Improving the Exchange of Lessons Learned in Security Incident Reports: Case Studies in the Privacy of Electronic Patient Records

Ying He, Chris Johnson, Yu Lyu, Arniyati Ahmad

Improving the Exchange of Lessons Learned in Security Incident Reports: Case Studies in the Privacy of Electronic Patient Records

Ying He[1], Chris Johnson[2], Yu Lu[3], and Arniyati Ahmad[3]

School of Computing Science, University of Glasgow, UK
[1] yingh@dcs.gla.ac.uk
[2] christopher.johnson@glasgow.ac.uk
[3] {y.lu.3,a.ahmad.1}@research.gla.ac.uk

Abstract. The increasing use of Electronic Health Records has been mirrored by a similar rise in the number of security incidents where confidential information has inadvertently been disclosed to third parties. These problems have been compounded by an apparent inability to learn from previous violations; similar security incidents have been observed across Europe, North America and Asia. This paper presents the results of an empirical study that evaluates the utility and usability of conventional text-based security incident reports with a graphical formalism based on the Goal Structuring Notation. The two methods were compared in term of the users' ability to identify a number of lessons learned from investigations into previous incidents involving the disclosure of healthcare records. These lessons included both the causes of the incident but also the participants' ability to understand the reasons why particular recommendations were proposed as ways of avoiding future violations. Even using a relatively small sample, we were able to obtain statistically significant differences between the two approaches. The study showed that the graphical approach resulted in higher accuracy in terms of number of correct answers generated by participants. However, subjective feedback raised further questions about the usability of both approaches as the readers of security incident reports try to interpret the lessons that can increase the security of patient data.

Keywords: Lessons Learned, Security Incident, Electronic Patient Record, Generic Security Template, Empirical Study

1 Introduction

According to Symantec, the healthcare accounted for 42% in the total number of attacks on electronic information systems in 2012 [1]. At 36% in 2013, healthcare continues to be the sector responsible for the largest percentage of disclosed data breaches by industry [2]. Almost identical breaches have occurred across Europe, North America and Asia [3]. Learning from the incident enables the organisation to extract meaningful information from incidents, and use this information to

improve security management systems [4]. Effective communication mechanism is needed to synthesis the information from the incident into the security incident management system so as to prevent a similar incident.

Popular communication mechanisms include formal reports, less formal meetings, newsletters, emails, as well as presentations to management [4]. However, the detailed incident reports that are produced in the post-incident activity [5] have not been given enough attention. Those reports contain comprehensive information, which is typically classified into two types, business impact and remediation information [5]. Business impact information involves how the incident is affecting the organisation in terms of mission impact, financial impact, etc. For example, "The missing external hard drive is believed to contain numerous research-related files containing personally identifiable information and/or individually identifiable health information for over 250,000 veterans, and information obtained from the Centres for Medicare & Medicaid Services (CMS), Department of Health and Human Services (HHS), on over 1.3 million medical providers"[6], Remediation information information mainly refers to the suggested remediation actions, plans, procedures, and lessons learned. For example, "We recommend that the Assistant Secretary for Information and Technology revise VA Directive 6601 to require the use of encryption, or an otherwise effective tool, to properly protect personally identifiable information and other sensitive data stored on removable storage devices when used within VA."[6].

As for a purpose of sharing, it is suggested to avoid sharing business impact information with outside organisations unless there is a clear value proposition or formal reporting requirements. When sharing information with peers and partner organisations, incident response teams should focus on exchanging remediation information [5]. The remediation information information reported describes (1) the security issues, e.g. "The position sensitivity level for the IT Specialist was inaccurately designated as moderate risk, which was inconsistent with his programmer privileges and resulted in a less extensive background investigation", (2) the security objectives violated during this process, e.g. "Position Sensitivity Level Assessments were Not Adequately Performed", and (3) the recommendations, e.g."We recommend that the Under Secretary for Health direct the Medical Centre Director to re-evaluate and correct position sensitivity levels and associated background investigations for positions at the Birmingham VAMC " [6]. Those granular information are inter-related, however, they are scattered documented in a pure textual based report that makes it difficult for the readers to identify the relationships among them. This issue has been compounded by the lengthy security incident report, which is usually around hundred of pages [6]. The stakeholders responsible for protecting patient data lack the time and the motivation to spend the many hours needed to read and digest existing reports. This creates significant problems within the wider scope of security management systems; it can be difficult to accurately assess the likelihood or consequences of future attacks when managers are unaware of previous incidents.

