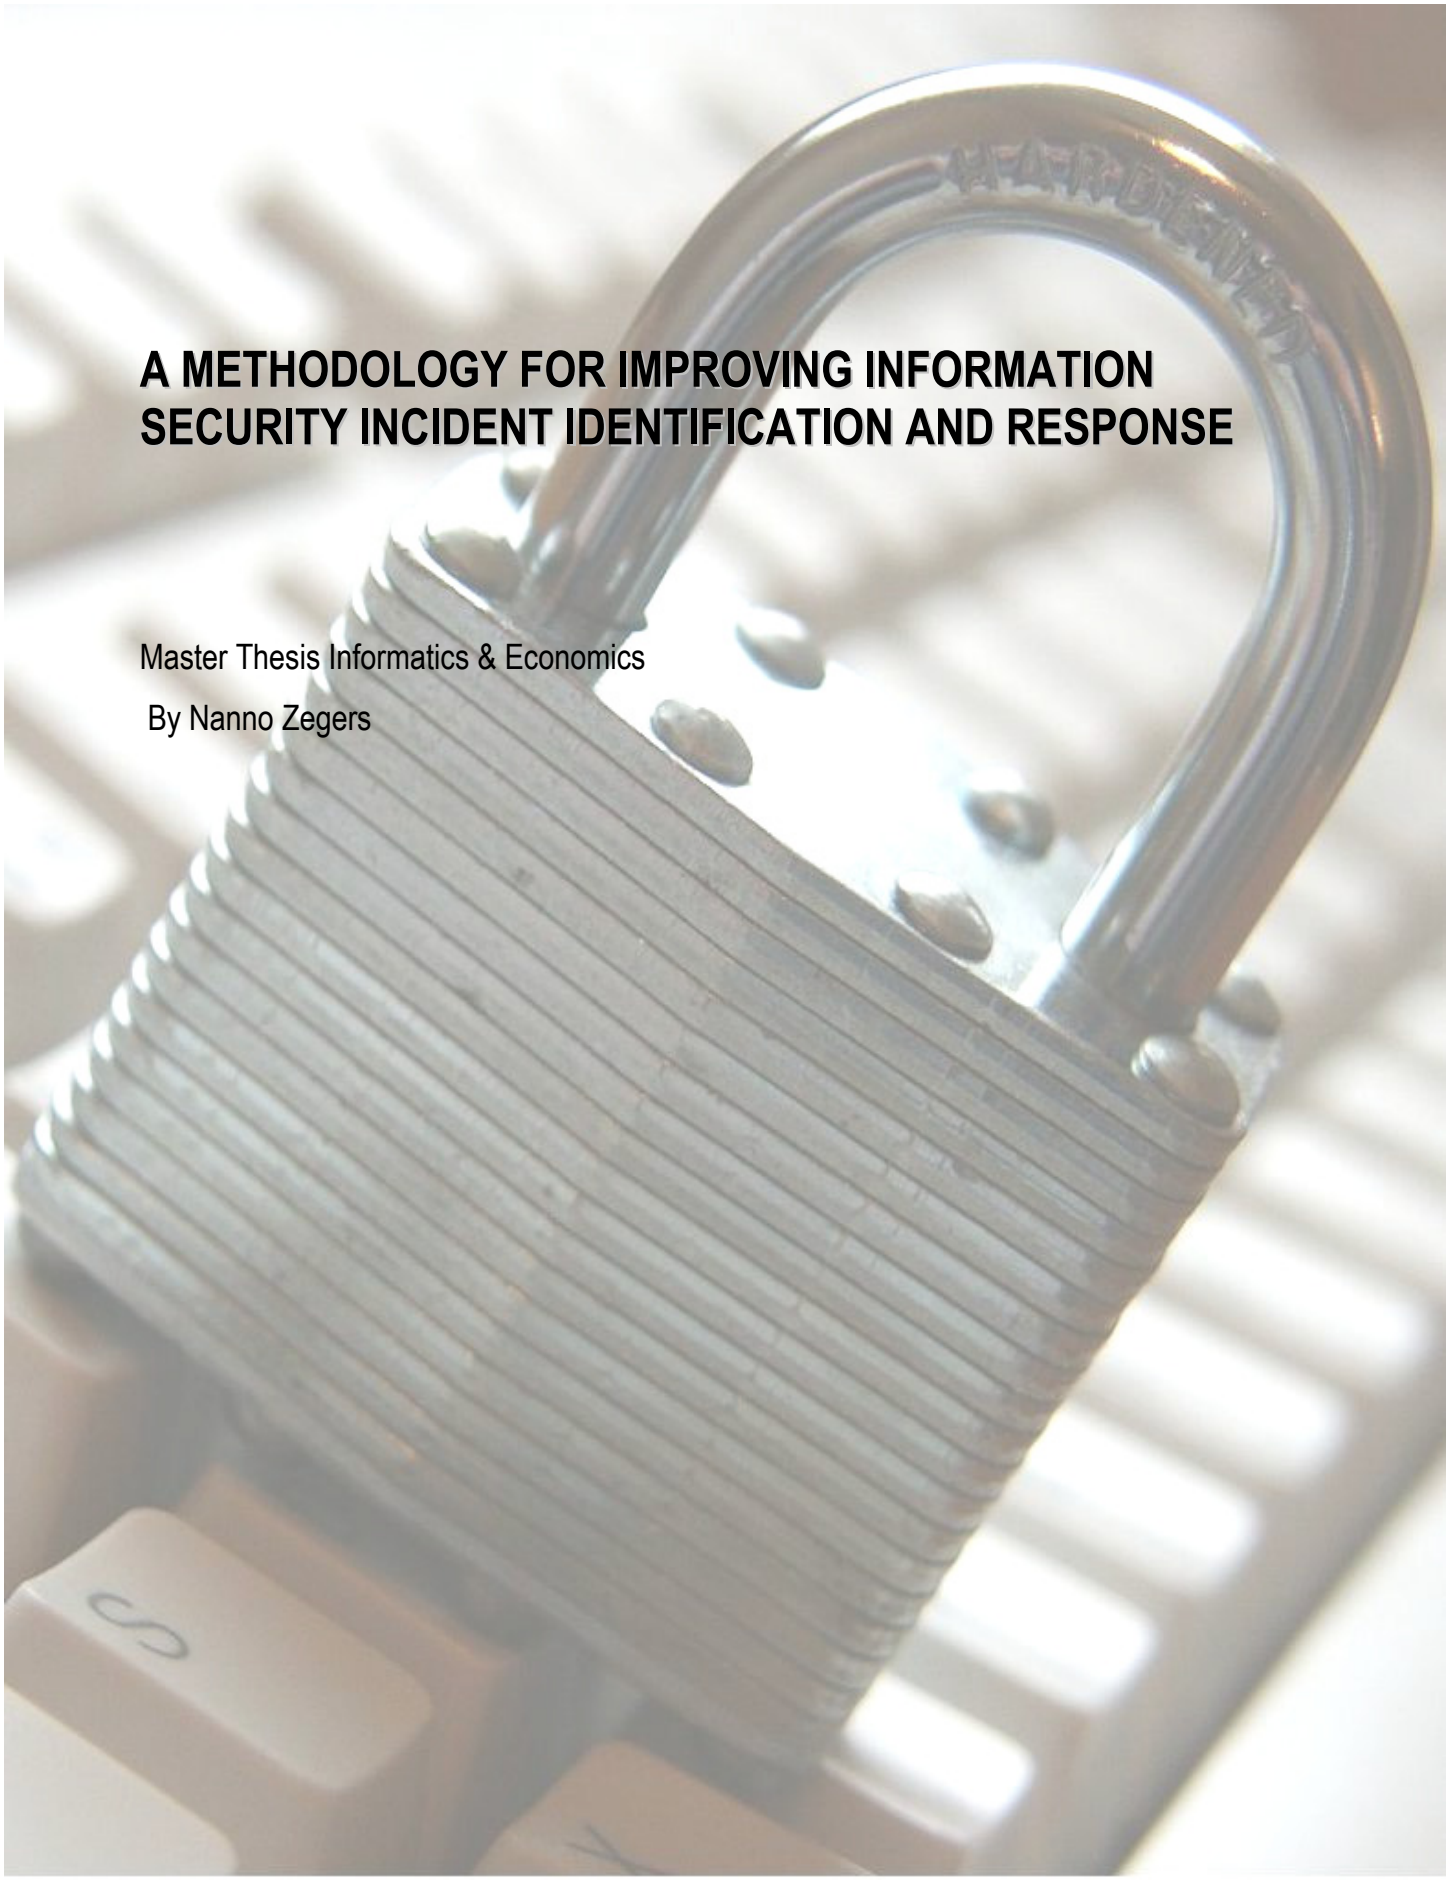

A METHODOLOGY FOR IMPROVING INFORMATION SECURITY INCIDENT IDENTIFICATION AND RESPONSE

Master Thesis Informatics & Economics

By Nanno Zegers

A METHODOLOGY FOR IMPROVING INFORMATION SECURITY INCIDENT IDENTIFICATION AND RESPONSE

Master Thesis Informatics & Economics

Master programme Economics and ICT

Faculty of Economics, Erasmus University Rotterdam

Nanno Zegers (260234)

Summer 2006

Co-supervised by Dr. ir. Jan van den Berg, Erasmus University Rotterdam

Co-supervised by Drs. Eelco Stofbergen, EY EDP Audit Rotterdam

Erasmus Universiteit Rotterdam

ERNST & YOUNG
Quality In Everything We Do

TABLE OF CONTENTS

Chapter 1

INTRODUCTION

1.1 Introduction

The business value of information has increased dramatically over the last few decades. Information systems have pervaded the business world in a rapid pace and have become critical assets in many organisations. Many organisations have become largely dependent on information and information systems to support their core business processes. Unfortunately, many threats to information and information systems exist today, which threaten the reliability of information (systems) and consequently business continuity. Some examples of common threats that information is faced with are hackers, computer viruses and human errors. Next to these common threats, one should not neglect the (fortunately less common) threats caused by natural disasters, such as earthquakes and floods, which can have devastating effects on the reliability of information.

Because of this dependence of businesses on information and the threats to that information, information security has become very important in most organisations. Information security forms the line of defence against threats to information (systems). That it is a hot topic is demonstrated by the various recent surveys on the topic and their results, e.g. [PWC05; EY05; DTT05; FBI05; CSO05].

