On The Enhancement of Data Quality in Security Incident Response Investigations

GEORGE GRISPOS

SUBMITTED IN FULFILMENT OF THE REQUIREMENTS FOR THE DEGREE OF
Doctor of Philosophy

SCHOOL OF COMPUTING SCIENCE

COLLEGE OF SCIENCE AND ENGINEERING
UNIVERSITY OF GLASGOW

MAY 2016

Abstract

Security incidents detected by information technology-dependent organisations are escalating in both scale and complexity. As a result, security incident response has become a critical mechanism for organisations in an effort to minimise the damage from security incidents. To help organisations develop security incident response capabilities, several security incident response approaches and best practice guidelines have been published in both industry and academia. The final phase within many of these approaches and best practices is the 'feedback' or 'follow-up' phase. Within this phase, it is expected that an organisation will learn from a security incident and use this information to improve its overall information security posture. However, researchers have argued that many organisations tend to focus on eradication and recovery instead of learning from a security incident.

An exploratory case study was undertaken in a Fortune 500 Organisation to investigate security incident learning in practice within organisations. At a high-level, the challenges and problems identified from the case study suggests that security incident response could benefit from improving the quality of data generated from and during security investigations. Therefore, the objective of this research was to improve the quality of data in security incident response, so that organisations can develop deeper insights into security incident causes and to assist with security incident learning.

A supplementary challenge identified was the need to minimise the time-cost associated with any changes to organisational processes. Therefore, several lightweight measures were created and implemented within the case study organisation. These measures were evaluated in a series of longitudinal studies that collected both quantitative and qualitative data from the case study organisation.

Acknowledgements

There are a number of people I would like to thank for their various contributions. Without their help and encouragement this research project and thesis would not have been possible. I am indebted to many colleagues and friends for their support in its development. I would like to extend a specific thank you to the following individuals.

First, I would like to thank my parents Michael and Maria Grispos for their patience, understanding, love, support and encouragement throughout my studies. I will eternally be grateful for all you have done to get me to this stage. I would like to thank Nicola Parpas, for not only reading pages and pages of my work in the past few years, but also for support, her belief in my abilities and her encouragement to always follow my heart and dreams.

I am indebted to my supervisors Dr. Tim Storer and Dr. William Bradley Glisson for patiently steering me and guiding me throughout the course of this work. I am grateful for their constructive comments they provided on not only various drafts of this thesis, but also several paper critiques throughout the course of this project.

I would also like to thank the School of Computing Science at the University of Glasgow, the Fortune 500 Organisation and the A.G. Leventis Foundation for supporting this research.

Lastly, I would like to extend my thanks to the individuals within the Fortune 500 Organisation, who have assisted me during my research, particularly with undertaking interviews and helping with various aspect of my experiments.

Declaration

I declare that this thesis has been composed by myself, that the research presented embodies the results of my own work and that it does not include work forming part of a thesis presented for a degree in this or any other University.

The author's original work presented in this dissertation has contributed to a number of publications that have been co-authored with Dr. Timothy Storer and Dr. William Bradley Glisson:

- **G. Grispos**, W.B. Glisson, and T. Storer (2015). Security Incident Response Criteria: A Practitioner's Perspective. 21st Americas Conference on Information Systems (AMCIS 2015), August 13-15, 2015, Puerto Rico, USA.

- **G. Grispos**, W.B. Glisson, and T. Storer (2014). Rethinking Security Incident Response: The Integration of Agile Principles. 20th Americas Conference on Information Systems (AMCIS 2014), August 7-10, 2014, Savannah, Georgia, USA.

Table of Contents

List of Tables

List of Figures

Chapter 1

Introduction

This chapter introduces the research background, the motivation guiding this work, as well as the thesis statement and research questions. The chapter is divided into five sections. Section 1.1 presents background information that guided the research project, while Section 1.2 provides the motivation for examining incident learning and the quality of data within security incident response. Section 1.3 defines the thesis statement and the research questions. Section 1.4 discusses the research contributions and Section 1.5 provides an overview of each chapter in the thesis.

1.1 Background

Recent industrial reports [1, 2] have suggested that organisations are detecting an increasing number of security attacks on their infrastructure. The 2015 PricewaterhouseCoopers Information Security Breaches Survey [1], reported that 90% of large organisations and 74% of small business in the United Kingdom had detected a security attack in the past twelve months. The report goes on to state that 59% of respondents expect the number of security attacks to increase in the next year [1]. This increase comes at a great financial cost. The 2014 Ponemon Cost of Cybercrime [2] report estimates that financial losses attributed to security attacks cost United States-based organisations an average of $12.7 million, an increase of 9% since the previous year.

In an effort to reduce the likelihood and impact of security attacks, many organisations implement information security programs [3, 4]. The purpose of an information security program is to protect the confidentiality, integrity and availability of application data, information systems and computer networks, while ensuring that any legal and regulatory requirements are also fulfilled [3]. A number of information security standards outline security controls, which can be implemented within an organisation as part of an information security program. The objective of implementing these security controls is to prevent or mitigate further security attacks [3]. However, information security programs and implemented security controls cannot guarantee complete protection of an organisation's assets. Consequently, many information security standards also advocate the development of security incident response capabilities. The purpose of these capabilities is to help organisations manage the investigation and recovery from security attacks which have become incidents [5,6].

The objective of security incident response is to minimise the damage from a security incident, and to allow an organisation to ultimately learn about the cause of the incident and how it could be prevented in the future [7]. Several organisations, such as the National Institute of Standards and Technology [8], the European Network and Information Security Agency [9] and the International Organization for Standardization [10], have published guidance on security incident investigation and recovery techniques. In addition to these industry practices, numerous academic researchers have also developed several security incident approaches [7,11,12].

A typical security incident response process consists of six phases: *preparation*, which leads to the *detection* of an incident, followed by its *containment* which, in turn, allows security incident response teams to *eradicate*, *recover* and then, potentially, provide *feedback* information into the preparation stage. The final phase within many security incident response approaches is the 'feedback' or 'follow-up' phase [7,8,13]. It is within this phase where an organisation attempts to learn from a security incident with the aim of improving its overall information security posture [7, 8, 13]. Within security incident response, incident learning is usually accomplished through a series of formal reports, meetings and presentations to management after the closure of an investigation [13]. Lessons learned from a security investigation can include information about enhancements to existing security controls, the identification of additional tools for investigation and analysing whether changes are required to existing security incident response process and procedures [7,14].

Incident learning has been used in several domains including aviation, safety and healthcare, with the aim of preventing similar incidents in the future [15]. In order to undertake investigations within these industries, investigators require that any necessary data is made available to examine the underlying causes of an incident [15]. Stephenson [16] argues that this is also true within security incident response, where detailed data can help establish root causes that have contributed to an incident. However, Stephenson [16] adds that in many organisations, obtaining detailed data about a security incident can be difficult because security investigations are costly and often require great expertise to conduct. Similar problems have also been noted in investigations within the transport safety domain where concerns were raised about the quality of data derived from incident investigations [17].

1.2 Motivation

Although security incident response best practices stress the importance of post-incident learning, researchers have observed that many organisations find it difficult to learn from security incidents [18–20]. A contributing factor to this problem is that many organisations tend to focus on eradication and recovery and less on security incident learning [18–20]. As a result, there is the possibility, that similar to the transport safety domain, the quality of data derived from security investigations may be unfit for in-depth security incident learning. Without access to 'good' quality data, a security incident response team could find it difficult to undertake a lessons learned meeting to engage in security incident learning.

An exploratory case study was undertaken within a Fortune 500 Organisation, with the aim of investigating the quality of data within security incident response. The results from this study identified that security incident response could benefit from improving the quality of data generated from and during security incident investigations, which in turn would help enhance post-incident learning. Although numerous practices and tools [21–24] have been developed to help improve the quality of data in various processes, such as databases, very little research has examined what practices and tools can improve the quality of data in security incident response. Therefore, specific lightweight practices were proposed as one approach to improving the quality of data in security incident response.

There are two specific characteristics of lightweight practices that make them a prime candidate to help enhance the quality of data within security incident response. Lightweight practices and methodologies have been proposed in the software development literature as approaches that strive for simplicity and that are easy to follow, with very few rules [25, 26]. Time is a critical factor in security incident investigations [7, 18]. As a result, security incident handlers may not necessarily respond positively to data quality improvement practices, which take time and effort to use during an investigation when eradication and recovery is the focal point. Therefore, lightweight practices that are simple to use could fit into a time-focused environment, such as security incident response. Lightweight practices and methodologies are also people-orientated rather than process-orientated and attempt to work with people's nature, rather than against it [26]. With individuals considered to be one of the key factors in the success or failure of security investigations [27, 28], this lightweight practice characteristic, with its focus on people, lends itself strongly to security incident response. Based on these motivations, this research investigated the efficacy of integrating a set of lightweight practices into security incident response processes. Their use was observed to investigate if the use of the practices helped to enhance the quality of data generated from and during security incident investigations.

1.3 Thesis Statement

Organisations need to strengthen the quality of data produced during security incident response investigations in order to enhance learning from security incidents. Therefore, the hypothesis is as follows:

The quality of data generated during a security incident investigation can be enhanced through the application of specific lightweight measures.

The lightweight measures were:

- A process for categorising security incidents according to a well-defined taxonomy, allowing a security incident response team to monitor trends in incident types.
- A security incident investigation record template that focused investigation efforts on data that provided value to a security incident response team.
- A security incident investigation dashboard to enhance data quality transparency and to help a security incident response team to identify where data quality needs to be improved.
- The integration of retrospectives into a security incident response process as a method of validating data that has been collected within security investigations and to collect and enhance data that has been missed.
- The use of a root cause analysis framework to guide the security incident investigation process and to generate higher quality data for subsequent lessons learned.

The following research questions concern the above hypothesis:

RQ1: What data is generated by a real-world security incident response team?

RQ2: What challenges and problems do a security incident response team face when attempting to learn from information security incidents?

RQ3: What effect did the application of the described measures have on the data generated by the security incident response process?

1.4 Research Contribution

The research presented in this thesis makes several contributions to the body of knowledge that include:

Exploratory Case Study of Security Incident Response in a Fortune 500 Organisation

The exploratory case study of the security incident response landscape in a Fortune 500 Organisation contributes to the understanding of the variety of challenges and problems that can hinder security incident learning within an organisation. An in-depth analysis of the organisation's security investigation records from the perspective of *accuracy, completeness* and *consistency* identified that the quality of data within security incident response is an area which needs to be addressed to help security incident learning. The findings support and enhance the results of previous studies of organisations through an analysis of the organisation's security investigation records, documented processes and semi-structured interviews with relevant practitioners. The case study results provided the initial body of data, which later guided the research through the identification of challenges and problems with incident learning within security incident response. Undertaking a case study with a real-world security incident response team and investigating multiple data sources allowed the research to be based on real problems encountered within industry.

Evaluation of a Security Incident Response Taxonomy and Investigation Record Template

The experiment contributes to the body of knowledge in two ways. First, the experiment evaluates the effect of employing a revised security investigation record that is used to collect data that has value to a security incident response team with regards to incident learning. Second, the experiment also evaluates the efficacy of a taxonomy of security incident types with the purpose of removing ambiguity surrounding incident type identification. The experiment also contributes to the body of knowledge with regards to identifying culture and management challenges associated with collecting information of value by a security incident response team.

Evaluation of the Integration of Retrospectives into a Security Incident Response Process

The experiment evaluates the effectiveness of using retrospectives as a method for validating and enhancing data collected during security investigations. In addition, the experiment also contributes to the understanding of how to design and integrate retrospectives into a security incident response team. The results from the experiment also suggest that retrospectives could be implemented in a 'process culture' organisation, in order to receive feedback from a security incident response process. Furthermore, the experiment evaluates the use of meta-retrospectives (a retrospective of retrospectives) as a method for tracking changes that have been implemented back into an organisation following a security investigation.

Empirical Experiment involving Security Incident Response Dashboard

The experiment evaluates the effect of using dashboards to enhance data quality transparency in security incident response. The experiment contributes to knowledge by examining how dashboards can be used to identify where data in security incident investigation records needs to be enhanced. The experiment also contributes to the knowledge of how to develop and implement a security incident response dashboard within organisations. The lessons learned from developing and implementing the dashboard provides insights into the integration of security incident response dashboards within organisations. The results from the experiment also highlight culture challenges with regardless to deploying security incident response dashboards in a regulatory-intensive organisation.

Evaluation of a Root Cause Analysis Framework

The contribution of this experiment is threefold. First, the application of the framework to historical investigation records investigates the extent to which a root cause analysis method can help produce enhanced lessons learned when used within security incident response. Second, the design and implementation of the root cause analysis framework contributes to the knowledge by providing a method for conducting a root cause analysis within security investigations. The lessons learned from the implementation of the framework contributes to the knowledge by providing insights into the process and where problems lie, which can be improved upon in future initiatives. Third, the experiment contributes to the knowledge by providing an approach for enhancing data quality through a root cause analysis method during a security investigation. Moreover, the results from the application of the framework to 'live' investigations contributes to the body of knowledge by examining organisational culture changes required for an effective root cause analysis.

1.5 Thesis Overview

The objective of this research was to investigate if the integration of lightweight practices into security incident response can help improve data quality from a security incident response process. A thesis chapter breakdown is provided detailing the various domains of research conducted to achieve this objective.

Chapter two presents the research methods that were used in the construction of this thesis. The research methods used included a review of the relevant literature, exploratory and explanatory case studies within a Fortune 500 Organisation, interview surveys and action research.

Chapter three examines the relevant literature concerning information security management, security incident response processes and the challenges of undertaking security incident learning within organisations. The chapter also reviews the term 'data quality' and examines how 'data quality' can affect security incident learning. The chapter concludes by setting the scope for the remainder of the thesis.

Chapter four presents an exploratory case study of the security incident response landscape in a Fortune 500 Organisation. The case study identifies the challenges that hinder security incident learning within an organisation. The chapter includes a discussion of the organisation's security incident response landscape, an analysis of its incident response documentation, as well as the organisation's information security incident response database. The chapter also presents the results from interviews conducted within the organisation. The results from the exploratory case study identified initial challenges, which guided the rest of the research described in this thesis.

Chapter five presents an experiment to evaluate the use of a security incident taxonomy and revised security investigation record template within the Fortune 500 Organisation. The purpose of the experiment was to evaluate if the taxonomy and investigation record template help to collect data, which has value to the security incident response team within the organisation with regards to incident learning.

Chapter six presents an experiment to evaluate the integration of retrospectives into the Fortune 500 Organisation's security incident response process with the aim of validating data generated during security investigations and enhancing data collection that was missed during the initial investigation. The chapter also presents meta-retrospectives as a method for tracking whether changes identified during lessons learned are in fact implemented within an organisation.

Chapter seven presents an experiment to evaluate the use of a dashboard with the security incident response team in the Fortune 500 Organisation. The purpose of the experiment presented in this chapter was to determine if the dashboard helps to enhance data quality transparency, i.e. identify security investigation records where data quality needs to be improved.

Chapter eight presents an experiment to evaluate the use of a root cause analysis framework, as a method for enhancing the quality of data during a security investigation. The aim of this experiment was twofold. First, the experiment was used to evaluate if using a root cause analysis can help to produce enhanced lessons learned derived from a security investigation. Second, the experiment was used to demonstrate how a root cause analysis can help enhance the quality of data captured from an investigation.

Chapter nine presents the conclusions to the research questions detailed in the introduction and discusses further work.

Chapter 2

Research Methodology

This chapter describes the research methods used in the construction of this thesis. These methods included literature reviews, case studies and action research. The chapter is structured as follows. Section 2.1 provides the justification of the research approach and presents the research methods used in this thesis. Sections 2.2 through 2.6 explain the use of literature review, case study research, interviews, documents analysis, observation and action research respectively, and their application to the research in this thesis. Section 2.7 summarises the chapter.

2.1 Justification of Research Approach

Several researchers [29, 30] have argued that in order to understand the work of information technology professionals it is necessary to study the tools used by these professionals, as well as the surrounding social and cognitive processes. In the context of software engineering, Easterbrook, et al. [29] has developed a taxonomy of methods for this purpose, with each method suited to answering different forms of research questions. These methods include controlled experiments, case studies, survey research, ethnographies and action research [29]. Oates [30] supplements these research methods by arguing that a literature review is important in order to relate research findings to previous work. Oates [30] goes on to present a detailed framework for conducting information system research that is derived from a literature review and consists of six research strategies that include survey, design and creation, experiment, case study, action research and ethnography.

The research methods proposed to investigate the tools, as well as the surrounding social and cognitive processes within software engineering and information systems lend themselves strongly to security incident response research. Researchers within security incident response [18, 27] have called for further research into the socio-organisational perspectives of security incident response, as well as the relationship between security incident response teams and their wider organisations. Therefore, out of the research methods proposed above, four have been used in the construction of the work described in this thesis: *literature review, case studies* and *action research*.

The initial research described in the thesis was guided by a literature review. The purpose of the literature review was twofold. First, the literature review was used to identify the high-level research questions that the proposed research will look to address and helped to define the conceptual framework. Second, the literature was used to relate the research findings presented in this thesis to previous work. This allowed the author to discuss strengths and weakness in previous work, identify key issues and challenges troubling the specific research community, along with identifying theories that might explain the research findings [30].

The case study is a research method which involves an in-depth and detailed examination of an organisation, a department, an information system, a specific team or project ('the case'), as well as its related contextual conditions [30]. In this research, the case study method was used to investigate why security incident learning is a problem within organisations. The case study research method was chosen over other research methods including the survey approach and the controlled experiment. While the survey approach allows a researcher to obtain a wide but shallow view of many instances of the case under investigation, the method is unlikely to assist the researcher who wishes to obtain a richer context about a particular problem [30, 31]. In contrast, the case study method is better suited when a researcher wants "to obtain a rich and detailed insight into the 'life' of a case including its complex relationships and processes" [30]. Moreover, during a controlled experiment researchers have to separate a phenomenon from its context, in order to establish if research outcomes have been caused by changing the independent variable [30]. While variables can be added and removed during a controlled experiment, a researcher will find that they will have limited control over variables during longer experiments [32]. This can make the controlled experiment method unsuitable for studying an organisation because of the limited control over the variables within an organisation. For example, business processes can change, employees can leave the organisation and new employees can be hired. Therefore, the case study research method is more appropriate for situations where researchers have little control over events and where results produced from the study better reflect the real-world [30].

Although the case study approach was selected for use in this research, it must acknowledged that several researchers have discussed the weakness of the approach in the literature. These concerns include difficulties in generalising results [30–32]; difficulties in obtaining access to people, documents and settings [30]; the presence of the researcher affecting how people behave in their natural setting [30]; and case studies being labelled as being too long, difficult to conduct and producing a massive amount of documentation [32].

In addition to the case study method, 'action research' was used to investigate real-world problems while simultaneously examining the experience of solving these problems within a studied organisation [33]. The action research method was chosen for this research project because it can be used to concentrate research efforts that are relevant to improving people and organisations in the real-world [33]. The action research method was used in the research process to guide the implementation and investigation of lightweight measures into security incident response as a means of improving the quality of data. However, one of the biggest criticisms of action research is that it can resemble consultancy [30]. As a result, researchers need to ensure knowledge is either created or validated from the action research process [30]. The use of the literature review, case study and action research methods will be discussed in more detail below.

2.2 Literature Review

The relevant literature was examined from a variety of sources including industry whitepapers and best practice documentation, journals, conference proceedings and relevant textbooks. The main body of relevant literature is discussed in Chapter three. However, relevant literature is also discussed prior to each experimental chapter where appropriate.

Chapter three presents an overview of information security within organisations and introduces security incident response and its position within information security. More specifically, the chapter presents the theoretical foundations from the literature, including how lessons learned are typically developed from security investigations and how incident learning from such investigations is a problem which has been identified in several organisations. The chapter also introduces the concept of data quality and examines relevant literature concerning data quality improvement processes.

Relevant literature is also discussed in Chapters six and seven. Within Chapter six, relevant literature introduces retrospectives and discusses how they have been used in the software development community. In Chapter seven relevant literature introduces dashboards and provides an overview of previous research focused on developing dashboards within information security.

2.3 Case Study

The case study is a research method that involves an in-depth and detailed examination of an organisation, a department, an information system, a specific team or project ('the case'), as well as its related contextual conditions [30]. Yin defines case study research as:

> "an empirical inquiry that investigates a contemporary phenomenon (the 'case') in-depth and within its real-world context, especially when the boundaries between phenomenon and context may not be clearly evident" [32].

Oates [30] adds that the objective of a case study is to "obtain a rich, detailed insight into the 'life' of the case and its complex relationships and processes". Oates goes on to define three types of case studies [30]:

- **exploratory case studies**, which define the questions or hypotheses to be used in subsequent case studies;
- **descriptive case studies**, which can result in a detailed analysis of a particular phenomenon and its context; and
- **explanatory case studies**, which try to explain why certain events have occurred as they did or why particular outcomes were obtained.

Alternatively, Hancock and Algozzine define three stages in designing case studies [34]: *defining a 'case'*, *selecting a case study design* and *choosing to use theory in the design*. Another prospective of a case study process is provided by Yin [35] as:

1. Case study design: objectives are defined and the case study is planned
2. Preparation for data collection: procedures and protocols for data collection are defined
3. Collecting evidence: execution with data collection on the studied case
4. Analysis of collected data
5. Reporting

In the case study undertaken for this research, the first stage is where the initial negotiation and agreement with a Fortune 500 organisation took place to determine their involvement as the 'case'. Therefore, the research described in this thesis involved a 'special' single case study [30, 35, 36]. The Fortune 500 organisation is a multi-national organisation within the financial services sector and therefore under a large amount of regulation, particularly towards the reporting and investigation of information security incidents. This makes the Fortune 500 organisation an ideal candidate to explore security incident learning challenges within organisations. The organisation was selected based on previous academic collaborative relationships with the author's supervisor. The Fortune 500 organisation is representative of other organisations in the financial category in terms of interest in security incident response, regulatory obligations towards security incident reporting, size, bureaucracy and resource constraints. The opportunity to undertake research with the Fortune 500 organisation allowed the project to be based on real-world problems and to examine potential solutions in a real-world environment. The Fortune 500 organisation case study can be considered a 'special case' in the sense that a unique chance arose to study security incident learning in a large organisation [36].

An internship with the organisation was agreed and a confidentiality agreement was signed that restricted disclosure of the name of the organisation. With this in mind, the names of any organisational documents and processes discussed in this and other chapters have been altered and the results of any data collected is presented anonymously. The research agreement consisted of first examining the organisations security incident response practices towards security incident learning, and then second making recommendations on how to strengthen processes for security incident learning data. Although the organisation was supportive of the research initiative, there was no guarantee that they would accept the recommendations or implement the changes into the production environment. The final *gatekeeper*, who would decide on the improvement initiatives was the organisation's Head of Information Security (hereafter referred to as the *industrial sponsor*). Once these issues had been agreed upon, the project moved to its second stage.

An initial meeting took place where the author explain the theoretical perspective surrounding security incident learning challenges within organisations to the industrial sponsor. The organisation, as part of the research project, agreed to allow the author access to individuals and resources within the organisation to undertake the research described in the thesis. An exploratory case study was then conducted at this time to acquire a more in-depth understanding of the problem. An exploratory case study is used to define the questions or hypotheses to be used in subsequent case studies [30]. The *unit of analysis* for this exploratory case study was the organisation's Information Security Incident Response (ISIR) team. This means that the exploratory case study can be considered to be a holistic case study [35].

The objective of the exploratory case study was to define the research questions in the context of why security incident learning is a problem within organisations. The results from the exploratory case study were then used to define and guide the research questions and hypotheses in the remainder of the research described in this thesis [30]. The case study undertaken within the organisation can be classified as longitudinal study, because it was conducted over a long period of time [30]. The exploratory case study was undertaken from May 2013 to December 2013. Longitudinal studies are useful when attempting to examine chronological timelines of events or changes in real-world organisations over time, as was the case in the Fortune 500 organisation case study [35].

A variety of different data sources were used in the generation of data in the case study.

Oates [30] defines data generation as the "means by which you produce empirical data or evidence, which can be either quantitative or qualitative". Three data generation methods were used during the course of the exploratory case study, which are presented in Table 2.1 [30].

Generation Method	Definition
Interview	A particular kind of conversation between two people where, at least at the beginning of the interview if not all the way through, the researcher controls both the agenda and the proceedings and will ask most of the questions.
Documents	that already exist prior to the research and documents that are made solely for the purpose of the research task.
Observation	Watching and paying attention to what people actually do, rather than what they report they do.

Table 2.1: Data Generation Methods

The third stage implements the data generation methods into the organisation. As a result, of using two or more data generation methods, 'Method Triangulation' was used during the case study [30]. At the end of the exploratory case study time period, the data collected from the three data generation methods were evaluated as part of the fourth stage. The results from the exploratory case study are then reported in the final stage.

At this point it must be acknowledged that despite the advantages of the case study method, there are concerns in the literature regarding its reliability and validity [37–39]. The research presented in this thesis uses three different data generation methods (interview, documents and observation), which Oates defines as "Method Triangulation" [30]. Furthermore, the research in the thesis also uses three research strategies (literature review, case study and action research), which Oates defines as "Strategy Triangulation" [30]. The implementation of "Strategy Triangulation", along with the use of multiple data generation methods, can help validate or highlight differences in research findings and data collected, as well as helping to enhance research validity and the reliability of collected data [30].

2.4 Interviews

Interviews were used during the course of this research to gather qualitative data about various aspects of the research project. Interviews were selected over questionnaires for two reasons. First, they provide the interviewer with more control over the question being asked and second, interviews are better suited to keeping the interviewee focused and on track to completion [30].

In total, three sets of interviews were undertaken, all of which were conducted within the Fortune 500 Organisation. The interviews utilised a semi-structured interview approach, which means that there were a set number of interview questions that were read to each participant [30]. However, if the participant wished to talk about other issues related to the interview question, this was allowed to go on to completion. When the participants completed voicing their thoughts and answered the question that had been put forth, they were directed to the next question. The instrument for all three interviews consisted of a combination of open-ended and closed questions [40].

To mitigate researcher bias in terms of reliability and viability, the interview instruments were validated by two security professionals [40]. An information security manager and a senior security analyst validated the instruments by taking the interviews and providing feedback. The feedback from these individuals ranged from simplifying open-ended questions to adding response options to closed questions. In all three cases, this validation was only conducted once due to time constraints. At the start of the interviews, participants were read a statement thanking them for participating, explaining the reason for the research and reassuring respondent anonymity. The interviews were conducted in conference rooms and participant's desks within the organisation. All responses to the individual questions were initially recorded by hand. The hand written results for all three sets of interviews were digitally recorded soon after the interview were completed, typically within an hour. The results were then examined by hand to identify trends, patterns, and anomalies.

The first set of interviews was undertaken as part of the exploratory case study within the Fortune 500 Organisation. The purpose of these interviews was to obtain an in-depth understanding of how security incident response is perceived and undertaken within the organisation from the perspective of practitioners. The interviews were also used to explore challenges to conducting security incident response within the organisation and examine data gathering and incident learning from the overall security incident response process perspective. The survey instrument for the exploratory case study can be found in Appendix A. The questions within the instrument were derived from themes identified in industrial white-papers [41–43] and academic papers [18–20, 44] related to security incident learning challenges in organisations. These interviews were conducted between November and December 2013. Initially, the interviews were conducted with three individuals identified through the organisation's security incident response process as the 'Primary Incident Handlers'. A further twelve individuals were then identified and interviewed based on answers from the initial respondents' interviews. All fifteen individuals are members of the organisation's information security unit and have at some point, been involved in the investigation and handling of a security incident. The University of Glasgow, College of Science and

Engineering Ethics Committee approved the interview instrument under approval number CSE01330.

Two further sets of interviews were also undertaken during the explanatory case study. The questions used in these two sets of interviews can be found in Appendix B and Appendix C. These interviews were used to ascertain the impact of the lightweight measures implemented within the organisation and to further explore certain phenomena identified from the quantitative data analysis. As a result, the questions within these two instruments were derived from themes that emerged from the quantitative data generated through the use of the lightweight measures. The interviews were conducted with seven individuals. These seven individuals were at the time of the interviews, identified as the organisation's 'security incident response team' and consisted of six 'primary incident handlers' and the security incident response policy owner. Although the policy owner did not actively participate in or use the implemented measures, their opinion was sought in their capacity as an information security manager who is ultimately responsible for the security incident response process within the organisation. These two sets of interviews were approved by the University of Glasgow Ethics Committee under approval numbers 300140061 and 300140162.

2.5 Document and Observation Analysis

In order to acquire an understanding of the Fortune 500 organisation's context during the exploratory case study, the organisation's internal documentation repository was examined. The primary advantage of using document-based data is that it can be obtained quickly and cheaply because documents already exist and are readily available in the organisation [30]. However, researchers need to evaluate documents to ensure that the author of the particular document has not introduced any bias and that the content of the document can be trusted [30]. Obtaining access to the documentation repository was negotiated within the confidentiality agreement discussed above, and involved the author examining the repository from May 2013 to December 2013 to identify and analyse *found documents* [30] related to security incident response processes.

The author was granted access to documentation that would normally be available to individuals within the organisation's Information Security Incident Response team. Documents, which were considered sensitive, confidential and only available to Management, were outside the scope of the analysis. The 'found documents' used in the case study were all signed-off by management before they are stored in the documentation repository. In this sense, the author considered the content of the 'found documents' to be authentic and trustworthy [30]. The documents were analysed using theme analysis, which is a qualitative technique that allows a researcher to examine different topics covered in a variety of documents [30]. In this case, the topics examined were related to security incident response processes (such as the tasks incident handlers are expected to undertake within the organisation) and learning from security incidents.

Observation is a data generation method that can be used by researchers to investigate what people *really* do, as opposed to what they report they do, when queried [30]. In the exploratory case study, observation was concentrated on individuals considered to be the Information Security Incident Response team in the organisation. While the document analysis was used to investigate how security incident response is *expected* to be undertaken within the organisation, observation was used to explore if incident handlers deviate from this process, i.e. how security incidents are *really* managed and handled in the organisation. This involved the author shadowing the security incident handlers in an *overt manner*, when an incident was reported to the team. Field notes were then documented on how the incident was managed and handled by the specific individual(s). The author did not participate in the management or handling of the security incident and therefore can be considered a *complete observer* [30]. It must be noted that observations were limited to security incidents where the author was present in the organisation, and in some cases the author could not observe incident meetings for security incidents considered sensitive in nature.

2.6 Action Research

Davison, et al. [45] define Action Research (AR) as a research method, which attempts to solve a real-world problem, while simultaneously studying the experience of solving the problem. Baskerville adds that AR “assists in practical problem solving and expanding scientific knowledge” [33]. The AR method was used to guide the implementation of the lightweight measures in the Fortune 500 Organisation from February 2014 until March 2015. The implementation of these measures was guided using the AR cycle discussed by Baskerville [33] which consists of five phases: diagnosing, action planning, action taking, evaluating and specifying learning.

Diagnosing is concerned with the identification of the problems that are the underlying causes of an organisation’s desire for change [33]. The exploratory case study provided information describing the ‘diagnosis’ and helped to identify the initial challenges in security incident learning within the organisation. These challenges, along with the exploratory case study are discussed in Chapter four.

Baskerville [33] describes *action planning* as researchers and practitioners collaborating on planning organisational actions that should relieve or improve the problems identified in the diagnosis stage. Within this research project, the author held numerous meetings with the industrial sponsor within the Fortune 500 Organisation. Initially, the meetings involved the author presenting the findings from the exploratory case study and highlighting potential enhancements identified within the organisation's security incident response process. The author then explored and identified opportunities to improve security incident response within the organisation. Opportunities were identified based on available access to specific resources under the management of the industrial sponsor. Several recommendations were made to the organisation and some of these recommendations were accepted. These recommendations involved the lightweight measures described in previous sections. The accepted recommendations were implemented into the production environment. Chapters five through eight, report on the individual recommendations implemented within the organisation. These implemented changes were the 'action taking' phase in Baskerville's AR cycle. In Chapter five, two main modifications were recommended, a security incident classification taxonomy and a revised security incident response investigation record template. Within Chapter six, retrospectives and meta-retrospectives (a retrospective of retrospectives) were proposed and integrated within the security incident response process, while in Chapter seven; a dashboard was designed and implemented within the organisation. In Chapter eight, a framework was designed and developed to assist and guide root cause analysis within the organisation's security incident response process. It must be noted that during the course of the research project, organisational changes occurred in the security incident response team, and the Information Security team within the organisation. Shortly after the implementation of the lightweight measures, the Information Security team was restructured, which involved merging existing teams into new units. The security incident response team was not affected by this reconstructing, however the team was expanded on two occasions after the lightweight measures were implemented. Two new employees joined the security incident response team at the end of 2015.

In order to *evaluate* the implemented changes, quantitative and qualitative data analysis was undertaken. Quantitative analysis involved analysing a variety of artefacts created during the course of the case study, which included security investigation records, as well as data and logs created from security incident response activities. In addition, qualitative data was collected through semi-structured interviews. The use of both quantitative and qualitative data helped to validate or highlight differences in findings and data collected [30].

The results from the quantitative and qualitative analysis provided the input into the activity of *specifying learning* from the implemented recommendations [33]. This involved examining the results from the quantitative and qualitative analysis and identifying knowledge which could be considered important to the security incident response and information security communities for dealing with similar future research settings [33]. Figure 2.1 presents a timeline of the research activities involved in the construction of this thesis. The timeline includes the implementation of the exploratory and explanatory case studies, the modifications implemented within the organisation, as well as the various data generation methods used to

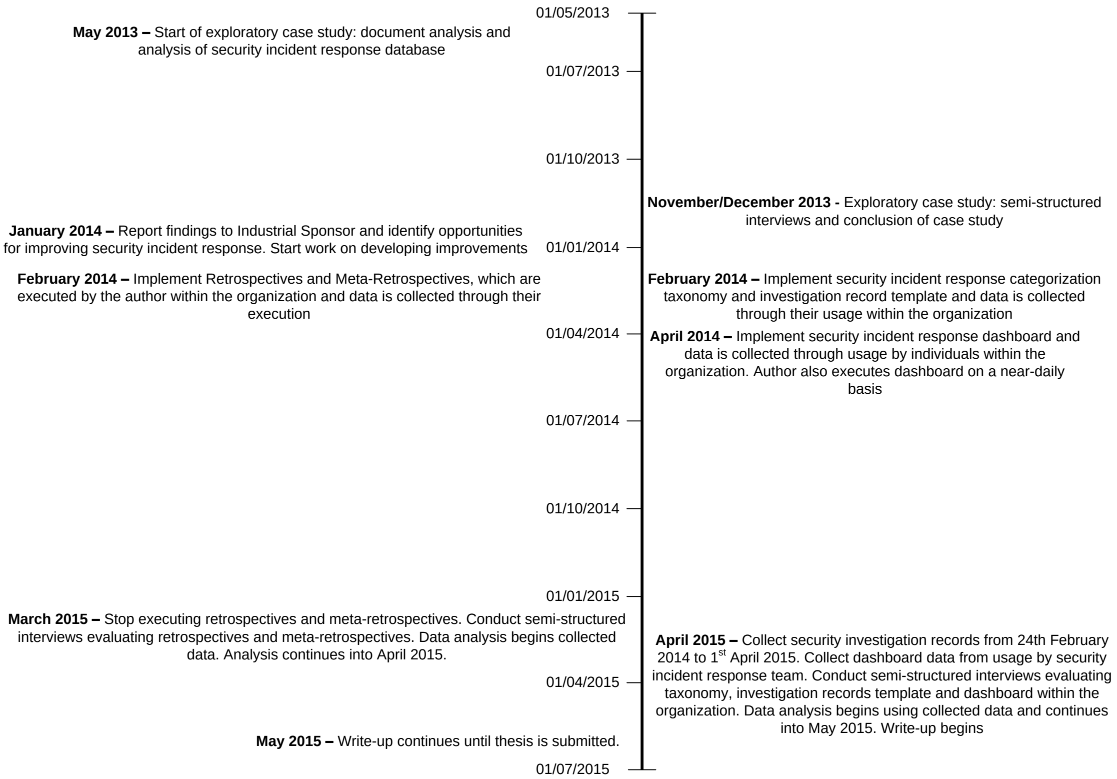

May 2013 – Start of exploratory case study: document analysis and analysis of security incident response database

November/December 2013 - Exploratory case study: semi-structured interviews and conclusion of case study

January 2014 – Report findings to Industrial Sponsor and identify opportunities for improving security incident response. Start work on developing improvements

February 2014 – Implement Retrospectives and Meta-Retrospectives, which are executed by the author within the organization and data is collected through their execution

February 2014 – Implement security incident response categorization taxonomy and investigation record template and data is collected through their usage within the organization

April 2014 – Implement security incident response dashboard and data is collected through usage by individuals within the organization. Author also executes dashboard on a near-daily basis

March 2015 – Stop executing retrospectives and meta-retrospectives. Conduct semi-structured interviews evaluating retrospectives and meta-retrospectives. Data analysis begins collected data. Analysis continues into April 2015.

April 2015 – Collect security investigation records from 24th February 2014 to 1st April 2015. Collect dashboard data from usage by security incident response team. Conduct semi-structured interviews evaluating taxonomy, investigation records template and dashboard within the organization. Data analysis begins using collected data and continues into May 2015. Write-up begins

May 2015 – Write-up continues until thesis is submitted.

Figure 2.1: Timeline of Industrial Research Project

collect data in the studied organisation.

2.7 Summary

This chapter presented the research methods that were used in the construction of this thesis. The research methods used included a review of the relevant literature, explanatory and explanatory case studies within a Fortune 500 Organisation, interview surveys and action research.

Chapter 3

Literature Review

This chapter presents a review of the relevant literature and provides the theoretical background that the thesis draws upon. The chapter is structured as follows. Section 3.1 defines information security and introduces background security concepts such as security threats, vulnerabilities and risk management. Section 3.2 first examines the evolution of information security management within organisations from the 1980's to the present time, then introduces the concept of Information Security Management Systems (ISMS) and presents an overview of information security standards. Section 3.3 introduces information security incident management and examines the various definitions of the term 'security incident' available in the literature. The section also presents and reviews the phases within a typical security incident response process and then reviews how security incident response is undertaken in various organisations. Section 3.4 presents related work on post-incident learning, which includes an analysis of the literature to identify how post-incident learning is undertaken in organisations, as well as methods for learning from security incidents. Section 3.5 introduces the concepts of data, information, knowledge and wisdom in context of this research. The section also defines the term 'data quality' and speculates on the extent to which data quality could be a problem within security incident response investigations. Section 3.6 introduces organisational culture presents an overview of different types of organisational culture. Furthermore, the concept of organisational learning is discussed, as well as previous work on organisational learning in the security incident response domain. Section 3.7 summarises the theoretical motivations, which have guided the remainder of the research described in this thesis.

3.1 Information Security Fundamentals

The widespread use of electronic data processing, along with the emergence of business conducted through the Internet has fuelled the need for methods to protect business information and information systems. As a result, many organisations have recognised the importance of implementing effective information security programs [46,47]. Information security concerns the protection of information and information systems from unauthorised access, modification or destruction. The International Organization for Standardization (ISO) state that:

> “information security involves the application and management of appropriate security measures that involves consideration of a wide range of threats, with the aim of ensuring sustained business success and continuity, and minimising impacts of information security incidents” [48].

The ISO go on to define three properties of secure information: *confidentiality*, *integrity* and *availability*. Confidentiality refers to “the protection of sensitive information from authorised disclosure” [3]. Integrity is defined as “the accuracy, completeness, and validity of information in accordance with business values” [3]. Availability relates to “information being available when required by the business process now and in the future” [3]. Collectively, confidentiality, integrity and availability are referred to as the CIA Triad [49]. In summary, an organisation’s information security objectives should be to protect the confidentiality, integrity and availability of business information and information systems through the implementation of countermeasures, while minimising damages from threats that exploit vulnerabilities [3, 50]. Workman, et al. [51] define a lapse within information security as a failure of one of the three properties of secure information in the CIA Triad.

An information security lapse can result from a potential *vulnerability* being exploited by a *threat*. A *vulnerability* is defined as “a weakness in the security system that might be exploited to cause loss or harm” [52]. For example, a vulnerability can be viewed in terms of a mistake or error in software code that can allow an attacker to gain access to an affected information system or computer network [53]. Other examples of vulnerabilities include unpatched applications or operating systems, unrestricted wireless access points, open ports on a firewall and unenforced password policies [54]. Industry vendors such as Cisco [55], Microsoft [56] and Symantec [57] regularly publish notices and advisories highlighting potential security vulnerabilities for a variety of information systems, applications and network devices.

A security *threat* refers to “any potential danger that is associated with the exploitation of a vulnerability” [54]. Alternatively, Pfleeger and Pfleeger define a threat as “a set of circumstances that has the potential to cause loss or harm” [52]. Security threats can be unexpected and have the potential to cause undesired effects that can negatively impact information or an information system. There can be many threats to information or an information system, including human-initiated threats and computer-initiated threats [52].

An organisation can minimise the impact of potential security threats exploiting vulnerabilities by implementing a suitable *security countermeasure* or *control* [54]. A security countermeasure is defined as “a control, measure, technique or procedure that is put in place to prevent a threat agent from exploiting a vulnerability” [54]. Examples of security countermeasures include strong password management, firewalls, access control mechanisms, encryption, and security-awareness education [54]. The relationship between threats, controls and vulnerabilities can be summarised as “a *threat* is blocked by *control of* a *vulnerability*” [52].

While organisations implement security countermeasures with the aim of minimising the impact of security threats exploiting vulnerabilities, there is still the possibility that a threat could damage, destroy, or disclose information or information systems. This possibility is known as *information security risk* [58]. Information security risk is defined as “the potential that a given threat will exploit vulnerabilities of an asset or group of assets and thereby cause harm to the organisation” [59]. Managing risk is considered an important element to sustaining a secure environment [54]. Organisations can assess and measure their security risk through a risk assessment [54]. A risk assessment enables an organisation to identify threats and vulnerabilities that have the potential to negatively impact their business operations [54]. There are several tools and methods available for conducting a risk assessment. These include the National Institute of Standards and Technology’s Guide for Conducting Risk Assessments [60], the Facilitated Risk Analysis Process (FRAP) [61], and the Operationally Critical Threat, Asset, and Vulnerability Evaluation (OCTAVE) approach [62]. An organisation can use the outcomes from a risk assessment to decide how it will allocate resources to implement security controls that will reduce the likelihood and/or the potential impact of the threats being realised [54]. A risk assessment is normally integrated into a wider information security management program.

3.2 Information Security Management

In an effort to help organisations fulfil information security requirements, the process of implementing security has been formalised through *information security management* [63]. The purpose of information security management is to implement appropriate countermeasures in order to minimise the impact that security-related threats and vulnerabilities might have on an organisation. Several researchers have examined the changing role of information security management within organisations in the past few decades [64–67]. Von Solms’ analysis separated the evolution of information security management within organisations into three ‘waves’ [66]. Von Solms argues that first generation information security management practices existed up until the early 1980’s and can be characterised as the ‘Technical Wave’, which focused on technical information security approaches [66]. Information security management during this wave was concerned with mainframes and organisations attempted to solve security problems through the built-in functions of the mainframe operating system. This was usually accomplished using access control lists, user-ids and through the use of passwords on the mainframes themselves [66].

Von Solms called the second generation of information security management, which lasted from the early 1980's to the mid 1990's, the 'Management Wave'. During this period, senior management within organisations realised that information security was no longer just a technical issue, but also required the development of security policies and procedures, as well as involving managers and executives in the security-decision making process [66]. The third generation, referred to as the 'Institutional Wave', began in the mid 1990's and continued into the 2000's [66]. This wave is characterised by the demand for information security standardisation and certification within organisations [66]. Therefore, it is during this period that many organisations looked to implement information security best practices and standards such as the International Organization for Standardization and the International Electrotechnical Commission (ISO/IEC) 17799/27002 standards [66].

Von Solms expands and updates his earlier work by arguing that a fourth wave emerged at the turn of the millennium called the 'Information Security Governance Wave' [67]. Von Solms explains that the forces behind this wave include the emergence of information security corporate governance, as well as growing legal and regulatory security requirements dictating the security of information and information systems [67]. Corporate governance is the "set of policies and internal controls by which organisations, irrespective of size or form, are directed and managed" [68]. As a result, information security governance can be viewed as a subsection of an organisation's general corporate governance program.

3.2.1 Information Security Management Systems (ISMS)

Humphreys [69] argues that within many organisations, it is ultimately the Board of Directors who are responsible for the protection of an organisation's information assets. It is this group of individuals who need to ensure that an appropriate risk assessment process is in place and that an effective system of security controls are implemented to mitigate identified threats. A growing list of legal and regulatory requirements means that a failure to comply can result in large financial penalties and possibly regulatory repercussions for an organisation [70]. Therefore, to help the Board of Directors fulfil these legal and regulatory commitments in a manner where due diligence can also be performed, many organisations are turning to Information Security Management Systems (ISMSs) [70].

An Information Security Management System (ISMS) is defined as "the policies, procedures, guidelines, and associated resources and activities, collectively managed by an organisation, in the pursuit of protecting its information assets" [48]. The primary objective of an ISMS is to implement, review and improve an organisation's information security so that it can achieve its business goals. [48, 71]. Similarly, Eloff and Eloff define an ISMS as "used for establishing and maintaining a secure information environment" [72].

Two main approaches have been identified in the literature which define how an organisation can structure an effective and efficient ISMS, the ISO/IEC 27001 standard [5] and the European Union Agency for Network and Information Security (ENISA) ISMS Framework [73]. The predecessor to the ISO/IEC 27001 standard, the British Standard (BS) 7799-2:2002 [74] introduced the four-step management Plan-Do-Check-Act (PDCA) cycle as a method for designing, implementing and reviewing an ISMS. The use of the PDCA cycle has been subsequently described in ISO/IEC security standards, which have since superseded 7799-2:2002. Within the PDCA cycle, the *Plan* phase is concerned with designing the ISMS, assessing information security risks and then selecting the appropriate security controls. Security controls identified in the plan phase are then implemented during the *Do* phase. The objective of the *Check* phase is to review and evaluate the efficiency and effectiveness of the ISMS. In the *Act* phase, any changes identified during the check phase are implemented to improve the performance of the ISMS. It is however, worth noting that the latest version of ISO/IEC 27001 [5], which was published in 2013, no longer emphasises the PDCA cycle. Instead, organisations are free to use other improvement processes such as Six Sigma's Define, Measure, Analyse, Improve and Control (DMAIC) as a method to continuously evaluate and improve their ISMS [5, 75].

Alternatively, ENISA describes an approach for implementing an ISMS using the ISMS Framework [73]. Similarly to the ISO/IEC standard, the ENISA approach consists of six steps [73]: (i) definition of security policy, (ii) definition of ISMS scope, (iii) risk assessment, (iv) risk management, (v) selection of appropriate controls, and (vi) statement of applicability. At the heart of the ENISA ISMS Framework are steps three and four, the risk assessment and management actions [73]. It is the activities, which take place within these steps, where the objectives of the ISMS are transformed into the implementation of controls that aim to prevent threats and vulnerabilities [73]. Regardless of the approach taken to implement and maintain an ISMS, organisations look to implement appropriate information security controls to mitigate identified risks. These security controls are selected and derived from existing information security standards or guidelines with the aim of meeting industry-specific requirements.

3.2.2 Information Security Standards and Guidelines

Often referred to as "voluntary standards, frameworks or sets of best practices" [76], several information security standards and guidelines have been developed and published by internationally recognised organisations. The target of a standard may be an organisation, a product or a service, while the criteria within the standard are likely to be primarily objective as opposed to subjective [77]. A standard will typically state what an organisation ‘must’ do, rather than ‘may’ or ‘might’ do with regards to implementing security controls [77]. An information security standard is further defined as “guidance to organisations on how to design, implement, and maintain policies, processes, and technologies to manage risks to its sensitive information assets” [54]. Security standards can be readily audited against, and security controls that do not conform to a standard should be easily identifiable. Several information security standards exist including the ISO/IEC 27002 standard [6], the National Institute of Standards and Technology (NIST) Federal Information Processing Standards Publications [78] and ‘Special Publications’ [79], and the Payment Card Industry Data Security Standard (PCI-DSS) [80]. The following subsections review these information security standards.

The ISO/IEC 27002 Standard

ISO/IEC 27002 is an international information security standard which was last updated in 2013. The standard provides recommendations with regards to initiating, implementing and maintaining secure systems. Unlike ISO/IEC 27001, which is a formal specification, ISO/IEC 27002 is an advisory document, which recommends information security controls that can be implemented within the process of establishing and maintaining an ISMS [6]. ISO/IEC 27002 consists of fourteen security domains that cover security control information on areas such as security policies, asset management, human resource security, business continuity management and operations security [6]. The standard also addresses requirements concerned with developing procedures that define security incident response planning and preparation, the handling of forensic evidence and learning from security incidents. ISO/IEC 27002 is considered to be a standard that can be applied to all types and sizes of organisations, irrespective to the security threats and risks they face [6]. This supports Siponen and Willison’s [81] argument that such standards do not pay enough attention to different organisations’ security requirements. In practice, the broad range of security controls covered in the standard can also provide an organisation with some flexibility to adopt only the security controls that they require. However, this can complicate compliance testing for certification because an organisation may implement specific security controls and not the entire standard.

The Federal Information Processing Standards and Special Publications

The National Institute of Standards and Technology (NIST) have developed a range of information security publications in response to the United States (U.S.) Congress enacting the Federal Information Security Management Act of 2002 (FISMA) [82] as part of the E-Government Act of 2002 [83]. FISMA dictates that U.S. federal agencies identify, develop, document, and implement an agency-wide information security program [84]. To assist federal agencies implement information security programs and audit methods, NIST have created the Federal Information Processing Standards Publications (FIPS PUBS) [78] and 'Special Publications' [79].

FIPS PUBS are issued by NIST in accordance with FISMA and are mandatory security requirements for U.S. federal agencies [84]. FISMA requires that all federal agencies comply with the requirements within these standards [84]. These publications contain information on everything from baseline security requirements for information systems, to local area network security, cryptography and encryption requirements, as well as data transfer requirements [78]. In addition to FIPS PUBS, NIST have also created a range of information security publications, which are published as recommendations and guidance documents and are to referred to as 'Special Publications' [79]. Federal agencies are also required to follow any NIST Special Publication referred to within a FIPS PUB [85]. For example, to apply FIPS PUB 200 [86], agencies first need to classify and determine the security category of their information systems as required by FIPS PUB 199 [87]. Agencies can then select suitable security controls from NIST Special Publication 800-53 [85] to satisfy their minimum security requirements as required within FIPS PUBS 200 [86].

Although NIST has based these 'Special Publications' on principles for securing federal government information systems, other organisations can also use these publications to establish a minimum security-control baseline within their specific environments [88]. This is because these 'Special Publications' cover a wide range of information security topics including wireless security, access controls, media sanitisation, contingency planning, encryption and key management, electronic mail security, mobile device security and server security [86].

The Payment Card Industry Data Security Standard

The Payment Card Industry Data Security Standard (PCI-DSS) [80] was developed by a number of major credit card companies. PCI-DSS was created to assist the adoption of consistent data security controls for organisations around the world that manage, handle and store payment card processing information [80]. The current version of the PCI-DSS (Version 3.1) was released in April 2015 and changes are made to the standard every three years [89]. The PCI-DSS identifies twelve security requirements, which are organised into six groups called 'control objectives'. These six control objective groups include building and maintaining a secure network; protecting card-holder data; maintaining a vulnerability management program; implementing strong access control measures; regularly monitoring and testing networks; and maintaining an information security policy [90]. While PCI-DSS is not legally binding within the European Union, the laws of some States in the U.S. either directly refer to PCI-DSS or make equivalent provisions. For example, the State of Washington has incorporated the standard into state law, which specifically states that compliant organisations are shielded from liability in the event of a data breach or a security incident [91].

3.2.3 Information Security Standards in Practice

Information security standards and guidelines can play an important part in managing and certifying organisational information security [92]. Siponen and Willison [81] state that organisations which adopt security standards can "demonstrate their commitment to secure business practices; apply for security certification, accreditation, or a security-maturity classification attesting to their compliance to a set of rules and practices" [81]. Effectively, information security standards and guidelines provide a baseline for an organisation's overall information security management strategy.

Several case studies have been used to examine how organisations implement information security standards and guidelines, as well as evaluating the impact of these publications on an organisation's wider security posture [81,93,94]. Wiander [93] analysed the implementation of the ISO/IEC 17799:2005 standard in four separate organisations using semi-structured interviews. The results from the study showed that individuals perceived that the implementation of the standard increased the overall understanding of information security within the studied organisations [93]. However, interviewees also reported that they had experienced difficulties in implementing the standard, with the readability of the standard being the main problem [93]. While Wiander [93] has reported that the implementation of information security standards has increased the overall understanding of information security within an organisation, Siponen has been more critical of their use. Siponen [94] argues that the foundations of information security standards and guidelines are not universally validated and are simply based on personal experience. Therefore, these information security management standards and guidelines should not be treated as 'gold standards', but rather a library of material on information security management for organisations [94].

Separately, Siponen and Willison [81] evaluated four information security standards (British Standard 7799, ISO/IEC 17799, the Generally Accepted Information Security Principles, and the Systems Security Engineering Capability Maturity Model) and argued that these standards did not pay enough attention to the differences between organisations because their information security requirements could be very different. For example, a smaller organisation could lack the demand for a dedicated security incident management team and place more emphasis on anti-virus solutions and firewalls, while a larger organisation could place equal emphasis on all aspects of information security.

3.3 Information Security Incident Management

Security policies and controls alone will not guarantee total protection of organisation data or information systems. After policies or controls from published security standards have been implemented, vulnerabilities are likely to remain within an organisation. Vulnerabilities can also emerge from previously unknown threats and make any implemented security controls ineffective. As a result, successful security attacks can still affect an organisation. In some cases, these attacks can go on to become information security ‘incidents’, which can have adverse effects on an organisation’s business operations. It is therefore useful for an organisation to develop and implement plans and procedures for managing security incidents when they do occur within their organisational landscape. While some of the above security standards include procedures for handling security incidents, there are also numerous other specific security incident response approaches, which can be used to manage and handle security incidents within organisations.

3.3.1 Definition of Security Incident in the Literature

An analysis of the literature has highlighted numerous definitions for the term ‘information security incident’. The International Organization for Standardization (ISO) define an information security incident as a “single or a series of unwanted or unexpected information security events that have a significant probability of compromising business operations and threatening information security” [48]. An information security event within this context is “an identified occurrence of a system, service or network state indicating a possible breach of information security policy or failure of controls, or a previously unknown situation that may be security relevant” [48].

Alternatively, Cichonski, et al. [8] describe a security incident from the perspective of violating security policies. They define a security incident as “a violation or imminent threat of violation of computer security policies, acceptable use policies, or standard security practices” [8]. Ahmad, et al. [18] add that “an information security incident occurs when there is a direct or indirect attack on the confidentiality, integrity and availability of an information system”. Howard and Longstaff [95] take a different approach to defining a security incident as part of their computer security taxonomy. They describe a security incident as “a group of attacks that can be distinguished from other attacks because of the distinctiveness of the attackers, attacks, objectives, sites, and timing” [95].

Information security incidents can also be defined in terms of regulatory impact. For example, within the Health Insurance Portability and Accountability Act of 1996 [96], a security incident is defined as “the attempted or successful unauthorised access, use, disclosure, modification, or destruction of information or interference with system operations in an information system" [97]. Article 13a of the European Union's Telecommunications Directive (2009/140/EC) take an alternative view on defining a security incident from a telecommunication's perspective. For the purpose of Article 13a, The European Union Agency for Network and Information Security (ENISA) have defined a security incident as "a breach of security or a loss of integrity that could have an impact on the operation of electronic telecommunications networks and services" [98]. Within the United Kingdom (UK), the Information Commissioner's Office (ICO) defines a security incident in terms of its impact on data loss. The ICO define a security incident as "a breach of security leading to the accidental or unlawful destruction, loss, alteration, unauthorised disclosure of, or access to, personal data transmitted, stored or otherwise protected in connection with the provision of a public electronic communications service" [99].

With various definitions available in the literature, there is a risk that an organisation could consume valuable resources on security incidents that may not be considered incidents within their distinct landscape. For example, a 'security incident' in a telecommunications company may not be an 'incident' in a financial services organisation and vice versa. This situation is further complicated by the applicability of multiple definitions within specific domains. An organisation in the UK which implements the ISO/IEC 27001/27002 security standard, would be inclined to follow the ISO/IEC definition of the term 'security incident'. However, the same organisation is also likely to process and store personal customer data and is likely to also need to take into consideration the ICO's definition of the same term. In this scenario, an organisation's security incident response team could be faced with a situation where a security problem is *not* an incident according to the ISO/IEC definition, but *is* an incident when examined using the ICO's definition. The wide range of definitions that describe a security incident means that organisations need to define the term within the context of their business operations. Regardless of the definition(s) used in a specific organisation, security incident response capabilities are increasingly becoming a regulatory requirement in a variety of industries [82, 96, 100–102].

3.3.2 Motivation for Security Incident Response Capabilities

Increased regulatory pressure has been applied to organisations in a variety of industries, mandating security incident handling requirements, including post-incident learning [82, 96, 100–102]. For example, within the healthcare industry, the introduction of the 'Security Rule' [100] to the Health Insurance Portability and Accountability Act (HIPAA) [96] dictates organisations implement administrative, physical, and technical safeguards to protect patient information. Section 164.308 of the Security Rule specifically requires organisations to

> "identify and respond to suspected or known security incidents; mitigate, to the extent practicable, harmful effects of security incidents that are known to the covered entity or business associate; and document security incidents and their outcomes" [100].

Further regulations from the Department of Health and Human Services mandates that organisations "implement policies and procedures to prevent, detect, contain, and correct security violations" [97]. As noted above, HIPAA security incident response requirements place a strong emphasis on documenting security incidents and the outcomes from any investigation [97]. It is anticipated that this documented information can then be used to help produce lessons learned at a later stage. This documented information can come from information within an investigation record, as well as digital evidence collected throughout the investigation itself. Digital evidence is defined as "any digital data that contain reliable information that supports or refutes a hypothesis about the incident" [103]. The collection and documentation of digital evidence and other investigation information such as conversations between security incident handlers can help create a more efficient and less error-prone approach to handling a security incident. However, the collection of digital evidence and investigation information can be affected by an organisation's inclination to focus on the eradication and recovery from an incident instead of collecting digital evidence [18, 20]. Similar to HIPAA, the Federal Information Security Management Act of 2002 [82], the Gramm-Leach-Bliley Act of 1999 [101], the European Union's Article 13a (as part of a regulatory framework for electronic communications) in the European Union [98], and Sarbanes-Oxley Act of 2002 [102] all require organisations to have policies and procedures to detect, report and respond to security incidents.

These laws not only mandate that organisations have procedures to manage and handle security incidents but also include requirements with regards to reporting specific security incidents to relevant authorities. In the United States (U.S.), federal law requires governmental agencies to report the outcomes of security incidents to the U.S. Computer Emergency Readiness Team [8]. Meanwhile in the European Union (E.U.), Article 13a of the European Commission's Telecommunications Directive (2009/140/EC) requires network service providers to report 'significant' security breaches and losses of integrity to national authorities [104]. The proposed E.U. Network and Information Security Directive [105], extends these reporting obligations to 'market operators' who are responsible for critical national infrastructures, across the energy, banking, health, transport, financial services and food sectors [106].

The British Government has extended these security incident reporting requirements to include the sharing of lessons learned [107]. Through the Cyber-security Information Sharing Partnership, organisations are encouraged to disseminate and exchange information about lessons learned from security incidents with other organisations in the same industry [107].

Organisations need to define how security incident response teams will respond and manage information security incidents in order to meet these legal and regulatory requirements. To help organisations define these requirements, several security incident response approaches have been proposed in the literature.

3.3.3 Security Incident Response Processes

Numerous security incident response processes and best practice guidelines have been published in industry [8–10, 13, 108] and academia [7, 11, 12, 109, 110] describing how organisations can investigate and manage security incidents. These processes are typically described in terms of a number of successive phases. The structure and phases contained in various security incident response approaches are summarised in Table 3.1.

The table shows that there are a number of differences in the phases presented in these approaches. These differences suggest that there is currently no single de-facto approach that can truly be classified as an industry-standard for handling security incidents [7]. For example, three of the approaches (CERT/CC Incident Response, Good Practice Guide for Incident Management and the Security Incident Tasks), do not include a 'follow-up' phase as part of their approach. In contrast, some approaches contain phases that are not present in others. For example, the Incident Response Process [11] is the only approach with an 'data collection' phase, while the CERT/CC Incident Response [108] process is the only approach with a 'protection' phase. However, the table also shows that there are some consistencies with the approaches as well. For example, five out of the ten approaches contain similar phases: preparation, detection, containment, eradication, recovery and follow-up. For the purpose of this research, these five phases are considered to be part of a typical security incident response approach, as shown in Figure 3.1. Numerous tasks and activities are typically undertaken in each of these five phases.

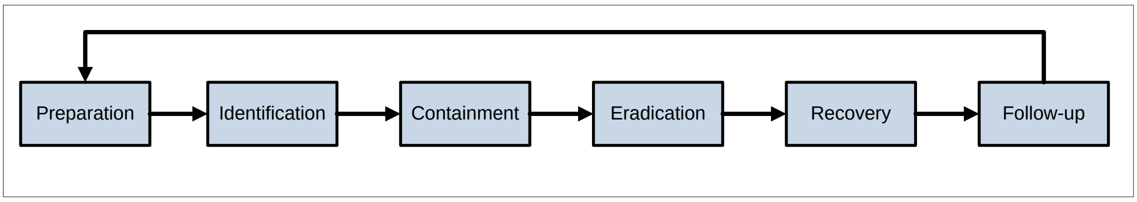

Figure 3.1: Typical Security Incident Response Cycle

Preparation Phase

The preparation phase is concerned with the creation of a security incident team and providing this team with the necessary tools and resources to manage and handle security incidents. Management within an organisation can play an important role in the creation of a

Model	**Reference**	**Preparation**	**Protection**	**Detection**	**Initial Response**	**Triage**	**Analysis**	**Data Collection**	**Response**	**Containment**	**Eradication**	**Recovery**	**Follow–Up**
CERT/CC Incident Response	[108]	✓	✓	✓		✓			✓				
ISO/IEC 27035 Standard	[10]	✓		✓		✓			✓	✓	✓	✓	✓
The NIST 800-61 Model	[8]	✓		✓			✓			✓	✓	✓	✓
Good Practice Guide for Incident Management	[9]			✓		✓	✓		✓				
SANS Incident Response	[13]	✓		✓						✓	✓	✓	✓
The Incident Response Process	[11]	✓		✓	✓	✓	✓	✓				✓	✓
The Incident Response Process	[7]	✓		✓						✓	✓	✓	✓
Approach for Detecting and Reacting to IT Systems Abuse	[109]	✓											✓
Incident Response Model	[12]	✓		✓						✓	✓	✓	✓
Security Incident Tasks	[110]			✓			✓		✓				

Table 3.1: Security Incident Response Processes

security incident response team through providing support to the team, communicating team objectives and purpose to the wider organisation, as well as assisting in the recruitment of staff [111]. Management can also help ensure that a team has access to any tools and resources that they may require including an incident record tracking system, any tools for use during the investigation or analysis, as well as copies of operating systems and applications to help with restoration.

Detection Phase

This phase is also referred to as the ‘identification’ phase and is associated with the identification, detection or reporting of security incidents within an organisation. There are several approaches an organisation can take to identify and detect a security incident. These can include monitoring intrusion detection systems and other network monitoring appliances for anomalies, as well as monitoring third-party news websites for information about security incidents, which may have affected other organisations [9]. Security incidents can also become known to a security incident response team through a security incident reporting system [15]. A security incident reporting system provides a secure mechanism for employees within an organisation to report security incidents directly to the security incident response team [15]. A security incident response team also needs to determine if the reported event constitutes a security incident.

Containment Phase

The primary objective of this phase is that actions are taken by a security incident response team before a security incident overwhelms resources, increases damage, or spreads to other networks or systems within an organisation [7]. Examples of containment strategies can include shutting down a system, disconnecting a computer host from a network and disabling a user account [8]. While these strategies can help to contain the spread of a security incident, they can also lead to the destruction of digital evidence. Therefore, an important task within this phase is to also collect any digital evidence with regards to an incident prior to the commencement of these containment strategies [8].

Eradication Phase

After a security incident has been contained, the next phase involves a security incident response team implementing a solution to prevent the incident from escalating any further. Examples of eradication solutions can include disabling services or firewall ports that may have been exploited and deleting malware that has been installed on compromised hosts [8].

Recovery Phase

Within the recovery phase, a security incident response team, together with other technical personnel, will look to restore affected systems and users to normal business operation [7]. This can be achieved in numerous ways including rebuilding systems, replacing compromised data, changing passwords and installing patches to affected hardware and software [7].

Follow-up Phase

The primary aim of the follow-up/post-incident phase is to establish lessons learned and disseminate these, as appropriate, to relevant stakeholders within an organisation and possibly to external regulators. Lessons learned are defined "as knowledge or understanding gained by experience" [112]. In general, these lessons should be used to improve an organisation's information security management and security incident response processes in the longer term. The experiences during an investigation may be positive, such as an activity or task within the security investigation lifecycle that worked well, or negative, such as a failed security control [112]. The main activities that can evolve from this phase include the completion of investigation reports, dissemination of lessons learned as well as implementing improvements to information security management and security incident response processes.

While seven out of the ten studied security incident response approaches contain a follow-up phase, researchers have argued that many organisations do not pay enough attention to this phase and instead, tend to focus on eradication and recovery [18, 19, 44]. One explanation of this problem is that security incident response processes, which are linear in nature, could experience a similar problem as another linear approach, the Waterfall model in software engineering. In the Waterfall model, when a project begins to run out of time and money, testing usually the last phase in the approach, is either cut short or not performed [113]. As a result, the quality of software is reduced and risks can be introduced into a project. The same problem appears to be prevalent within security incident response, where resources appear to be exhausted during the eradication and recovery phases. Hence, very little resources may be left to execute the follow-up phase. Reduced activity within this phase can translate into less incident learning taking place for a particular security incident.

3.3.4 Security Incident Response Practices in Organisations

While some parts of the literature have focused on developing various processes for implementing security incident response capabilities within organisations, other researchers have identified and discussed several problems with these processes through case studies within organisations. Hove, et al [114] argue that many organisations find it difficult to implement established security incident response processes. This is evident in an analysis of empirical case studies conducted in various organisations.

Tan, et al. [20] explored the factors that have influenced information security managers to not conduct investigations following a security incident. Tan, et al. [20] reported that their studied organisation had no clear definition for the term ‘security incident’. As a result, incident handlers did not realise what security problems were actually ‘incidents’ and were slow to react to real security incidents. Tan, et al. [20] also observed that incident handlers were not encouraged to view security investigations as means of improving the organisation’s overall security. Instead, the main focus of investigations was to ‘restore and recover’ normal business operations. Furthermore, Tan, et al. [20] found that participants were unaware of the benefits of collecting digital evidence. Management within the organisation noted that they would only decide to prosecute after the problem was fixed. As a result, evidence that would assist in the prosecution of any offenders, may either be contaminated or destroyed [20].

Werlinger, et al. [27] conducted an exploratory study to investigate the security incident activities of practitioners in various organisations. The purpose of the study was to examine and analyse the tasks, skills, strategies and tools that practitioners use to diagnose security incidents [27]. Practitioners argued the importance of integrating the input from several monitoring tools in order to gain a proper overview of the security incidents occurring within an organisation. However, Werlinger, et al. [27] concluded that current security incident response practices and tools do not appropriately support the highly collaborative nature of investigations and that practitioners were often required to write custom tools in order to perform specific investigative tasks.

Casey [28] outlines a case study of a network attack, which required the victim organisation to cooperate with law enforcement agencies to apprehend the attacker. Casey argues that many organisations may not be fully exploiting their digital forensic capabilities and are likely to be undermining the value of forensic evidence [28]. Casey adds that organisations should consider integrating evidence handling into their incident response capabilities and educating system administrators about the need to report and investigate even the most trivial of incidents. This suggests that even with a moderate amount of preparation, the victim organisation could have been better prepared to prosecute the identified attacker. This is evident in the methods used in the preservation and collection of evidence after the attack. Although the studied organisation developed guidelines defining digital evidence collection, one individual collected file listings instead of files themselves. Furthermore, the individual did not keep a log of his actions and could not determine which files were obtained from which system. As a result, the evidence gathered from seven systems involved in the attack could not be included in the criminal case that followed [28].

Hove, et al. [114] studied three large organisations with the purpose of investigating the plans and procedures for handling security incidents within the studied organisations. Two of the studied organisations had a dedicated security incident response team, while the third organisation did not, but used the services of an external security incident response team when the need arose. Hove, et al. [114] identified differing purposes for responding to security incidents within the organisations. One of the studied organisations argued that although it is important to restore affected systems, it was equally important to ensure the incident was properly investigated to avoid a similar attack occurring in the future. One the other hand, one of the other studied organisations preferred to apply a temporary solution to the incident so that they can resume normal business operation and minimise the impact of the incident. Hove, et al. [114] concluded that although the organisations have plans and procedures in place, based on industry best practices, many other procedures were missing from the studied organisations. For example, in two of the organisations security incident reporting procedures were not established, while the respondents in another organisation indicated that they did not have enough staff to respond to incidents efficiently [114]. Other issues identified by the participants in the case study included difficulties in collecting incident information from multiple sources, as well as deciding how much information to distribute about a particular incident [114].

Line, et al. [115] examined how six electricity distribution service operators within the power industry planned and prepared for information security incidents. The case study revealed that many of the surveyed organisations had little or no documentation regarding the investigation of security incidents [115]. In addition, personnel within the studied organisations confirmed a lack of incident response training and that responsibilities for security incidents were not adequately established within their specific organisations. Furthermore, Line, et al. [115] reported that the majority of their studied organisations did not have a clear definition for a security incident. In a separate case study, Line [116] added that none of the surveyed distribution service operators had established a security incident response team. Line also noted there appeared to be lack of cooperation between information technology and power automation staff with regards to security incidents within their organisations [116].

Metzger et al. [117] presented the experiences of the security incident response team at the Leibniz Supercomputing Centre (LRZ-CSIRT) in Germany, where a holistic approach based on ISO/IEC 27001 had been developed for security incident response. Metzger et al. [117] explained that LRZ-CERT's approach to security incident response included a combination of strong incident reporting capabilities, automatic response and analysis, as well as a process-oriented approach for intervention. However, Metzger et al. also noted several challenges within LRZ-CSIRT including a lack of sufficient personnel to operate forensic tools and evidence collection. Metzger et al. also noted that some individuals were not reporting security incidents because they were not sure what to report as 'incidents' [117].

Orderlokken [118] studied Norwegian public and private organisations to examine how these organisations performed security incident response. Orderlokken reported that public organisations were found to have inferior incident reporting and training compared to their private counterparts [118]. Orderlokken added that only half of the studied organisations were following standards for information security. The case study also revealed that less than half of the studied public organisations recorded statistics of the number of incidents, which impacted their organisation. When statistical information was recorded, Orderlokken reported that this was often inaccurate due to a lack of implemented process, a lack of training and weak definitions of security incidents [118].

Moller [119] examined the organisational and technical challenges of the German Computer Security Incident Response Team (CSIRT) while extending their incident response services to include grid environments. Moller argues that it is important that a CSIRT ensures that its services are known to its constituency and that system administrators within the constituency trust a CSIRT with their confidential data [119].

Several themes emerge from the case studies concerning post-incident learning which include current incident response processes not reflecting the concurrent lifecycle of real-world security incident handling, current processes not providing enough insight into the causes of an incident and current processes not maximising the benefits of digital forensic capabilities. As a result, post-incident learning could be a challenge in these studied organisations. For example, Casey [28] and Tan, et al. [20] have highlighted that organisations may not be maximising corporate forensic capabilities. A lack of forensic readiness could affect both an organisation's ability to take subsequent legal action and also limit the amount of post-incident learning. Without access to detailed information including forensic data, a security incident response team could find it difficult to learn from a security incident. The next section will be used to explore the follow-up phase, including the activities involved in this phase, as well discussing post-incident learning challenges within organisations.

3.4 Post-Incident Learning

Incident learning can be defined as "the collection of organisational capabilities that enable an organisation to extract useful information from incidents of all kinds and to use this information to improve organisational performance over time" [120]. While information security incidents are unwanted events, they do at the same time present an opportunity for an organisation to learn more about the risks and vulnerabilities which can exist in an organisation's systems and processes [121]. Researchers have argued that organisations do not pay enough attention to incident learning [18, 19, 44]. These researchers go on to claim that organisations are more concerned with eradication and recovery and in some cases, have failed to learn from security incidents [18, 19, 44].