Graphical techniques can address some of these limitations. The Generic Security Template (G.S.T.) has been developed [3, 7] to help readers understand

the lessons learned from previous security incidents. In particular, it extends the Goal Structuring Notations (GSN) [8] to provide an overview of previous data breaches. The intention is to map out the security objectives, security issues and recommendations that are embedded in the many pages of text that are used in conventional reports. More information on the GSN and the G.S.T. is provided in section III. Fig. 1 provides an excerpt from one of these diagrams. It is based on a report into the disclosure of personal information about 250,000 veterans and over 1.3 million medical providers by the US Veterans Affairs Administration (VA) [6]. This incident report provides the case study that is used throughout this paper. The leaf nodes in this diagram are used to gather together the recommendations that were intended to avoid future incidents. The internal nodes are used to show how each of these findings supports higher level goals and sub-goals intended to ensure that systems meet an acceptable level of security, defined in terms of the US Government's Federal Information System Controls Audit Manual (FISCAM) [9]. Further information about the graphical technique is provided in [3, 7, 10]. The use of graphical overviews is intended to make it easier to identify recommendations that can be transferred from a previous incident to prevent similar breaches from occurring in other organisations.

Previous work has shown that GSN can be used to map common lessons from data breaches in healthcare organisations in healthcare organisations in both the United States and in China [3]. Although these incidents occurred in very different contexts, the security concerns and the consequences for patient confidentiality show remarkable similarities. This previous work provided initial case studies but did not, present empirical support for the benefits of using graphical representations compared to text-based reports of security incidents. This paper, therefore, presents a controlled experiment to investigate whether graphical approaches can be used to augment conventional, text-based documents. The remainder of the paper is structured as the following, section 2 reviews the related work, section 3 briefly introduces the G.S.T., section 4 outlines the experiment design, section 5 presents the experiment procedure, section 6 prepares data to analyse the results, section 7 analyses the results, and section 8 summarises the paper.

2 Related Work

There is a natural reluctance to share details of previous security breaches - reports may inspire new attacks or publicize vulnerabilities. However, a growing number of regulatory agencies now provide detailed reports that are intended to help avoid any recurrence of previous failures. Security management systems have also been introduced into many healthcare organisations to ensure previous security incidents inform threat and risk assessments [11]. Improving situation awareness, in particular about security breaches, help persuade end users of the importance of existing policies and procedures. There are further benefits from the wider dissemination of incident reports. Security engineers can learn