To counter threats, and preserve the reliability of information and information systems, organisations implement various security controls. An example would be to have a fire extinguisher present in a server room, so damages resulting from a possible fire can be prevented or minimised. Adversely, threats to information (systems) are generally caused by current vulnerabilities in an organisation's security programme[1] [BABI05; PELT05]. Vulnerabilities are gaps in the security controls of an organisation through which a threat can circumvent the implemented security controls and materialise into a security incident, as depicted in Figure 1.1.

[1] It is not the case that vulnerabilities 'create' threats to information (systems). Rather, many threats exist. It is just that vulnerabilities in an organisation's security programme make it possible for threats to exert themselves and become (security) incidents. If there are no vulnerabilities in the security programme, threats are consequently not an issue, because everything is covered.

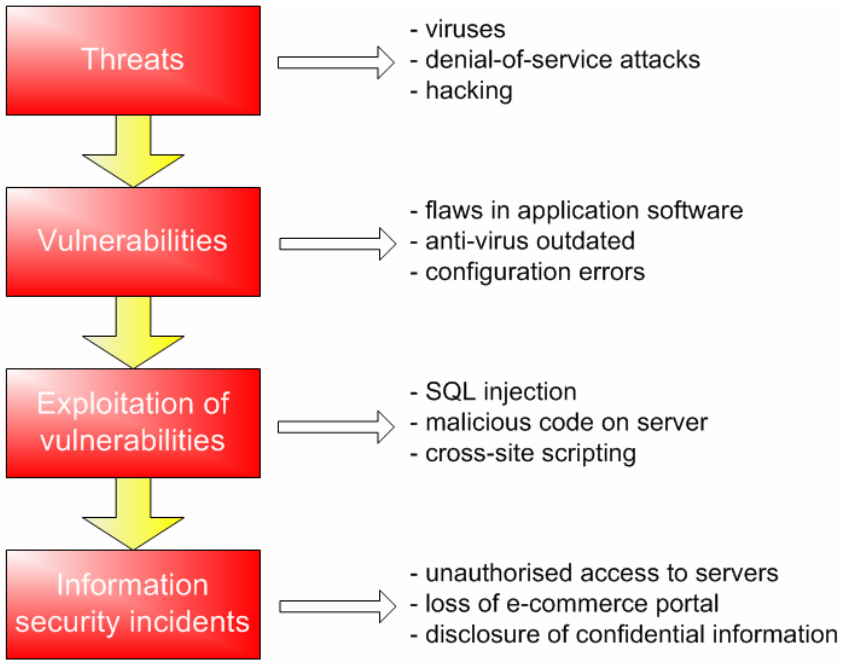

Figure 1.1 - How threats exploit vulnerabilities to cause information security incidents [adapted from ISIM06]

Information security incidents[2] (breaches of information security) can cause substantial damages to the reliability of information and information systems. This has become particularly true in today's globally interconnected world in which a seemingly small incident can have widespread consequences (in a worst case scenario). A lot of organisations are now interconnected with their suppliers and/or customers. If, for example, a computer virus outbreak is not contained in a timely fashion, the virus may spread to connected suppliers and/or customers. This example shows that the impact of an incident can easily transcend organisational borders. Furthermore, incidents can lead to great damages to an organisation. These damages are not limited to financial losses due to e.g. the incapacitation of critical business processes; other damages can include loss of reputation and the possibility of claims.

[2] The terms information security and security are used interchangeably in this thesis.

1.2 Inducement

In order to minimise damages of security incidents, organisations should establish some form of incident response as part of their incident management. According to IT executives who have dealt with public security breaches, organisations should implement predefined response programmes [RADC05]. Predefined response programmes state how to contain and resolve specific types of incidents. A predefined response can aid the efficient resolution of incidents when they occur. An example would be to have documented a response plan to handle a theft of sensitive customer information.

History has shown that in some organisations, despite the presence of security management, information security incident management is not optimal. Some incidents may not be recognised as such, because people may not be aware of potential indicators pointing towards an incident ('flawed' incident identification). Furthermore, it may be that the response to an incident is not predefined, but is instead determined after an incident occurs.

For example, an American healthcare organisation failed to report a mistake that exposed a handful of patient lab records. They started an investigation before reporting the incident, while they were obliged by law to report such incidents immediately upon discovery [RADC05].

In another example, AOL[3] acknowledged an "issue" that allowed some of its members to gain access to online financial portfolios of other members. But they downplayed the incident, saying no personal identifying information such as usernames or credit card numbers was ever compromised. An AOL user thought otherwise and reported the case to the U.S. Federal Trade Commission [ROSE04].

Additionally, in May 2006, electronic data containing the personal information of as many as 26.5 million veterans and some spouses has been stolen from the home of a Department of Veterans Affairs (VA) employee who violated agency policy by leaving the office with the information. The data, which resided on a laptop that was stolen in a burglary, contained the names, Social Security numbers and dates of birth of the veterans. The VA Secretary was extremely upset that he wasn't notified

[3] America Online Inc., the largest American internet service provider

earlier about the theft of the personal information of 26.5 million veterans. He was very concerned about the timing of the department's response once the burglary became known [SCMA06].
From these examples it seems that in some organisations there is room for improvement in the way they manage, and respond to, information security incidents. If an organisation does not know the appropriate (base) response to security incidents in advance, this may lead to untimely and improperly handled security incidents and possibly unnecessary damages resulting from such incidents.

Besides the issue of not having predefined incident response plans, another issue arises when incident identification falls short. If employees, or systems, do not recognise an incident as such, this may have damaging consequences. This can happen if potential indicators of an incident are not known (by systems and/or people); an example of such an indicator is internal port scanning, which may indicate the presence of a computer worm somewhere in the network. It can thus happen that the incident is first noticed when its damaging consequences surface. As a result, incident response may now only be able to take actions that minimise further damage. Furthermore, it may be that the help desk function does not identify a reported security incident properly and in a timely fashion. An example of ad hoc incident identification and response (meaning no indicators, no classification of incidents and no predefined incident responses) is depicted in Figure 1.1.

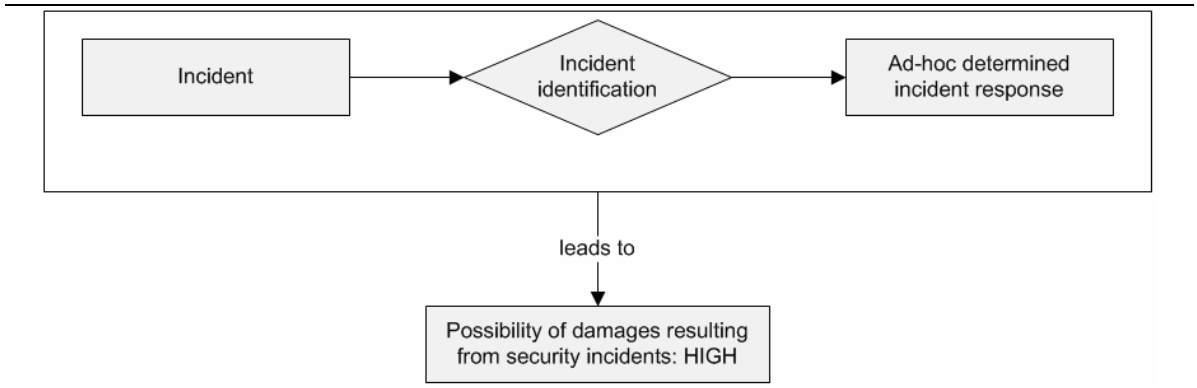

Figure 1.1 - Example of ad hoc incident response.

To optimise incident response (i.e. minimising negative consequences of an incident), people and/or systems have to be able to correctly and swiftly identify and resolve security incidents. Most organisations know that they need predefined incident response plans in place, as described in the various best practice methods [CERT]. But the people and systems in organisations also need to be able to identify incidents as such to carry out a response in the first place. Documenting possible information security incidents, how to recognise them (i.e., indicators) and how to respond to them may prove useful for enhancing security incident identification and response. To this end this research will provide a methodology for the aforementioned. This methodology should be established in such a fashion that it can help improve the efficiency and effectiveness of an organisation's information security incident identification and response. An example would be when a helpdesk can use the methodology to assess and identify an incident faster and more specifically. As a result, incident response can be more adequately tuned to the (type of) incident and any possible damages can be quickly contained. Figure 1.2 depicts a possible new process of information security incident identification and response with the support of the methodology.

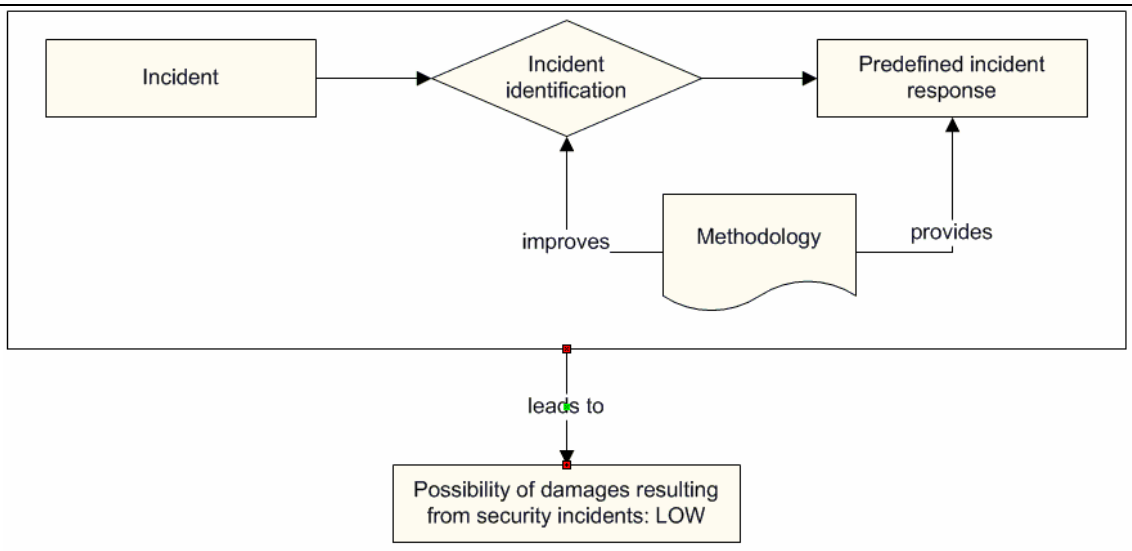

Figure 1.2 – Improved incident identification and response.