3.4.1 Post-Incident Learning in Organisations

Several researchers have examined security post-incident learning within organisations [18–20, 122, 123]. In a study involving the petroleum industry, Jaatun, et al. [122] explained that while learning from security incidents was considered important, organisations found it difficult to implement in practice. Jaatun, et al. [122] go on to argue that organisations must be prepared for incident learning and this includes obtaining managerial commitment and the willingness to commit resources to facilitate learning from security incidents. This is a view that is shared by Tan, et al. [20], who also noted that their studied organisation were not prepared to gather data or learn from security incidents.

Ahmad, et al. [18] reported that within their studied organisation, 'high-impact' security incidents are reviewed within 24 hours of the system services being restored, with a post-incident report being produced at the end of the review. However, the studied organisation does not have a structured process for reviewing 'low-impact' security incidents [18]. Nonetheless, Ahmad, et al. [18] noted that, although the organisation in their study closely followed industry best practices, the organisation's inclination was to focus on containment, eradication, and recovery. As a result, the security incident response team and organisation as a whole did not fully exploit their ability to learn from security incidents [18].

Rollason-Reese [123] presented a case study that analysed the lessons learned by a security incident response team within a public university in the United States. The team responds to a diverse range of incidents and uses a five-step response process that comprises of: alert, analysis, response, recovery and maintenance. Rollason-Reese highlights several important lessons within his study. These include educating university personnel on the importance of reporting potential security threats; organisations developing their own definition of 'security incident'; and the importance of the incident record as a means of documenting activities, decisions and evidence gathered during an investigation [123].

While the above researchers have argued that organisations need to do more to effectively learn from a security incident, Shedden, et al. [44] argue that current best practices and approaches do not provide enough guidance and support as to how this can be achieved. Shedden, et al. [19] state that "researchers and practitioners must accept that informal activities will occur below the surface in security incident response" and hence, security incident response approaches should be less formal and cater for informal learning approaches. Shedden, et al. [19] reported that within their studied organisation, incident handlers were undertaking informal learning through conversation and observation. While Shedden, et al. [19] propose that security incident learning should be informal within organisations, very few informal tools have been proposed in the literature to help organisations undertake security incident learning.

3.4.2 Methods for Conducting Post-Incident Learning

Best practice approaches such as NIST 800-61 [8] and SANS [13] suggest that security incident response teams hold lessons learned meetings at the conclusion of investigations. The purpose of these meetings is that a security incident response team explores various ways of improving both their own security incident response process, as well as the wider organisational security posture [8, 13]. This can involve establishing what has transpired in a recently concluded investigation and what can be done better in future investigations [8, 13]. The output of these meetings typically includes an executive summary that includes the cost of the investigation, its impact and where possible, the investigation results [8, 13]. The ISO/IEC 27035 standard [10] adds that post-incident meetings are also used to identify new and review existing security controls, discuss whether further forensic analysis is required, as well as communicating investigation results from the meeting to a trusted community [10]. Ideally, the objective of any lessons learned is that they can be used to improve the preparation for future incidents, be used for leverage for increased security budgets, improve education awareness programs, as well as strengthening security within the wider organisation [8, 10, 13, 124].

The majority of the security incident response approaches identified in the literature place a strong emphasis on identifying the technical causes of a security incident during the development of any lessons learned [7, 8, 10, 12, 13, 109]. However, recently researchers [18–20] have argued that organisations need to look beyond the immediate technical causes of an incident and start to examine the underlying root causes. Ahmad, et al. [18] add that often, the root cause of a security problem may not necessarily be a technical problem. A root cause is defined as “the most basic cause that can be reasonably identified and that management has control to fix” [125]. Root Cause Analysis (RCA) is a method of problem solving used for identifying the root causes of faults or problems [15]. RCA has been used in several industries such as health and safety, aviation and medical care in assisting in the prevention of future incidents [15].

Although there have been increased calls for the use of RCA within security incident response, relatively little work has examined the effectiveness of these tools and methods for conducting a RCA within this domain. However, Johnson [126] has discussed and demonstrated how Violation and Vulnerability (V2) diagrams can be used to assist RCA using an incident at Allfirst Bank as an example case study. Similarly, Stephenson [127] has also proposed an RCA approach which uses coloured Petri nets to model attacks of security processes within an organisation, and presents an example of how a worm or virus infection has spread through an organisation's network structure. Although these two methods have been identified in the literature, their scope and applicability to security incident response investigations within organisations has not been established. When lessons learned are created at the conclusion of an investigation, irrespective of them containing a root cause, this information typically needs to be disseminated back into the wider organisation [18].

3.4.3 Dissemination of Post-Incident Learning

The work of a security incident response team is usually completed by the issuance of a report, detailing the investigation findings and any lessons learned [18]. This information can then be distributed within an organisation or reported to external regulatory bodies. Several security incident response approaches [8, 10, 13] also propose that any lessons learned are disseminated to the wider organisation. However, very little research has focused on effective methods for sharing and exchanging lessons learned in the aftermath of security investigations. Traditionally, lessons learned can be circulated and exchanged through a series of formal reports, executive summaries and informal meetings [7, 13, 18]. Communicative notes from post-incident meetings can then be disseminated to relevant individuals within an organisation [13]. These notes can contain responses to incidents, disagreements over incident handling procedures, suggestions and enhancements to security controls, policies and incident response procedures [13].

There are limitations to the use of these approaches for the dissemination of lessons learned within security incident response [128, 129]. As many of the approaches consist of text reports, He [14] argues that a long incident report can discourage readers from determining the key insights into a particular incident. He [14] adds that any causal factors identified during an investigation may not be clear in the linear format of a text report. In an effort to address these concerns, He, et al. [129] developed a Generic Security Template (GST) using graphical Goal Structuring Notation to present and disseminate lessons learned in a graphical structured manner. Through a series of experiments, He [14] has demonstrated how the GST can be used to present lessons learned from security investigations within the health-care industry. The intention of the experiments were to present security objectives, issues and recommendations that are embedded within the pages of text report in a graphical manner.

3.5 Data, Information, Knowledge, Wisdom

The terms 'data', 'information', 'knowledge' and 'wisdom' are often grouped together in a popular model called the Data-Information-Knowledge-Wisdom (DIKW) hierarchy [130].

The DIKW hierarchy is often discussed or used to define the terms data, information and knowledge within the domains of information management and information systems [130]. Ackoff offers the following definitions of the terms as used in the DIKW hierarchy [131]:

- *Data* are the "symbols that represent properties of objects, events and their environment."
- *Information* "consists of processed data, the processing directed at increasing its usefulness. Information systems generate, store, retrieve, and process data."
- *Knowledge* is "know-how, and is what makes possible the transformation of information into instructions."
- *Wisdom* is the "ability to increase effectiveness."

Rowley [130] further explored the definitions of the above terms and identified a range of ambiguous and/or conflicting definitions. Rowley summarised from her analysis that the assumption from the DIKW hierarchy is that *data* can be used to create *information*, which in turn, can then be used to create *knowledge* and knowledge can be used to create *wisdom* [130].

While Ackoff and Rowley have provided generic definitions for the terms 'data', 'information', 'knowledge' and 'wisdom', their application to security incident response has not been well discussed in the literature. However, Knapp [132] has attempted to discussed the terms in the context of Service Desk incidents. Knapp states that *data* are the raw facts about an incident, which are not organised or structured in a meaningful way, for example which workstation is affected, its manufacturer and model number [132]. Knapp explains that data becomes *information* in an incident when it has been organised and is capable of answering questions such as "who?, what?, when? and where?". For example, in an incident information, is who owns the workstation, where it is located and when it was installed [132]. With regards to *knowledge*, Knapp explains that this is "the application of information along with people's experiences, ideas and judgements", for example knowledge in an incident would include how to resolve the problem affecting the incident [132]. Finally, Knapp states that *wisdom* is the "judicious application of knowledge" and provides an example of how incident handlers can use data, information and knowledge to make 'wise' decisions when they face a particular incident [132].

Baskarada and Koronios [133] state that in terms of the DIKW hierarchy, many researchers use the terms data quality and information quality interchangeably, when it comes to discussing the quality dimension of the two elements. As this research is concerned with the quality of data in security incident response, the next section will provide an overview of data quality along with a working definition of 'good quality' data within security incident response.

3.5.1 Data Quality

Klein and Rossin [134] argue that there is no single definition for the term 'data quality'. An analysis of the literature indicates that a consensus on the definition of 'data quality' has not yet emerged and this is perceptible in the following definitions:

- "the degree to which data meets the expectations of data consumers, based on their intended uses of the data" [135].
- "is the measure of the agreement between the data views presented by an information system and that same data in the real world" [136].
- "data that are fit for use by data consumers" [137].
- "fitness (to the purpose) of use" [138, 139]

The quality of data is therefore directly related to the perceived purpose of the data and that high-quality data meets the expectations of the intended users to a greater degree than that of low-quality data [135]. One factor in how well it meets the expectations of the data user is how those users perceive the data to represent what it purports to represent [135]. Within information security and particularly security incident response, there is very little reference to the term 'data quality'. Therefore, for the purpose of this thesis a definition of data quality within security incident response is derived from the concepts discussed above. The proposed definition of 'good' quality data within security incident response is *data that is derived from a security incident response investigation that is fit to facilitate security incident learning within an organisation.*

The concept of data quality can also be further defined in terms of a set of dimensions, which are considered to be quality properties or characteristics of data [140]. Wang et al. [141] define a data quality dimensions as "a set of data quality attributes that most data consumers react to in a fairly consistent way". In 1996, Wand and Wang [142] noted there was no general agreement on data quality dimensions. Wand and Wang [142] defined a data quality classification, of which they define four dimensions: completeness, unambiguousness, meaningfulness and correctness. Levitin and Redman [143] present and discuss a list of six 'quality dimensions' from the perspective of database data. These six 'quality dimensions' include content, scope, level of detail, composition, consistency, and reaction to change [143]. Alternatively, Ballou and Pazer [144] divide data quality into four dimensions:

accuracy, timeliness, completeness and consistency. The accuracy dimension is concerned with the difference between the correct data value and the value actually used. Timeliness focuses on data that is in error because it is out-dated and differs from the original value. Completeness is concerned with ensuring that no data is missing, while consistency implies that some form of data representation standard exists throughout the data values [144]. However, like the definition of data quality, researchers have yet to agree on a standard for data quality dimensions.

Several practices have been developed for improving the quality of 'poor data', one set of practices include a three-step process which involves data profiling, data cleansing and data defect prevention [21]. *Data profiling* involves analysing the data and looking for anomalies within records in a file or database [21]. After the extent of 'dirty data' is known, *data cleansing* is undertaken where incorrect, corrupt or inaccurate data is replaced, modified, or deleted. The final stage is to determine how to prevent 'dirty data' in the future and this done using a process called *data defect prevention* [21]. This process typically involves the owners of any systems/processes implementing procedures to prevent 'dirty data' from being produced when data is actually created [21]. Numerous tools [22–24] have been developed to help organisations implement these practices. However, very little research has examined if these practices and tools can be extended and used to monitor and improve the quality of data in security incident response investigations. A lack of evaluated practices and tools for improving the quality of data in security incident response means that it can be difficult to measure if the lessons learned derived from a security investigation are correct, if there are possible quality issues within the investigation data.

3.5.2 Data Quality within Security Incident Response

Incident learning and the dissemination of information has been used in a variety of industries in the prevention of future incidents [15]. Johnson [15] notes that investigators in these domains rely on necessary data being made available to help identify underlying root causes. Stephenson [16] argues this is also true for security incident response investigations, where detailed data is necessary to help a security incident response team to analyse root causes. However, Stephenson [16] states that obtaining this detailed security data can be difficult within many organisations. Regardless of the methods used for incident learning and the dissemination of lessons learned, if organisations are more focused on eradication and recovery, then there is the potential that poor quality data is emerging from security incident response processes. Without enriched quality data, a security incident response team could find it difficult to undertake a lessons learned meeting or perform a root cause analysis.

There have been documented concerns with regards to the quality of data used in investigations within the transportation safety domain [17]. A report [17] published by the United

States Bureau of Transportation highlights concerns with the reliability and accuracy of incident information in this domain. The data quality issues can be summarised in the following quotes [17]:

- “there needs to be better information and it needs to be of a higher quality”
- “there needs to be better data on results”
- “accuracy is a challenge because of budgetary problems and different interests”
- “it is difficult to get accurate, undiluted information on human error and performance”

While these concerns have been raised in the safety domain, the problem of poor quality data could provide one explanation as to the problems encountered by organisations when attempting to learn from information security incidents.

3.6 Organisational Culture and Learning

Organisational culture is a widely used term and several different people have attempted to explain the concept [145]. Watson [146] states that the concept of organisational culture was derived from a metaphor of the organisation as ‘something cultivated’. Researchers have suggested that organisational culture is a concept of basic assumptions, in an organisation, developed around their handling of employees, promoted values or statements of belief that have worked well in the past to be considered valid in the specific organisation [147–149]. Culture can provide an organisation with a sense of identity and determines rituals, beliefs, meanings, values, norms and language, effectively the way in which ‘things are done around here’ [150]. Schein argues that “the most intriguing aspect of culture as a concept is that it points us to phenomena that are below the surface, that are powerful in their impact but invisible and to a considerable degree unconscious” [151]. Schein goes on to use an analogy that culture is to a group, what personality or character is to an individual in the sense that an individual’s personality and character guides and constrains their behaviour, culture guides and contains the behaviour of members of a team or an organisation [151]. Schein emphasises [152] the invisible levels of corporate culture including underlying values, assumptions, beliefs, attitudes and feelings, while Deal and Kennedy [150] emphasise the more visible levels of culture, such as heroes, rites, rituals, legends and ceremonies. There have been suggestions in the literature that often change strategies will focus on changing visible culture levels [147]. This is a view which is shared by Deal and Kennedy [150], who state that it is the visible culture attributes that will shape the behaviour of employees within an organisation.

3.6.1 Types of Organisational Culture

O'Donnell and Boyle [147] argue that there is not just one single type of organisational culture and that different organisations can have distinctive cultures. Furthermore, a single organisation can have more than one type of culture [147]. Hence, numerous studies [150, 153–155] have been undertaken examining the culture and norms that take can shape an organisation. Hofstede [153] undertook an initial study of IBM, a large multinational organisation where employees in 64 countries were studied in order to identify *national cultures*. Hofstede [153] explained that this was necessary because there are national and regional cultural groupings that can affect the behaviour and culture within organisations. The results from the IBM study allowed Hofstede to identify four dimensions of culture: *power distance*, *intolerance for uncertainty and ambiguity*, *individualism versus collectivism* and *masculinity versus femininity* [153]. In a later study, Hofstede,et al [154] examined the culture within 20 organisational units located in 10 different companies in Denmark and the Netherlands. One of the results from the study was a six-dimensional model to identify cultures within organisations, which is presented in Table 3.2 [154, 155].

In their book, Deal and Kennedy [150] propose an alternative organisation cultural model which focuses on how quickly employees within an organisation receive feedback, the way its employees are rewarded, and the level of risks taken within the organisation itself. The model is based on four types of culture within an organisation [150]:

- **Work-hard, play-hard** – in a work-hard, play-hard organisational culture, employees will take few risks but any feedback on their performance is almost immediate.
- **Tough-guy macho culture** – in this type of organisation, there exists an 'all-or-nothing' culture where individualism prevails and employees will work hard to become stars within the organisation. Teamwork is not valued in this type of organisation.
- **Process culture** – in this type of culture, feedback is slow and the risks taken by employees are low. Stress can arise within these types of organisation because of internal politics, bureaucracy and problems with systems currently used in the organisation.
- **Bet-the-company culture** – in these types of organisations, decisions are high risk and employees will have to wait for a long period of time before knowing if their actions actually paid off.

Cameron and Quinn [156] argue that a major feature in a successful organisation is that it will have a distinctive, readily identifiable organisational culture. This is an opinion that is shared by various several researchers [157–159] in the organisational culture domain, who

Type of Culture	Description
Means-oriented vs. Goal-oriented	In a means-oriented organisation, the key feature is the way in which work has to be carried out. In a goal-oriented organisation, employees focus their attention on achieving specific internal goals, even if these involve substantial risks.
Employee-oriented vs. Work-oriented	In employee-oriented organisations, employees feel that personal problems are taken into account and that the organisation will take responsibility for their welfare. In a work-oriented organisation, there is heavy pressure to execute tasks, even if this is at the expense of employees.
Local vs. Professional	In local organisations, employees will identify with the individuals with whom they work. In a professional organisation the identity of an employee is determined by his profession and/or the content of the job.
Open System vs. Closed System	In an ‘open’ organisation, new employees are made to feel welcome almost immediately and it is believed that almost anyone would fit in the organisation. In a ‘closed’ organisation it is the opposite.
Easy-going work discipline vs. Strict work discipline	An easy-going work culture reveals loose internal structure within an organisation, a lack of predictability, and little control and discipline. A strict work discipline reveals that people are very cost-conscious, punctual and serious.
Internally-driven vs. Externally-driven	In an internally-driven culture, employees perceive their work towards the outside world as totally given. In an externally-driven culture the only emphasis is on meeting the customers requirements; results are important and an ethical attitude is dominant.

Table 3.2: Hofstede’s Organisational Culture Model

recognised that organisational culture can have an effect on the performance of the organisation, as well as its long-term effectiveness. Egan, et al [160] state that in order for an any organisation to remain competitive, its employees need to adapt and learn to handle changes in their organisational environment. As a result, Egan, et al [160] suggest that embedding a *learning culture* can help an organisation ensure that knowledge is shared between both individuals and teams to improve the organisation as a whole. However, Handy states that installing a learning culture in a classic bureaucratic organisation, which are usually process-orientated (such as the studied organisation) can be a difficult challenge [161].

3.6.2 Organisational Learning

Organisational learning is "the process of forming and applying collective knowledge to problems and needs" [162]. Alternatively, Huber defines "organisational learning as a process where organisations aim to incorporate and disseminate valuable experience and knowledge across its communities of practice over time" [163]. Organisational learning can take place through a variety of ways including through direct experience, through interpreting the experiences of others and by encoding knowledge and information into organisational memory [164]. The objectives of organisational learning is to correct errors, reduce the time it takes for employees to undertake corrective actions and to help make strategic decisions within the organisation itself [165]. An organisation which becomes skilled at creating, organising, storing, retrieving and transferring knowledge learnt within the organisation and then modifying behaviours in order to reflect the new insights gained is known as a 'learning organisation' [166].

Organisational learning theories are concerned with how organisations adapt their behaviour and learn in their specific environment [167]. The origins of organisational learning theory can be traced back to research undertaken by Argyris and Schon [165]. Argyris and Schon argued that there are two types of learning that can take place within an organisation: *single-loop learning* and *double-loop learning* [165]. Single-loop learning is an approach that is used in most organisations on a daily basis, and works on the principle that employees will detect and correct errors and deviations from policies, procedures or expected norms. [18, 165]. In contrast, double-loop learning involves employees questioning the very principles such as the policies and procedures the organisation itself functions on and examines if these are the cause of the error or deviation [18, 165].

Argote [168] agues that "understanding how groups or teams learn to work effectively together provides micro foundations for understanding organizational learning because groups are the building blocks of most organizations". Argote goes on to state that teams can learn and develop knowledge in several different ways including eliciting or sharing information that a team member already possesses or generating new information through collaboration and interaction [168].

Gill [162] presents an similar approach for describing organisational learning by arguing that an organisation can learn through different levels including individual learning, team learning and whole organization learning. Individual learning can occur when individual employees acquire knowledge, develops their skills and adopts new attitudes and beliefs that will help the organisation to improve [162]. Team learning occurs when members of a group of employees discover together how to contribute and improve the performance of the group. Effectively, the employees within the group learn from each other and apply that knowledge to improving the purpose of the group [162]. In contrast to individual learning, any knowledge which is gained through learning resides with the team as a whole and not any single individual within the team [169]. Whole organisation learning is concerned with the process whereby employees and teams can learn, grow and change as result of experiences they have encountered with the organisation [162, 169]. Gill argues that one of the most important conditions for 'whole organisation learning' to occur is that managers eliminate boundaries and allow a free flow of information across the whole organisation [162]. While Gill proposed these three levels of learning, there is very little work in the literature examining how these levels of learning would be affect security incident response learning nor how it would influence learning in an information security team. Cooke [120] has discussed barriers that need to be overcome with regards to safety-critical organisations becoming "learning organisation", but this was not discussed from a security incident perspective. In Cooke's discussion he argues that the culture within these organisation need to change so that people are dealt with fairly, incidents are openly discussed, and corrective actions are implemented in a cross-functional team environment [120].

3.6.3 Organisational Learning and Security Incident Response

Researchers [18, 44, 122, 170] have argued that even though many incident response approaches include a 'post-mortem' or 'follow-up' phase, there is little evidence to suggest that many organisations actively engage in organisational learning or look how to improve their incident response processes. Melara et al. [171] enhanced Cooke's 'Incident Learning System' [120] to take into consideration an insider threat attack on an organisation. By analysing pre-cursor 'events' as well as 'incidents' in-depth using double-loop learning, Melara et al. reported that a serious insider threat attack could have been averted using organisational learning [171].

Baskerville, et al. [170] have stated that organisations need to modify their behaviour as to the purpose of security incident response. Baskerville, et al. [170] argue that traditional security incident response approaches focus on detecting losses and reacting quickly, efficiently, and effectively in recovering from security incidents. Instead, Baskerville, et al. have proposed that the focus of security incident response is shifted to developing organisational learning in order to prevent and deter security incidents [170].

Shedden, et al. add that "if organisations were to appropriately learn from and manage their security incident response capability, they would be able to leverage opportunities to learn from incidents to their best advantage and realise the benefits of a robust process and fortified security strategy" [44]. Shedden, et al. analysed the literature on security incident response and organisational learning literature and suggested that organisations could look to integrate double-loop learning into their security incident response process [44]. The proposition from Shedden, et al. is that double loop learning would help organisations learn more appropriately and that any learning would focus on underlying issues in security and organisational structures and not just quick short-term changes to prevent incident recurrence.

To help assist organisations with security incident response in the petroleum industry, Jaatun [122] incorporated organisational learning theories in their Incident Response Management (IRMA) method. IRMA's incident learning included both single-loop and double-loop learning. Single-loop learning was included so that any response was based on "the difference between expected and obtained outcomes", while double-loop learning was focused on questioning and changing "governing variables related to technology, organisation and human factors that lead to the outcome" [122]. While the literature has indicated that there is a gap in how local learning (single-loop learning) from security incident response teams can be translated into double-loop learning with the wider organisation, very little research has examined how this can actually be achieved.

3.7 Research Context

The emergence of information security governance, along with a growing list of legal and regulatory information security requirements has resulted in many organisations adopting Information Security Management Systems (ISMSs) and information security standards. An important component in the implementation of these security standards is the assessment and improvement of new or existing security controls to improve the overall information security environment within an organisation.

While activities such as risk assessments and the implementation of information security controls can help prevent security incidents, no information security solution is a guarantee against successful attacks. As a result, an organisation needs to develop plans and procedures to not only manage the eradication and recovery of a security incident, but to also define how lessons learned will be derived from a security investigation. To help organisations minimise the effects of a security incident along with managing an organisation's return to an acceptable security posture, numerous security incident response approaches have been published by various organisations. However, researchers have identified several problems with these approaches and have stated that many organisations do not find it easy to implement them. As security incidents increasingly impact organisations, it is imperative that organisations have the ability to investigate, report and, ultimately, improve overall security efforts based on previous security incidents. Although security incidents are unwanted problems for an organisation, they do present an opportunity to improve upon existing or identify new security controls based on the lessons learned derived from a security incident investigation.

Incident learning is typically undertaken in the follow-up phase in a security incident response process. Within this phase, it is recommended that a security incident response team are recommended to conduct lessons learned meetings, where the primary objective should be to improve security incident handling procedures and security controls. There have also been calls for organisations to look beyond the immediate causes of security incidents and examine the underlying root causes. However, previous case studies in organisations have identified that many organisations tend to focus efforts more on eradication and recovery and are not exploiting their ability to learn from security incidents. This imbalanced focus can result in the loss of opportunities to investigate why a potential security incident has not been prevented by existing security controls, or what further security controls improvements could prevent similar future incidents. Researchers add that many organisations are simply not prepared to learn from security incidents.

With many organisations focusing more on eradication and recovery, there is the potential that poor quality data could be emerging from the security incident response processes used within these organisations. In the safety domain, a report has cited poor quality data as one factor hindering incident investigations in the transport safety sector. There could be a similar problem in the security incident response domain. Organisations are coming under increased legal and regulatory pressure to share lessons learned from security incidents within industry communities. Therefore, it is becoming increasingly important that research examines how incident learning can be improved within the security incident response domain and whether poor quality data is a factor, which is hindering security incident learning.

Chapter 4

Security Incident Response in Practice: An Exploratory Case Study

Chapter three outlined several security incident response processes used within organisations and established that researchers have reported that organisations focus more on eradication and recovery than security incident learning. The next step was to investigate how an organisation performs security incident response and to examine the challenges that emerge when a security incident response team attempts to produce lessons learned. This chapter reports on an exploratory case study of the challenges a security incident response team in a Fortune 500 Organisation faces with regards to lessons learned development.

The chapter is divided into the following sections. Section 4.1 outlines the objectives of the case study in greater detail. Section 4.2 presents an analysis of the organisation's internal documentation repository, with a specific focus on documents related to security incident response processes. Section 4.3 presents an analysis of the organisation's security incident response database. Section 4.4 presents the results of a survey conducted with practitioners who have been involved in the management and handling of security incidents within the organisation. Section 4.5 introduces initial problems identified from the case study, which limit the organisation's ability to learn from security incidents and Section 4.6 summarises the chapter.

4.1 Exploratory Case Study Objectives

As part of the industrial research project, the author undertook an exploratory case study within a Fortune 500 Organisation. The purpose of the exploratory case study was to identify and explore the challenges that a security incident response team face when attempting to undertake security incident learning. The research objectives of the exploratory case study were to:

1. Investigate the documented security incident response processes within the Fortune 500 Organisation, which define how security incident response should be undertaken within the organisation.
2. Compare observed security incident response in practice with documented processes.
3. Examine how the organisation's security incident response team attempt to learn from security incidents.
4. Identify what challenges the security incident response team encounter when attempting to learn from security incidents within the organisation.

In order to answer these research objectives, the case study involved two stages of data gathering [30]. The first stage involved analysing relevant documentation and the organisation's security incident response database. This was done in order to determine how management expect security incident response to be undertaken within the Fortune 500 organisation. The second stage involved undertaking semi-structured interviews with practitioners within the organisation, who have been involved in various stages of an investigation or management of a security incident. This was done to attain an understanding of the security incident processes and incident learning challenges from the practitioner's perspective. Finally, results from the document analysis and case study were related to the available case study literature in order to validate findings and support generalisation.

In previous case studies [18, 19, 44], which have been used to investigate security incident response learning challenges, researchers have used interviews as the primary method of data collection and documentary evidence to support their findings. However, to the extent of the author's knowledge, this is the first case study that has analysed multiple in-depth data points (i.e. relevant documentation, an organisation's security incident response database, as well as semi-structured interviews) in order to understand security incident learning challenges in an organisation. Hence, it can be difficult to compare the findings from the Fortune 500 Organisation case study with those from previous work.

4.2 Security Incident Response Document Analysis

In order to acquire an understanding of the Fortune 500 organisation's context, the organisation's internal documentation repository was analysed to identify documentation related to security incident response processes. Security incident response within the organisation is managed by the Information Technology Service Incident Response (ITSIR) team and the

Information Security Incident Response (ISIR) team. These two incident response teams have very different objectives, goals and approaches to handling incidents.

ITSIR focuses on incidents that impact availability and accessibility, while ISIR specifically investigates information security incidents, including those, which have a regulatory impact which have been identified within the organisation. The organisation has developed a range of documentation, which governs how these incident response teams achieve these objectives. This documentation includes the Information Technology Service Incident Response Process, the Incident Review Process and the Information Security Incident Response Process. These will be discussed in the following subsections.

4.2.1 IT Service Incident Response Process

The Information Technology Service Incident Response (ITSIR) team is a dedicated team residing in the Operations and Information Technology Services unit. The team works full-time in an incident response capacity. The size of the team can quickly expand to dozens of incident handlers, engaging anyone in the organisation from their normal role in order to remedy an incident. The ITSIR team's goal is to ensure the continuous availability and accessibility of any service that is provided by the organisation. There are three components that interact in the identification, escalation and investigation of such incidents. These components include Helpdesk Services, the ITSIR Management Process and the Incident Review Process.

Incidents that affect the availability of information resources are reported through the organisation's Helpdesk. A Helpdesk incident tracking system is used to monitor and log the progress of such incidents. If the Helpdesk can provide a solution, the incident is logged in the tracking system as a 'Service Event'. However, if the Helpdesk cannot find a solution, the 'Service Event' ticket is logged and then assigned to the ITSIR team as a 'Service Incident'. The incident is then classified by the level of service impact, which determines how long before the ITSIR team must provide a solution. There are four levels of classification: critical, high, medium and low. Incidents classified as 'critical' must be responded to immediately and service must be restored within two hours. Incidents classified as 'low' can be responded to within 24 hours, while service must be restored within 12 weeks.

The ITSIR Management process is then initiated, which involves holding an incident meeting where key roles and actions are assigned to resolve the issue. Depending on the service issue, a technical meeting can also take place. The purpose of this meeting is to diagnose the root cause of the incident and to make a technical recommendation for service restoration. In the event that the root cause diagnosis is determined to contain a security threat to the organisation, the ITSIR Manager will then appoint an information security manager from within the organisation to the role of Security Coordinator. This coordinator is then responsible for creating and implementing mitigation strategies to minimise risk to the organisation; compiling a threat review report and defining how the incident could affect the confidentiality and integrity of the organisation's information resources. Incident meetings are then held periodically until a solution is found and normal service availability is restored.

4.2.2 IT Service Incident Response Review Process

In the event that a service incident has been classified by the ITSIR team as either 'critical' or 'high priority', the team initiates the Incident Review Process. This is a formal post-incident process, which facilitates the identification and assignment of actions required to prevent the re-occurrence of the service incident. A number of meetings take place, where a review coordinator invites business units that are affected by the incident to be involved in the review process. The purpose of the meetings is to confirm the root cause of the incident, establish the business impact and to identify rectifying actions in an attempt to mitigate the risk of the incident reoccurring.

Individuals in the affected business units are assigned rectifying actions to complete and are expected to fulfil their actions and update an incident record in the incident review tracking system. When all actions are completed, the review record is then sent to senior management who will either approve the review for closure or reject the proposed actions. If the actions are accepted, the review process is then closed and all affected business units are notified of the closure through the incident review tracking system. If the actions are not accepted, the issues around the rejection are raised during monthly meetings with Operations and Information Technology Services senior management. New remedy actions to close the review are proposed, discussed and documented in the meeting minutes. Once all the remedy actions set out in these meetings are accepted, the incident is then considered closed.

4.2.3 Information Security Incident Response Process

The organisation's Information Security unit is responsible for implementing everyday operational security, enforcing security controls during the development process and investigating information security incidents. The Information Security Incident Response (ISIR) team is an ad-hoc team of individuals who are a part of the Information Security unit. This team facilitates the identification and assignment of actions required to prevent the recurrence of an issue, which have been determined to be a 'security incident'. The ISIR team follows a customised security incident response approach. The approach, as shown in 4.1 is comprised of four phases: incident detection and reporting; recording, classification and assignment; investigation and resolution; and incident closure.

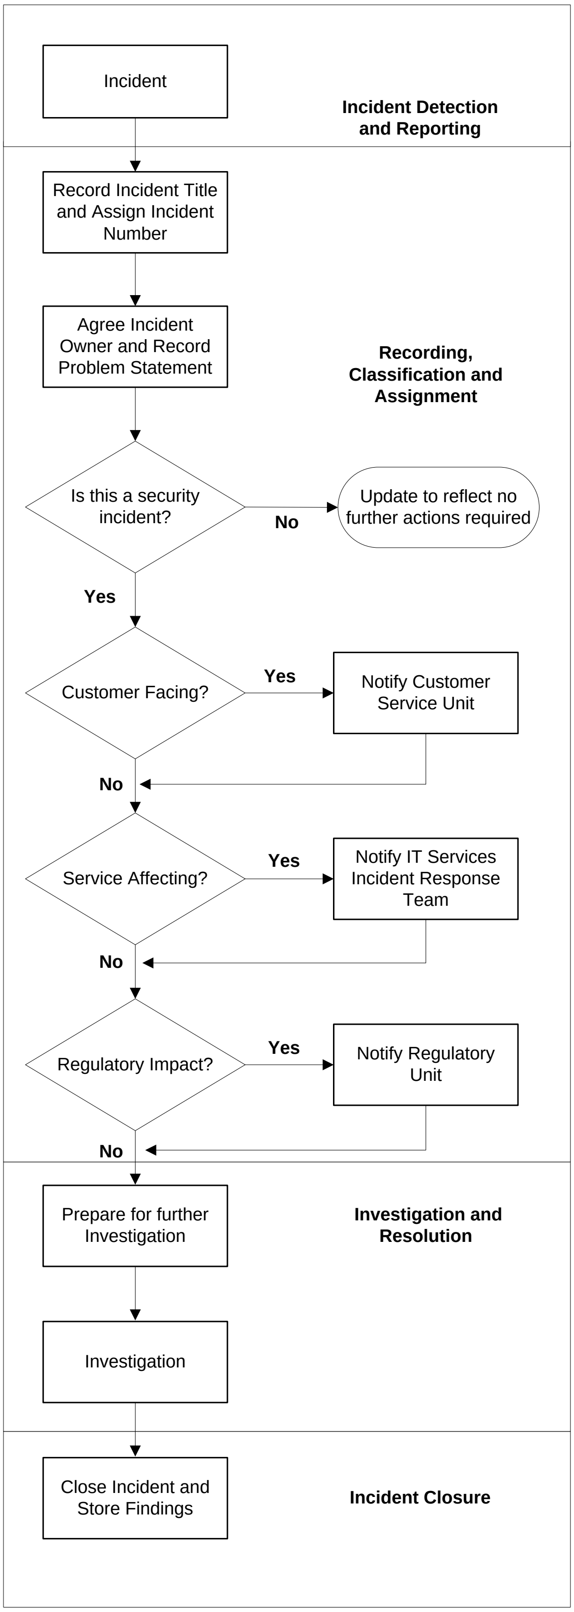
Incident
Incident Detection and Reporting
Record Incident Title and Assign Incident Number
Agree Incident Owner and Record Problem Statement
Recording, Classification and Assignment
Is this a security incident?
No
Update to reflect no further actions required
Yes
Customer Facing?
Yes
Notify Customer Service Unit
No
Service Affecting?
Yes
Notify IT Services Incident Response Team
No
Regulatory Impact?
Yes
Notify Regulatory Unit
No
Prepare for further Investigation
Investigation and Resolution
Investigation
Close Incident and Store Findings
Incident Closure

Figure 4.1: The Security Incident Response Process

The incident detection and reporting phase is concerned with the reporting of an 'incident' to the ISIR team. This can typically come from one of the following sources: a direct request from senior management within the organisation; a request from a member or management of the Information Security unit; a request from the Legal Services unit; or a request from the Human Resources department within the organisation. During the recording, classification and assignment phase, the ISIR team will determine if a security incident has really occurred. However, the term 'security incident' is not defined in the documented process.

If the 'incident' is security-related, an investigation record is created in the Information Security Incident Response database. An incident handler together with senior management of the Information Security unit will agree on the problem statement. Depending on the impact of the security incident, different stakeholders could be involved in the subsequent management and investigation. For example, security incidents that are determined to have an impact on availability or accessibility are referred to the Information Technology Service Incident Response (ITSIR) team and the ITSIR Management Process is invoked. Likewise, if the security incident is determined to have a regulatory impact, a governance process is invoked together with the organisation's risk unit, although the incident is still managed by the ISIR team.

The third phase of the process (investigation and resolution), identifies the evidence and information that is required to conduct an investigation. At this point, the ISIR team holds an incident meeting where the root cause is supposed to be established and remedy actions associated with the incident are assigned to individuals. These individuals are expected to fulfil their actions and update the incident handler upon their completion. The final phase, incident closure, involves two stages. First, relevant stakeholders are notified that all assigned actions have been completed and the security investigation record is updated to reflect the closure of the incident. The second stage requires that the incident handler stores any findings and lessons learned acquired from the investigation in the ISIR database. At this point, the security incident is closed.

As noted in the discussion above, the organisation has several documented processes and procedures highlighting how security incident response should be conducted within the organisation. However, researchers [172–174] have observed that a common trend is that employees do not comply with their organisations' information security policies and procedures. As a result, actions actually undertaken are often very different to what is stated within documented processes [172–174]. Therefore, the next step in the case study was to investigate if security investigations are managed and handled as per documented processes within the organisation. This involved analysing the organisation's Information Security Incident Response (ISIR) database and undertaking semi-structured interviews with relevant individuals who are a part of the ad-hoc security incident response team.

4.3 Security Incident Response Database Analysis

The purpose of the database analysis was twofold. First, the analysis was used to identify the type of security threats that the ISIR team need to investigate within the organisation. Second, the analysis was used to investigate the content of the security investigation records. This was done in order to determine what information is actually recorded about security investigations conducted by the ISIR team and if this information follows the documented process requirements. The results from this analysis are discussed below.

4.3.1 Security Investigation Record Structure

The Information Security Incident Response (ISIR) database is hosted on a IBM Lotus Notes server within the organisation. Within this database, individual security investigation records are stored as separate documents. Each document includes a copy of the security investigation record template, as shown in Figure 4.2. The investigation record template consists of three parts, shown as parts A - C in Figure 4.2. The labels have been added to the record template to aid with the discussion below.

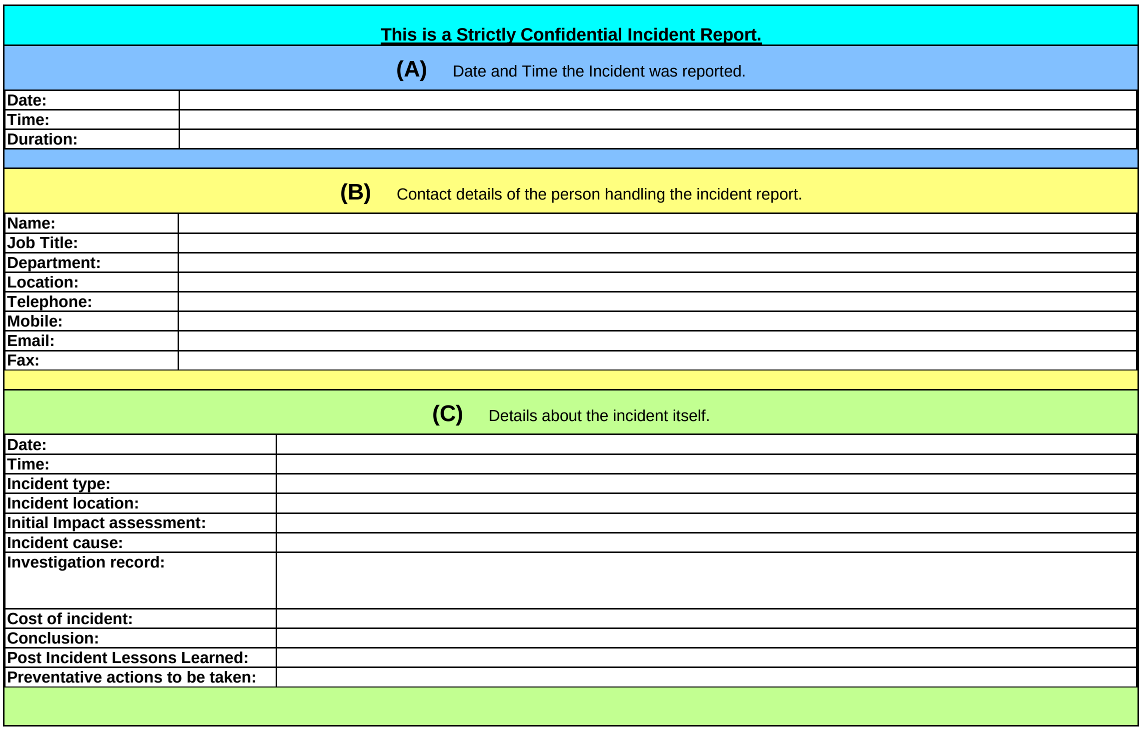

This is a Strictly Confidential Incident Report.	
(A) Date and Time the Incident was reported.	
Date:	
Time:	
Duration:	
(B) Contact details of the person handling the incident report.	
Name:	
Job Title:	
Department:	
Location:	
Telephone:	
Mobile:	
Email:	
Fax:	
(C) Details about the incident itself.	
Date:	
Time:	
Incident type:	
Incident location:	
Initial Impact assessment:	
Incident cause:	
Investigation record:	
Cost of incident:	
Conclusion:	
Post Incident Lessons Learned:	
Preventative actions to be taken:	

Figure 4.2: The Security Incident Response Record Template

Section A of the template prompts incident handlers to record information concerning the date and time the security incident was reported. In addition, there is a third field in this section called 'Duration', which is used to document how long a particular security investigation took to complete. Section B of the template concerns contact details about the individual who is *managing and handling* the investigation within the ISIR team. Information which is recorded includes the incident handler's name and job title; the name of their department and its location; their telephone and mobile phone numbers; email address and fax number.

Section C within the template provides fields where incident handlers are expected to document information about the investigation itself. Although no confirmation is provided in the record template or within the ISIR process, the 'Date' and 'Time' fields within this section appear to be used to document the date and time the investigation was started. The purpose of the 'Incident Type' and 'Incident Location' fields is to document the type of investigation and its location in the organisation. The 'Initial Impact assessment' and 'Incident Cause' fields are used to document any initial assessment of how the incident has affected the organisation and what caused the incident to occur. However, the ISIR process does not elaborate on what information should be documented in these fields, nor is any categorisation taxonomy provided. The 'Investigation Record' field provides a space for the ISIR team to document and record investigation events as and when they occur. At the conclusion of an investigation, the incident handlers can complete the remaining fields at the bottom of Section C. The 'Cost of Incident' field can be used to record the resources expended on an investigation, while the 'Conclusion' field provides a space for the incident handler to document concluding remarks from the investigation. The final two fields 'Post Incident Lessons Learned' and 'Preventive Actions to be Taken' are used to document and record any lessons learned identified from the investigation, as well as any actions which need to be taken post-incident.

As of August 2013, a total of 188 security investigations were recorded in the ISIR database. These investigation records were documented in the database from November of 2003 to August of 2013. The security investigation records stored in the database were then analysed from three perspectives, 1) analysis of security incidents, 2) categorisation analysis, 3) completion of individual fields within the records and completion of the investigation records.

4.3.2 Analysis of Security Incidents

Initial observations from the analysis of the Information Security Incident Response (ISIR) database demonstrated that no consistent categorisation taxonomy was applied to the investigation records in the database. As a result, similar security investigations that could have been categorised under a single common category, were in fact categorised under several different 'categories'. Based on the above observations, a categorisation taxonomy was applied to the security investigation records within the database. This taxonomy is based on a combination and expansion of concepts proposed by Lindqvist and Jonsson [175], as well as Schieber and Reid [176]. The taxonomy was validated by an information security manager within the organisation. This taxonomy consists of:

- **Denial of Service** - an attack or an attempt to make an information system or network resource unavailable to its intended users within the organisation.

- **Digital Forensics and E-Discovery** - any digital forensics work or E-discovery requests to the ISIR team.

- **Compromised Information Asset** - an attempted or successful compromisation of an information system, network device, application or user account.

- **Unlawful Activity & Regulatory Incidents** - computer-related incidents of a criminal nature, likely involving law enforcement or having a regulatory impact.

- **Malware** - malicious software typically affecting one or more information systems, network devices or applications.

- **Email Incidents** - spoofed email, spam, and other security-related email investigations.

- **Policy, Process or Procedure Violation** - an investigation involving a violation of one or more the organisation's information security policies, processes or procedures.

- **Masquerading** - an investigation focusing on an attacker who pretended to be an employee or authorised user of a system in order to gain system access, information or greater access privileges.

The author read the 188 investigation records and one category from the taxonomy above was applied, based on the content of the investigation record. Table 4.1 presents a breakdown of the security investigations in the ISIR database based on the proposed taxonomy. Several observations can be drawn from this table. The analysis of the ISIR database shows that the number of security incident investigations within the organisation is increasing. From 2011 to 2013, the number security investigations handled by the ISIR team increased more than threefold. The ISIR team handled 17 security investigations in 2011 and had already handled 56 investigations in the first eight months of 2013. More specifically, during the same period of time, there was also a large increase in the number of digital forensic investigations and E-discovery requests being recorded in the database. This category accounted for the largest number (41%) of investigations recorded by the ISIR team. The increase in digital forensic investigations and E-discovery requests is not unexpected as many of the organisations business transactions are undertaken digitally and are processed and stored within information systems.

Furthermore, the number of investigations classified as 'masquerading', dropped to zero during the period 2007-2013, while 19 investigations categorised as 'masquerading' occurred

Incident Type/Year	2003	2004	2005	2006	2007	2008	2009	2010	2011	2012	2013
Denial of Service Incidents	0	0	0	1	0	0	1	0	2	0	0
Digital Forensics and E-Discovery	0	0	0	1	0	0	0	0	4	27	46
Compromised Information Assets	0	0	0	1	1	0	4	0	3	4	2
Unlawful Activity & Reg. Incidents	1	2	3	1	1	0	0	0	3	3	4
Malware	0	3	1	1	0	1	1	1	1	0	0
Email Incidents	0	0	0	1	0	0	0	0	0	1	1
Policy, Process, Procedure Violation	0	0	0	3	5	3	8	4	4	12	3
Masquerading	0	1	1	17	0	0	0	0	0	0	0
Total	**1**	**6**	**5**	**26**	**7**	**4**	**14**	**5**	**17**	**47**	**56**

Table 4.1: Analysis of Security Incident Response Database

during the period 2004-2006. There were 17 such investigations in 2006 alone. Exploratory conversations with individuals within the ISIR team identified that many of these investigations involved examining spoofed websites on the World Wide Web which were found to be mimicking several of the organisation's customer services. Since 2006, the responsibility to investigate this specific type of security problem has been delegated to another team within the organisation and therefore, such investigations are no longer documented in the database.

4.3.3 Security Investigation Record Analysis

In addition to analysing the security incidents handled by the ISIR team within the organisation, the 188 security investigation records in the database were also analysed from the perspective of the quantity and quality of information documented within the records.

Categorisation Analysis

No formal security incident response categorisation taxonomy exists within the organisation. However, the use of the word 'incident' in many of the investigation records strongly suggests that the majority of the records stored in database are considered to be 'security incidents'. Exploratory consultations with practitioners within the organisation suggested that a large number of the investigation records which were classified as 'security incidents' were not all 'incidents', but a combination of 'security incidents and security events'. In this context, the practitioners suggested that 'security events' are security investigations which do not require regulatory reporting and therefore are 'less important' within the organisation. The practitioners went on to argue that the ISIR database does not provide a true indication of the number of actual 'incidents' which are affecting the organisation, whose number are much lower than those reported in the database. The European Union Agency for Network and Information Security (ENISA) [177] and the British Government [178] have both provided guidelines outlining how information security incidents should be reported to respective legislators and regulators. The lack of a formal incident response categorisation taxonomy could mean that the Fortune 500 organisation may find it difficult to fulfil its regulatory obligations towards ‘real’ security incident reporting. While this particular information in the investigation record could appear trivial, it is this piece information which is used in the organisation to determine which investigations need to be reported as ‘security incidents’ and which do not require to be reported. Furthermore, recording and documenting a security investigation type using a security incident response categorisation taxonomy means that management within the organisation can compile accurate statistics on previous investigations and access the severity or impact of incidents when they occur [175]. An interesting observation from the analysis of the organisation’s documentation, as well as the expectation from management is that incident handlers are required to document and record ‘Incident Type’ information, yet no formal security incident response categorisation taxonomy was provided to the incident handlers.

An analysis of the ‘Incident Type’ field within the investigation record template was also undertaken. This analysis examined what category information was actually documented within the investigation records in the database. The results of this analysis showed that similar investigations are often classified under several different categories in the database. For example, the following ‘categories’ have all been used to categorise investigation records where potential data loss has been an issue: potential data exposure, potential data leakage, potential security breach, exposure of live data, email to the wrong person, and loss of data. This observation reiterates the lack of a consistent security incident response categorisation taxonomy within the organisation.

Completion of Individual Fields within Investigation Records

The analysis of the investigation records was extended to examine the extent of information, which was documented in the 22 fields in the record template used by the ISIR team. This analysis involved reading the 188 investigation records and identifying what information was being record and what information was not being recorded, as per the organisation’s process requirements. Table 4.2 presents the results of this analysis.

The results show that only one field, the ‘Investigation Record’ field, was completed in all 188 investigation records. This field describes the actions taken during the investigation. A further seven out of the 22 fields were completed in at least 94% of the analysed investigation records. In contrast, the ‘Duration’ and ‘Cost of Incident’ fields were found to be *incomplete* in 87% of the investigation records. Information in the ‘Duration’ field was recorded in only

Field Name	Number of Records (% of Overall)
Date	181 (96%)
Time	176 (94%)
Duration	24 (13%))
Name	186 (99%)
Job Title	185 (98%)
Department	184 (98%)
Location	180 (96%)
Telephone	180 (96%)
Mobile	62 (33%)
Email	164 (87%)
Fax	39 (21%)
Date	178 (95%)
Time	167 (89%)
Incident Type	145 (77%)
Incident Location	97 (52%)
Initial Impact Assessment	68 (36%)
Incident Cause	155 (82%)
Investigation Record	188 (100%)
Cost of Incident	25 (13%)
Conclusion	54 (29%)
Post-Incident Lessons Learned	50 (27%)
Preventive Actions To Be Taken	28 (15%)

Table 4.2: Analysis of Fields within Security Incident Response Database

24 records, while information in the ‘Cost of the Incident’ field was only documented in 25 investigation records. From the perspective of the documentation of lessons learned, 50 investigation records (27%) contained information within the ‘Post-Incident Lessons Learned’ field, while only 28 investigation records (15%) contained information within the ‘Preventive Actions to be taken’ field. This finding shows that in nearly three quarters of the analysed records, no lessons learned were documented, even though this is a documented requirement in the ISIR process.

An additional observation from the analysis of the investigation records was that the ISIR team was only documenting contact information for the incident handler assigned to the investigation. Information about individuals who are *reporting* a security incident is not documented, nor is it required according to the ISIR process. As a result, very little information about the individuals who have reported security incidents is actually documented in the investigation records.

There is also a level of ambiguity within the investigation record, with regards to specific information that is required to be documented. For example, the ‘Date’ field is first used in Section A of the record template. Therefore, it can be assumed that this ‘date’ field is concerned with the date and time the security incident was reported to the team. However, a 'Date' field is also present in Section C, which is concerned with details about the incident itself. Further examination of the ISIR process document revealed that the documented process does not specify what 'date' information should be documented within Section C. There are several potential uses of this field including the date the investigation record was opened, the date the incident was first identified or discovered, or the date the first mitigating actions were undertaken. This means that some ambiguity could arise between incident handlers if any analysis of this information is undertaken as part of the lessons learned aspect of the investigation.

The results show that there is also a problem with *incomplete* and *inconsistent* investigation records within the ISIR database. Only 1 out of the 188 investigation records analysed from the database was considered to be 'complete' from the perspective of all 22 fields. This means that within 187 investigation records, one or more fields within the record template were missing information. 15 records were missing information between one and three fields; 41 records were missing information between four and six fields and 117 records were missing information between seven and ten fields. Furthermore, 14 records were missing information in eleven or more fields. This finding suggests that the incident handlers are not completing the entire investigation record during the investigation of security incidents.

If we consider Ballou and Pazer's data quality dimensions of [144], the results from the investigation record analysis can also be discussed from an *accuracy, completeness and consistency* perspective. From the *accuracy* perspective, analysis of the 'date field' in Section A from the record template revealed that some of the information in this field is inaccurate and does not represent the actual information the field is supposed to represent. The problem is further inflated because the documented process does not specify what information should actually be recorded in the template.

With regards to *completeness*, the analysis of all the investigation records has shown that only 1 out of the 188 records were 'complete' from the perspective of all 22 fields. However, one could argue that not all the missing information is important or will be used for incident learning at a later stage. Abandoning some of the information collection requirements (e.g. fax number field) could mean that the incident handlers focus on gathering only relevant information which will help them learn better about a particular incident. While this could be true for some of the fields, it is worth noting that when these fields were included in the record template, this was may have been done to fulfil potential regulatory requirements. As a result, even though this information may not help the security incident response team learn from an incident, the additional information could be needed for regulatory purposes.

Challenges with the *consistency* dimension is evident in the analysis of the 'Incident Type' field within the investigation record template. Over the past decade, security incident handlers in the organisation have been inconsistent with the names they have used to label particular types of investigations. For example, several different label names have been given for an investigation where potential data loss has been an issue. As a result, incidents of the same type have been categorised using a variety of different names. This can make analysing metrics, and identifying trends and incident type patterns particularly difficult within the organisation. In summary, the *accuracy, completeness and consistency* challenges identified through the analysis of the investigation records suggests that the quality of data within the records could be one factor hindering incident learning within the Fortune 500 organisation.

4.3.4 Security Incident Response Metric Analysis

As a security incident response team matures, it can be beneficial for an organisation to assess how well they are conducting their security incident response operations [179]. West-Brown, et al. [108] and Wiik, et al. [180] argue that metrics provide an accurate way of quantifying the performance of security incident response teams. West-Brown, et al. [108] define two metrics as examples, *response time* and *total time to resolve*. West-Brown, et al. [108] defined response time as the period of time from the first report of a security incident to the implementation of the first mitigating actions. They define the total time to resolve an incident as the time from first reporting and initiation of an investigation to closure. These two metrics were calculated as part of the case study analysis using information in the investigation records within the ISIR database.

52 out of the 188 investigation records contained information to calculate a response time. In the remaining 136 records, although a reporting date and time was documented in the record, no information was available regarding the date and time of the first mitigating actions. The calculated minimum response time was two minutes and the maximum response time was 325 minutes in the 52 records considered. The mean average response time was 56.30 minutes. Figure 4.3 presents the percentage of security investigations which were responded to over a given period of time. The figure shows that the ISIR team documented taking mitigating actions within 30 minutes for approximately 60% of the recorded security investigations and within two hours for 90% of the investigations.

62 of the investigation records contain data concerning the total time to resolution. The remaining 126 records had information missing relating to the closure of the investigation. Out of the 62 investigation records, the minimum time to resolve an investigation was half a day and the maximum time was 130 days. Therefore, the mean average documented time to resolve an investigation was just under 12 days. Figure 4.4 presents the percentage of investigations that were resolved from the first reporting to closure over a given period of time.

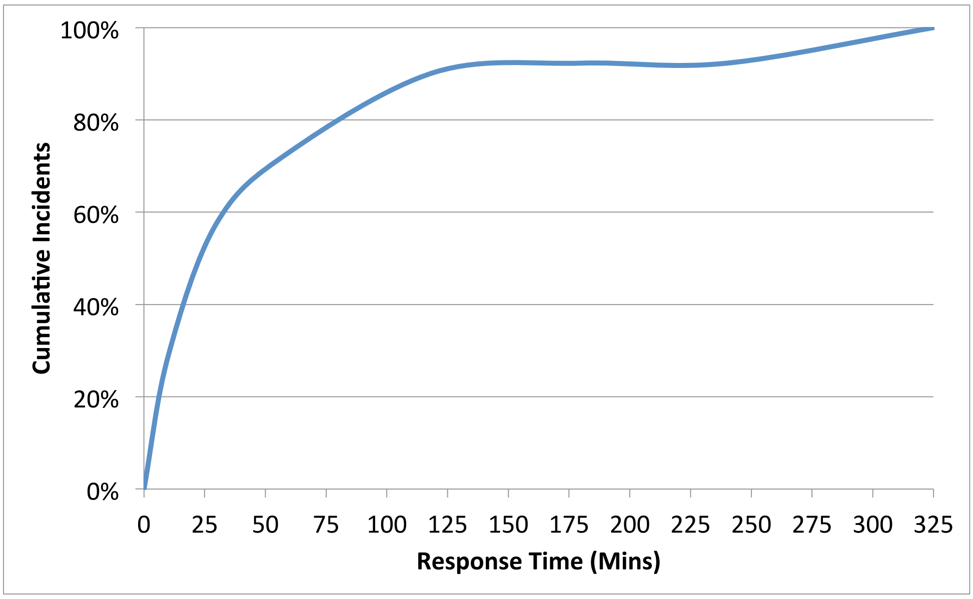

Figure 4.3: Cumulative Response Time Analysis for 52 records

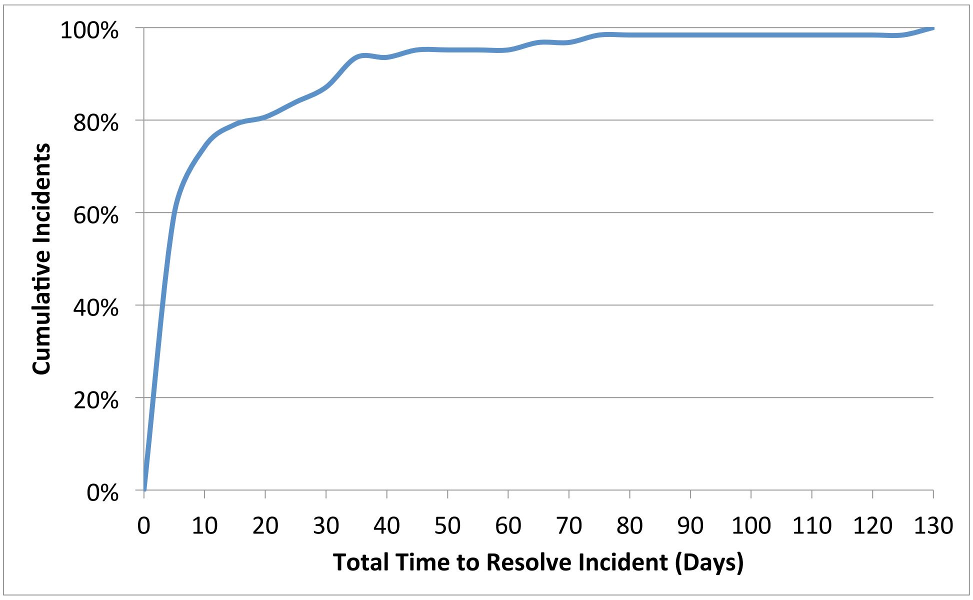

Figure 4.4: Cumulative Total Time to Resolve for 62 records

The results of this metric calculation show that the ISIR team documented information about resolving 20% of the investigations in half a day and 80% of the analysed investigations are resolved within 20 days. The results also show that a small percentage (1.6%) of the investigations took longer than three months (120 days) to complete.

Although West-Brown, et al [108] note that such metric information is useful to analyse the historical performance of a security incident response team, there are no published recommended times to evaluate such teams. However, for the purpose of this case study, the missing data that could have been used in the calculation of the metrics highlight a potential issue with the quality of data within the investigation records. The metric calculations have highlighted how much information is actually missing from the investigation records, which could have been used in the above calculations. This problem then raises the question as to whether the calculated metrics actually provide a reliable measurement of security incident response performance within the organisation, if so much information is actually missing from the investigation records. Further, given the variability in recording information needed to calculate metrics, it is not clear whether the data that was available provides an accurate measure of the organisation's performance.

4.4 Security Incident Response Interview Analysis

In addition to analysing the organisation's Information Security Incident Response (ISIR) process and database, in-depth semi-structured interviews were also conducted within the organisation. The interviews were conducted between November and December 2013 and involved a variety of individuals who have in the past, been involved in or supervised security incident investigations in the organisation. The objective of the interviews was to examine how security incident response is perceived and undertaken within the organisation by practitioners. The interviews were also used to explore any challenges to conducting security incident response within the organisation and also examine data gathering and incident learning.

4.4.1 Interviewee Demographics

15 semi-structured interviews were conducted within the organisation during the exploratory case study. The interview sample consisted of individuals in a variety of information security roles and who have a diverse range of work experience within a technical role. The initial interview questions were used to establish the interviewee's current role in the organisation and quantify his/her number of years of experience in information technology.

The answers from these questions revealed that the interviewees have a maximum of 39 and a minimum of 2 years experience within an information technology role. The mean average experience of the interviewees was thirteen and a half (13 ½) years. The individuals identified themselves as information security managers, senior security analysts or security analysts who assume various roles within the organisation. The following subsections present an analysis of the results from the interviews. These results examine the individual's perception of the ISIR process and security incident response in general, data gathering within the ISIR process, challenges to conducting security incident response within the organisation, incident learning and dissemination.

4.4.2 Security Incident Response Process

In order to comprehend the challenges faced by the ISIR team, initial questions focused on examining the practitioner's perspective on the ISIR process. Some of the findings from these questions confirmed initial results presented in previous sections. The findings are discussed below.

The interviewees confirmed that the organisation uses a customised linear document-centric security incident response approach which consists of four phases 1) incident detection and reporting; 2) recording, classification and assignment; 3) investigation and resolution; and 4) incident closure. The results from this question conclusively established that such a process does exist within the organisation, with twelve individuals stating "Yes" responses and three individuals stating "Do Not Know" responses. Although twelve individuals indicated that the organisation has a security incident response process, only five people could recall the process itself. This suggests that immediate mental recollection of the details in the documented process was limited to a few individuals. These individuals included two information security managers, one senior security analyst and two security analysts.

Five out of the twelve individuals, who indicated that a documented security incident response process exists within the organisation, suggested that the process is not always followed. When queried as to the reasons for deviating from the process, answers included time constraints, a lack of staff to run the entire process, a lack of support for handling specific security investigations and a rigid, document-centric approach to processing a security investigation, which contains too many steps and is considered to be inappropriate for certain types of investigations.

The general indication from the participants is that the documented ISIR process provides structure to the ISIR team. The documentation offers insight into how the ISIR team can resolve security investigations and provides clarity on the escalation path to other business units and senior management within the Information Security team. However, the results also revealed that the current ISIR process is not dynamic when considering quick resolution, and lacks detail regarding the level of information that is required for individual security investigations. The two findings identified above could provide one explanation as to why investigation records are not always completed. If the documented process is not always followed and lacks detail regarding the level of information required within investigation records, then security incident handlers may not capture enough information to produce lessons learned at a later stage in the process.

4.4.3 Perception of Security Incident Response

Interviewees were queried as to what the term ‘security incident’ meant to them. A wide variety of answers were received, which included “a breach of security policy”, “a degradation or circumvention of security controls”, “data loss”, “financial losses” and “a threat to service availability”. The variety of answers received from this query indicates that there is a lack of consensus on how a security incident is viewed within the organisation. This can also indicate that the organisation does not have a unified definition for the term ‘security incident’. This result supports the initial observations that not all the investigation records in the database are true ‘security incidents’ and that no documented definition for the term ‘security incident’ exists within the ISIR process.

Although all the participants have been involved at some stage in the management of a security investigation, when asked specifically if the organisation has a security incident response team, thirteen out of the fifteen respondents indicated that there is no dedicated team. Actually, the participants indicated that the organisation’s ISIR team can be best described as an ad-hoc security incident response team, where individuals are brought together to investigate particular incidents.

Interviewees were then asked which phases of the security incident response process they have been actively involved in within the organisation. The results of this inquiry are summarised in Table 4.3. The phases provided to the respondents were derived from the literature [7, 8, 11] and individuals were allowed to identify their participation in more than one phase.

The table indicates that the majority of participants reported that they were involved in the identification, eradication, investigation, and recovery phases of security incident response. The lack of perceived involvement in the preparation and follow-up phases reflects Shedden, et al. [44] and Jaikumar’s [181] contention that security incident response teams are viewed as ‘fire-fighters’ within organisations, whose role it is to detect, eradicate and recover the organisation from security incidents.

When the participants were queried about improving the documented ISIR process, ten out

Phase Name	No. of Participants
Preparation	5
Identification	9
Containment	7
Eradication	9
Investigation	13
Recovery	10
Follow-up	4

Table 4.3: Security Incident Response Involvement

of the fifteen individuals indicated that the process could be improved, while five individuals provided “Do Not Know” answers. The ten individuals provided numerous potential enhancements to the process including the process should contain more information on how to handle specific security incidents, the process needs to be more available to a wider audience, lessons learned could be more detailed, and previous security incidents need to be evaluated to determine how best to handle future incidents.

4.4.4 Data Gathering within the Process

When asked if the organisation collected information during a security investigation, the majority of the answers returned were positive. Fourteen out of the fifteen individuals indicated that information about security investigations is collected and stored during the response lifecycle. One individual indicated that he/she did not know if the practice took place. The respondents indicated that security information collection was usually assigned and performed by the Primary Incident Handler (PIH). This individual is given the responsibility to collect and record this information within the ISIR team.

The information typically recorded during a security investigation includes investigation meeting notes, actions to be taken for remediation, copies of any logs and emails associated with the investigation, as well as communication between the PIH and management. The security information documented is usually tailored to specific investigations and there does not appear to be a uniform approach to capturing specific information. Based on the individuals answers, it appears the general practice is to capture information which is required to help with the eradication and recovery, but which may not necessarily facilitate incident learning at a later stage. For example, individuals noted that meeting minutes, email trails and the documentation of actions for remediation can be found in nearly all the investigation records.

The interviews confirmed that the above information was not the only information being collected about an investigation. The ISIR team also gathers forensic data from various sources. This forensic data can include logs, emails, hard disk drive images, and physical memory dumps. This type of information can be used as evidence in legal proceedings, if the need arises [182].

Harris [54] states that if forensic evidence is going to be used in subsequent legal action, then an organisation needs to have documented procedures to ensure standardisation when this information is gathered and collected. In the interview, when individuals were asked if there was a process to collect this information, only five individuals indicated that there was such a process within the organisation. One individual indicated that no such process exists and nine 'did not know' if such a process existed within the organisation. Documented processes do exist within the organisation, which describe methods for acquiring data from hard drives, the storage of information in a secure location, as well as an E-Discovery process. However, the results again indicate that recollection of these processes was limited to only a few participants.

Casey [182] argues that if an organisation is going to collect security or forensic data for use in possible legal action, then a chain of custody process should be considered. The five respondents who provided positive answers to the existence of a security data collection process were then queried if the chain of custody practice was performed. Two individuals specified that this practice exists and was done for all incidents. However, two other individuals noted that a chain of custody process does exist but was not performed all the time. There was one 'Do Not Know' answer. The organisation does have a defined chain of custody process, which specifies how an incident handler acquires and stores forensic data to show a continuous chain of custody. The two individuals who noted that the chain of custody was not performed consistently suggested that further guidance was required as to when this should be undertaken.

The respondents indicated that the organisation uses two databases to store information related to security investigations. One database is used to record intrusion detection system alerts and events, while the ISIR database is used to record all other security investigation that involve information security. However, the respondents indicated that access to these databases was limited to only a subset of individuals within the information security unit. These individuals includes managers and information security analysts who manage the organisation's response to a security incident. The participants indicated that there are opportunities for process improvements with regards to data gathering and documentation. The enhancements proposed by the individuals included: making the ISIR database more searchable, providing additional guidance to the incident handlers with regard to what information to record, as well as implementing a 'lightweight' security investigation record. In this context, the individual used the term 'lightweight' to refer to less fields to complete and capturing more relevant information in the investigation record.

4.4.5 Challenges to Conducting Security Incident Response

One of the goals of the interviews was to establish the challenges faced by a security incident response team within an organisation with regards to security incident response. The responses from the interviews highlighted three main challenges for the organisation's ISIR team. First, the process for employee reporting and escalation of security incidents needs to be clearly defined, established, implemented and, periodically, refined as processes and technology evolve. Second, the establishment of comprehensive data access for ISIR team members to perform investigations. Third, resolving conflict between the security incident response team and the business units responsible for the availability of customer-facing assets.

The respondents were asked how security incidents are reported within the organisation. Three individuals stated that there is a documented process for reporting security incidents, six indicated that no such process exists within the organisation, and six provided a 'Do Not Know' answer. The respondent's answers suggest that the majority of security incidents are reported informally either verbally or via email, usually to a known contact within the information security unit. The respondents generally agreed that the reporting of security incidents could be improved within the organisation. The respondent's main concern with the current approach is that when certain members of the information security team were unavailable, reported incidents can take longer to reach the ISIR team. Other potential improvements proposed include an incident reporting hotline and a dedicated email address for reporting security incidents directly to the ISIR team.

Ten out of the fifteen respondents indicated that the ISIR team often has difficulties conducting in-depth security investigations due to a lack of access to security data. There were a variety of answers describing the obstacles preventing detailed investigations including limited physical access to security data, short data retention times, logs not containing enough detailed information and limited support from third-parties involved in security investigations.

The respondents also indicated that conflicts arise between the security incident response team and various other stakeholders within the organisation. From time-to-time, reported security incidents which affect the availability of customer-facing assets can lead to a disagreement over returning the asset back to the production environment and performing a more complete security investigation. This conflict originates from the fact that the organisation relies on the continuous availability of both customer-facing and back-office applications. The respondents also stated that conflicts occur when there is a lack of physical access to systems and applications to extract security data for security investigations. However, the majority of respondents indicated that any conflicts are resolved at a management level.

4.4.6 Incident Learning and Dissemination

Researchers [18, 20, 120] have suggested that many organisations are not maximising their post-incident learning potential and tend to focus on only improving technical processes in an attempt to prevent reoccurrence. When the interviewees were asked if the organisation performs any ‘post-incident’ activities, ten out of the fifteen respondents indicated that depending on the type of security investigation, several activities can take place. These activities are in line with the findings of previous researchers [18, 20, 120], which include implementing security controls to prevent reoccurrence, producing reports for management and education awareness through policy reiteration. There were four ‘Do Not Know’ answers and one individual stated that no ‘post-incident’ activities took place within the organisation.

In order to investigate the extent to which further learning was taking place during the organisation’s ISIR process, the respondents were queried to determine if a root cause analysis was performed post-incident. Seven out of the ten respondents indicated that this was the case within the organisation. The three remaining respondents indicated that they have not been involved in incidents where a root cause analysis was required. One respondent indicated that a root cause analysis should be done for each incident, but they have been involved in incidents where this activity was not performed.

All ten respondents who indicated that the organisation performs ‘post-incident’ activities suggested that there is the potential to enhance and extend these post-incident activities in order to focus on improving the effectiveness of internal policies, procedures, controls and training. A number of recurring themes were mentioned as potential enhancements. These included a deeper analysis during security investigations, improving methods and data to assist in the development of lessons learned, implementing security controls focused more on preventing incident reoccurrence, as well as, increasing the dissemination of lessons learned.

Cooke [120] argued that the work of a security incident response team is usually completed by the issuance of a report, detailing the investigation findings and any lessons learned. Cooke adds that this information can be distributed within an organisation or reported to external regulatory bodies [120]. When the interviewees were asked if any post-incident information was distributed or disseminated to any groups or departments within the organisation, six out of the fifteen respondents said ‘Yes’. Two individuals said that post-incident information was not distributed or disseminated within the organisation and seven individuals ‘did not know’ if this took place.

The six respondents indicated that there are two methods for disseminating post-incident information within the organisation. First, through an electronic announcement on the organisation’s Intranet and second through a monthly pack of statistics prepared for management. However, when queried if a process exists for distributing this information, the respondents were unanimous that there is no formal process for distributing this information within the organisation. The respondents did indicate that the decision to distribute incident information is held with the operational security information security manager, who will decide what information is disseminated and to whom.

The respondents were also asked if any post-incident information was distributed or disseminated outside of the organisation, for example to regulatory bodies. Nine out of the fifteen respondents indicated that post-incident information was distributed or disseminated outside of the organisation. One individual suggested that this did not take place and five individuals noted they 'Do Not Know' if the practice was taking place. When asked if there was a process in place to govern this practice, all nine respondents indicated that they 'did not know' if such a process existed. One respondent did note that if such actions were required, the Regulatory Compliance Unit would disclose the security occurrence to the relevant regulators.