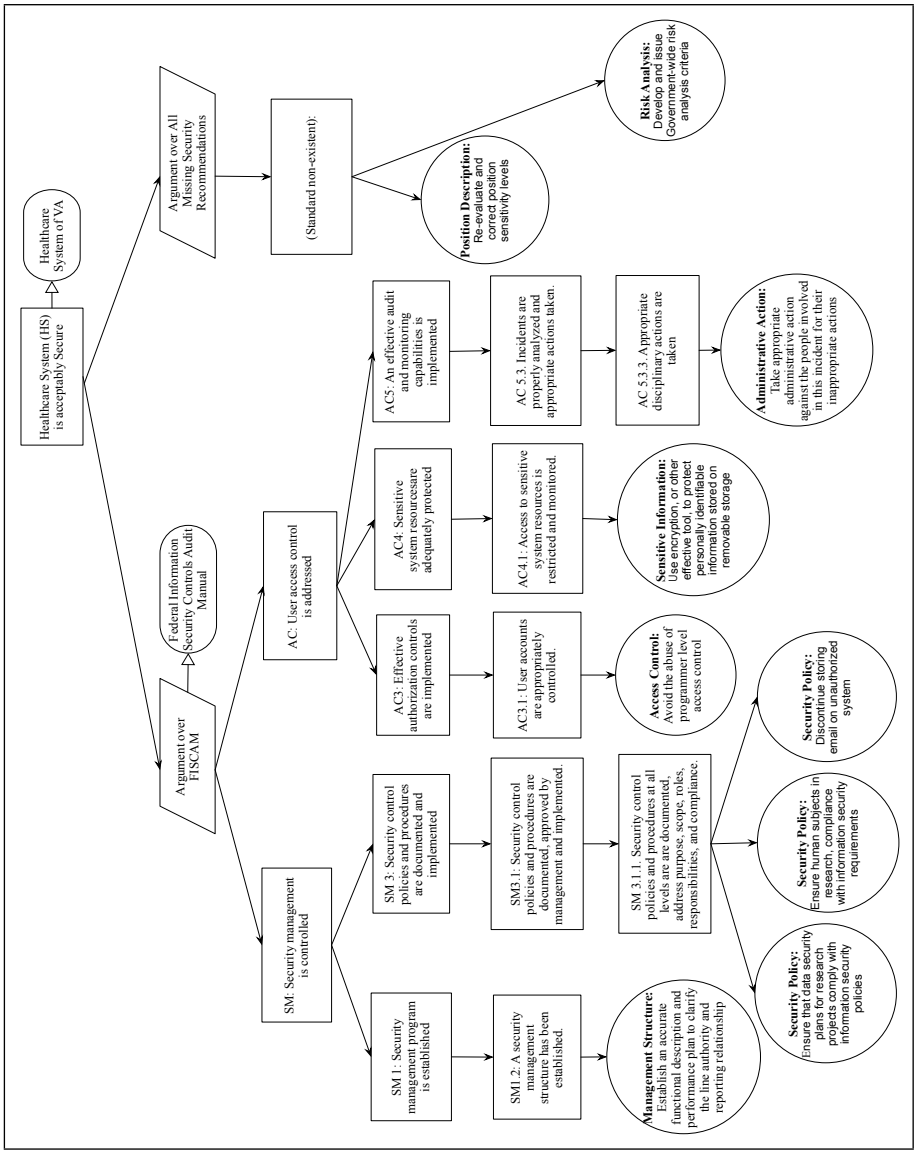

Fig. 1. Generic Security Template - VA dataloss 2007

important lessons about the analysis, containment, eradication, and recovery from previous attacks.

The introduction has argued that existing, text-based reports can be supported through the use of graphical notations that provide an overview of many dozens of pages of detailed prose. Fig.1 uses the Goal Structuring Notations to summarize key findings from an enquiry into a loss of confidential patient data from the US Veterans Affairs Administration [6]. The aim is to present the security objectives in a structured and coherent manner. It is also hoped that this use of a semi-formal notation will encourage greater consistency and correctness [12, 13]. However, the notation introduces unfamiliar syntax and semantics. There is a danger that our use of these techniques can prevent stakeholders from understanding the arguments in security incident reports [14–16]. This paper, therefore, presents a controlled experiment to evaluate the utility of graphical representations for security incident reports.

3 The Generic Security Template

As mentioned, the G.S.T. extends the Goal Structuring Notations (GSN) [8] to provide an overview of previous security breaches. GSN is the dominant approach in the UK defence sector, increasingly being used in safety-critical industries to improve the structure, rigor, and clarity of design requirements. A particular strength is that it also links the evidence to show that particular requirements have been met. The same approach has more recently been extended to document security requirements [3, 7]. There are four principal notations used in the GSN, A *Goal* is a claim, the statements that the goal structure is designed to support. *Evidence* exists to support the truth of the claimed goal, which can be documented by providing a solution in GSN. *Strategy* is inserted between goals at two levels of abstraction, to explain how the top-level goal is addressed by the aggregation of the goals presented at the lower level. *Context* is used to declare supplementary information and provide adequate understanding of the context surrounding the claim (or strategy). Usually it presents concepts clarification introduced in the claim (or strategy) [8].

The G.S.T. has customised the GSN. Instead of collecting evidence to support design and development requirements, it collects lessons (i.e. security causes and recommendations) from previous security incidents. These lessons are defined to be the knowledge or understanding gained by experience [17]. In the G.S.T., it refers to the security issues that cause a security breach, and the security recommendations intended to avoid any recurrence. The evidence of compliance with the security objectives is presented in the form of a specific security standard or guideline applied to the organisation where the security incident happened. This has reflected the granular information described in section 1. Generic, is defined as "characteristic of or relating to a class or group of things; not specific". In other words, the intention is to create a GSN diagram that conveys the lessons learned from specific previous security breaches at a level of abstraction that

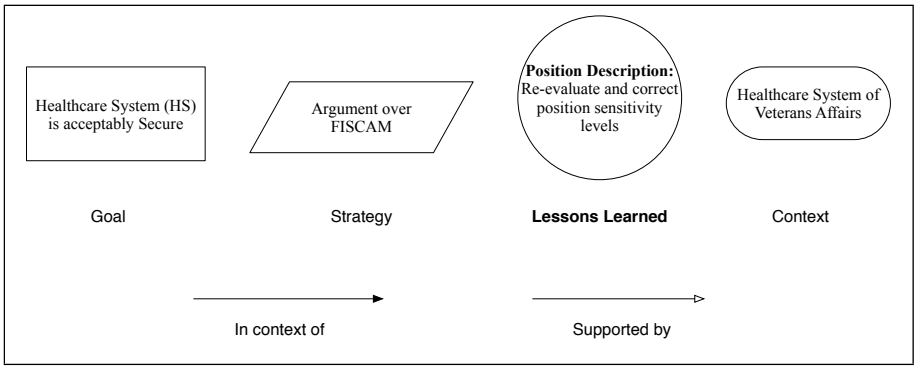

Fig. 2. Customised GSN Notations

helps others to use them to improve the security of other systems. Fig.2 presents the notations used in the Generic Security Template.