1.3 Research objective

From what was discussed in the previous sections, the following research objective is derived:

> *The research goal is to formulate a methodology that can be used to improve information security incident identification and response in organisations.*

To reach the goal of establishing such a methodology for improving information security incident identification and response in organisations, the following research questions need to be answered:

> *1. What is information security and what is its role in organisations?*
> *2. What is security incident management and incident response and what are current best practices?*
> *3. How can information security incident identification and response be improved?*
> *4. How does the methodology fit into current information security incident management processes?*
> *5. How can the methodology be applied in practice?*
> *6. Can the methodology improve security incident identification and response in practice?*

Furthermore, it is to be researched what the possible implications of the methodology are for IT auditing, especially regarding the audit of incident management processes. This leads to the following additional research question:

> *7. What are the implications of the methodology for IT auditing?*

1.4 Research methodology

To answer the questions mentioned above, and ultimately the research question, several things need to be done. The first is an extensive study of existing literature on the subjects that are part of the thesis. Among them are information security, information security incidents, information security management and information security incident management.

Furthermore, experts (e.g., IT auditors and IS specialists) will be asked to comment on and evaluate the methodology designed in this research to assess it's added value. These experts have extensive practical information security knowledge and their feedback forms an important input to the research. This feedback will be used for possible adjustment of the methodology and as a basis of subjects for future research.

1.5 Thesis Layout

The layout of this master thesis follows the research questions outlined in paragraph 1.3. Each research question will be answered in a separate chapter. In the final chapter, chapter 7, a short summary will be given as well as the conclusion of the research. Furthermore, some recommendations for future research will be given.

In chapter 2 we will take a look at information security in general. The importance of information security for organisations is discussed, as well as various aspects of information security. Additionally, some key fundamentals for proper information security are examined.
Chapter 3 describes the various aspects involved with security incidents, security incident management and security incident management best practices.
In the fourth chapter, a preliminary classification methodology of security incidents and associated proper response is constructed. The classification consists of information security incidents classified by type, linked to indicators and predefined responses.

In chapter 5, example processes are presented regarding the practical application and use of the information security incident classification methodology.

Chapter 6 discusses feedback on the methodology developed in this thesis. The feedback was acquired through one-on-one interviews with experts.

In chapter 7 the possible impact of the methodology on IT-auditing, specifically with regard to incident management, is examined.

Finally, in chapter 8, a conclusion is presented, along with some ideas and suggestions for future research.

Chapter 2

INFORMATION SECURITY

Research question 1: What is information security and what is its role in organisations?

In this chapter we will first take a look at the need for information security in safeguarding information and the systems providing it. The different reliability aspects of information security as well as the various control measures that protect those aspects are examined. Then we will take a look at some requirements for good information security and some information security standards and best practices.

2.1 The changing role of information

The field of Information Technology (IT) was formed by the introduction of the computer in the 1950s. From then on, IT continued to have a major impact on the way businesses work. Most of the major business changes of the last 50 years were driven by newly found technologies. Roughly, the evolution of IT can be divided into three eras [APPL03].

In the early days, IT was mainly used to automate tasks previously done by people. Since computers could do certain simple tasks much faster and efficiently than people, they took over those tasks.

In the second era, brought about by the introduction of the microcomputer, the focus of IT was on facilitating individual decision making and enhancing individual productivity.

The third era, which continues even today, is known as the network age (or internetworking age). Everything and everybody is connected, which can be seen by businesses that are becoming integrated with their suppliers and customers through electronic channels (vertical integration). The target for IT use became electronic integration and learning. This was also the first time that the importance of the information part of the concept of IT superseded the technology part. As opposed to the first two eras, in which the technology part of IT was a business driver, the information part of IT has now become a business driver. This greatly enhanced the business value of IT, as can be seen in Figure 2.1.

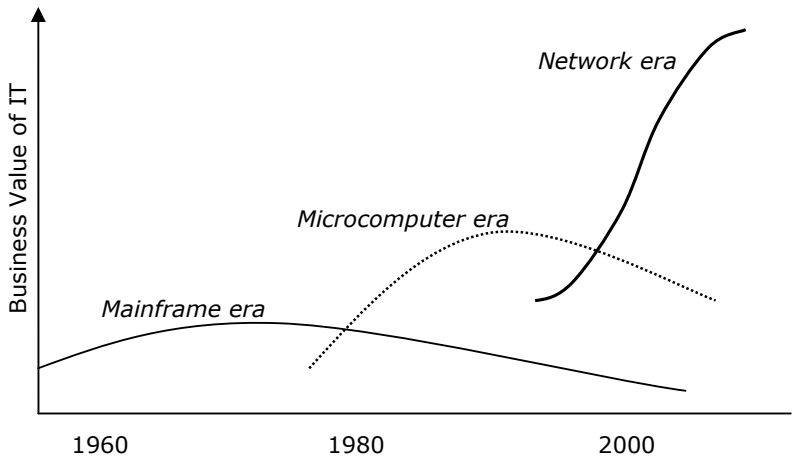

Figure 2.1 Business value of IT during the three IT eras [GIBS05]

Organisations today rely more and more on information and consequently information systems for providing the needed information to conduct business. More and more companies' core (or critical) processes heavily depend on information and information systems for their operation. By some, information has also been named the fourth factor of production, next to established factors nature, money and labour [OVER00].

2.2 The need for information security

The need for information security for an organisation is driven by the fact that there are threats to information (systems). These threats are discussed in paragraph 2.2.1. Furthermore, a distinction can be made between the internal and the external need for ensuring the reliability of information, which will be discussed in paragraph 2.2.2.

2.2.1 Threats to information

Information security would not be an issue if there were no threats to information (systems). A threat is an unwanted, intentional or accidental, event that may cause harm to an organisation's assets,

through a breach in reliability. Damages that could arise when threats materialise (i.e., turn into an incident) are loss of money, loss of business, loss of reputation, claims, etc. [BERG04]

Threats to information can be categorised in various ways. A common method is to classify threats based on their origin: accidental, intentional or natural (the first two are frequently considered acts of men; the latter are regarded as acts of God). Examples of accidental threats are human errors, power failures, equipment failures, etc. Intentional threats are such things as fraud, information theft and sabotage. Natural threats are, for example, earthquakes, lightning strikes and floods. Another classification divides threats in intentional and unintentional threats, where unintentional threats encompasses both the aforementioned accidental and natural threats.
Another more extensive division of causes of security threats is given by Mollema [MOLL05]. His classification of so called 'inherent exposures' consists of bad behaviour, mismanagement, mistakes, systems failure, technical infrastructure failure, social/political unrest and acts of God.

Generally, effects of threats can be divided into five distinct groups: interception, interruption, modification, fabrication and destruction [BABI05]. [BISH05] refers to a similar division of threats: disclosure (unauthorised access), deception (acceptance of false data), disruption (interruption/prevention of correct operation) and usurpation (unauthorised control).

2.2.2 Internal and external need for information security

A distinction between the internal and external need for information security within an organisation can be made. The most apparent need for information security comes from within the organisation itself, the internal need. As discussed in the introductory chapter, organisations are more and more dependent on information for the operation of their (core) business processes. If the reliability of information is breached, it could have devastating effects for the organisation: loss of customers, loss of image, financial losses, etc. Information security has become essential for assuring reliability of information that provides uninterrupted, reliable business conduct.

In addition to the internal need for information security, there is also an external need for information security in organisations. There are many external stakeholders that demand that an organisation secures its information. Possible stakeholders are customers, stockholders and suppliers. Furthermore, organisations have to abide by the laws, rules and regulations that apply in the countries where they conduct their business.

An example of local law is the Sarbanes-Oxley Act (SOx), which imposes requirements with respect to internal control and reporting on companies doing business in the USA[1].

The SOx law was a response to the extravagant conduct of managers and directors in firms such as WorldCom, Enron, and Arthur Andersen. Their market values were artificially inflated (whether intentional or not) to the point where they no longer had any relation to reality. This led to a downslide in stock markets that began in March 2000, which ultimately led to the collapse of the New Economy[2]. Investors needed to be protected from such extravagant conduct of organisations and thus existing accounting practices needed to be restructured. [BLOE06]

2.3 Definition of information security

When looking for definitions of information security, it can be found that a lot of people that try to define information security have difficulties defining what it exactly is. The U.S. Committee on National Security Systems[3] defines information security as:

> *"[Information security entails] the protection of information (systems) against unauthorised access to or modification of information, whether in storage, processing or*

[1] All businesses that have $10 million in shares in the United States or 500 shareholders of whom 60 percent are American are affected by this law. Also, all companies listed on American exchanges are under its jurisdiction. […] Foreign-based firms must comply with Sarbanes-Oxley, if they wish to continue listing financial instruments for trading on U.S. exchanges or if they intend to do so in the future. [BLOE06]

[2] The New Economy is a term that was coined in late 1990s to describe the evolution of the world economy into a high technology-based economy, arising largely from new developments in the Internet, telecommunications and computer sectors.

[3] The CNSS provides a forum for the discussion of policy issues, sets national policy, and promulgates direction, operational procedures, and guidance for the security of United States national security systems. For more information, see http://www.cnss.gov.

> *transit, and against the denial of service to authorised users or the provision of service to unauthorised users, including those measures necessary to detect, document, and counter such threats."*

Although the abovementioned definition encompasses a lot of the aspects important to information security, it overlooks the fact that information security is a process. Another definition, by [OVER00], does emphasise the process aspect of information security. This definition is stated below.

> *Information security is a process that deals with taking and maintaining a coherent set of measures to protect the reliability of an information facility. [OVER00]*

So, organisations should take a process approach to information security, instead of just implementing some technical security measures and leaving it at that; information security is not a product, nor is it a technology. Information security is a process, which consists of many elements including security policies, procedures and training. Information security includes, among others, control measures (see paragraph 2.4.2), a degree of security awareness, incident management (see chapter 3) and business continuity. [PALM05]

The information security process is the method an organisation uses to implement and achieve its security objectives[4]. These security objectives have to be in line with, and support the realisation of, the business objectives of the organisation. The information security process should be a continuous effort in order to provide adequate security for the organisation.