4.5 Initial Challenges Identified from Case Study

The exploratory case study has established that the organisation's Information Security Incident Response (ISIR) process dictates that security incident handlers document and store investigation findings, including lessons learned at the closure of an investigation. However, the analysis of the security investigation records in the ISIR database has revealed that less than 30% of the examined records actually contained lessons learned. Numerous challenges and problems, which have been identified during the exploratory case study, could explain this finding. These challenges and problems can be summarised as a lack of consistent and defined security incident response taxonomy, access to enriched security data, absence of organisational learning from security investigation records, and a lack of tools and methods to conduct more comprehensive incident learning.

4.5.1 Lack of Consistent Security Incident Taxonomy

The responses from the semi-structured interviews have indicated that the ISIR team disseminates security incident response-related information through electronic announcements on the local Intranet and by producing a monthly pack of statistics for management. However, the results from the analysis of the ISIR database showed that no consistent security incident response taxonomy was being applied within the organisation. As a result, multiple category types of security investigations are being used to describe similar investigations. Lindqvist and Jonsson [175] argue that without a consistent security incident response taxonomy, an organisation cannot compile accurate statistics on previous investigations or have a 'grading' system for analysing the severity or impact of incidents when they occur.

With no defined security incident response taxonomy within the organisation, nearly all of the analysed security investigation records in the ISIR database were categorised as ‘security incidents’. Practitioners within the organisation have suggested that many of these ‘security incidents’ are not true incidents. Coupled with a lack of a consistent view on how a security incident is viewed within the organisation, this can result in complications with regards to the regulatory reporting of information security incidents [104, 183].

4.5.2 Access to Enriched Security Data

The final phase of many security incident response approaches focuses on post-incident activities. Two objectives from these post-incident activities include a security incident response team learning from a security incident along with the integration of any lessons learned back into the wider organisation’s security posture [8, 184, 185]. Therefore, security incident investigations can provide information that can be instrumental in avoiding a recurrence of a security incident. However, a security incident response team cannot establish ‘lessons learned’ from a security incident without access to detailed incident security data. Access to detailed incident security data is a problem, which has been identified throughout the exploratory case study.

In the analysis of the security investigation records in the ISIR database, only one out of the 188 analysed records were considered ‘complete’. This means that 187 records were missing information from one or more fields from the record template. Extended analysis showed that 117 records were missing information from between seven and ten fields from within the record template. Further evidence of this problem was evident in the analysis of the documented lessons learned in the examined record. The analysis of the investigation records showed that lessons learned were only documented in 27% of the examined records, while only 15% of the records contained actions to be taken after the closure of an investigation.

The interviews with individuals within the organisation provided some insight on why these fields were missing information. The answers from the interviews showed a lack of perceived involvement in the preparation and follow-up phases. This suggests that the ISIR team are viewed as ‘fire-fighters’ within the organisation, whose role it is to detect, eradicate and recover the organisation from security incidents. Individuals within the interview survey further supported this argument by indicating that the general practice within the team is to capture information which is required to help with the eradication and recovery from a security incident. As a result, because security incident learning is not considered to be a priority, incident handlers do not capture information which can help facilitate incident learning at a later stage.

4.5.3 Absence of Organisational Learning from Security Investigations

Line et, al. [121] argue that while information security incidents are unwanted situations within an organisation, at the same time, they present an opportunity to learn about the risks and vulnerabilities which can exist within an organisation's systems and processes. However, researchers have argued that organisations do not pay enough attention to organisational learning from information security incidents [18, 19, 44]. These researchers have gone on to state that organisational learning concepts such as double-loop learning are critical to ensure that organisations learn appropriately so that underlying security and incident response structures are modified to prevent incident recurrence [18, 19, 44].

The data collected within the studied organisation has revealed that the Information Technology Service Incident Response (ITSIR) team includes an incident review as part of the follow-up activities for any critical or high-impact incident. During this review process, a root cause analysis is performed, which is used to determine actions to prevent a reoccurrence. Furthermore, Information Security Incident Response (ISIR) documentation directs security incident handlers to store any findings and lessons learned achieved throughout the investigation for future reference. In the security incident response interviews, individuals identified that a range of post-incident activities take place including producing reports for management, education awareness and security control implementation. The responses from the interviewees indicated that a root cause analysis was usually undertaken within the security incident response team. However, these findings support the results from previous work that organisations do not pay enough attention to organisational learning concepts [18,19,44].

An observation from the case study is that the security incident response team do not appear to utilise any of the organisational learning models or processes identified in the literature. While the organisation's security incident response team acknowledge that a variety of post-incident tasks take place, many of these tasks can be categorised as *information dissemination* and not organisational learning. In order for organisational learning to take place, the security incident response team (and by contrast the wider information security team) should be focusing on double-loop learning [165]. Double-loop learning would involve the security incident response team questioning and challenging the underlying rules, principles and knowledge of the organisation [165]. While the interviewed individuals have suggested that a root cause analysis is undertaken in the organisation, a double-loop response would begin with the root cause analysis and then be followed by an investigation into the process or policy which resulted in the security incident to occur [18]. In the health and safety domain, Cooke [120] proposes the implementation of a Incident Learning System as part of implementing continuous improvements. However minimal work has examined this solution in a security incident response context.

4.5.4 Lack of Tools and Methods for Deeper Incident Learning

Both the ITSIR and ISIR documented processes recommend that their respective incident response teams undertake Root Cause Analysis (RCA) as part of their post-incident activities. This finding suggests that the organisation is not only focused on the technical issues related to security incidents, but is also concerned with examining the organisational and human factors which may have contributed to the incident [122].

However, the results from the interviews within the organisation established that there is potentially a lack of agreement as to when to undertake a RCA or how to actually perform a RCA within the ISIR team. Furthermore, when individuals were queried as to how post-incident activities can be improved within the organisation, these individuals have called for a more detailed analysis of security incidents, as well as improved methods and data to assist in the development of lessons learns, including RCA.

The results from the case study support previous findings [122, 126] that security incident response teams need further support and clarity surrounding the benefits of RCA, when to undertake and how to perform a RCA within their specific organisations. Stephenson [127] argues that there are not many structured approaches for conducting a RCA of security incidents. This is an opinion which is shared by Johnson [126] who adds that there are relatively few established tools and techniques to support the RCA of security incidents. The case study findings, combined with the literature, supports the needs to empirically investigate tools and methods for conducting a RCA of information security incidents. However, ensuring that the correct tools and methods are available may not sufficient for an effective RCA. Several researchers [15, 186, 187] have indicated that there are several factors which can influence how effective an RCA can be within an organisation. The quality of a RCA is highly dependent on the accuracy of the input data as well as the capability of the investigative team to appropriately use this data [15, 186]. Furthermore, researchers have argued that in some RCAs, only one source of error is emphasised, and once this has been found, an investigative team will discontinue its analysis [186]. Johnson also argues that the investigator has an important part to play in the RCA and their knowledge of the incident and surrounding environment can play an important part in a ‘good’ RCA [15]. RCAs can be time consuming and therefore organisations need to support investigators with adequate resources and management involvement to ensure that the investigator will be able to carry out the task effectively [187].

4.6 Summary

This chapter described an exploratory case study of the security incident response landscape in a Fortune 500 Organisation. The case study has identified that the studied organisation has implemented a customised security incident response process which consists of four phases: incident detection and reporting; recording, classification and assignment; investigation and resolution; and incident closure. However, analysis of the relevant investigation records and results from the survey interviews has shown that security investigations within the organisation can vary from the documented process for a variety of reasons including time constraints and a lack of staff to run the entire process.

Evidence of this variation can also be seen in the analysis of the security investigation records in the Information Security Incident Response (ISIR) database. The results from the analysis of the ISIR database revealed that less than 30% of the investigation records in the database contained lessons learned. This result differs from the organisation's requirements within the ISIR process, which mandates that a root cause analysis is undertaken and lessons learned are documented within investigation records. Furthermore, the interviewees also indicated that detailed root cause analysis are not always performed within the organisation. Several potential explanations emerged from the case study results and interview surveys which potentially explain these findings. These include, a lack of discrimination between low impact security 'events' and significant security incidents which justify further analysis; limited access to detailed security incident data; capturing information which helps eradication and recovery and which does not necessarily help facilitate incident learning; and a lack of tools and methods to assist in the development of lessons learns.

At a high-level, these challenges suggest that improving the quality of data generated from the organisation's security response investigations needs to be addressed. The results from the case study suggests that the data currently generated from the organisation's security incident response process does not appear to help facilitate incident learning either during or after the closure of an investigation. Therefore, alternative solutions are required to enhance the quality of data generated by the ISIR process to not only assist in the eradication and recovery of a security incidents, but also assist with incident learning at a later stage.

Chapter 5

Enhancing Data Collection from Security Investigations

Chapter four presented the results from an exploratory case study of a Fortune 500 Organisation which identified several opportunities to enhance the quality of data within the organisation's security incident response process. This chapter describes an experiment to address two of these opportunities through the introduction of a revised security investigation record template and a well-defined security incident categorisation taxonomy. The experiment contribution is twofold. First, the experiment evaluates the effect of employing a revised security investigation record template that is used to collect data that has value to a security incident response team with regards to incident learning. Second, the experiment evaluates the efficacy of a taxonomy of security incident types with the purpose of removing ambiguity surrounding incident type identification.

The chapter is structured as follows. Section 5.1 reviews the data quality challenges identified in the exploratory case study that are relevant to this experiment. Section 5.2 describes the alterations proposed to the organisation as enhancements to their existing security incident response process. These alterations included a revised security investigation record template and a well-defined security incident categorisation taxonomy. Section 5.3 discusses how the alterations were implemented within the organisation. Section 5.4 presents the results from the quantitative and qualitative analysis undertaken within the organisation to determine the impact of the alterations on the organisation's security incident response process. Section 5.5 discusses the implications of the results and Section 5.6 summarises the chapter.

5.1 Review of Data Generation Challenges

The results from the exploratory case study of the Fortune 500 Organisation presented in Chapter four identified that there were several opportunities to enhance the quality of data within the organisation's security incident response process. An analysis of the organisation's Information Security Incident Response database (hereafter refereed to as the security incident response database) showed that only 1 out of the 188 investigation records in the database contained information in all 22 fields in the investigation record template. Furthermore, the results from interviews conducted within the organisation showed that there does not appear to be a uniform approach to capturing specific information within the organisation's security incident response process. When expanded upon, interviewees indicated that the general practice is to capture information, which helps eradication and recovery from a security incident, but does not necessarily facilitate incident learning at a later stage.

Further examination of the security incident response database showed that at the time of the case study, all the analysed investigation records, which were assigned a category classification, were all classified as 'security incidents'. However, exploratory consultations with security professionals within the organisation revealed that many of these investigation records were not 'security incidents' but a combination of 'security incidents' and 'security events'. These professionals went on to argue that the current state of the database did not provide a true reflection of the number of actual 'incidents' affecting the organisation, which in fact, were a lot lower than those reported in the database.

The results from the exploratory case study also showed that no consistent security incident taxonomy (i.e. identifying specific types of security incidents) was being applied to security investigation records within the organisation. Although some form of classification was being applied, at the time of the study, no pre-defined taxonomy was available to help guide the security incident response team with regards to security incident classification. As a result, security incident handlers generate different category titles, often for the same type of incident. The consequence is that the security incident response team find it difficult to identify trends regarding the number of security incidents occurring within the organisation.

5.2 Proposed Changes

Based on the data quality challenges described above, two alterations to the Fortune 500 Organisation's security incident response process were proposed: 1) the introduction of a well-defined security incident taxonomy to remove ambiguity with regards to incident classification and to allow the security incident response team to monitor incident types; and 2) a revised security investigation record template to help focus investigation efforts on data that provides value to the security incident response team. These two alterations are discussed in more detail below.

5.2.1 Security Incident Taxonomy

The objective of the security incident taxonomy was twofold. First, the taxonomy was projected to improve the consistency of the classification of security investigations. Second, it was anticipated that the taxonomy would help the organisation's security incident response team to monitor security incident trends within the organisation. At a high-level, the taxonomy consists of two main categories: 'security events' and 'security incidents', and sub-categories beneath each main category. Effectively, the sub-categories define different types of 'security events' and 'security incidents'.

In order to define what the terms 'security events' and 'security incidents' would mean within context of the taxonomy, customised definitions were proposed to the organisation. The definitions were developed using an analysis of the relevant literature, along with feedback from one of the organisation's information security managers. The proposed definitions were as follows:

An **information security event** is an identifiable anomaly, alert, report or request, which involves a change in the security state of any information system and/or computer network which warrants attention from the Information Security team.

An **information security incident** is a single or series of adverse events, which has resulted in one or more of the following:

- a violation of an information security policy or policies;
- has resulted in the loss of confidentiality regarding customer or organisational data;
- has resulted in the loss of integrity regarding any information system, computer network, customer or organisational data;
- has resulted in the loss, disruption or denial of service availability;
- has resulted in financial losses due to a change in the security state of any information system and/or computer network.

In addition to defining the two main categories, different sub-categories were also developed and proposed to the organisation. These sub-categories are used together with the 'security event' and 'security incident' categories. A specific sub-category is mutually exclusive to one of the two main categories. The sub-categories were developed using input from industry white-papers to identify current best-practice, as well as the results from the analysis of the

Event Type	Definition
Audit	Any event which involves a systematic evaluation of the security of any organisational information system, device or computer network
Customer Dispute	Any event which has resulted from a dispute involving a customer
Data Subject Access Request	Any event which has resulted from a data subject access request within the organisation
Data Loss Event	Any event which involves the disclosure or loss of customer or organisational information
E-Disclosure	Any event where digital information stored within the organisation is requested as part of litigation support
Equipment Theft/Loss	Any event where an organisation device or system has been lost or stolen
Human Resources Investigation	Any event involving a Human Resources investigation
Internal Usage Investigation	Any event involving an investigation of internal usage of resources
Policies, Process or Procedure Event	Any event resulting from a violation of the organisation's acceptable usage policies, processes or procedures
Regulatory Investigation	Any event involving a regulatory investigation
Security Assistance	Any event where security assistance is required from the security incident response team

Table 5.1: Proposed Security Event Categories

organisation's security incident response database. Feedback was also received from the same information security manager used in the category definitions. The sub-categories proposed to the organisation consisted of eleven types of 'security events' and seven types of 'security incidents'. Tables 5.1 and 5.2 summarise the 'security event' and 'security incident' sub-categories proposed to the organisation, along with the definition of each sub-category.

Several observations can be made with regards to the feedback and final decision made by the information security manager involved in the above proposals. The first of these observations is that the proposal to provide a taxonomy based on 'security events' and 'security incidents' is consistent with recommendations from the literature. Several best practice guidelines [8,9,185] have proposed organising a security incident response taxonomy based on security events and security incidents. However, none of these guidelines have evaluated the effectiveness of this approach in a real-world security incident response environment.

A second observation from the information security manager's feedback is the use of mutually exclusive sub-categories within the proposed taxonomy. On the one hand, the advantage of using a security incident taxonomy with mutually exclusive sub-categories is that an or-

Incident Type	Definition
Data Exposure	Any incident involving the disclosure or loss of customer or organisational information
Fraudulent Activity	Any incident relating to fraudulent activities involving customer or organisational information
Malware Incidents	Any incident where a virus, worm, Trojan horse, root-kit, key-logger, spyware, rogue security software or any other malicious software affects any organisational system, device, network or data
Policies, Process or Procedure Incident	Any incident resulting from a violation of the organisation's acceptable usage policies, processes or procedures
Service Outage	Any incident which has resulted in the unavailability of any organisational information system, application or computer network
Unauthorised Access to Information	Any incident where access to customer or organisational information was gained by an unauthorised individual, information system or computer network
Unauthorised Modification of Information	Any incident where customer or organisational information has been destroyed, corrupted or modified without authorisation

Table 5.2: Proposed Security Incident Categories

ganisation can compile accurate statistics on previous investigations [175]. However on the other hand, limiting a taxonomy to only mutually exclusive sub-categories can also be seen as a potential weakness. Information security incidents are becoming increasingly complex and will often include more than one attack vector [1]. Furthermore, multiple systems or parts of an organisation could be affected in a different way by a single security incident. Limiting a security incident taxonomy to mutually exclusive sub-categories has limited the depth of statistics that can be compiled on the type of security investigations in the organisation. For example, if a security investigation has identified that malware was sending customer information outside of the organisation, in the current taxonomy this would have to be classified as either a 'malware incident' or a 'data exposure' incident. Using more than one sub-category means that the organisation has the ability to compile more in-depth statistics on numerous attack for a particular type of incident, something which is not possible using mutually exclusive sub-categorises.

A third observation from the information security manager's feedback was the decision to limit the taxonomy to particular types of security investigations. Numerous additional categories could have been added to the taxonomy including 'configuration vulnerability', 'theft of resources' or 'physical attack'. Furthermore, no 'Other' option was added to the taxonomy, which means that all security investigations undertaken within the organisation would have to be classified using one of the eleven security event types or one of the seven security incident types. This means there is the potential that a security incident handler could have to use a ‘best fit approach’ for classifying a security investigation, which does not fit completely into the description of the propose categorises.

5.2.2 Security Investigation Record Template

In addition to a well-defined taxonomy, a revised security investigation record template was also proposed to the organisation. The objective of the proposed investigation record template was to focus investigation efforts and data collection on information that provides value to the security incident response team. The investigation record template was developed after collecting requirements from both the security incident response team and one of the information security managers within the organisation. The information security manager made the final decision, as to which fields ‘represent value’ to the security incident response team and was therefore included in the investigation record template. This is the same manager who provided feedback during the development of the taxonomy.

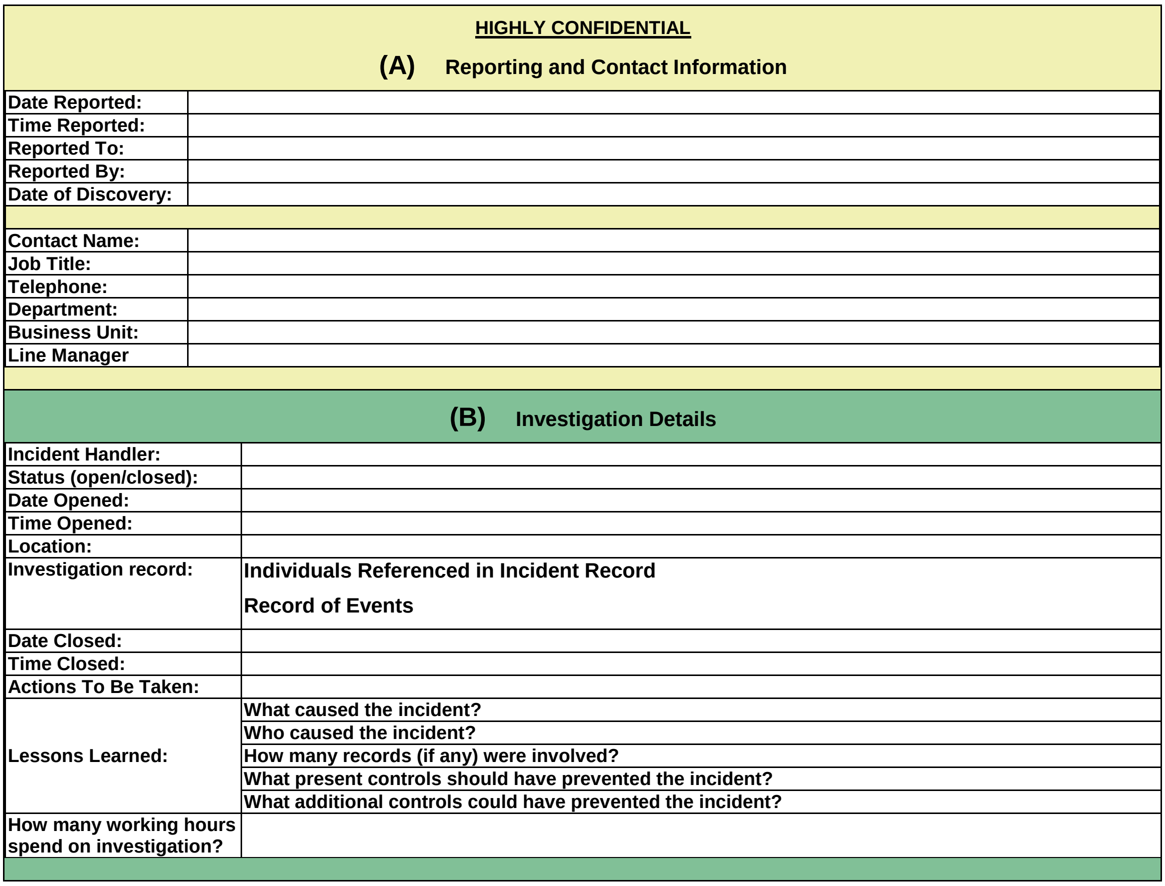

HIGHLY CONFIDENTIAL	
(A) Reporting and Contact Information	
Date Reported:	
Time Reported:	
Reported To:	
Reported By:	
Date of Discovery:	
Contact Name:	
Job Title:	
Telephone:	
Department:	
Business Unit:	
Line Manager	
(B) Investigation Details	
Incident Handler:	
Status (open/closed):	
Date Opened:	
Time Opened:	
Location:	
Investigation record:	Individuals Referenced in Incident Record Record of Events
Date Closed:	
Time Closed:	
Actions To Be Taken:	
Lessons Learned:	What caused the incident?
	Who caused the incident?
	How many records (if any) were involved?
	What present controls should have prevented the incident?
	What additional controls could have prevented the incident?
How many working hours spend on investigation?	

Figure 5.1: Proposed Security Investigation Record

The investigation record template, as presented in Figure 5.1, consists of 26 fields within two sections: Section A - Reporting and Contact Information and Section B - Investigation

Details. Labels have been added to the record template to aid the discussion below. Section A provides fields to document information about the initial report of the security problem (either a 'security event' or a 'security incident') to the security incident response team. This information includes the date and time the problem was reported; to whom the problem was reported to; the name of the individual who is reporting the problem; and the date that the problem was first discovered. The bottom half of Section A, contains fields to document the contact details of the individual with knowledge about the problem under investigation. The purpose of these fields is to assist the security incident response team and management to quickly identify whom to contact about further information regarding the specific investigation.

Section B of the investigation record template, contains fields for the documentation of information related specifically to the security investigation itself, which includes the documentation of 'Lessons Learned'. The purpose of the 'Incident Handler' field is to provide for accountability to the investigation record. Information documented within this field allows the security incident response team and management to know who is responsible for the record and potentially, the security investigation. The 'Date Opened' and 'Time Opened' fields are used to record the date and time the security investigation record was first opened. Similarly, the 'Status' and 'Location' fields are used to document whether the record is currently 'Open' or 'Closed' as well as the location of the security problem within the organisation's infrastructure. The 'Investigation Record' field is used to document the actions undertaken during a specific investigation. The 'Date Closed' and 'Time Closed' fields are used to document the date and time the investigation record was completed and closed within the team.

In order to encourage security incident learning, the five 'Lessons Learned' fields prompt the incident handler to document security lessons through a series of five questions. The five questions, which are posed to the incident handler as part of the 'Lessons Learned' field are:

- What caused the incident?
- Who caused the incident?
- How many records (if any) were involved?
- What present controls should have prevented the incident?
- What additional controls could have prevented the incident?

The purpose of the last field within the record (How Many Working Hours Spent on Investigation?) is to provide the incident handlers with an opportunity to identify and document how long the incident handler dedicated to the investigation. A detailed summary and description of each field in the investigation record template is presented in Appendix D. The security incident taxonomy and revised investigation record template recommendations were made to the organisation and were accepted and implemented into the organisation's security incident response process.

5.3 Implementation of Recommendations

The recommendations were implemented through modification of the organisation's Information Security Incident Response (ISIR) process document and the security investigation record template stored in the security incident response database. This document defines how security incident response is undertaken within the organisation, while the database is used to store investigation records which describe security investigations managed by the security incident response team. Within the ISIR process document, a new sub-section was added to the Introduction section of the document. This sub-section was used to define the terms 'security event' and 'security incident', as they are presented in Section 5.2.1. In order to implement the security incident taxonomy into the ISIR process, changes were required to define how security investigations are initially identified and classified within the security incident response team. These changes involved modifying the ISIR process diagram in the process document. This process diagram defines how security incident investigations are undertaken within the organisation. Figure 5.2 presents the new process diagram with the changes made to the process shown in yellow. Labels have been added to the process to assist with the discussion below. An overview of the activities within the individual phases of the modified ISIR process diagram is discussed below.

Step 1. Event Reported – this step covers the reporting of a 'security event' to the security incident response team. The modification to the ISIR process means that initially, all security problems are identified as 'security events' within the team until further information is known. The reporting of a 'security event' can come from several sources including directly from Senior Management or the Information Security team's management structure; a request from the Legal Services business unit; a request from the Human Resources department; or any organisational employee.

Step 2. Document Record Title and Assign Reference Number – a new investigation record is then created using the template and assigned a reference number using a pre-defined format *YYYY-XX Record Title*. *YYYY*, refers to the year that the record was initiated, while *XX* is the number assigned to a security investigation within the organisation. *Record Title* refers to a free-text field used to describe the investigation in one or two words. For example, 'E-Discovery Request - Grispos', describes an E-Discovery request concerning an individual called 'Grispos'.

At this stage of the investigation, the following information must be documented within an

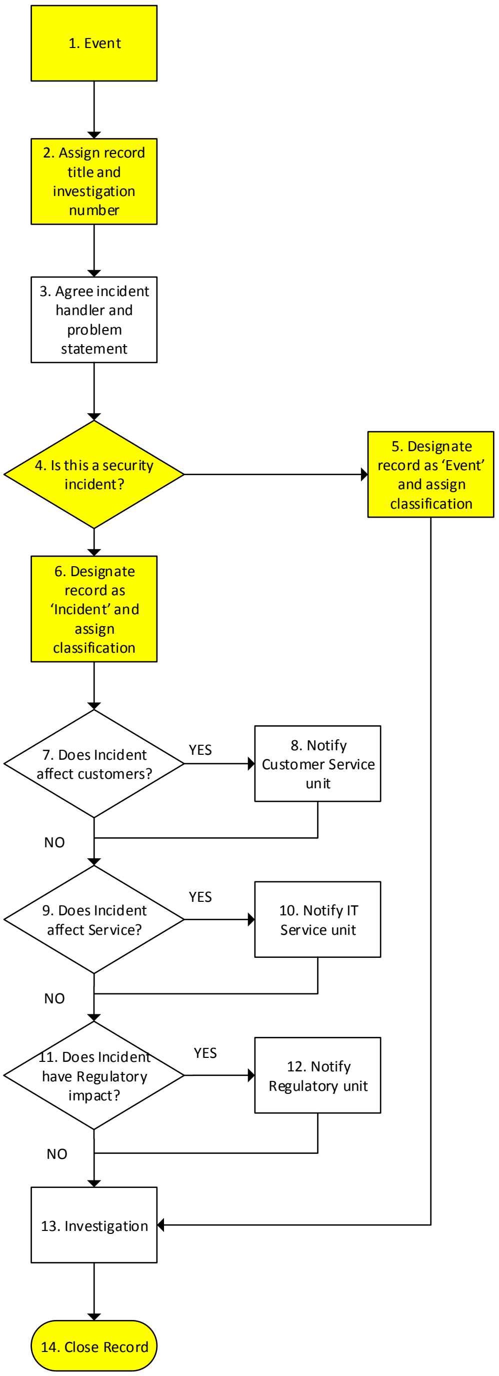
1. Event
2. Assign record title and investigation number
3. Agree incident handler and problem statement
4. Is this a security incident?
5. Designate record as 'Event' and assign classification
6. Designate record as 'Incident' and assign classification
7. Does Incident affect customers?
YES
8. Notify Customer Service unit
NO
9. Does Incident affect Service?
YES
10. Notify IT Service unit
NO
11. Does Incident have Regulatory impact?
YES
12. Notify Regulatory unit
NO
13. Investigation
14. Close Record

Figure 5.2: Modifications to Security Incident Response Process

investigation record: the date and time of the initial reporting; the name of the individual to whom the event was reported to; the name of the individual who has reported the event; and the date when the event was discovered or detected within the organisation. In addition, the following information about the individual submitting the event report must also be documented: contact name; job title; telephone; department name; business unit; and the name of the individual's line manager. The security incident response team should also document the following information about the 'security event' as the investigation progresses: incident handler name; the date the investigation record was opened; the time the investigation record was opened; the location of the event; and the investigation record status (Open/Closed). While this information is collected for problems defined as 'security events', if a particular 'event' becomes an 'incident', then this information is used for the investigation of the particular 'security incident'.

Step 3. Agree incident handler and record problem statement – the security incident response team identifies who the Primary Incident Handler (PIH) will be for the investigation. This is the individual who will coordinate the organisation's response to the particular security problem. This information can change during the course of the investigation.

Step 4. Is this a Security Incident? – during this stage of the process, the PIH or the security incident response team assesses whether the reported 'security event' is a 'security incident' using the definitions within the documented process. If the 'security event' *is not* an 'incident', the process continues to Step 5. If the 'security event' *is* an 'incident', the process continues to Step 6.

Step 5. Update the investigation record (Event) – within this step of the process, the PIH reflects within the investigation record, that the investigation is a 'security event'. The PIH is then required to assign an event sub-category to the particular investigation record. For a 'security event', the next stage in the process is Step 13.

Step 6. Update the investigation record (Incident) – the PIH updates the security investigation record template to reflect that the investigation is a 'security incident'. The PIH is then required to assign one of the incident sub-categories to the particular investigation record.

Steps 7 – 12. Who does the incident affect? – during these steps the PIH, together with the security incident response team, are prompted to make three decisions. These decisions determine if the 'security incident' affects any specific part of the organisation. If the incident affects any customers, the security incident response team must notify the Customer Service team, if the incident affects the availability of service, then the Information Technology Service team must be informed. Likewise, if the 'security incident' has regulatory implications, then the security incident response team must inform the relevant regulatory business unit.

Step 13. Security Investigation – this stage involves the PIH investigating and resolving the 'security incident' or handling the 'security event' to completion.

Step 14. Close Investigation Record – this phase is invoked when no further actions are required from the investigation discussed in Step 13. The PIH is then required to document the following information within the record: the date and the time that the investigation record was closed; any actions which need to be taken after the closure of the investigation record; document any lessons learned by answering the five questions within the template; and document the number of hours taken to resolve/handle the 'security event' or 'incident'.

The security investigation record template presented in Figure 5.1 was implemented within the organisation's security incident response database. From the date of implementation, all security investigations within the organisation used the revised security investigation record template.

In order for the security incident response team to use the security incident taxonomy, the categories and sub-categories were implemented as two lists within the Lotus Notes document. These two lists are an application-specific feature within the Lotus Notes database and were used to display to the incident handlers the categories and sub-categories from the taxonomy. The incident handlers then selected, based on their preference, which category and sub-category they wanted to assign to a specific investigation. The changes to both the ISIR process document and the security investigation record template were implemented in February 2014. No further changes were made to either the ISIR process document or the security investigation record template from February 2014.

5.4 Data Collection and Analysis

One of the goals of this experiment was to reduce ambiguity with regards to the classification of security investigations and to assist security incident handlers document information which provides value to the security incident response team. In order to evaluate if the proposed recommendations discussed in Section 5.3 fulfilled these objectives, quantitative and qualitative data are collected from the security incident response database, as well as the individuals involved in the experiment.

Quantitative data was collected through an analysis of the relevant security investigation records within the security incident response database. All the investigation records from February 2014 to March 2015 were captured for review. The data analysis aspect of this experiment was carried out in April 2015. At the time of the analysis, there were a total of 371 security investigation records within the security incident response database. 42 of these records were considered 'open' and were excluded from the analysis. These records were excluded because the specific investigations documented within the records were considered to be 'on-going'. Therefore, not all the fields within these particular investigation records would have been completed at the time of the analysis. A further five investigation records were designated by management to be ‘highly-sensitive’ and their content was only accessible by the information security manager. As a result, 324 investigation records were then analysed from the security incident response database. It must be noted that 59 of these records were documented in April 2014 as part of one large ‘multi-record’ investigation. This is the result of the organisation detecting a security problem and then opening multiple investigation records for different individuals affected by this problem. The 324 investigation records were analysed from four different aspects. The results from both the qualitative and quantitative analysis will be presented together under each specific aspect.

The first aspect, focused on analysing the investigation records from the perspective of identifying what specific ‘security event’ and ‘security incident’ investigations were undertaken within the organisation during the experiment. This was done in order to determine what types of investigations were actually handled by the security incident response team. The second aspect focuses on examining the actual category and sub-category assigned to the investigation records. The aim of this analysis was to first evaluate if the security incident taxonomy was being used and second, how accurately the taxonomy was being applied to the investigation records. The third aspect of the analysis focuses on the security investigation record template. The purpose of this analysis is to determine first, if the 26 fields within the investigation record template were being used and second, to examine if the information documented in these fields can be considered ‘actionable information’. In this context, ‘actionable information’ is “information (that can be) used to take actions that mitigate against future threats, or help address existing compromises” [188]. The fourth aspect of the analysis was to examine if the data generated from the revised security investigation record template can be used to enhance the metric information calculated in the exploratory case study, as discussed in Section 4.3.4.

In addition to analysing the security incident response database, qualitative data was also collected through follow-up interviews with individuals within the organisation. The purpose of the interviews was twofold. First, the interviews were used to gather the practitioners’ perspective on the alterations introduced into the organisation and second, to further explore specific phenomena identified from the results of the quantitative data analysis.

5.4.1 Analysis of Security Events and Incidents

The 324 investigation records were read by the author to determine the type of security investigations documented in the security incident response database from February 2014 to March 2015. Table 5.3 presents a breakdown of the security investigations in the security incident response database based on the sub-categories discussed in Section 5.2.1. Note that only the sub-category types with one or more investigation records in the database are presented in the table.

Security Investigation Type	**Feb 2014 – Mar 2015**
Data Loss Event	59
Data Subject Access Request	14
E-Disclosure	20
Equipment Theft/Loss	5
Human Resources Investigation	2
Internal Usage Investigation	3
Policies, Process or Procedure Event	67
Regulatory Investigation	2
Security Assistance	83
Data Exposure	63
Fraudulent Activity	4
Policies, Process or Procedure Incident	2
Total	**324**

Table 5.3: Analysis of Database: Feb. 2014 – Mar. 2015

The analysis of the security incident response database shows that not all the sub-categories proposed to the organisation, as part of the security incident taxonomy, were used during the experiment. The security incident response team used nine out of the eleven 'security event' sub-categories and three out of the seven 'security incident' sub-categories. This means that six out of the eighteen sub-categories were not used during the experiment. These included the 'Audit' and 'Customer Dispute' event sub-categories and the 'Malware Incidents', 'Service Outage', Unauthorised Access to Information' and 'Unauthorised Modification of Information' incident sub-categories.

The analysis of the security incident response database revealed that the number of security investigations undertaken within the organisation during the experiment has increased dramatically. Section 4.3.2 was used to analyse the number of investigations documented in the security incident response database from November 2003 to August 2013. 47 investigations were undertaken in 2012 and 56 investigations were undertaken in 2013. However, from February 2014 to March 2015, the security incident response team undertook over 300 security investigations. In fact during the first three months of 2015, 61 investigations were recorded in the security incident response database. This already exceeds the number of investigations throughout both 2012 and 2013.

The analysis also shows that the most common type of security investigation now handled by the security incident response team are 'Security Assistance' and 'Policies, Process or Procedure Events'. These two types of investigations accounted for 46% of the total number of security investigations documented in the database from February 2014 to March 2015. The results from the exploratory case study showed that there were on average five 'Policies, Process or Procedure Violations' investigations a year since 2006. However, the new analysis from the database shows that now 67 ‘Policies, Process or Procedure Events’ and two ‘Policies, Process or Procedure Incidents’ investigations were recorded in the database, over thirteen times more investigations. While there has been an increase in the number of ‘Policies, Process or Procedure’ investigations, the results from the analysis have also shown a decrease in the number of E-Discovery requests, when compared to the exploratory case study. Within the exploratory case study, 27 investigations were classified as ‘Digital Forensics and E-Discovery’ in 2012 and 46 such investigations in 2013. However, only 20 investigations classified as ‘E-Disclosure’ were found in the database for the period February 2014 to March 2015.

5.4.2 Category and Subcategory Assignment Analysis

The security incident taxonomy required incident handlers to decide if a particular investigation was a ‘security event’ or a ‘security incident’ and then to choose a sub-category based on this selection. This section presents the results of an analysis to evaluate if first, the taxonomy was used within the organisation and second, how accurately the taxonomy was applied to the 324 investigation records. The results of this analysis showed that all 324 (100%) of the examined investigation records contained both a category and sub-category classification.

As all 324 investigation records contained both a category and sub-category classification, this prompted further analysis into the investigation records. The purpose of the extended analysis was to examine if the correct category and sub-category classification was being applied to the investigation records. This analysis involved examining the 324 investigation records to first identify the assigned category and sub-category and then using the information in the investigation record to determine the scope and focus of the investigation. Based on the information in the investigation record, the category and sub-category assigned within the record were then compared with the definitions and/or categories within the documented process. If the information from the record did not match with the definitions and/or categories in the documented process then this record was considered to be ‘*miscategorised*’.

For example, consider an investigation record that describes how the organisation’s data loss prevention system audited an email, which was successfully sent to a third-party containing unencrypted customer data. If this investigation record were classified as an ‘Event - Data Loss Event’, it would be considered an example of a ‘miscategorised’ investigation record. Since there has been a loss of confidentiality and integrity regarding customer data, according to the definition of the term ‘security incident’ in the organisation, this investigation should have been classified as an ‘incident’ rather than an ‘event’.

25 (7.7%) out of the 324 investigation records were ‘miscategorised’, either from the perspective of an incorrect category, an incorrect sub-category or both. 22 out of the 25 inves-

tigation records were found to have an incorrect category. That is, 22 investigation records were categorised as a 'security event' by their respective incident handlers. However, the investigation record appeared to describe what has been defined as a 'security incident' in the organisation's documented process.

In addition, three investigation records were found to have an incorrect sub-category, all of which involved 'security events'. One record was categorised as an 'E-Disclosure' event, yet the record was found to describe an investigation involving the Human Resources department. As a result, a more appropriate category may have been 'Human Resources Investigation' and not 'E-Disclosure'. The remaining two investigation records were categorised as 'Security Assistance'. However, the information documented within these records suggested that they could also have been categorised under the 'Data Loss Event' sub-category. It must be noted that the 22 investigation records, which were found to contain an incorrect category, also by extension contained an incorrect sub-category. This is because different sub-categories exist for each category and therefore in reality, 25 investigation records were found to have an incorrect sub-category. The follow-up interviews were then used to identify reasons why these 'miscategorised' investigation records were assigned the specific category and sub-category as found in the documented record. Further questions from the interview also focused on identifying any problems or benefits of using the security incident taxonomy within the organisation.

5.4.3 Participant's Perspective on Classification Taxonomy

During the interview, participants were presented with five examples of 'miscategorised' investigation records. The five 'miscategorised' investigation records were selected at random by the author during the quantitative data analysis. The five investigation records shown to the individuals are included below. The purpose was to establish potential reasons as to why the miscategorisation had occurred. Note these examples have been anonymised to protect the identity of the Fortune 500 Organisation.

Example 1

Example 1 was found to be documented as a 'Security Event - Data Loss Event‘, as shown in Figure 5.3. One individual agreed that the investigation record described a 'security event'. However, six individuals disagreed and argued that the investigation should have been labelled as a 'security incident'.

The individual who argued that this investigation was a 'security event' stated that this was because previous security evaluations of the third-party involved had determined that sufficient security controls were in place to prevent additional data leakage. Therefore, based

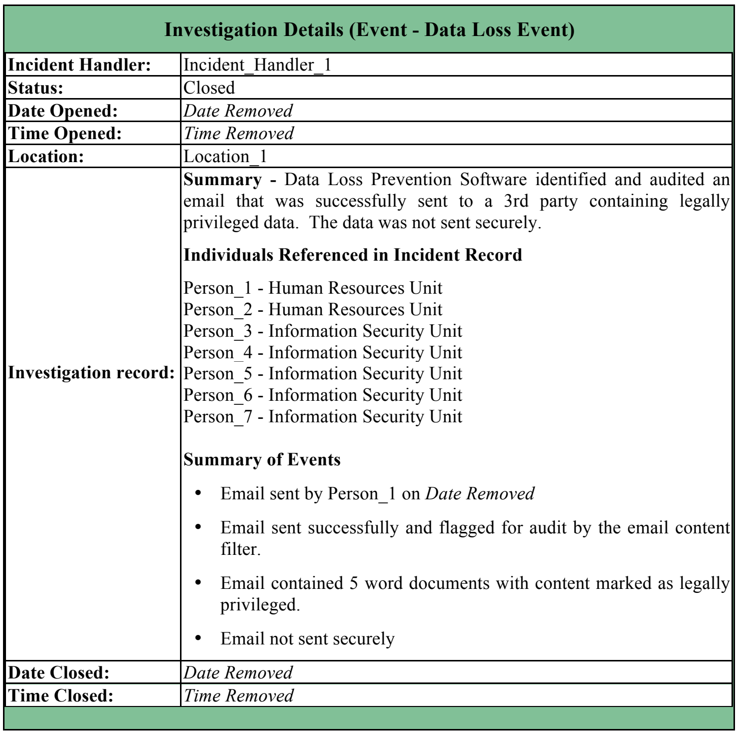

Investigation Details (Event - Data Loss Event)	
Incident Handler:	Incident_Handler_1
Status:	Closed
Date Opened:	*Date Removed*
Time Opened:	*Time Removed*
Location:	Location_1
Investigation record:	**Summary -** Data Loss Prevention Software identified and audited an email that was successfully sent to a 3rd party containing legally privileged data. The data was not sent securely. **Individuals Referenced in Incident Record** Person_1 - Human Resources Unit Person_2 - Human Resources Unit Person_3 - Information Security Unit Person_4 - Information Security Unit Person_5 - Information Security Unit Person_6 - Information Security Unit Person_7 - Information Security Unit **Summary of Events** • Email sent by Person_1 on *Date Removed* • Email sent successfully and flagged for audit by the email content filter. • Email contained 5 word documents with content marked as legally privileged. • Email not sent securely
Date Closed:	*Date Removed*
Time Closed:	*Time Removed*

Figure 5.3: Miscategorised Investigation Record 1

on this information, the individual labelled the investigation as a ‘security event’. The individual went on to add that this label would still apply, even if the investigation uncovered that the third-party should not have had the data in the first place. Note that the information regarding the third-party security controls is not mentioned in the investigation record. This is an example of information about an investigation, which is not always documented by security incident handlers.

The six individuals who argued that the investigation should be labelled as a ‘security incident’ provided two main reasons. These reasons included the data was sent successfully and in the ‘clear’, and that the data contained ‘legally privileged’ information which is considered to be ‘sensitive’. However, two out of the six individuals also noted that in certain circumstances, this investigation could have also been labelled as a ‘security event’. For example, one individual implied that although encryption was not used (which is actually against organisational policy), if the third-party involved was supposed to receive this information, then they would label this investigation as an event. This is because a breach of data had not actually occurred. Another individual added that if TLS (Transport Layer Security) was used and the third-party was *not* supposed to receive this information, then they would label this investigation as a ‘security event’. When asked to explain further, the individual stated that if the secure transport mechanism was in place and the third-party had secure controls to prevent additional data leakage, then they were confident that this would have been an ‘event’.

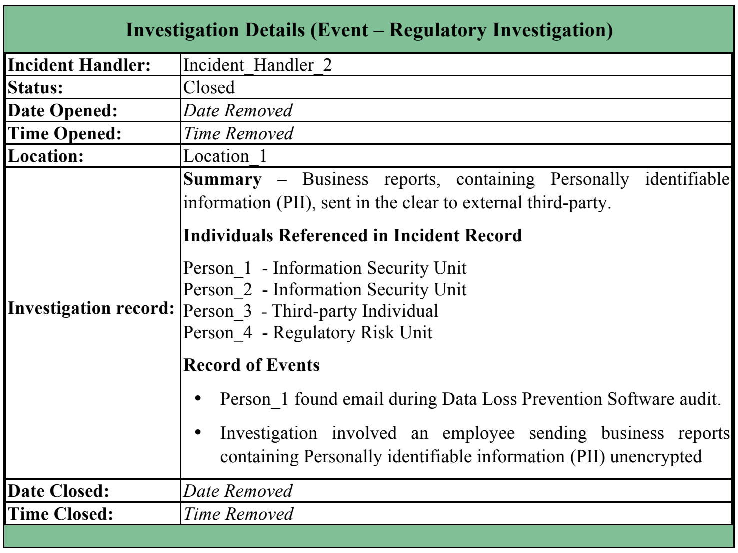

Investigation Details (Event – Regulatory Investigation)	
Incident Handler:	Incident_Handler_2
Status:	Closed
Date Opened:	*Date Removed*
Time Opened:	*Time Removed*
Location:	Location_1
Investigation record:	**Summary** – Business reports, containing Personally identifiable information (PII), sent in the clear to external third-party. **Individuals Referenced in Incident Record** Person_1 - Information Security Unit Person_2 - Information Security Unit Person_3 - Third-party Individual Person_4 - Regulatory Risk Unit **Record of Events** • Person_1 found email during Data Loss Prevention Software audit. • Investigation involved an employee sending business reports containing Personally identifiable information (PII) unencrypted
Date Closed:	*Date Removed*
Time Closed:	*Time Removed*

Figure 5.4: Miscategorised Investigation Record 2

Example 2

Example 2 (presented in Figure 5.4) was labelled as a 'security event' by the Primary Incident Handler (PIH). Two individuals agreed that this investigation record described a 'security event'. However, four individuals argued that the investigation described a 'security incident' and one individual was undecided.

The two individuals who suggested that the investigation was a 'security event' argued that this was because the Personally Identifiable Information (PII) which was sent to the third-party, was neither confidential nor 'enough' for it to be considered a 'security incident'. While the organisation does not have a uniform definition for the term PII, discussions with the two individuals revealed that PII in this case can be defined in terms of the National Institute of Standards and Technology definition:

> "any information about an individual maintained by an agency, including, but not limited to, education, financial transactions, medical history, and criminal or employment history and information which can be used to distinguish or trace an individual's identity, such as their name, social security number, date and place of birth, mother's maiden name, biometric records, etc., including any other personal information which is linked or linkable to an individual" [189].

When the individuals were asked to define 'enough PII' for a 'security event' to be considered a 'security incident', both individuals answered that this would vary depending on the regulatory requirement that was concerned with the investigation.

The four individuals who indicated that the investigation was a 'security incident' provided two main reasons as their justification. First, the data was sent in the 'clear' without the use any encryption and second, the PII was not blocked by any of the organisation's data loss prevention systems. Two out of the four individuals added that because there was no TLS between the two end-points, there was the potential for interception on route and hence, the investigation has to be classified as an 'incident'. These two individuals went on to state that if TLS had been implemented between the two end-points, then this investigation could have been documented as a 'security event'. The individual who was undecided if the investigation was a 'security event' or a 'security incident' stated that further information was required within the investigation record before they could make a decision.

Example 3

Example three as presented in Figure 5.5, was defined by its PIH as a 'security event'. Six individuals agreed that this investigation should be labelled as a 'security event', but one individual argued that the investigation described a 'security incident'. The one individual who argued that the investigation was a 'security incident' stated that this was because the investigation record described a 'regulatory investigation'. This individual was implying that a 'regulatory investigation' means that it must be a 'security incident' and not a 'security event' because of its regulatory nature. This is consistent with the definition of a 'security incident' as defined by the Information Commissioner's Office (ICO) [99].

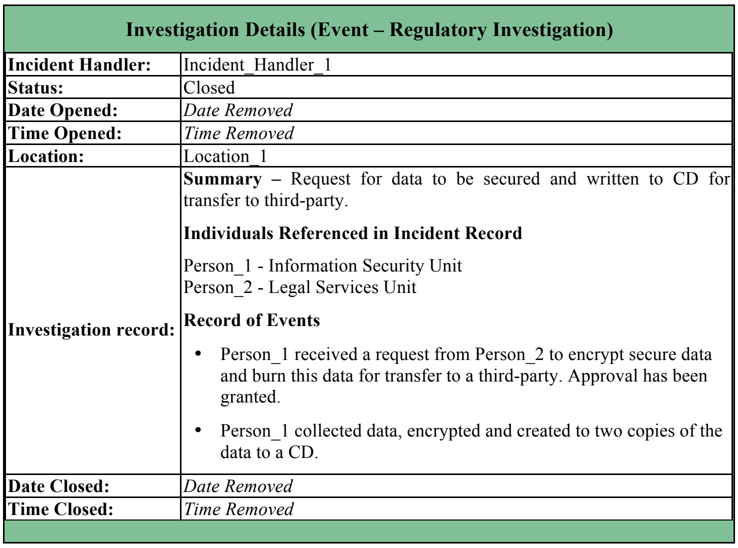

Investigation Details (Event – Regulatory Investigation)	
Incident Handler:	Incident_Handler_1
Status:	Closed
Date Opened:	*Date Removed*
Time Opened:	*Time Removed*
Location:	Location_1
Investigation record:	**Summary** – Request for data to be secured and written to CD for transfer to third-party. **Individuals Referenced in Incident Record** Person_1 - Information Security Unit Person_2 - Legal Services Unit **Record of Events** • Person_1 received a request from Person_2 to encrypt secure data and burn this data for transfer to a third-party. Approval has been granted. • Person_1 collected data, encrypted and created to two copies of the data to a CD.
Date Closed:	*Date Removed*
Time Closed:	*Time Removed*

Figure 5.5: Miscategorised Investigation Record 3

The six individuals who argued that the investigation described a 'security event' stated that this label was a "best-fit approach". The individuals went on the state that this type of investigation was actually neither a 'security event', nor a 'security incident'. Instead, the six individuals suggested that the investigation should be labelled as 'Security Assistance' and nothing more. In fact, one of the six individuals implied that this type of investigation should not really be documented in the security incident response database.

Example 4

Example four, as shown in Figure 5.6, was labelled by its PIH as a ‘security event’. Three individuals agreed and indicated that the investigation did indeed describe a ‘security event’. However, two individuals argued that the investigation described a ‘security incident’, while two individuals were undecided. Two out of three individuals who suggested that the investigation was a ‘security event’ indicated that this was because even though the data in question had been accessed by another part of the organisation, it had still remained within the parent organisation’s physical infrastructure. The third individual who suggested that the investigation was a ‘security event’ indicated that this was based on their understanding of the documented definitions, that an ‘event’ becomes an ‘incident’ when it is confirmed that a security problem has occurred. The individual argued that the security problem and not yet been confirmed and therefore, this investigation was a ‘security event’. The individual also added that the investigation should be extended to examine the access control issue and not only the data loss problem identified in the investigation record.

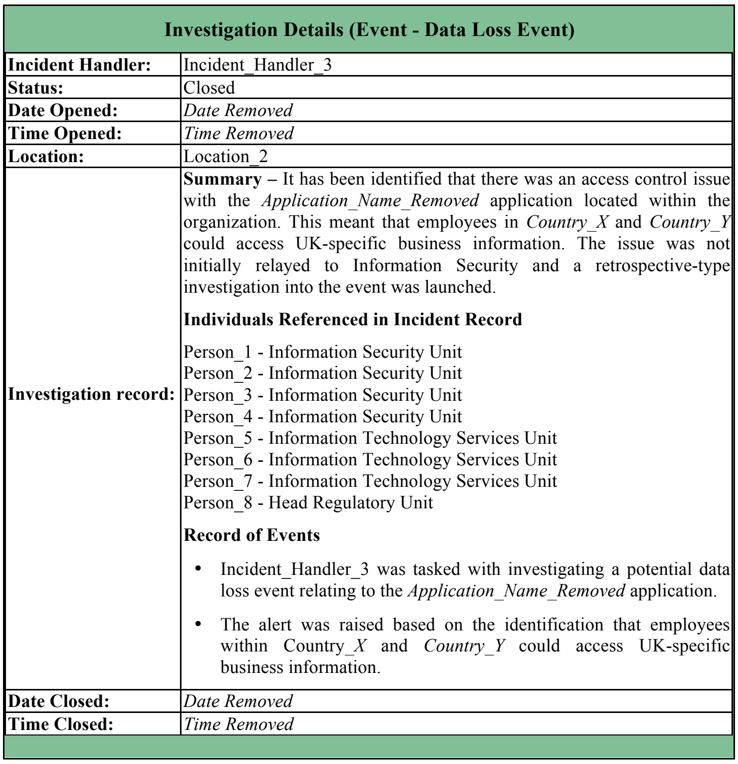

Investigation Details (Event - Data Loss Event)	
Incident Handler:	Incident_Handler_3
Status:	Closed
Date Opened:	*Date Removed*
Time Opened:	*Time Removed*
Location:	Location_2
Investigation record:	**Summary** – It has been identified that there was an access control issue with the *Application_Name_Removed* application located within the organization. This meant that employees in *Country_X* and *Country_Y* could access UK-specific business information. The issue was not initially relayed to Information Security and a retrospective-type investigation into the event was launched. **Individuals Referenced in Incident Record** Person_1 - Information Security Unit Person_2 - Information Security Unit Person_3 - Information Security Unit Person_4 - Information Security Unit Person_5 - Information Technology Services Unit Person_6 - Information Technology Services Unit Person_7 - Information Technology Services Unit Person_8 - Head Regulatory Unit **Record of Events** • Incident_Handler_3 was tasked with investigating a potential data loss event relating to the *Application_Name_Removed* application. • The alert was raised based on the identification that employees within Country_*X* and *Country_Y* could access UK-specific business information.
Date Closed:	*Date Removed*
Time Closed:	*Time Removed*

Figure 5.6: Miscategorised Investigation Record 4

The two individuals who indicated that the investigation record described a ‘security incident' stated that this was the case because unauthorised individuals accessed sensitive information, regardless of the fact that they work within the parent organisation. The two individuals who were unsure how to label the investigation stated that they would require additional information before confirming the classification. The individuals then went on to state that if the parent-company employees were supposed to view this information then the investigation could be labelled as a 'security event'. However, both individuals added that if access to the data were successful, this investigation would likely become a 'security incident' as it would have repercussions under local data protection laws. One of the two individuals also noted that there was a potential 'grey area' with regards as to whether the data involved had actually left the organisation or if it had remained within the organisation's control. This individual was referring to the fact that although the data has left the UK-part of the organisation, it had not actually left the parent company and as a result, was still within the 'organisation' itself. Hence, the individual requested that further clarity be made on defining 'organisation' in this context.

Example 5

Example five as shown in Figure 5.7, was categorised as a 'security event' by its PIH. Two individuals agreed that this was a 'security event', four suggested that the investigation described a 'security incident' and one individual was undecided. The two individuals who indicated that the investigation was a 'security event' argued that the information involved was not confidential and according to the Information Commissioner's Office (ICO), there was not 'enough' Personally Identifiable Information (PII) for it to be considered an 'incident'.

It is interesting to note that the ICO guidelines [99] do not actually define what the term 'enough' entails, but the document does provide guidance on what should be reported to the ICO. One of the two individuals added that because of this assumption, they were inclined to classify the investigation as a 'security event'. However, the same individual added that if the PII was 'enough' then the investigation should be classified as a 'security incident'.

The four individuals who answered that this investigation should be a 'security incident' indicated this was the case because PII was disclosed to individuals who were not supposed to receive this information. However, one of the four individuals did note that if the information in the memo was publicly available, this investigation could be classified as an 'event'. The undecided individual wanted more information documented in the investigation record before they could make a decision.

Two main themes emerge from the above analysis of 'miscategorised' investigation records. The first theme is that there still appears to be a lack of agreement on the definition of a 'security event' and a 'security incident' within the organisation. The results from the

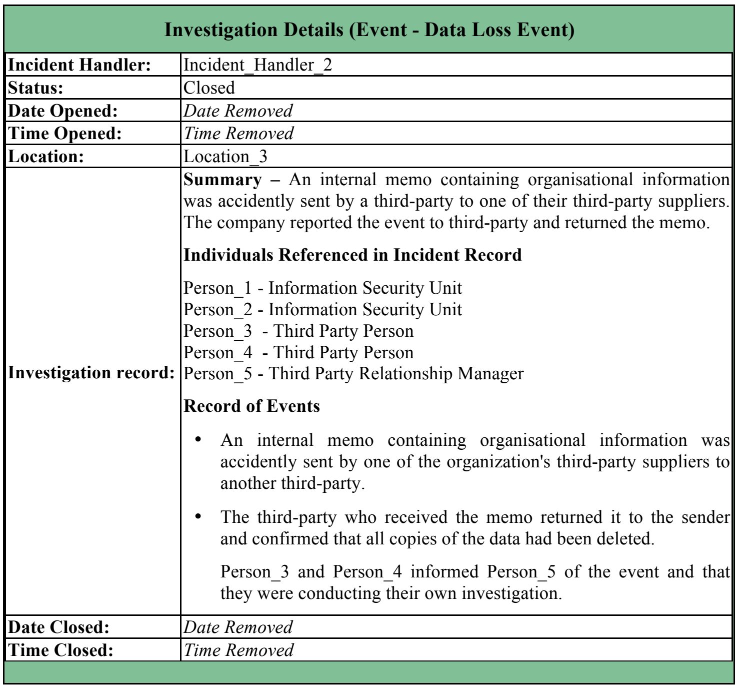

Investigation Details (Event - Data Loss Event)	
Incident Handler:	Incident_Handler_2
Status:	Closed
Date Opened:	*Date Removed*
Time Opened:	*Time Removed*
Location:	Location_3
Investigation record:	**Summary** – An internal memo containing organisational information was accidently sent by a third-party to one of their third-party suppliers. The company reported the event to third-party and returned the memo. **Individuals Referenced in Incident Record** Person_1 - Information Security Unit Person_2 - Information Security Unit Person_3 - Third Party Person Person_4 - Third Party Person Person_5 - Third Party Relationship Manager **Record of Events** • An internal memo containing organisational information was accidently sent by one of the organization's third-party suppliers to another third-party. • The third-party who received the memo returned it to the sender and confirmed that all copies of the data had been deleted. Person_3 and Person_4 informed Person_5 of the event and that they were conducting their own investigation.
Date Closed:	*Date Removed*
Time Closed:	*Time Removed*

Figure 5.7: Miscategorised Investigation Record 5

analysis have shown that the individuals had conflicting views on the classification of the five examples presented to them. As a result, the organisation could still be consuming valuable resources on investigations, which have been classified as a ‘security incident’ by incident handlers, but may not really be considered an actual ‘incident’.

The second theme to emerge from the analysis is that some incident handlers appeared to underestimate the impact of specific security problems. As seen within Example 1, the investigation record stated that sensitive information was sent to a third-party without the use of encryption. One of the interviewed individuals argued that they were confident that the third-party had sufficient security controls in place to prevent additional leakage. As a result, the individual suggested that nothing more could be done with regards to the data leakage and therefore, this was a ‘security event’. However, the other interviewees disagreed. These individuals stated that regardless of the security control situation, data containing ‘legally privileged’ information was transmitted without the use of encryption and therefore this investigation was a ‘security incident’. The conflicting views on estimating the impact of a particular security problem could lie with how the problem is initially identified at the start of an investigation. If the incident handler does not correctly identify the ‘problem statement’ and scope of the actual security problem, then the proceeding investigation could end up being too narrow. In addition, the investigation may not actually uncover the root causes of the actual security problem.

5.4.4 Classification Taxonomy: Problems and Benefits

In addition to asking practitioners to identify why investigation records were assigned a specific category and/or sub-category, participants were also queried on potential benefits and problems with the security incident taxonomy. Four individuals indicated that they did not encounter any problems with the taxonomy. However, two of these individuals noted that a future problem, which could occur, is that a particular investigation may not fit into any of the proposed sub-categories. As a result, these individuals suggested that additional sub-categories could be required in the future. The other two individuals who indicated that they did not experience any problems with the taxonomy stated that they did not foresee any future problems.

Three individuals indicated that they had experienced a problem with the implemented taxonomy. One individual suggested that the definition of a 'security event' be expanded to include investigations where a security problem can exist, but the existing security controls *have* been effective. When queried further, the individual added that they have encountered scenarios where if a security control 'has worked', the security problem is not reported to the team, even though there could be other issues related to the problem to investigate. Another challenge identified by these three individuals was that in some cases, the provided sub-categories did not adequately describe their investigation. As a result, these individuals noted that they had to use a 'best fit' category for these investigations. The individuals proposed that the organisation re-examine the taxonomy on a regular basis, possibly once a year. It is also worth noting that the three individuals who experienced problems also suggested against using mutually exclusive sub-categories, as dictated in the implemented taxonomy. The individuals preferred that multiple sub-categories be used in the taxonomy. As discussed at the end of Section 5.2.1, the use of mutually exclusive sub-categories was a managerial decision made at the start of the experiment. Simon [190] states that the work of a manager includes making decisions for the good of an organisation and communicating these decisions to other individuals within an organisation. However, Simon adds that in order for a manager to make the right decisions, they need to know about the environment in which they work and how their decisions will impact other employees and their work [190]. The incident handlers opinion against using mutually exclusive sub-categories suggests that they were either not consulted or over-ruled by managers with regards to their demands for non-exclusive sub-categories. The use of multiple sub-categories, along with a potentially 'Other type' in the security incident taxonomy, could have allowed for deeper analysis and trend identification.

When the individuals were asked if they had experienced any benefits in using the implemented taxonomy, all seven respondents answered 'Yes'. Five out of the seven respondents indicated that the main benefit derived from using the taxonomy was that it helped the security incident response team and management understand the true extent of the threat landscape within the organisation. The five individuals added that without a well-defined taxonomy it can be difficult to identify the true frequency of particular ‘security events’ and ‘incidents’ within the organisation. One of the five individuals added “*that the addition of the definitions and sub-categories have enhanced both the reliability and consistency of the data generated from the incident response process*”. One individual identified a different benefit. This individual suggested that the definitions and sub-categories removed uncertainty and confusion within the team on how to approach a particular investigation, especially at the start of the organisation’s security incident response process. Furthermore, this individual noted that if anyone was to query why a particular investigation had been identified as either a ‘security event’ or a ‘security incident’, the team could point to a documented definition to justify their selection.

When the participants were asked if the security incident taxonomy assisted or hindered the overall investigation process, all seven individuals agreed that the taxonomy assisted the process. The individuals highlighted numerous benefits including removing ambiguity at the start of an investigation, helping to define clear escalation paths, reducing the time consumed on ‘security event’ investigations and assisting with reporting true ‘security incidents’ to regulators outside the organisation.

5.4.5 Classification Taxonomy Anomalies over Time

Examining the security investigation records, which were ‘miscategorised’ from a chronological perspective, presents an alternative view on the results. Note that the data collected for the analysis of this section for February 2014 and March 2015 does not encompass the whole of these months. This is because the data collection for February 2014 begins when the taxonomy was implemented within the organisation and the data collection ends in March 2015 when the analysis was undertaken. Figure 5.8 presents the results from the chronological analysis. Three observations were made from this analysis.

The first observation from the chronological analysis is the high number of ‘miscategorised’ investigation records at the start of the experiment. The results show that in February 2014, three investigation records contained incorrect categories and four investigation records contained incorrect sub-categories. These results suggest that initially, the security incident response team took some time adjusting to the implemented taxonomy within the organisation. As a result, a high number of ‘miscategorised’ records were documented at the start of the experiment. However, as the experiment progressed the number of ‘miscategorised’ records decreases to about one record per month during the period March 2014 to October 2014. No ‘miscategorised’ records were identified in May 2014.

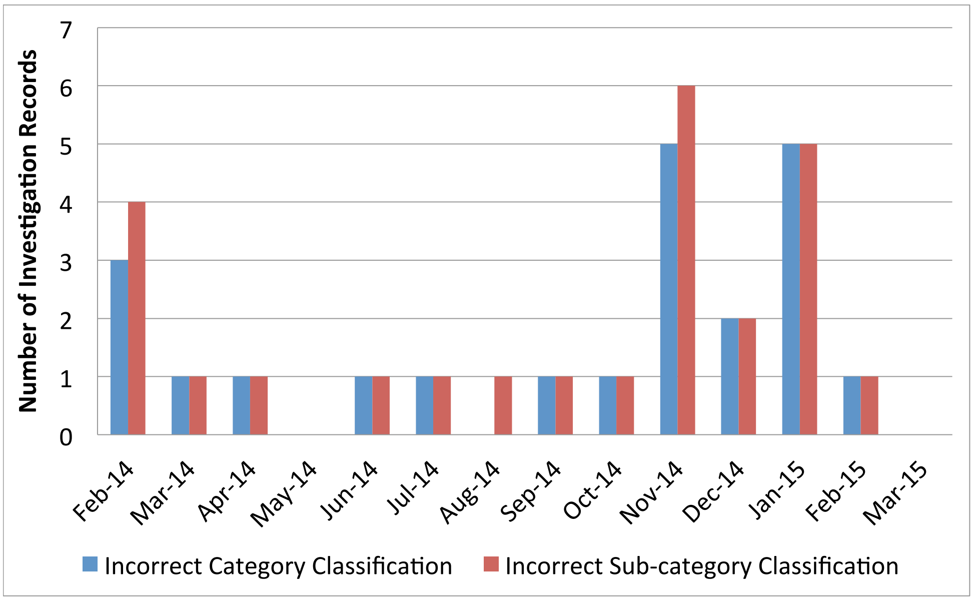

Figure 5.8: Chronological Miscategorised' Investigation Record Analysis

The second observation from the chronological analysis is the increase in the number of 'miscategorised' records during the month of November 2014. A total of five investigation records had an incorrect category and six records had an incorrect sub-category during this month. A possible explanation for this increase is that a new incident handler joined the security incident response team during this period. This incident handler may have taken some time to familiarised themselves with the implemented taxonomy. As a result, this incident handler may have used incorrect categories or sub-categories for investigations, where the distinction was not very clear, or they did not receive enough training with regards to the implemented taxonomy.

The third observation identified from the chronological analysis is an increase in the number of 'miscategorised' records during the month of January 2015. A total of five investigation records with incorrect categories and five records with incorrect sub-categories were identified during this month. There are several possible explanations for this increase including a reduced number of incident handlers, increased work-load during vacation schedules and variations in the experience of newly appointed incident handlers. After January 2015, the number of 'miscategorised' records drops to one investigation record with an incorrect category and one investigation record with an incorrect sub-category, both of which were documented February 2015. No 'miscategorised' records were found in March 2015.

5.4.6 Analysis of Information within Investigation Records

The security investigation record template implemented within the organisation contained 26 fields. This part of the analysis focused on examining which of the 26 fields within the 324

Field Name	Investigation Records With Completed Fields (% of Overall)
Date Reported	324 (100%)
Time Reported	324 (100%)
Reported To	323 (99.7%))
Reported By	324 (100%)
Date of Discovery	324 (100%))
Contact Name	321 (99.1%)
Job Title	314 (96.9%)
Telephone	231 (71.3%)
Department	310 (95.7%)
Business Unit	290 (89.5%)
Line Manager	225 (69.4%)
Incident Handler	324 (100%)
Status	324 (100%)
Date Opened	324 (100%)
Time Opened	324 (100%)
Location	306 (94.4%)
Investigation Record	324 (100%)
Date Closed	322 (99.4%)
Time Closed	322 (99.4%)
Lessons Learned Field 1	288 (88.8%)
Lessons Learned Field 2	286 (88.2%)
Lessons Learned Field 3	269 (83%)
Lessons Learned Field 4	264 (81.4%)
Lessons Learned Field 5	265 (81.7%)
Actions To Be Taken	324 (100%)
Hours Working	324 (100%)

Table 5.4: Analysis of Completed Fields within Investigation Records

investigation records were used by the security incident response team and what information was recorded within these fields. This analysis involved two steps. First, the investigation records were examined to determine whether the fields within each record contained information. Second, the fields were then examined to identify the extent to which the information within the 26 fields was considered to be ‘actionable information’. Actionable information is “information (that can be) used to take actions that mitigate against future threats, or help address existing compromises” [188]. An example of ‘actionable information’ in the ‘Location’ field is the name of the business unit or team within the organisation affected by a particular ‘security event’ or ‘incident’. In the same field, an example of ‘non-actionable’ information would be either a blank entry in the field or any other information, which does not describe a location within the organisation. Table 5.4 presents the results of the analysis which examined how much information was documented in the 324 investigation records.

The results of this analysis show that 11 out of the 26 fields were completed in all 324 investigation records. Furthermore, four fields were found to have been completed in over 99% of the investigation records. These four fields were the ‘Contact Name’, ‘Date Closed’, ‘Time Closed’ and ‘Reported To’ fields. In the five ‘Lessons Learned’ fields, information was documented in between 81% and 88% of the analysed investigation records. In contrast, information from six fields was found to be missing in ten or more of the analysed investigation records. The results show that the ‘Telephone’ and ‘Line Manager’ fields were the least documented fields within the investigation record. The ‘Telephone’ field was not documented in 93 records, while the ‘Line Manager’ field was not documented in 99 investigation records. Further analysis showed that five out of the six fields, which were missing information in ten or more fields within the record template, were those in the ‘Reporting and Contact Information’ section. These fields are used to document the name and contact information of the individual who has reported a security problem to the security incident response team.

The results presented above have shown that information was documented in the majority of the fields in the investigation record template. However, the initial analysis has identified that some of this information was ‘non-actionable’. The next step in the analysis was to examine the information in the actual fields to determine how much of the information can be considered to be ‘actionable’ information. This involved separating the 324 investigation records into ‘event investigations’ (255) and ‘incident investigations’ (69) depending on the sub-category assigned by the incident handlers. Table 5.5 presents the results of this analysis.

Several observations can be drawn from the analysis of the ‘actionable’ information within the investigation records. The results show that ten fields within the investigation records labelled as ‘security incidents’ contained more ‘actionable’ information than those records labelled as ‘events’. This is particularly evident in the five ‘Lessons Learned’ fields. While 85.5% of the ‘security incident’ records contained ‘actionable’ information from Lessons Learned field 1, only 43.1% of the ‘security event’ records contained ‘actionable’ information within this field. Similarly, 55% of ‘security incident’ records contained ‘actionable’ information from Lessons Learned field 3, only 35.2% of the ‘security event’ records contained ‘actionable’ information within this field. Furthermore, 100% of the ‘security incident’ records contained ‘actionable’ information in the ‘Location’ field, but 92.9% of the ‘security event’ records contained ‘actionable’ information in this field.

Five fields within the ‘security event’ investigation records contained more ‘actionable’ information than the same fields in the ‘security incident’ records. However, in all five cases the difference in ‘actionable’ information between the two types of investigation records was minimal. For example, while 99.6% of the ‘security event’ records contained ‘actionable’ information in the ‘Contact Name’ field, while 97.1% of the ‘incident records’ contained ‘actionable’ information in this field. Likewise, in the ‘Date Closed’ and ‘Time Closed’ fields,

Field Name	Event Investigations With Actionable Info. (% of Type Total)	Incident Investigations With Actionable Info. (% of Type Total)
Date Reported	255 (100%)	69 (100%)
Time Reported	255 (100%)	69 (100%)
Reported To	254 (99.6%)	69 (100%)
Reported By	255 (100%)	69 (100%)
Date of Discovery	255 (100%)	69 (100%)
Contact Name	254 (99.6%)	67 (97.1%)
Job Title	252 (98.8%)	62 (89.8%)
Telephone	181 (70.9%)	50 (72.4%)
Department	254 (99.6%)	65 (94.2%)
Business Unit	235 (92.1%)	64 (92.7%)
Line Manager	169 (66.2%)	56 (81.1%)
Incident Handler	255 (100%)	69 (100%)
Status	255 (100%)	69 (100%)
Date Opened	255 (100%)	69 (100%)
Time Opened	255 (100%)	69 (100%)
Location	237 (92.9%)	69 (100%)
Investigation Record	255 (100%)	69 (100%)
Date Closed	254 (99.6%)	68 (98.5%)
Time Closed	254 (99.6%)	68 (98.5%)
Lessons Learned Field 1	110 (43.1%)	59 (85.5%)
Lessons Learned Field 2	110 (43.1%)	56 (81.1%)
Lessons Learned Field 3	90 (35.2%)	38 (55%)
Lessons Learned Field 4	5 (1.9%)	6 (8.6%)
Lessons Learned Field 5	3 (1.1%)	8 (11.5%)
Actions To Be Taken	255 (100%)	69 (100%)
Hours Working	255 (100%)	69 (100%)

Table 5.5: Analysis of Actionable Information within Investigation Records

'actionable' information was recorded in 99.6% of the 'security event' records and 98.5% of the 'security incident' records. In summary, the results have shown that the security incident response team are more likely to document 'actionable' information for a 'security incident' over a 'security event' investigation.