Fig.2 presents the notations used in the Generic Security Template. In particular, rather than using the evidence derived from validation and verification to support safety arguments. The G.S.T. uses the findings from previous data breach incidents (i.e. leaf nodes of Fig. 1) to support security arguments in terms of the compliance with the security guideline (i.e. internal nodes of Fig. 1). The other concepts remain the same between both application areas.

4 Experiment Design

4.1 Experiment Objective and Scope

A controlled experiment was conducted to evaluate whether the use of graphical techniques helps improve the comprehension of the lessons from previous security incident reports compared to conventional text-based approaches. The aim was not to show that graphical techniques could replace conventional reports; in contrast the focus was in the use of our extended GSN approach to provide a map or overview of complex text-based reports. Accuracy, efficiency and task load are compared quantitatively in this experiment and the following hypotheses are proposed for the comparison.

H1: Participants will be better able to identify the Lessons learned (security issues, security recommendations) in security incident report with the help of the G.S.T. than using Text-based Document alone;

H2: Participants will be better able to identify the compliance with the security objectives in security incident report with the help of the G.S.T. than using Text-based Document alone;

H3: The time taken to complete the designed task will be less using the G.S.T. than that using the text-based Document alone;

H4: The task load will be lower using the G.S.T. than using the text-based Document alone.

Ease of use is compared qualitatively based on the feedback obtained from participants.

4.2 Experiment Variables

Dependent Variables We evaluate the usability [18] in terms of the accuracy, efficiency, ease-of-use and task load compared to the conventional, text-based approach. *Accuracy*, is measured by assessing the quality of the security causes, recommendations and the compliance with the security objectives from the security incident. *Efficiency*, is measured by the time it takes to complete the experiment task. *Ease of use*, is evaluated by the feedback obtained from the post-experiment questionnaire. *Task load*, is measured by the application of NASA's Task Load Index to assess workload [19].

Independent Variables *Generic Security Template (G.S.T.)*, we used the same G.S.T. across all participants. This presents findings from the US Veteran's Affairs administration 2007 Dataloss Incident [6]. *Text-based approach*, we developed an executive summary (reduced to four pages) and a simplified security guidelines (reduced to three pages) from the FISCAM. More details on the experiment material preparation are provided in 4.3.

Controlled Variables *Participants*, the participants were post-graduate and undergraduate students with different education background. *Tasks*, the experiment itself lasted for maximum one hour. Participants had to identify causes, recommendations and the relationships with security objectives using either a conventional text-based document or using the graphical overview plus the existing report.

Extraneous Variable *Experience with GSN*, is defined as an extraneous variable in this experiment. People who have experience with GSN will have an obvious advantage in comprehending the security incident with the help of the G.S.T. People who have experience with GSN were excluded from this experiment.

4.3 Experiment Materials

Security Incident Report. Security Incident Report. The technical context of the task focused on a data loss incident involving the Veterans Affairs' Administration [6]. The original report was around 80 pages long and hence we could not use it directly within the time available for the experiment. We also felt that our more focused approach was more appropriate for an initial study that could, in turn, inform future empirical work over a longer period of time and with a larger number of participants. We, therefore, provided both groups with the executive summary from the VA report reduced to four pages. As is

mentioned, the evidence of compliance with the security objectives is presented in the form of a specific security standard or guideline applied to the organisation where the security incident happened. Therefore, a simplified version of security guidelines (reduced to three pages) cited from the FISCAM that are relevant to this incident are also provided as a part of the security incident report.

The G.S.T. The G.S.T. used in this experiment is created from the above mentioned security incident report only. It is an abstraction and extraction of the desirable information and did not bring any other information that can bias the results of the experiment.

The Questionnaire. We developed separate tasks description for the two groups and a post-experiment questionnaire, to provide subjective insights into perceived workload. A slightly different version of this post-experiment questionnaire was developed for the group using the graphical overview of the security incident. They were asked to provide information about the usability of the approach by completing subjective questionnaire.