Another important issue is that the information security process entails a way of thinking that must permeate an organisation and its culture to be effective. [PALM05]

Additionally, it must be noted that information security is not the same as computer security in any of its forms (e.g. computer and network security). Rather, these are all subsets of information security. Information security covers all infrastructures that facilitate the use of information (e.g. processes,

[4] Information Security – Security Process. Federal Financial Institutions Examination Council. http://www.ffiec.gov/ffiecinfobase/booklets/information_security/01_security_process.htm

systems, services, technology), not just information itself. Information security is also not confined to computer systems and information in an electronic or machine-readable form. Information security covers all aspects of safeguarding all types of information or data.

2.4 Aspects of information security

In this section several aspects of information security are reviewed. First, the various reliability aspects of information security are discussed. Then, different security controls are examined.

2.4.1 Reliability aspects of information security

Basically, information security deals with protecting three different "reliability" aspects of information: confidentiality, integrity and availability, which can be remembered by the mnemonic "CIA" and are frequently referred to as the CIA triad [PELT05]. These three widely accepted attributes of information security are stated below [PELT05, BISH05].

Confidentiality

Confidentiality is the concealment of information or resources and is defined by ISO-17799 as "ensuring that information is accessible only to those authorised to have access to it." So, to attain confidentiality, a business needs to keep secret information secret. This also means that only certain people should know about the existence of certain information in the first place, on a need to know basis.

Integrity

Integrity refers to the trustworthiness of information or resources and is defined by the ISO-17799 standard as "the action of safeguarding the accuracy and completeness of information and processing methods." When a user requests any type of information from the system, the information will be complete, correct and up to date. Also a distinction between the trustworthiness of the content of the information and its origin can be made (i.e. data integrity and origin integrity, [BISH05]).

Availability

ISO-17799 defines availability as "ensuring that authorised users have access to information and associated assets when required." Availability requires measures to ensure timeliness and continuity of information, so that business processes don't come to a halt.

2.4.2 Security controls

To counteract threats, organisations implement various countermeasures. These measures can be of two different classifications: what can organisations do and when can they do them. The first taxonomy, what organisations can do, consists of the following controls: physical, organisational and logical [BERG04, PATH05]. They are described below.

Physical controls

This term is commonly used for describing protection measures that try to protect information process equipment from physical threats. Goal is preventing damage from e.g. malfunction, unauthorised access, physical damage and theft. Examples of physical measures are gates, locks and tourniquets.

Logical controls

These measures are implemented to protect software and information (or data). Examples of these measures are authentication and authorisation controls (access controls), encryption, security certificates and virus protection. The goal is to prevent damage from e.g. unauthorised access, mistakes and fraud.

Organisational controls

These controls provide a context in which the system of physical and logical controls can work, they provide measures to complement the system of physical and logical controls and they (should) support the realisation of all security goals. Some examples are segregation of duties and security policies.

Nowadays, legal controls are by some considered a fourth category of controls [PATH05]. These controls consist of legislation with regard to privacy, intellectual property rights and of course laws such as the Sarbanes-Oxley Act (SOx).

The second taxonomy of measures is based on when they assert themselves and is displayed in Figure 2.4. Preventive controls try to prevent threats from materialising into incidents. Detective controls are the first type of repressive controls; they follow an information security incident, to prevent (further) damaging consequences of that incident. Corrective controls are the second type of repressive controls; they follow an information security incident that has already had damaging consequences (i.e. breach of one or more reliability aspects).

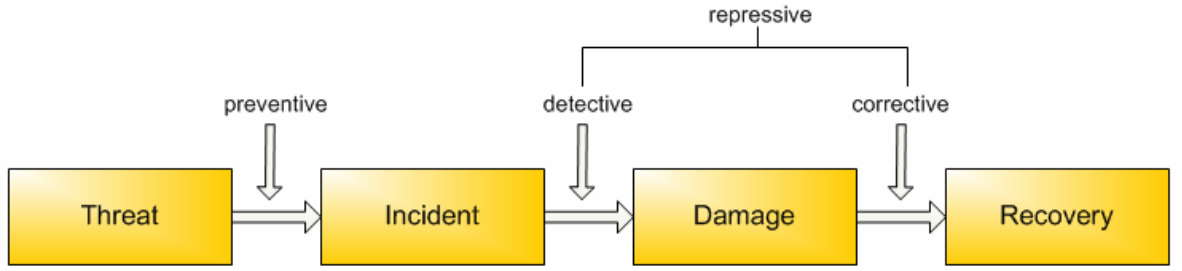

Figure 2.4 Security controls from threat to recovery [BERG04]

As can be seen in the picture, preventive controls stop a threat from developing into an incident. Detective and corrective controls however, act upon the materialisation of a threat.

2.5 Requirements for good information security

Just having the aforementioned security controls in place doesn't guarantee good information security. The most important requirement is that an information security policy should be in place, which is discussed in paragraph 2.5.1. Furthermore, for an organisation to have good information security its security programme should incorporate several key characteristics [BABI05], which are discussed in paragraph 2.5.2.

2.5.1 The security policy

The security policy is probably the most important aspect of information security, because it serves as a foundation for all the other security components. It is an enterprise-wide policy that steers the execution of information security in order to prevent breaches of reliability of information (systems) and to control the consequences of possible violations [STOF04]. The security policy is essential for a proper operation of the information security process. Its main goals are shaping and directing information security, get management support and affirmation, setting responsibilities within information security and satisfying legal, social and commercial requirements [STOF04]. A good example of the importance of security policies is provided by [PELT05]: "in the documentation for their Cisco[5] PIX product, the folks at Cisco even refer to the security policy as the centre of security".

[STOF04] found several definitions of security policies, and these are the main characteristics of security policies he distilled from those definitions:

- Organisational policy
- Propagated by board-level management
- Possess strategic goals, procedures and guidelines
- Aimed at securing information processing (information security)
- Steering and supporting middle management and business processes
- Specify how to react to breaches of reliability of the information facilities

[5] Cisco Systems, Inc. is the worldwide leader in networking for the Internet through hardware, software, and service offerings.

2.5.2 Aspects of a strong information security programme

There are 6 characteristics which provide a suitable foundation by which management can design and implement security that is aligned with the business objectives and supported throughout the organisation [BABI05].

Aligned

For every organisation it is important to align information security to the business objectives a company has. This means that a company can not look at another company and just copy their security programme, or copy from a standard manual of some sort. It is important to keep security and the business objectives aligned at all times to ensure (information) security aids the company in reaching its objectives. Furthermore, security has to be recognised as a critical element of organisational performance by executive management, not just as a sinkhole.

Enterprise-wide

For a security programme to be successful, it needs to be deployed on an enterprise-wide scale. Actually, not only the security needs of the entire organisation need to be taken into account, but also those of the suppliers and the customers of the organisation. If any business unit is less secure than other units, this may undermine the secure state of those other business units.

Continuous

An organisation's security programme is not a one time effort which provides continuous protection. Rather, the opposite is true. A security programme needs to be continually maintained to provide sufficient protection. This is largely due to the rapid technological changes that organisations face nowadays and the fact that new technologies are often implemented without analysing the security risks[6] they pose.

[6] A [security] risk can be defined as the probability that a threat against a vulnerability would produce an impact. [STEP05]

Proactive

Having a proactive security programme means that known threats and risks should be mitigated and a predefined plan for incident response should be in place. A proactive approach to information security requires the organisation to provide sufficient resources for enabling learning of vulnerabilities, monitoring common attack patterns and limiting unauthorised activity.

Validated

No matter how good the design, planning and alignment of a security programme, if it is not independently validated it's not worth anything. The critical security elements all need to be reviewed, just as the business goals. The extent to which security needs to be validated independently depends on the risks associated with the security component undergoing the validation. Components associated with high risks need to be independently tested, by an unbiased party. When there are only low risks involved, testing can be done internally (e.g. self-testing). Nonetheless, aside from the risk, all components need to be validated both thoroughly and repeatedly.

Formal

Formalising a security programme with policies, guidelines, standards, etc. makes it more effective. Policies, standards and guidelines provide fundamental direction on security issues. Not only does this formal programme need to be in place, it is also required that everyone in the organisation is informed about it. Ultimately, a formal programme requires documentation and confirmation of organisation, process and technology.

2.5.3 Limits to information security

When looking at the possible damages of security incidents, it becomes clear that they need to be resolved as efficiently and effectively as possible in order to minimise damages and prevent similar incidents in the future. Even better would be to prevent incidents from happening by making sure that there are no security vulnerabilities, but that isn't an option. This is because there are certain limits to information security. That a security programme has some weaknesses does not mean it is

not a good security programme; these vulnerabilities are inherent to any security programme. This is due to multiple factors.

First, there is the non-perfectible nature of information security, which means that no one can ever eliminate all risk of improper or unpredictable use of any information [SCHN00].

Secondly, there is no need to perfect information security, because the level of information security required in any particular situation should correspond with the value of the information and the possible loss that might result from improper use [SCHN00].

A third factor, which introduces a lot of risk to organisations, is the increasing gap between spending on IT and spending on security; spending on IT in general rises significantly every year, while the spending on security really lags behind. This is called "the digital security gap" by [BABI05] and is illustrated in Figure 2.3. Instead of the lines moving parallel to each other, they diverge over time. This means that the percentage of total spending on IT that is spent on security is dropping. This results in less secure systems and consequently a higher likelihood that (security) incidents will occur.

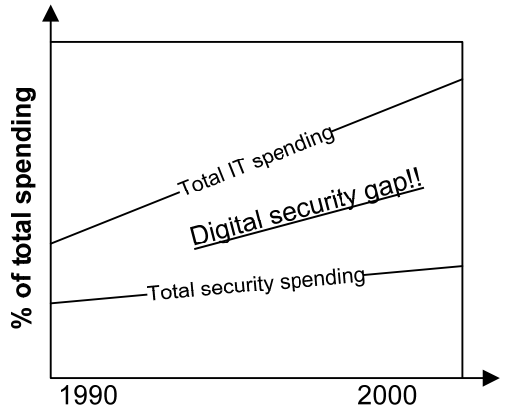

Figure 2.3 The digital security gap [BABI05].