5.4.7 Incomplete Investigation Records over Time

The results from the analysis discussed in Section 5.4.6 have shown that a number of investigation records contained fields where 'non-actionable' information was documented. This section of the analysis focused on examining these investigation records in order to determine when they occurred in the context of the experiment. 149 out of the 324 (46%) investigation records were found to have one or more fields containing ‘non-actionable’ information. Figure 5.9 provides a breakdown of the number of investigation records containing ‘non-actionable’ information against the total number of records created for each month of the experiment.

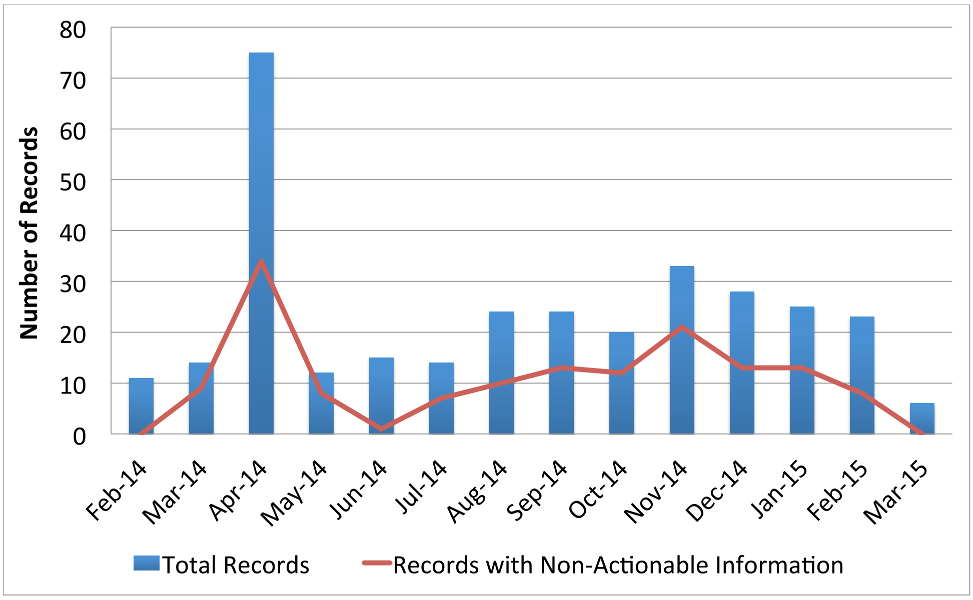

Figure 5.9: Non-Actionable Information Investigation Record Analysis

The figure shows that highest number of investigation records with ‘non-actionable information’ was observed during April 2014. 34 such records were identified during this period. Furthermore, the investigations records documented during April 2014 included one very large ‘multi-record’ investigation. As result, the number of investigation records with ‘non-actionable information’ could be much lower if the ‘multi-record’ investigation is excluded from the analysis. If the results for April 2014 are excluded, then the month with the highest number of investigation records with ‘non-actionable information’ is November 2014. 21 records contained ‘non-actionable information’ during this month.

The results also show that May 2014 has the highest percentage of investigation records with ‘non-actionable information’. 8 (67%) out of the 12 investigation records in this month, were identified as containing ‘non-actionable information’ in one or more fields. Data was not collected during the entire months of February 2014 (start of experiment) and March 2015 (end of experiment). Therefore, if the data for February 2014 and March 2015 are excluded from the analysis, June 2014 is considered the month with the lowest number of investigation records containing ‘non-actionable information’. Only one out of the fifteen (7%) records created during this period was found to contain ‘non-actionable information’. After this period, the number of investigation records with ‘non-actionable information’ increases each month until November 2014. After November 2014, the number of investigation records with ‘non-actionable information’ begins to decrease until the end of the experiment.

The 149 incomplete records were then analysed to determine which fields were the most prevalent with regards to ‘non-actionable information’. The results from this analysis showed that 140 out of the 149 records contained ‘non-actionable information’ in one or more fields from the ‘Reporting and Contact Information’ section of the investigation record template. This section consisted of the ‘Contact Name’, ‘Job Title’, ‘Telephone’, ‘Department’, ‘Business Unit’, and ‘Line Manager’ fields. The follow-up interviews with practitioners were then used to investigate why these particular fields contained ‘non-actionable information’.

5.4.8 Investigation Record Template Participant Analysis

Initial questions about the investigation record template focused on the participant’s opinion on whether the revised template implemented within the organisation captures all relevant information about a ‘security event’ and a ‘security incident’. All seven participants were unanimous in that the investigation record template captured all relevant information about a ‘security event’. However, only six out of the seven individuals agreed that the revised template captured all relevant information about a ‘security incident’. The one individual who disagreed noted that they would like to see a field to record information related to the remedy being applied to prevent a recurrence of the problem, as well as if the remedy has been successful or unsuccessful. Another field the individual suggested which they would like to see implemented was a drop-down menu with a list of security controls, which have worked/failed depending on the particular investigation. The first recommendation proposed by this individual supports the findings from the literature that security incident response teams (and by extension organisations) are more focused on eradication and recovery when it comes to security investigations [18, 44]. While this was not examined further, the proposal from the individual suggests that there is likely pressure from management within the organisation on incident handlers to demonstrate that their investigations are conclusive and that any problems will not reoccur. Hence, the proposal of an addition field to the investigation record is likely to satisfy demands from management about statistics regarding incident eradication and recovery.

However, the second recommendation from the individual provides a different view to the above assumption. What this request suggests is that potential discussions are taking place in the security incident response team about which security controls are working/failing within the organisation. For this type of information to be identified and documented, incident handlers would need to conduct in-depth investigations, and not just focus on eradication and recovery [15]. Alternatively, like the first suggestion, the request for the security control drop-down list could also be to satisfy management demands regarding security control implementation.

5.4.9 Investigation Record Template: Problems and Benefits

Subsequent interview questions were used to query individuals about any problems or benefits experienced using the revised investigation record template. The participant's answers suggested that they encountered some problems with documenting and completing various fields within the investigation record template. These fields include the 'Date Reported' field, the 'Lessons Learned' fields, the 'Contact Information' section fields, the 'Date of Discovery' field and the 'How Many Hours Working' field.

Four individuals indicated that they had encountered problems completing the 'Lessons Learned' fields within the investigation record template. These four individuals stated that they found it difficult to identify and document security controls to prevent a future recurrence of the problem. The individuals attributed this problem to a lack of awareness surrounding the security controls actually implemented within the organisation. The individuals argued that without knowing enough about security controls actually implemented within the organisation, it can be difficult to suggest improvements. Currently, security analysts in other parts of the Information Security unit will identify and implement security controls based on requests from information security managers. This is done because of a segregation of duties policy, which exists within the organisation. One individual suggested that either the security incident handlers are provided with additional training regarding what security controls are implemented within the organisation or that the security incident response team includes an individual who has knowledge of these controls. The proposal is that this individual would help the team to identify what controls are currently implemented and what recommendations can be made to management to strengthen existing controls. This problem provides an explanation as to why very little 'actionable information' was documented in the 'Lessons Learned' fields. Segregation of duties is a classic security method to manage conflicts of interest, fraud and the amount of power held by any one individual [54]. In the security incident response context, it is there to prevent individuals from *not* applying security controls which are required to prevent a recurrence of a security incident. However, the individual's answers have suggested that this segregation of duties policy has impacted the completeness of information in the security investigation record. In this case the organisation needs to trade-off the segregation of duties policy with the data capture requirement for the security incident response team. An alternative would be to remove the requirement for the security incident response team to document security control failure, and prompt the individuals implementing the control to gather information about which control failed and what modifications were required to improve it. This way, the information captured about security control failure and the segregation of duties policy is both sustained.

Three individuals stated that they had encountered problems with identifying and documenting contact information about the employee who has reported a security problem to the security incident response team. These individuals noted that because the reporting of security problems within the organisation can come from various sources (including through word of mouth), often the incident response team will only have the employee's forename and surname available. When this information is used to lookup the employee in the organisation's electronic address book, a common problem that is encountered is that many fields within the address book are either missing or contain incomplete information. This finding shows that the quality of data in the organisation's address book directly impacts the security incident response team's ability to document the information they require for their process.

Two individuals answered that they had encountered difficulties in identifying information for the 'Date of Discovery' and 'Date Reported' fields within the investigation record template. When asked to elaborate further, these individuals stated that employees are sometimes unaware of how or when to report a security problem to the incident response team. One of these participants added that in their experience, employees can often take a long time to report a security problem. As a result, when the report is eventually made to the security incident response team and employees are queried on the 'Date of Discovery', they can often reply that they are unsure or cannot remember the exact date when the problem was first discovered.

Another difficulty experienced by two individuals concerns the identification of information within the 'Date Reported' field. The individuals argued that because security problems are often reported via word of mouth, this can usually involve as many as three or four individuals being notified of the problem before it actually reaches the security incident response team. As a result, when attempts are made to obtain the true 'Date Reported' information, the team finds it difficult to track all the individuals who the problem was reported to, in an effort to obtain the actual date. Instead, the team has to document the date that the team receives the actual report. Interviewees suggested that the information within the 'Date Reported' field could be more accurate if employees are made aware to report security problems directly to the incident response team. Although 100% 'actionable information' was identified in both the 'Date of Discovery' and 'Date Reported' fields, the answers from the individuals above suggest that some of the dates stored in these fields do not report the true dates that they are supposed to represent.

Two individuals stated that they had problems completing the 'How Many Hours Working' field within the record template. The problem raised by both individuals is that when multiple incident handlers are working on the same investigation it can become difficult to calculate the number of hours each incident handler has contributed to the investigation. Hence, these individuals revealed that in such scenarios, it could be difficult to identify the exact number of hours that both incident handlers consumed on the investigation.

All seven participants agreed that they had experienced a benefit in using the revised investigation record template. The main benefit identified by the majority of the interviewees was that when an investigation record is fully complete, it can provide a comprehensive picture of a specific security problem resolved by the team. The interviewees added that a complete investigation record can also provide a rich source of information for the Information Security unit's intelligence life-cycle. This is a process where the Fortune 500 Organisation collects and processes security information from various sources and produces intelligence for its management to make strategic information security decisions. One individual added that more information was now being captured using the revised investigation record template, when compared to the previous investigation record template. This individual went on to state that this enhanced information allowed the security incident response team to identify trends and themes surrounding the threat landscape, as well as which parts of the organisation are affected the most by particular security problems. All seven individuals answered that the revised security investigation record template assisted the overall investigation process. The main reason provided is that when compared to the original template, the revised investigation record template captures more relevant information using a standardised approach.

While all seven interviewees appeared to suggest that they had experienced some benefit with using the revised investigation record template, concerns were also raised that incident handlers encountered difficulties in capturing specific data for the investigation record. However, an observation from the individual's responses is that the data capture problems were nearly always caused by someone else's (lack of) actions within the organisation. This indicator suggests that there is a blame culture within the organisation [191]. This is a problem which has been identified in the financial services industry in previous case studies [192,193] and will be discussed in more detail in Section 5.5.2.

5.4.10 Security Incident Response Metric Analysis

In Section 4.3.3 of the exploratory case study, two metrics were calculated the *response time* and the *total time to resolve*. Response time is defined as the period of time from the first report of a security problem to the implementation of the first mitigating actions [108]. The total time to resolve is defined as the time from when the security problem was reported to when the investigation was closed [108]. A limitation identified in the exploratory case study was that the calculated metrics only incorporated a small proportion of the investigation records in the organisation's security incident response database. For the response time analysis, 28% of the investigation records contained information for this calculation, while 36% of the investigation records contained information for the total time to resolution calculation. As a result of the small number of investigation records included in the metric calculations, there is the possibility that the results from the exploratory case study do not provide an accurate representation of the organisation's security incident response performance. Therefore, in order to investigate this claim, this section will revisit the above metric calculations and present the results of new metric calculations using data collected from the experiment (the post-implementation analysis). The purpose of this new analysis was to investigate whether the new metrics calculated, using a more complete dataset of security investigation records, presented a more accurate representation of the organisation's security incident response performance.

Response Time Analysis

To calculate the response time metric in the post-implementation analysis, information was used from the 'Date Reported', 'Time Reported', 'Date Opened' and 'Time Opened' fields. Although 'actionable information' was documented in all four fields within the 324 investigation records, the information within 21 investigation records was disregarded. This was because the response time values calculated from these 21 investigation records were a negative number. This means that the information recorded in the date and time opened fields was *before* the date and time the problem was actually reported. Therefore, information from 303 (94%) out of the 324 records was used in the calculation of this metric.

The results from the post-implementation response time calculation show that the minimum response time was zero minutes and the maximum response time was 39,840 minutes (27 days). Therefore, the mean average response time calculated was 2,516 minutes (1.74 days). The results also show that using the documented investigation record information, the security incident response team took mitigating actions within 60 minutes for 37% of the security investigations and within 8520 minutes (5.91 days) for about 90% of the investigations. Figure 5.10 compares the distribution of response times for the investigation records used in the post-implementation calculation with the distribution of response times from the exploratory case study calculation.

Figure 5.10 shows that in comparison with the response time calculations from the exploratory case study, the total duration of the longest response time is several times longer in the post-implementation calculation. The longest response time in the exploratory case study was 325 minutes, while the longest response time in the post-implementation calculation was 39,840 minutes. Only 28% of the investigation records in the post-implementation calculation were responded to within 325 minutes. While these results suggest a degrading security incident response team performance, the results from the post-implementation calculation could also provide a more accurate indication of the team's true performance. While 52 (of 188, 27.6%) investigation records were used in the exploratory case study calculation, 303 (of 324, 93.5%) investigation records were used in the post-implementation calculation. With more data being available in the post-implementation calculation, it is plausible that

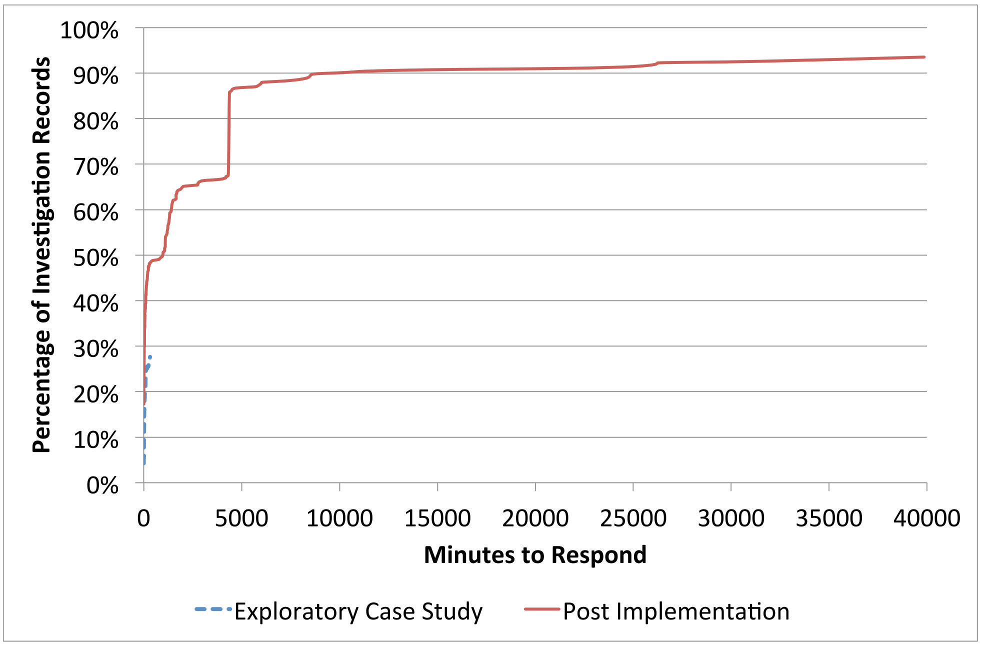

Figure 5.10: Response Time Analysis for 303 records

the new results from the response time analysis provide a better reflection of the security incident response team's 'response time' performance.

In order to further investigate the plausibility of the new data, the analysis then focused on the number of investigation records responded to between 1 and 325 minutes for both the exploratory case study and post-implementation calculations. This was the maximum time for response during the exploratory case study calculations. This was done to obtain a better understanding of the differences between the two metric calculations with regards to the response times and plausibility of the new results. The post-implementation analysis revealed that 18% of the analysed investigation records had a response time of 'zero' minutes, which is considered implausible within the organisation. Therefore, this analysis excluded the investigation records where the response time was zero minutes. Figure 5.11 presents the results of this analysis.

The figure shows that up to about 100 minutes, the two curves are very similar. In fact, it can be argued that the shape of the two curves shows that the enhanced investigation record template implemented within this experiment has not significantly altered the data recorded for the analysis of this metric between 0-5 hours. Given that response times for greater than 325 minutes in the post-implementation investigations follow a similar trend, this suggests that this portion of the data can be considered reliable.

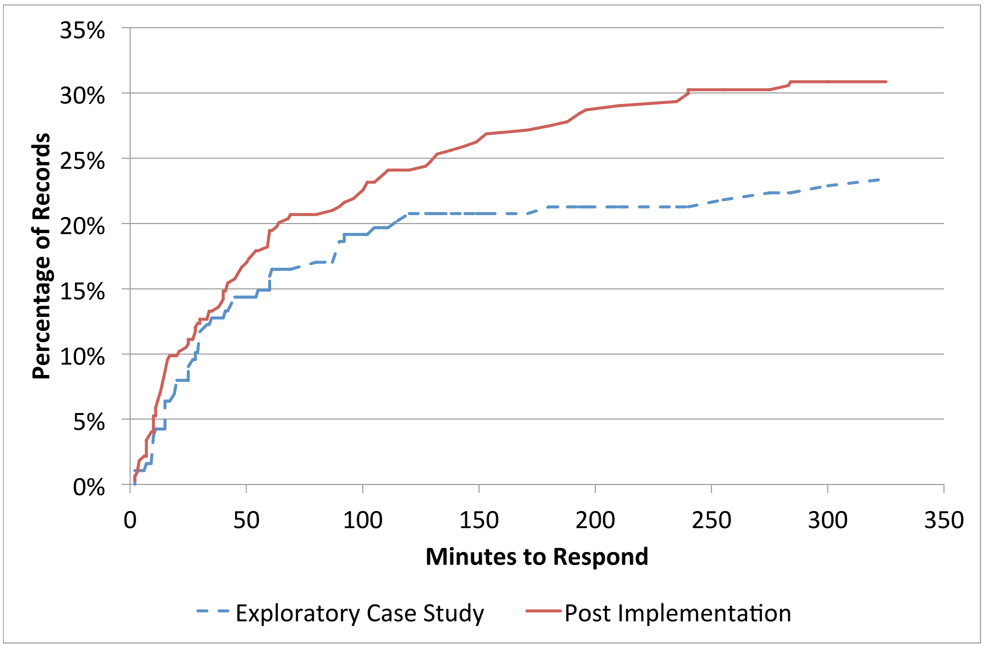

Figure 5.11: Response Time Analysis Within 325 Minutes

Total Time to Resolve Analysis

In the post-implementation analysis, the Total Time to Resolve (TTR) metric was calculated using information recorded in the ‘Date and Time Reported’ fields and the ‘Date and Time Closed’ fields. Information from 317 (98%) out of the 324 records was used in this metric calculation. Seven records were disregarded because the TTR calculated from these records was a negative value. This means that the information recorded in the ‘Date and Time Reported’ fields was *after* the information recorded in the ‘Date and Time Closed’ fields. The identification of these seven records suggested that any metric calculated using the information in the remaining 317 records should be treated with caution, as data quality issues could exist within these investigation records.

The analysis of the 317 records showed that the minimum TTR was half a day and the maximum TTR was 208 days. Therefore, the mean average TTR was 23.41 days. A comparative analysis was then undertaken to compare the TTR results from the exploratory case study calculation, with the TTR results from the post-implementation calculation. Figure 5.12 presents the results from this comparative analysis.

The results of the comparative analysis showed that investigations analysed in the post-implementation calculations were taking longer to resolve when compared to investigations analysed in the exploratory case study. In the exploratory case study, 60% of investigations were resolved within five days, while in the post-implementation calculation only 35% of investigations were resolved in five days. Similarly, in the exploratory case study calculation, 81% of investigations were resolved within 20 days but in the post-implementation

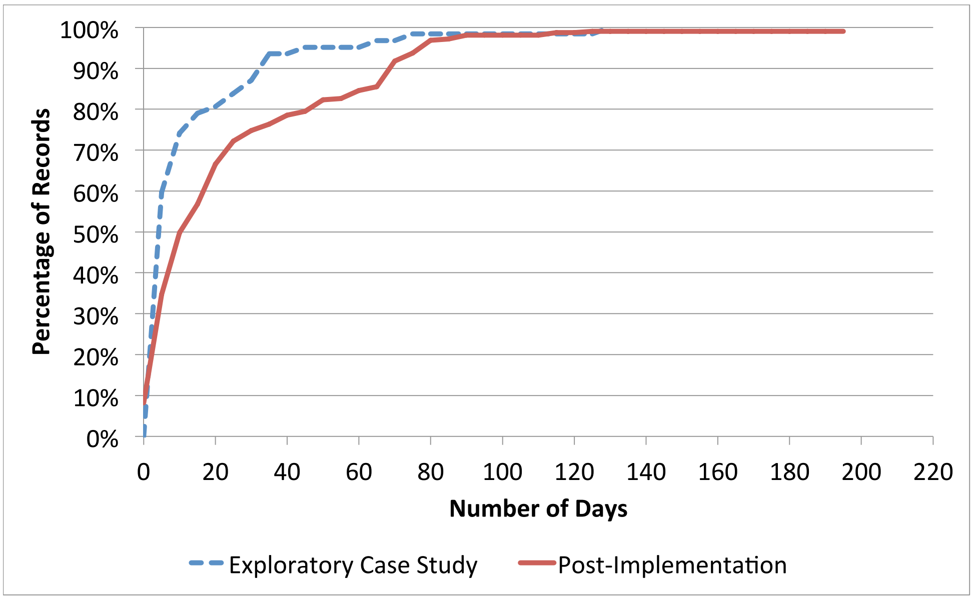

Figure 5.12: Total Time to Resolve Comparison Analysis

calculation, only 67% of investigations were resolved in 20 days.

Similar to the response time analysis, an extended analysis was also undertaken examining the TTR calculations from the exploratory case study and post-implementation calculations. In the exploratory case study, the results showed that 100% of the investigations were resolved in 130 days. Therefore, the extended TTR analysis focused on the number of investigations from the post-implementation calculation, which were resolved in 130 days. The results of this analysis were then compared to the findings from the exploratory case study calculations and the results are presented in Figure 5.13.

In summary, the results suggest that the security incident response team is now taking longer to resolve security investigations in the post-implementation calculation. However, one possible explanation to this increase in resolution time is that the security incident response team is now focusing more attention on enhancing data capture. In particular, more attention could now be paid to recording information about the timing of an investigation more accurately. Therefore, while these investigations appear to be taking longer to 'resolve', more detailed information is now being captured, which shows a clearer picture of the resolution times of investigations within the organisation.

Methods of Detection Analysis

As noted in Section 5.4.8, practitioners were queried about the potential benefit of using the revised security investigation template implemented in the organisation. Individuals who were interviewed stated that one of the benefits from the revised investigation record tem-

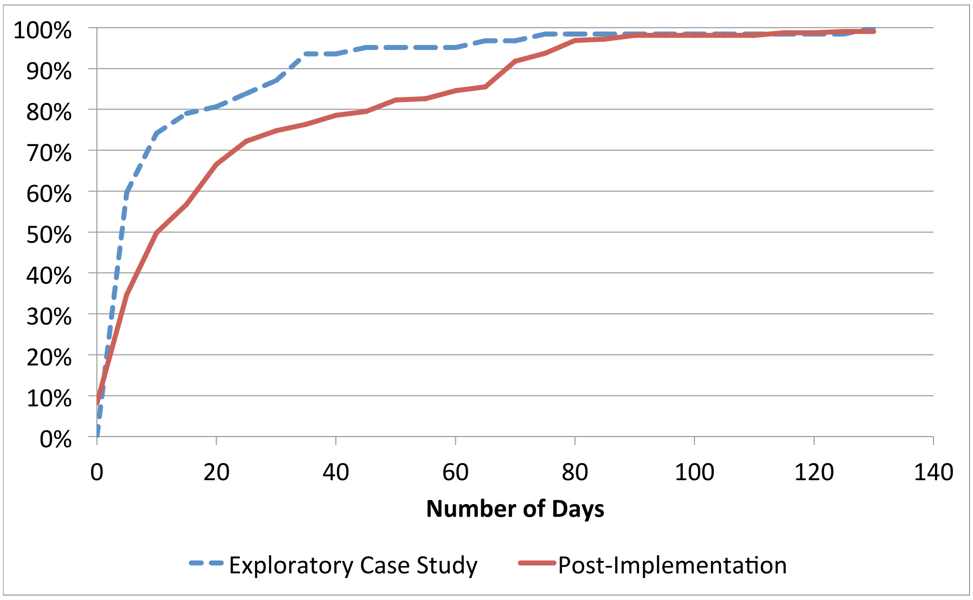

Figure 5.13: Total Time to Resolve Analysis over 140 days

plate was that it provided a better source of information for the organisation's information security intelligence life-cycle. An example provided by one of the individuals was that the investigation records now allow the security incident response team to determine the methods used to detect/report security problems within the organisation.

Industrial surveys [41, 43, 194] have suggested that given the workload of an average security incident response team, a key metric for an organisation is to determine the methods of detection and reporting within security incident response. The purpose behind this metric calculation is to determine the level of automation in the identification and detection of security problems and how many of these problems require manual discovery or reporting [194]. The calculation of this metric was not possible in the exploratory case study because the organisation was not capturing the required information. However, the metric can now be calculated using information from the post-implementation data, which has been recorded in the 'Reported By' field. 'Actionable information' within this field was available in all 324 investigation records. Table 5.6 presents an overview of the results. The table shows that five main methods of identification/detection/discovery/reporting were found within the 324 investigation records.

The table shows that 162 security problems were brought to the attention of the security incident response team through direct employee reporting to the team. Internal notifications and third-party notifications to the team accounted for ten security problems. This means that 172 out of the 324 (53%) investigation records were as a result of 'manual reporting'. These forms of reporting included via email, verbally or through the telephone directly to the team. Information from 152 out of the 324 (47%) investigation records showed that some level of

Detection/Identification/Reporting Method	**Number of Investigations**
Detected by Data Loss Prevention System	54
Detected by Email Content Filter System	98
Employees Reporting to SIRT	162
Internal Notifications to SIRT	5
Third-party Notifications to SIRT	5
Total	**324**

Table 5.6: Total Time to Resolve Comparison Analysis

automated detection/reporting does exist within the organisation. 98 security problems were identified to the team via the organisation's email content filter and 54 security problems were identified to the team by the data loss prevention system.

5.5 Discussion

The results from the data analysis have raised several points for discussion with regards to using a revised investigation record template to collect data of value to a security incident response team and removing classification ambiguity using a well-defined security incident taxonomy. These points of discussion include the organisational factors influencing data capture, the importance of security education within security incident response teams, consolidated data quality improvement, incentives for security incident reporting, and difficulties in classifying investigations upfront.

5.5.1 Organisational Factors Influencing Data Capture

The data analysed from the experiment has revealed several organisational factors that has influenced the capture of specific data in the investigation record, as well as the classification of security events and incidents. The consequences of these factors are particularly evident in the amount of none-actionable information in the Lessons Learned field, as well as the answers provided by incident handler's, who stated that their failure to capture security control information was because it is 'someone else's job'. These observations from the experiment results suggests the prominence of a 'blame culture' within the organisation.

According to Khatri, et al. a blame culture is a "set of norms and attitudes within an organisation characterised by an unwillingness to take risks or accept responsibility for mistakes because of a fear of criticism or management admonishment" [191]. Khatri, et al. go on to state that "an organisation does not purposefully choose a blame culture, but rather, such a culture evolves out of a bureaucratic management style that is highly rule-oriented, compliance-driven, and focused on assigning blame or accountability to individuals even for system-level failures" [191]. A blame culture will force individuals to protect themselves by shifting blame and more importantly, hindering continuous improvement [191, 193]. While these findings suggest that a blame culture could be one reason why information was not correctly captured, it would also be appropriate to examine how the 'culture of information security teams' relates to the findings.

While the fear of blame is certainly a barrier as to why certain information was not captured, there could also exist other deep-seated socio-cultural problems within the security incident response team [195]. These can include security incident handlers not understanding their responsibilities, relationship to the wider role of information security within the organisation, and generally not understanding their task and function [196]. It is also worth remembering that the requirements regarding information which was not captured or investigations which were not correctly classified, was all documented in the security incident response policy and related documentation. Hence, managers also need to be held responsible for the failure in compliance of the policies as they are traditionally responsible for the implementation of the policy in their teams and units.

5.5.2 Security Education within Incident Response Teams

As discussed in Chapter three, several security incident response approaches have been published providing guidance to security incident response teams [7–9, 12, 13, 109, 185]. Some of these approaches highlight that extended incident response preparation and training can provide new security incident handlers with the necessary skills to undertake their work; broaden the abilities of existing incident handlers and generally ensure that a security incident response team's skill set is up-to-date with emerging threats, trends and technologies [108]. This means that security education within security incident response teams can cover an array of topics including knowledge transfer, other security processes, every day activities of organisational units who collaborate with the team, as well as the results of potential risk assessments.

The importance of security education and training to enhance the quality of data in security investigation records has been highlighted at several points in the experiment. First, during the analysis of the investigation records examining the information documented in the 'Lessons Learned' fields, it was identified that much of the information recorded was 'non-actionable'. When incident handlers were queried about this, they argued that they had very little knowledge about the security controls actually implemented within the organisation and therefore, found it difficult to suggest improvements. In addition, when incident handlers were shown five examples from the 25 records, which were considered to be 'miscategorised' during the follow-up interviews, the results suggest that there were still varied translations of the term ‘security event’ and ‘security incident’ within the organisation.

A potential explanation points to a lack security education and knowledge transfer. The definitions and categories were developed using input from one of the organisation’s information security managers. It was anticipated that the manager would transfer any knowledge with regards to the definitions, the sub-categories and their translations directly to the security incident response team. The results from the experiment suggest that this knowledge transfer may not have completely taken place. Evidence for this suggestion comes from the follow-up interviews, where the incident handlers provided varied interpretations of terms ‘security event’ and ‘security incident’, even though these terms were defined in the documented process within the organisation.

Another example where enhanced security education could have benefited the quality of data generated using the revised investigation template was during the identification of security controls. The analysis of the investigation records showed that only eleven records contained information about security controls to prevent recurrence. Within the follow-up interviews, practitioners indicated that one of the reasons why this field was not completed was because incident handlers did not know enough about the security controls implemented within the organisation. As a result, these individuals argue that before they can propose security control improvements, they need to know more about what is currently implemented within the organisation. This again points to a wider security education issue. Although many of the incident handlers identify themselves as ‘security analysts’ or ‘senior security analysts’, specific interview questions established that their primary job role is within security incident response. There are however, specific individuals within the organisation whose job role includes the implementation and review of security controls. Therefore, there is a need to transfer knowledge from these individuals to the security incident response team. This need has only emerged as a result of reviewing and improving the quality of data captured with regards to security control implementation. There are two possible solutions. Either the security incident response team receive training in the security controls implemented in the organisation, or an individual who is aware of the security controls implemented within the organisation is integrated into the security incident response team.

5.5.3 Consolidated Data Quality Improvement Initiatives

The data analysis of the security investigation records showed that 140 (43%) out of the 324 investigation records were missing information from the ‘Reporting and Contact information’ section of the investigation record. The purpose of these fields is to document the name and contact information of the individual who has reported a security problem to the security incident response team. Three individuals interviewed during the experiment stated that they had encountered difficulties in actually obtaining contact information about the employee who had reported a security problem to the security incident response team. This is because when this information is used to lookup the employee in the organisation's electronic address book, a common problem that is encountered is that many fields within the address book are either missing or contain incomplete information.

While the security incident response team are not the asset owners of the organisation's electronic address book, their use of the information in this asset directly impacts the quality of data in their own security investigation records. The findings from the investigation records analysis and interviews suggest that other organisational processes and data sources can impact the security incident response team. However, research in this thesis specifically focused on improving the quality of data generation within security incident response and data quality problems outside this domain fall outside the scope of the thesis. Therefore, a suggestion which arises from this discussion is that organisations should look to address data quality issues in all processes which input data either directly or indirectly into a security incident response process and not just those which are used by the team during their investigations.

5.5.4 Lack of Incentives for Incident Reporting

Two participants from the follow-up interviews noted that they had encountered difficulties in identifying the 'Date of Discovery' for particular security investigations. The participants explained that this was a problem because employees who reported security problems, often found it difficult to recall when the problem was first discovered. The participants suggested that employees may not recognise that a particular security problem could have regulatory implications for the organisation and therefore, may not prioritise its reporting. In addition, if a particular security problem has been on-going for a period of time it may be difficult for employees to recall the exact date when the issue was first discovered.

Jaatun, et al. [122] argue that employees within organisations need to be aware of their responsibilities to send alerts about potential security incidents to relevant teams as soon as possible. However, as the results from the experiment show, in reality this can become a challenge for security incident response teams. One solution to this problem could be the introduction of an incentive to report security problems in a timely manner. This incentive can be either positive (e.g. an increased financial bonus) or negative (e.g. a decreased financial bonus) depending on the enforcement of the incentive [197].

5.5.5 Difficulties Affecting Upfront Investigation Classification

The results from the analysis of the 'miscategorised' investigation records have highlighted the complexity of security incident response when incident handlers attempt to classify a security investigation upfront. During the follow-up interviews, when incident handlers were shown examples of ‘miscategorised’ investigation records, they requested more information from the investigation record before they could assign a classification. Within the first example of a ‘miscategorised’ investigation record (Example 1), one of the incident handlers answered that they would have used a different classification if the third-party involved had sufficient controls in place to prevent additional data leakage. Likewise in Example 4, one of the incident handlers stated that they would have requested further information about who should have access to the information in the parent organisation. In both these examples, the requested information would have required some form of investigation to take place. Without this information, incident handlers have argued that they would have found it difficult to classify these investigations upfront, which is what is required in a typical security incident response process.

5.6 Summary

This chapter presented an experiment to evaluate the use of a well-defined security incident taxonomy and revised security investigation record template within the Fortune 500 Organisation. The purpose of the experiment was to evaluate if the taxonomy and the investigation record template help to collect data, which has value to the security incident response team within the organisation.

The results from the experiment have shown that the introduction of a well-defined security incident taxonomy can help reduce any investigation classification uncertainty at the start of an investigation. However, the analysis of the investigation records showed that some uncertainty still exists in the organisation with regards to investigation classification. 25 out of the 324 analysed investigation records from the experiment were ‘miscategorised’ by the security incident response team. When the issue was discussed during the follow-up interviews, security incident handlers revealed that various translations of the terms ‘security event’ and ‘security incident’ still exist within the organisation. One potential explanation as to why this has occurred is because the information security manager who helped to define the terms did not transfer their knowledge across to the security incident response team who use the definitions on a daily basis.

In addition to evaluating a well-defined security incident taxonomy, the experiment also evaluated a revised security investigation record. The results from this part of the experiment are encouraging. The results show that the addition of the revised investigation record template has enhanced the overall investigation process and incident handlers are now documenting more detailed information about a security investigation. Practitioners have suggested that the revised investigation record template also provides a more comprehensive picture of the investigation undertaken, when compared to the information from the previous investigation record template.

However, the results from the experiment have also identified that even when a security incident response team is given the correct tools to collect enhanced information for incident learning, other factors can prevent this from occurring. It is only through the process of improving the investigation record template and enhancing the information captured during an investigation that this has become apparent. The results from the analysis of the ‘actionable information’ within specific fields revealed that contact information about employees who have reported or detected a security problem was not well documented. When queried about this problem in the follow-up interviews, incident handlers argued that this was because there was a data quality problem with the organisation’s electronic address book. This is used by the team to locate employee information when they report a security problem. The finding suggested that organisations that are attempting to improve data quality within security incident response, need to also address data quality issues with other processes which can input data into security incident response.

Chapter 6

Using Retrospectives to Validate and Enhance Security Investigations

Chapter three established that many security incident response approaches include a ‘follow-up’ phase, where an organisation can reflect and learn from a security investigation. However, little research has been conducted to examine tools and techniques for this purpose. Chapter four presented evidence that one of the areas where the Fortune 500 Organisation’s security incident response process could be improved was within its ‘post-incident’ phase. One of these opportunities was the development of lessons learned and the capture of additional information to help with security incident learning. This chapter presents an experiment where a lightweight measure, *the retrospective*, was implemented within the Fortune 500 Organisation’s security incident response process. The purpose of the experiment was twofold. First, the experiment was used to evaluate if a retrospective can help a security incident response team validate information that has been collected from a security investigation. Second, the retrospective was evaluated to determine if it can be used to collect additional information that may have been missed during an investigation without requiring substantial extra resources. It is anticipated that the information collected from the retrospective would help enhance security incident learning within an organisation.

This chapter is structured as follows. Section 6.1 introduces retrospectives and discusses how they have been used within software development. Section 6.2 describes how retrospectives and a ‘retrospective of retrospectives’ (*meta-retrospective*) were adapted for use within the organisation’s security incident response team. Section 6.3 outlines the experiment design and defines the experiment’s research questions. Sections 6.4 and 6.5 present the data collected from the retrospectives and meta-retrospectives respectively which were implemented in the Fortune 500 Organisation. Section 6.6 examines the impact of the retrospectives through a qualitative and quantitative data collection. Section 6.7 discusses the main findings from the experiment and Section 6.8 summarises the chapter.

6.1 Retrospectives within Software Development

The last principle within the agile manifesto proposes that “at regular intervals, the team reflects on how to become more effective, then tunes and adjusts its behaviour accordingly” [198]. The basis for this principle is that no software development process is perfect and that agile development teams will encounter new and unique situations whenever they take on software projects [199]. As a result, agile software development teams are encouraged to continually inspect, reflect and adapt their development processes to match their changing situations and environments [199–203].

A common technique used by agile software development teams to inspect, reflect on and adapt processes and practices is the *retrospective* [204]. Retrospectives are typically held at the end of a Sprint (a development iteration) [204]. Pham and Pham define retrospectives as a “meeting during which a Scrum team will go through what worked and what did not work during the Sprint they have just finished and determine whether there is anything they can learn from their experience that will make the process even better for the next Sprint” [205]. Although this definition suggests that retrospectives are a part of the Scrum methodology, retrospectives also play important roles in eXtreme Programming [206] and Crystal Clear [207]. In addition to finding ways to improve, agile teams can also use retrospectives to understand the reasons behind missed targets, findings ways to improve responses to customers and rebuilding damaged relationships [204]. Numerous methods have been proposed for gathering data during a retrospective, including using colour coded dots, ‘mad sad glad’ and satisfaction histograms [204]. However, a retrospective can also be implemented by simply asking an agile team three questions [208]:

- What worked well during this iteration and that we want to continue doing?
- What did not work well during this iteration that we should stop doing?
- What should we start doing to improve?

The answers to these three questions can then be collected and analysed to determine what actionable changes can be implemented, which would allow an agile team to continue to the next iteration with an incrementally-improved process [208].

Several studies have examined the impact of retrospectives on agile teams [200, 209, 210]. Maham [200] studied how an agile Scrum team performed retrospectives after a three-week Sprint. The results from the study showed that the team members highlighted what had worked well and what could be improved on from the previous Sprint. Areas of improvement were then prioritised and implemented in the next Sprint [200]. Maham noted that the success of the practice resulted in the team performing ‘release retrospectives’ at the end of each software version release [200]. The purpose of the ‘release retrospective’ was that it provided the software development team with a high-level view of “where they had been and how they had gone to their current location” [200].

Tiwari and Alikhan [209] reported that the agile team within their study found retrospectives to be ineffective and boring. To help the team realise the value of retrospectives, Tiwari and Alikhan changed the scope and method of the retrospectives and decided to include customers and clients in the practice instead of just the agile team [209]. The perceived benefit of this modification was that it would provide customers and clients with an opportunity to hear first-hand about the team’s performance in the previous iterations [209]. Tiwari and Alikhan reported that after a few retrospectives, the customers and clients began to contribute to the retrospective and that these individuals noted that they felt like they were a part of the agile team and could contribute towards the development of their product [209].

McHugh, et al [210] studied three agile teams to examine if agile practices can enhance trust among team members. The results from this study showed that all three teams agreed that retrospectives provide transparency and visibility regarding the achievement of Sprint goals. In addition, the studied team members noted that retrospectives provided their teams with an opportunity to seek clarification from each other when delays occur and what caused these delays [210].

Multiple agile teams can often be working on the same product or project. Often, each team will do their own retrospective and then look to conduct a *retrospective of retrospectives* [211]. Gonçalves and Linders describe a retrospective of retrospectives as a method of improving collaboration between the various teams [211]. Retrospectives of retrospectives can also be viewed as a tool for sharing information between agile teams [211].

The success of retrospectives within software development motivated a proposal (from the author) to employ retrospectives within security incident response. The intention was that the retrospective would provide two benefits. First, the retrospectives would provide security incident handlers with an opportunity to ‘take a step back’ from the complex environment of a security investigation and allow them to reflect on the information documented during the investigation. Second, the retrospectives could help a security incident response team capture additional information about a security investigation.

As noted in Chapter three, one of the goals of the ‘follow-up’ phase within typical security incident response approaches is that an organisation implements security controls and process improvements that it identifies during the investigation [7,8,185]. However, these approaches do not define how an organisation can identify or evaluate if the improvements identified during the ‘follow-up’ phase are actually implemented or if the modifications were effective. Therefore, a retrospective of retrospectives (hereafter referred to as a *meta-retrospective*) was also introduced into the organisation’s security incident response process. Based on the approach proposed by Gonçalves and Linders [211], a meta-retrospective was used to evaluate if security controls and incident response-related process improvements identified in the retrospectives were actually implemented within an organisation.

6.2 Retrospectives Within Security Incident Response

This section describes how retrospectives and meta-retrospectives were adapted for use within security incident response and then implemented within the Fortune 500 Organisation. The retrospective consisted of six questions. The objective of the six questions was to establish 'What worked well?', 'What did not work well?' and 'What should we start doing or improve?'. The six questions used in the retrospectives are shown in Figure 6.1.

1. Which information assets did you need to investigate in this security event/ incident?
2. Which information asset could you not investigate in this security event/ incident?
3. Who did you need to communicate with during this security event/ incident?
4. Who could you not communicate with during this security event/ incident?
5. What information security controls could have prevented the security event/ incident from occurring?
6. What process changes would help you investigate a similar security event/incident in the future?

Figure 6.1: Security Incident Response Retrospective Questions

Questions 1 and 3 are used to identify 'What worked well' regarding any information assets and individuals involved in the investigation. Information assets are defined as previously (in Chapter 3) for the purpose of this experiment. Questions 2 and 4 are used to identify 'What did not work well' during the investigation. Question 2 is used to identify from the incident handler what information assets could not be included in the investigation, while Question 4 is used to identify which individuals within the organisation, the incident handler had problems communicating with during an investigation. Questions 5 and 6 are used to identify, ‘What should we start doing or looking to improve’. These two questions focus on identifying security control improvements to prevent recurrence and security incident response-related process improvements to assist security incident handlers with a similar investigation in the future. Question 5 probes the identification and documentation of one or more security controls, which could have prevented the security event/incident from occurring. Question 6 provides an opportunity for the incident handler to identify security incident response-related process improvements to assist in future investigations.

In software development, retrospectives are typically held at the end of a development iteration [204]. However, unlike agile software development where work is broken down into iterations, security incident response investigations generally ‘pause’ at the end of an investigation (i.e. the end of the process life-cycle). Therefore, the retrospectives within security incident response in the organisation were conducted at the end of each security investigation. This was typically within one to three days after the closure of the investigation record by an incident handler. In practice, just under 92% of the retrospectives were held within this time frame and the longest time between the closure of an investigation and a retrospective was seven days. This time period was proposed and agreed upon with the organisation’s Head of Information Security.

The meta-retrospective was also adopted for use within security incident response and consisted of asking two questions as shown in Figure 6.2. The purpose of the meta-retrospective within security incident response is to evaluate if the security controls and/or security incident response-related process improvements identified during the retrospectives are implemented within an organisation. Therefore, a meta-retrospective is only undertaken when a security incident handler identifies a security control and/or process improvement during the initial retrospective. If the control or improvement had not been implemented, then a query was made as to why the change had not happened. Meta-retrospectives were undertaken with the same incident handler(s) involved in the initial retrospective. It was planned to conduct a meta-retrospective between three to four weeks after the initial retrospective was undertaken. 100% of the meta-retrospectives were undertaken within this time period. This time period would allow the organisation to implement the improvements identified in the retrospective. The specific time period was used based on the requirements of the organisation’s Head of Information Security.

Typically in a software development context, a retrospective is performed with all the members of an agile team present at the same time [204, 208]. As a result, an agile team can view the outcomes from the retrospective and collectively decide on how to improve in the next iteration [208]. However, security incident response within the organisation is typically undertaken by only one security incident handler. Therefore, a briefing was used to discuss the collective results from the retrospectives and meta-retrospectives at the end of the experiment. The briefing also provided the security incident handlers and their managers an opportunity to reflect on the data collected from the retrospectives and meta-retrospectives. The next section will discuss the experiment design explaining how the retrospectives and meta-retrospectives were implemented within the Fortune 500 Organisation.

1. Have the security controls you identified in the retrospective been implemented? If No, why not?
2. Have the process improvements you identified in the retrospective been made? If No, why not?

Figure 6.2: Meta-Retrospective Questions

6.3 Experiment Design

An experiment was devised based on the on-going case study in the Fortune 500 Organisation, where the retrospectives and meta-retrospectives were implemented within the organisation's security incident response team. Three research questions guided the experiment:

1. Do retrospectives help a security incident response team to identify and document additional information about a security investigation, which may otherwise not be documented within the corresponding investigation record?
2. Do retrospectives help a security incident response team to identify and document security controls and security incident response-related process improvements?
3. To what extent can a meta-retrospective highlight how many security controls and security incident response-related process improvements are actually implemented within an organisation?

An experiment was designed as follows. The author performed face-to-face retrospectives with the Primary Incident Handler(s) (PIHs) identified from the particular investigation record as being involved in a security investigation. This involved the author asking the PIHs the six questions as shown in Figure 6.1. Recall from Chapter four, that a PIH is an individual who facilitates and coordinates the organisation's response to a security event/incident. During a retrospective, PIHs were allowed to open and view the corresponding investigation record. Responses to the retrospective questions were initially recorded by hand and then digitally documented. Each retrospective lasted approximately ten minutes and was conducted at the PIH's desk.

A total of 324 retrospectives were undertaken between February 2014 and March 2015. These retrospectives were conducted with six individuals identified as the PIHs from the investigation records. 321 out of the 324 retrospectives were conducted with only one PIH. One retrospective was conducted with two PIHs and two retrospectives were conducted with three PIHs. The security incident response team uses multiple PIHs when a large-scale security investigation arises within the organisation. The data collected from the 324 retrospectives is presented in Section 6.4.

The organisation was given the opportunity to implement any security controls and/or security incident response-related process improvements identified in the retrospective. Any changes to security controls and security incident response-related processes were initiated based on the recommendations from the PIHs. These controls were implemented together with the support of the individuals within the Information Security unit who are responsible for security control and security-related process implementation. A meta-retrospective was then undertaken to examine if the control or process improvement that had been identified in the original retrospective was implemented within the organisation. If the control or process improvement had not been implemented, PIHs were asked to explain why. A total of 48 meta-retrospectives were undertaken within the organisation. Each meta-retrospective lasted approximately five minutes and was conducted at the PIH's desk. The PIH's responses were initially recorded by hand and then digitally documented. The results of the meta-retrospectives are presented in Section 6.5.

The briefing was then undertaken with the PIHs who participated in the retrospectives and meta-retrospectives, as well an information security manager. This provided the security incident response team together with their managers an opportunity to reflect on the data generated from the retrospectives and meta-retrospectives. After the briefing, the final stage of the experiment involved undertaking a qualitative and quantitative data collection in order to analyse the effect and impact of the retrospectives on the security incident response team. Qualitative data was collected using semi-structured interviews [30]. The purpose behind the interviews was to gather the security incident response team's opinion on the use of a retrospective as a method for identifying and documenting further information from security investigations. The interview participants consisted of the PIHs who had direct experience with the retrospectives as well as the information security manager who had participated in the briefing. The interviews were conducted in March 2015 and the interview instrument consisted of a combination of open-ended and closed questions [30]. The follow-up interview questions can be found in Appendix B.

In addition to the follow-up interviews, quantitative data was also collected through the analysis of the relevant security investigation records. The analysis involved comparing the answers provided from each question in the retrospectives with the information documented in the corresponding investigation record. The purpose of this analysis was to examine if more or less information was being identified using the retrospective. This analysis was undertaken for all 324 retrospectives and corresponding investigation records.

6.4 Retrospectives Data

This section describes the data collected from the 324 retrospectives undertaken within the organisation.

Question 1: Which assets did you need to investigate in this security event/ incident?

This question was answered in all 324 retrospectives. A total of 502 information assets were identified from this question. Further analysis of the data shows that 37 different information assets were identified from the question. Table 6.1 provides an overview of these information assets which have been grouped together according to their type or purpose to the organisation. A complete list of the assets which were identified from this question can be found in Appendix E.

Asset Group Name	Occurrences
Desktop, personal and laptop computers	6
Email assets and associated logs	280
Intranet and Internet-based assets	5
Network devices, servers and logs	15
Organisation-specific assets	7
Security-specific assets	162
Third-party assets	9
Various data repositories and databases	18
Total	**502**

Table 6.1: Information Assets Identified as Investigated in Retrospectives

The results from Table 6.1 show that the organisation's security incident response team require access to a range of information assets both within and outside the organisation. The data indicates that a large number of security investigations involved the organisation's

Asset Name	Total	Reason Provided
Deleted email messages	1	Could not recover deleted emails because incident handler together with relevant business unit could not determine when to restore the affected email account.
Email attachments	3	Email attachments were missing because of a historical problem where certain attachments were deleted several years ago.
Encrypted file contents	4	Lack of decryption keys because individual left organisation and decryption keys were deleted.
Laptop computer	1	Incident handler wanted to perform deeper analysis but was not requested to undertake such an analysis.
Live email account	1	Could not access live email account because it was disabled after the individual left the organisation.
Lotus Notes mail file	9	Data retention period had expired at the time of the security investigation.
Organisation-specific system	1	System owner was not found during the course of the investigation.
Virtual machines	1	Could not get physical access to the virtual machines to perform a deeper analysis.
Windows Registry settings	1	Lack of available tools

Table 6.2: Information Assets which could not be Investigated

‘Email assets and associated logs’ group. This group of assets were identified in 280 (55%) out of the 502 information assets identified from this question in the retrospective. This result reflects the organisations use of Email as the primary method of communication both within the organisation and to third-parties outside of the organisation.

Question 2: Which information asset could you not investigate in this security event/incident?

Question two from the retrospectives was answered 18 times by the incident handlers. In the remaining 306 retrospectives, the answer provided to this question was “None”. The incident handlers identified that they had encountered problems investigating nine individual assets during their security investigations. In total, 22 information assets were identified from the 18 retrospectives, with three of these assets repeatedly identified in multiple retrospectives. Table 6.2 provides a breakdown of the information assets, which the incident handlers could not investigate, along with the reason provided during the retrospective.

The table shows that the incident handlers identified several different information assets which they could not investigate because of various reasons. These assets include virtual machines that had since been deactivated, expired email accounts and missing email attachments due to expired retention periods. One observation from Table 6.2 is that in a number of cases, the ability to investigate an asset is largely down to factors outside the control of the security incident response team. For example, there were nine investigations where the incident handlers required access to a Lotus Notes mail file. However, the data retention period for these mail files had expired and therefore, was no longer available for examination. It is also worth noting that in one retrospective, the incident handler indicated that the reason why they could not investigate the contents of the Windows Registry was due to limited tool access. This supports previous findings [27] that security incident response teams often need to develop their own tools and use tacit knowledge to perform specific tasks within security incident response.

Question 3: Who did you need to communicate with during this security event/ incident?

Question three was answered in all 324 retrospectives and the incident handlers indicated that they needed to communicate with a total of 737 individuals or teams within the organisation during their security investigations. Table 6.3 provides an overview of the individuals or teams within the organisation, whose assistance was required during security investigations. These security incident response contacts have been categorised according to the common function they provide to the organisational structure. A complete list of the individuals or organisational teams identified by the incident handlers in question three can be found in Appendix E.

Individual/Organisational Team Identified	Total
Customer-facing organisational units	21
Email and information technology services units	129
Individuals, managers and team leaders affected by incidents	323
Legal and regulatory requirements units	58
Physical and information security units	154
Software development and support units	10
Third-parties and contractors	42
Total	**737**

Table 6.3: Individuals and Organisational Teams Identified in Retrospectives

The data collected from question three shows that in addition to the individuals affected by a security event or incident, their managers and team leaders (323 cases), the most common teams within the organisation that the incident handlers needed to communicate with during security investigations are the Physical and Information Security units (154 cases) and the Email and Information Technology Services units (129 cases). The findings from this retrospective question support previous findings [43, 110, 212], that security incident response teams need to collaborate with various individuals and organisational teams (often in various disciplines) when investigating security events and incidents. In addition, the results also suggest that the incident handlers needed to communicate with individuals external to the organisation including third-party vendors, Information Technology support teams and external legal services (Third-parties and contractors).

Question 4: Who could you not communicate with during this security event/ incident?

The purpose of fourth question was to identify which individuals or teams within the organisation the incident handlers had problems communicating with during a security investigation. The incident handlers identified five cases where communication was a problem during an investigation. A total of four individuals were identified through this question, with one individual being identified in two different retrospectives.

In four retrospectives, the individuals were information asset owners and their assistance was required for gathering data from their information asset. In the fifth retrospective, the individual's assistance was required as he/she was a managerial figure for an individual affected by a security event. The incident handler indicated that the communication problems were caused by outdated contact information (four cases) and the individual concerned being away on holiday (one case).

Question 5: What information security controls could have prevented the security event/incident from occurring?

Question five was used to identify whether one or more security controls, if implemented, could have prevented the particular security event/incident from occurring. The incident handlers identified a total of 36 security controls from 30 retrospectives. Moreover, there were 179 “None” responses and 115 “Not-applicable” responses from the individuals in response to this question.

Within 25 retrospectives, the incident handlers identified that a single security control could have prevented the security event/incident, while in four retrospectives, the incident handlers noted that two security controls would be required. There was one retrospective were three security controls were identified from this question. The security controls identified in each retrospective can be classified as belonging into one of three groups: technical security controls, administrative security controls or a combination of these two types [54]. Within five

ISO/IEC 27002 Domain	Total	Example
Access Control	9	Define administrative permissions for virtual machines used in test development servers
Asset Management	2	Check all hardware for CD media prior to sending for recycling
Communications Security	1	Enhanced logging on Network File System share
Cryptography	2	Transport Layer Security to be implemented between organisation and third-party involved in incident
Human Resources Security	4	Education re-enforcement surrounding Secure Remote Access Service tokens and PIN numbers
Information Security Policies	4	Creation of new lock-down standard for web server security
Operations Security	10	Block access to specific file-upload portal on web gateway
Supplier Relationships	4	Third-party involved to implement technical and procedural controls
Total	**36**	

Table 6.4: Security Controls Identified in Retrospectives

retrospectives, the incident handlers indicated that a combination of administrative and technical security controls were required, while in 16 retrospectives, only technical controls were recommended. In nine retrospectives, only administrative security controls were proposed.

In order to protect the confidentiality of the organisation's implemented security controls, the controls identified by the incident handlers in response to Question 5 are mapped to one of the fourteen domains within the International Organization for Standardization/ International Electrotechnical Commission (ISO/IEC) 27002 information security standard [6]. Table 6.4 presents the number of controls identified in eight of the domains from the standard, along with an example security control identified during the retrospectives. This mapping shows that 19 out of the 36 security controls identified from question 5 are related to 'Access Control' and 'Operations Security'. This result reflects the organisational location of the security incident response team within the Information Security unit, which is responsible for implementing everyday operational security. The results also show that the incident handlers prefer to focus on restricting access to information assets, through the enforcement of enhanced access controls.

Question 6: What process changes would help you investigate a similar security event/incident in the future?

Question six specifically attempts to identify security incident response-related process improvements, which would help the incident handlers investigate similar security problems in the future. The incident handlers identified a total of 29 process improvements from 28 retrospectives. Within 27 retrospectives, the incident handlers identified a single process improvement, while in one retrospective, the incident handlers identified two process improvements. The 29 process improvements can be categorised into one of four types. Four of the process improvements would involve the creation of a new process within the organisation, e.g. the creation of a process for managing lost and stolen laptops within the organisation. Eleven of the process improvements involved enhancing existing processes within the organisation, which includes adding contact information for asset owners to the organisation's electronic address book. Eight of the process improvements identified by the incident handlers involved the introduction of new tools and/or methods to assist the security incident response team with future investigations. Six process improvements involved changes to existing processes owned by other teams within the organisation, but affect the security incident response team. An example of such a process improvement is modifying existing email message recovery procedures so that attachments are included with all email messages.

6.5 Meta-Retrospectives Data

Recall that questions five and six from the retrospectives were used to identify security controls and security incident response-related process improvements. A total of 65 security controls and security incident response-related process improvements were identified from 48 retrospectives. These improvements were identified as follows.

- Within 20 retrospectives, the incident handlers identified only a security control and no process improvements. 23 security controls were identified in these 20 retrospectives.
- Within 18 retrospectives, the incident handlers identified only a process improvement and no security controls. 18 process improvements were identified in these 18 retrospectives.
- Within 10 retrospectives, the incident handlers identified both a security control and a process improvement. 13 security controls and 11 process improvements were identified in these 10 retrospectives.

48 meta-retrospectives were undertaken in order to determine whether the security controls and process improvements identified in the retrospectives had been implemented and if not, determine why an enhancement was not made. The results from the meta-retrospectives inquiries were categorised in the following manner:

- Change Made (CM) – the proposed security control or security incident response-related process improvement was implemented within the organisation at the time of the meta-retrospective.
- Change On-going (CO) – the implementation of the proposed security control or security incident response-related process improvement was still on-going at the time of the meta-retrospective.
- Change After meta-retrospective (CA) – the implementation of the proposed security control or security incident response-related process improvement was initiated during or after the meta-retrospective.
- Escalated to Management (EM) – the implementation of the proposed security control or security incident response-related process improvement was escalated to senior management within the Information Security unit for progression.
- No Changes (NC) – the proposed security control or security incident response-related process improvement was not implemented.

Table 6.5 presents the results of the 48 meta-retrospectives. The table shows that 42 out of the 65 security control and process improvements identified in the retrospectives were either implemented in the organisation or their implementation was considered ‘on-going’ at the time of the meta-retrospective. Two enhancements (one security control and one process improvement) resulted in no changes being implemented at the time of the meta-retrospective. However, in both cases the PIHs started to take actions to implement these enhancements shortly after the meta-retrospective. One possible explanation is that the act of the undertaking a meta-retrospective may have prompted the PIHs to remember to take action with regards to the implementation of these enhancements.

The results from the meta-retrospectives also identified that a security incident response team can face numerous challenges when attempting to implement security controls and process-related improvements. During the meta-retrospectives, the PIHs indicated that 15 out of the 65 enhancements could not be implemented. As a result, these enhancements had to be escalated to senior management within the Information Security unit. In all 15 cases, senior management either assisted in the implementation of the proposed enhancement or continued to champion the enhancement on behalf of the PIHs. This shows that although

Meta-retrospective Type	**CM**	**CO**	**CA**	**EM**	**NC**	**Total**
Security control-only	15	5	1	2	0	**23**
Process improvement-only	5	2	1	6	4	**18**
Combination	10	5	0	7	2	**24**
Total	**30**	**12**	**2**	**15**	**6**	**65**

Key: *CM = Change Made; CO = Change On-going; CA = Change After meta-retrospective; EM = Escalated to Management; NC = No Changes*

Table 6.5: Security Control and Process Improvement Implementation Status

a security incident response team may have autonomy within an organisation to identify, manage and handle security incidents, the team may not necessarily have the authority to change security controls to prevent recurrence.

Six out of the 65 security control and process-related improvements identified in the retrospectives resulted in 'No Changes' being made within the organisation. In all six cases, this involved process-related improvements and not security controls. The reason provided was that the security incident response team does not have authority over all the processes within the organisation. When the security incident response team propose changes to processes owned by other organisational teams, it is dependent upon that particular team to decide whether it will modify its process or not to satisfy the security incident response team. For example, a PIH identified that enhanced logging for a web-based system would assist with investigating the system. However, while the system's owners acknowledged the potential benefit, a business decision was made not to implement more detailed logs and the risk associated with this decision was accepted by the organisation.

6.6 Measuring the Impact of Retrospectives

Recall, the aim of the experiment was to determine if retrospectives can be used by a security incident response team to identify and document enhanced information at the conclusion of a security investigation. In order to examine this further, qualitative and quantitative data was collected from both the security incident response database and the individuals involved in the experiment. Quantitative data was collected through analysing the relevant investigation records and comparing the information documented within these records with the information identified in the corresponding retrospective. This analysis was performed for all 324 retrospectives. Qualitative data was collected through follow-up interviews with practitioners who actively participated in the retrospectives, as well as those who had direct experience with the retrospectives as security incident response process owners within the organisation.

The aim of the data collection was twofold. First, it was used to evaluate if a particular retrospective identified more or less information when compared to its associated security investigation record. Second, the follow-up interviews were conducted with the aim of gaining a practitioner's perspective on the use of retrospectives within security incident response.

6.6.1 Comparison of Retrospectives and Investigation Records

The comparison of the information identified from the retrospectives with the information documented within the investigation records showed that more 'information items' were identified using the retrospectives than those found in the corresponding investigation records. The term 'information item' is used in this context to describe any information asset, individual, group, security control or process improvement identified in either the retrospective or investigation record.

The results show that 148 (46%) out of the 324 retrospectives contained more information about an investigation when compared with the information documented in the corresponding record. 151 (47%) out of the 324 retrospectives identified the same information about an investigation as the information documented in the relevant record. Finally, 25 (7%) investigation records contained more information than what was actually identified using the retrospectives. A detailed comparison of the 'information items' identified in each question from the retrospective with the number of 'information items' documented in the corresponding investigation record is shown in Table 6.6. Note for question three, that results have been provided for a) where more information was identified using the retrospectives; and b) where more information was identified in the investigation records.

Question	Investigation Record	Retrospective	Total
Question 1	424 assets	78 additional assets	502
Question 2	11 assets	11 additional assets	22
Question 3a	601 individuals/groups	94 additional individuals/groups	695
Question 3b	34 additional individuals/groups	42 individuals/groups	76
Question 4	1 individual/group	4 additional individuals/groups	5
Question 5	11 security controls	25 additional security controls	36
Question 6	3 improvements	26 additional improvements	29

Table 6.6: Retrospective vs. Investigation Record Analysis Results

Retrospective Question 1

Within question one, the PIHs identified 502 information assets using the retrospectives. In comparison, the analysis of the corresponding investigation records revealed that 424 information assets were documented within the records. 65 retrospectives contained more information about information assets being required for inclusion in a security investigation than what was documented in the investigation records. In terms of actual assets being identified, 78 additional information assets were identified from the 65 retrospectives. Furthermore, 259 retrospectives and investigation records contained the same information assets, while no investigation record was found to contain more information than a retrospective.

Retrospective Question 2

The information analysed concerning question two showed that there were eleven retrospectives where the PIHs identified more information about the assets, which could not be investigated. Within these eleven retrospectives, the PIHs identified 22 information assets. The analysis of the corresponding investigation records revealed only half the number of assets (11) than those in the retrospectives. The comparative results also showed that seven retrospectives and investigation records contained the same information about assets which could not be investigated.

Retrospective Question 3

The analysis of the information collected from question three showed that there were 80 cases where the PIHs identified more individuals and groups within the organisation using the retrospectives than those documented within the investigation records. An additional 94 individuals and groups were identified in the 80 retrospectives, which were not documented within the corresponding investigation records. However, the analysis also showed that there were 25 instances where more information was documented within the investigation record when compared with the information identified using the retrospective. As a result, an additional 34 individuals and groups were identified in the investigation records from these cases.

Retrospective Question 4

Regarding question four, answers to this question were only provided by the PIHs on five occasions. Although this question was answered only five times, the analysis showed that more information about individuals and groups where communication was a problem was identified using the retrospectives than the investigation records. This proved to the case in four out of the five retrospectives. In the remaining case, the information identified from the retrospective was also found in the corresponding investigation record.

Retrospective Question 5

With regards to question five, the analysis of the investigation records showed that eleven security controls were documented within the investigation records. When compared to the information in the retrospectives, 24 retrospectives contained more information than the corresponding investigation records. In these retrospectives, a further 25 security controls were identified. Six retrospectives and investigation records contained the same security control, while no investigation record contained more information about the security controls than the retrospective.

Retrospective Question 6

For question six, the comparative analysis showed that a total of 29 security incident response-related process improvements were identified using the retrospectives. When the investigation records were analysed, only three process improvements were documented within the records. This means that 26 additional process improvements were identified using the retrospectives. The results suggest that PIHs are identifying process improvements during their investigation but these improvements are not being documented in the investigation records. This may because PIHs are be under time pressures to conclude their investigation and move onto the next security problem. As a result, they may not have enough time during the initial documentation of the investigation to identify process improvements which could help the security incident response team. Further analysis was required to explore this assumption and this was examined in the follow-up interviews.

6.6.2 Follow-Up Interview Analysis

The follow-up interviews were undertaken in March 2015. The objective of the follow-up interviews was to assess the practitioner's perspective on using retrospectives as a method for validating and enhancing the quality of data within security incident response. Seven individuals were interviewed and the follow-up interview questions can be found in Appendix B. The answers to the questions are summarised below.

Description of Participants

Initial questions established the participant's current role in the organisation and provided a brief idea of his/her role within the security incident response team. The answers from these questions revealed that the participants identified themselves as information security managers, senior security analysts and security analysts. One individual was a trainee information security analyst. These individuals assume various roles within the security incident response team. These roles include managers who own and enforce the organisation's security incident response process, as well as analysts who identify, investigate and eradicate information security events and incidents. Six out of the seven individuals indicated that they had participated in the retrospectives. The one individual, who did not participate, is the security incident response policy owner, who had an overview of how the retrospectives were implemented and undertaken within the organisation.

When the participants were asked if they were involved in any post-investigation activities within the organisation, six out of the seven respondents indicated that they are involved in various activities after an investigation has been closed. These activities can include identifying and improving security controls, reviewing and modifying existing security policies and standards, analysing the risk associated with security incidents and escalating security control recommendations to wider stakeholders. Two participants added that in addition to the above activities, retrospectives were also considered to be a part of the organisation's post-investigation activities. This result shows that the retrospectives have not yet been fully recognised by all the individuals within the security incident response team as a post-investigation activity, even though a retrospective was undertaken at the completion of all the investigations during the experiment. One individual indicated that they had not participated in any post-investigation activities because their role did not explicitly involve such activities.

Perceptions of Briefing

The interviewees were then queried on their participation and perception of the briefing discussion. Six out of the seven respondents participated in the briefing, with one individual unable to attend due to time constraints. The participants' responses indicated that there generally advantages with regards to the briefing and the data discussed. The positive answers indicated that the briefing provided an alternative view of the security incident response data, which was useful to identify patterns, themes and potential bottleneck areas. One of the respondents indicated that the data assisted in highlighting where new processes and guidelines for interacting with specific assets and individuals were required. This individual stated "*the retrospectives have highlighted that we have a lot of conversations and interactions with the Lotus Notes team, nearly on a daily basis, yet we don't have an actual process which defines how this should be done, what data is actually available nor what we should do when something goes wrong.*" Another individual, whose job role includes examining how information security risk changes within the organisation noted that "*the data tells me we should be looking to implement security processes at a much lower level than we are at the moment, however this information is usually not shared with the relevant units and therefore the changes are never actually implemented.*" The one manager who did participate in the briefing, added that the data from the briefing has helped to highlight where management needs to focus their attention with regards to the most common threats to the organisation.

Two out of the six respondents who participated in the briefing indicated that there are various opportunities to enhance the briefing experience. One individual suggested they would have liked to see the data from question six from the retrospectives presented into quarter periods instead of the entire year. Another individual suggested that the data from question one from the retrospectives was presented focusing on the types of investigations affecting the specific assets. For example, the number of 'security events', which required the PIHs to have access to the email server and the number of 'security incidents' which required the PIHs to have access to the email server. The individual who made this suggestion stated that this type of analysis could help management improve access to specific assets, which are considered to be more important in overall security investigation objectives.

Using Retrospectives to Identify 'What Worked Well'

The interviewees were then queried about using the retrospectives to identify 'What Worked Well'. Question 5 from the follow-up interview, was used to determine the benefits of using the retrospectives to identify 'What Worked Well' from the perspective of information assets investigated and people communicated with during an investigation. The predominant answer from the majority of the interviewees was that this part of the retrospective has helped to capture additional information, which can be used to identify frequency of asset use and stakeholder involvement. One individual added that the information generated from this part of the retrospective has also helped the security incident response team to identify which teams within the organisation they must enhance and maintain relationships with in the future. The suggestion was that building stronger relationships with the teams and individuals identified in this part of retrospective would assist with resolving investigations quicker and more effectively. One information security manager also argued that this part of the retrospective provided incident handlers with a second chance to document investigation information, which can then be "rolled-up" and used to identify trends and gaps in current practices.

Question 6 from the follow-up interview was used to determine the disadvantage of using the retrospectives to identify 'What Worked Well' during an investigation. None of the respondents indicated that there was a disadvantage to using the retrospective to identify and collect this information. All seven participants stated that this part of the retrospective provided an additional avenue for the PIHs to identify and document information, which may otherwise not be documented within an investigation record. One information security manager added that there can be no disadvantage to this part of the retrospective, provided that the information collected is used by the PIHs to become more proactive in future investigations.

Question 7 was used to determine the overall impact of using the retrospectives to identify assets investigated and people communicated with during an investigation. The majority of respondents stated that this part of the retrospective has had a positive influence on the security incident response process. The respondents indicated that this part of the retrospective can be used to enhance information capture with regards to important assets and stakeholders required for investigations. However, one individual argued that this part of the retrospective could be better suited for more complex 'security incidents' instead of 'simple events', which do not involve several assets and/or individuals during an investigation. When asked to expand on this answer, the individual added that performing a retrospective for 'simple events' does not justify the time required to actually undertake the activity.

Using Retrospectives to Identify 'What Did Not Work Well'

Questions 8 to 10 queried the interviewee about using the retrospective to identify 'What Did Not Work Well'. Question 8 was used to determine any benefits of using retrospectives to identify 'What Did Not Work Well' during an investigation. Five interviewees stated that this part of the retrospective has helped to identify and capture information, which would otherwise not be documented within an investigation record. One of the interviewees argued that this information was not usually documented within an investigation record because it is perceived to be "negative" information, which the incident handlers do not have time to document. The interviewee went on to state that capturing this "negative" information can help with the identification of gaps with regards to obtaining access to key information assets and stakeholders, which could be required for future investigations.

With regards to Question 9, the seven interviewees indicated that there were no disadvantages to using a retrospective to identify and document information about which assets and individuals 'did not work well' during an investigation. Three interviewees added that there could be no disadvantage in identifying and documenting information about an investigation which would normally not be included in a typical security investigation record.

Question 10 was used to examine the overall impact of the retrospective question 'What Did Not Work Well' with regards to information assets, which could not be investigated, and individuals who the PIHs could not communicate with during an investigation. All the interviewees agreed that this part of the retrospective has provided an opportunity for PIHs to stop and reflect about 'what went wrong' in the investigation and to document additional information, which can help the team, improve its process. One interviewee added that this part of the retrospective should be undertaken immediately after the conclusion of an investigation so that any identified problems can be eradicated or improved upon before the next security investigation. Although this was the objective at the start of the experiment, it was not always possible to execute the retrospective immediately after the closure of an investigation. This was because incident handlers were assigned to new investigations, which then took priority over performing a retrospective of previous investigations.

Identifying Information Security Controls

Questions 11 to 13 queried the interviewees about using the retrospective to identify security controls that can be improved upon after a security investigation. The respondents generally agreed that security controls *should* be identified and documented within an investigation record. One of the information security managers reiterated that this is actually a process requirement and that the investigation record includes a field for documenting this information. However, the manager added that in reality this does not occur. These comments support the results from the analysis presented in Section 6.6.1, which showed that security control information is not always documented in the investigation record. When asked why this was the case, the majority of the incident handlers stated that information was not documented because of limited knowledge surrounding current security controls. While the incident handlers acknowledge that its is important to document this information, they have also stated that their limited knowledge about the security controls implemented in the organisation has affected their ability to identify security control improvements.