4.4 The Pilot Study

Two security experts reviewed the design of the experiment pilot studies then helped to identify issues that had not been identified during the preparation of the materials. In the first pilot study, participants had to identify security issues, recommendations and compliance with the security objectives; writing them down using freestyle text. This was to simulate how security incident reports are analysed in practice, where people normally have no tools assisting them throughout this process. The feedback from the participants showed that the task was very mentally demanding and they were not able to complete it within one hour. We corrected this problem by introducing a table that provided guidance on the security issues and recommendations. The participants are required to fill in the blanks cells in the table. For the measurement of compliance with the security objectives, we have used multi-choice questions as the measurement mainly focuses on the relationships between the security objectives and the recommendations for prevention. Two more participants conducted a pilot test of the new experiment design. They were able to finish the tasks and stated that the level of mental effort was acceptable.

4.5 Experiment Task Design

In Group A, the experiment materials included the textual incident report (reduced executive summary and reduced security guidelines from FISCAM), the graphical G.S.T. and a task description. The pilot study had confirmed the arguments presented in the opening sections of this paper; that it can be difficult for readers to identify the causes, recommendation and the compliance with the

security objectives of previous security incidents from existing textual reports. We, therefore, created tasks that guided the participants' analysis:

Task 1: Identify security issues and recommendations from the security incident report with the help of the G.S.T. They had to complete missing information from a table that provided partial information about the causes and recommendations. Table 1 is an exempt of the table. Issue Category and Description are provided. The participants need to fill in the blank about the recommendation description.

Issue Category	Issue description	Recommendations description
Access Control Related	The IT Specialist was improperly given access to multiple data sources.	

Table 1. An exempt of the security issue and recommendation table

Task 2: Answer multiple-choice questions on compliance of the security objectives. This removed the additional contextual support of the tabular format used in task one and provided a stepping stone towards the open ended analysis of security incident reports that proved problematic in the pilot studies.

In Group B, the experiment materials included the textual incident report without the G.S.T. but participants had the same task descriptions as the first group.

5 Experiment Procedures

5.1 Experiment Treatment

There was only one treatment in the experiment using a between groups (Group A and B) design. The empirical comparisons are between one group using a conventional text-based document and another using the graphical overview as well as the existing report.

5.2 Participants

Twenty-four subjects were assigned to either of the two experimental conditions using the textual report only or using both the textual report and the graphical overview. Group A consists of one undergraduate student and eleven postgraduate students, within which three of them have information security experience; Group B has one undergraduate student and eleven postgraduate students, within which three of them have information security experience. Each of the group have three females and nine males.

5.3 Training of the Participants

A pre-scripted familiarisation tutorial was provided before the experiment. Participants from both Group A and B attended the same tutorial session. This was to ensure that they received equal knowledge related to the handling of security incidents. The participants were introduced to the Goal Structuring Notations and G.S.T.

5.4 Experiment Execution

The experiment was conducted on a one-to-one mode to provide any support needed during the whole process including the familiarization tutorial session, the experiment session and the post-experiment questionnaire session. During the familiarization tutorial session, the participant had unlimited time to study the material and to have any question clarified. The participants were allowed to refer to the tutorial document or notes. The participants were instructed to inform the experiment conductor if they had any trouble in understanding the questions. During the post-experiment questionnaire session, an informal interview was conducted to make sure their attitudes were consistent with the answers they have provided. They are also requested to write down their subjective feedback on the G.S.T.

6 Results - Prepare the data

6.1 Scoring Scheme for the Experiment Tasks

Sample answers for the experimental tasks were agreed on by two independent security experts.

6.2 Preparation for Task 1 - Open-ended questions

For Task 1, the answers expected were qualitative. The marking was based on the description of security issues and recommendations expected from the sample answers. The answers for each task were marked by two further independent experts (Rater A and B) using an agreed scoring scheme. The participants' answers were classified into four categories, which are "Correct", "Incomplete", "Wrong" and "Blank". A correct answer completely described the recommendation to support the given issue; incomplete answers show that the participant had a partial understanding of the recommendation, but lacked comprehension of an important aspect of it. Wrong answers showed that the participant did not understand a particular recommendation. Blank, no answer was provided at all. The following paragraph provides an example from task one:

The report identifies the security concern: "The IT Specialist was improperly given access to multiple data sources". An answer is marked as, *Correct*, if the participant states that the recommendation associated with this issue was to "Consider the conditions under which programmer level access may be granted

for research project". A correct answer completely describes the recommendation to support the given issue; *Incomplete*, if the answer is stated as "Ensure the access control is appropriately granted". Incomplete answers show that the participant had a partial understanding of the recommendation, but lacked comprehension of an important aspect of it; *Wrong*, if the answer provided is not relevant to a particular recommendation. *Blank*, if no answer was provided at all.

