believe may be reusable and useful to other researchers who conduct security assessments as part of their science or practice activities. Much work remains in order to make this closed vocabulary into an automated or semi-automated system for assisting with security audits, but we hope that this initial work helps to pave the way for that eventual goal.

6. References

- [1] Burnard, P. “A method of analyzing interview transcripts in qualitative research.” *Nurse Education Today*, 1991, 11, 461-466.
- [2] Burnard, P. “Searching for meaning: a method of analyzing interview transcripts with a personal computer.” *Nurse Education Today*, 1994, 14, 111-117.
- [3] Champlain, J.J., *Auditing Information Systems*, John Wiley, Hoboken, NJ, 2003.
- [4] Chatterjee, K., Morton, S. and Mukherji, A., “Strategic Audit Policies Without Commitment”, 1999, mimeo.
- [5] Clarke, A. *Evaluation research: An introduction to principles, methods, and practice*. Sage, London, 1999.
- [6] Gordon, L.A., Loeb, M.P. and Sohail, T. “A Framework for Using Insurance For Cyber-Risk Management”, *Communication of the ACM*, 2003, 46(3), 81-85.
- [7] Herrmann, D.S., *Using the Common Criteria for IT Security Evaluation*, Auerbach Publications, Boca Raton, Fla., 2003.
- [8] Hunton, J.E., Bryant, S.M., and Bagranoff, N.A., *Core Concepts of Information Technology Auditing*, Wiley, Hoboken, NJ, 2004.
- [9] Kotulic A.G. and Clark, G.J., “Why there aren’t more information security research studies”, *Information & Management*, 2004, 41, 597-607.
- [10] Krauss, L.I., *SAFE: Security audit and field evaluation for computer facilities and information systems*, AMACOM, New York, 1980.
- [11] Myers, M.D. “Qualitative Research in Information Systems”, *MISQ Discovery*, June 1997.
- [12] Myers, M.D. & Walsham, G. “Exemplifying interpretive research in information systems: an overview”, *Journal of Information Technology*, December 1998, 13 (4), 233-234.
- [13] Neuendorf, K.A. *The content analysis guidebook*. Sage, Thousand Oaks, CA, 2002.
- [14] Peltier, T.R., *Information Security Risk Analysis*, Auerbach Publications, Boca Raton, Fla., 2001.
- [15] Rubin, H.J. & Rubin, I.S. *Qualitative Interviewing: The Art of Hearing Data. Second Edition*. Sage, Thousand Oaks, CA, 2005.
- [16] Senft, S., Daniel, P., Ph.D. Manson, Gonzales, C., and Gallegos, F., *Information Technology Control and Audit, Second Edition*, Auerbach Publications, Boston, MA, 2004.
- [17] Vidalis, S., “A Critical Discussion of Risk and Threat Analysis Methods and Methodologies”, School of Computing Technical Report CS-04-03, 2004.

[18] Winkler, I. "Audits, assessments & test (oh, my)", *Information Security Magazine*, 2000.

[19] Walsham, G. *Interpreting Information Systems in Organizations*. Wiley, Chichester, 1993.