The incident handler's interview answers also suggest that the retrospective could be one way of solving this problem. The incident handlers indicated that the retrospective provided a 'safety-net mechanism' to help document security controls improvements. More specifically, one of the PIHs suggested that the problem with identifying security controls could be countered by making this part of the retrospective, a group activity. The reasoning behind this suggestion was that individuals within the security incident response team will have varied levels of understanding of the security controls implemented within the organisation. The individual added that a group retrospective specifically focusing on security control improvements, would help PIHs with limited knowledge about security controls to identify improvements. Although the majority of the retrospectives undertaken within the organisation were conducted with a single individual, a group retrospective for identifying security controls emulates how retrospectives are undertaken in other domains, such as agile software development [200, 209, 210]. Within agile software development, the Agile Manifesto encourages autonomy and providing individuals the environment and support they need to reach specific objectives [198]. Similarly, undertaking a group retrospective could provide a security incident response team with the right environment to expand their knowledge and help with the identification of further security controls.

Identifying Security Incident Response-Related Process Improvements

Questions 14 to 16 from the follow-up interview, were used to query the interviewees about using the retrospective to identify security incident response-related process improvements. The interviewees were unanimous in that the retrospectives have assisted with the identification of process improvements that can help the security incident response team. Two individuals indicated that this type of information was important to identify and document. However, the individuals added that the organisation's security incident response process does not require the PIHs to document this information. The two individuals went on to state that the retrospective has assisted with the identification and documentation of process improvements, which would otherwise not have been documented within the organisation.

One of the information security managers provided an alternative view on the matter. The manager suggested that they did not expect any security incident response-related process improvements to be recorded within the investigation record. This is because PIHs are more likely to be focused on eradicating and recovering from security problems, rather than deciding how to improve the way they work. The manager added that looking for ways to improve is important and that these specific questions in the retrospective have helped with identifying and documenting process improvements, which would otherwise not be documented.

Overall, the interviewees were in agreement that the retrospectives had a positive influence on the identification and documentation of security incident response-related process improvements. The main reason provided by the respondents was that this part of the retrospective provided a mechanism, which prompted incident handlers to stop and think about how they can improve the way they conduct security investigations. Six out of the seven interviewees indicated that there was no disadvantage to using the retrospectives to identify and document security incident response-related process improvements. However, one individual noted that this part of the retrospective was 'less important' than identifying security controls and suggested that instead, a monthly retrospective is undertaken to identify process improvements. There are both advantages and disadvantages to executing a monthly retrospective. The advantage is that a security incident response team can obtain a wider perspective on the various processes that they have identified as needing to be improvement to assist with future investigations. However, the disadvantage is that undertaking a monthly retrospective can also mean that there is the possibility that PIHs could forget details about investigations and therefore, fewer security process improvements will actually be identified.

Other Factors

Question 17 is used to determine what other factors have contributed to the successful or unsuccessful attempt in using retrospectives to identify and document security controls and security incident response-related process improvements. Five interviewees acknowledged that the incident handlers limited knowledge of the organisation's security controls and security incident response-related processes was a key factor in their inability to identify improvements. These individuals added that if the PIH does not know about a specific control or the existence of a particular process, it could be difficult to suggest improvements.

Two individuals added that the success of a retrospective was dependent on how quickly it was undertaken after the closure of an investigation. The individuals indicated that the retrospectives should be undertaken swiftly at the conclusion of an investigation. However, both individuals concluded that this might not be feasible. This is because when the security incident response team has to handle multiple investigations at once, the priority is to close the investigation and move onto the next problem. One information security manager suggested that a security incident response team needs to 'buy-in' into the idea of using retrospectives as a method for documenting further information and improving security controls and processes. The manager added that "*a security incident response team, together with their managers, need to want to improve and without the will to want to improve, there is no point in undertaking the retrospectives*."

6.7 Discussion

One of the themes which emerges from the experiment is that the completion of a security investigation and the closure of its corresponding investigation record does not necessarily mean that an incident handler has documented all the information about the particular investigation. The comparative analysis of the retrospectives and security investigation records has shown that further information about an investigation can be identified and documented during a retrospective.

Experiment Research Question 1

The results from the experiment have shown that retrospectives can help a security incident response team to identify and document additional information about a security investigation. The results from the comparative analysis has shown that integrating retrospectives into security incident response can help with the documentation of additional information which may not necessarily be captured during the initial investigation. 148 out of the 324 retrospectives (46%) undertaken during the experiment resulted in more information being identified and documented during a retrospective when compared to the corresponding investigation record. Further analysis showed that the retrospectives were particularly useful in capturing 'negative information' such as what assets could not be investigated and which individuals the PIHs encountered problems communicating with during an investigation. However, the results of the comparative analysis also showed that a retrospective is not a replacement for the investigation process itself. 25 security investigation records were found to contain more information about the investigation, when compared to the information identified using the retrospective. This was particularly evident from question three in the retrospectives, where 34 individuals and groups within the organisation were documented within the investigation record, but were not identified during the retrospective.

The results from the follow-up interviews have shown that retrospectives can also provide an additional avenue for the identification and documentation of information, which may otherwise not have been recorded. The interviewee's comments indicated that information captured from the 'What Worked Well' questions could provide the security incident response team and management with an indication of the number of security investigations involving key information assets and stakeholders. As highlighted in the follow-up interviews, the advantage of capturing this information is that it can help a security incident response team to identify which individuals and groups it should maintain relationships with to resolve investigations more quickly and efficiently in the future.

Experiment Research Question 2

The results show that retrospectives can help a security incident response team to identify and document information security controls and security incident response-related process improvements. A comparison of the information from the retrospectives and the information in the investigation records has shown that the retrospectives have helped to identify and document 25 additional security controls and 26 additional security incident response-related process improvements. With many organisations looking to implement information security standards such as ISO/IEC 27001/27002, it is becoming increasingly important to be able to audit what security controls are implemented within an organisation. Having a rich dataset of security investigation records from a detailed data generation process can assist with this auditing process. The results from the meta-retrospectives have shown that the activity can be used to audit which security controls have actually been implemented, as well as which controls still need to be implemented within an organisation.

Although the identification and documentation of security controls is a process requirement, within the Fortune 500 Organisation, individuals noted in the follow-up interviews that this information is typically not recorded. The reason provided is that PIHs are under time pressures and often have limited knowledge surrounding existing security controls. While the PIHs acknowledged that the documentation of this information is important, they have argued that unless they know about existing security controls, it can be difficult to suggest improvements. A solution proposed by one PIH is to make this part of the retrospective a group activity where individuals with a good understanding of the security controls can help those members who have less knowledge within this area.

Experiment Research Question 3

The results from the meta-retrospectives show that 42 out of the 65 security control and security incident response-related process improvements identified using the retrospectives were either implemented in the organisation or their implementation was considered 'on-going' at the time of the meta-retrospective. A further two improvements were implemented after the specific meta-retrospective were undertaken. In both of these cases, no changes were implemented prior to the meta-retrospective and the changes within the organisation were only instigated during the meta-retrospective.

The results from the meta-retrospectives have also showed that a security incident response team faces numerous challenges when attempting to implement security controls and process-related improvements. During the meta-retrospectives, incident handlers reported that 15 out of the 65 security control and process improvements had to be escalated to management within the Information Security unit. These managers then either assisted with the implementation of the security control/process improvement or continued to champion its implementation on-behalf of the PIH. Furthermore, six out of the 65 security control/process improvements identified in the retrospectives resulted in no changes within the organisation. This is because the incident handlers do not have authority over all the processes within the organisation. These findings suggest that organisations need to assist security incident response teams implement security controls/process improvements in order to improve the wider security posture and security incident response processes.

Process Culture and Retrospectives

The comparative analysis of the retrospective results with the investigation records suggests that further information about an investigation can be identified and documented using a retrospective. However, the results of this analysis have also emphasised challenges with process and management demands with regards to security investigations. This is particularly evident in answers from question one through to question four, which were used to identify what assets could/could not be investigation, and which individuals could/could not be communicated with during a particular investigation. The answers from questions two and four provide specific evidence that a 'process culture' exists within the organisation, which Deal and Kennedy [150] argue is common in financial services organisations. The downside of a process culture is that stress can arise because of internal politics, bureaucracy and problems with systems and processes currently used in an organisation [150]. As the retrospectives have highlighted, management expect incident handlers to conduct through and conclusive investigations, but this is not always possible and the security incident response process does not take into consideration deviations from the ‘normal’ investigative events. The process is expected to work all the time, but the retrospectives have highlighted that incident handlers may not have access to some assets or individuals required for their investigations. In other words, the organisation places less emphasis on adaptability and changes to the process, which can mean that assets may not be investigated, and therefore incident causes may not be identified [213].

However, the retrospectives can also be used as a feedback mechanism for a security incident response process. Another problem with a process culture is the lack of immediate feedback from the process, which can result in employees not knowing if the process actually works [150]. The results from the experiment suggest that retrospectives could be one way a security incident response team can measure their work by receiving feedback from the questions they ask during the retrospectives. Focusing the retrospective questions on the incident response process, as was demonstrated in the experiment, allows an incident response team to obtain feedback on the process they use for incident investigations.

Impact on Information Dissemination and Incident Learning

Researchers [13, 14] have argued that organisations find it difficult to disseminate information from security incidents. These researchers go on to state that information and knowledge from a security incident investigation is either not documented, and when it is documented, information often does not reach management [13, 14]. Ying adds that there is no systematic or standardised way to disseminate or manage information dissemination in security incident response [129]. While a detailed report could be produced for management, these individuals may find it difficult to digest the information in the report because of the ‘interrelated information it will contain [14]. Furthermore, these reports are likely to be written for an administrative purpose rather than an ‘engineering’ purpose [214]. The results from the comparative analysis of the retrospectives with the investigation records suggest that the retrospective could be a lightweight mechanism to support information dissemination in incident response for ‘engineering’ purposes. As the experiment has shown, the retrospective can be tailored to collect information that will be useful to a security incident response team and its managers. Therefore, unlike a report the information will not serve just an administrative purpose, but also provide a perspective on how security controls and security incident response-related processes can be improved.

While the retrospectives can be used improve information dissemination and feedback from a security incident response process, the information collected using retrospective can also be used for wider organisational learning. By extension, this can be considered one of the objectives of the experiment that the additional information captured using the retrospectives would be used by a security incident response team to learn about a security incident. However, an observation made during the experiment was that the security incident response team in the Fortune 500 organisation did not frequently use this information for the purpose of extending their incident learning capabilities. Even though the team had this additional information at their disposal, very few actions were taken to extend learning from particular incidents. A further observation was that managers often requested information from the incident response team. This information often came from the retrospectives themselves, which was then used for metric reporting to higher-level managers. The question which needs to be asked is does the additional information captured from the retrospectives help the security incident response team learn about an incident or does it help management learn more about how the process is used by the team? One could argue that some form of learning is taking place by management who were using the information to learn if the process was working, but this would not necessarily help the organisation learn about security incidents or how to prevent them. Furthermore, it would appear that minimal support was given to the security incident response team in order for them to learn about security incidents using the collected information. These are problems and issues which can arise in organisations that have a process culture, where internal politics, bureaucracy and problems with systems currently used in the organization can hinder improvements and feedback in the organization itself [150].

6.8 Summary

This chapter described an experiment to evaluate the use of retrospectives as a method for validating and enhancing data collected during a security investigation. Furthermore, the experiment evaluated the use of meta-retrospectives within security incident response as method for identifying what security controls and process improvements are actually implemented within an organisation. The results from the experiment have shown that retrospectives can be used to 'inspect and adapt' security investigations processes. It is through the process of 'inspection and adaptation' that the retrospectives have been used to first validate the information documented within an investigation record and second, enhance the information that has been missed during the investigation. Practitioners interviewed within the studied organisation have stated that the retrospectives provided a 'safety net' and an additional avenue for incident handlers within the organisation to document information which may otherwise not have been documented in the investigation record. A comparative analysis with the corresponding investigation records has also shown that retrospectives can help to identify and document "negative" information, which can be used to enhance the quality of data generated from a security investigation.

Chapter 7

Security Incident Response Dashboard

Chapter three established that detailed security investigations increase the potential for an organisation to identify information that could be used to prevent future incidents. However, before an organisation can learn from a security incident, they need to collect detailed data from their investigations. The exploratory case study in Chapter four demonstrated that many of the security investigation records in the Fortune 500 Organisation did not contain enough information. For example, 126 out of 188 investigation records did not contain date and time information with regards to response-time metric calculations. Furthermore, the exploratory case study highlighted that the information documented within the security investigation records did not reflect the state of an investigation as understood by security incident handlers. In these cases, there was the potential to capture additional information about the investigation, as many fields in the relevant security investigations records were incomplete.

This chapter describes an experiment in which a security incident response dashboard was implemented within the Fortune 500 Organisation. The purpose of the dashboard was twofold. First, it was hypothesised that the dashboard would enhance the transparency of data quality issues within security investigation records in the organisation. Second, it was hypothesised that the dashboard would assist the security incident response team to identify and correct incomplete security investigation records. The chapter is structured as follows. Section 7.1 introduces dashboards and provides an overview of previous research focused on developing dashboards within information security. Section 7.2 provides an overview of the experiment and the research questions that guided the experiment. Section 7.3 discusses how the various dashboard requirements were identified and collected from individuals within the Fortune 500 Organisation. Section 7.4 describes the dashboard architecture including the hardware and software components required to develop and deploy the dashboard application. Section 7.5 discusses the deployment of the dashboard within the Fortune 500 Organisation's security incident response team. Section 7.6 reports on the impact of the dashboard through a qualitative and quantitative data analysis. Section 7.7 discusses the main findings from the experiment and examines the challenges and limitations of the experiment. Section 7.8 concludes and summarises the chapter.

7.1 Information Security Dashboards

In the past few years, information security vendors have attempted to integrate a variety of visual indicators into information security dashboards as a method of reporting organisational security performance. A dashboard is defined as a "visual display of information needed to achieve one or more objectives which fits entirely on a single computer screen so it can be monitored at a glance" [215]. Marty [216] argues that one of the key points not covered in this definition is that a dashboard must be constructed for a specific audience. Marty goes on to state that there are three main types of dashboards within information security [216]:

- **Operational dashboards**, which are used to track core processes, metrics, and statuses within an organisation. This type of dashboard provides low-level information at a glance and is often used for real-time security monitoring activities.
- **Tactical dashboards**, which are used to track processes of organisational groups, computer networks and security states of systems. These types of dashboard help to analyse security conditions and summarise data to analyse security problems.
- **Strategic dashboards**, which are used to monitor the execution of strategic objectives and are usually used by senior security executives. Information from these dashboards is typically presented in the form of trend visualisations, while real-time information is less common.

Marty [216] adds that, no matter what type of dashboard is employed within an organisation, an important requirement is 'comparison'. In this context, 'comparison' is the capability to see trends and changes over time that should attract the attention of the viewer who uses a dashboard [216]. At the time of writing, several commercial information security dashboards were available including Splunk [217], IBM Tivoli Compliance Insight Manager [218] and Security Wizard's Radar [219]. In addition to these commercial dashboards, several researchers have also proposed and developed a variety of dashboards for use within an information security context [220–222].

Sun, et al. [220] proposed and develop a dashboard with the aim of helping system administrators understand the state of system security within their organisation through the use of security metrics. The dashboard was developed using Java and a prototype was presented in their paper, however the prototype was not evaluated outside of a lab-based environment [220].

Novikova and Kotenko [223] proposed and developed a prototype visualisation component for a Security Information and Event Management (SIEM) system which helps to capture and store large amounts of data from monitoring activities. The visualisation offered attack modelling graphics and security evaluation calculations using data gathered from the SIEM system. These same researchers then expand on this earlier work and developed *VisSec-Analyzer* [221], a dashboard that presents information about network security monitoring activities. The tool was demonstrated in a lab-based environment, but was not implemented within a commercial setting.

Sobesto, et al. [222] developed DarkNOC, which is a dashboard for managing and monitoring collections of honeypots and other data collecting devices on a computer network. DarkNOC was deployed on a university network, which is used for honeypot research purposes. The results from the study showed that DarkNOC provided researchers with information about targeted systems, attacks and their origin, as well as an overview of the honeypot activities within the subnet [222].

Dashboards have also been proposed for use with data generated from a security incident response process [224, 225]. Jacobs and Rudis [224] provide guidelines and recommendations for using visualisations, including a dashboard to display information from security incidents and data breaches. Jacobs and Rudis argue that the goal in collecting and visualising security breach data "is to support the decision-making process within security leadership" [224]. They go on to state that collecting and visualising data on breaches helps to reduce uncertainty between "what you know and what you need to know" and this is particularly important in responding to data breaches [224].

As part of a larger project into exploring security incident response data generated by national Computer Security Incident Response Teams (CSIRTs), Madnick, et al. [225] developed the Exploring Cyber Incident Relations (ECIR) Data Dashboard. The purpose of the dashboard was to provide a comprehensive view of a data set collected from national level CSIRTs about the state of their security incidents affecting their countries [225]. The dashboard used publicly available data and presented visualisations of this information in three parts: demographic information about the countries, information technology data and cyber-security data about the incidents affecting the specific countries [225]. A prototype of the dashboard was developed but the focus of the study was primarily publicly available data from national-level security incident response teams and not security incident response teams within organisations. Based on the literature, the following points informed the design of the dashboard application within the Fortune 500 Organisation:

- An operational-type dashboard was chosen as the preferred design. This is because the dashboard will be used to present real-time information about the state of security incident response within the organisation.
- Marty's [216] concept of 'comparison' will be integrated into the visualisations created for the dashboard application. The visualisations in the dashboard, when viewed by security incident handlers within the organisation, will present trends and changes to various aspects of the security incident response process over time.
- Ultimately, the goal of the dashboard application will be to provide the security incident team within the organisation with information which will help reduce any uncertainty between "what you know and what you need to know" about a particular security investigation [224].

Although numerous security dashboards have been proposed in the literature, very little research evaluates the use of a security incident response dashboard as a tool for identifying data quality issues in security investigation records. Therefore, based on the points identified from the literature, a security incident response dashboard was developed and implemented in the Fortune 500 Organisation in order to meet these requirements. The next section will describe the experiment, which guided this research.

7.2 Experiment Design

An experiment was designed based on the on-going case study in the Fortune 500 Organisation which involved a security incident response dashboard being developed, implemented and evaluated within the organisation. The purpose of the dashboard was twofold. First, it was hypothesised that the dashboard would enhance the transparency of data quality issues within security investigation records in the organisation. Second, it was hypothesised that the dashboard would assist the security incident response team to identify and correct incomplete security investigation records. Three research questions guided this experiment:

1. Does a dashboard assist a security incident response team to identify investigation records which are incomplete and require further detailed information?
2. What fields are considered important to a security incident response team with regards to security investigation record closure?
3. Does a dashboard enhance collaboration within a security incident response team?

In order to answer these research questions, an experiment was designed which involved four stages: 1) identifying and collecting the dashboard requirements; 2) developing the dashboard application; 3) implementing the dashboard within the organisation; and 4) evaluating the dashboard through a data analysis period. The following subsections will describe each stage of the experiment.

7.3 Identifying and Collecting Requirements

Guidance was sought from various books [215, 216, 224] and white papers [226–228] on the topic of developing dashboards and visualising security data as discussed in Section 7.1. In addition, discussions were held with the security incident response team and information security managers within the organisation. These discussions were held to gather the practitioners' requirements regarding the dashboard design. A recurring theme from these discussions was that a 'lightweight' dashboard layout would be essential to any final design. The consensus from these individuals was that "less visualisations would provide more information".

Two information security managers (hereafter referred to as Manager 1 and Manager 2) were consulted over the design of the dashboard. Manager 1 is responsible for security intelligence generation and management within the organisation, while Manager 2 is responsible for everyday operational security management including the security incident response team. As noted in Chapter four, one of the problems identified from the exploratory case study was that 187 out of the 188 investigation records were missing information from one or more fields in the record template. In these cases, there is the potential for security incident handlers to capture further information about a security investigation. It was hypothesised that the dashboard would help identify those records considered to be 'incomplete' and help a security incident handler to identify the fields which needed to be corrected so that the investigation record would be considered 'complete'. In order to define a 'complete' and 'incomplete' security investigation record, discussions were held with both information security managers to gather their requirements.

After discussions with both managers, it was proposed that a security investigation record would be considered 'incomplete' if information from one or more of eleven fields from the security investigation record was missing. The eleven fields considered by the managers to be the criteria to measure an 'incomplete' investigation record, along with their descriptions are presented in Table 7.1. The security investigation record described in this discussion was the same investigation record implemented in the organisation as discussed in Chapter 5.

Several observations can be made with regards to the selection of the eleven fields by management. An initial glance at the choice of fields shows that the 'Investigation Record' was

Field Name	Description
Subject	Subject line which describes investigation record
Category	Category and Sub-Category type of security occurrence
Date Reported	The date the security occurrence was reported to the Security Incident Response Team (SIRT)
Time Reported	The time the security occurrence was reported to the SIRT
Date Discovered	The date the security occurrence was discovered within the organisation
Status	Status of the investigation record (Open/Closed)
Date Opened	The date the security investigation record was opened by the SIRT
Time Opened	The time the security investigation record was opened by the SIRT
Date Closed	The date the security investigation record was closed by the SIRT
Time Closed	The time the security investigation record was closed by the SIRT
Working Hours	The number of working hours that have consumed on a particular security occurrence

Table 7.1: Proposed Incomplete Investigation Record Criteria

excluded from the list of criteria to measure an ‘incomplete’ investigation record. On one hand this could be considered an important field, as it would describe information on how an actual investigation has taken place. On the other hand, the manager’s choice to exclude the field from the selection of ‘closed fields’ highlights that particular metric information regarding the time to eradicate and recover from an incident have been favoured instead. The choice of these metrics support the findings from other case studies [18, 19, 44] that organisations are more focused on eradication and recovery, instead of incident learning. This idea is further supported by the lack of consideration of the ‘Lessons Learned’ and ‘Actions to be Taken’ fields from the criteria. While the information security managers made a decision to exclude certain fields from their criteria, a problem that could arise is that information from these excluded fields could actually be used to learn more about a security incident. However, its capture has not been formally specified in a policy or procedure to the incident response team and therefore, it may be missing from the organisation’s security investigation records.

Manager 2 requested that specific metrics related to the security incident response process are integrated into the dashboard application. The manager requested metrics which would indicate a) how long it takes the security incident response team to resolve a security problem; b) how long a particular type of security problem affected the risk profile of the organisation; and c) how long the incident handlers were taking to manage investigations based on the number of working hours. Three metrics were proposed to meet the manager’s requirements and taken from the relevant literature [179, 229]. These three metrics are:

1. **Time To Resolution** - this metric is calculated as the number of working hours be-

tween when a security problem was reported to the security incident response team and the time to when the investigation record related to the problem was closed. Working hours in this context is defined as 09:00 to 17:00, Monday to Friday.

2. **Time To Reduce Risk** - this metric is calculated as the number of calendar hours between when a security problem was reported to the security incident response team and the time when the investigation record related to the problem was closed. Calendar hours in this context are defined as 00:00 to 23:59, Monday to Sunday.

3. **Time Working on Investigation** - this metric is calculated as the number of working hours consumed by the security incident response team on a particular security problem and is calculated using the information documented in the ‘Working Hours’ field in the revised security investigation record.

The metrics were accepted by Manager 2 and were integrated into the design of the dashboard Graphical User Interface (GUI). Both Manager 1 and Manager 2 also requested further metrics. The calculated results from these metrics would need to provide a graphical visualisation within the dashboard, showing the frequency of security events and incidents, which have occurred within the organisation over a particular period of time. Both managers agreed that four graphical visualisations would need to be integrated into the dashboard GUI. Two visualisations would present frequency information for ‘security events’ and two visualisations would present frequency information for ‘security incidents’. The time periods suggested for the graphical visualisations was 30 days for one part of the visualisation and the previous 12 months for the second part of the visualisation.

During the development stage of the application, the Head of Information Security within the organisation requested that an additional graphical visualisation was integrated into the dashboard. This visualisation would provide information about the number of ‘Open’ and ‘Closed/Incomplete’ security investigation records at a particular point in time. This would involve calculating the number of records in each case and then storing the results so that the graphical visualisation can be generated, presenting the number of respective records over time. The implementation of this visualisation was included into the design of the dashboard GUI.

7.4 Dashboard Architecture

This section describes the various systems which interacted with the dashboard application, how the application collected data and how this data was used to produce the dashboard GUI. The dashboard application consists of a Microsoft Excel 2010 macro-enabled file that was hosted on a network file share within the organisation. The network share is accessible to all individuals within the Information Security unit, including the security incident response team.

The dashboard application was designed and developed using Visual Basic for Applications (VBA version 7) within the Microsoft Excel 2010 file. VBA was the choice of development framework for two reasons. First, VBA has been used in the past within the organisation to develop various security tools for the Information Security unit. Second, VBA was the only development framework readily available. The VBA code for the dashboard application has subsequently been hosted on Github[1]. Some data processing routines have been removed from the published code to protect the organisation's anonymity. The dashboard application consists of three main parts: 1) the back-end 'Raw Data' worksheet which interacts with an IBM Lotus Notes server to extract data; 2) the network component for saving data from the dashboard back to the network file share; and 3) the GUI which presents the various visualisations. The various components and systems, which interact with the dashboard are discussed in more detail below.

7.4.1 IBM Notes Security Incident Response Database

At a high-level, the dashboard application collects information from the security incident response database, calculates various metrics or creates visualisations and then presents these in the dashboard GUI. The security incident response database is hosted on an IBM Lotus Notes server within the organisation. The database itself is stored within an IBM Notes Document Library [230]. Within this document library, individual security investigation records are stored as separate documents.

7.4.2 Dashboard Application

The core of the dashboard application is a 'Raw Data' worksheet, which is used to store the individual security investigation records within the Microsoft Excel file. When the dashboard application is executed, the 'Raw Data' worksheet is cleared and a network session with the IBM Lotus Notes server is initiated. After a connection has been established, the security investigation records are retrieved from the document library on the server. As the individual documents (i.e. investigation records) are retrieved from the server, text saved in the various fields in the investigation record is stored under individual columns within the 'Raw Data' worksheet. These columns represent the individual fields within the investigation record template, with each field in the record having its own column in the Excel worksheet. Therefore, each row from the worksheet corresponded to a particular investigation record.

[1] `http://github.com/grisposgeorge/SIR-Dashboard`

While the data from the investigation records was stored to the ‘Raw Data’ Excel worksheet, several data consistency checks were undertaken. These data checks focused on the format of any dates and times in the investigation record. These were done in order to provide data consistency when using this information for metric calculations. Dates were corrected to the format *dd/mm/yyyy*, while times were stored in the format *hh:mm*. If a date or time was in an incorrect format, the information was corrected and stored in the right format within the worksheet. If any information from a particular investigation record field was missing, the word ‘Unknown’ was used as a substitute. The intention was that security incident handlers, prior to executing the dashboard GUI would view the word ‘Unknown’ within a cell and then attempt to correct this information either in the ‘Raw Data’ worksheet or in the security investigation record itself. Figure 7.1 shows several security investigation records as they stored in the ‘Raw Data’ Excel worksheet. Note, the incident handler’s identities have been redacted to ensure anonymity.

7.4.3 Data Saved to Network File Share

After the dashboard application completed retrieving and sorting the security investigation records from the database into the ‘Raw Data’ worksheet, a ‘snapshot’ of the worksheet was taken. A ‘snapshot’ is a ‘Raw Data’ worksheet, containing the information from the security investigation records, as a separate Microsoft Excel file. This Excel file was saved in a sub-folder in the network share hosting the dashboard application. The ‘snapshot’ file was named using the date the snapshot was taken. Figure 7.2 shows a set of snapshots as stored on the local network file share.

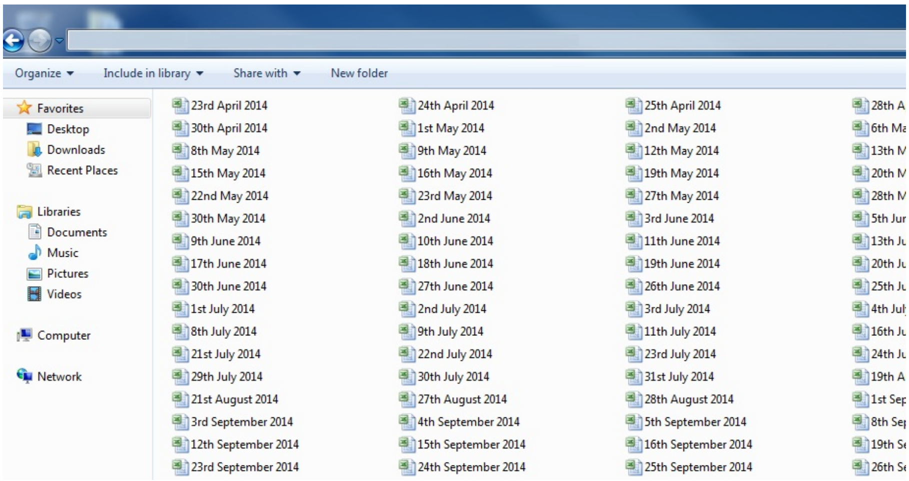

23rd April 2014
30th April 2014
8th May 2014
15th May 2014
22nd May 2014
30th May 2014
9th June 2014
17th June 2014
30th June 2014
1st July 2014
8th July 2014
21st July 2014
29th July 2014
21st August 2014
3rd September 2014
12th September 2014
23rd September 2014

24th April 2014
1st May 2014
9th May 2014
16th May 2014
23rd May 2014
2nd June 2014
10th June 2014
18th June 2014
27th June 2014
2nd July 2014
9th July 2014
22nd July 2014
30th July 2014
27th August 2014
4th September 2014
15th September 2014
24th September 2014

25th April 2014
2nd May 2014
12th May 2014
19th May 2014
27th May 2014
3rd June 2014
11th June 2014
19th June 2014
26th June 2014
3rd July 2014
11th July 2014
23rd July 2014
31st July 2014
28th August 2014
5th September 2014
16th September 2014
25th September 2014

28th A
6th Ma
13th M
20th M
28th M
5th Jur
13th Ju
20th Ju
25th Ju
4th Jul
16th Ju
24th Ju
19th A
1st Sep
8th Sep
19th Se
26th Se

Figure 7.2: Snapshots Stored on Network File Share

‘Snapshot’ file generation was implemented into the dashboard application and therefore, snapshots were created every time the dashboard was used by individuals within the organisation. Furthermore, the author also executed the dashboard application on a working day

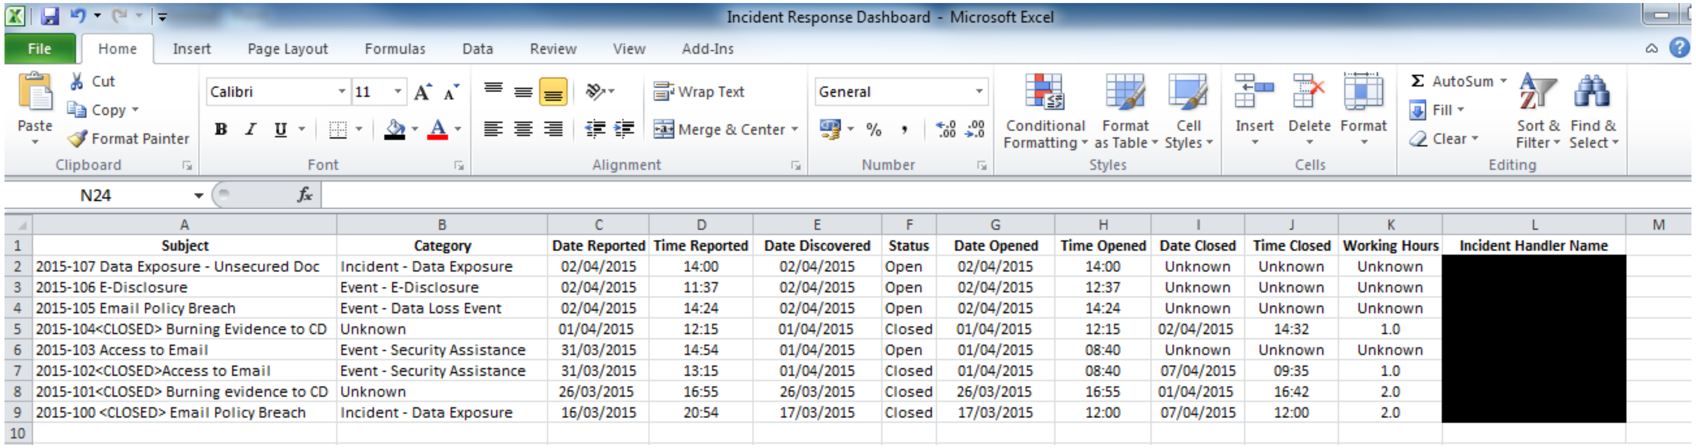

	A	B	C	D	E	F	G	H	I	J	K	L	M
1	Subject	Category	Date Reported	Time Reported	Date Discovered	Status	Date Opened	Time Opened	Date Closed	Time Closed	Working Hours	Incident Handler Name	
2	2015-107 Data Exposure - Unsecured Doc	Incident - Data Exposure	02/04/2015	14:00	02/04/2015	Open	02/04/2015	14:00	Unknown	Unknown	Unknown	[redacted]	
3	2015-106 E-Disclosure	Event - E-Disclosure	02/04/2015	11:37	02/04/2015	Open	02/04/2015	12:37	Unknown	Unknown	Unknown	[redacted]	
4	2015-105 Email Policy Breach	Event - Data Loss Event	02/04/2015	14:24	02/04/2015	Open	02/04/2015	14:24	Unknown	Unknown	Unknown	[redacted]	
5	2015-104<CLOSED> Burning Evidence to CD	Unknown	01/04/2015	12:15	01/04/2015	Closed	01/04/2015	12:15	02/04/2015	14:32	1.0	[redacted]	
6	2015-103 Access to Email	Event - Security Assistance	31/03/2015	14:54	01/04/2015	Open	01/04/2015	08:40	Unknown	Unknown	Unknown	[redacted]	
7	2015-102<CLOSED>Access to Email	Event - Security Assistance	31/03/2015	13:15	01/04/2015	Closed	01/04/2015	08:40	07/04/2015	09:35	1.0	[redacted]	
8	2015-101<CLOSED> Burning evidence to CD	Unknown	26/03/2015	16:55	26/03/2015	Closed	26/03/2015	16:55	01/04/2015	16:42	2.0	[redacted]	
9	2015-100 <CLOSED> Email Policy Breach	Incident - Data Exposure	16/03/2015	20:54	17/03/2015	Closed	17/03/2015	12:00	07/04/2015	12:00	2.0	[redacted]	
10													

Figure 7.1: Worksheet Showing Retrieved Investigation Records

basis. Snapshots generated by the author were stored separately from the snapshots generated by the individuals within the organisation. The purpose behind the ‘snapshots’ was twofold. First, the action of taking daily snapshots of the security incident response database, provided the author with a view of the number of investigation records which were considered ‘Open’ or ‘Closed/Incomplete’ for the given day. Second, the collection of snapshots stored in the network file share was used to calculate metrics about the number of ‘Open’ and ‘Closed/Incomplete’ records and data for a visualisation within the dashboard GUI.

7.4.4 Graphical User Interface Display

The Graphical User Interface (GUI), is a worksheet within the Microsoft Excel file where the visualisations are organised to present a summary of security incident response activities within the organisation. Figure 7.3 shows the dashboard GUI. The following sub-sections discuss the various parts of the GUI.

Traffic light-based listing of security events and incidents

The traffic light-based visualisation, as shown in Figure 7.4 provides the dashboard user with information about the state of the last 30 security investigation records in the security incident response database. The visualisation takes into consideration that an investigation record can be in one of three possible states: Open (*red*), Closed (*green*) and Incomplete (*yellow*).

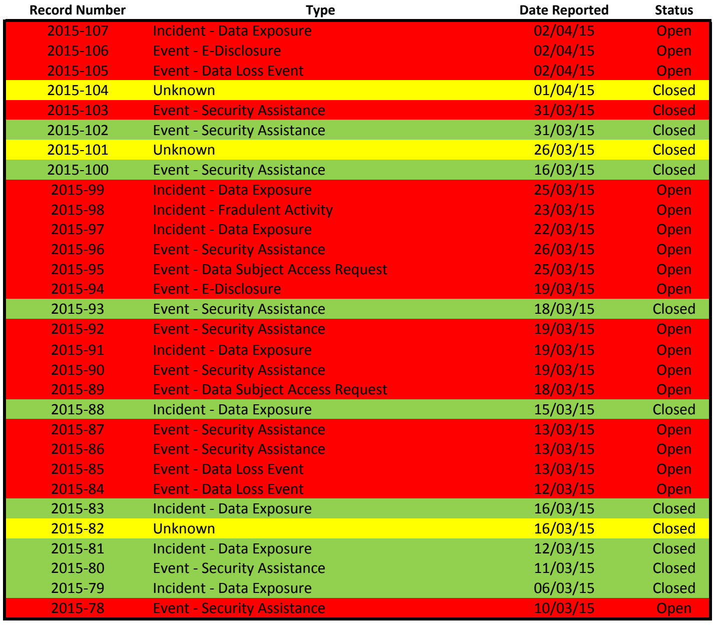

Record Number	Type	Date Reported	Status
2015-107	Incident - Data Exposure	02/04/15	Open
2015-106	Event - E-Disclosure	02/04/15	Open
2015-105	Event - Data Loss Event	02/04/15	Open
2015-104	Unknown	01/04/15	Closed
2015-103	Event - Security Assistance	31/03/15	Closed
2015-102	Event - Security Assistance	31/03/15	Closed
2015-101	Unknown	26/03/15	Closed
2015-100	Event - Security Assistance	16/03/15	Closed
2015-99	Incident - Data Exposure	25/03/15	Open
2015-98	Incident - Fradulent Activity	23/03/15	Open
2015-97	Incident - Data Exposure	22/03/15	Open
2015-96	Event - Security Assistance	26/03/15	Open
2015-95	Event - Data Subject Access Request	25/03/15	Open
2015-94	Event - E-Disclosure	19/03/15	Open
2015-93	Event - Security Assistance	18/03/15	Closed
2015-92	Event - Security Assistance	19/03/15	Open
2015-91	Incident - Data Exposure	19/03/15	Open
2015-90	Event - Security Assistance	19/03/15	Open
2015-89	Event - Data Subject Access Request	18/03/15	Open
2015-88	Incident - Data Exposure	15/03/15	Closed
2015-87	Event - Security Assistance	13/03/15	Open
2015-86	Event - Security Assistance	13/03/15	Open
2015-85	Event - Data Loss Event	13/03/15	Open
2015-84	Event - Data Loss Event	12/03/15	Open
2015-83	Incident - Data Exposure	16/03/15	Closed
2015-82	Unknown	16/03/15	Closed
2015-81	Incident - Data Exposure	12/03/15	Closed
2015-80	Event - Security Assistance	11/03/15	Closed
2015-79	Incident - Data Exposure	06/03/15	Closed
2015-78	Event - Security Assistance	10/03/15	Open

Figure 7.4: Traffic Light-based Visualisation

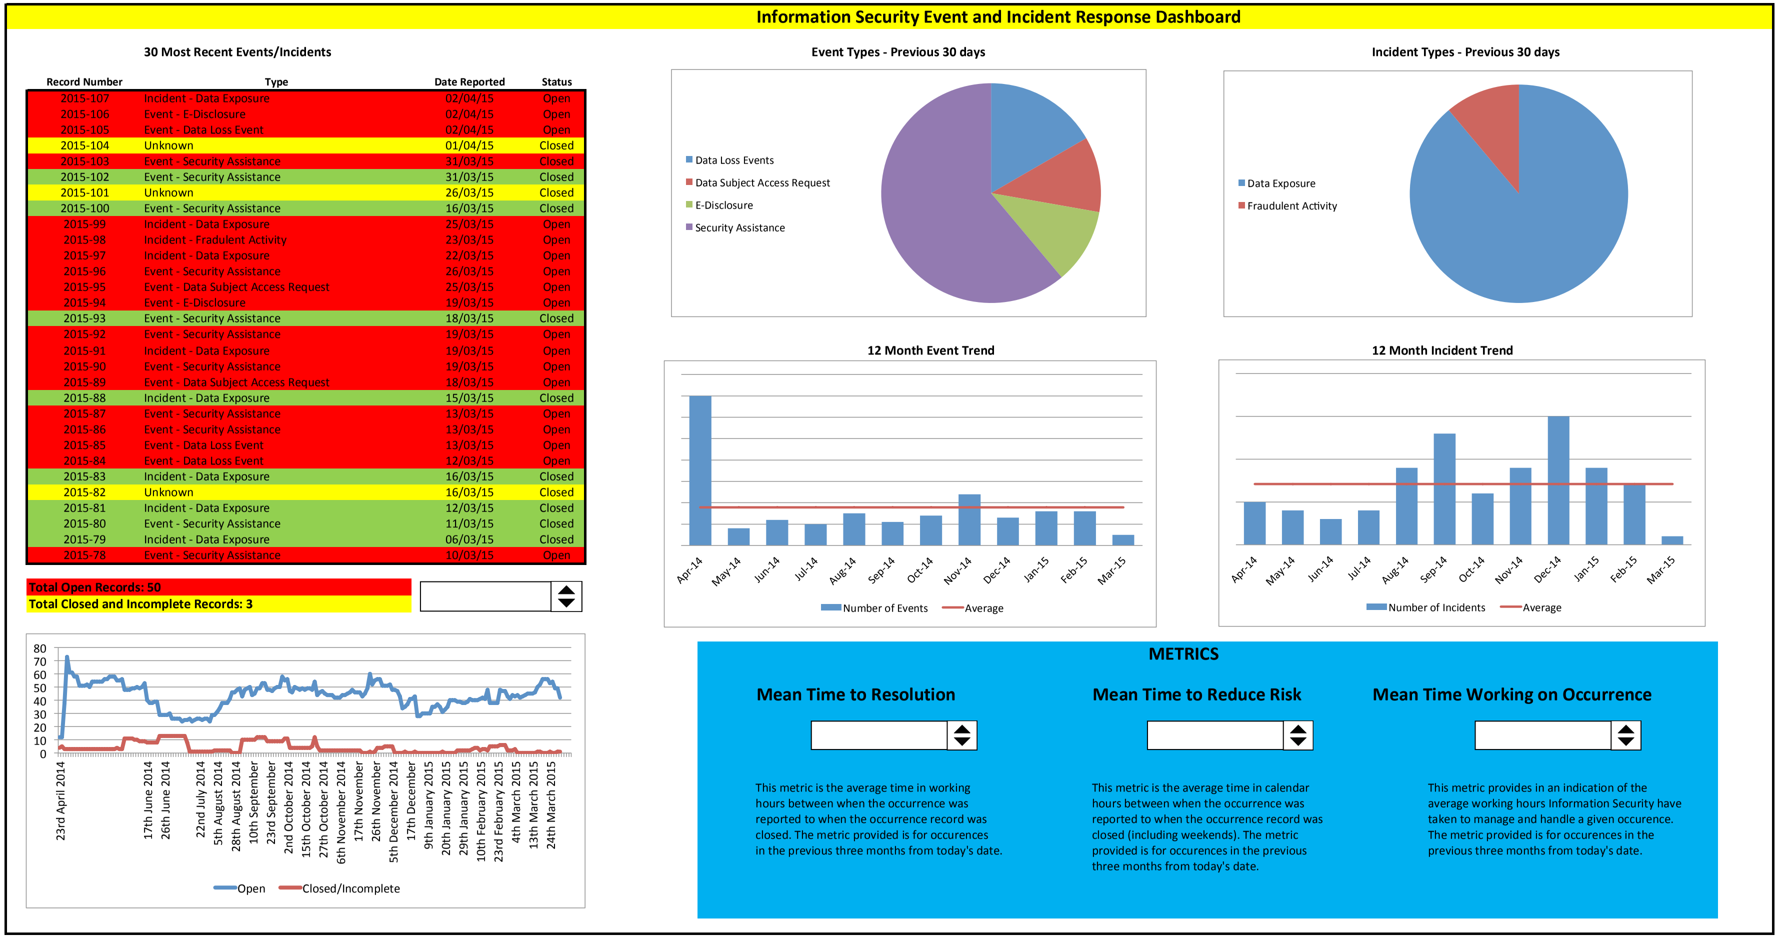

Information Security Event and Incident Response Dashboard

30 Most Recent Events/Incidents

Record Number	Type	Date Reported	Status
2015-107	Incident - Data Exposure	02/04/15	Open
2015-106	Event - E-Disclosure	02/04/15	Open
2015-105	Event - Data Loss Event	02/04/15	Open
2015-104	Unknown	01/04/15	Closed
2015-103	Event - Security Assistance	31/03/15	Closed
2015-102	Event - Security Assistance	31/03/15	Closed
2015-101	Unknown	26/03/15	Closed
2015-100	Event - Security Assistance	16/03/15	Closed
2015-99	Incident - Data Exposure	25/03/15	Open
2015-98	Incident - Fradulent Activity	23/03/15	Open
2015-97	Incident - Data Exposure	22/03/15	Open
2015-96	Event - Security Assistance	26/03/15	Open
2015-95	Event - Data Subject Access Request	25/03/15	Open
2015-94	Event - E-Disclosure	19/03/15	Open
2015-93	Event - Security Assistance	18/03/15	Closed
2015-92	Event - Security Assistance	19/03/15	Open
2015-91	Incident - Data Exposure	19/03/15	Open
2015-90	Event - Security Assistance	19/03/15	Open
2015-89	Event - Data Subject Access Request	18/03/15	Open
2015-88	Incident - Data Exposure	15/03/15	Closed
2015-87	Event - Security Assistance	13/03/15	Open
2015-86	Event - Security Assistance	13/03/15	Open
2015-85	Event - Data Loss Event	13/03/15	Open
2015-84	Event - Data Loss Event	12/03/15	Open
2015-83	Incident - Data Exposure	16/03/15	Closed
2015-82	Unknown	16/03/15	Closed
2015-81	Incident - Data Exposure	12/03/15	Closed
2015-80	Event - Security Assistance	11/03/15	Closed
2015-79	Incident - Data Exposure	06/03/15	Closed
2015-78	Event - Security Assistance	10/03/15	Open

Total Open Records: 50

Total Closed and Incomplete Records: 3

METRICS

Mean Time to Resolution

This metric is the average time in working hours between when the occurrence was reported to when the occurrence record was closed. The metric provided is for occurences in the previous three months from today's date.

Mean Time to Reduce Risk

This metric is the average time in calendar hours between when the occurrence was reported to when the occurrence record was closed (including weekends). The metric provided is for occurences in the previous three months from today's date.

Mean Time Working on Occurrence

This metric provides in an indication of the average working hours Information Security have taken to manage and handle a given occurence. The metric provided is for occurences in the previous three months from today's date.

Figure 7.3: The Dashboard Graphical User Interface

An ‘Open’ record is one where an investigation was likely to still be on-going or the investigation has been completed and the incident handler(s) have not updated the status of the record to ‘Closed’. A ‘Closed’ record is where the status field within the record has been set to ‘Closed’ and that information has been documented in all the fields shown previously in Table 7.1. A ‘Closed/Incomplete’ record is one where the status field within the record has been set to ‘Closed’ but one or more fields from Table 7.1 have not been documented within the investigation record.

Pie and Bar Chart Metric Display

The pie and bar charts provide the dashboard user with a visualisation of the proportion of ‘security events’ and ‘incidents’ documented within the database, based on the different sub-categories. The pie charts display the number of security events and incidents recorded in the database in the past 30 days, while the bar charts display the number of security events and incidents recorded in the database in the past 12 months. Figure 7.5 and Figure 7.6 present examples of the pie and bar charts implemented in the dashboard GUI. Note that numerical values for the pie and bar charts have been removed to protect the anonymity of the organisation.

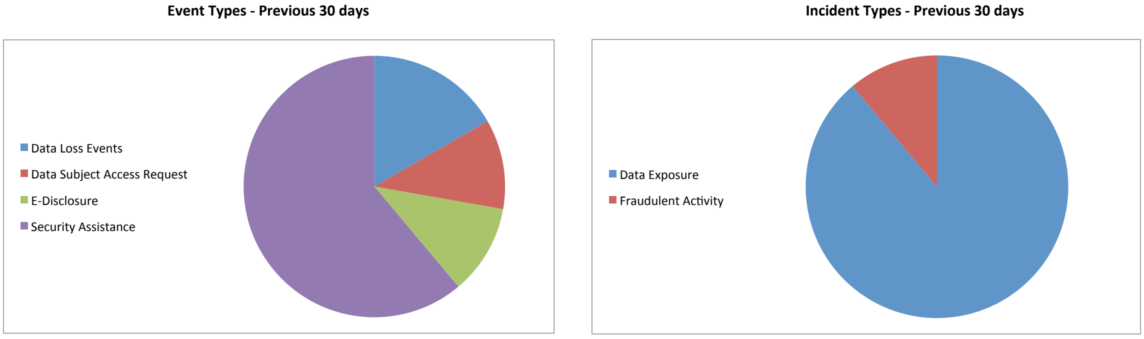

Figure 7.5: Pie Chart Visualisation

The sub-categories used in the pie charts are those which are presented in Section 5.2.1. Information for both the pie and bar charts was calculated by analysing the ‘Date Reported’ field in the security investigation record template. Within the bar charts, a calculation was made to determine how many records were reported in each specific month. In addition, the mean average number of security events/incidents over the twelve month period was also calculated and this is shown as a red line in the two bar charts.

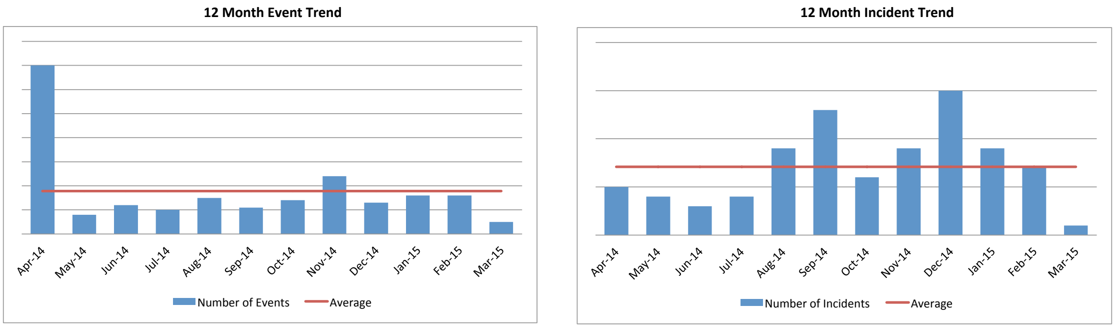

Figure 7.6: Bar Chart Visualisation

Trend Chart Metric Display and Drop-Down Menu

In order to provide the dashboard users with the necessary information to identify incomplete fields, a drop-down menu was added to the dashboard. The purpose of the drop-down menu was reduce the scope of the ‘Raw Data’ worksheet and present to the dashboard user with only the investigation records which are either ‘Open’ or ‘Closed/Incomplete’. Upon selection of the desired option, a new worksheet is then presented to the dashboard user showing the relevant information. In the case of the ‘Open’ records menu option, the dashboard displays all those investigation records within the database, which are considered to be ‘Open’. Similarly, when the ‘Incomplete records’ menu option is selected, a separate worksheet is presented to the dashboard user and all those records considered to be ‘Closed/Incomplete’ are shown. Incomplete fields in this worksheet are identified to the dashboard user with the word ‘Unknown’ in the respective cell. Figure 7.7 shows an extract from the worksheet highlighting the ‘Unknown’ fields to incidents handlers from the ‘Closed/Incomplete’ investigation records.

In addition to providing functionality to identify investigation records which are ‘Open’ or ‘Closed/Incomplete’, a line-chart visualisation was also introduced in the dashboard GUI, as shown in Figure 7.8. This chart visualises the number of ‘Open’ and ‘Closed/Incomplete’ records over a period of time. The data for this chart is calculated using the information stored in the ‘snapshot’ files. The computation of this information involves calculating the number of ‘Open’ and ‘Closed/Incomplete’ records from each snapshot on a particular day. This action is repeated for each ‘snapshot’ file in the network file share and the resulting data identifies the total number of ‘Open’ and ‘Closed/Incomplete’ records from all the snapshot files. The results of this calculation is then used to create the visualisation shown in Figure 7.8.

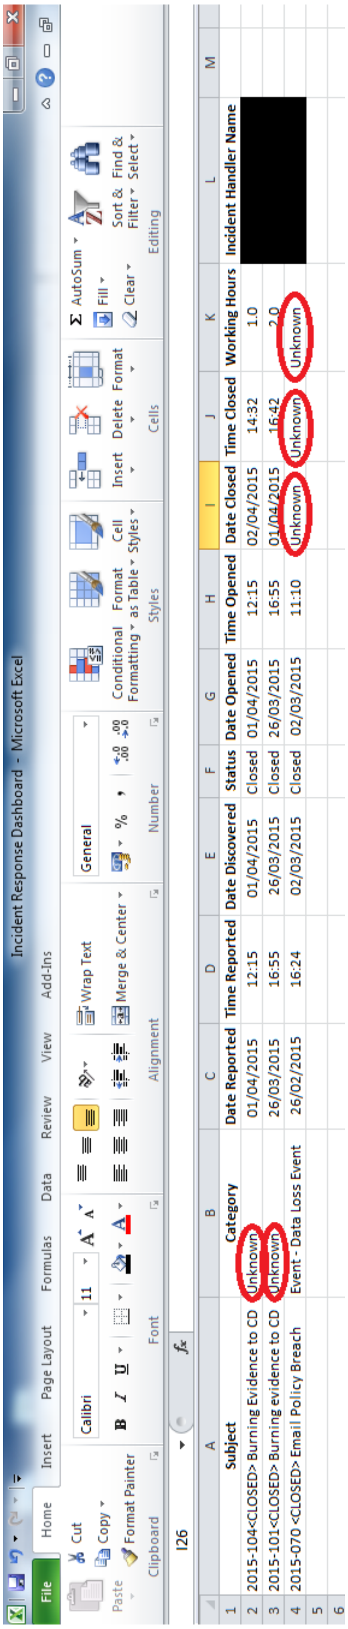

Subject	Category	Date Reported	Time Reported	Date Discovered	Status	Date Opened	Time Opened	Date Closed	Time Closed	Working Hours	Incident Handler Name
2015-104<CLOSED> Burning Evidence to CD	Unknown	01/04/2015	12:15	01/04/2015	Closed	01/04/2015	12:15	02/04/2015	14:32	1.0	[redacted]
2015-101<CLOSED> Burning evidence to CD	Unknown	26/03/2015	16:55	26/03/2015	Closed	26/03/2015	16:55	01/04/2015	16:42	2.0	[redacted]
2015-070 <CLOSED> Email Policy Breach	Event - Data Loss Event	26/02/2015	16:24	02/03/2015	Closed	02/03/2015	11:10	Unknown	Unknown	Unknown	[redacted]

Figure 7.7: Dashboard Extract Showing ‘Incomplete’ Fields

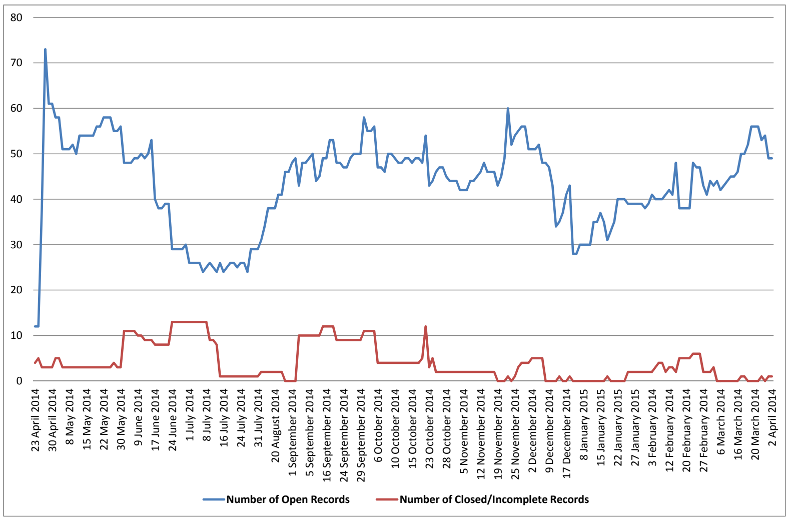

Figure 7.8: Trend Chart Visualisation

Security Incident Response Metric Display

This part of the dashboard GUI is comprised of three drop-down menus. These drop-down menus present the dashboard user with the option to calculate one of three metrics: *mean time to incident resolution*, *mean time to reduce risk* and *mean time working on occurrence*. The options within each drop-down menu allow the dashboard user to select how they would like to apply the metric calculation. The dashboard user can choose to apply the metric by *category* (all events *or* all incidents); by *sub-category* (e.g. all Malware incidents; all data loss events; all E-discovery requests etc.); or *all occurrences* (all security events *and* all security incidents). These options were repeated within each drop-down menu. Figure 7.9 provides an example of the drop-down menu for the mean time working on occurrence metric.

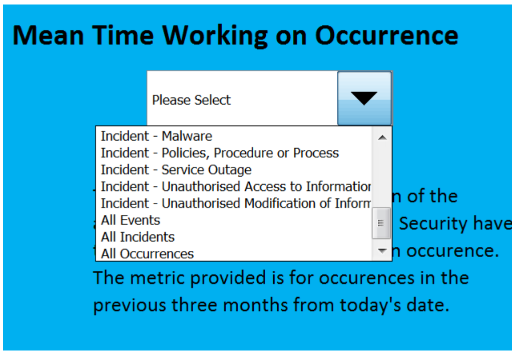
Mean Time Working on Occurrence

Please Select
Incident - Malware
Incident - Policies, Procedure or Process
Incident - Service Outage
Incident - Unauthorised Access to Informatior
Incident - Unauthorised Modification of Inforn
All Events
All Incidents
All Occurrences

n of the
Security have
n occurence.
The metric provided is for occurences in the previous three months from today's date.

Figure 7.9: Example of Dashboard Drop-down Menu

The metrics were calculated using the information in the ‘Raw Data’ sheet for the given occurrence in the past three months. Only investigation records, which were denoted as ‘Closed’ and ‘complete’, were included in the metric calculations. When an end-user selects the type of occurrence they would like a metric for, the dashboard returned the calculation within a message box. Figure 7.10 illustrates an example where the dashboard user has selected ‘All Events’ from the drop-down menu for the ‘Mean Time Working on Occurrence’ metric.

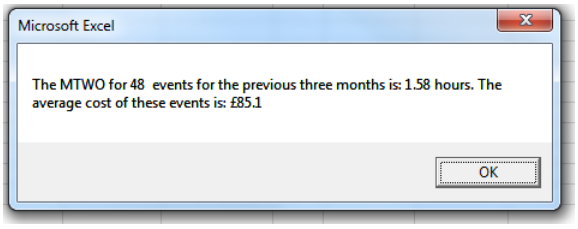

The MTWO for 48 events for the previous three months is: 1.58 hours. The average cost of these events is: £85.1

Figure 7.10: Metric Calculation for Mean Time Working on Events

7.5 Dashboard Implementation

The dashboard application was demonstrated to the organisation in April 2014. Cosmetic changes were made to the dashboard GUI, which included providing a name to the dashboard as seen in Figure 7.3. The application was then made available to the security incident response team and their relevant managers in the middle of April 2014. The dashboard was used in the organisation from the middle of April 2014 until the end of March 2015.

7.6 Data Analysis

As stated earlier, the purpose of the dashboard was twofold. First, it was hypothesised that the dashboard would enhance the transparency of data quality issues within security investigation records in the organisation. Second, it was hypothesised that the dashboard would assist the security incident response team to identify and correct incomplete security investigation records. This section discusses the analysis of data collected from the experiment to investigate these two aims.

At the time of the analysis, the dashboard application had been in use for just under 12 months. The analysis consisted of quantitative data collected through the use of the dashboard within the organisation, as well as qualitative data collected through follow-up interviews with practitioners within the organisation. Quantitative data was collected from both the ‘snapshot’ files generated through the incident handlers use of the dashboard, as well as through the author executing the application on a working day basis. The set of snapshots generated by the incident handlers will be referred to as Snapshot Data Set 1, while the set of snapshots generated by the author will be referred to as Snapshot Data Set 2. The discussions below will make explicit which data set was used in that particular analysis.

In addition to examining the quantitative snapshot data, qualitative data was also collected in the form of semi-structured interviews. The purpose of the interviews was to gather the practitioners' perspective on the dashboard application, as well as to explore their motivations for using the application during everyday security investigations. As discussed earlier, that the interviews were undertaken with seven individuals, six 'Primary Incident Handlers' (PIHs) and one individual who is the Security Incident Response Policy Owner. The interviews queried individuals on several aspects of the dashboard including potential benefits and problems, features most/least used, as well as determining what impact (if any) the dashboard application had on the overall security incident response process.

7.6.1 Dashboard Application Usage

During the application development stage, an attempt was made to create a log that would be used to document the name of individuals who used the dashboard application. This however, proved unsuccessful due to limitations with the development environment. Instead, application usage was examined using the 'snapshot' files. Each time a snapshot was created, the name of the individual executing the dashboard along with the date of snapshot creation were embedded into the snapshot file's meta-data. The meta-data from each snapshot file are then used to analyse application usage. Snapshot Data Set 1 was used for this part of the analysis to examine incident handler dashboard usage.

The results from the application usage analysis show that the dashboard was used a total of 155 times within the organisation. The lowest number of login attempts in a single month was five in April 2014. However, if we consider that the application was only implemented in the middle of April 2014 and therefore was not available to the team for a full month, then May 2014 has the lowest number of login attempts in a single month. The dashboard was only used six times in this month. The highest number of login attempts in a single month was 18 in November 2014. Therefore, the dashboard application was used on average 12.91 times each month, between the middle of April 2014 and the end of March 2015. Figure 7.11 shows the monthly dashboard usage within the organisation.

The graph shows that initially, the security incident response team did not use the dashboard frequently at the start of the experiment. However, as the experiment progressed, application usage continued to increased within the organisation. The dashboard was used five times in April 2014 and six times in May 2014. The results show that application usage within the first two months was 'below the mean average' for the period of the experiment. From this point forward, usage of the application increased and was used nine times in June 2014 and then increased 'above average' to 16 times in July 2014.

A sharp fall in application usage was then recorded in August 2014, with the dashboard only being used nine times in this month. A possible explanation for this result is that incident

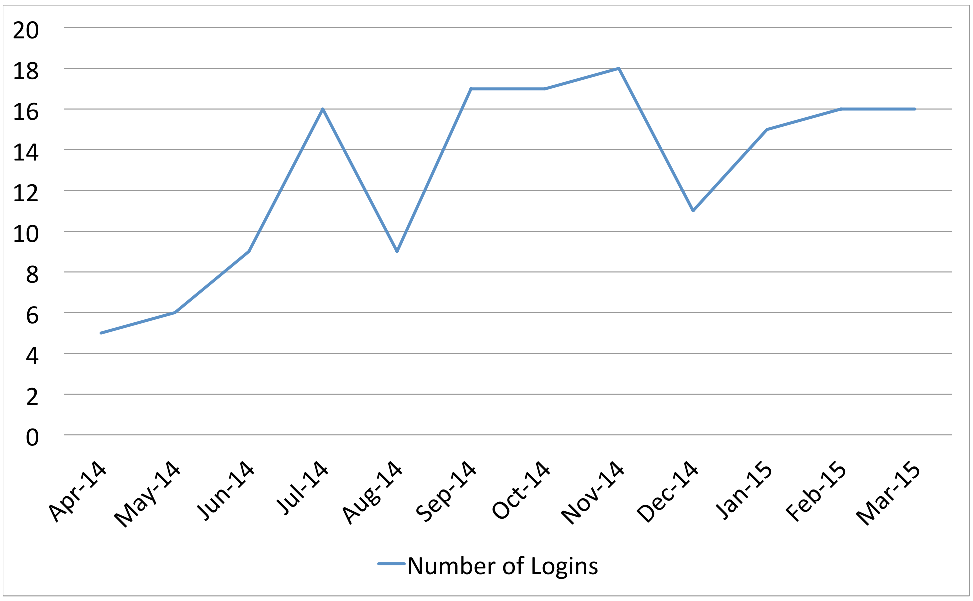

Figure 7.11: Number of Dashboard Application Logins per Month

handlers may have been away on summer vacation. Dashboard usage returns to ‘above average’ in September 2014 and this trend continues until March 2015. From September 2014 to the end of March 2015, the dashboard was used between 15-17 times per month. An exception can be seen in December 2014, where the application was only used eleven times, which is only slightly below the average usage. This drop in usage can again be attributed to the winter vacation period and closure of the organisation over the holiday period.

Examining the number of dashboard logins relative to the size of the security incident response team, presents an alternative view of the information. Figure 7.12 shows the percentage of individuals within the security incident response team who used the dashboard application each month. Recall that the results from the analysis of the dashboard usage showed that the number of login attempts during April 2014 and June 2014 was between five to nine login attempts. Further analysis of the data showed that this translated to 75% - 100% of the individuals within the security incident response team actually using the dashboard during this period. Additional team members joined the security incident response team, first in September 2014 and then in November 2014. In these two months, dashboard usage fell to between 50% and 80% of the team. This decrease in percentage of individuals using the dashboard continued and in December 2014 and January 2015, only 33% of the individuals in the team used the application. After this period, an increase is observed again from February 2015 onwards until the end of the experiment, where 83% of the team used the dashboard.

Alternatively, examining the dashboard logins for each individual per month provides another insight on the information. Figure 7.13 highlights the dashboard usage for the six

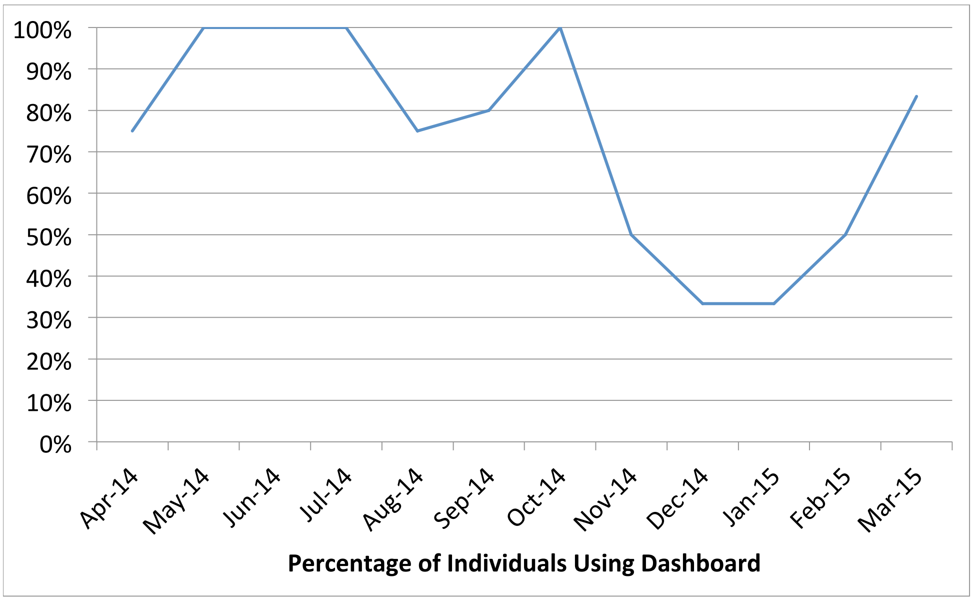

Figure 7.12: Percentage of Individuals Using Dashboard per Month

individuals who accessed the dashboard during the experiment. Several observations can be drawn from this analysis. The graph shows that dashboard usage only increased for three out of the six incident handlers. The graphs for incident handlers one, four and six all show increasingly dashboard usage. However the results for incident handlers two, three and five all showed decreasing dashboard use.

One individual (Incident Handler 1) used the dashboard application every month since it was implemented within the organisation. Incident Handler 6 only joined the security incident response team in September 2014 and therefore, has only used the dashboard application every month since they became a member of the team. The second individual who joined the team in November 2014 (Incident Handler 5) has only used the dashboard application on two occasions, both in the same month that they joined the team. This incident handler later noted in their follow-up interview that their role within the organisation was the forensic examination of systems and therefore they did not need to refer to the dashboard to perform these tasks.

While the initial results indicate that the security incident response team was using the dashboard application, further analysis examined how the number of application logins equate to the total number of investigations undertaken by the team. This analysis involved comparing the total number of investigations each month the dashboard was used with the number of collective logins from the security incident response team. The results of this analysis is shown in Figure 7.14. The results show that initially in April 2014, the number of logins (5) were far less than the number of investigation records (75) at a ratio of one login for every 15 investigation records. This result can be explained as the start of the experiment and that

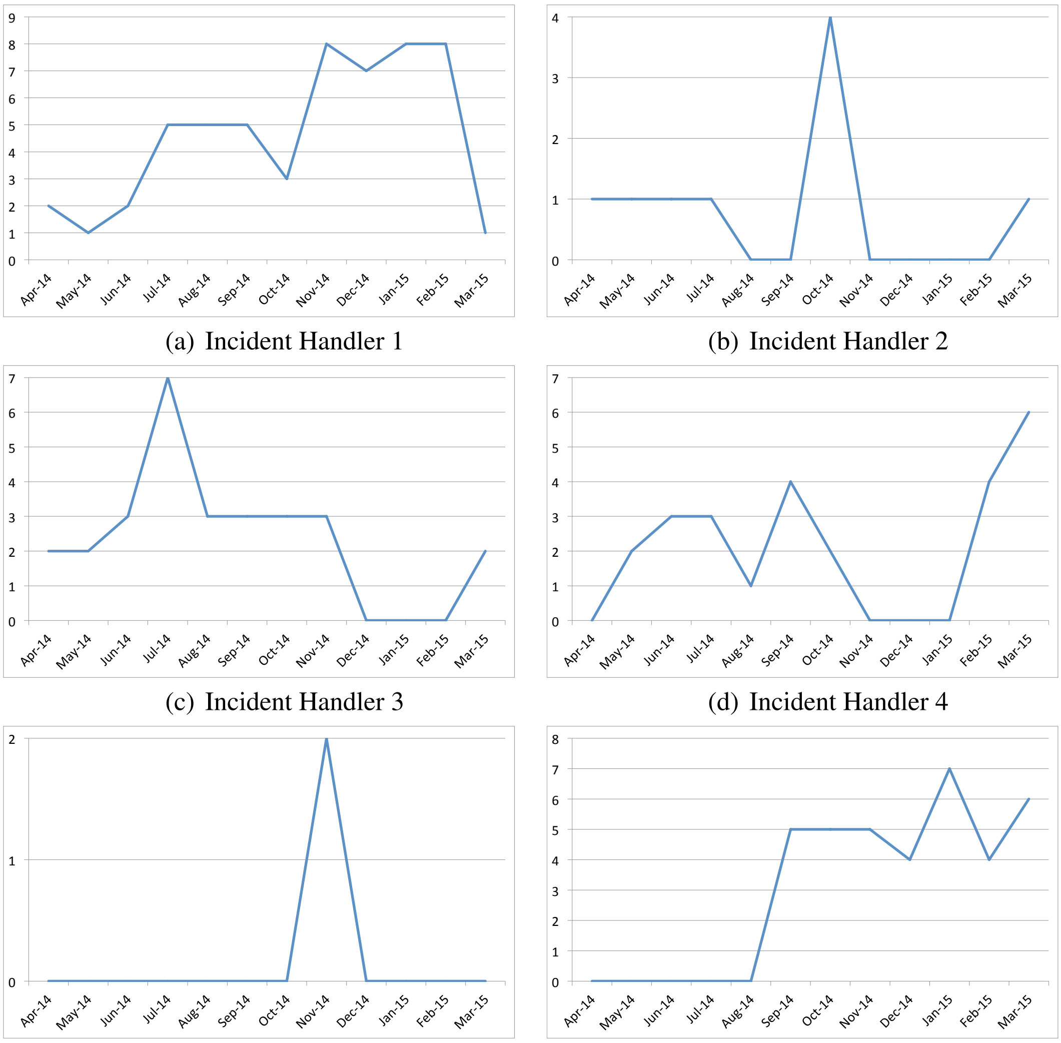

(a) Incident Handler 1

(b) Incident Handler 2

(c) Incident Handler 3

(d) Incident Handler 4

(e) Incident Handler 5

(f) Incident Handler 6

Figure 7.13: Individual Dashboard Usage per Month

the individuals had not yet accustomed to using the dashboard application. Recall also that the investigations during April 2014 included one very large ‘multi-record’ investigation. Therefore, the application was used less frequently even though there were a high number of investigations. During the remaining period of the experiment, the ratio of logins to investigation records decreased. In May 2014, there was one login to every two investigation records, while in the months of July 2014 and March 2015, there were more logins to the dashboard application than the number investigation records. In fact, during March 2015 the dashboard was used 16 times, while only six records were opened during the month. In some months (August 2014, November 2014 and December 2014), the number of login attempts was about half the number of investigations. These findings suggest that the apart from April 2014, the number of dashboard logins by the security incident response team, were approximate to the number of investigations in the security incident response database.