Each participant was free to use his or her own words to describe the recommendations in this part of the study. The group identifiers were removed so that Rater A and B marked the answers without knowing whether or not the participants had access to the G.S.T. diagram.

6.3 Preparation for Task 2 - Multi-choice questions

Task 2 used multi-choice questions to examine the participant's ability in understanding the compliance with the security objectives. Less subjectivity was involved in interpreting the answers. There can be more than one correct choice for each question and participants were asked to select all of the responses they believe were relevant to the questions. Below is an example.

What are the security recommendations for addressing the security objective "User Access Control"?

a. Develop and implement policies describing the conditions under which programmer level access may be granted for research purposes.

b. Effective procedures are implemented to determine compliance with authentication policies.

c. Attempts to log on with invalid passwords are limited. Use of easily guessed passwords (such as names or words) is prohibited.

d. None of the above

Correct answer: a, b

The sample answers were prepared by the independent security expert A. Each answer was classified as, *Correct*, *Correct but broad*, *Incomplete*, *Incomplete and broad*, *Wrong*, and *Blank*. A *correct* answer contained and only contained all the acceptable choices (e.g. a, b); *Correct but broad* contained all the acceptable choices, but also incorrect choices (e.g. a, b, c); *Incomplete* answers contained only some of the acceptable choices but not all (e.g. a). *Incomplete and broad* answers contained some of the acceptable choices and also other choices. (e.g. a, c); *Wrong* answers contained none of the acceptable choices (e.g. c).There was only one blank answer out of 144 responses; therefore we ignore this in the subsequent analysis.

7 Results - Analysis

7.1 Results for Accuracy (Task 1)

Out of a total number of 168 answers to the seven questions by 24 participants, three were left blank with one in Group A and two in Group B. During the debrief, the participants stated that, for the blank response, they could understand the questions but could not find the answer in the given materials. We ignore these blank answers in the subsequent analysis. Inter-rater reliability was checked for each question in Task 1, recall that these open ended questions were assessed by two independent raters. The results are listed in Table 2. Questions 1, 2 have achieved "almost perfect agreement"; Questions 3, 4, 5, and 6 have achieved "substantial agreement"; Question 7 has achieved "Fair agreement" [20].

	Inter-rater reliability check for each question						
Question No.	***1***	***2***	***3***	***4***	***5***	***6***	***7***
Kappa Value	0.85	0.80	0.75	0.57	0.72	0.78	0.50

Table 2. Inter-rater reliability for each question (Rater A and B)

A third independent security expert was invited to decide whether he agreed with Rater A or Rater B. The third security expert came to a 65.3% agreement with the Rater A and a 34.7% agreement with Rater B. Their interpretation was definitive for our analysis; in other words where there was disagreement between the first two assessments, the third rater decided which score was correct.

7.2 Comparing the performance of Task 1

Since the results are categorical data, we use cross-tabulation analysis to analyse the results. As is shown in Table 3, the results from the Cross-tabulation analysis (Table 3) show that 62.7% of the responses from Group A were correct, which is 17.6% higher than Group B. This might seem a relatively low level of accuracy. However, it is important to recall that our marking scheme was careful to distinguish between complete, perfect responses and partially correct or incomplete answers. The total percentage of Incomplete and Correct answer is 81.9% in Group A, which is 13.8% higher than Group B. As is shown in Table 4, the Chi-Square Test ($P = 0.048 < 0.05$) shows that these results are statistically significant. Therefore, hypothesis H1 "Participants will be better able to identify the recommendations and causes in security reports with the help of a graphical method than using text alone" is supported.

This result again shows that Group A has demonstrated a slightly higher level of comprehension than Group B. Therefore, hypothesis H1 "Participants will be better able to identify the recommendations and causes in security reports with the help of a graphical method than using text alone"is supported.