An ideal situation would be when security spending would follow total IT spending in a way that it is at least at a fixed percentage, e.g., security spending at 15% of total IT spending. Unfortunately, what is happening is that organisations introduce a lot of new technologies in their businesses, but do not invest in securing those new technologies in order to preserve the reliability of information (systems).

2.6 Information security standards and best practices

There are several security standards and best practice models available for information security. There seems to be a growing interest among organisations in information security standards and certification and they are increasingly looking to adopt standards [EY05]. Standards can not only provide a framework for implementing effective information security practices, they can also make sure that information security and organisational objectives are properly aligned. Furthermore, organisations recognise that standards demonstrate to clients and customers their commitment to good information security practices. Four of the more widely used standards will be briefly discussed here, namely ITIL, COBIT, The Standard of Good Practice for Information Security, and ISO17799 (BS7799).

2.6.1 ITIL

The Information Technology Infrastructure Library (ITIL) is a framework of best practices that promote quality computing services in IT sector. ITIL was first developed by the British Central Computer & Telecommunications Agency, which merged with the UK Office of Government Commerce (OGC) in 2001. The OGC took on further development of the framework.

ITIL presents a broad set of management procedures, which apply to all aspects of IT infrastructure, with which an organisation can manage its IT operations. So, its focus is not solely on information security, but rather on the management of IT processes. Nevertheless, ITIL, and more specifically the ITIL Security Management process, is widely used for the implementation of information security within an organisation.

ITIL is currently embodied in the ISO 20000 standard. In December 2005, the OGC issued notice of an ITIL refresh (i.e., an update), ITIL v3, which is planned to be available late 2006.

2.6.2 COBIT

The Control Objectives for Information and related Technology (COBIT) is a set of best practices for information (IT) management. COBIT provides a set of generally accepted measures, indicators, processes and best practices to assist organisations in maximising the benefits derived through the use of IT and developing appropriate IT governance and control in an organisation.

The COBIT process model consists of 34 high-level control objectives (processes) that are spread out over the following four process areas: Plan and Organise (PO), Acquire and Implement (AI), Deliver and Support (DS), and Monitor and Evaluate (ME). Each high-level control objective consists of multiple detailed control objectives (also defined as processes), for a grand total of 215 detailed control objectives (COBIT 4.0). For each process, a 5-stage maturity model is defined, which shows how mature an organisation is with regard to that specific process (see [MEUL05] for more information on (security) maturity).

The COBIT framework links IT to the business goals, organises IT activities into a generally accepted process model, identifies the major IT resources to be leveraged, and defines the management control objectives to be considered. The business orientation of COBIT consists of linking business goals to IT goals, providing metrics and maturity models to measure their success, and identifying the associated responsibilities of business and IT process owners.

Although the focus of COBIT is not on information security alone, there are many references to security and security related topics throughout the framework.

2.6.3 The Standard of Good Practice for Information Security

The Standard of Good Practice for Information Security addresses information security from a business perspective, providing a practical basis for assessing an organisation's information security arrangements. It focuses on the arrangements that should be made by leading organisations to keep the business risks associated with critical information systems under control in today's dynamic and competitive environment [ISFS05].

The Standard is aimed at organisations that recognise that information security is a key business issue. Good practice detailed in the Standard will usually be incorporated into an organisation's information security plans by various key individuals, including information security managers, business managers, IT managers and IT audit managers.

The Standard is based on in-depth research and the extensive knowledge and practical experience of Information Security Forum (ISF) members. It is updated at least every two years in order to reflect the most up-to-date thinking in information security and to include the latest 'hot topics'. Other international and national standards (such as ISO 17799) are also used as sources of information.

2.6.4 ISO 17799

ISO 17799 establishes guidelines and general principles for initiating, implementing, maintaining, and improving information security management in an organisation. It consists of 127 security controls which are taken from the first part of BS7799, the British Standard for information security. The objectives in ISO17799 provide general guidance on the commonly accepted goals of information security management. ISO 17799 contains best practices of control objectives and controls in the various areas of information security management, such as security policies, organisation of information security and business continuity management. The control objectives and controls in ISO 17799 are intended to be implemented to meet the requirements identified by a risk assessment. Furthermore, it is intended as a common basis and practical guideline for developing organisational security standards and effective security management practices.

2.6.5 Usage of best practices and standards

The usage of standards is growing, according to the survey by [EY05]. Twenty-five percent of surveyed organisations are using ISO 17799 (i.e. they are certified), and another 30% of them are planning to do so. Usage of ITIL is also gaining, even though security is not its primary focus. Almost a quarter of survey respondents are currently applying it, and an additional 22% are planning to do so. The adoption of COBIT and The Standard lags considerably behind that of ITIL and ISO17799. 18% of respondents have adopted COBIT, and another 18% is planning to adopt it. The Standard gets the least support of the four standards discussed. Only 12% has adopted it, and a meagre 8% is going to do so [EY05].

2.7 Chapter summary

In this chapter the role of information security in organisations was discussed. It was shown that information plays a huge role in organisations these days and that protecting that information has become vital for the business continuity of most organisations. It was mentioned that information security is to be seen as a continuous process, not some separate set of activities. Furthermore, a range

of information security aspects was examined, such as the reliability aspects of information, different security controls and requirements for a successful security programme. It was shown that there are various factors that impose limits on information security, such as security spending that is lagging behind total IT expenditures. Finally, several best practices and standards were addressed.
In the next chapter we will take a look at the incident management processes of the best practices that were examined in this chapter.

Chapter 3

SECURITY INCIDENT MANAGEMENT

Research question 2: What is security incident management and incident response and what are current best practices?

This chapter tries to answer the above mentioned question. First of, we need to have a clear definition of what information security incidents are. Then, we will look at how security incident management and security incident response as a part of that incident management are organised in various best practice frameworks. After discussing a best practice, a short critical reflection will be given, which summarises some points that could be improved.

3.1 Defining security incidents

Without threats ever materialising and vulnerabilities ever being exploited, there would not be any security incident. However, this is clearly not the case: security incidents occur frequently within organisations [e.g. E&Y, 2005; DTT, 2005; FBI/CSI, 2005]. Results from a survey by the ISF show that incidents erode companies' profits, depress the value of the business and compromise future earnings [ISFS05].

3.1.1 Definition of security incident

For the purpose of this thesis a clear definition of the term security incident[1] needs to be given. A (rather general) definition of security incidents is given by CERT/CC[2]:

[A security incident is] the act of violating an explicit or implied security policy.

[1] The terms security incident and information security incident are used interchangeably in this thesis, given the context in which they are used, namely information security.
[2] Computer Emergency Response Team Coordination Center (CERT/CC) - Incident Reporting Guidelines. http://www.cert.org/tech_tips/incident_reporting.html

Another definition that emphasises information security incidents is given by [ITIL99, p. 9], and is stated below.

> *Information security incidents are those events that can cause damage to confidentiality, integrity or availability of information or information processing.*

The abovementioned definitions of security incidents do not make a distinction between intentional and unintentional acts, or acts of men and acts of God (see section 2.2); an incident is a security incident when it impacts (information) security. For instance, a flood that disables several workstations and subsequently compromises availability of information (processing) for one or more employees is considered an information security incident.

There may be some cases in which it is not really clear whether an incident is really a security incident, but ITIL states that a general rule for dealing with uncertainty in such cases is to treat the incident as a security incident (i.e. better safe than sorry) [ITIL99].

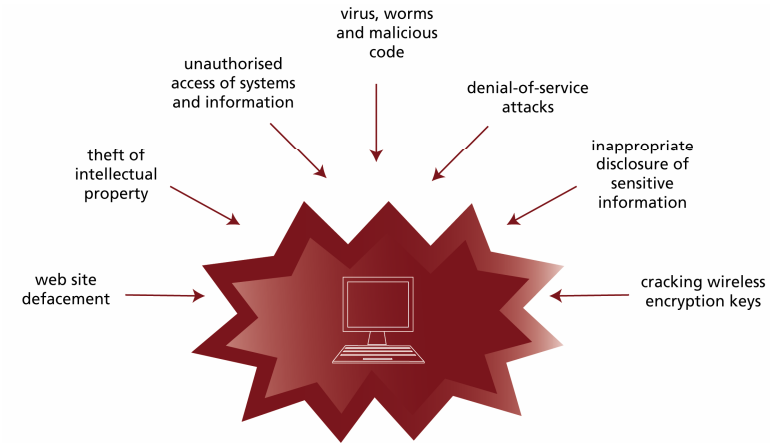

Figure 3.1 - Examples of information security incidents [ISIM06]

3.1.2 Aspects of security incidents

Security incidents can be described by many different aspects related to them. Some examples are nature of the incident, intentionality, associated risks and (possible) consequences, and (potential) indicators. In this section a variety of these associated aspects is briefly examined.

Nature

The nature of an information security incident specifies its origin (see also section 2.2, where the nature of threats was briefly discussed). There are several ways to classify information security incidents by their nature, but a commonly used method when classifying by nature is to divide information security incidents into two groups: acts of men and acts of God. Acts of God are related to things as natural disasters (floods, fire) or power failures. Acts of men however, are security impeding acts caused by people, for example when a hacker is hacking into a company's database to extract or expose sensitive data.

Intentionality

Intentionality of a security incident is another criterion by which information security incidents can be described (see also section 2.2, where threat intentionality was discussed). The distinction here seems easy to make: an information security incident is either caused intentionally or unintentionally. However, there are cases in which it is not overly clear whether a security incident is caused intentionally or accidental. For example, an employee unknowingly opens an e-mail attachment that contains malicious code of some sort, e.g. a virus. In this case one could argue that the incident was caused unintentionally, because the employee did not know the e-mail was infected (lack of security awareness). But, it could also be argued it was caused intentionally, because the one who created the malicious code and released it did that on purpose.

The categorisation of incidents by intentionality can be linked to the classifying of information security incidents by their nature, i.e. acts of men and acts of God. Acts of men can be either intentional or unintentional, where acts of God are always unintentional (at least from the scientific point of view, religious views might differ).