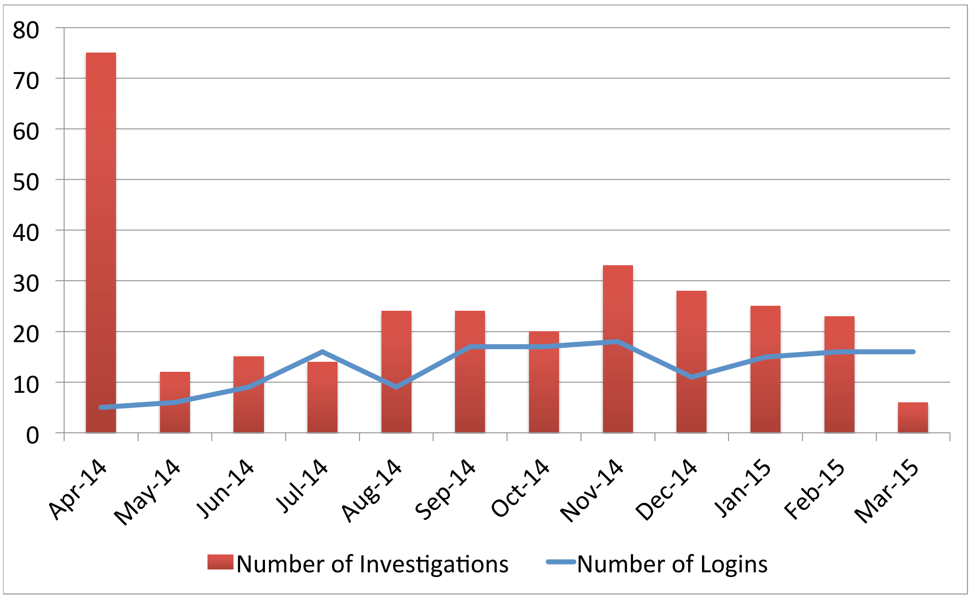

Figure 7.14: Dashboard Logins vs. Number of Investigation Records

The results presented above have shown that the dashboard was used within the organisation. However, the results do not indicate what features were actually used by the individuals. The follow-up interviews were used to examine this aspect. Within the interviews, participants were asked which features of the dashboard application they used most often. Four interviewees stated that the drop-down menu options, which provide information about investigation records which are ‘Open’ or ‘Closed/Incomplete’ were the most used part of the dashboard. The reason provided was because this part of the application helped the security incident response team to identify data quality issues within the security incident response database. This is an opinion that is shared by the Security Incident Response Policy Owner, who agreed that the information about investigation records which are ‘Open’ or ‘Closed/Incomplete’ did appear to help the team the most during the experiment. The manager added that the metric information also provided by the dashboard, has helped to provide visual indicators about the number of threats impacting the organisation.

One individual answered that the most used part of the dashboard application was the ‘Raw Data’ spreadsheet. This individual suggested that this part of the dashboard has helped to provide information that could easily be ‘sliced’ by managers who wanted the data for the organisation’s information security intelligence lifecycle. Recall that this is a process where the organisation collects data and produces intelligence for its management to make strategic information security decisions. One individual stated that they did not use the dashboard enough to determine which features were used the most because their role involved forensic investigations.

Five of the interviewees argued that the incident response performance metrics and to a lesser extent the bar and pie charts, were the least used features in the dashboard application. The reason provided was because the participants believed that this type of information is more suited for managers and security executives and not the security incident response team. Two individuals indicated that they did not use the dashboard enough to determine which features were used the least within the application.

7.6.2 Benefits of Using the Dashboard

When the practitioners within the organisation were queried if they see any benefit to using the dashboard, all seven individuals answered that they see the added value of having the dashboard application within the security incident response team. The interview participants identified three main benefits. Four individuals reported that the dashboard has helped to improve the quality of data documented in the security investigation records through visualising missing information using the Graphical User Interface (GUI). Two individuals indicated that the dashboard has helped to provide information summarising the security incident response landscape for the monthly pack of statistics for senior management. One individual added that another benefit was that the dashboard has assisted with the extraction of information for the organisation’s information security intelligence lifecycle.

In order to investigate how the dashboard GUI has helped to improve the quality of data documented in the security investigation records, an analysis of the data in Snapshot Data Set 2 was undertaken. This analysis focused on identifying and calculating the number of ‘Closed/Incomplete’ records as well as the time taken to correct these incomplete records. The analysis from Data Set 2 showed that 61 ‘Closed/Incomplete’ records were identified from the snapshot files. The minimum number of ‘Closed/Incomplete’ investigations records at any given point during the experiment was zero, while the maximum number of

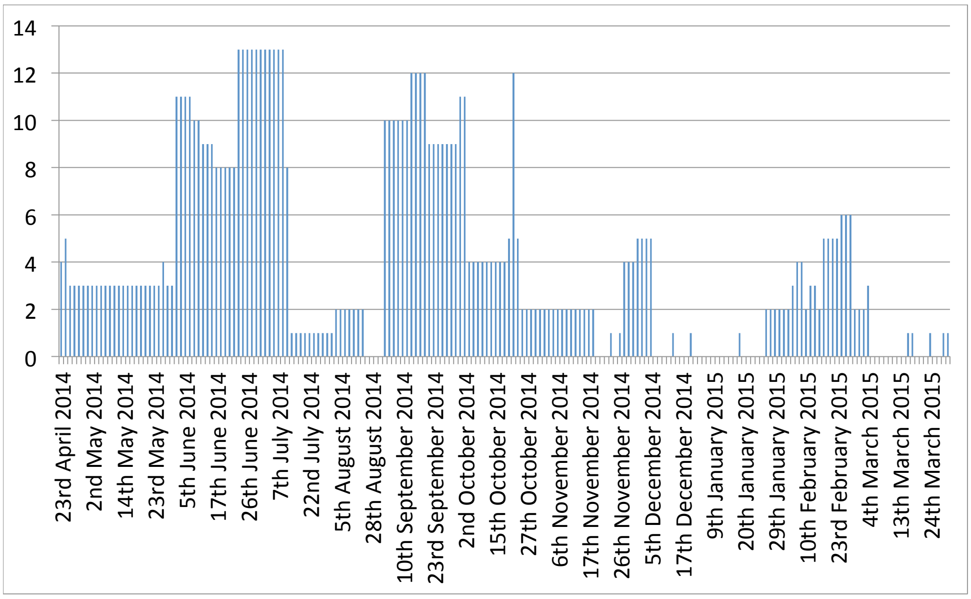

Figure 7.15: Closed/Incomplete Investigation Records Analysis

‘Closed/Incomplete’ records was thirteen. The mean average number of ‘Closed/Incomplete’ records using Data Set 2 was 3.95 records. Figure 7.15 shows the number of ‘Closed/ Incomplete’ records over the period of the experiment. The figure shows that from the end of April 2014 to the start of July 2014, there was a gradual increase in the number of investigation records considered to be ‘Closed/Incomplete’. The number of records increased from four at the end of April 2014 to thirteen records in July 2014. What this increase in records suggests is that although the security incident response team was using the dashboard application, the team was slow to identify and correct ‘Closed/Incomplete’ records.

After this initial increasing trend, the figure also shows that the number of ‘Closed/Incomplete’ records declined from thirteen at the start of July to zero records at the start of September 2014. During September 2014, it was observed that the number of ‘Closed/Incomplete’ records fluctuated between 10 and 12 records. From October 2014 to the end of December 2014, there was a large decrease in the number of ‘Closed/Incomplete’ records, with the mean average during this period being 2.83 records. From January 2015 to the end of March 2015, it can be seen that the number of ‘Closed/Incomplete’ records continues to be about two records. An exception to this trend can be seen in February 2015, where the number of ‘Closed/Incomplete’ records increased to eight and then decreased swiftly back to about two records at the end of this month. From January 2015 to the end of March 2015, the mean average number of ‘Closed/Incomplete’ records was 1.48 records.

The results above suggest that initially the security incident response team was slow to react to identifying and correcting security investigation records which were considered to be ‘Closed/Incomplete’. Therefore, the next stage in the analysis was to examine how long the security incident response team took to correct a ‘Closed/Incomplete’ investigation record. This analysis involved examining the 61 ‘Closed/Incomplete’ investigation records within Snapshot Data Set 2. To calculate the time taken by the security incident response team to correct an ‘Incomplete’ record, involved first identifying the date the investigation record was first visible in the dashboard as ‘Incomplete’ and then, examining the snapshot data to determine the date when the record was considered to be ‘Complete’. The difference between these two dates was considered the time taken to correct an ‘Incomplete’ investigation record. The results of this calculation show that the maximum number of days it took the security incident response team to correct an ‘Incomplete’ record was 57 days, while the minimum time was one day. The mean average number of days to correct a ‘Closed/Incomplete’ record was 14.47 days (just over two weeks). Figure 7.16 presents the number of days that were required to correct the 61 ‘Incomplete’ investigation records. The investigation records in the graph are presented in the order that they appeared as ‘Incomplete’ within the snapshot data.

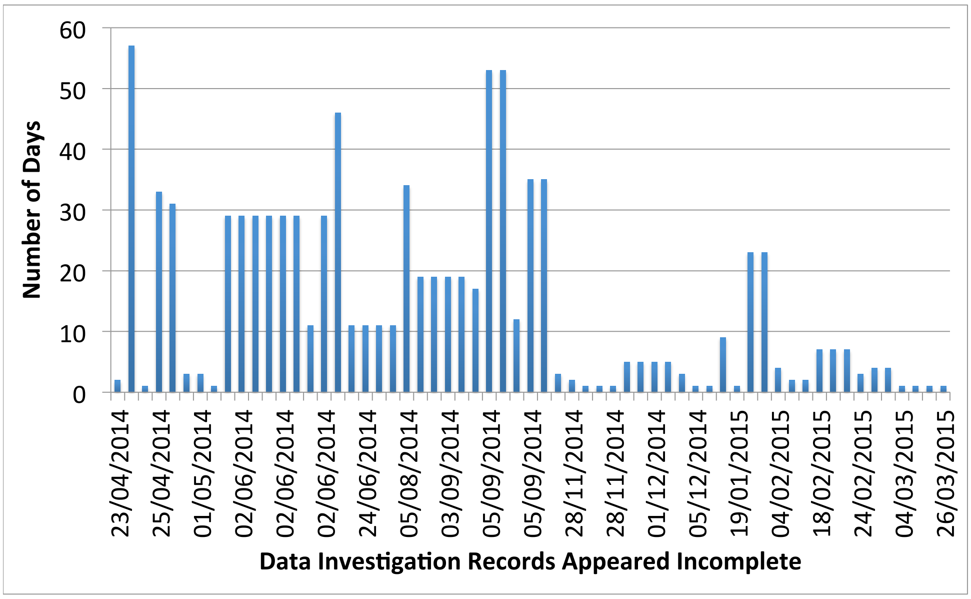

Figure 7.16: Time to React to Incomplete Investigation Records

Figure 7.16 can be interpreted in two parts. The graph shows that initially, the security incident response team was slow to correct an ‘Incomplete’ investigation record. From April 2014 to September 2014, 32 records were found to be ‘Incomplete’. During this period, the mean average time for the security incident response team to correct an ‘Incomplete’ records was 23.44 days. However, as the experiment progressed the time taken to correct an ‘Incomplete’ record decreased. From October 2014 to the end of March 2015, 29 investigation records were identified as ‘Incomplete’ and the mean average time was 4.59 days. This findings show that the mean average time to correct an ‘Incomplete’ record decreased by nearly 81% in the second half of the experiment. However, two investigation records deviated significantly from this average during this second half of the experiment. These two records were identified as 'Incomplete' at the end of January 2015 and were both closed and corrected within 23 days.

Alternatively, the dashboard data can be analysed to examine the relationship between the number of logins by the security incident response team and the mean average number of days to correct an 'Incomplete' investigation record. This analysis involved analysing Data Set 2 and in particular, the 61 'Incomplete' investigation records. The investigation records were first sorted by the month they had appeared 'Incomplete'. The average number of days to correct the 'Incomplete' records was then calculated for each month. The average number of days to correct 'Incomplete' records for each month was then compared with the number of logins in that month by the security incident response team. The results of this analysis is presented in Figure 7.17.

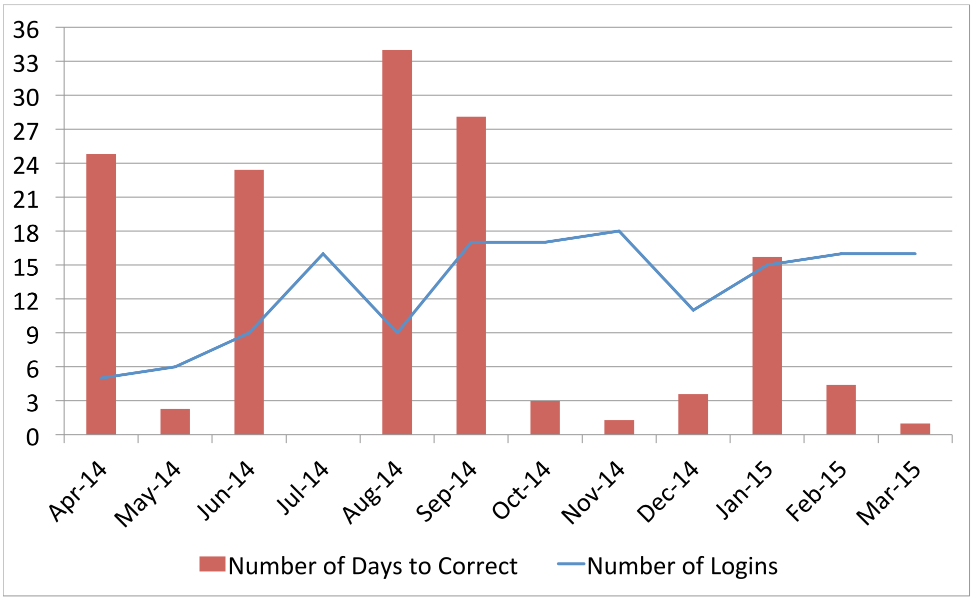

Figure 7.17: Average Time to Correct Record vs. Number of Logins

The results show that in the months where a low number of dashboard logins were recorded (April, June and August 2014), the mean average number of days to correct an 'Incomplete' record increased. In April 2014, five logins resulted in a mean average of 25 days to correct an 'Incomplete' record. Similarly in June 2014, nine dashboard logins resulted in a mean average of 23 days. In contrast, months where a high number of logins were recorded (October - December 2014 and February - March 2015), the average number of days to correct an 'Incomplete' record decreased. For example, in October 2014, 17 logins resulted in an average of three days to correct an 'Incomplete' record and February 2015, 16 logins resulted in an average of four days to correct an 'Incomplete' record. Therefore, the results suggest that the number of logins to the dashboard application had an impact on the number of days it took the security incident response team to correct an ‘Incomplete’ investigation record.

The analysis above was then expanded to examine the findings from the perspective of individual security incident handlers. This expanded analysis focused on the six incident handlers and examined the number of incident handler logins to the dashboard, with the average number of days each incident handler took to correct an ‘Incomplete’ investigation record. The results of this analysis is shown in Figure 7.18. The figure shows that the three incident handlers (Incident Handlers 1,4 and 6) who used the dashboard application more frequently, took less time to correct an ‘Incomplete’ record. In contrast, Incident Handler 2, who used the dashboard less frequently, took longer to correct an ‘Incomplete’ record. However, the results for Incident Handler 3 go against both these trends. This individual extensively used the dashboard however; they still had the highest average number of days to correct an ‘incomplete’ investigation record. Incident Handler 5 had no ‘Incomplete’ records and therefore, no calculation was possible for this individual.

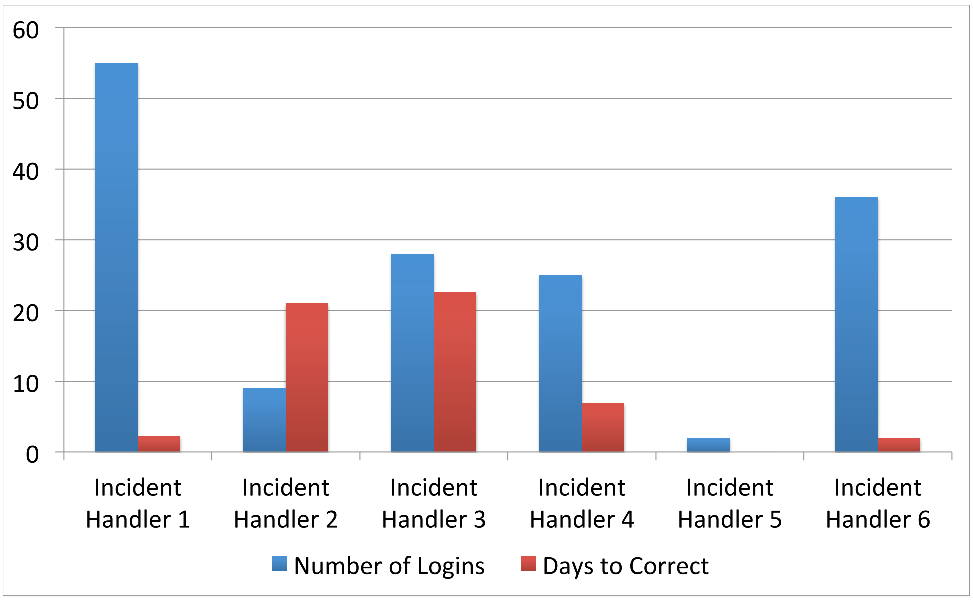

Figure 7.18: Time to Correct Record vs. Number of Logins per Incident Handler

7.6.3 Complications of Using the Dashboard

Five out of the seven interviewees stated that they did not encounter any problems with the dashboard application. However, one of these five individuals did state that the output data presented by the application should be periodically reviewed. The reason provided by this individual was that in order to maintain the value of the application, the data presented to the security incident handlers should be “*real-time relevant information and within the scope of the threat environment currently affecting the organisation*.”

One individual stated that they had a problem using the dashboard application. This problem involved the application running slowly when it was used from remote locations for example, when working from home. This is a problem which cannot be avoided as the connection speeds between the individual's home and the organisation are likely to influence the rate at which the application retrieves records from the security incident response database. However, it may have affected the frequency with which the individual accessed the application. One individual stated that they did not use the dashboard enough to encounter any problems.

7.6.4 Impact of the Dashboard on Security Incident Response Process

All seven individuals unanimously agreed that the dashboard has assisted the overall security incident response process. A variety of answers were provided as potential justifications. These included the dashboard providing a graphical interface to identify 'Incomplete' and 'Open' investigation records; as well as the fields required to complete them; centralised management of information related to open/closed/complete/incomplete security investigation records; visualisations of the threat environment using security incident response data; and clear indicators showing what records have been open for a long period of time (for example records escalated to management) and need to be closed.

When the interview participants were asked if the dashboard application helped with collaboration within the security incident response team, six out the seven participants answered 'Yes'. The six respondents provided a variety of answers on how the dashboard application assisted in collaboration efforts. Three individuals stated that they used the dashboard to discuss which fields need to be completed within an 'Incomplete' investigation record. In fact, one of these six individuals noted they would often execute the dashboard looking for 'Incomplete' records, which belong to other individuals within the team and then notify the individual which records and fields needed to be corrected.

One individual stated that they had used the dashboard application to identify which investigation records needed to be transferred to another incident handler while the initial incident handler was away on vacation. Another individual answered that the dashboard application provided the security incident response team with a platform to discuss investigation records. This individual stated that "*the dashboard is a useful tool to show between ourselves, what records and fields need our attention because as a team, we have become more transparent in what records need more information to be completed, basically now the whole team can see those records which need attention.*" This comment from the individual shows that the dashboard application has resulted in increased data awareness and that transparency has been introduced into the security incident response database from the perspective of 'Open' and ‘Closed/Incomplete’ records. This finding is in line with the findings of previous researchers [231–233], who have evaluated dashboards in software development contexts.

In order to further evaluate the impact of the dashboard on the security incident response process, a quantitative data analysis was undertaken. This analysis focused on the number of records opened during the experiment, as well as a combined perspective of Open/Closed/ Incomplete investigation records. The analysis was conducted using data from Snapshot Data Set 2. The results of the open record analysis examined the number of records which were opened by the security incident response team on each day of the experiment. The results of this analysis are presented in Figure 7.19. Note, the results of the analysis do not include data for the 11th April 2014. On this day, the security incident response team opened 59 investigation records. These investigation records were in reaction to a security incident involving various employees, and each employee was assigned a specific record. These 59 records were excluded because they obscure the identification of any trend with regards to the open investigation records. The results of the analysis show that the workload of the security incident response team has increased during the course of the experiment. In the first half of the experiment, from the middle of February 2014 to the end of August 2014, the results show that the security incident response team created on average 0.58 new investigation records per day. From the start of September 2014 to the end of the experiment at the end of March 2014, the team created 1.01 new investigation records per day, nearly double the number of investigation records. Therefore, the results show that the workload of the security incident response team has increased significantly during the course of the experiment.

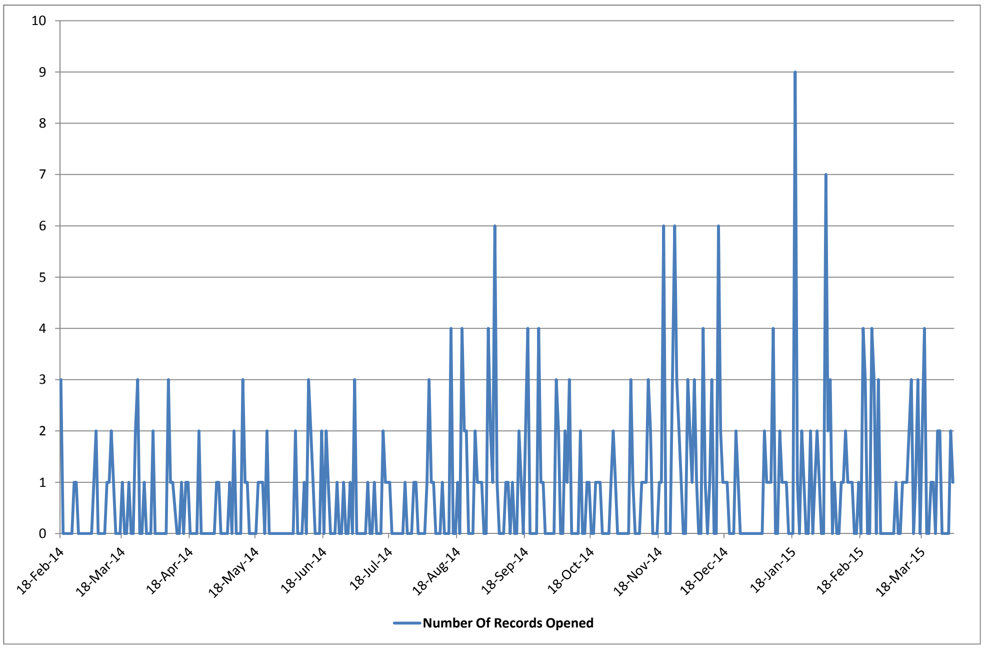

Figure 7.19: Open Investigation Records Analysis

In addition to examining the number of investigation records opened by the security incident response team, a combined perspective of the open/closed/incomplete records presents an alternative view on the information. This analysis involved examining the cumulative number of records that were considered ‘Open’ and the number of ‘Closed/Incomplete’ records on a given day. Figure 7.20 presents the results of this combined perspective analysis. Several observations can be identified from this analysis. The figure shows that initially, as the number of ‘Open’ records decreased (i.e. investigation records were completed), the number of ‘Closed/Incomplete’ records increased. This can be best seen from the data for the 30th May 2014 where the number of ‘Open’ records decreased from 56 to 48 and the number of ‘Incomplete’ records increased from three to eleven. This result suggests that the security incident response team was under pressure with an increased workload and as a result the number of ‘Incomplete’ investigation records increased. Further evidence for this can be seen during the months of June and July 2014. The figure shows that the number of ‘Incomplete’ investigation records during this period went from eight on the 23rd June 2014 to thirteen on the 8th July 2014 and then back to eight records on the 9th July 2014. Only 24 records were considered ‘Open’ during this period. Therefore, as the security incident response team was under less pressure, they attempted to correct investigation records which were identified as ‘Incomplete’ within the dashboard application.

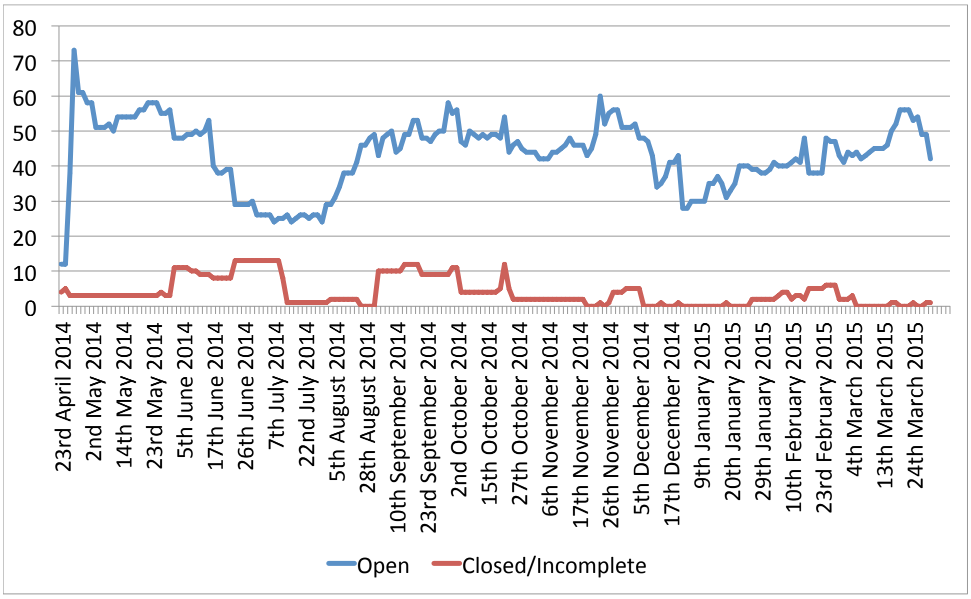

Figure 7.20: Time to React to Incomplete Investigation Records

From the start of October 2014, approximately five months after the dashboard was implemented, the data shows that the security incident response team improved on identifying and closing ‘Incomplete’ records. On the 2nd of October 2014, the number of ‘Open’ records went from 56 to 47 records and the number of ‘Incomplete’ records went from eleven to four. The data also shows that from the end of October 2014, the security incident response team's reaction to 'Incomplete' records improved even further. From this date, the number of 'Incomplete' records does not surpass six for the rest of the experiment. In fact, the mean average of 'Incomplete' records measured during this period was 1.60 records, while the mean average number of 'Open' records was 43.27 records. In summary, this result, together with a reduced time to correct an 'Incomplete' record (from 23.44 to 4.59 days) within the same period, suggests that the dashboard application has had a positive impact on the team's ability to correct 'Incomplete' investigation records in a timely manner.

7.7 Discussion

Overall, the results from the evaluation of the security incident response dashboard shows that for a first iteration design, the outcomes were mainly positive. This section answers the research questions presented at the start of the chapter and discusses the use of the dashboard, the data challenges encountered, as well as the experiment limitations.

Experimental Research Question 1

The experiment shows that a dashboard application has provided a mechanism for the security incident response team in the organisation to identify investigation records which are considered to be 'Incomplete' and require further detailed information. The results from this experiment have shown that 61 security investigation records, which were identified as 'Closed/Incomplete', were corrected by the incident handlers and are now considered 'Complete'. The findings from the exploratory case study presented in Chapter four showed that security incident handlers within the organisation did not appear to go back and correct investigation records that were missing information. Therefore, in the context of the experiment, the dashboard has provided an opportunity for security incident handlers to discuss and analyse 61 'incomplete' investigation record and improve the capture of information in these records.

The analysis of the results from the snapshot files and follow-up interviews supports the suggestion that the dashboard application has assisted with the identification of investigation records which were considered to be 'Incomplete'. Analysing the results further, the data shows that the actual number of 'Incomplete' records has decreased as the experiment progressed. In the initial months of the experiment, the number of 'Incomplete' records was approximately twelve records per day and by the end of the experiment this had dropped to approximately two records per day. Furthermore, as the experiment progressed, the time it took for the security incident response team to identify an 'Incomplete' investigation record decreased from 23.44 days to 4.59 days.

Experimental Research Question 2

In response to the second experimental research question "What fields are considered important to a security incident response team with regards to security investigation record closure?", the answer can be seen in the requirements collected from individuals as discussed in Section 7.3. Two information security managers were involved in the identification of requirements including the definition of an 'Incomplete' investigation record. These managers selected eleven fields from the organisation's security investigation record template, which would need to be completed by incident handlers. These fields included information about the classification of the investigation, dates and times security events/incidents were reported and discovered, dates and times investigations were open and closed, as well as the number of hours consumed working on a specific investigation.

Experimental Research Question 3

The results from the follow-up interviews undertaken during the experiment have suggested that the dashboard has encouraged collaboration between individuals within the security incident response team. Practitioners within the organisation argued that the dashboard application has assisted with collaboration efforts in various ways. These include the incident handlers using the dashboard to discuss which investigation record fields need to be corrected within an 'Incomplete' investigation record, as well as identifying and transferring 'Open' investigation records when one incident handler goes away on vacation and another incident handler has to take over the investigation.

Use of the Dashboard Application

The analysis of the snapshot files has shown that all the individuals within the security incident response team, including their respective managers, have used the dashboard application. Although some individuals have used the dashboard more than others, specific job roles and tasks dictated which individuals accessed the dashboard more frequently than others. One of the security incident handlers identified that their job role focuses on the forensic aspects of investigations and therefore, they did not use the dashboard as frequent as other incident handlers.

Through the analysis of the results from of the experiment, it was also found that security incident handlers periodically executed the dashboard and identified 'incomplete' investigation records and then informed the relevant individuals that their records needed to be corrected. While this was not a requirement and the security incident handlers were free to execute the application when desired, having an individual periodically take on the identification and management of incomplete records appears to have assisted the team. When asked about these actions in the follow-up interview, these individuals confirmed that they would execute the dashboard and specifically identify incomplete records, including those that may not belong to them. These individuals confirmed that they would then share the dashboard findings with the incident handlers who needed to correct their 'Incomplete' records.

Experiment and Data Challenges

While the results from the experiment suggest that the dashboard has provided an opportunity for incident handler to improve the capture of information in security investigation records, numerous challenges and limitations were also observed during the experiment. Although data consistency checks were performed when data was fetched from the IBM Lotus Notes server, when the author periodically analysed the 'Raw Data' worksheet, there was still evidence of 'dirty' data. This data caused numerous problems with the dashboard application, for example run-time errors because information was not being stored in the correct format within the investigation record. The majority of the run-time errors were cause by dates and times being stored in a variety of different formats.

Problems were also observed within the 'Status' field, where words other than 'Open' or 'Closed' were documented within this field. These data quality challenges point to problems with the security investigation record template stored within the security incident response database. The problem is primarily caused by the lack of data checking during the input of information into the record template. As a result, incident handlers can input information in a wrong format, which is then accepted by the record template, because it is based on free-text fields with no data checking. However, when the dashboard fetches this information from the security incident response database, if some form of 'data cleaning' is not undertaken during the processing of this information, the dashboard uses 'dirty data' in the metric calculation and visualisations. As a result, either a run-time error is presented to the user or the graphical visualisations are not generated. In these cases, once the data had been corrected the dashboard continued to function as normal.

The experiment was also limited with regards to the dashboard being designed using requirements from individuals who were no longer employed within the organisation. Recall that two information security managers provided input into the requirements for the design of the dashboard. One of these managers left the organisation half way through the experiment and was replaced by another individual. However, various parts of the dashboard, such as the metrics and security event/incident visualisations, which were created to present information for managerial decisions, were not considered to be as important features by the in-coming manager. This could explain why these features of the dashboard were used the least within the security incident response team. Therefore, if the original manager had remained throughout the experiment, these features could have been used more often during the experiment.

Another limitation which needs to be acknowledged is the extend to which the dashboard itself has improved with the capture of information within security investigation records. While the data collected from the experiment has suggested that when the dashboard was deployed in the organisation, 'incomplete' investigation records were corrected, it can be difficult to isolate if the dashboard was the primary instigator of this change. An alternative view is that the dashboard has assisted with improving group and investigation record awareness within the security incident response team [231, 234]. In this context, awareness refers to the "understanding of who you are working with, what is being worked on, and how your actions affect others [231]. Evidence to support this argument can be seen in the answers from the interviews with individuals in the organisation. Individuals during the interviews suggested that the dashboard allowed the entire team to identify which investigation records need more information and in a sense know who is working on what investigation and what stage the investigation currently resides.

7.8 Summary

This chapter described an experiment to evaluate the use of a dashboard application within a real-world security incident response team. The application was evaluated to determine if a dashboard can help a security incident response team to improve the quality of data within 'incomplete' security investigation records. The experiment also evaluated if a dashboard application can help provide a platform for collaboration within security incident response teams. The initial findings from the experiment show that the dashboard application has assisted with the identification and correction of 'Incomplete' security investigation records. The results have shown that individuals within the security incident response team corrected 61 security investigation records, which were identified as 'Closed/Incomplete. In the context of the experiment, these are 61 investigations records, where further information has been recorded, which if historical trends had continued, may not have been corrected within the organisation. In addition to capturing further information, the results suggest that the dashboard has assisted in reducing the time taken to correct an incomplete investigation record. In the initial months of the experiment, the security incident response team took 23.44 days to correct an 'Incomplete' record, while at the end of the experiment, this decreased to just over four days.

Within the follow-up interviews, practitioners indicated that the dashboard application has helped with collaboration efforts within the security incident response team. Practitioners argued that the dashboard has assisted with collaboration efforts in several different ways.

These include using the GUI interface to discuss which fields need to be corrected within an ‘incomplete’ investigation record and using the dashboard to identify and discuss the transfer of open investigation records from one incident handler to another because of vacation schedules.

Chapter 8

Root Cause Analysis Framework

Chapter three presented evidence from the literature where researchers have argued that security incident response teams should look beyond the immediate causes of security incidents using methods such as root cause analysis. However, there is little empirical research actually investigating the effectiveness of such methods for security incident response. Chapter four identified practitioners' requirements for improved methods and tools to assist in the development of lessons learned in the Fortune 500 Organisation ('the organisation'). This chapter presents a root cause analysis framework for the analysis of security incidents within the organisation. The framework was developed and applied to a historical study of investigation records within the organisation, which highlighted the need for increased data quality in security investigation records in order to examine underlying security issues. The framework was then applied to three 'live' security investigations within the organisation. The results have shown that applying the framework to 'live' security investigations helped to identify underlying root causes as well as improve the quality of data generated from a security investigation.

The chapter is structured as follows. Section 8.1 reviews the motivation for developing the framework and provides an overview of the research method used in the experiment, while Section 8.2 describes the development of the framework. Section 8.3 describes how the framework was applied to a historical set of security incident response investigation records, along with the implications of these results. Section 8.4 describes how the framework was applied to three 'live' security investigations within the organisation, along with the questions asked to obtain a more in-depth root cause analysis. Section 8.5 presents the results from semi-structured follow-up interviews undertaken in the organisation which attempted to gather practitioner opinions on the framework and its use within security incident investigations. Section 8.6 presents an analysis of the 'live' three security investigations undertaken and discusses the significance of the study. Section 8.7 summarises the chapter.

8.1 Motivation and Research Method

Recall (Section 3.4.2) that researchers [16, 126, 127] have argued that organisations should look beyond the immediate causes of a security incident and examine underlying root causes using various tools and techniques. Previously, Johnson described how Violation and Vulnerability (V2) diagrams can be used to assist root cause analysis [126]. While Johnson demonstrated how the V2 diagrams can be used in an analysis of fraudulent transactions involving the Allfirst Bank, he observes that “much remains to be done” to extend and tailor the technique described to support security investigations in a range of different domains [126]. Separately, Stephenson [127] proposed a methodology that uses coloured Petri nets to model attacks of security processes within an organisation. Stephenson [127] then demonstrated how to model a virus infection that has spread through an organisation’s network structure. While these researchers have demonstrated various root cause analysis tools and techniques, very little empirical research investigates the integration and effectiveness of such tools and techniques in real-world security incident response investigations.

In addition to the literature, the interview results from the exploratory case study presented in Chapter four showed that security incident handlers within the organisation called for improved tools and techniques to assist in the development of lessons learned. While these practitioners specifically mentioned root cause analysis as one potential tool, there is currently little evidence to suggest that such a tool would work in the organisation. Therefore, based on increasing calls from academia to examine the underlying root causes of security incidents, coupled with practitioners requests for improved tools and techniques to assist in the development of lessons learned, a root cause analysis framework was designed and evaluated in the organisation. The objective of the research was twofold. First, the research was used to investigate if a root cause analysis framework can help to identify underlying causes of a security incident. Second, the research was used to examine if the root cause analysis framework helps to enhance the quality of data during a security investigation. Three research questions guided this experiment:

1. Can a framework be designed to help guide root cause analysis within security incident investigations?

2. Can it be demonstrated through the use of the framework, that a root cause analysis method does help to identify underlying causes of a security incident?

3. Can it be demonstrated through the use of the framework, that enhanced data can be generated from security incident investigations within an organisation?

In order to answer these research questions, an experiment was designed which involved four stages: 1) design the root cause analysis framework; 2) undertake a study using the framework to evaluate historical investigations within the organisation; 3) use the framework in real-world 'live' security investigations; and 4) evaluate the framework through semi-structured interviews with practitioners.

The first stage of the experiment was to design the framework. Inspiration for developing the framework was sought from various books [15, 235], technical white-papers [236–239], as well as previous academic research [240, 241] into root cause analysis in the nuclear, chemical, manufacturing, aviation and healthcare domains. In addition to analysing the literature, discussions were held with the industrial sponsor in the organisation. These discussions were used to present ideas and recommendations for developing a framework for use within the organisation's security incident response team. The framework was validated by two academics with experience within the information security research community, as well as the Head of Information Security within the organisation. The validation involved the author explaining the framework and feedback being received from the above individuals. The framework, including the suggested changes, is described in more detailed in Section 8.2.

The second stage of the experiment involved applying the validated framework to historical security investigation records in the organisation. The purpose of this analysis was to investigate if a) a root cause can be determined using the framework and only the information documented in a security investigation record and b) in the investigation records where no root cause was established, to determine what questions need to be asked with regards to obtaining more information about the root cause. The historical analysis involved reading the security investigation records and applying the framework to identify a root cause. It must be noted that any root causes described in the investigation records were not included in the analysis and only information about the investigation itself was used for the analysis. In the event a root cause could not be established or an answer to a question could not be determined from the record, the analysis was stopped and the position in the framework was documented along with the question which could not be answered. The purpose was to determine what questions would need to be asked during an investigation to enhance the quality of data for root cause analysis. The historical analysis and the results from the study will be discussed in Section 8.3.

In the third stage of the experiment, the framework was applied to three 'live' security investigations. A 'live' investigation in this case is a security investigation that was reported, investigated and closed by the security incident response team, at the time of the experiment. This involved executing the framework together with one or more incident handlers within the organisation's security incident response team. The author undertook the responsibility of executing the framework and collected the required information from the relevant security incident handler(s). Effectively, the author executed a parallel investigation to that being conducted by the incident handlers. The purpose of this third stage was to evaluate if the framework can be used to develop enhanced lessons learned. In addition to obtaining answers from incident handlers with regards to the questions from the framework, additional information was requested about the answer. This information included who provided the answer to the question, how the answer was derived and what evidence could confirm the answer to the relevant question. The results from the application of the framework to the three 'live' investigations are presented in Section 8.4.

The fourth stage of the experiment involved a qualitative data collection through the use of semi-structured interviews within the organisation. The purpose of the interviews was to determine the practitioner perspective on the benefits/weaknesses of the framework, as well to determine the impact that such a framework could have on the security incident response process. The results from the semi-structured interviews are presented in Section 8.5.

8.2 Developing the Framework

The framework is an extension of the 5-Whys root cause technique. The 5-Whys technique was originally developed for use within the Toyota Motor Corporation's Production System [242]. The technique is an iterative question-asking approach that can be used to determine the root cause of a defect or problem [242]. The foundations of the technique involves repeating the question "Why?" five times, with the answer from each "Why?" question effectively forming the basis of the next question. It is expected that by repeatedly asking "Why?", an investigator 'peels' away layers of issues and symptoms that can lead to a root cause [243]. Typically, five iterations are required to resolve a problem using the 5-Whys technique [243].

The 5-whys technique was selected as a basis for the framework for two reasons. First, researchers have cited the 5-Whys technique as an approach which can be adapted and applied to problems in a variety of domains including manufacturing [244], engineering [245] and healthcare [246]. At the time of writing, no one has attempted to evaluate the 5-Whys method as a root cause analysis technique within security incident response. The second reason the 5-Whys technique was selected was because the approach does not require a considerable amount of training, can be completed without statistical analysis and is considered to be less stressful on participants [243, 247]. This was considered important, following the 'lightweight' practitioner requirements identified in the exploratory case study in Chapter four.

8.2.1 Framework Overview

The framework as shown in Figure 8.1 consists of three parts labelled as Part A-C on the figure below. Labels have been added to aid the discussion. Part A is concerned with providing guidance as to when to execute the framework and undertake a root cause analysis within a security investigation. A prerequisite of the framework is that a security incident response team will have some information about an investigation. The first decision within the framework prompts a security incident response team to determine from its security investigation information, if there 'Has there been an (un)deliberate attempt to violate the confidentiality, integrity and/or availability of an information asset?'. For the purpose of this discussion, an information asset is used in the context of the definition proposed in Chapter 3. The term '(un)deliberate' is used to emphasise that both purposeful and accidental violations should be considered. The purpose behind this question is to help security incident response teams decide when to undertake a root cause analysis. More specifically, a root cause analysis on a security investigation is performed when there has been a breach in either confidentiality, integrity and/or availability of an information asset. This means that investigations such as an E-Discovery Request, which do not involve a breach in confidentiality, integrity and/or availability, would result in a 'No root cause analysis required' decision. As a result, an organisation will not consume valuable resources on investigations, which do not require a root cause analysis.

If the security information from an investigation identifies that there is a breach in either the confidentiality, integrity and/or availability of an information asset then the next step (Part B) is to determine if a third-party provider is involved. According to industrial surveys by Ernst & Young [248] and Deloitte [249], organisations continue to outsource their information technology facilities. As a result, there is an increasing probability that security investigations could involve third-parties. While these third-parties may not own the confidential information targeted in an attack, they might own the underlying hardware and therefore their assistance could be required to conduct any in-depth investigation, including a root cause analysis. Part B prompts a security incident response team to query if a third-party is involved in the investigation. If a third-party is involved, then depending on any pre-existing Service-Level-Agreements (SLAs), the security incident response team needs to determine if it can obtain the investigation data it requires to undertake a root cause analysis. In this context, an SLA is an agreement between an organisation and a third-party where the services provided and expectations of the organisation are formally documented. Although not explicitly part of the framework, if a security incident response team identifies that the SLA is inappropriate (i.e. it does not permit requests of assistance from a third-party for a RCA), then this could be identified during a retrospective as a potential issue for process improvement (as discussed in Chapter 6). If the answer to the third-party involvement is 'Yes' then the next part of the framework (Part C) is initiated.

If the answer is 'No' then a request must be made to the third-party involved to undertake a root cause analysis on-behalf of the requesting organisation. Upon completion of the third-party root cause analysis, the results can be examined and an organisation can ask 'If the

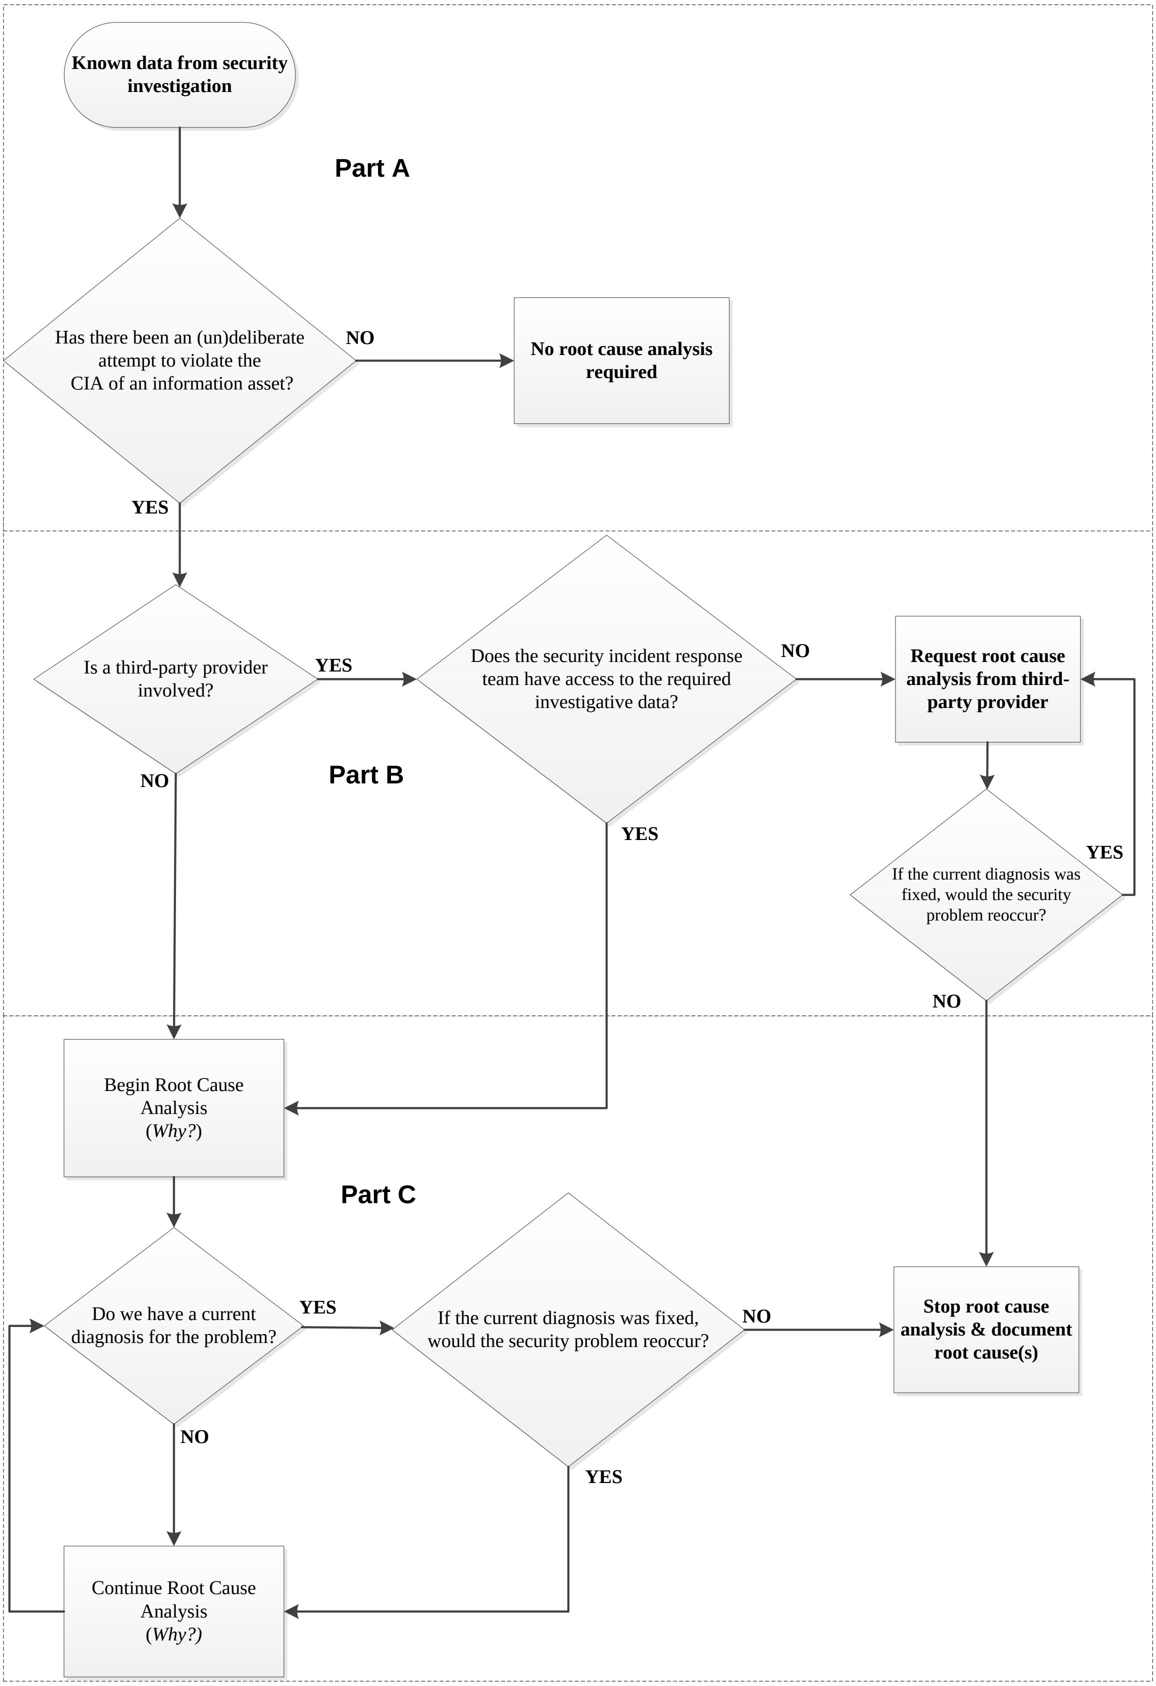
Known data from security investigation
Part A
Has there been an (un)deliberate attempt to violate the CIA of an information asset?
NO
No root cause analysis required
YES
Is a third-party provider involved?
YES
Does the security incident response team have access to the required investigative data?
NO
Request root cause analysis from third-party provider
NO
Part B
YES
If the current diagnosis was fixed, would the security problem reoccur?
YES
NO
Begin Root Cause Analysis (Why?)
Part C
Do we have a current diagnosis for the problem?
YES
If the current diagnosis was fixed, would the security problem reoccur?
NO
Stop root cause analysis & document root cause(s)
NO
YES
Continue Root Cause Analysis (Why?)

Figure 8.1: The Root Cause Analysis Framework

current diagnosis was fixed, would the security problem reoccur?' If the answer is 'Yes' then a root cause analysis is again requested from a third-party. If the answer to the query is 'No' the framework stops and the root cause(s) are documented for the organisation's future reference.

In the event a third-party is not involved or a security incident response team has access to the required security data to undertake a root cause analysis, Part C can be initiated. This part of the framework is where the question "Why?" is asked and the answer from each question effectively forms the basis of the next question. When an answer from a "Why?" question is formulated, a security incident response team needs to consider: "Do we have a current diagnosis for the problem?" If the answer is 'No', the next "Why?" question is asked using the answer from the previous question. However, if the answer is 'Yes' then the next question asked is: "If the current diagnosis was fixed, would the security problem reoccur?" If the answer to this query is 'Yes', then the root cause analysis continues and the next "Why?" question is asked. However, if the answer to this query is 'No', it is considered that the current diagnosis, if fixed, would prevent a security problem from reoccurring. With this indication, the iteration of framework stops and prompts the documentation of the identified root cause(s).

The framework does not specify a particular number of 'Whys?' that must be asked, or a limit on how many can be asked. Unlike other root cause analysis approaches, which do not define when to stop a root cause analysis, the framework stops when a solution is identified by a security incident response team, which if fixed, would prevent a security problem from reoccurring. This feature of the framework helps to guide the conclusion of a root cause analysis, when a solution(s) has been found and without consuming additional resources. The next section will describe how the 'Why?' questions can be documented using a tabular worksheet.

8.2.2 Framework Tabular Sheet

A pragmatic decision was made to use a tabular sheet for the purpose of documenting the root cause analysis using the framework. This decision was based on literature which argued that visualising a root cause analysis through a table can help to identify any links between incident causes and the ultimate root cause(s) [243]. The tabular sheet used within the framework is an extension of a Microsoft Excel table proposed by Bulsuk [250], who used it to conduct a 5-Whys analysis. The table used in the framework is shown in Figure 8.2.

Prior to undertaking the actual root cause analysis, a security incident handler should look to define the problem statement that the analysis will look to address as well as the name of the investigation. The root cause analysis itself is documented in a series of "Why?" columns.

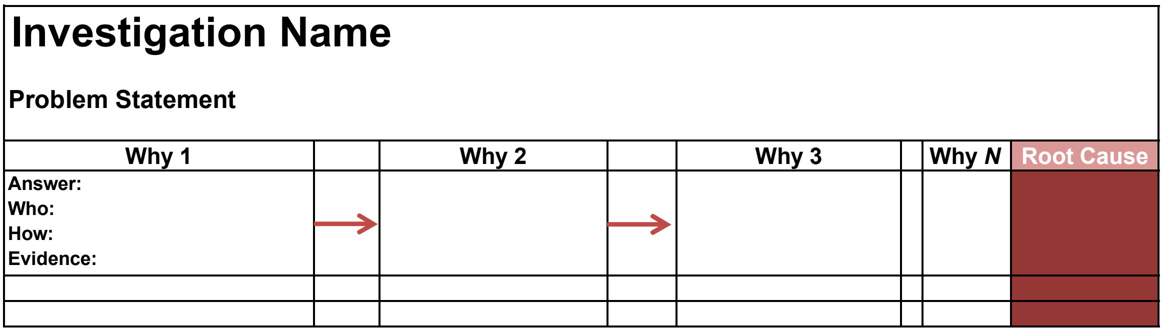

Investigation Name

Problem Statement

Why 1		Why 2		Why 3		Why *N*	Root Cause
Answer: Who: How: Evidence:	→		→				

Figure 8.2: Table Used to Document Root Cause Analysis

Each "Why?" column is used to document one or more answers (causes) to a particular "Why?" question. The 'Root Cause' column contains the final root cause(s) identified from the various "Why?" questions. The first "Why?" question attempts to directly address the problem statement. If more than one answer to a "Why?" question exists, then each answer is documented on a separate line under the particular column. Each answer to a "Why?" question is then handled separately in the next iteration and forms the basis of the next "Why?" question. In subsequent iterations, "Why?" is asked for each answer under the previous column and the answer to this question is placed under the next "Why" column (in this case "Why 2"). The iteration continues until a root cause is found, which if fixed, would prevent the security problem from reoccurring.

Additional provenance information is also documented for each answer to a "Why?" question. This information includes "Answer" which is the actual answer to the "Why?" question asked; "Who?", the job title or name of the individual who provided the answer to the "Why?" question; "How?", what method did the individual use to provide the answer; and "Evidence" what evidence can the individual provide justifying their answer to the "Why?" question. The purpose of documenting this additional information was to enhance the credibility and validation of the root causes identified using the framework. If an information security manager or any other individual within a management role wanted to audit the root causes identified, this additional information could help with this validation. The next section will be used to discuss how a historical study of the organisation's security investigation records was undertaken as well as the results from this historical study.

8.3 Historical Study

The second stage in the experiment involved applying the framework to historical security incident response investigation records in the organisation. There are two parts to this section. The first part presents a worked example, which is used to illustrate how this historical study was undertaken. The second part presents and analyses the results from the historical study.

Note, the security investigation records used in the historical study have been anonymised to protect the identity of the studied organisation.

8.3.1 Historical Study Worked Example

The worked example describes an investigation about a virus infection within the organisation. Below is a narrative summary of the content of the actual security investigation analysed as part of the worked example:

1. The Intrusion Detection System (IDS) alerted the Helpdesk that a laptop was infected with the myDoom.F virus, the laptop was then cleaned.

2. The Helpdesk were then notified by employees via email that files were being deleted by the virus on the Local Area Network (LAN).

3. Information Security Analyst 1 advised the Helpdesk that this should not be the case and that the virus should not be deleting any files.

4. Information Security Analyst 2 confirmed that files were indeed being deleted, further analysis identified that a variant of the myDoom.F virus does indeed delete files.

5. A scan of the LAN was performed for Internet Protocol (IP) addresses with open port TCP/UDP 1080 (the port the virus leaves open if it is on a laptop/computer).

6. Three laptops that were detected as infected were cleaned. It was documented in the investigation record that two out of three laptops were not running the latest .dat antivirus update.

7. Subnets with affected nodes were then isolated from the LAN.

8. Discussions were held around how the virus arrived on organisational LAN. It was suggested that the virus arrived via email servers, which did not have latest .dat antivirus updates.

9. Antivirus signatures (.dat files) updates 'rolled-out' across the organisation and these updates were now blocking the virus.

10. Further virus infections reported overnight, network shares were then disabled to prevent further infection and to help identify the nodes responsible for causing the new infections.

11. Various restoration actions were then undertaken.

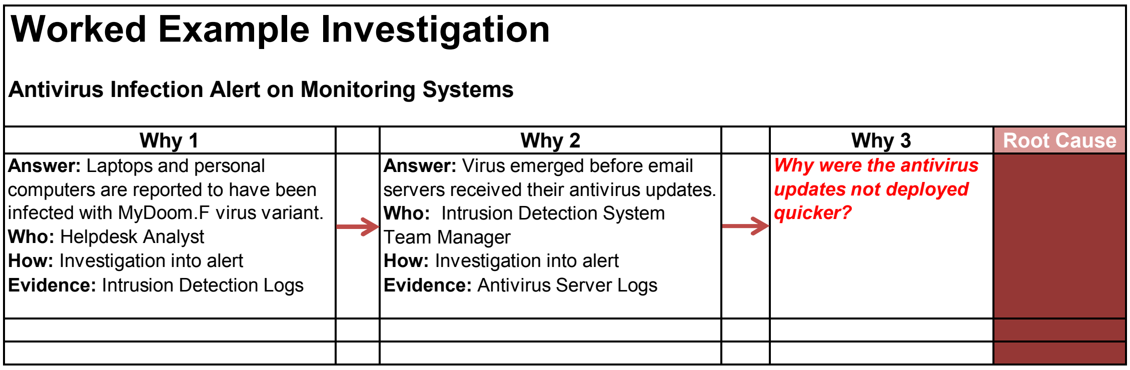

Worked Example Investigation

Antivirus Infection Alert on Monitoring Systems

Why 1		Why 2		Why 3	Root Cause
Answer: Laptops and personal computers are reported to have been infected with MyDoom.F virus variant. **Who:** Helpdesk Analyst **How:** Investigation into alert **Evidence:** Intrusion Detection Logs	→	**Answer:** Virus emerged before email servers received their antivirus updates. **Who:** Intrusion Detection System Team Manager **How:** Investigation into alert **Evidence:** Antivirus Server Logs	→	***Why were the antivirus updates not deployed quicker?***	

Figure 8.3: Worked Example Root Cause Analysis Table

The framework was then applied using only the information available in the security investigation record.

'Has there been an (un)deliberate attempt to violate the confidentiality, integrity, or availability of an information asset?' - **Yes**

'Is a third-party involved?' - **No**

Figure 8.3 presents the framework worksheet for the worked example. The table was constructed using information only in the relevant security investigation record, which was summarised above. As Figure 8.3 shows, the problem statement extracted from the investigation record was: "Antivirus infection alert on monitoring systems". This problem statement was then used as the input to the first 'Why?' question: "Why was the alert notifying the Helpdesk team?" One cause (the answer) was identified: "Laptops and personal computers have been infected with a variant of the myDoom.F virus". This answer was provided by a Helpdesk Analyst, through their investigation into the alert and supported by logs from the IDS. The cause identified in the first 'Why?' question is then mapped under the "Why 1" column in the table. The second "Why?" question uses this answer as the basis for the next question: "Why was the virus present on the laptops/computers?". The cause identified from the second question is "The virus likely arrived via email before the email servers received their antivirus updates". This answer is provided by the IDS Team manager through their investigation into the alert and supported by logs from the antivirus server. The cause is then mapped to the "Why 2" column and is then used as the basis of the next question, "Why were the antivirus updates not deployed quicker?". The answer to the third "Why?" question can not be found in the investigation record and as a result, this record is considered to have stopped at the third "Why?". The actions described in the worked example above were then repeated in a historical study involving 530 investigation records in the organisation's security incident response database, which were recorded in the database from November 2003 to November 2014.

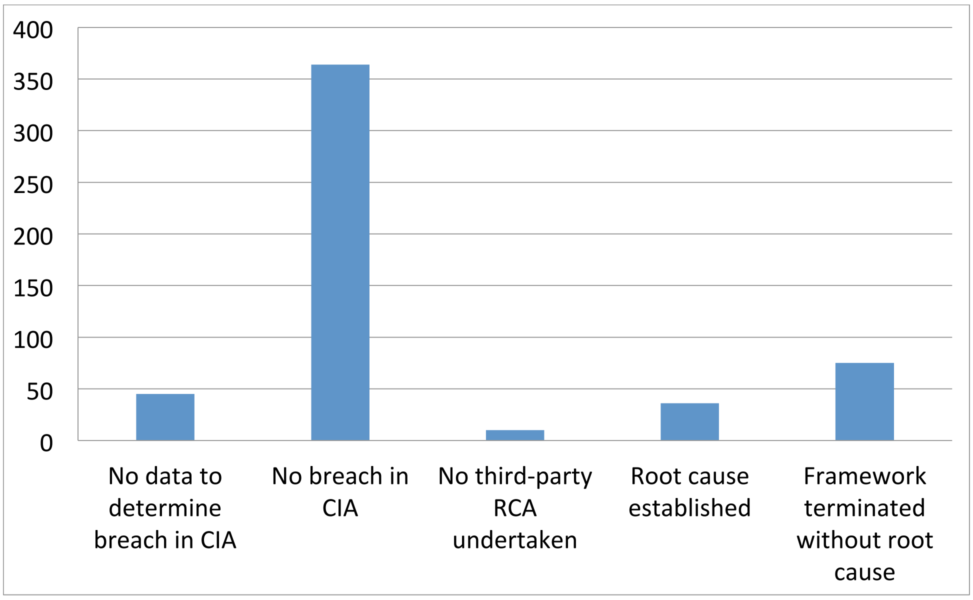

Figure 8.4: Summary of Results from Historical Root Cause Analysis

8.3.2 Historical Study Results

The historical study was undertaken between November and December 2014. The study involved examining the 530 investigation records and applying the framework to these records. The purpose of the study was to investigate if these investigation records could have benefited from a more in-depth analysis to look beyond the immediate causes of the incident and examine underlying root causes. A deliberate decision was made not to involve security incident response employees and only use the information available in the security investigation records. This was done because the purpose of the study was to identify the quality of information in security investigation records and if this information can be used to examine underlying root causes. This section presents the results of the historical study analysis. As shown in Figure 8.4, the results show that out of the 530 investigation records, 364 records were found to describe problems which did not result in the breach of confidentiality, integrity and/or availability of an information asset. Therefore, the analysis of these 364 records terminated at the first check. A further 45 records lacked sufficient information to determine whether an attempt to breach the confidentiality, integrity and/or availability of an information asset occurred, so again the analysis of these 45 records terminated at the first check.

10 out of the 530 analysed investigation records terminated at the third check-box. This is because the investigation records indicated that a third-party was involved in the security incident, but the security incident response team did not have access to the required data to undertake a root cause. In this case, according to the framework, the third-party would have had to undertake the root cause analysis. However, no documented information within these

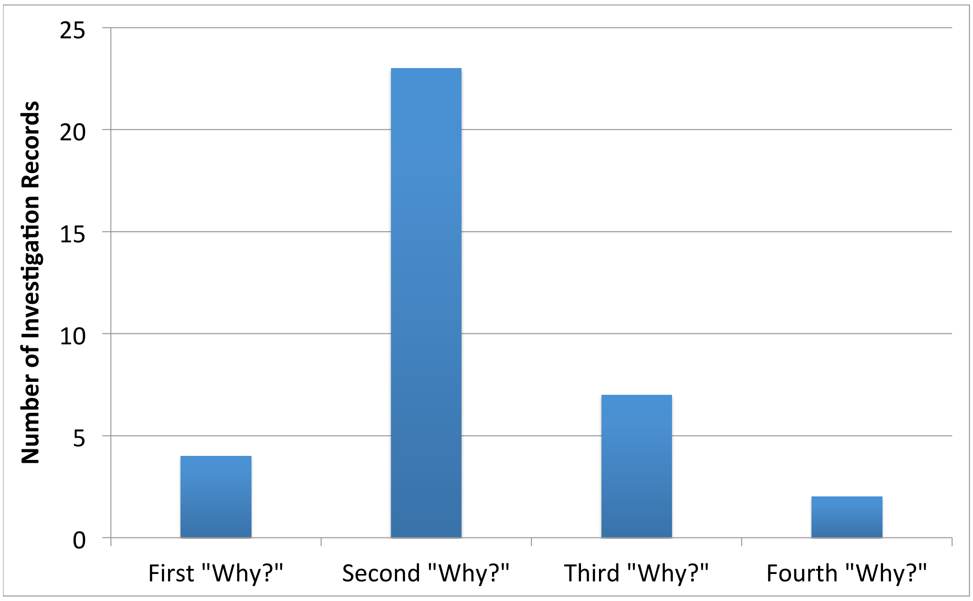

Figure 8.5: Summary of Terminated Iterations for Discovered Root Causes

ten investigation records suggested that a root cause was undertaken by these third-parties.

A root cause was identified in 36 (7%) out of the 530 analysed investigation records. In these 36 records, the framework terminated because a root cause was found, which if fixed, would have prevented the security problem from reoccurring. These 36 causes were identified at various iterations in the framework. Figure 8.5, summarises at which "Why?" question the 36 records terminated because a root cause was found. In four out of the 36 records, the root cause was found in the first "Why?" question. In these four cases, two causes were identified as laptops stolen from the employee's homes and two causes were identified as laptops being lost while employees were away on business for the organisation. Based on the assumption that the employees were permitted to remove the laptops from the organisation, in all four cases the root cause was neither the employees nor the organisation's fault but the result of a particular set of circumstances. In 23 out of the 36 records, a root cause was identified in the second "Why?" question. This means that for nearly 64% of the 36 investigation records the root cause was identified in the second iteration of the framework. Within seven out of the 36 investigation records, a root cause was identified in the third "Why?", while within two investigation records the root cause was identified in the fourth "Why?" question.

The results from the above analysis were then compared with the root cause information documented in the relevant investigation records. This involved examining the 36 investigation records and comparing the root cause established by the security incident handler with the results from the historical analysis. The results of this comparison showed that within nine out of the 36 investigations a different root cause analysis was identified using the framework. In these nine cases, it was observed from the analysis of the investigation records that

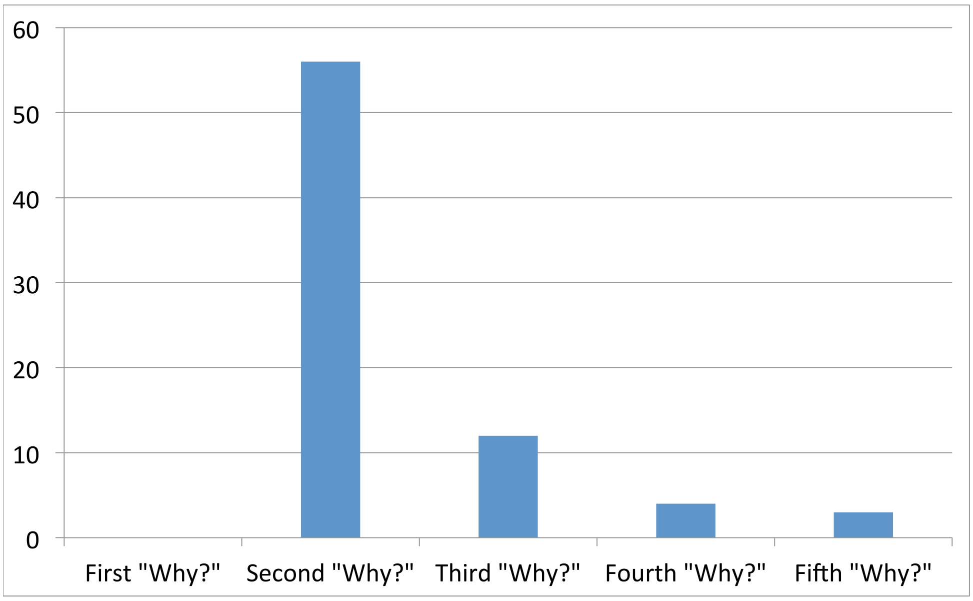

Figure 8.6: Summary of Terminated Iterations for Undiscovered Root Causes

the incident handlers stopped their investigation prior to identifying the actual underlying cause. For example, one investigation record described an incident about organisational data being copied onto an USB flash drive. The underlying cause documented in the investigation record was a ‘Breach of Policy’. However, the root cause analysis findings for this investigation record indicted that the laptop which was used to copy the data did not have its USB ports disabled and this was because the laptop was not ‘locked-down’. This finding suggests that security incident handlers within the organisation could be focusing on documenting the consequences of the incident rather than the causes of the security incidents.

While a root cause was identified using the framework in 36 security investigation records, there were 75 records where a root cause could not be established using only the information in the investigation records. This was because the answers to “Why?” questions could not be found in the investigation record itself and the answer from the previous “Why?” question would not have prevented the security problem from recurring, if it was fixed. This finding highlights that if more rigorous data was captured during an investigation then a more in-depth root cause analysis could have taken place, should the need arise. Figure 8.6 summarises the number of investigation records which did not complete the framework, along with the position in the framework the analysis was terminated.

No investigation records stopped at the first “Why?”. However, 56 out of the 75 (nearly 75%) of the investigation records stopped at the second “Why?” question. This means that in nearly three quarters of the examined records, an answer to the second “Why?” question could not be found and the framework was terminated. Recall that nearly 64% of the discovered root causes were identified at the second “Why?” question. This result suggests that if the incident handlers had recorded a little more information, then more root causes could have been identified in the historical study. Furthermore, 12 out of the 75 records stopped at the third "Why?" question and four out of the 75 records stopped at the fourth "Why?" question. Finally, three investigations record stopped at the fifth "Why?" question.

8.3.3 Historical Study Findings

The significance of the findings from the historical study are discussed below. These findings are discussed from the perspective of investigation record data quality for root cause analysis, documenting security investigations for root cause analysis and the impact of third-parties on root cause analysis.

Investigation Record Data Quality for Root Cause Analysis

The results from the historical study have showed that the quality of the data in 120 of the analysed investigation records may not be detailed enough for an in-depth Root Cause Analysis (RCA). This is visible at two points during the historical study. 45 security investigation records lacked sufficient information to determine whether there had been an attempt to breach the confidentiality, integrity and/or availability of an information asset. Therefore, in these 45 cases, the analysis of the investigation record stopped at the first check and could not proceed to undertake a full RCA. Furthermore, 75 security investigation records did not contain enough information to provide an answer to one of the "Why?" questions asked during the application of the framework. In these 75 records, the analysis using the framework could have proceeded further if more information was document during the initial investigation. If more information was available in these 120 investigation records, then further analysis using the framework could potentially have identified an underlying root cause within these investigation records.

Documenting Security Investigations for Root Cause Analysis

The analysis of the 530 security investigation records showed that typically, the documented investigation records describe the consequences rather than the causes of a security incident, i.e. explaining what happened, but not why it happened. An extract from one investigation record documented that "individual customer records were found on an unauthorised local area network share, the manager was contacted and asked to remove the data immediately to avoid regulatory repercussions" and then "data removed" is one example where consequences rather than causes have been documented. This is a problem which has also been identified in the safety domain [239]. Livingston, et al. [239] add that this can be corrected when an organisation adopts investigation techniques that explicitly identify root causes and the mind set with incident handlers is that there is the potential to learn from past slips and avoid similar incidents in the future. This observation from the content of the security investigation records supports previous findings that organisations are more concerned with eradication and recovery and less on security incident learning [18, 19, 44].

Third-Party Involvement

Organisations are progressively outsourcing their Information Technology (IT) demands to third parties [248, 249]. However, outsourcing IT services can also introduce vulnerabilities into an organisation's infrastructure and therefore, alternative strategies to conduct security incident response in these settings could be required [251]. This can also introduce several additional challenges for an organisation's security incident response team tasked with investigating incidents involving third-parties.