		Task			Total
Group A	Count	***Wrong***	***Incomplete***	***Correct***	
	% within Group	15	16	52	83
Group B	Count	18.1%	19.3%	62.7%	100.0%
	% within Group	27	18	37	82
Total	Count	32.9%	22.0%	45.1%	100.0%
	% within Group	42	34	89	165

Table 3. The performance of Task 1 using Cross-tabulation

	Chi-Square Tests		
	Value	***df***	***Asymp. Sig. (2-sided)***
Pearson Chi-Square	6.068a	2	.048
Likelihood Ratio	6.129	2	.047
Linear-by-Linear Association	6.032	1	.014
N of Valid Cases	165		

Table 4. Chi-Square Tests performance of Task 1 using Cross-tabulation

7.3 The Results for Accuracy (Task 2)

The results from the cross-tabulation analysis show that the participants from Group A achieved a 33.3% accuracy rate, which is 9.7% higher than Group B. The total percentage of Correct, Broad, Incomplete, and Incomplete but broad answer is 87.5%, which is 18.1% higher than Group B. As is shown in Table 5, the Chi-Square Test ($P = 0.038 < 0.05$) shows that these results are statistically significant. This multi-choice results were not due to coincidence. Therefore, hypothesis H2 "H2: Participants will be better able to identify the compliance with the security objectives in security incident report with the help of the G.S.T. than using Text-based Document alone" is supported in Task 2.

	Chi-Square Tests		
	Value	***df***	***Asymp. Sig. (2-sided)***
Pearson Chi-Square	10.140a	4	.038
Likelihood Ratio	10.449	4	.034
Linear-by-Linear Association	2.995	1	.084
N of Valid Cases	144		

Table 5. Chi-Square Tests performance of Task 2 using Cross-tabulation

7.4 The Results for Efficiency (Time)

The mean total time used by Group A was almost equal that in Group B; 47.3 versus 47.8 minutes. The total time taken across all tasks is not statistically significant ($P = 0.932 > 0.05$). Therefore, we can accept the null hypothesis that "the mean time taken to complete our experimental tasks using a textual security incident report and a textual report with a graphical overview are not significantly different."Hypothesis H3 is not supported. One interpretation of these results is that significant time is required to understand security incidents, irrespective of whether they are presented in graphical or textual format. However, this would require further empirical support to determine whether or not other graphical notations might lead to significant differences in the time taken to understand security incident reports. It is also important for further work to consider the learning effects that might be expected through repeated use of the graphical maps.

7.5 The Results for Task Load Index (TLX)

We used NASA's Task Load Index [19] to assess workload using a post-evaluation questionnaire. The t-test results show a significant difference ($P = 0.047 < 0.05$) in the first dimension of the task load index regarding "how mentally demanding was the whole task". With a mean value of task load, 12.75 versus 15.50, participants expressed a lower subjective level of workload in terms of "mentally demand"when using the G.S.T. The results for the other four dimensions of the Task Load Index are not significantly different. However, a more sustained analysis is required to replicate these findings across a wider range of workload measures and with a larger sample of potential users.

7.6 Subjective Feedback

In Group A, approximately half of the participants expressed some difficulty in understanding the text based Security Incident Report. Half of the participants reported that they have no difficulty in completing task 1 of Group A: identifying security elements from the security incident report with the help of the G.S.T. Group B reported a slightly higher level of understanding of the Security Incident Report. However, less than half of the participants suggested that they have no difficulty in completing task 1 of Group B: identifying lessons learned from the security incident report, and the rate is much lower than that of Group A. These subjective findings are consistent with the quantitative results in section 6.3.

The participants' answers to the open questions regarding the overall experience of using the graphical overviews suggested that a longer training session might have helped them to better prepare for the tasks. Several participants mentioned that they had experienced learning effects; their confidence in answering the questions increased as they worked their way through the questions. This finding from Group A reveals generally positive feedback for the G.S.T. Group B did not use the G.S.T. during the experiment. They were asked to

review the G.S.T. after the experiment and provide the feedback by completing Questionnaire Section 6 designed for Group B. Almost all of them suggested that they would have no difficulty in understanding the G.S.T. and agreed that the G.S.T. can help them better comprehend existing security incident reports. Two thirds of the participants reported their willingness to use the G.S.T. if they are requested to do a similar task in future. "It will help to understand terminologies security elements easily, less confusing, very structured and don't have to waste time, most importantly very easy to understand with less information". In summary, the participants overall experience with the G.S.T. is positive, however, questions remain about the ability of participants to apply the lessons from the report within their own organisation rather than answering directed questions about the contents of a security report.