Associated risk

The associated risk of an information security incident differs per individual incident. There are security incidents that will have no impact on the business continuity of an organisation, but are more of a nuisance and a drain on resources because they have to be resolved. On the other hand, there are security incidents that do pose serious risks with regard to an organisation's business continuity.

An example of a low risk security incident would be when a user's personal computer gets infected with spyware, which can cause a slight drop in productivity and may necessitate intervention by a member of the help desk, who removes the spyware from the system. There could also be a breach of confidentiality, because the spyware may send usage statistics to an (unknown) third party.

When a hacker would succeed in breaking into a credit card company's database and extracting sensitive customer data from it, the company faces several possible negative consequences. They would probably lose reputation, which in turn means that they would lose customers and thus revenue. Furthermore, they would likely be held responsible for the loss of privacy sensitive data and could be sued by afflicted customers. All these negative consequences, when added up, can seriously threaten the organisation's business continuity.

Because of the possible extent to which an information security incident can damage an organisation (see the example mentioned above), it becomes apparent that insight in the possible associated risks of an information security incident is necessary in order to adequately manage these risks.

Indicators

Besides the aforementioned aspects of security incidents, there is another aspects related to security incidents that is important for incident response. This aspect is the potential indicator(s) of a security incident. An indicator can be defined as[3]:

> *"something that helps us to understand where we are, where we are going and how far we are from the goal. […] it can be a sign, a number, a graphic and so on. It must be a*

[3] http://hostings.diplomacy.edu/baldi/malta2001/statint/Statistics_Int_Affairs-27.htm.

> *clue, a symptom, a pointer to something that is changing. Indicators are presentations of measurements. They are bits of information that summarize the characteristics of systems or highlight what is happening in a system."*

Indicators may thus point out that something has happened, is happening or is going to happen. With regard to the abovementioned definition of indicators, a security incident indicator may point out the fact that a security incident has taken place, is taking place or will be taking place. This information may prove valuable for organisations when they try to determine whether a certain (system) state or event signifies a security incident. These security incident related indicators could thus serve as a part of a security incident classification.

Another topic in which indicators are commonly used nowadays in is terrorism prevention. There are 7 general indicators[4] that may point towards possible terrorist activity. Among these are surveillance (e.g., checking the target area for the terrorist activity), elicitation (gaining information about people, places, etc.), security testing (probing the target, e.g. gathering information on law enforcement response times) and suspicious people who do not belong (e.g., people who do not fit in due to their demeanour or behaviour).

Now, some examples of security incident indicators will be given. Some of these indicators fall along the lines of the terrorism indicators (e.g., surveillance) mentioned above. For instance, scanning of ports on computers can hint towards someone who is possibly attempting to break into company systems, for whatever reason. In another example, a user might experience suspicious processor activity, which might indicate an unauthorised piece of software (e.g., spyware or malware) running on the user's computer. As a concluding example, the company's website suddenly becomes unavailable due to extraordinarily high server load, which might point in the direction of a DoS[5] attack.

[4] http://www.njcounterterrorism.org/pdfs/7-Terrorism-Signs.pdf.

[5] DoS stands for denial-of-service, which is an attack on a computer system or network that causes a loss of service to users, typically the loss of network connectivity and services by consuming the bandwidth of the victim network or overloading the computational resources of the victim system.

From these examples it becomes clear that indicators can play a role in the successful identification of security incidents. If such indicators are linked to security incidents, an organisation's incident liaison[6] can identify security incidents faster and more accurately. When these incidents are then also linked to appropriate incident responses, indicators form the first step towards successful incident resolution.

3.2 Incident management in best practices and standards

As was already mentioned in the previous section, information security incidents can depress the value of the business and compromise future earnings [ISFS05]. Given the fact that information security incidents will happen sooner or later and the possible impact these incidents may have on the availability, integrity and confidentiality of information (which, in turn, could possibly result in damages to the organisation), it is paramount for organisations to effectively and efficiently resolve these incidents. In other words, information security incidents need to be properly managed. In the following paragraphs the incident management components of various best practices are reviewed, namely those of ITIL, COBIT and The Standard of Good Practice for Information Security (see section 2.6 for a brief overview of these best practices). These best practices provide information on how to manage (information security) incidents. The aim is to discover how these best practices organise the incident management process, what their key components are and how a classification method for security incidents can contribute to these best practices.

3.2.1 ITIL

As was stated in section 3.1, ITIL defines information security incidents as "those events that can cause damage to confidentiality, integrity or availability of information or information processing."
The incident related processes within ITIL Security Management that have the greatest linkage with the research topic are Incident Control and Problem Management. There are several more processes

[6] The incident liaison is an organisation's central incident handling function, usually a help desk function of some sort. See section 3.2.1 for more detailed information.

besides the two aforementioned ones that are important for security management as a whole, but incorporating them here would be beyond the scope of this research. The linkage between Security Management, Incident Management and Problem Management is depicted in Figure 3.1.

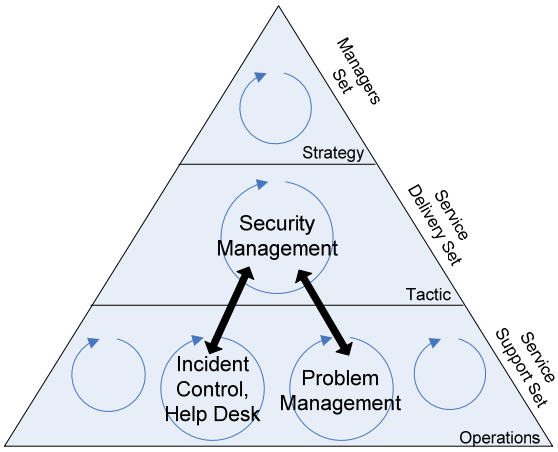

Figure 3.1 – Linkage between security management, incident control and problem management in ITIL (adapted from [ITIL99])

In this figure it can be seen that the Security Management process is situated on a tactical level. This is where e.g. SLAs are drawn up and service is provided in accordance with the security measures in those SLAs. The operational level is where the beneficial processes for service delivery (i.e. Security Management) are grouped, among which are Incident Control and Problem Management.

In ITIL [ITIL99], Incident Control (IC) (or Help Desk) is defined as a single contact point within the organisation, whose core process is usually incident management. Their aim is to provide continuity of services to the users of those services, which is also the most important goal of incident management. Furthermore, Incident Control is concerned with how quickly a solution to an incident is given, whereas the quality of the solution is part of the Problem Management process [ITIL99, p. 45-46].

Incident Control basically deals with three core activities: administration, monitoring and management of incidents. Incident Control is the so-called 'owner' of all security incidents and they are the single point of assistance for security incident reporting and first-line help. The input of Incident Control mostly consists of users' reports and automated monitoring reports. Incidents are categorised and for each category there exists some predefined set of activities that have to be carried out by Incident Control in order to contain any possible negative consequences. Reported security incidents are prioritised based on their impact and linked to the part(s)[7] of the information infrastructure they affect.

It is stated in ITIL that it is essential to recognise an information security incident as such. For information security incidents a different procedure may apply in comparison with non security – related incidents. It is obvious that only by applying correct procedures, incidents will be dealt with appropriately. If users (and/or systems) are not able to properly recognise security incidents as such (see section 1.2), a problem may arise when those users (and/or systems) are responsible for the reporting of information security incidents (incident reporting means notifying the appropriate contact point the moment an incident is noticed).

The importance of security incident reporting is explained by ITIL by the fact that it is the first step towards the resolution of a security incident. If a security incident is not reported in a timely fashion, or not reported at all, it may escalate and begin to (further) exert its associated damaging consequences. Incident reporting should be facilitated for all users and, on exception, accepted anonymously, in order to lower the reporting threshold some users might experience. A standard procedure for security incident reporting should be described in a manifest of some sort (e.g. a security handbook) and communicated to all employees.

[7] The smallest unit that is managed individually: a Configuration Item (CI). All the CI's together form the entire IT infrastructure. An overview of all the CI's is called a Configuration and Asset Management Database (CMDB). Some examples of CI's are software, hardware and documentation.

ITIL advises to include some examples of, and procedures for, security incidents in an SLA. Furthermore, it is proposed that organisations should also consider classifying security incidents by confidentiality and agree on a procedure about communication of security incidents (both internal and external).

ITIL states that the proper handling of information security incidents is comprised of several parts. First, to successfully solve security incidents, it is imperative that proper corrective action is taken, i.e. correcting what needs to be corrected. Furthermore, measures should be taken to prevent reoccurrence of the same (type of) information security incident. Some actions that may be taken are revision of involved security measures (where/how did current measures fail?) and imposing sanctions on the one(s) responsible for a security incident.

Finally, the importance of security incident registration is made clear in ITIL by the notion that such registration offers the possibility of analyses on past incident responses, e.g. effectiveness of applied countermeasures and the possible extent of damages. Such analyses form inputs to problem management, as will be described in the following paragraph.

Problem management

Problem management in ITIL is focused on discovering and tracing the causes of incidents. Once the cause(s) of an incident is determined, appropriate prevention measures can be formulated and implemented, to ensure that that specific security incident will not reoccur. The process is meant to manage incidents, establish links between incidents, systematically solve incident causes and finally eliminate incidents. Solutions for incidents are documented and controlled by Incident control. Known errors and requests for changes (RFCs) form the input for the change management process, which deals with implementing the needed changes for the prevention of incident reoccurrence.

It may be the case that there are problems related to an incident. In that case, it may be appropriate to follow a special procedure. Three specific issues are important: First, minimise the extent to which

knowledge of the incident is spread by considering the people that are somehow involved with the incident. Secondly, think about the people who need to be involved in order to properly solve the security incident. And finally, the solution to the incident should not introduce new problems (e.g. new security vulnerabilities).

Summary and discussion

The ITIL incident management process can be summarised as follows. Incident Control, a central point of contact for (security) incidents, deals with the administration, monitoring and management of incidents. Incidents should be categorised and sets of activities for categories of incidents should be predefined. Incident recognition, reporting and registration are seen as vital to the resolution of incidents. Problem management focuses on discovering and tracing the causes of security incidents and preventing reoccurrence of security incidents. Solutions for incidents are documented and controlled by Incident Control for future reference.