The results from the historical study identified one such challenge. The analysis from the study showed that for 10 out of the 530 investigations, the investigation records suggested that the security incident response team did not have access to the required data. As a result, the security incident response team may not be able to perform a comprehensive investigation to identify underlying root causes. At this point, depending on the organisation's SLA with the third-party, either no RCA would have been conducted or the organisation would have had to depend on the third-party to undertaken the RCA. However, an organisation may not control how this RCA is undertaken or the quality of investigation that may follow. Furthermore, third-parties can provide services to a variety of other customers, often with competing demands and therefore may not be able to provide specific information to a security incident response team who are looking to perform a root cause analysis.

In summary, the results from the historical analysis have highlighted that if more information was documented during an investigation, enhanced underlying root causes could have been identified from the investigations. The next section will describe the application of the framework to 'live' security incidents to investigate the impact of asking further questions and documenting more information during a security investigation.

8.4 Analysis of Security Incidents using Framework

This section presents the results from the analysis of three 'live' security investigations, where a parallel investigation to that being conducted by the security incident handlers was undertaken. This involved the author applying the framework together with one or more incident handlers within the organisation. The framework was applied to three investigations during the course of the experiment. The author undertook the responsibility of applying the framework and collected the required information from the incident handler(s). Information which was requested from the incident handler(s) included answers to "Why?" questions, who provided the answer to the question, how the answer was derived (e.g. 'investigation analysis' or 'through discussion') and what evidence could confirm the answer to the relevant question. The details of the analysis have been anonymised, hence the names of individuals, processes and systems have been altered or generalised to protect the identity of the studied organisation.

8.4.1 Investigation 1

This security incident involved unauthorised access to a restricted folder within a local area network file share belonging to the Human Resources (HR) Department. The restricted folder contained details for a number of employees and the investigation was initiated because access to the folder was granted to the 'Authenticated Users Group'. This means that all users in the organisation with a valid username and password could access the folder. The following is a narrative summary derived from the actual investigation record of the events leading to the incident:

1. A specific HR folder (a child folder) resided within a parent folder, which inherited access control permissions from the parent folder.

2. A request was made to limit access to the child folder but the folder owner was informed this was not possible because the child folder inherited permissions from the parent folder and therefore access was based on the parent folder's permissions.

3. Instead a new folder was created at the same-level as the parent folder which would have its own permissions that only the individual could access.

4. The folder was created with restricted permissions, and the access permissions were verified and authenticated by the individual. At this point, only the individual who requested the folder could access it.

5. Data was copied from the HR folder to this new folder by the Windows Server Support Team. After the data transfer was completed, the permissions were tested and only the folder owner could access the folder.

6. A month later, whilst auditing folder permissions, the individual's line manager identified that the access control list for the new folder had been modified and as a result was now accessible to anyone within the organisation.

7. A report was made to correct the access control list and set the permissions to the correct access level. However, for an unknown period of time, anyone within the organisation had access to the folder.
8. A report was made to the security incident response team.

The following is a narrative summary of the actual investigation record which highlights the actions taken during the investigation as documented in the relevant investigation record:

1. The incident was reported to the security incident response team and a security incident investigation was raised.
2. Discussions were held with the individual on the Access Control team who created the folder and assigned the restricted permissions.
3. This individual confirmed that the permissions were validated by the requester and his line manager. The individual also confirmed that the folder permissions were rechecked once the data had been copied and this check showed that the access control restrictions were still valid.
4. An incident meeting took place to establish what cause the security incident. In the meeting gaps were identified in two processes, which could result in persons gaining access to folders to which they do not have the correct permissions.

Framework Application

The results from the application of the framework to Investigation 1 can be found in Figure 8.7. Labels have been added to the RCA table sheet to assist with the discussion below. The problem statement for this investigation was “Inadvertent access to a restricted folder containing sensitive information”. The first “Why?” question came directly from the problem statement, “Why was access possible on a folder which was supposed to be restricted?”. One answer (ANS 1) was identified from this question: “Folder permissions were modified from ‘Restricted All’ to ‘Full Control All’ for the Authenticated Users group”. The answer was provided by the Primary Incident Handler (PIH) through their investigation of the folder’s permissions, as well as through discussions with the Access Control team.

The next “Why?” question focused on the answer from ANS 1 and the question was: “Why were the folder permissions changed so that the authenticated users group had full control?”. This resulted in two answers (ANS 2.1 and ANS 2.2). The answer for ANS 2.1 identified that “The folder had been deleted from the network file share and then restored from a backup, sometime later”. The answer for ANS 2.2 revealed that “When a folder is restored at the

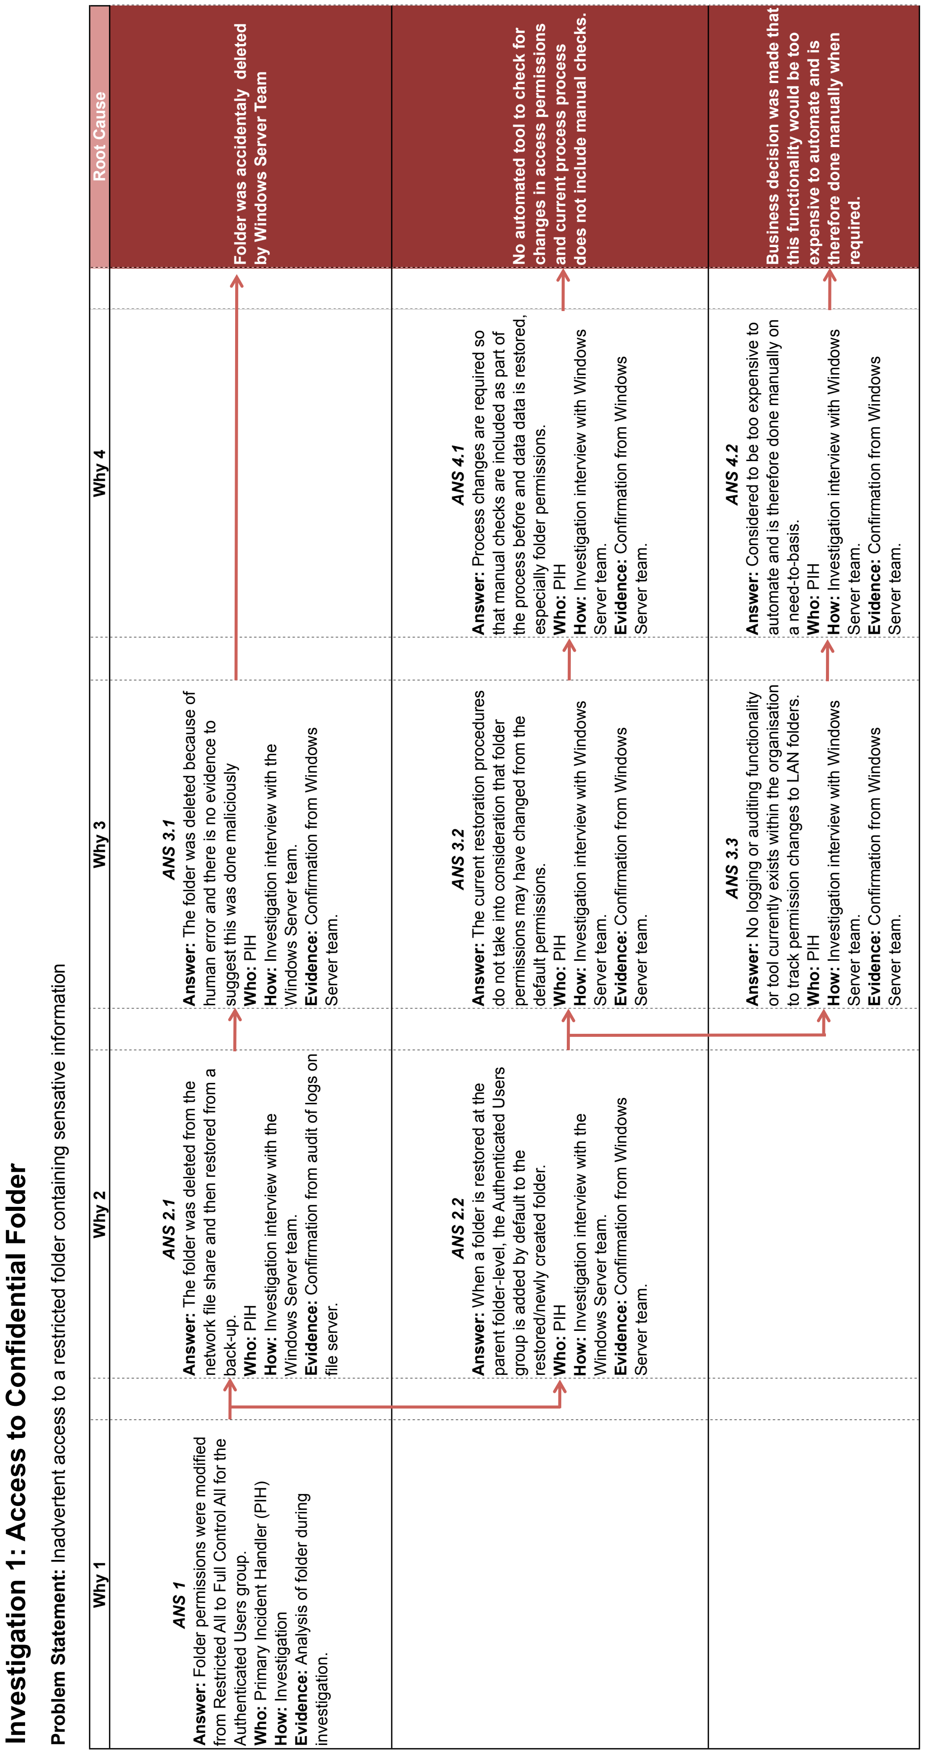

Investigation 1: Access to Confidential Folder

Problem Statement: Inadvertent access to a restricted folder containing sensative information

Why 1	Why 2	Why 3	Why 4	Root Cause
ANS 1 **Answer:** Folder permissions were modified from Restricted All to Full Control All for the Authenticated Users group. **Who:** Primary Incident Handler (PIH) **How:** Investigation **Evidence:** Analysis of folder during investigation.	***ANS 2.1*** **Answer:** The folder was deleted from the network file share and then restored from a back-up. **Who:** PIH **How:** Investigation interview with the Windows Server team. **Evidence:** Confirmation from audit of logs on file server.	***ANS 3.1*** **Answer:** The folder was deleted because of human error and there is no evidence to suggest this was done maliciously **Who:** PIH **How:** Investigation interview with the Windows Server team. **Evidence:** Confirmation from Windows Server team.		**Folder was accidentaly deleted by Windows Server Team**
	ANS 2.2 **Answer:** When a folder is restored at the parent folder-level, the Authenticated Users group is added by default to the restored/newly created folder. **Who:** PIH **How:** Investigation interview with the Windows Server team. **Evidence:** Confirmation from Windows Server team.	***ANS 3.2*** **Answer:** The current restoration procedures do not take into consideration that folder permissions may have changed from the default permissions. **Who:** PIH **How:** Investigation interview with Windows Server team. **Evidence:** Confirmation from Windows Server team.	***ANS 4.1*** **Answer:** Process changes are required so that manual checks are included as part of the process before and data data is restored, especially folder permissions. **Who:** PIH **How:** Investigation interview with Windows Server team. **Evidence:** Confirmation from Windows Server team.	**No automated tool to check for changes in access permissions and current process process does not include manual checks.**
		ANS 3.3 **Answer:** No logging or auditing functionality or tool currently exists within the organisation to track permission changes to LAN folders. **Who:** PIH **How:** Investigation interview with Windows Server team. **Evidence:** Confirmation from Windows Server team.	***ANS 4.2*** **Answer:** Considered to be too expensive to automate and is therefore done manually on a need-to-basis. **Who:** PIH **How:** Investigation interview with Windows Server team. **Evidence:** Confirmation from Windows Server team.	**Business decision was made that this functionality would be too expensive to automate and is therefore done manually when required.**

Figure 8.7: Framework Worksheet Results for Investigation 1

parent folder-level, the Authenticated Users group is added by default to the restored/newly created folder". Both of these answers came from the PIH, through their investigation and interviews with the Windows Server team, who had examined the logs on the file server concerned with the investigation.

The next "Why" questions were derived from ANS 2.1 and ANS 2.2. The first "Why?" question in the third iteration focused on ANS 2.1 and was: "Why was the folder deleted?". One answer (ANS 3.1) was derived from this question, "The folder was deleted because of human error, but there is no evidence to suggest that this was done maliciously". At this point, no further "Why?" questions were asked with regards to the answer labelled ANS 3.1 and this was identified as a root cause. The second "Why?" question in this iteration focused on ANS 2.2: "Why were the incorrect permissions and groups applied to the folder upon restoration?". Two answers (ANS 3.2 and ANS 3.3) were identified in response to this question. ANS 3.2 focused on the folder restoration procedure within the organisation: "The current restoration procedures do not take into consideration that folder permissions may have changed from the default". This means that when folders are restored the access control permissions need to be reconfigured for the specific folder. ANS 3.3 continued with the theme of the folder restoration: "No logging or auditing functionality or tool currently exists to track permission changes to LAN folders". This means that the team, which restores the folders to the parent-level, needs to manually determine what permissions are required and then apply them. All three answers (ANS 3.1 - ANS 3.3) were provided by the PIH, who interviewed the Windows Server team as part of their investigation.

The next "Why?" questions focused on the two answers (ANS 3.2 and ANS 3.3) from the previous iteration, where the root cause had not yet been established. The first "Why?" question in this iteration focused on ANS 3.2 and asked: "Why do the current restoration procedures not take into consideration that folder permissions may have changed from the default?". One answer (ANS 4.1) was identified in response to this question, "Changes to processes are required which include manual permission checks before data is restored and after data is restored, particularly focusing on folder permissions". This answer was identified as a root cause and no further "Why?" questions were asked with regards to the answer labelled ANS 4.1. The second "Why?" question in this iteration focused on ANS 3.3 and the "Why?" question asked was: "Why is no logging or auditing undertaken to track permission changes on LAN share folders?". The answer (ANS 4.2) to this question identified another root cause: "A business decision was made that this would be too expensive to automate and is therefore done manually on a need-to-basis." Both answers to the questions were provided by the PIH through their interviews with the Windows Server team.

8.4.2 Investigation 2

This security incident involved investigating auto-forward rules within the organisation's Lotus Notes email client system. An audit was undertaken to review email auto-forward rules, which were detected by the organisation's Data Loss Prevention (DLP) software. Auto-forward rules can be used to automatically forward or redirect email messages sent from one email account to another email account [252]. The following is a narrative summary derived from the actual investigation record of the events leading to the incident:

1. The Information Security unit initiated an audit review of the email auto-forward rules that had been detected by the DLP software in the past year.
2. Over 20,000 potential cases were identified from the initial review and an in-depth analysis focused on these cases to determine if any breach of organisation's information security policy had occurred.
3. The results from the analysis showed that 59 cases were identified where a potential breach of policy had occurred and a separate investigation was initiated for each case.

A master investigation record was used by the incident handlers to document the actions undertaken during the 59 security investigations. The following is a narrative summary of the actual investigation record which highlights the actions taken during the investigation as documented in the in the relevant investigation record:

1. The two incident handlers who were assigned to the investigation, contacted all those affected by the auto-forward cases with the instruction that unless the auto-forward had been approved for a policy exception, all the rules must be deleted.
2. The investigation continued by contacting Email Support asking them for assistance so that all non-exempt auto-forwards can be deleted by the team, as soon as possible.
3. Technical solutions were then discussed to block all auto-forwards.

Framework Application

The results from the application of the framework for Investigation 2 can be found in Figure 8.8. The problem statement derived for this investigation was: "Data Loss Prevention (DLP) software identified approximately 22,000 cases of emails being auto-forwarded to external email addresses". The first "Why?" question came directly from the problem statement, "Why were auto-forward emails being detected by the DLP software?". One answer (ANS 1) was identified from this question, "Employees created the auto-forward rules within their

Lotus Notes email clients". The answer to this why question was provided by one of the Primary Incident Handlers (PIHs) involved in the investigation, through their analysis of the DLP logs.

The next "Why" question was derived from ANS 1, "Why were employees creating auto-forward rules?". Three answers (ANS 2.1 - 2.3) were identified in response to this "Why?" question. ANS 2.1 identified that "Employees used their organisation email accounts for personal activities". ANS 2.2 identified that "Lotus Notes was not configured (i.e. 'locked down') to prevent the creation of auto-forward rules on the client". ANS 2.3 identified that "Contractors travel between different sites and want to access their organisational email on their personal devices and therefore set-up auto-forward rules to forward emails to their personal email accounts". All three answers were identified by the PIHs using a variety of methods. ANS 2.1 was identified during the analysis of logs from the DLP software, ANS 2.2 was identified from observations that there is a lack of documentation within the organisation and ANS 2.3 was identified through interviews by the PIHs with the contractors involved in the investigation.

The next "Why?" questions focused on the three answers (ANS 2.1 - 2.3) from the previous iteration. The first "Why?" question in this iteration focused on ANS 2.1 and was used to ask: "Why are employees using their organisational email address for personal usage?". Two answers (ANS 3.1 and ANS 3.2) were identified from this question. The first answer (ANS 3.1) was "Convenience". This is because employees attend various external vendor training events and use their organisational email address at these events for registration. This answer suggests that there is a lack of awareness surrounding potential data leakage from personal email addresses. At this point, no further "Why?" questions were asked with regards to the answer labelled ANS 3.1 and this was identified as a root cause. ANS 3.1 was identified by the PIHs through their interviews with the affected individuals. The second answer (ANS 3.2) was that "Local Information Security policies allow limited personal use of organisational communication resources, but organisation-issued email accounts are not specifically mentioned". ANS 3.2 was identified by the PIHs during their investigation, as well as through their observation of relevant information security policies. The second "Why?" question in this iteration focused on ANS 2.2. The "Why?" question asked was "Why is there no Lotus Notes lock-down standard?". One answer (ANS 3.3) was identified from this question,"It is not considered practical to have a lock-down standard for Lotus Notes clients because there are frequent version changes and if problems arise, technical solutions are often better suited to prevent a reoccurrence of a problem". This answer was provided by a Compliance Officer in the organisation, who was interviewed by the PIH. The result from this "Why?" question provided another root cause, "Lack of lock-down standard for Lotus Notes clients, which could have prevented auto-forward rules from being created". No further queries were submitted with regards to the lock-down standard. The

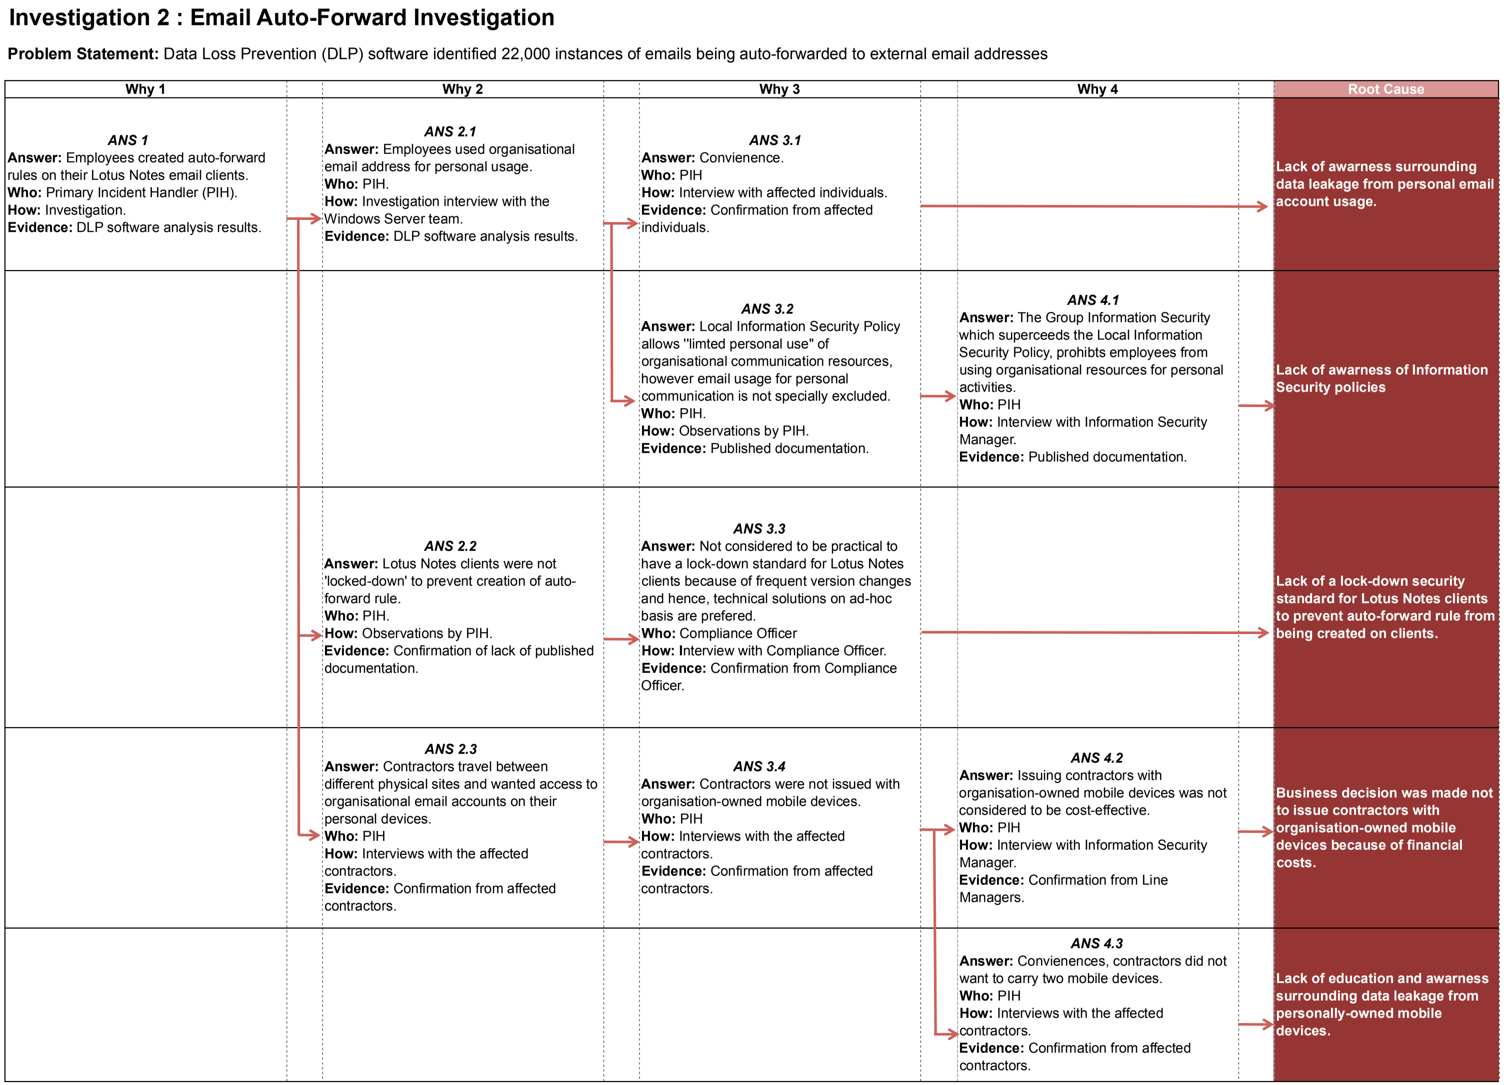

Investigation 2 : Email Auto-Forward Investigation

Problem Statement: Data Loss Prevention (DLP) software identified 22,000 instances of emails being auto-forwarded to external email addresses

Why 1	Why 2	Why 3	Why 4	Root Cause
ANS 1 **Answer:** Employees created auto-forward rules on their Lotus Notes email clients. **Who:** Primary Incident Handler (PIH). **How:** Investigation. **Evidence:** DLP software analysis results.	***ANS 2.1*** **Answer:** Employees used organisational email address for personal usage. **Who:** PIH. **How:** Investigation interview with the Windows Server team. **Evidence:** DLP software analysis results.	***ANS 3.1*** **Answer:** Convienence. **Who:** PIH **How:** Interview with affected individuals. **Evidence:** Confirmation from affected individuals.		**Lack of awarness surrounding data leakage from personal email account usage.**
		ANS 3.2 **Answer:** Local Information Security Policy allows "limted personal use" of organisational communication resources, however email usage for personal communication is not specially excluded. **Who:** PIH. **How:** Observations by PIH. **Evidence:** Published documentation.	***ANS 4.1*** **Answer:** The Group Information Security which superceeds the Local Information Security Policy, prohibits employees from using organisational resources for personal activities. **Who:** PIH **How:** Interview with Information Security Manager. **Evidence:** Published documentation.	**Lack of awarness of Information Security policies**
	ANS 2.2 **Answer:** Lotus Notes clients were not 'locked-down' to prevent creation of auto-forward rule. **Who:** PIH. **How:** Observations by PIH. **Evidence:** Confirmation of lack of published documentation.	***ANS 3.3*** **Answer:** Not considered to be practical to have a lock-down standard for Lotus Notes clients because of frequent version changes and hence, technical solutions on ad-hoc basis are prefered. **Who:** Compliance Officer **How:** Interview with Compliance Officer. **Evidence:** Confirmation from Compliance Officer.		**Lack of a lock-down security standard for Lotus Notes clients to prevent auto-forward rule from being created on clients.**
	ANS 2.3 **Answer:** Contractors travel between different physical sites and wanted access to organisational email accounts on their personal devices. **Who:** PIH **How:** Interviews with the affected contractors. **Evidence:** Confirmation from affected contractors.	***ANS 3.4*** **Answer:** Contractors were not issued with organisation-owned mobile devices. **Who:** PIH **How:** Interviews with the affected contractors. **Evidence:** Confirmation from affected contractors.	***ANS 4.2*** **Answer:** Issuing contractors with organisation-owned mobile devices was not considered to be cost-effective. **Who:** PIH **How:** Interview with Information Security Manager. **Evidence:** Confirmation from Line Managers.	**Business decision was made not to issue contractors with organisation-owned mobile devices because of financial costs.**
			ANS 4.3 **Answer:** Convienences, contractors did not want to carry two mobile devices. **Who:** PIH **How:** Interviews with the affected contractors. **Evidence:** Confirmation from affected contractors.	**Lack of education and awarness surrounding data leakage from personally-owned mobile devices.**

Figure 8.8: Framework Worksheet Results for Investigation 2

third “Why?” question in this iteration focused on ANS 2.3. The “Why?” question asked was “Why are contractors sending organisational email to their personal devices?”. One answer (ANS 3.4) was identified from this question, “This is because contractors are not issued with organisation-owned mobile phones.” The answer was provided by one of the PIHs who interviewed contractors affected by the investigation.

The next “Why?” questions focused on the answers (ANS 3.2 and 3.4) from the previous iteration, where a root cause had not yet been identified. The first “Why?” question in this iteration focused on ANS 3.2 and asked: “Why are organisation-issued email accounts not explicitly excluded from the Local Information Security policies?”. One answer (ANS 4.1) was identified, “The organisation’s Group Information Security Policy, supersedes the Local Information Security Policy, and this explicitly prohibits users from auto-forwarding group information to personal accounts as well as using group accounts for non-group activities”. In this context, ‘Group’ refers to the parent company, which owns the studied organisation. Therefore, the root cause identified from this query was that there is a “Lack of awareness surrounding information security policies which prohibit personal usage of organisational resources”. No further inquiries were made in reference to ANS 4.1. The answer was identified through interviews with one of the organisation’s information security managers, as well as through an analysis of the organisation’s relevant documentation. The second “Why?” question in this iteration focused on ANS 3.4 and asked: “Why were contractors not issued with organisation-owned mobile phones?”. Two answers (ANS 4.2 and ANS 4.3) were identified from this question. The first answer (ANS 4.2) was that “Issuing contractors with organisation-owned mobile phones was not considered cost effective”. This answer was provided through interviews with an information security manager and confirmed by the contractor’s line manager. This answer provided another root cause, “A business decision was made not to issue contractors with organisation-owned mobile phones because of financial costs”. No further enquires were made regarding this question. The second answer (ANS 4.3), identified was that “Convenience, contractors did not want to have two mobile devices”. This too was identified as a root cause and suggested a “Lack of education and awareness surrounding potential data leakage from personally-owned mobile devices”. No further questions were asked regarding why contractors auto-forwarded emails to their personal mobile devices. Both answers were identified by the PIHs through interviews with affected contractors and with an information security manager.

8.4.3 Investigation 3

A security incident was raised in relation to a vulnerability being identified in a number of Lotus Notes databases within the organisation. The vulnerability allows any authenticated user to access administrative ‘views’ of the databases and read potentially confidential data of varying classification. The following is a narrative summary derived from the actual investigation record of the events leading to the incident:

1. A report was made via the Helpdesk to the Information Technology Service Incident Response (ITSIR) team about a potential vulnerability in three Lotus Notes databases.
2. The ITSIR team initiated the ITSIR process and held an incident meeting where it was decided that this was an information security issue and not an IT service issue.
3. ITSIR team manager notified an information security analyst of a potential vulnerability in the Lotus Notes databases.
4. The information security analyst in turn notified the security incident response team who started the investigation.

The following is a narrative summary of the actions taken during the investigation as documented in the in the actual investigation record:

1. The security incident response team notified information security management of the potential vulnerability and then initiated the security incident response process and a primary incident handler was assigned to the investigation.
2. An incident meeting was held with the primary incident handler, the Lotus Notes Development team and various information security analysts involved within the Access Control team.
3. The vulnerability was demonstrated at the incident meeting, which bypassed implemented access controls and allowed employees to access administrative pages, where confidential information could be viewed.
4. A fix was agreed between the primary incident handler and the Lotus Notes Development team.

Framework Application

The results from the application of the framework for this investigation can be found in Figure 8.9. The problem statement for this investigation was “Unauthorised access to administrative areas within Lotus Notes Databases”. The first “Why?” question came directly from the problem statement, “Why was unauthorised access possible within the Lotus Notes Databases?”. One answer (ANS 1) was identified from this question: “Implemented access controls for the administrative areas of the databases were bypassed by employees”. The

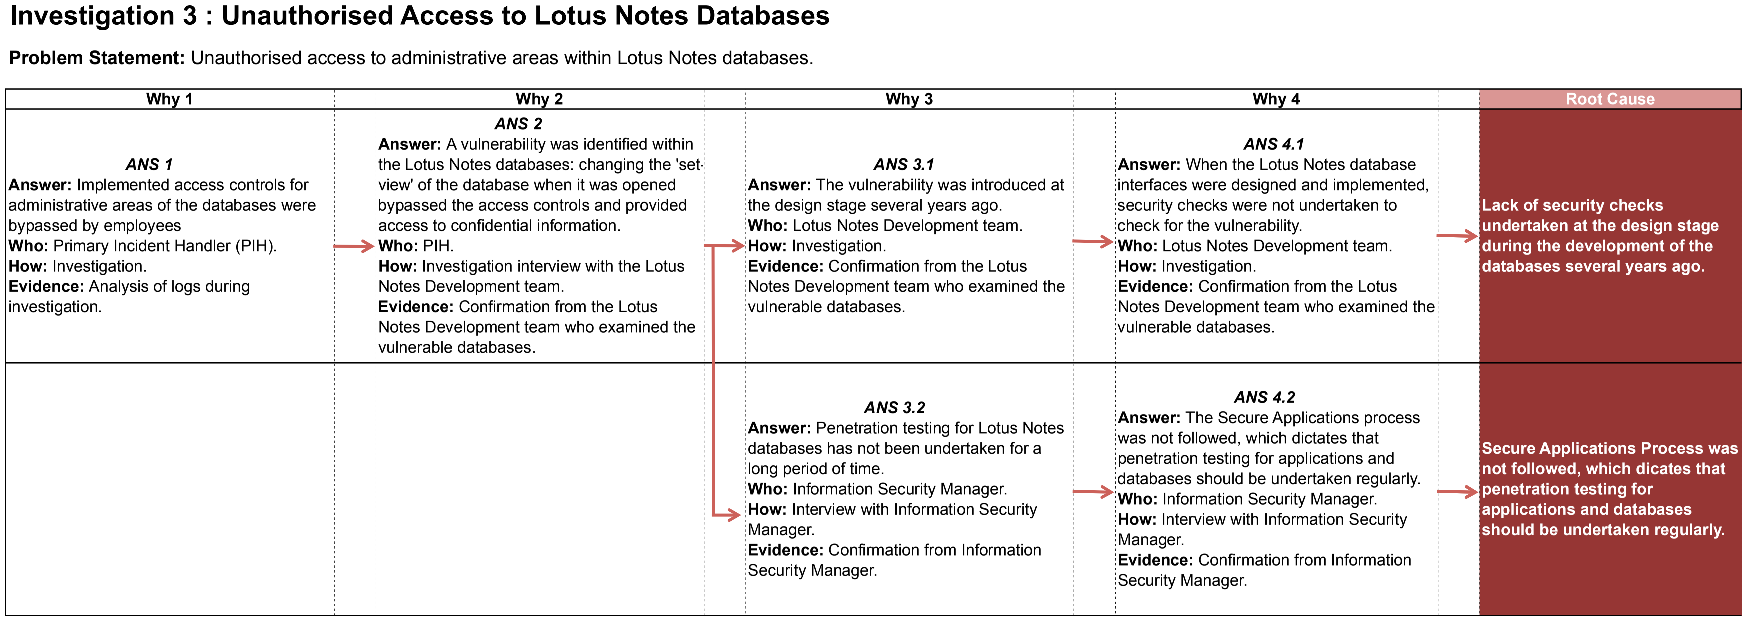

Investigation 3 : Unauthorised Access to Lotus Notes Databases

Problem Statement: Unauthorised access to administrative areas within Lotus Notes databases.

Why 1	Why 2	Why 3	Why 4	Root Cause
ANS 1 **Answer:** Implemented access controls for administrative areas of the databases were bypassed by employees **Who:** Primary Incident Handler (PIH). **How:** Investigation. **Evidence:** Analysis of logs during investigation.	***ANS 2*** **Answer:** A vulnerability was identified within the Lotus Notes databases: changing the 'set-view' of the database when it was opened bypassed the access controls and provided access to confidential information. **Who:** PIH. **How:** Investigation interview with the Lotus Notes Development team. **Evidence:** Confirmation from the Lotus Notes Development team who examined the vulnerable databases.	***ANS 3.1*** **Answer:** The vulnerability was introduced at the design stage several years ago. **Who:** Lotus Notes Development team. **How:** Investigation. **Evidence:** Confirmation from the Lotus Notes Development team who examined the vulnerable databases.	***ANS 4.1*** **Answer:** When the Lotus Notes database interfaces were designed and implemented, security checks were not undertaken to check for the vulnerability. **Who:** Lotus Notes Development team. **How:** Investigation. **Evidence:** Confirmation from the Lotus Notes Development team who examined the vulnerable databases.	**Lack of security checks undertaken at the design stage during the development of the databases several years ago.**
		ANS 3.2 **Answer:** Penetration testing for Lotus Notes databases has not been undertaken for a long period of time. **Who:** Information Security Manager. **How:** Interview with Information Security Manager. **Evidence:** Confirmation from Information Security Manager.	***ANS 4.2*** **Answer:** The Secure Applications process was not followed, which dictates that penetration testing for applications and databases should be undertaken regularly. **Who:** Information Security Manager. **How:** Interview with Information Security Manager. **Evidence:** Confirmation from Information Security Manager.	**Secure Applications Process was not followed, which dicates that penetration testing for applications and databases should be undertaken regularly.**

Figure 8.9: Framework Worksheet Results for Investigation 3

answer was provided by the Primary Incident Handler (PIH) through their own analysis of logs files from the affected Lotus Notes databases.

The next “Why?” question focused on the answer (ANS 1) from the previous iteration, “Why did the access controls fail?”. This resulted in one answer (ANS 2) being identified, “A vulnerability was identified within the Lotus Notes databases which involved changing the ‘set-view’ of an open database and this bypassed the access controls and provided access to confidential information”. The evidence for this answer came from the PIH’s investigation and interviews with the Lotus Notes Development team, who examined the vulnerable databases.

The next “Why” question focused on the existence of the vulnerability as identified in ANS 2 and therefore, the “Why?” question asked was “Why did the vulnerability exist within the databases?”. Two answers (ANS 3.1 and ANS 3.2) were identified from this query. The first answer (ANS 3.1) focused on the design of the databases themselves, “The vulnerability was introduced during the design stage of the database interfaces several years ago”. This answer was reported by Lotus Notes Development team who examined the historical interface designs affected by the incident. The second answer (ANS 3.2) focused on the lack of penetration testing looking for potential security weaknesses, “Penetration testing for the Lotus Notes databases has not been undertaken for a long period (undefined) of time”. This answer was identified through an interview with one of the information security managers who was assigned to assist with the investigation.

The next “Why?” questions focused on the answers (ANS 3.1 and 3.2) from the previous iteration. The first “Why?” question in this iteration focused on ANS 3.1 and asked “Why did the vulnerability exist within the interface?”. This resulted in one answer (ANS 4.1) being identified, “When the Lotus Notes database interfaces were designed and implemented, security checks were not undertaken to check for the vulnerability”. This answer resulted in the identification of a root cause, “Lack of security checks undertaken at the design stage during the development of the databases several years ago”. The answer was reported by the Lotus Notes Development team who examined the historical interface designs affected by the incident. The second “Why?” question in this iteration focused on ANS 3.2 and asked “Why had penetration testing not been undertaken on the database for a long period of time?”. The answer (ANS 4.2) to this question was that “The Secure Applications process was not followed, which dictates that penetration testing for applications and databases should be undertaken regularly”. This answer resulted in the identification of another root cause. This answer was identified through an interview with one of the information security managers who was assigned to assist with the investigation.

8.5 Framework Evolution using Semi-Structured Interviews

A series of follow-up interviews were undertaken within the organisation which were used to assess practitioner perception of the framework. The interview questions focused on individual's usage of the framework during 'live' investigations, factors which contributed to the successful or unsuccessful use the framework, perceived strengths and weaknesses of this approach, the impact of the framework on security incident response, as well as additions to the framework to improve overall root cause analysis.

8.5.1 Individuals Participation

Individuals were asked if they had participated in an investigation where the framework was used to perform a root cause analysis. Three out of the seven individuals indicated that they had participated in an investigation where the framework was used and four individuals answered that they did not participate in an investigation where the framework was used.

8.5.2 Perceived Strengths and Weaknesses of the Framework

All seven individuals agreed that the framework provided a consistent and methodical way of undertaking a root cause analysis, which improves both on approaches used in the organisation and those highlighted in the best-practices. Three out of the seven individuals added that the standardised framework also helps with repeatability and provides guidance to the security incident response team about when to undertake a root cause analysis.

However, three individuals suggested that a variety of potential weaknesses *could* exist within the framework. Two individuals argued that incident handlers *"might become lazy because a root cause analysis also requires analytical problem solving skills and the structured framework could mean that if pieces do not fall into place, will the incident handler(s) know what to do when they need to deviate or will they just follow the process?"*. The issue raised by these interviewees is that incident handlers who use the framework might just follow the process and not apply analytical thinking, when a deviation is required. As a result, these incident handlers could just stop and note that a root cause cannot be found using the framework. This is a valid concern, and while the framework provides a structured approach to help guide a root cause analysis, some analytical skills are still required, especially if a deviation from framework is required.

Another problem identified by one individual is that not every 'fix' in a root cause analysis will prevent reoccurrence and in fact not all problems can be solved, therefore some incidents could still occur regardless of the root cause found. When asked to provide an example, the individual stated that any security problem caused by 'human error' would likely result in a root cause that cannot be fixed. This answer suggests that there may be some inconsistency in the understanding of the term 'root cause'. This is because according to the definition in this work, the identification of a root cause implies the availability of fixes to prevent reoccurrence.

8.5.3 Impact of the Framework on Security Incident Response

Five out of the seven individuals indicated that the framework provides some form of benefit to the overall security incident response process. These five individuals suggested that the framework, when integrated into a security incident response process, helps guide when a root cause analysis is required, outlines how to determine one or more root causes and then provides justification as to when the analysis should be concluded. Four out of the five individuals noted that one of the most important parts of the framework was the justification of when to stop the root cause analysis. This is because security incident handlers could interpret when to stop differently and the standardised decision-making process in the framework removes any ambiguity surrounding this issue.

Two out of the seven individuals suggested that the framework both *assists* and *hinders* the overall security incident response process. The individuals argued that because the framework involves more work, it could slow down a security incident response process. However, the individuals added that the framework could also help derive one or more root causes which may not be identified using a 'normal' incident response process. One of the two individuals added that the framework can also help security incident handlers who are unsure of how and when to do a root cause analysis, which means a more enhanced security investigation and 'deeper' lessons learned.

8.5.4 Improving the Framework

Four out of the seven survey participants said they would not change the current framework and could not think of anything that would improve it at the present time. The remaining three participants proposed various modifications to the designed framework. Two individuals suggested that the framework might need to be modified with regards to inserting a 'break' in the loop surrounding "If the current diagnosis was fixed would the security problem reoccur?'. This was raised as a concern because the individuals argued that the 'fix' is likely to be applied by other teams within the information security unit and the incident handlers are unlikely to complete the process waiting for other teams to implement changes. While this is a valid suggestion, the purpose of the framework is to help identify what fix could be applied to prevent re-occurrence and not to ensure that a fix is actually applied. Applying fixes to security problems would vary between organisations and the framework does not require that a fix is applied, merely that a fix is identified.

The third individual suggested that the framework should be modified so that a relationship manager is interacting with a third-party provider and not the security incident response team. This is an interesting suggestion that supports previous observations and findings [27, 212] that security incident response is becoming a multidisciplinary affair, that involves the integration of technical security professionals, information technology specialists and relevant managers. Within the specific organisation, this suggestion has been made because policies mandate that only relationship managers should be interacting with third-parties who have contract obligations with the organisation.

8.5.5 Other Factors

Further interview questions attempted to determine what factors the respondents perceived as contributing to the successful or unsuccessful use the framework for root cause analysis within the security incident response team. Four individuals provided ‘do not know’ answers because they were not involved in any root cause analysis using the framework.

Three individuals answered that *time* and *access to data* were important factors which need to be taken into consideration when attempting to undertake a root cause analysis. Two out of the three individuals argued that time is important because a root cause analysis can take time to complete and the security incident response team cannot stop eradication and recovery tasks to undertake a full analysis. Time was also discussed from the perspective that the security incident response team’s workload can become very large and as a result, incident handlers quickly move on to the next incident and may not find the time to undertake a root cause analysis. The interviewee’s comments shows that the root cause analysis is viewed as an after-event to the actual investigation process. This is not necessarily true because the framework was designed to run in parallel to a security investigation being undertaken by an incident handler. In fact, the framework can actually be used to guide an investigation being undertaken and not just executed at the end of an investigation.

The comments from the three individuals have also highlighted a potential management culture issue, particularly with regards to leadership [253]. The implementation of a root cause analysis in the organisation implies that management supports and encourages the process so that new understandings into incident causes can be identified [254]. However, the individual’s comments indicate that further leadership from management is required to encourage incident handlers to undertake a root cause analysis. Encouraging and leading incident handlers to undertake a root cause analysis will also provide managers with information on problems that they can change and improve [255]. One out of the three individuals added that access to data is important for a root cause analysis, in the sense that data needs to be both consistent and available to the team undertaking a root cause analysis.

8.6 Discussion

The significance of the results from the 'live' security investigations using the framework are discussed from four perspectives: selection of security incidents for root cause analysis, learning without undertaking a root cause analysis, root cause analysis 'buy-in' and the development of organisational culture towards root cause analysis.

8.6.1 Selection of Security Incidents for Root Cause Analysis

While the framework defines that a Root Cause Analysis (RCA) is only undertaken when there has been a breach in the confidentiality, integrity and/or availability (CIA) of an information asset, an observation from the application of the framework was that security incident handlers disagreed on this proposed criteria. These incident handlers suggested that even though a security incident has resulted in a breach of CIA, it must be 'severe' enough to do a full RCA. During the application of the framework to 'live' security incidents, there were several investigations where it was clear that there had been a breach of CIA. However, incident handlers who were approached to undertaken an analysis using the framework considered that these investigations were not 'serious enough' to undertaken a full RCA. As a result, the number of security incidents that were available to apply the complete framework were reduced. Hence, only three 'severe' incidents were analysed using the complete framework.

While the organisation's information security incident response process dictates that a 'root cause' is established for all security incidents, the term 'root cause' is not defined in the process. Informal discussions with security incident handlers confirmed that the security incident response team has not defined the term. A definition of the term 'root cause', as well as when to undertake a root cause analysis within the organisation's security in incident response process could have helped to remove any ambiguity within the security incident response team.

8.6.2 Learning Without a Full Root Cause Analysis

While the discussion above proposes that organisations should define when to undertake a RCA, there is still the potential to learn from security incidents without a full RCA being undertaken. Recall Cooke's definition of incident learning as "the collection of organisational capabilities that enable the organisation to extract useful information from incidents of all kinds and to use this information to improve organisational performance over time" [120]. Based on this definition, "useful information" and "improve organisational performance" suggests that learning from security incidents can be as simple as calculating metrics as described in Sections 4.3.3 and 6.4.9 and improving security incident response performance. During the experiment, an observation was made that security incident handlers received requests from information security managers for various information about the security incident response landscape. This information included the number of security incidents of particular types, the sources of specific incidents and which business units were affected the most by incidents over a period of time.

These observations, when coupled with the practitioner's concerns about the RCA framework suggest that it is still possible to learn about a security incident without undertaking a full RCA. Although underlying root causes will not be identified, analysing previous security incidents can still provide an organisation with information about its security threat landscape, which it can then use to improve its wider security posture.

8.6.3 Root Cause Analysis 'Buy-In'

The recommendation to conduct an experiment involving the framework was approved by the organisation's Head of Information Security. However, observations and informal discussions during the experiment showed that not all the security incident handlers and information security managers had completely 'bought-in' to the idea of undertaking a root cause analysis for particular incidents. The results from Investigation 2, as described in Section 8.4.2, were requested at the conclusion of the experiment and included as an appendix in the actual investigation record. However, informal discussions with security incident handlers, in the months that followed the investigation, revealed that none of the improvements identified from the root cause analysis were actually implemented. Several reasons were established including disagreements over the proposed improvements, as well as a lack of champion to promote improvement implementation. From the perspective of time and resources required, a root cause analysis can become an expensive activity for an organisation. Therefore, if an organisation is going to undertake a root cause analysis, security incident response teams and information security managers need to be convinced of the benefits of spending time and resources to do so. A challenge here is that implementing proposed improvements (and accepting associated costs) may be the only way of determining if the proposed fixes are effective.

8.6.4 Culture Shift for Root Cause Analysis

One of the individuals in the follow-up interviews stated that if a security problem is caused by human error, the likely outcome would be a root cause that cannot be fixed. The basis of this comment suggests that there are potential culture conflicts when it comes to conducting a root cause analysis in the organisation, particularly with regards to fixing human issues which have caused an incident. Several researchers [253–255] have argued that there is a strong correlation between undertaking an effective root cause analysis in an organisation, and the development of an organisation's culture to ensure that the outcomes of the analysis result in change. In order for a root cause analysis to be a success, it requires a culture of openness, honesty, transparency, accountability, as well as a willingness to embrace the fact that humans can make mistakes, which we need to learn from in order to avoid repeating them in the future [253].

The answer from the individual that if an incident is caused by human error it cannot be fixed, suggests the presence of a culture of blame in the organisation [191, 195]. By transferring blame from the security team to the human element, the incident response team can avoid blaming colleagues and potential allies in the Information Security unit. The presence of a blame culture could also point to managerial culture issues, who could be quick to blame employees for even the smallest of errors, which lead to security incidents. In turn, incident handlers are likely to shift the blame to the human outside of the security process to avoid being blamed for the security incident themselves.

One solution to this problem could be a culture shift towards a 'just culture' [253]. The European Organisation for the Safety of Air Navigation define a 'just culture' as a "culture in which front-line operators and others are not punished for actions, omissions or decisions taken by them which are commensurate with their experience and training, but where gross negligence, wilful violations and destructive acts are not tolerated" [256]. To translate this into security incident response, a 'just culture' would mean that employees are not punished for actions, omissions or decisions which are identified using a root cause analysis. However, if incidents are a result of gross negligence, wilful violations or destructive acts, then employees could face disciplinary matters. The shift towards a 'just culture' would mean that security incident handlers would not be afraid of examining 'human error' as a root cause, as only very serious 'human error' incidents would result in repercussions. In turn, employees could also be more inclined to report less serious incidents, which can result in identifying and fixing a correctable failure by management [256].

8.7 Summary

This chapter described an experiment to evaluate a root cause analysis framework as a method for conducting a root cause analysis of information security incidents. The framework was applied to historical security investigation records in a Fortune 500 Organisation with the purpose of investigating if the information documented within the records assist in conducting a root cause analysis using the framework. The results from the study identified numerous challenges with regards to undertaking a root cause analysis using the organisation's security investigation records. These challenges included poor data quality for in-depth analysis, a need to document more information describing *why* security incidents have happened rather than *what* has happened, as well as the impact of third-parties in the root cause analysis process. The principle finding from the historical study was the quality of data in the security investigation records could be improved if detailed data about the security incidents was documented during the initial investigations.

In order to demonstrate the value of documenting more detailed data during a security investigation, the framework was then applied to three 'live' security investigations. The results from this analysis showed that several underlying causes of security incidents could be identified using the framework. In addition, the application of the framework to these investigations help produce enhanced data from the respective security investigations.

Chapter 9

Conclusions and Future Work

This thesis proposed that the quality of data generated during a security incident investigation can be enhanced through the application of specific lightweight measures. After analysing the results from an exploratory case study of a Fortune 500 Organisation, four experiments presented in Chapters 5 - 8, were used to evaluate if specific lightweight measures can help improve the quality of data generated during security investigations within the specific organisation. These lightweight measures focused on enhancing data quality transparency, capturing data that has been missed during an investigation, generating higher quality data using root cause analysis and focusing investigation efforts on data that will provide value to a security incident response team. The next three sections of this chapter address the research questions that were presented in Chapter one. Section 9.4 discusses the extended research contributions. Section 9.5 discusses the interconnection of the experiments presented in the thesis and how the implemented lightweight measures collectively assist with improving the quality of data in a security incident response investigation. Section 9.6 examines the scope and validity of the research conducted, while Section 9.7 presents an overview of an ideal security incident learning system in the studied organisation. Section 9.8 presents areas for future work and Section 9.9 concludes the chapter.

9.1 Thesis Research Question 1

The answer to the first research question "What data is generated by a real-world security incident response team?" can be derived from the exploratory case study presented in Chapter 4. The data generated by a security incident response team was determined through an analysis of the organisation's security incident response database, relevant documentation and interviews conducted in the Fortune 500 Organisation. The analysis of the relevant security incident response documentation and the security incident response database showed that the organisation's security incident response team generates various data during a security investigation. This data can include the date and time a security incident was reported to the team, contact details of security incident handlers, the physical location of a security incident, investigative notes, calculations regarding financial costs of incidents and data describing any lessons learned.

The interviews undertaken within the exploratory case study revealed that the data generated and collected by the security incident response team is usually tailored to specific investigations. The individuals added that in addition to the above data, the security incident response team also collects data with regards to incident meeting minutes, email trails and data describing actions for remediation. The results from the interviews also identified that forensic data is collected from various sources. This forensic data can include logs, emails, hard disk drive images and physical memory dumps. This type of data can then be used as evidence in legal proceedings, if the need arises.

9.2 Thesis Research Question 2

The answer to the second research question "What challenges and problems do a security incident response team face when attempting to learn from information security incidents?" can be derived from Chapters 4 and 5. In Chapter 4, the results from the exploratory case study of the Fortune 500 Organisation identified several challenges to security incident learning within the organisation. Practitioners within the organisation indicated that limited physical access to security data, short data retention times, logs not containing enough detailed information and limited support from third-parties involved in security investigations could all impact security incident learning. Within Chapter 4, practitioners also argued that security incident response teams could benefit from additional support and clarity surrounding tools and techniques for incident learning. Individuals within the organisation have also called for improved methods, such as root cause analysis, to assist in the development of lessons learns.

The results from the experiment presented in Chapter 5 also revealed an additional incident learning challenge for a security incident response team. This challenge focuses on the problem of data quality in the wider organisation. The results from the analysis of the information documented in the security investigation record template, showed that the fields pertaining to the 'Reporting and Contact Information' section were missing information. When queried about this in the follow-up interviews, security incident handlers indicated that this missing information was because of a data quality problem with the organisation's electronic address book. The security incident response team use this address book to complete the 'Reporting and Contact Information' section of the investigation record. However, the address book itself has data consistency issues. As a result, information is often missing for particular employees. This finding shows that security incident learning can be affected by data quality issues involving other organisational processes and data cycles that a security incident response team may have little or no control over.

9.3 Thesis Research Question 3

The answer to the third research question, “What effect did the application of the described measures have on the data generated by the security incident response process?” can be derived from the experimental evaluations presented in Chapters 5 through 8. In Chapter 5, the analysis of the security investigation records in the organisation showed that all 324 (100%) records, examined as part of the experiment, now contained both a category and sub-category to describe the particular investigation. When compared to investigation records from the exploratory case study where all the records were classified as a ‘security incident’ and no formal taxonomy was available, means that all the records in the organisation now have a category and sub-category to describe the particular investigation. Chapter 5 also discussed how a revised security investigation record template was used to enhance data capture with regards to information, which provides value to a security incident response team. In the exploratory case study in Chapter 4, only 1 out of the 188 analysed security investigation records were found to have been complete from the perspective of all 22 fields within the record template. This means that within 187 investigation records, one or more fields within the record template were missing information. However, after the evaluation of the revised record template (as presented in Chapter 5), the results have shown that 11 out of the 26 fields within the revised record template were completed in all 324 investigation records. Furthermore, four fields were completed in over 99% of the investigation records.

Retrospectives were evaluated in Chapter 6 as a method for validating information that has been collected during a security investigation and enhancing information may have been missed. The results have shown that retrospectives provide an additional avenue for security incident handlers to identify and document information, which may not have been possible to document in the initial investigation. 148 out of the 324 retrospectives undertaken in the organisation were used to capture more data when compared to the initial investigation. This data has included the identification of additional information assets involved in an investigation, individuals and groups within the organisation whose assistance was required during an investigation, as well as enhanced security controls and improvements to various security-related processes within the organisation. Chapter 7 described an experiment to evaluate how a security incident response dashboard can be used to identify where data quality needs to be improved within investigation records and enhance data quality transparency. The results from this experiment showed that 61 investigation records were corrected and completed by security incident handlers during the experiment. If historical trends had continued from the exploratory case study, then these 61 records may not have been corrected without the dashboard application. Chapter 8 described the evaluation of a root cause analysis framework as a method for generating higher quality data for subsequent lessons learned. The results of the evaluation showed that underlying root causes of security incidents can be identified using the framework. In addition, the framework has also helped to enhance the quality of data produced from the application of the framework to the three 'live' security investigations.

9.4 Contributions

The main objective of this research was to examine if the quality of data generated during a security incident investigation can be enhanced through the application of specific lightweight measures. These lightweight measures focused on enhancing data capture that was missed during an investigation, generating higher quality data using root cause analysis and focusing investigation efforts on data that will provide value to a security incident response team. The remainder of this section discusses the contributions and conclusions of this thesis. The main contributions of this thesis are as follows:

In **Chapter 4**, a exploratory case study of a security incident response team in Fortune 500 financial services organisation was presented, which was used to validate the thesis statement. Security incident response documentation and investigation records were examined and interviews were conducted in the studied organisation. The objective was to establish the challenges a security incident response team face with regards to lessons learned development. The organisation's security investigation records were examined from the perspective of Ballou and Pazer's [144] data quality dimensions of *accuracy, completeness* and *consistency* and data quality issues were identified in the investigation records. Furthermore, several other initial challenges were also identified from the case study, which helped to guide the reminder of the work presented in this thesis. These initial challenges included no consistent security incident taxonomy, problems with access to enriched security data, absence of organisational learning from security investigations, and a lack of tools and methods for deeper incident learning. At a high-level, these challenges suggested that improving the quality of data generated from the organisations security response investigations needs to be addressed.

Chapter 5 described an experiment that attempted to address two of the data quality challenges identified in the exploratory case study through the introduction of a well-defined security incident categorisation taxonomy (Section 5.2.1) and a revised security investigation record template (Section 5.2.2). The objective of the experiment was twofold. First, the experiment was used to evaluate the efficacy of the taxonomy with the purpose of removing ambiguity surrounding incident type identification. Second, the experiment was used to examine if the investigation record assists with the enhancement of data collection that provides value to the organisation's security incident response team. The above mechanisms were then implemented within the organisation and evaluated to determine their effect on the quality of data in the organisation's security incident response process.

The analysis of organisation's security investigation records (Section 5.4.2) revealed that all 324 (100%) of the examined investigation records contained both a category and sub-category classification. However, when a more in-depth analysis of the investigation records was undertaken, it was identified that 25 of the records were incorrectly classified, according to the taxonomy. Interviews with individuals in the organisation revealed two main contributing factors which could explain the incorrect classifications: a lack of agreement on the definition of a 'security event' and a 'security incident' within the organisation, and incident handlers underestimating the impact of specific security problems. Moreover, the interviews also identified a potential managerial culture issue, which could also explain the incorrect classifications. The taxonomy, described in Section 5.2.1, was developed together with one of the organisation's information security managers who had requested the use of mutually exclusive sub-categories. However, interviews with the incident handlers showed that they would have preferred to select multiple categories during their investigation, instead of mutually exclusive sub-categories. What this finding suggested was that incident handlers were either not consulted or were over-ruled by managers with regards to their demands for non-exclusive sub-categories. Effectively, management undertook a decision to use mutually-exclusive sub-categories, which did not reflect the demands of the incident handlers in the organisation. In this sense, the manager has decided what information has value to the security incident response team, which provides an indication of a means-orientated culture by management responsible for the security incident response team.

The experiment also evaluated a revised security investigation record. The results show that the addition of the revised investigation record template has enhanced the overall investigation process and incident handlers are now documenting more detailed information about a security investigation. However, several observations were made from the analysis of the results from the experiment. The first of these observations is that the culture within an organisation can impact on the quality of data in security investigation records. In the specific organisation, it was identified that there is a potential 'blame culture' in the organisation. This was evident in the experiment when incident handlers stated that the reason they did not capture information requested by management was because it was 'someone else's job' (i.e. individuals who implement security controls). In this case, individuals attempted to protect themselves by shifting blame on why certain information was not capture and therefore, hindered continuous improvement.

A second observation from the experiment was the impact that a highly regulated environment can have on the quality of data documented by incident handlers. In the experiment, it was identified that incident handlers documented more information for a security 'incident' investigation when compared to a security 'event' investigation. One explanation for this is that the financial services industry is highly regulated and the reporting of security incidents is likely to be a priority of the organisation. However, this does not mean that the organisation itself should refrain from attempting to learn from security events. There have been suggestions in the literature that an organisation can learn just as much from low-level security events than high-level security incidents and therefore should focus on learning regardless of the level of severity.

A third observation that was made during the experiment is related to the amount of information captured and the quality of this information. Section 5.4.7 discussed how even though security incident handlers were now collecting more information in the organisation's investigation records, a lot of this information would potentially be of little use towards incident learning. A lesson that can be learnt from this experiment is that in order to enhance information capture towards incident learning, an organisation must first establish what objectives it will want to achieve when it attempts to learn about a security event or incident. Only then can the organisation decide what information it should attempt to capture during investigations. In this experiment, management decided on which information may be useful for incident learning. However, it can be argued that some of the fields (e.g. how many hours working on the investigation) were less important when it comes to learning from an incident and wider organisational learning, but more suited to providing metrics about resolution and recovery.

A final observation that was made during the experiment was the need for consolidated data quality improvement initiatives when attempting to improve security incident response data quality. The results from the experiment revealed that incident handlers often found that other organisational processes and data sources impacted the security incident response team. As a result, the data recorded in the investigation records was incomplete because it was missing in other business processes. Therefore, a suggestion that arises from this thesis is that organisation should look to address data quality issues in all processes, which input data into a security incident response process, and not just those which are used by the team during their investigations.

In **Chapter 6**, an experiment was undertaken in the Fortune 500 organisation which involved integrating retrospectives into the organisation's security incident response process. The purpose of the experiment was to evaluate if a retrospective can assist with validating information collected from an investigation. Furthermore, the experiment also evaluated if the retrospective can also be used to collect additional information that may have been missed during an investigation, without requiring substantial extra resources. The results from the experiment suggest that the completion of a security investigation and the closure of its corresponding investigation record does not necessarily mean that an incident handler has documented all the information about a specific investigation. In particular, the results show that 148 (46%) out of the 324 retrospectives contained more information about an investigation, when compared with the information documented in the corresponding record. 151 (47%) out of the 324 retrospectives identified the same information about an investigation as the information documented in the relevant record. Finally, 25 (7%) investigation records contained more information than what was actually identified using the retrospectives.

However, the results from the experiment also revealed challenges with process and management demands. This was particularly evident with the data collected from questions two and four from the retrospectives, which identified the presence of a 'process culture' in the organisation. In this case, the experiment identified a problem with the organisation's security incident response process, which does not take into consideration that incident handlers may not have access to some assets or individuals required for their investigations. Incident handlers are required to follow the process, but it does not take into consideration adaptability and changes to the process for security incident response. Ultimately, this can mean that assets are not investigated and root causes of incidents may not be found. The results from the experiment also suggest that retrospectives could be one way an organisation receives feedback on how its security incident response process works, which is often discussed as another problem with a 'process culture' in the literature.

While the results from the experiments suggest that additional information can be captured using retrospectives, observations and interviews within the organisation indicate that the information captured is more likely to be useful to management rather than the security incident team. While one could argue that some form of learning was taking place by management who were using the information to learn if the process was working, this same information could also have been used to learn about security incidents or how to prevent them. However, observations suggested that minimal support was given to the security incident response team in order for them to learn about incidents using the collected information. This is a typical problem, which can arise in an organisation in which a process culture is present and where internal politics, bureaucracy and problems with processes can hinder improvements and feedback.

In **Chapter 7**, a security incident response dashboard was implemented in the Fortune 500 organisation in order to enhance the transparency of data quality issues and to assist the security incident response team to identify and correct incomplete security investigation records. During the design of the dashboard, input was obtained from the one of the organisation's information security managers regarding the identification of fields that would result in an 'incomplete' investigation record in the organisation. The choice of fields from the manager shows that particular metric information regarding the time to eradicate and recover from a security event or incident have been favoured over information to assist with organisational learning. The recommendations made also indicated that it is managers and not the security incident response team who decide what information has value in this context.

Nonetheless, the dashboard application was deployed in the organisation and was made available to the security incident response team in order to identify and correct 'incomplete' investigation records. The results from this experiment have shown that 61 security investigation records which were identified as incomplete were corrected by the security incident handlers. The 61 records were missing information from one or more fields designated by the information security manager at the design of the dashboard. While the data collected from the investigation records suggests that the dashboard has improved the completeness of the 61 records, it can be difficult to isolate if the dashboard was the primary reason for this change. An alternative view is that the dashboard has assisted with improving group and investigation record awareness so that security incident handlers are now more aware of who they are working with, what is being worked on and how their actions affect others in the security incident response team. Evidence to support this argument can be seen in the interviews with participants who indicated that the dashboard allowed the team to identify which investigation records need more information, what each incident handler was working on and the progress of specific investigations.

Chapter 8, presents a root cause analysis framework for the analysis of security incidents within the organisation. The approach is an extension of the 5-Whys root cause technique that includes two components: the framework (Section 8.2.1) and the worksheet (Section 8.2.2), which used to document the outcomes from the root cause analysis. The framework was first applied to a historical set of investigation records within the organisation. The results from the historical analysis highlighted that if more information was documented during investigations, then enhanced underlying causes could have been identified from the investigations. In addition, numerous challenges with regards to undertaking a root cause analysis using the organisations security investigation records were identified. These challenges included poor data quality for in-depth analysis, a need to document more information describing why security incidents have happened rather than what has happened, as well as the impact of third parties in the root cause analysis process.

The framework was then applied to three 'live' security investigations, where a parallel investigation to that being conducted by the security incident handlers was undertaken. The results of this analysis showed that several underlying causes of security incidents can identified using the framework. However, follow-up interviews undertaken within the organisation which were used to assess practitioner perception of the framework revealed interesting perspectives on the use of framework in security incident response. In particular, individuals in the follow-up interviews stated that if a security problem is caused by human error, the likely outcome would be a root cause that cannot be fixed. This answer suggests the presence of a blame culture in the organisation. By transferring the blame from the security team to the human element, the incident response team can avoid blaming colleagues and potential allies in the Information Security unit. The presence of a blame culture could also point to managerial culture issues, who appear to quickly blame employees for even the smallest of errors that lead to security incidents and not security processes or failures in security controls. An alternative could be installation of a 'just culture', so that only more serious incidents result in employees facing disciplinary matters. This will mean that security incident handlers would not be afraid of examining 'human error' as a root cause, as only very serious 'human error' incidents would result in repercussions.

A further observation from the application of the framework to the live investigation was that incident handlers who were approached to undertaken an analysis using the framework indicated that some investigations were not 'serious enough' to undertaken a full root cause analysis. When incident handlers were asked about this in the follow-up interviews, the answers provided suggested that not everyone had completely 'bought-in' to the idea of undertaking a root cause analysis for particular incidents. Furthermore, in the months which followed the experiment, revealed that none of the improvements identified from the root cause analysis were actually implemented. The comments made during the interviews, coupled with these observations suggested a management culture issue, particularly with regards to leadership. The implementation of a root cause analysis implies that management supports and encourages the process in order to develop new understandings by drilling down precisely and narrowly into the causes of an incident. One recommendation is that management encourage incident handlers to undertake a root cause analysis because the results from this analysis will provide managers with practical guidance on problems that they can change and improve within the organisation.

9.5 Interconnection of Experiments

The four experiments presented in Chapter 5-8 have demonstrated how various lightweight measures can be used to improve the quality of data within different stages of a security incident response process. However, there are also interconnections between the experiments and the results have shown that collectively, the lightweight measures can be used to improve the quality of data throughout a security incident investigation. These interconnections mean that each implemented lightweight measure can complement other lightweight measures in a variety of ways, although a risk is that this can complicate the evaluation process. These interconnections are discussed here.

The taxonomy and security investigation record template evaluated in Chapter 5, provide a security incident response team with an environment to focus investigation efforts on data that will provide value both at the start and throughout the investigation lifecycle. The taxonomy helps with defining the security incident type at the start of the investigation, while the record template ensures that the data collected throughout the investigation will provide value to the security incident response team.

The root cause analysis framework evaluated in Chapter 8 can be used to guide a security investigation so that higher quality data is generated for lessons learned. This is achieved through a root cause analysis method based on the 5-Whys technique. The method prompts security incident handlers to ask "Why?" during a security investigation and in the process, enhances the quality of data captured during an actual investigation. Difficulties in eliciting answers during a root cause analysis are potential opportunities for process improvement that can be identified during a subsequent retrospective.

The security incident investigation dashboard, as presented in Chapter 7, can be used throughout a security investigation to enhance data quality transparency. The dashboard can also be used to assist a security incident response team to identify where the quality of data within security investigation records needs to be improved. The results from the evaluation of the dashboard have shown that the security incident response team corrected a number of 'incomplete' investigation records with more information.

Finally, the retrospectives as evaluated in Chapter 6, provided security incident handlers with a method to confirm the data documented during an investigation. The retrospectives have also provided a 'safety net' to capture additional information, which may have been missed during the initial investigation.

9.6 Scope and Validity

Scope and validity issues associated with the research need to be specifically acknowledged at this point. The first point of discussion is the ability to generalise from the Fortune 500 Organisation case study, which can be considered a 'single case' study. The view that one cannot generalise on the basis of a single case study is an issue that has been well discussed in the literature [36, 39, 257]. As noted in Section 2.3, the Fortune 500 Organisation case study can be considered a 'special case' in that a unique chance arose to study security incident learning in a large organisation [30]. In this sense, the Fortune 500 Organisation is considered to representative of other organizations in the financial services category in terms of interest in security incident response, regulatory obligations towards security incident reporting, size, bureaucracy and resource constraints. This makes the organisation ideal in which to study security incident response learning challenges.

Furthermore, Walsham suggests that four main types of generalisations are possible from case studies: concepts, theory, implications and rich insight [258]. The results from the

Fortune 500 Organisation case study can be generalised from the rich insight perspective. In particular, the case study can provide a rich insight into a variety of topics including using specific practices to improve the quality of data in security incident response, managerial and cultural issues related to security incident learning, and the need for more thoughtful security incident data gathering.

It should be acknowledged that the lightweight practices were implemented in the same organisation, where the exploratory case study was undertaken to determine if the quality of data was a potential problem within security incident response. While this is not negative, it does raise the question of whether the lightweight practices are applicable in other organisations. The practices, which were implemented in the Fortune 500 Organisation, can be generalised from the perspective that they can be used in similar organisations, which have a similar interest in security incident response. More specifically, the security incident response taxonomy and enhanced investigation can be customised for the specific organisation in which they will be used. As a result, the information that is collected by incident handlers will provide value to that specific organisation. The retrospectives and the dashboard can also be customised to suit the requirements of the organisation in which they will be used. The questions from the retrospectives can be customised to capture specific information, which should be validated, while the dashboard can also be customised to highlight specific information, which has been missed in initial investigations. In both cases, the customisations will be specific to the implementing organisation, but the practices themselves can remain the same. Similarly, while the foundational components of the root cause analysis framework are transferable to other organisations, the questions being asked during the framework can be customised to the specific organisation.

The case study presented in this thesis has also presented a rich insight into the culture issues related to improving the quality of data in security incident response. While many of the culture challenges identified during the experiments would likely be also found in similar financial services-type organisations, some of these culture challenges could be specific to the studied organisation. These are likely to be the specific culture challenges related to management within the organisation and their impact on the security incident response team.

It should also be acknowledged that the author worked visibly within the Fortune 500 Organisation to undertake the exploratory case study and to implement the lightweight measures. This raises the question of the author's impact on the results of the work conducted. In empirical studies such as the one documented in this thesis, it can become difficult to know how much influence the author has had over the success of an experiment. The fact that the author was integrated into the organisation's security incident response team and the presence of the author in the team and asking questions about the quality of data within security investigations, needs to be acknowledged in reference to potentially impacting the study. Finally, it should also be noted that the results from the experiments to evaluate the lightweight measures, implemented within the organisation, could impact each other. For example, the experiment to evaluate the dashboard application could have influenced the results from the revised security incident investigation record template experiment. While specific fields were used to identify an ‘incomplete’ record in the dashboard experiment, the notification of these fields to security incident handlers could have prompted the completion of other fields in the record. This in turn could have influenced the results from the investigation record template experiment.

9.7 Successful Security Incident Learning

It should be stressed that the work in thesis should be treated as a proof-of-concept. Hence, the insights and challenges identified in the thesis need to be investigated further in order to understand the implications of the practices and culture challenges in other organisations. However for the sake of completion, this section will provide some suggestions based on the results of this thesis in order to overcome some of the barriers to security incident learning identified during the research. In this sense, the suggestions are based on what successful security incident learning would look like in the Fortune 500 Organisation. As the results from Chapters 5 and 6 have shown, the building blocks of security incident learning should start with collecting and validating the *right* information during a security investigation. In this sense, information captured should include data that will provide the opportunity for an organisation to learn from a security incident, and increase its chances to prevent a reoccurrence. For this to be a success, managers need to work closely with security incident handlers and individuals responsible for securing the organisation in order to identify what information would be need to initiate improvement changes. Furthermore, the organisation should not only look at its data quality in security and incident response-related processes, but extend this to any process which provides data to security-related investigations. As discussed in Section 5.5.3, the Fortune 500 Organisation’s security incident response team could not improve the quality of data in its investigation records because it relied on the data from another process outside of its control.

Once the right information has been captured, in order for the security incident response team to maximise their incident learning capabilities, the organisation should look to incorporate double-loop learning. Double-loop learning would involve the team questioning and challenging the underlying rules, principles and knowledge of the organisation. In other words, double-loop learning would question the culture and processes within the organisation itself. It is then important to implement corrective actions and follow-up any recommendations made by the security incident response team. This is particularly true for actions to eliminate causes of incidents which could span many parts of the organisation and which could involve various different processes. As discussed in Chapter 6, the retrospectives can be one way of evaluating if outstanding recommendations have been corrected in the organisation. Finally, the work of the security incident response team will be completed when it documents its detailed recommendations and findings either in the Information Security Incident Response database or through the issuance of an incident report.