8 Conclusions and Future Work

There have been numerous empirical studies to evaluate the utility and usability of graphical notations, including Entity-Relationship diagrams [21], UML[22] [23] etc. However, as far as we are aware, there have been no previous studies to assess the strengths and weaknesses of graphical notations to help transfer the lessons learned from previous security incidents. These studies are urgently needed as both the Obama administration and the European Commission have recently published proposals to support the mandatory reporting of security incidents across national critical infrastructures, including healthcare. In this paper, we have presented the results derived from an initial study into the use of Goal Structuring Notation (GSN) to represent and reason about the recommendations made in a report of a data confidentiality breach involving the US Veterans' Affairs Administration. We were able to show significant benefits from the use of a graphical technique in answering a number of comprehension questions when compared to the more conventional use of text-based incident reports. However, we could not demonstrate any significant benefits in terms of the time taken to complete our experimental tasks, nor could we demonstrate significant benefits when participants were asked to identify the compliance with the security objectives provided by multiple-choice questions.

It is important to stress that this was a preliminary study. The sample size was relatively small due to practical reasons: (1) the approach is new and people have little experience with security incident analysis; (2) the tasks were mentally demanding; (3) participation was voluntary. However, our work did yield important insights into the difficulties that engineers face when trying to understand the implications that previous security incident reports have for their own organisations.

Acknowledgment

The first author would like to thank the China Scholarship Council (CSC) for funding this research work.

References

1. Symantec: Internal security threat report 2011 trends. Volume 17. (2012)
2. Symantec: Internet security threat report 2013. Volume 18. (2013)
3. He, Y., Johnson, C.: Generic security cases for information system security in healthcare systems. In: Proceedings of the 7th IET International Conference on System Safety, incorporating the Cyber Security Conference, IET (2012) 1–6
4. Hadgkiss, J.: Computer security incident response teams: Exploring the incident learning capability. (2004)
5. Hadgkiss, J.: Computer security incident handling, step-by-step. (1997)
6. Administration, U.V.A.: Administrative investigation loss of va information va medical center birmingham, al. Volume Report No. 07-01083-157. (2007)
7. He, Y., Johnson, C., Renaud, K., Lu, Y., Jebriel, S.: An empirical study on the use of the generic security template for structuring the lessons from information security incidents. In: Proceedings of the 6th International Conference on Computer Science and Information Technology, IEEE Press (2014) 178–188
8. Kelly, T.P.: Arguing safety-a systematic approach to safety case management. (1998)
9. Dacey, R.F.: Federal Information System Controls Audit Manual (FISCAM). DIANE Publishing (2010)
10. He, Y., Johnson, C., Lu, Y., Lin, Y.: Improving the information security management: An industrial study in the privacy of electronic patient records. In: Proceedings of the 27th IEEE International Symposium on Computer-Based Medical Systems (CBMS 2014), IEEE Press (2014)
11. Commissioner, E.: Directive 2009/140/ec of the european parliament and of the council of 25 november 2009. (2009)
12. Craigen, D.: Formal methods technology transfer: Impediments and innovation. In: CONCUR'95: Concurrency Theory. Springer (1995) 328–332
13. Hinchey, M.G.: Confessions of a formal methodist. In: SCS. (2002) 17–20
14. Finney, K., Fedorec, A.: An empirical study of specification readability. Teaching and Learning Formal Methods, Academic Press, New York (1996)
15. Finney, K.: Mathematical notation in formal specification: Too difficult for the masses? Software Engineering, IEEE Transactions on **22**(2) (1996) 158–159
16. Carew, D., Exton, C., Buckley, J.: An empirical investigation of the comprehensibility of requirements specifications. In: Empirical Software Engineering, 2005. 2005 International Symposium on, IEEE (2005) 10–pp
17. Weber, R., Aha, D.W., Becerra-Fernandez, I.: Intelligent lessons learned systems. Expert Systems with Applications **20**(1) (2001) 17–34
18. Folmer, E., Bosch, J.: Architecting for usability: a survey. Journal of systems and software **70**(1) (2004) 61–78
19. Hart, S.G., Staveland, L.E.: Development of nasa-tlx (task load index): Results of empirical and theoretical research. Human mental workload **1**(3) (1988) 139–183
20. Landis, J.R., Koch, G.G.: The measurement of observer agreement for categorical data. biometrics (1977) 159–174
21. Shoval, P., Shiran, S.: Entity-relationship and object-oriented data modeling—an experimental comparison of design quality. Volume 21., Elsevier (1997) 297–315
22. Glezer, C., Last, M., Nachmany, E., Shoval, P.: Quality and comprehension of uml interaction diagrams-an experimental comparison. Volume 47. (2005) 675–692
23. Razali, R., Snook, C., Poppleton, M., Garratt, P., Walters, R.: Usability assessment of a uml-based formal modelling method. In: 19th Annual Psychology of Programming Workshop (PPIG'07), Citeseer (2007) 56–71