Although the incident management process in ITIL, and ITIL as a whole for that matter, is a best practice, there are some shortcomings to it. First, ITIL gives an overview of what they believe are requirements for proper security, incident and problem management. However, they do not provide a clear visualisation of the processes, e.g. in the form of a schematic process overview of some sort. This makes the processes unnecessarily abstract and difficult to grasp.

Another shortcoming is that although the significance of classification of security incidents is emphasised, there is no concrete method on how to set up (i.e., design) such a classification, or by what aspect incidents should be categorised. This is probably due to the fact that ITIL mostly focuses on what organisations should do, not how they should do them. In light of this 'shortcoming' of ITIL, the methodology for setting up a security incident classification, which is the aim of this research, should provide organisations with a clear, operational method of how to classify security incidents.

3.2.2 COBIT

As described in section 2.6, COBIT [COBI05] is a framework for IT governance and control, which consists of 34 high-level control objectives (i.e. processes) that are spread out over the following four

process areas: Plan and Organise (PO), Acquire and Implement (AI), Deliver and Support (DS), and Monitor and Evaluate (ME).

The main high-level control objectives that are related to security incident management are part of the DS process area, namely "Manage service desk and incidents" (DS8) and "Manage problems" (DS10). These objectives do not merely focus on incidents, but for the sake of the research subject the focus in this section will be on the incident aspects of those objectives. One other high-level control objective is of importance, namely "Ensure systems security" (DS5). This control objective provides input to the "Manage service desk and incidents" objective.

To start managing information security incidents, they first have to be defined. Besides the definitions mentioned in section 3.1, COBIT describes a detailed control objective for "Security incident definition" (DS5.6), which can be found in the "Ensure systems security" objective (DS5) of the framework. "Security incident definition" (DS5.6) forms an input to "Manage service desk and incidents" (DS8). DS5.6 states that organisations should make certain that the aspects of potential security incidents are clearly defined and communicated so security incidents can be appropriately treated by the incident and/or problem management processes. Aspects include a description of what is considered an information security incident and for example its level of impact (i.e. on an impact scale of some sort). COBIT prescribes that for each impact level, required actions and people who need to be notified have to be defined.

Manage service desk and incidents

COBIT states that a service desk function should be established to register, communicate, dispatch and analyse all reported incidents. Monitoring and escalation procedures need to be in place, so reported incidents can be properly classified (e.g. by type, group or domain), and prioritised (e.g. based on the incident's severity). Incidents should be classified according to a business and service priority and forwarded to the proper problem management team. Furthermore, a function and system should be established that allows tracking and logging of incidents. The people involved with the incident should be kept informed of its status.

For incidents that can not be immediately resolved, service desk procedures need to be established so that incidents are appropriately escalated and if appropriate workarounds are provided. When an incident has been resolved, the root cause, if identified, should be recorded by the service desk. This will aid in preventing future reoccurrence of the same (type of) incident(s). Incident reports generated in the "Manage service desk and incidents" (DS8) process form an input to the "Manage problems" process (DS10). Additionally, it should be confirmed that the action that has been taken to correct the incident is in line with the SLA(s) and does not introduce new security problems. Finally, in order to continually improve incident management, service desk performance and response time needs to be evaluated.

Manage problems

The "Manage problems" control objective of COBIT (DS10) deals with the management of problems that were identified as part of incident management (DS8). Processes should be implemented that facilitate the reporting and classification of problems, where the method of classification is similar to that of incidents. Problems should be categorised into related groups or domains (e.g. hardware and software), so they can be allocated to appropriate support staff.

Problem management should provide for audit trail services that allow tracking, analysing and determining the root cause of problems. Sustainable solutions to problems, which address the root cause, should be identified and initiated. The progress regarding problem resolution should be monitored against SLAs. When the impact of a problem becomes (more) severe, problem management should escalate the problem.

In addition, problems should be properly closed by putting in place a procedure to close problem records after confirmation of successful elimination of the problem (e.g. security incident) or after agreement on an alternative problem handling method.

To further ensure effective management of problems and incidents, COBIT states that the processes related with problem management should be integrated.[8] Organisations should monitor how much effort is put into resolving errors and issues rather than enabling business improvements and improve the aforementioned processes, where and if necessary, to minimise problems.

Summary and discussion

In essence, the service desk function of COBIT functions much like the Incident Control function of ITIL. They both form a central (security) incident liaison which provides registration, communication, dispatching and analyses of reported security incidents. COBIT also states the need for classification and prioritisation of reported incidents.

The "Manage problems" objective in COBIT focuses on recording, tracking and resolving problems; investigating the root cause of all significant problems; and defining solutions for identified problems. Additionally, COBIT states that, besides incidents, problems should be categorised as well.

The incident management process described in COBIT suffers from the same 'shortcomings' as does ITIL: the importance of classifying security incident is mentioned on more than one occasion, but how an organisation should realise such a classification is not discussed. This is due to the fact that concrete methods for doing the things prescribed in various best practices are intentionally left out, because this falls seemingly outside the scope of the reviewed best practices of ITIL and COBIT. Furthermore, COBIT also omits a graphical representation of the various processes, which would make process structures and dependencies easier to distinguish.

[8] Related processes are e.g. change management and configuration management.

3.2.3 ISF information security incident management

Besides the aforementioned incident management processes of ITIL and COBIT, there is also an incident management process described in a report by the ISF [ISIM06], the organisation that established The Standard of Good Practice for Information Security [ISFS05]. Recently, in April 2006, the ISF publicised a report containing the results of a study on information security incident management, which yielded a process for information security incident management.

The first three stages of that process, namely Identification, Response and Recovery, consist of steps specific to minimising the impact of and resolving an information security incident. The final stage, Post-incident review, involves follow-up activities which relate to the information security incident. It must be noted that the stages of the information security incident management process, and the steps contained within those stages, are not necessarily fully sequential.

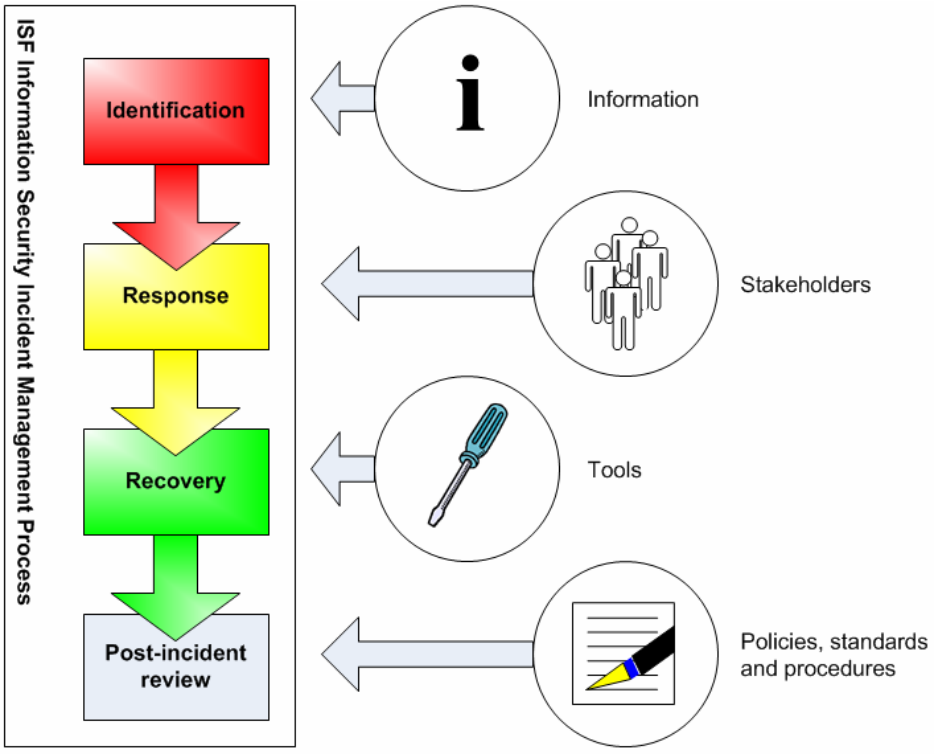

Figure 3.2 – The ISF information security incident management process (adapted from [ISM06])

Identification

The Identification stage is usually triggered following the detection of a possible information security incident. This trigger may be via automated means (e.g., an IDS) or by a human (e.g., an employee reporting odd application behaviour to a help desk).

The primary objective of the Identification stage is to determine whether an information security incident has occurred. A wide variety of information is gathered and assessed (often referred to as triage), which will typically result in the classification of the information security incident which indicates its type and severity.

The nature of an information security incident may require the careful handling of any associated information (e.g., the information may be considered potential evidence and, as a result, may be subject to computer forensic examination at a later time).

Undertaking the Identification stage is critical as proceeding to the Response stage will typically incur more and significant costs for the organisation. However, should the potential information security incident be considered genuine, the transition to the Response stage should take place as fast as possible to help minimise business impact and to assist in the resumption of normal business operations as soon as possible.

The report also revealed that there are various scenarios which, although not information security incidents as such, also require a response similar to that triggered by an information security incident. Such scenarios include detection of network and computer footprinting and scanning activity by malicious parties, attempted compromise of networks or computer systems and discovery of vulnerabilities that could be exploited by an attacker. Although these scenarios do not constitute an information security incident (as information has not been compromised), they still require similar activities to be performed, particularly those that are carried out during the Identification and Response stages.

Response

The Response stage is triggered following the positive identification of an information security incident. This stage can often be the busiest, involving various stakeholders within the organisation

and third parties required to help resolve the information security incident and minimise business impact.

The effectiveness of the steps performed during the Response stage is dependent on the information obtained about the information security incident during the Identification stage (e.g. type of information security incident, its scope and possible impact).

During the early part of the Response stage, action should be taken to quickly contain the information security incident. Containment is a critical activity and is a method of isolating the area affected by the information security incident to prevent further damage to the infrastructure and the organisation. Containment may include disconnection of computer systems from the network, shutting down servers or restricting access to applications.

Once the information security incident has been contained, the next objective is to eliminate the cause of the information security incident. Eliminating the cause of an information security incident depends on the circumstances but typically involves a combination of disabling the threat and fixing vulnerabilities exploited by an attacker.