While the above process adjustments would help enhance security incident learning in the organisation, specific culture issues also need to be addressed. One of these culture concerns is the management culture issue identified at various stages in the thesis. For example, managers selected what information provides value to the security incident response team through their choice of fields in the investigation record template and their choice of mutually exclusive sub-categories in the implemented taxonomy. In order for the security incident response team to be accountable for their actions, they need to be trusted to select which information they demand is needed in order to learn from a particular security incident. However, at this point it is worth stating that in a process-based culture, such as the Fortune 500 Organisation which is highly-regulated, changing and removing the process culture could be a difficult task. This is because these types of organisations grow in terms of knowledge, scale of operations and efficiency irrespective of the contribution of its members. Therefore, changing the management culture could actually have little or no impact on the underlying organisational culture.

Another culture issue, which would be important to address is the ‘blame culture’ which appears to exist in the organisation. The results from the experiments suggest that a blame culture in the organisation appears to result in individuals protecting themselves by shifting blame and more importantly, hindering continuous improvement. In an effort to correct this culture of blame, management need to accept that individuals can make mistakes, and that sometimes it is better to learn from these mistakes rather than punish individuals through disciplinary affairs. The installation of a strong culture of transparency and accountability would help focus the security incident response team’s efforts in not only eradicating and recovering from security incidents, but also learning from underlying issues which have caused the incident to occur.

9.8 Future Work

The research reported in this thesis has identified several areas for future work. In the short term, there is a need for further work to examine the quality of data within security incident response in other organisations, in a variety of industrial domains. The exploratory case study presented and discussed in Chapter 4 can be repeated in other organisations to strengthen or dispute the theoretical argument presented in the thesis. That is the quality of data in security incident response needs to be enhanced to improve the generation of lessons learned. This future work can also be used to evaluate the lightweight measures presented in the thesis in other organisations in order to enhance and refine their application as methods and tools to enhance the quality of data within security investigations. Experience reports and results from the evaluation of the lightweight measures would help alleviate some of the issues associated with undertaking a single-case research project as discussed in Section 9.5.

While the implementation of these lightweight measures in other organisations will help enhance the use of these measures, there is a need to further explore the business perspective for improving the quality of data within security incident response. As highlighted in earlier chapters, there was some resistance to implementing the lightweight measures that would be used to improve the quality of data within security investigations. Therefore, future work should look to explore the actual and/or perceived Return on Investment (ROI) on improving the quality of data within security investigations. The result of this analysis can be examined in several different types of organisations. The purpose of this work would be to investigate which organisations would be more willing to undertake the investment of improving the quality of data and which would struggle to justify the time and costs associated with such an improvement initiative. The results can then be used to assist adopters examine ROI from the measures within their specific industry domain.

Previous chapters in the thesis discussed how the security incident response team within the Fortune 500 Organisation indicated that data quality issues involving the organisation's electronic address book had implications on data quality within security incident response. This finding shows that there is a clear need to research the impact of improving the quality of data within other business processes that input data into a security incident response process. This future work can be extended to explore the possibility of applying some or all of the lightweight measures described in this research in other teams, apart from security, within an organisation. This future work can also examine the legislative perspective with regards to enhancing the quality of data in other business process, with a particular emphasis on the privacy and data retention legislative requirements that are being imposed on organisations.

In the future, research needs to be conducted into the barriers and aids that third-parties introduce into an organisation's security incident response process. The results from the implementation of the lightweight measures, particularly the root cause analysis framework, highlighted the impact of third-parties on a security incident response team's investigation. Therefore, research to examine the reduction and/or elimination of barriers presented by third-parties to security investigations could provide valuable information to organisations.

In the longer term, future research needs to examine the restructuring of security incident response teams. The results from the experiments presented in this thesis have shown that the success of some security investigations often depends on the team's relationship with other business units and teams within an organisation. This is also true for investigations that involve third-parties. Therefore, future work needs to examine if the structure and practices of interaction between security incident response teams and these entities needs to be reconsidered. Effectively, this could result in multidisciplinary security incident response teams, which would require the integration of technical security professionals, information technology specialists, relevant asset stakeholders, third-party contractors, along with individuals from an organisation's legal department.

9.9 Summary

The research presented in this thesis examined how specific lightweight measures can be used to enhance the quality of data produced during security incident response investigations. The lightweight measures were implemented in an industrial case study and their impact was examined with regards to improving the quality of data in various phases of a security investigation. The high-level finding from these evaluations showed that the lightweight measures have had a positive affect on the quality of data generated during and after security investigations. The experience and knowledge gained throughout the research provides foundational work to generalise the use of this approach in other organisations and in a variety of industries.

Appendix A

Exploratory Case Study Interview Questions

Thank you, for participating in this interview survey. The aim of this exercise is to obtain a more in-depth understanding of the security incident response process within the organisation and to investigate the perception of organisational learning and knowledge gathering from within this process. This is not a test and you are not being examined. Hence, there are no right or wrong answers; it is your opinion that is being sought. Please be assured that this interview survey will be conducted with your anonymity ensured and that I will not record or disclose any personal information.

I ask you to please not discuss the interview survey with anyone else in the organisation as this may invalidate the results.

1. What is your current job title/role?
2. How many years have you worked in IT?
3. Briefly describe the key areas of your job function/role?
4. From your perspective, what is meant by the term ‘security incident’?
5. Does the organisation have a security incident response team? YES/NO/Don’t Know (DNK)
   (a) If YES, what are the overall goals of the security incident response team?
   (b) If NO, are there plans to develop one in the future? YES/NO/DNK
      (i) If YES, what is the projected time frame?
      (ii) If NO, is there a reason for not developing a security incident response team?

6. Does the organisation have a documented security incident response process? YES/NO/ DNK
   (a) If YES, can you briefly describe this documented process and the activities involved in the process? In the event participant cannot recall the exact process, he/she will be asked to recall to the best of their ability any phases of the process to stimulate discussion.
   (b) If YES, in your opinion, does this documented process ensure that the goals of the security incident response team are met?
   (c) If YES, in your opinion what are the good points in this documented process?
   (d) If YES, in your opinion what are the bad points in this documented process?
   (e) If YES, have you ever found it necessary to deviate from this documented incident response process? YES/NO/DNK
      (i) If YES, why?
   (f) If NO, are there plans to develop one in the future? YES/NO/DNK
      (i) If YES, what is your projected time frame?
      (ii) If NO, is there a reason for not developing such a documented process in the organisation?

7. Is there an individual in the organisation who is accountable for the documented security incident response process not being followed? YES/NO/DNK
   (a) If YES, what is his/her title? Is this person accountable for legislative compliance as well? YES/NO/DNK
      (i) If NO, who in the organisation is accountable for legislative compliance?
   (b) If NO, why not?

8. In your opinion, what are an individual's responsibilities for each the following categories in the security incident response lifecycle?
   (a) Preparation
   (b) Identification
   (c) Containment
   (d) Eradication
   (e) Investigation
   (f) Recovery
   (g) Follow-up (Post-incident)

9. Are you involved in any of the above categories of the security incident response lifecycle? YES/NO

    (a) If YES, which categories?

10. In your opinion, do you think that the documented security incident response process could be improved? YES/NO/DNK

    (a) If YES, how and why?

    (b) If NO, why not?

11. Does the organisation have a documented procedure outlining how security incidents are reported within the organisation? YES/NO/DNK

    (a) If YES, what is this procedure?

    (b) If YES, in your opinion, is this security incident reporting procedure effective? YES/NO/SOMETIMES/DNK

        (i) If YES, why?

        (ii) If NO, why not?

        (iii) If SOMETIMES, when is it effective?

        (iv) If SOMETIMES, when is it not effective?

    (c) If NO, how are security incidents reported?

    (d) Are employees within the organisation educated on how to report security incidents? YES/NO/DNK

12. Does the organisation record and store information related to security incidents? YES/NO/DNK

    (a) If YES, can you briefly describe this practice and what information is recorded for security incidents?

    (b) If YES, in your opinion, is the practice used for recording information related to security incidents effective? YES/NO/SOMETIMES/DNK

        (i) If NO, why not?

        (ii) If SOMETIMES, when is it effective?

        (iii) If SOMETIMES, when is it not effective?

    (c) If YES, where are these incidents recorded?

    (d) If YES, who has access to these incident reports (groups and departments)?

    (e) If YES, is this procedure used to document and recorded all types of security incidents? YES/NO/DNK

(i) If NO, which incidents types are not recorded using this procedure?

(ii) If NO, where are these incidents recorded?

(f) If NO, why not?

13. Is there a documented process for collecting 'forensic data' related to security incidents such as (forensic images, logs etc.)? YES/NO/DNK

(a) If YES, can you briefly describe this process?

(b) If YES, does this process include a 'chain of custody' document or is an audit trail maintained? YES/NO/DNK

(i) If YES, can you briefly discuss this process?

(ii) If NO, why not?

(c) If NO, are there plans to develop one in the future? YES/NO/DNK

(i) If YES, what is your projected time frame?

(ii) If NO, is there a reason for not developing such a documented process in the organisation?

14. In your experience, do situations occur where an incident handler does not have access to all the information he/she requires (forensic images, logs etc.), when investigating security incidents? YES/NO/DNK

(a) If YES, what types of situations?

(b) If YES, what information was required?

(c) If NO, do you for see this as a problem in the future? YES/NO/DNK

(i) If YES, in what way?

(ii) If NO, why do you think this will not be a problem in the future?

15. In your experience, do conflicts arise between different stakeholders within the organisation when investigating a security incident? YES/NO/DNK

(a) If YES, what are these conflicts?

(b) If YES, how are these conflicts resolved?

16. In your experience, what constitutes the closure of an information security incident within the organisation?

(a) Can you recall an instance(s) where an incident has not been closed? YES/NO/DNK

(i) If YES, why was the incident not closed?

(b) Can you recall an instance(s) where an incident has been reopened? YES/NO/DNK

    (i) If YES, why was the incident reopened?

17. From your experience, does the organisation perform any 'post-incident' activities after a security incident has been closed? YES/NO/DNK

    In the event participant answers NO or DNK, they will be asked questions 17a-d seeking their opinions on what these activities should involve, the idea is to promote discussion around the topic of post-incident activities.

    (a) If YES, what activities take place 'post-incident'?

    (b) If YES, in your opinion do you think these activities are sufficient? YES/NO/DNK

        (i) If NO, what 'post-incident' activities do you think should occur?

    (c) If YES, do these activities include a root cause analysis? YES/NO/DNK

        (i) If NO, why not?

    (d) If YES, do these activities include a risk analysis based on the incident? YES/NO/DNK

        (i) If NO, why not?

    (e) If NO, why not?

18. Does the organisation distribute any 'post-incident' information to any group(s) or department(s) within the organisation? YES/NO/DNK

    (a) If YES, is there a formal process for disseminating this 'post-incident' knowledge? YES/NO/DNK

        (i) If YES, can you briefly describe this process?

        (ii) If NO, why not?

    (b) If YES, is there an individual in the organisation who is responsible for disseminating 'post-incident' knowledge within the organisation?

        (i) If YES, what is his/her title?

        (ii) If NO, why not?

    (c) If YES, to which groups or departments?

    (d) If YES, what type of information is circulated?

    (e) If NO, why not?

19. Does the organisation distribute any 'post-incident' information or incident notifications outside of the organisation, for example to regulatory bodies? YES/NO/DNK

(a) If YES, is there a formal process for disseminating this 'post-incident' knowledge? YES/NO/DNK

(i) If YES, can you briefly describe this process?

(ii) If NO, why not?

(b) If YES, is there an individual in the organisation who is responsible for disseminating 'post-incident' knowledge outside the organisation?

(i) If YES, what is his/her title?

(ii) If NO, why not?

(c) If YES, where is this information circulated?

(d) If YES, what type of information is circulated?

(e) If NO, why not?

20. Are you aware of any individual in the organisation who is responsible for analysing previous security incidents to identity trends, anomalies, and patterns? YES/NO/DNK

(a) IF YES, what is his/her title?

(b) If NO, how are trends, anomalies, and patterns identified from security incidents?

21. Does the organisation consult or contract any individuals/ contractors/businesses to help detect, investigate, eradicate or recover from a security incident? (i.e. forensic analysis work) YES/NO/DNK

(a) If YES, can you briefly describe an instance where an individual/contractor/ business has been included in the handling of an incident?

(b) If YES, do individuals/contractors/businesses follow the same security incident response process as employees? YES/NO/DNK

(c) If NO, do you think this will occur in the future?

22. Does the organisation have a documented secure development process (SDLC) or policies around the secure development of systems and applications? YES/NO/DNK

In the event participant answers DNK, they will be asked question 22d to promote discussion around the idea or integrating lessons learned into the application development environment.

(a) If YES, can you briefly describe the organisations secure development process?

(b) If YES, in your opinion what are the good points in this documented process?

(c) If YES, in your opinion what are the bad points in this documented process?

(d) If YES, in your opinion, do you think that lessons learned from previous security incidents should play a role in the secure development process? YES/NO/DNK

(i) If YES, why so?

(ii) If NO, why not?

(e) If NO, why not?

23. Does the organisation have an information security education program? YES/NO/DNK

(a) If YES, can you briefly highlight the purpose of the program?

(b) If YES, do you know if any information from previous security incidents is included in the security education program? YES/NO/DNK

(i) If YES, can you briefly highlight what this information contains?

(ii) If NO, are there plans to include such information in the future? YES/NO/DNK

(c) If NO, are there plans to develop one in the future? YES/NO/DNK

(i) If YES, what is your projected time frame?

(ii) If NO, is there a reason for not developing an information security education program in the organisation?

24. Were any of the survey questions vague or difficult to follow?

25. Do you have any additional information you would like to add? Is there anyone else you suggest I talk to regarding any information in this survey?

Appendix B

Retrospectives Follow-up Interview Questions

Thank you, for participating in this follow-up interview survey. The purpose of this exercise is to obtain a more in-depth understanding of the affect and impact of the agile retrospectives on the security incident response process within the organisation. This is not a test and you are not being examined. Hence, there are no right or wrong answers; it is your opinion that is being sought. Please be assured that this interview survey will be conducted with your anonymity ensured and that I will not record or disclose any personal information.

I ask you to please not discuss the interview survey with anyone else in the organisation as this may invalidate the results.

1. What is your job title?

2. Briefly describe your role within the security incident response team?

3. Have you ever been involved in any security post-investigation activities within the organisation? YES/NO

   (a) If YES, what activities have you been involved in?

   (b) If NO, why not?

4. Did you participate in the briefing about the retrospectives? YES/NO

   (a) If YES, in your opinion what is the advantage of conducting the briefing?

   (b) If YES, in your opinion what is the disadvantage of conducting the briefing?

   (c) If NO, why not?

5. In your opinion, what is the advantage of using the retrospectives to identify 'what worked well' (i.e. assets investigated and individuals communicated with during an event/incident)?

6. In your opinion, what is the disadvantage of using the retrospectives to identify ‘what worked well’ (i.e. assets investigated and individuals communicated with during an event/incident)?

7. Overall, what influence did the retrospectives have on the identification of assets investigated and individuals communicated with during an event/incident?

8. In your opinion, what is the advantage of using the retrospectives to identify ‘what did not work well’ (i.e. assets which you could not investigate and individuals you could not communicate with during an event/incident)?

9. In your opinion, what is the disadvantage of using the retrospectives to identify ‘what did not work well’ (i.e. assets which you could not investigate and individuals you could not communicate with during an event/incident)?

10. Overall, what influence did the retrospectives have on the identification of assets not investigated and individuals that could not be communicated with during an event/incident?

11. In your opinion, what is the advantage of using the retrospectives to identify security controls that can be improved after a security event/incident?

12. In your opinion, what is the disadvantage of using the retrospectives to identify security controls that can be improved after a security event/incident?

13. Overall, what influence did the retrospectives have on the identification of security controls that can be improved after a security event/incident?

14. In your opinion, what is the advantage of using the retrospectives to identify incident response- related process enhancements that be improved after a security event/incident?

15. In your opinion, what is the disadvantage of using the retrospectives to identify incident response- related process enhancements that be improved after a security event/incident?

16. Overall, what influence did the retrospectives have on the identification of incident response-related process enhancements?

17. In your opinion, what other factor contributed to the successful or unsuccessful attempt to identify security control and incident response-related process enhancements using the retrospectives?

18. Were any survey questions difficult to follow? Are there any additional comments that you would like to make about the questions?

Appendix C

Follow-up Interview Questions

Thank you, for participating in this follow-up interview survey. The purpose of this exercise is to obtain a more in-depth understanding of the affect and impact of the agile retrospectives on the security incident response process within the organisation. This is not a test and you are not being examined. Hence, there are no right or wrong answers; it is your opinion that is being sought. Please be assured that this interview survey will be conducted with your anonymity ensured and that I will not record or disclose any personal information.

I ask you to please not discuss the interview survey with anyone else in the organisation as this may invalidate the results.

1. What is your current job title/role and briefly describe your role within the security incident response team?
2. How many years experience do you have within an information security role?
3. In your experience, what is the role of security incident response within the general Information Security team?
4. Have you observed any changes in the way security events and incidents are documented and recorded within the organisation in the past year? YES/NO
   (a) If YES, what differences?
   (b) If NO, why not?
5. Did you encounter any problems with categorising and/or classifying a security occurrence using the definition classification and categories in the documented security incident response process? YES/NO
   (a) If YES, please elaborate on these problems

(b) If NO, what problems do you foresee in using the definition classification and categories?

6. Did you encounter any benefits with categorising and/or classifying a security occurrence using the definition classification and categories in the documented security incident response process? YES/NO

   (a) If YES, what benefits

   (b) If NO, why not?

7. Did the addition of the definition classification and categories to the security incident response process assist or hinder the overall investigation process? Please explain why.

   The interviewees were shown five examples of security investigation records, which were completed by incident handlers over the duration of the case study since the data generation improvement imitative was implemented. Each investigation record shown will according to the definitions in documented process, either be considered a misclassified or a miscategorised record. Interviewees will then be asked Q8 and 9 for event of the five records.

8. In your opinion, does record *XZY* have the correct classification based on the definition criteria in the documented process? YES/NO

   (a) If YES, why?

   (b) If NO, why not?

9. In your opinion, does record XYZ have the correct category type based on the category criteria in the documented process? YES/NO

   (a) If YES, why?

   (b) If NO, why not?

10. In your opinion, does the current security incident record template capture all relevant information about a security event? YES/NO

    (a) If NO, what additional information needs to be captured?

11. In your opinion, does the current security incident record template capture all relevant information about a security incident? YES/NO

    (a) If NO, what additional information needs to be captured?

12. In your opinion, are there any fields within the security investigation record template which are difficult to complete? YES/NO

    (a) If YES, which fields and why?

13. Did you encounter any benefits with the completion of the current security investigation record template? YES/NO

    (a) If YES, what benefits?

    (b) If NO, why not?

14. Did the introduction of the new security investigation record template assist or hinder the overall investigation process? Please explain why?

15. Did the addition of the Dashboard assist or hinder the overall investigation process? Please explain why?

16. From your experience, did you encounter any benefits in using the dashboard? YES/NO

    (a) If YES, what?

    (b) If NO, why not?

17. From your experience, did you encounter any problems in using the dashboard? YES/NO

    (a) If YES, what problems?

    (b) If NO, what problems do you foresee in using the dashboard within the team?

18. Which feature(s) of the dashboard did you use most? Why?

19. Which feature(s) of the dashboard did you use the least? Why?

20. Did the dashboard have an impact on collaboration opportunities with fellow incident response handlers? YES/NO

    (a) If YES, how?

    (b) If No, could you elaborate why not?

21. From your experience, did you feel discouraged, stressed or annoyed when attempting to identify incomplete records using the dashboard? YES/NO/SOMETIMES

    (a) If YES, in what way?

    (b) If NO, why not?

   (c) If SOMETIMES, in what case(s)?

22. From your experience, do you feel that you were successful in using the dashboard to identify incomplete records YES/NO/SOMETIMES

    (a) If YES, in what way?
    (b) If NO, why not?
    (c) If SOMETIMES, in what case(s)?

23. From your experience, did you have to work hard to identify incomplete records using the dashboard? YES/NO/SOMETIMES

    (a) If YES, in what way?
    (b) If NO, why not?
    (c) If SOMETIMES, in what case(s)?

24. Did you participate in any of the 5-whys root cause analysis exercises? YES/NO

    (a) If YES to Q24, go to Q25 15, if NO, go to Q26

25. What factors contributed to the successful or unsuccessful use the 5-whys method for root cause analysis within the security incident response team?

    A Root Cause Analysis (RCA) approach was designed as shown below (see next page). The purpose of the approach is to support security incident response teams with regards to identifying when to undertake a RCA, how to start a RCA and how to end a RCA. (The respondents will have the approach explained step-by-step).

26. What do you see as the perceived strength of using this approach to evaluate how to start, perform and end a root cause analysis?

27. What do you see as the perceived weakness of using this approach to evaluate how to start, perform and end a root cause analysis?

28. Would the addition of the proposed RCA approach will assist or hinder the overall investigation process? Please explain why.

29. Would you change anything with regards to the designed approach? YES/NO

    (a) If YES, what?
    (b) If NO, why not?

30. Were any of the survey questions vague or difficult to follow? Are there any additional comments that you would like to make about the questions?

Appendix D

Description of Proposed Investigation Record Template Fields

Field Name	Description
Subject	Subject line provided by incident handlers which describes investigation record
Category	Security occurrence classification and category as assigned by the incident handler
Date Reported	The date the security occurrence was reported to the security incident response team
Time Reported	The time the security occurrence was reported to the security incident response team
Reported To:	The name of the individual to whom the security occurrence was reported to within the security incident response team
Reported By:	The name of the individual who has reported the security occurrence to the security incident response team
Date Discovered	The date the security occurrence was discovered within the organisation
Contact Name	The name of the individual who is the point of contact for the security occurrence and provide further information if required
Job Title	The job title of the individual named above
Telephone	The telephone number of the individual named above
Department	The name of the department of the individual named above
Business Unit	The name of the business unit of the individual named above
Line Manager	The name of the line manager for individual named above

Continued on next page

– table continued from previous page

Field Name	Description
Incident Handler	The name of the incident handler who is accountable for the particular investigation record
Status	Status of the investigation record (Open/Closed)
Date Opened	The date the security investigation record was opened by the security incident response team
Time Opened	The time the security investigation record was opened by the security incident response team
Location	The name of the location of where the occurrence has happened, either within or out-with the organisation
Investigation Record	Information about the investigation is recorded in this section
Date Closed	The date the security investigation record was closed by the security incident response team
Time Closed	The time the security investigation record was closed by the security incident response team
Actions to be Taken	Description of any actions that need to be taken by either the incident response team or the information unit in relation to the investigated occurrence
Lessons Learned	Four questions are asked in an attempt to produce lessons learned from the incident handler: what caused the incident? who caused the incident? how many records (if any) were involved? what present controls should have prevented the occurrence? what additional controls could have prevented the occurrence?
Working Hours	The number of working hours that have consumed on a particular security occurrence

Appendix E

Retrospectives Data

Please note, that the names of specific assets, individuals, documents and security processes identified in the retrospectives have been altered to protect the information security interests of the Fortune 500 Organisation.

Question 1: Which assets did you need to investigate in this security event/ incident?

Asset Name	Number of Occurrences
Active Directory server	2
Organisation-specific asset 1	1
Intranet web page	3
Telephone records	1
Building access card logs	1
CD containing data	1
Data Loss Prevention software	146
Desktop, personal and laptop computers	6
Dynamic Host Configuration Protocol (DHCP) server logs	1
Email journals, restores and logs	10
Email messages and memos	100
Organisation-specific asset 2	1
Organisation-specific database 1	1
Organisation-specific asset 3	1
Organisation-specific asset 4	1
Organisation-specific asset 5	1
Live email account	37

Continued on next page

– table continued from previous page

Asset Name	Number of Occurrences
Lotus Notes email server and mail files	132
Organisation-specific asset 5	1
Organisation-specific asset 6	2
Organisation-specific asset 7	1
Network file servers and file shares	5
Organisation-specific asset 8	3
Organisation-specific asset 9	3
SharePoint server	3
Secure Remote Access Service (SRAS) registry	4
Secure Remote Access Service (SRAS) token	3
Third-party assets and logs	9
Organisation-specific asset 10	1
Universal Serial Bus (USB) storage device	1
Various data pieces	9
Virtual machines and logs	1
Voice recordings	6
Web gateway logs	2
Organisation-specific asset 11	1
WiFi logs	1
Total	**502**

Question 3: Who did you need to communicate with during this security event/ incident?

Individual/Organisational Team Identified	Number of Occurrences
Access control team	3
Application development team	1
Application engineering team	5
Application and software projects Team	3
Audit team	3
Branch Managers	1
Building and physical security team	4

Continued on next page

– table continued from previous page

Asset Name	Number of Occurrences
Contractors	1
Customer-specific business unit 1	1
Customer-specific business unit 2	2
Customer-specific business unit 3	1
Customer-specific business unit 4	1
Customer-specific business unit 5	4
Customer-specific business unit 6	1
Customer-specific business unit 7	2
Customer-specific business unit 8	3
Customer-specific business unit 9	5
Customer-specific business unit 10	1
Data storage team	3
Desktop Support Services	5
Direct banking manager	1
District managers	2
ECM designer and change development team	1
Enterprise communications	1
External legal firms	3
Fraud unit	24
Group incident response team	3
Head of organisational compliance	1
Head of collections	1
Human resources	102
Individuals affects by events and incidents	144
IM Team	1
IT strategy team	1
Information security management	8
Information security team	136
Intranet support team	1
Laptop owners	2
Legal Services	16
Line managers	32
Lotus notes email team	106
Management services team	2

Continued on next page

– table continued from previous page

Asset Name	Number of Occurrences
Parent Organisation individuals	1
Organisational risk and compliance team	8
People Leaders	37
Regulatory risk team	3
Service provision team	2
Service restoration team	1
Systems support	1
Technology Services	1
Third-party desktop support team	1
Third-party WiFi support team	1
Third-party relationship managers	32
Third-parties who provide services to the organisation	7
Windows 7 support team	1
Windows Server support team	3
Total	**737**

Bibliography

[1] PricewaterhouseCoopers, "Global State of Information Security Survey," Online at: www.pwc.com/gx/en/consulting-services/information-security-survey/, 2015.

[2] Ponemon Institute, "2014 Global Report on the Cost of Cyber Crime," Online at: http://www8.hp.com/us/en/software-solutions/ponemon-cyber-security-report/.

[3] T. R. Peltier, *Information Security Fundamentals*. CRC Press, 2013, no. 1st Edition.

[4] F. Farahmand, S. B. Navathe, P. H. Enslow, and G. P. Sharp, "Managing Vulnerabilities of Information Systems to Security Incidents," in *Proceedings of the 5th international Conference on Electronic Commerce*. ACM, 2003, pp. 348–354.

[5] The International Organization for Standardization and The International Electrotechnical Commission, "ISO/IEC 27001:2013 – Information Technology – Security Techniques – Information Security Management Systems – Requirements," 2013.

[6] ——, "ISO/IEC 27002:2013 – Information Technology – Security Techniques – Code of Practice for Information Security Controls," 2013.

[7] S. Mitropoulos, D. Patsos, and C. Douligeris, "On Incident Handling and Response: A State-Of-The-Art Approach," *Computers & Security*, vol. 25, no. 5, pp. 351–370, 2006.

[8] National Institute of Standards and Technology, "Computer Security Incident Handling Guide – Special Publication 800-61, Version 2," NIST, Tech. Rep., 2012.

[9] The European Union Agency for Network and Information Security, "Good Practice Guide for Incident Management," ENISA, Tech. Rep., 2010.

[10] The International Organization for Standardization and The International Electrotechnical Commission, "ISO/IEC 27035:2011 - Information Technology. Security Techniques. Information Security Incident Management," 2011.

[11] K. Mandia, C. Prosise, and M. Pepe, *Incident Response and Computer Forensics*. McGraw-Hill, 2003.

[12] M. Vangelos, “Incident Response: Managing,” Encyclopedia of Information Assurance, pp. 1442–1449, 2011.

[13] The SANS Institute, “Computer Security Incident Handling Version 2.3.1,” 2003.

[14] Y. He, “Generic Security Templates for Information System Security Arguments: Mapping Security Arguments Within Healthcare Systems,” Ph.D. dissertation, University of Glasgow, United Kingdom, 2014.

[15] C. Johnson, *A Handbook of Incident and Accident Reporting*. Glasgow University Press, 2003.

[16] P. Stephenson, “Conducting Incident Post Mortems,” *Computer Fraud & Security*, vol. 2003, no. 4, pp. 16–19, 2003.

[17] United States Department of Transportation, “Safety Data Action Plan - Safety in Numbers,” Prepared by the Bureau of Transportation Statistics, Tech. Rep., 2000.

[18] A. Ahmad, J. Hadgkiss, and A. B. Ruighaver, “Incident response teams – Challenges In Supporting the Organisational Security Function,” *Computers & Security*, vol. 31, no. 5, pp. 643–652, 2012.

[19] P. Shedden, A. Ahmad, and A. B. Ruighaver, “Informal Learning in Security Incident Response Teams,” in *2011 Australasian Conference on Information Systems*, 2011.

[20] T. Tan, T. Ruighaver, and A. Ahmad, “Incident Handling: Where The Need for Planning Is Often Not Recognised,” in *Proceedings of the 1st Australian Computer Network, Information & Forensics Conference, Perth Australia*, vol. 24, 2003.

[21] C. Batini and M. Scannapieca, *Data Quality: Concepts, Methodologies and Techniques*. Springer, 2006, ch. Data Quality Dimensions, pp. 19–49.

[22] V. Raman and J. M. Hellerstein, “Potter’s Wheel: An Interactive Data Cleaning System,” in *Proceedings of the 27th International Conference on Very Large Data Bases*, vol. 1, 2001, pp. 381–390.

[23] F. Caruso, M. Cochinwala, U. Ganapathy, G. Lalk, and P. Missier, “Telcordia’s Database Reconciliation and Data Quality Analysis Tool,” in *Proceedings of the 26th International Conference on Very Large Data Bases*, 2000, pp. 615–618.

[24] H. Galhardas, D. Florescu, D. Shasha, and E. Simon, “AJAX: An Extensible Data Cleaning Tool,” in *ACM Sigmod Record*, vol. 29, no. 2. ACM, 2000, p. 590.

[25] K. Beck, *Extreme Programming Explained: Embrace Change*. Addison-Wesley Professional, 2000.

[26] M. Fowler, “The New Methodology,” *Wuhan University Journal of Natural Sciences*, vol. 6, no. 1-2, pp. 12–24, 2001.

[27] R. Werlinger, K. Muldner, K. Hawkey, and K. Beznosov, “Preparation, Detection, and Analysis: The Diagnostic Work of IT Security Incident Response,” *Information Management & Computer Security*, vol. 18, no. 1, pp. 26–42, 2010.

[28] E. Casey, “Case Study: Network Intrusion Investigation–Lessons in Forensic Preparation,” *Digital Investigation*, vol. 2, no. 4, pp. 254–260, 2005.

[29] S. Easterbrook, J. Singer, M.-A. Storey, and D. Damian, “Selecting Empirical Methods For Software Engineering Research,” in *Guide to Advanced Empirical Software Engineering*. Springer, 2008, pp. 285–311.

[30] B. J. Oates, *Researching Information Systems and Computing*. Sage Publications, 2005.

[31] G. G. Gable, “Integrating case study and survey research methods: an example in information systems,” *European journal of information systems*, vol. 3, no. 2, pp. 112–126, 1994.

[32] R. K. Yin, *Case Study Research: Design and Methods*. Sage Publications, 2014.

[33] R. L. Baskerville, “Investigating Information Systems With Action Research,” in *Communications of the AIS*, vol. 2, no. 19. Association for information Systems, 1999.

[34] D. R. Hancock and B. Algozzine, *Doing case study research: A practical guide for beginning researchers*. Teachers College Press, 2006.

[35] R. K. Yin, *Applications of case study research*. Sage, 2011.

[36] M. M. Kennedy, “Generalizing from single case studies,” *Evaluation Review*, vol. 3, no. 4, pp. 661–678, 1979.

[37] A. M. Riege, “Validity and reliability tests in case study research: a literature review with “hands-on” applications for each research phase,” *Qualitative market research: An international journal*, vol. 6, no. 2, pp. 75–86, 2003.

[38] N. Golafshani, “Understanding reliability and validity in qualitative research,” *The qualitative report*, vol. 8, no. 4, pp. 597–606, 2003.

[39] B. Flyvbjerg, “Five misunderstandings about case-study research,” *Qualitative inquiry*, vol. 12, no. 2, pp. 219–245, 2006.

[40] B. Kitchenham and S. L. Pfleeger, "Principles of Survey Research Part 3: Constructing a Survey Instrument," *ACM SIGSOFT Software Engineering Notes*, vol. 27, no. 2, pp. 20–24, 2002.

[41] FireEye, "The Need for Speed: 2013 Incident Response Survey," Online at: www2.fireeye.com/ismg-incident-response-survey.html, 2013.

[42] Ponemon Institute, "Threat Intelligence & Incident Response: A Study of U.S. & EMEA Organizations," Online at: /www.ponemon.org/blog/threat-intelligence-incident-response-a-study-of-u-s-emea-organizations, 2014.

[43] ——, "Cyber Security Incident Response: Are We as Prepared as We Think?" Online at: www.ponemon.org/blog/cyber-security-incident-response-are-we-as-prepared-as-we-think, 2014.

[44] P. Shedden, A. Ahmad, and A. Ruighaver, "Organisational Learning and Incident Response: Promoting Effective Learning Through The Incident Response Process," in *8th Australian Information Security Management Conference*, 2010, p. 131.

[45] R. Davison, M. G. Martinsons, and N. Kock, "Principles of Canonical Action Research," *Information Systems Journal*, vol. 14, no. 1, pp. 65–86, 2004.

[46] H. F. Tipton, *Official (ISC) 2 Guide to the ISSAP CBK*. Auerbach Publications, 2010.

[47] British Standards Institution, "BS-7799-1:1999 Code of Practice for Information Security Management," 1999.

[48] The International Organization for Standardization and The International Electrotechnical Commission, "ISO/IEC 27000 – Information Security Management Systems – Overview and Vocabulary," 2014.

[49] J. M. Stewart, M. Chapple, and D. Gibson, *Certified Information Systems Security Professional Study Guide*. John Wiley & Sons, 2012.

[50] B. Von Solms and R. Von Solms, "The 10 Deadly Sins of Information Security Management," *Computers & Security*, vol. 23, no. 5, pp. 371–376, 2004.

[51] M. Workman, W. H. Bommer, and D. Straub, "Security Lapses and The Omission of Information Security Measures: A Threat Control Model and Empirical Test," *Computers in Human Behavior*, vol. 24, no. 6, pp. 2799–2816, 2008.

[52] C. P. Pfleeger and S. L. Pfleeger, *Security in Computing*, 4th ed. Prentice Hall Professional Technical Reference, 2007.

[53] MITRE Corporation, "CVE Initiative – Terminology," Online at: www.cve.mitre.org/about/terminology.html.

[54] S. Harris, *CISSP All-in-One Exam Guide*. McGraw Hill Professional, 2013.

[55] Cisco, "Cisco Security Advisories, Responses, and Notices," Online at: tools.cisco.com/security/center/publicationListing.x.

[56] Microsoft, "Microsoft Security Bulletin Advance Notification," Online at: www.technet.microsoft.com/en-us/security/gg309152.aspx.

[57] Symantec, "Security Response – The Internet Security Threat Report," Online at: www.symantec.com/security_response/.

[58] M. Talabis and J. Martin, *Information Security Risk Assessment Toolkit: Practical Assessments Through Data Collection and Data Analysis*, 1st ed. Syngress Publishing, 2013.

[59] The International Organization for Standardization and The International Electrotechnical Commission, "ISO/IEC 27005:2011 – Information Technology – Security Techniques – Information Security Risk Management," 2011.

[60] National Institute of Standards and Technology, "NIST Special Publication 800–30 Revision 1: Guide for Conducting Risk Assessments," 2012.

[61] T. R. Peltier, "Facilitated Risk Analysis Process (FRAP)," in *Data Secuity Management*. Auerbach Publications, 2000.

[62] Carnegie Mellon University, "OCTAVE (Operationally Critical Threat, Asset, and Vulnerability Evaluation)," Online at: www.cert.org/resilience/products-services/octave.

[63] B. Blakley, E. McDermott, and D. Geer, "Information Security is Information Risk Management," in *Proceedings of the 2001 Workshop on New Security Paradigms*. ACM, 2001, pp. 97–104.

[64] A. Dutta and K. McCrohan, "Management's Role in Information Security in a Cyber Economy," *California Management Review*, vol. 45, no. 1, pp. 67–87, 2002.

[65] A. D. Veiga and J. H. Eloff, "An Information Security Governance Framework," *Information Systems Management*, vol. 24, no. 4, pp. 361–372, 2007.

[66] B. Von Solms, "Information Security: The Third Wave?" *Computers & Security*, vol. 19, no. 7, pp. 615–620, 2000.

[67] B. von Solms, "Information Security: The Fourth Wave," *Computers & Security*, vol. 25, no. 3, pp. 165–168, 2006.

[68] National Cyber Security Summit Task Force, "Information Security Governance : A Call to Action," Online at: www.dhs.gov/sites/default/files/publications/csd-informationsecuritygovernance-acalltoaction-2004.pdf, 20004.

[69] T. Humphreys and A. Plate, *Measuring the Effectiveness of Your ISMS Implementations Based on ISO/IEC 27001*. BSI, 2006.

[70] E. Humphreys, "Information Security Management Standards: Compliance, Governance and Risk Management," *Information Security Technical Report*, vol. 13, no. 4, pp. 247–255, 2008.

[71] The European Union Agency for Network and Information Security. The Information Security Management System Framework. Online at: www.enisa.europa.eu/activities/risk-management/current-risk/risk-management-inventory/rm-isms/framework.

[72] J. H. P. Eloff and M. Eloff, "Information Security Management: A New Paradigm," in *Proceedings of the 2003 Annual Research Conference of the South African Institute of Computer Scientists and Information Technologists on Enablement Through Technology*, 2003, pp. 130–136.

[73] The European Union Agency for Network and Information Security, "Risk Management: Implementation Principles and Inventories for Risk Management/Risk Assessment Methods and Tools," ENISA, Tech. Rep., 2006.

[74] British Standards Institution, "BS-7799-2:2002 Information Security Management. Specification with Guidance for Use," 2002.

[75] B. Shojaie, H. Federrath, and I. Saberi, "Evaluating the Effectiveness of ISO 27001:2013," in *9th International Workshop on Frontiers in Availability, Reliability and Security (FARES 2014). University of Fribourg, Switzerland, 11 Sep 2014*, 2014.

[76] M. Rhodes-Ousley, *Information Security: The Complete Teference*. McGraw-Hill Osborne, 2013.

[77] United Kingdom Government Department for Business, Innovation & Skills and Cabinet Office, "UK cyber security standards research," Online at: www.gov.uk/government/publications/uk-cyber-security-standards-research, 2013.

[78] National Institute of Standards and Technology, "Federal Information Processing Standards Publications (FIPS PUBS)," Online at: www.csrc.nist.gov/publications/PubsFIPS.html.

[79] ——, “Special Publications (800 Series),” Online at: www.csrc.nist.gov/publications/PubsSPs.html.

[80] Payment Card Industry, “Data Security Standard – Requirements and Security Assessment Procedures Version 3.0,” Online at: www.pcisecuritystandards.org/documents/PCI_DSS_v3.pdf, 2013.

[81] M. Siponen and R. Willison, “Information Security Management Standards: Problems and Solutions,” *Information & Management*, vol. 46, no. 5, pp. 267–270, 2009.

[82] United States Government, “Federal Information Security Management Act of 2002,” Online at: www.csrc.nist.gov/drivers/documents/FISMA-final.pdf, 2002.

[83] ——, “E-Government Act of 2002,” Online at: www.gpo.gov/fdsys/pkg/STATUTE-116/pdf/STATUTE-116-Pg2899.pdf.

[84] J. Moteff, “Computer Security: A Summary of Selected Federal Laws, Executive Orders, and Presidential Directives,” Online at: www.fas.org/irp/crs/RL32357.pdf, Congressional Research Service, 2004.

[85] National Institute of Standards and Technology, “Security and Privacy Controls for Federal Information Systems and Organizations,” Online at www.nvlpubs.nist.gov/nistpubs/SpecialPublications/NIST.SP.800-53r4.pdf.

[86] ——, “Minimum Security Requirements for Federal Information and Information Systems –FIPS PUB 200,” Online at: www.csrc.nist.gov/publications/fips/fips200/FIPS-200-final-march.pdf.

[87] ——, “Standards for Security Categorization of Federal Information and Information Systems – FIPS PUB 199,” Online at: www.csrc.nist.gov/publications/fips/fips199/FIPS-PUB-199-final.pdf.

[88] R. Ross, “Managing Enterprise Security Risk With NIST Standards,” *Computer*, vol. 40, no. 8, pp. 88–91, 2007.

[89] PCI Security Standards Council, “PCI Council Publishes Revision To PCI Data Security,” Online at: www.pcisecuritystandards.org/pdfs/15_04_15PCIDSS31PressRelease.pdf, April 2015.

[90] ——, “Payment Card Industry (PCI) Data Security Standard - Requirements and Security Assessment Procedures Version 3.1,” Online at: www.pcisecuritystandards.org/documents/PCI_DSS_v3-1.pdf, 2015.

[91] The House of Representatives of the State of Washington, “Financial Information Security Breaches - Credit and Debt Cards,” Chapter 151, Laws of 2010 - House Bill 1149.

[92] R. Von Solms, “Information Security Management: Why Standards Are Important,” *Information Management & Computer Security*, vol. 7, no. 1, pp. 50–58, 1999.

[93] T. Wiander, “Implementing the ISO/IEC 17799 Standard in Practice – Findings From Small and Medium-Sized Software Organisations,” in *Standardization and Innovation in Information Technology, 2007. SIIT 2007. 5th International Conference on*. IEEE, 2007, pp. 91–104.

[94] M. Siponen, “Information Security Standards Focus on The Existence of Process, Not Its Content,” *Communications of the ACM*, vol. 49, no. 8, pp. 97–100, 2006.

[95] J. D. Howard and T. A. Longstaff, “A Common Language for Computer Security Incidents,” Sandia National Laboratories, Tech. Rep., 1998.

[96] United States Government, “Health Insurance Portability and Accountability Act of 1996,” Online at: www.aspe.hhs.gov/admnsimp/pl104191.htm, 1996.

[97] ——, “Code of Federal Regulations Title 45: Public Welfare – Part 164 – Security and Privacy,” Online at: www.ecfr.gov/cgi-bin/text-idx?SID=1129cba6073938f1925220e48efe92c5&node=sp45.1.164.c&rgn=div6.

[98] The European Union Agency for Network and Information Security. (2014) Technical guidance on the incident reporting in article 13a - version 2.1. Online at: www.enisa.europa.eu/activities/Resilience-and-CIIP/Incidents-reporting/Technical%20Guidelines%20on%20Incident%20Reporting/technical-guideline-on-incident-reporting/at_download/fullReport.

[99] Information Commissioner’s Office, “Notification of PECR security breaches - Privacy and Electronic Communications Regulations,” Online at: www.ico.org.uk/media/for-organisations/documents/1583/notification-of-pecr-security-breaches.pdf, 2015.

[100] United States Department of of Health and Human Services, “HIPAA Administrative Simplification Regulation Text 45 CFR Parts 160, 162, and 164 - Security Standards: Administrative Safeguards,” 2007.

[101] United States Government, “Gramm-Leach-Bliley Act of 1999,” 1999.

[102] ——, “The Sarbanes–Oxley Act of 2002,” 2002.

[103] B. Carrier and E. H. Spafford, “An Event-based Digital Forensic Investigation Framework,” in *2004 Digital Forensic Research Workshop*, 2004, pp. 11–13.

[104] The European Parliament, “Directive 2009/140/EC of the European Parliament – Regulatory Framework for Electronic Communications,” The European Union, Tech. Rep., 2009.

[105] The European Commission, “Proposal for a Directive of the European Parliament and the Council: Concerning Measures to Ensure a High Common Level of Network and Information Security Across The Union - COM(2013) 48 Final,” The European Union, Tech. Rep., 2013.

[106] C. Johnson, “Contrasting Approaches to Incident Reporting in the Development of Security and Safety Critical Software,” in *2015 International Conference on Computer Safety, Reliability and Security (SAFECOMP 2015)*, 2015, accepted for SAFECOMP.

[107] United Kingdom Government, “Cyber-Security Information Sharing Partnership (CiSP),” Online at: https://www.cert.gov.uk/cisp/, 2013.

[108] M. J. West-Brown, D. Stikvoort, K.-P. Kossakowski, G. Killcrece, and R. Ruefle, “Handbook of Computer Security Incident Response Teams, Second Edition,” Software Engineering Institute, Carnegie Mellon University, Tech. Rep. CMU/SEI-2003-HB-002, 2003.

[109] M. Ryba, A. Poniewierski, J. Sulwiński, and M. Górnisiewicz, “The methodology for managing the abuse of it systems 1,” *EDPACS The EDP Audit, Control, and Security Newsletter*, vol. 40, no. 5, pp. 1–13, 2009.

[110] R. Werlinger and D. Botta, “Detecting, Analyzing and Responding to Security Incidents: A Qualitative Analysis,” in *Proceedings of the 3rd Symposium on Usable Privacy and Security (SOUPS 2007)*, 2007, pp. 149–150.

[111] J. Wack, “Establishing a Computer Security Incident Response Capability,” National Institute of Standards and Technology, Tech. Rep. 800-3, 1991.

[112] P. Werner and R. Perry, “The Role of Lessons Learned in the Investigate, Communicate, Educate-Cycle for Commercial Aviation,” in *Proceedings of the 35th Annual International Gold Coast, Queensland, Australia*, 2004, pp. 51–56.

[113] I. Sommerville, *Software Engineering, 9th Edition.* Addison Wesley, 2010.

[114] C. Hove, M. Tarnes, M. B. Line, and K. Bernsmed, “Information Security Incident Management: Identified Practice in Large Organizations,” in *2014 International Conference on IT Security Incident Management & IT Forensics (IMF)*. IEEE, 2014, pp. 27–46.

[115] M. B. Line, I. A. Tondel, and M. G. Jaatun, “Information Security Incident Management: Planning for Failure,” in *2014 International Conference on IT Security Incident Management & IT Forensics (IMF)*. IEEE, 2014, pp. 47–61.

[116] M. B. Line, “A Case Study: Preparing for the Smart Grids-Identifying Current Practice for Information Security Incident Management in the Power Industry,” in *2013 International Conference on IT Security Incident Management & IT Forensics (IMF)*. IEEE, 2013, pp. 26–32.

[117] S. Metzger, W. Hommel, and H. Reiser, “Integrated Security Incident Management – Concepts and Real-World Experiences,” in *IT Security Incident Management and IT Forensics (IMF), 2011 Sixth International Conference on*. IEEE, 2011, pp. 107–121.

[118] T. L. Orderløkken, “Security incident handling and reporting – a study of the difference between theory and practice,” Master’s thesis, Department of Computer Science and Media Technology, Gjøvik University College, 2005.

[119] K. Moller, “Setting up a Grid? CERT: experiences of an academic CSIRT,” *Campus-Wide Information Systems*, vol. 24, no. 4, pp. 260–270, 2007.

[120] D. L. Cooke, “Learning from Incidents,” in *21st System Dynamics Conference, NYC, New York*, 2003.

[121] M. B. Line, E. Albrechtsen, M. G. Jaatun, I. A. Tøndel, S. O. Johnsen, O. H. Longva, and I. Wærø, “A Structured Approach to Incident Response Management in the Oil and Gas Industry,” in *Critical Information Infrastructure Security*. Springer, 2009, pp. 235–246.

[122] M. G. Jaatun, E. Albrechtsen, M. B. Line, I. A. Tøndel, and O. H. Longva, “A Framework for Incident Response Management in the Petroleum Industry,” *International Journal of Critical Infrastructure Protection*, vol. 2, no. 1, pp. 26–37, 2009.

[123] R. L. Rollason-Reese, “Incident Handling: An Orderly Response to Unexpected Events,” in *Proceedings of the 31st Annual ACM SIGUCCS Fall Conference*. ACM, 2003, pp. 97–102.

[124] G. Grispos, W. B. Glisson, and T. Storer, “Rethinking Security Incident Response: The Integration of Agile Principles,” in *20th Americas Conference on Information Systems (AMCIS 2014), Savannah, USA*, 2014.

[125] M. Paradies and D. Busch, “Root Cause Analysis at Savannah River Plant,” in *Human Factors and Power Plants, 1988., Conference Record for 1988 IEEE Fourth Conference on*, June 1988, pp. 479–483.

[126] C. Johnson, “Using Violation and Vulnerability Analysis to Understand the Root Causes of Complex Frauds,” Department of Computing Science, University of Glasgow, Glasgow, Scotland, Tech. Rep. GIST Technical Report 2005-1, 2005.

[127] P. Stephenson, “The Application of Formal Methods to Root Cause Analysis 0f Digital Incidents,” *International Journal of Digital Evidence*, vol. 3, no. 1, pp. 1–5, 2004.

[128] United States Veterans Affairs Administration, “Administrative Investigation Loss of VA Information - Report No. 07-01083-157,” United States Veterans Affairs Administration, Tech. Rep., 2007.

[129] Y. He, C. Johnson, K. Renaud, Y. Lu, and S. Jebriel, “An Empirical Study on the Use of the Generic Security Template for Structuring the Lessons from Information Security Incidents,” in *Computer Science and Information Technology (CSIT), 2014 6th International Conference on*. IEEE, 2014, pp. 178–188.

[130] J. E. Rowley, “The wisdom hierarchy: representations of the dikw hierarchy,” *Journal of information science*, 2007.

[131] R. L. Ackoff, “From data to wisdom,” *Journal of applied systems analysis*, vol. 16, no. 1, pp. 3–9, 1989.

[132] D. Knapp, *A guide to service desk concepts*. Nelson Education, 2013.

[133] S. Baskarada and A. Koronios, “Data, information, knowledge, wisdom (dikw): a semiotic theoretical and empirical exploration of the hierarchy and its quality dimension,” *Australasian Journal of Information Systems*, vol. 18, no. 1, 2013.

[134] B. Klein and D. F. Rossin, “Data Errors in Neural Network and Linear Regression Models: An Experimental Comparison,” *Data Quality*, vol. 5, no. 1, p. 25, 1999.

[135] L. Sebastian-Coleman, *Measuring Data Quality For On-going Improvement: A Data Quality Assessment Framework*. A volume in MK Series on Business Intelligence, 2012.

[136] K. Orr, “Data Quality and Systems Theory,” *Communications of the ACM*, vol. 41, no. 2, pp. 66–71, 1998.

[137] R. Y. Wang and D. M. Strong, “Beyond Accuracy: What Data Quality Means to Data Consumers,” *Journal of management information systems*, vol. 12, no. 4, pp. 5–33, 1996.

[138] A. Maydanchik, *Data Quality Assessment*. Technics publications, 2007.

[139] G. K. Tayi and D. P. Ballou, "Examining Data Quality," *Communications of the ACM*, vol. 41, no. 2, pp. 54–57, 1998.

[140] K. Kerr, "The Institutionalisation of Data Quality in the New Zealand Health Sector," Ph.D. dissertation, University of Auckland, 2006.

[141] R. Y. Wang, M. Ziad, and Y. W. Lee, *Data Quality*, 1st ed. Springer, 2001.

[142] Y. Wand and R. Y. Wang, "Anchoring Data Quality Dimensions in Ontological Foundations," *Communications of the ACM*, vol. 39, no. 11, pp. 86–95, 1996.

[143] A. Levitin and T. Redman, "Quality Dimensions of a Conceptual View," *Information Processing & Management*, vol. 31, no. 1, pp. 81–88, 1995.

[144] D. P. Ballou and H. L. Pazer, "Modeling Data and Process Quality in Multi-input, Multi-output Information Systems," *Management science*, vol. 31, no. 2, pp. 150–162, 1985.

[145] D. Rollinson, *Organisational behaviour and analysis: an integrated approach*. Pearson Education, 2008.

[146] T. J. Watson, *Organising and Managing Work: Organisational, managerial and strategic behaviour in theory and practice*. Pearson Education, 2006.

[147] O. O'Donnell and R. Boyle, "Understanding and managing organisational culture," 2003.

[148] E. Martins and F. Terblanche, "Building organisational culture that stimulates creativity and innovation," *European journal of innovation management*, vol. 6, no. 1, pp. 64–74, 2003.

[149] E. H. Schein, "Organisational culture and leadership: A dynamic view," *San Francisco*, 1985.

[150] T. Deal and A. Kennedy, *Corporate cultures: the rites and rituals of corporate life*. Penguin business, 1988.

[151] E. H. Schein, *The corporate culture survival guide*. Jossey-bass, 1999.

[152] E. Schein, "Innovative cultures and adaptive organisations," *Sri Lanka Journal of Development Administration*, vol. 7, no. 2, pp. 9–39, 1990.

[153] G. Hofstede, "Culture's consequences: National differences in thinking and organizing," *Beverly Hills, Calif.: Sage*, 1980.

[154] G. Hofstede, B. Neuijen, D. D. Ohayv, and G. Sanders, "Measuring organizational cultures: A qualitative and quantitative study across twenty cases," *Administrative science quarterly*, pp. 286–316, 1990.

[155] G. Hofstede and B. Waisfisz. Organisational culture. [Online]. Available: http://geert-hofstede.com/organisational-culture.html

[156] K. S. Cameron and R. E. Quinn, *Diagnosing and changing organizational culture: Based on the competing values framework*. John Wiley & Sons, 2005.

[157] K. S. Cameron and D. R. Ettington, "The conceptual framework of organizational culture," *Higher education: Handbook of theory and research*, vol. 6, pp. 356–396, 1988.

[158] D. R. Denison, *Corporate culture and organizational effectiveness.* John Wiley & Sons, 1990.

[159] H. M. Trice and J. M. Beyer, *The cultures of work organizations*. JSTOR, 1993.

[160] T. M. Egan, B. Yang, and K. R. Bartlett, "The effects of organizational learning culture and job satisfaction on motivation to transfer learning and turnover intention," *Human resource development quarterly*, vol. 15, no. 3, pp. 279–301, 2004.

[161] C. Handy, *Understanding organizations*. Penguin UK, 1993.

[162] S. J. Gill, *Developing a learning culture in nonprofit organizations*. Sage, 2009.

[163] G. P. Huber, "Organizational learning: The contributing processes and the literatures," *Organization science*, vol. 2, no. 1, pp. 88–115, 1991.

[164] B. Levitt and J. G. March, "Organizational learning," *Annual review of sociology*, pp. 319–340, 1988.

[165] C. Argyris and D. A. Schön, *Organizational learning: A theory of action perspective*. Addison-Wesley Reading, MA, 1978.

[166] J. M. Sinkula, W. E. Baker, and T. Noordewier, "A framework for market-based organizational learning: Linking values, knowledge, and behavior," *Journal of the academy of Marketing Science*, vol. 25, no. 4, pp. 305–318, 1997.

[167] J. Van Niekerk and R. von Solms, "Organisational learning models for information security," in *The ISSA 2004 Enabling Tomorrow Conference*, vol. 30, 2004.

[168] L. Argote, *Organizational learning: Creating, retaining and transferring knowledge*. Springer Science & Business Media, 2012.

[169] D. F. Russ-Eft, H. S. Preskill, and C. Sleezer, *Human resource development review: Research and implications*. Sage Pubns, 1997.

[170] R. Baskerville, P. Spagnoletti, and J. Kim, "Incident-centered information security: Managing a strategic balance between prevention and response," *Information & management*, vol. 51, no. 1, pp. 138–151, 2014.

[171] C. Melara, J. M. Sarriegui, J. J. Gonzalez, A. Sawicka, and D. L. Cooke, "A system dynamics model of an insider attack on an information system," in *Proceedings of the 21st International Conference of the System Dynamics Society July 20-24*, 2003.

[172] M. Siponen and A. Vance, "Neutralization: New Insights Into the Problem of Employee Information Systems Security Policy Violations," *MIS quarterly*, vol. 34, no. 3, p. 487, 2010.

[173] S. Pahnila, M. Siponen, and A. Mahmood, "Employees' Behavior Towards IS Security Policy Compliance," in *System Sciences, 2007. HICSS 2007. 40th Annual Hawaii International Conference on*. IEEE, 2007, pp. 156b–156b.

[174] M. Siponen, S. Pahnila, and M. A. Mahmood, "Compliance with Information Security Policies: An Empirical Investigation," *Computer*, vol. 43, no. 2, pp. 64–71, 2010.

[175] U. Lindqvist and E. Jonsson, "How to Systematically Classify Computer Security Intrusions," in *Security and Privacy, 1997. Proceedings., 1997 IEEE Symposium on*. IEEE, 1997, pp. 154–163.

[176] D. Schieber and G. Reid. (2004) CSIRT Case Classification (Example for Enterprise CSIRT). Online at: www.first.org/_assets/resources/guides/csirt_case_classification.html.

[177] The European Union Agency for Network and Information Security, "Good practices on reporting security incidents," The European Union Agency for Network and Information Security, Tech. Rep., 2009.

[178] United Kingdom Government, "10 steps: Incident management – https://www.gov.uk/government/publications/10-steps-to-cyber-security-advice-sheets/10-steps-incident-management--11."

[179] A. Dorofee, G. Killcrece, R. Ruefle, and M. Zajicek, "Incident Management Capability Metrics Version 0.1," Software Engineering Institute, Carnegie Mellon University, Pittsburgh, PA, Tech. Rep. CMU/SEI-2007-TR-008, 2007.

[180] J. Wiik, J. J. Gonzalez, and K.-P. Kossakowski, “Limits to Effectiveness in Computer Security Incident Response Teams,” in *Twenty Third International Conference of the System Dynamics Society, Boston, Massachusetts:*, 2005.

[181] V. Jaikumar, “Build a Response Team,” Online at: www.computerworld.com/article/2577130/security0/build-a--response-team.html, 2002.

[182] E. Casey, *Digital Evidence and Computer Crime: Forensic Science, Computers and the Internet*. Academic press, 2011.

[183] The United States Securities and Exchange Commission, “CF Disclosure Guidance: Topic No. 2,” Available at: www.sec.gov/divisions/corpfin/guidance/cfguidance-topic2.htm, Tech. Rep., 2011.

[184] C. Alberts, A. Dorofee, G. Killcrece, R. Ruefle, and M. Zajicek, “Defining Incident Management Processes for CSIRTS: A Work in Progress,” DTIC Document, Tech. Rep., 2004.

[185] The International Organization for Standardization and The International Electrotechnical Commission, “ISO/IEC 27035, Information technology – Security techniques – Information security incident management,” 2001.

[186] H. H. Abujudeh and M. A. Bruno, *Quality and safety in radiology*. Oxford University Press, 2012.

[187] D. Guideline, “Root cause analysis guidance document,” 1992.

[188] The European Union Agency for Network and Information Security, “Actionable Information for Security Incident Response,” ENISA, Tech. Rep., 2014.

[189] National Institute of Standards and Technology. (2010) Guide to Protecting the Confidentiality of Personally Identifiable Information (PII) - NIST Special Publication 800-122 . Online at: http://csrc.nist.gov/publications/nistpubs/800-122/sp800-122.pdf.

[190] H. A. Simon, “Making management decisions: The role of intuition and emotion,” *The Academy of Management Executive (1987-1989)*, pp. 57–64, 1987.

[191] N. Khatri, G. D. Brown, and L. L. Hicks, “From a blame culture to a just culture in health care,” *Health care management review*, vol. 34, no. 4, pp. 312–322, 2009.

[192] A. I. Shahin and P. L. Wright, “Leadership in the context of culture: An egyptian perspective,” *Leadership & Organization Development Journal*, vol. 25, no. 6, pp. 499–511, 2004.

[193] A. Cowling and K. Newman, “Banking on people: Tqm, service quality and human resources,” *Personnel review*, vol. 24, no. 7, pp. 25–40, 1995.

[194] SANS, “Incident Response: How to Fight Back,” Online at: www.sans.org/reading-room/whitepapers/incident/incident-response-fight-35342, 2014.

[195] J. J. Waring, “Beyond blame: cultural barriers to medical incident reporting,” *Social science & medicine*, vol. 60, no. 9, pp. 1927–1935, 2005.

[196] T. Schlienger and S. Teufel, “Information security culture,” in *Security in the Information Society*. Springer, 2002, pp. 191–201.

[197] W. Bridges, “Get Near Misses Reported – Process Industry Incidents: Investigation Protocols, Case Histories, Lessons Learned,” in *Proceedings of the Center for Chemical Process Safety International Conference and Workshop*. New York: American Institute of Chemical Engineers, 2000.

[198] K. Beck, M. Beedle, A. van Bennekum, A. Cockburn, W. Cunningham, M. Fowler, J. Grenning, J. Highsmith, A. Hunt, and R. Jeffries, “Manifesto for Agile Software Development,” Online at: www.agilemanifesto.org/, 2001.

[199] J. Shore and S. Warden, *The Art of Agile Development*. O’Reilly Media, Inc, 2007.

[200] M. Maham, “Planning and Facilitating Release Retrospectives,” in *Agile, 2008. AGILE ’08. Conference*, Aug 2008, pp. 176–180.

[201] J. Babb, R. Hoda, and J. Norbjerg, “Embedding Reflection and Learning into Agile Software Development,” *Software, IEEE*, vol. 31, no. 4, pp. 51–57, July 2014.

[202] R. Hoda, J. Babb, and J. Norbjerg, “Toward Learning Teams,” *Software, IEEE*, vol. 30, no. 4, pp. 95–98, 2013.

[203] O. Hazzan and J. Tomayko, “The Reflective Practitioner Perspective In eXtreme Programming,” in *Extreme Programming and Agile Methods-XP/Agile Universe 2003*. Springer, 2003, pp. 51–61.

[204] E. Derby and D. Larsen, *Agile Retrospectives: Making Good Teams Great*. Pragmatic Bookshelf Raleigh, NC, 2006.

[205] A. Pham and P.-V. Pham, *Scrum in Action:: Agile Software Project Management and Development*. Cengage Learning, 2011.

[206] Extreme Programming (XP), “Fix XP When It Breaks,” Online at: www.extremeprogramming.org/rules/fixit.html.

[207] A. Cockburn, *Crystal Clear: A Human-Powered Methodology for Small Teams*. Pearson Education, 2004.

[208] K. S. Rubin, *Essential Scrum: A Practical Guide to the Most Popular Agile Process*. Addison-Wesley, 2012.

[209] G. Tiwari and Z. Alikhan, "From 'Team' to 'Wow Team': An Agile Team's Journey," in *Agile, 2011. AGILE '11. Conference*, 2011, pp. 296–301.

[210] O. McHugh, K. Conboy, and M. Lang, "Agile Practices: The Impact on Trust in Software Project Teams," *IEEE Software*, vol. 29, no. 3, pp. 71–76, 2012.

[211] L. Gonçalves and B. Linders, *Getting Value out of Agile Retrospectives*. Lean Publishing, 2015.

[212] G. Grispos, W. B. Glisson, and T. Storer, "Security Incident Response Criteria: A Practitioners Perspective," in *21st Americas Conference on Information Systems (AMCIS 2015), Puerto Rico, USA*, 2015.

[213] "Corporate sustainability and organizational culture," *Journal of World Business*, vol. 45, no. 4, pp. 357 – 366, 2010.

[214] Y. He and C. Johnson, "Generic security cases for information system security in healthcare systems," in *System Safety, incorporating the Cyber Security Conference 2012, 7th IET International Conference on*. IET, 2012, pp. 1–6.

[215] S. Few, *Information Dashboard Design: The Effective Visual Communication of Data*. O'Reilly Media, 2006.

[216] R. Marty, *Applied Security Visualization*. Addison-Wesley, 2008, no. 1st Edition.

[217] Splunk. (2015) Splunk App for Enterprise Security. Online: http://www.splunk.com/en_us/solutions/solution-areas/security-and-fraud/splunk-app-for-enterprise-security.html.

[218] IBM. (2007) IBM Tivoli Compliance Insight Manager. Online at: www-304.ibm.com/industries/publicsector/fileserve?contentid=182101.

[219] Computer Network Defence. (2015) Cyber Threat Intelligence Radar. Online at: http://www.securitywizardry.com/radar.htm.

[220] K. Sun, S. Jajodia, J. Li, Y. Cheng, W. Tang, and A. Singhal, "Automatic Security Analysis using Security Metrics," in *The 2011 Military Communications Conference*, Nov 2011, pp. 1207–1212.

[221] I. Kotenko and E. Novikova, "VisSecAnalyzer: A Visual Analytics Tool for Network Security Assessment," in *Security Engineering and Intelligence Informatics*. Springer, 2013, pp. 345–360.

[222] B. Sobesto, M. Cukier, M. A. Hiltunen, D. Kormann, G. Vesonder, and R. Berthier, "DarkNOC: Dashboard for Honeypot Management," in *USENIX: Large Installation System Administration Conference*, 2011.

[223] E. Novikova and I. Kotenko, "Analytical Visualization Techniques for Security Information and Event Management," in *Parallel, Distributed and Network-Based Processing (PDP), 2013 21st Euromicro International Conference on*. IEEE, 2013, pp. 519–525.

[224] J. Jacobs and B. Rudis, *Data-Driven Security: Analysis, Visualization and Dashboards*. Wiley, 2014.

[225] S. Madnick, X. Li, and N. Choucri, "Experiences and Challenges With Using CERT Data to Analyze International Cyber Security," Composite Information Systems Laboratory, Sloan School of Management, Massachusetts Institute of Technology, MIT Sloan School Working Paper 4759-09, 2009.

[226] EMC Corporation. (2014) The Critical Incident Response Maturity Journey. Online: http://www.emc.com/collateral/white-papers/h12651-wp-critical-incident-response-maturity-journey.pdf.

[227] Tenable Network Security. (2015) Speed Up Incident Response with Actionable Forensic Analytics. Online: http://www.infosecurityeurope.com/__novadocuments/76812?v=635611693680730000.

[228] McAfee and Intel Security. (2013) Creating and Maintaining a SOC. Online: http://www.mcafee.com/us/resources/white-papers/foundstone/wp-creating-maintaining-soc.pdf.

[229] The Center for Internet Security, "The CIS Security Metrics," CIS, Tech. Rep., 2010.

[230] IBM Notes. (2006) IBM Lotus Domino Designer 6 Templates Guide – Introducing the Discussion, Document Library, and TeamRoom Templates. Online: http://www-12.lotus.com/ldd/doc/domino_notes/Rnext/help6_templates_guide.nsf/b3266a3c17f9bb7085256b870069c0a9/4a277ff50de24b2f85256c77005b1b60.

[231] J. T. Biehl, M. Czerwinski, G. Smith, and G. G. Robertson, "FASTDash: A Visual Dashboard for Fostering Awareness in Software Teams," in *Proceedings of the SIGCHI conference on Human factors in computing systems*. ACM, 2007, pp. 1313–1322.

[232] B. List, R. Bruckner, and J. Kapaun, “Holistic Software Process Performance Measurement From the Stakeholders’ Perspective,” in *Database and Expert Systems Applications, 2005. Proceedings. Sixteenth International Workshop on*, Aug 2005, pp. 941–947.

[233] S. Larndorfer, R. Ramler, and C. Buchwiser, “Experiences and Results from Establishing a Software Cockpit at BMD Systemhaus,” in *Software Engineering and Advanced Applications, 2009. SEAA’09. 35th Euromicro Conference on.* IEEE, 2009, pp. 188–194.

[234] P. Dourish and V. Bellotti, “Awareness and coordination in shared workspaces,” in *Proceedings of the 1992 ACM conference on Computer-supported cooperative work.* ACM, 1992, pp. 107–114.

[235] U. Kjellén, *Prevention of accidents through experience feedback.* CRC Press, 2000.

[236] United States Department of Energy, “Root cause analysis guidance document,” Technical Report DOE-NE-STD-1004-92, US Department of Energy, Office of Nuclear Energy, Washington, DC, Tech. Rep., 1992.

[237] O. Svenson, “Accident Analysis and Barrier Function (AEB) Method,” Swedish Nuclear Power Inspectorate (SKI) and Netherlands Institute for Advanced Study in the Humanities and Social Sciences, Tech. Rep., 2000.

[238] United States Department of Energy, “Accident and Operational Safety Analysis, Volume 1: Accident Analysis Techniques,” Technical Report DOE-HDBK-1208-2012, US Department of Energy, Washington, DC, Tech. Rep., 2012.

[239] A. Livingston, G. Jackson, and K. Priestley, “Root Causes Analysis: Literature Review,” Health and Safety Executive, Tech. Rep., 2001.

[240] S. Sklet, “Comparison of Some Selected Methods for Accident Investigation,” *Journal of Hazardous Materials*, vol. 111, no. 1, pp. 29–37, 2004.

[241] R. Choularton, “Complex Learning: Organizational Learning from Disasters,” *Safety Science*, vol. 39, no. 1, pp. 61–70, 2001.

[242] T. Ōno, *Toyota Production System: Beyond Large-Scale Production.* Productivity Press, 1988.

[243] O. Serrat, “The Five–Whys Technique,” Online at: http://www.adb.org/sites/default/files/pub/2009/the-five-whys-technique.pdf, 2009.

[244] U. Murugaiah, S. Jebaraj Benjamin, M. Srikamaladevi Marathamuthu, and S. Muthaiyah, “Scrap Loss Reduction Using The 5-Whys Analysis,” *International Journal of Quality & Reliability Management*, vol. 27, no. 5, pp. 527–540, 2010.

[245] I. Reid and J. Smyth-Renshaw, “Exploring the Fundamentals of Root Cause Analysis: Are We Asking the Right Questions in Defining the Problem?” *Quality and Reliability Engineering International*, vol. 28, no. 5, pp. 535–545, 2012.

[246] R. J. Latino and A. Flood, “Optimizing FMEA and RCA Efforts in Healthcare,” *Journal of Healthcare Risk Management*, vol. 24, no. 3, pp. 21–28, 2004.

[247] M. Sondalini, “Understanding How to Use The 5-Whys for Root Cause Analysis,” Lifetime Reliability Solutions, Tech. Rep., 2012.

[248] Ernst & Young, “Outsourcing in Europe - An In-depth Review of Drivers, Risks and Trends in the European Outsourcing Market,” Online at: http://www.ey.com/Publication/vwLUAssets/Outsourcing_in_Europe_2013/$FILE/EY-outsourcing-survey.pdf, 2013.

[249] Deloitte, “2014 Global Outsourcing and Insourcing Survey,” Online at: https://www2.deloitte.com/content/dam/Deloitte/us/Documents/strategy/us-2014-global-outsourcing-insourcing-survey-report-123114.pdf, 2014.

[250] K. G. Bulsuk, “5-Why Excel Template,” Online at: http://sites.google.com/site/karnblog/5why/5-whyTemplate.xlsx?attredirects=0.

[251] G. Grispos, W. B. Glisson, and T. Storer, “Cloud Security Challenges: Investigating Policies, Standards, and Guidelines in a Fortune 500 Organization,” in *21st European Conference on Information Systems (ECIS 2013)*, 2013.

[252] Microsoft, “Use Rules in Outlook Web App to Automatically Forward Messages to Another Account,” Online at: https://support.office.com/article/Use-rules-in-Outlook-Web-App-to-automatically-forward-messages-to-another-account-1433e3a0-7fb0-4999-b536-50e05cb67fed.

[253] Montgomery County, District Attorney’s Office and University of Pennsylvania Law School, “Using root cause analysis to instill a culture of self-improvement,” University of Pennsylvania Law School, Tech. Rep., 2015.

[254] J. Carroll, J. Rudolph, and S. Hatakenaka, “Lessons learned from non-medical industries: root cause analysis as culture change at a chemical plant,” *Quality and Safety in Health Care*, vol. 11, no. 3, pp. 266–269, 2002.

[255] I. S. Sutton, "Use root cause analysis to understand and improve process safety culture," *Process Safety Progress*, vol. 27, no. 4, pp. 274–279, 2008.

[256] Eurocontrol, "Just culture policy," The European Union, Tech. Rep., 2012.

[257] P. F. Steinberg, "Can we generalize from case studies?" *Global Environmental Politics*, 2015.

[258] G. Walsham, "Interpretive case studies in is research: nature and method," *European Journal of information systems*, vol. 4, no. 2, pp. 74–81, 1995.