Recovery

The start of the Recovery stage typically coincides with the end of the Response stage, although some steps may be started before the cause of the information security incident has been identified and addressed. The Recovery stage primarily involves the steps required to resume normal business operations. Additionally, the Recovery stage is likely to involve different people to those participating in the Response stage. These typically include the individuals responsible for maintaining the infrastructure who perform many of the required recovery tasks, such as regaining network connectivity and restoring data.

Once all computer systems, applications and data affected by the information security incident have been restored and normal business operations have resumed, the information security incident can be closed. There may be additional activities relating to the information security incident that need to be performed after it is closed. These typically take place during the Post-incident review stage.

Post-incident review

The Post-incident review stage follows the recovery from and closure of an information security incident. It can consist of a range of activities to support follow-up action(s), help the organisation understand more about the information security incident and identify areas for improvement in both the information security incident management process and the information security plans.

Performing many of the tasks associated with post-incident review (sometimes referred to as the post-mortem) provides an organisation with a valuable opportunity to learn lessons about the strengths and weaknesses of their information security arrangements and therefore, make necessary improvements.

Summary and discussion

The incident management process as described above is similar to those described in ITIL and COBIT: there should be a single point of contact, incidents should be categorised and prioritised, root causes of incidents should be investigated, etc. However, the process described above is based on a recent survey (i.e., very up-to-date) which was carried out for the specific purpose of creating a process which can be used for setting an information security incident management capability. This process is therefore probably easier to apply in practice, because it provides a clear cut process that an organisation can follow. But the same shortcoming of the other two best practices holds here: there is no clear method on how to design a classification of security incidents. Besides for ITIL and COBIT, a methodology for the classification of security incidents and associated indicators and correct (i.e., damage and impact minimising) incident responses can provide a practical and useful addition to the ISF information security incident management process as well.

3.2.4 ISO 17799

The incident management process of ISO17799 does not differ enough from those of the best practices discussed in this chapter to warrant a separate section to discuss its details. Moreover, the same shortcomings apply to the incident management process of ISO17799 as they did for the other best practices. It is therefore expected that a method for the classification of security incidents contributes to the incident management process of ISO17799 as well.

3.3 Chapter summary

In this chapter security incidents have been defined and various associated aspects, such as the nature of an incident, were briefly discussed. Furthermore, the incident management processes of ITIL, COBIT were examined, as well as the ISF information security incident management process. Several shortcomings were discovered. One recurring shortcoming was that no method for the classification of security incidents was given by these best practices, although they recognised the importance of such a classification for the incident management process. This is where the goal of this research comes into play: to design a methodology for improving information security incident management. Such a methodology can possibly provide a practical contribution to the incident management processes of the aforementioned best practices. The methodology will be drawn up in the next chapter, and will be based on a method for the classification of security incidents.

Chapter 4

CLASSIFICATION OF SECURITY INCIDENTS

Research question 3: How can information security incident identification and response be improved?

Research question 4: How does the methodology fit into current information security incident management processes?

In the previous chapter it became clear that classification of information security incidents is important for the information security incident management process. Such a classification creates structure in the collection of all possible information security incidents that an organisation may face at some point in time. The first step then is to design the method by which information security incidents can be classified. Due to the fact that the method of classification of information security incidents alone will probably have limited added value for an organisation's information security incident management process, the classification will also incorporate incident indicators and incident responses. This is to maximise possible added value and create more of a process. Incident indicators and incident responses may have added value for the identification and response stages of the information security incident management process.

4.1 Introduction

The (unstructured) collection (also called a space[1]) of all possible information security incidents is a subset of the collection of all possible security incidents, which in turn is a subset of all possible incidents. This is depicted in Figure 4.1. In this section a methodology for classifying information security incidents will be described. Classification is needed to attain a structured overview of the collection of possible information security incidents that an organisation may face at some point in time.

[1] In mathematics, a space is a set with some particular properties and usually some additional structure. http://en.wikipedia.org/wiki/Space.

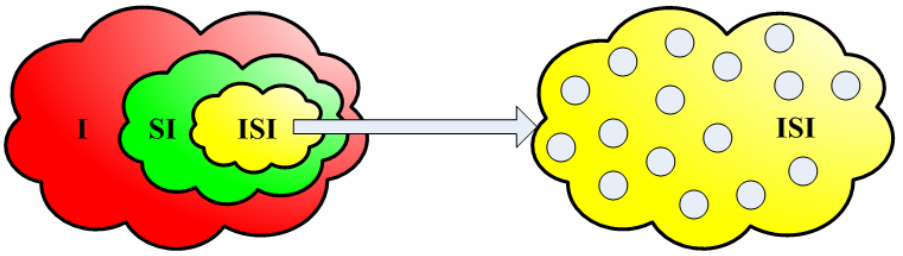

I = Incident Space

SI = Security Incident Space

ISI = Information Security Incident Space

○ = an Information Security Incident

Figure 4.1 – A representation of the information security incident space.

Classification can be done by various aspects related to information security incidents. In chapter 3 several of these aspects were described, e.g. it's impact on the business and it's type (e.g., hardware, software). These factors influence the response to such incidents. Since the classification needs to stay manageable, classification by too many factors will not do. Furthermore, the factors used to classify information security incidents need to be relevant for the purpose of being able to swiftly and accurately responding to those incidents. First, the contents of the response depends on the type of incident; secondly, the severity (i.e., impact on the business) determines the response priority (i.e., when will the incident be handled). Additionally, classifying by type will provide structure in the incident space (Figure 4.1). To this end, classification will be done by type and severity. In the subsequent sections the method for the classification of information security incidents will be described.

First, the typology used for classification by type will be discussed. Then, a method will be described for the determination of the severity of an information security incident. Then, the incident classification will be extended by linking to incident indicators and predefined responses.

4.2 Classification of information security incidents

4.2.1 Classifying incidents by type

Information security incidents can be classified by their type, following a specific typology (e.g., network, human, software). Classifying by type provides logical structuring of the information security incident space (see Figure 4.1). Additionally, the type of information security incident is necessary to determine the contents of the response to that incident, i.e., the response itself depends on the specific type of information security incident. Furthermore, incident classification by type provides users of that classification with a logical and easily understandable overview of incidents.
Based on the aforementioned arguments it is argued here that incidents should be classified by severity as well as by type, where the latter is used to give the resulting classification a logically structured format. To preserve this logical format when classifying, classification should first be done by type and then by severity.

When classifying by type, there are a variety of different classification methods available. All these methods organise information security incident into groups that have something in common. The first method classifies information security incidents by e.g. vulnerability, compromise, attack or exploit, and suspicious activity. A second way of classifying information security incidents, used in [DTIS04], is to group incidents into broad domains, e.g. virus infection and disruptive software, staff misuse of information systems, theft or fraud involving computers, and accidental incidents. A third method to classify information security incidents by type is provided by the ISF in their report on information security incident management [ISIM06], which divides information security incidents in the following categories: external attack, internal misuse and abuse, theft, system malfunction, service interruption, human error, and finally unforeseen effects of change.

For the classification of information security incidents by type, the ISF's categories will be used. The ISF categories are considered a best practice, resulting from discussions and meetings between its members (consisting of more than 270 leading companies and public sector organisations). A list of

examples of information security incidents that are ordered according to the ISF's typology can be found in Appendix C.

4.2.2 Classifying incidents by severity

Classification of information security incidents can further be done by severity, where severity is the overall negative impact of an incident on the business. This impact consists of various components:, among which are financial loss, loss of goodwill, loss of customers, etc.

The severity of an incident will be used for determining the handling priority of the incident. The four values for the severity of an incident are described below, with their corresponding handling priority:

S: Severe (i.e., very high); the incident may threaten business continuity if it occurs. Consequences of a very high severity incident could be: business processes are disturbed for a prolonged period of time, there is large damage to reputation, extensive loss of customers, etc. Handling priority: 1

H: High; the incident may have a high impact on the business. Consequences of a high severity incident could be: business processes are disturbed for some time, reputation could be damaged, there are big financial losses, etc. Handling priority: 2

M: Moderate; the incident may have a some impact on the business. It could be that business processes are disturbed for a short period of time, leading to some financial loss and there is perhaps some loss of goodwill and trust. Handling priority: 3

L: Low; the incident may have a low impact on the business. Losses are probably limited to costs made due to solving the incident. Handling priority: 4

Now the different handling priorities are discussed, which are directly linked to the severity of an incident (see above).

Handling priority: 1; These incidents need to be resolved immediately, so they are to be handled directly. Incidents with handling priority 1 have priority over incidents that have handling priority 2, 3 or 4.

Handling priority: 2; These incidents need to be resolved as quickly as possible, so they are to be handled directly. Incidents with handling priority 2 have priority over incidents that have handling priority 3 or 4.

Handling priority: 3; These incidents need to be solved within a short period of time to prevent damaging consequences and escalation to a higher category. Incidents with handling priority 3 have priority over incidents that have handling priority 4.

Handling priority: 4; These incidents are to be solved when there is time and manpower available.

4.2.2.1 Determining the organisational context

To put the severity of an information security incident into the proper perspective, it is necessary to determine the organisation specific context with regard to the need for information reliability (the three reliability attributes confidentiality, integrity and availability).

In his research, [STOF04] distinguished between a high *(H)*, moderate *(M)* and low *(L)* need for each of the three reliability aspects confidentiality *(C)*, integrity *(I)* and availability *(A)*, as described below.

H: High; High need for $C/I/A$. $C/I/A$ is crucial for business continuity.

M: Moderate; Moderate need for $C/I/A$. $C/I/A$ is moderately important for the business.

L: Low; Low need for $C/I/A$. $C/I/A$ is not very important for the business.

The reliability need of an organisation can be expressed in the form of a triple[2], i.e., reliability need of organisation A = (*C*,*I*,*A*). For instance, if organisation A has a high need for confidentiality and integrity and a medium need for availability, it's information reliability need can be described by the triple (*H*,*H*,*M*). The SPRINT method used to attain the values for each of the reliability attributes is described in Appendix B.1.

[2] A triple is a term from mathematics, depicting an n-tuple with n being 3.This means it depicts three elements, with the order of the elements carrying meaning (unlike a set, where the order of the elements is meaningless and where no element could exist more than once). http://en.wikipedia.org/wiki/Triple.

4.2.2.2 Determining the risk

In this section a method for determining the risk that incidents pose for reliability *(C,I,A)* is described. A SPRINT-based method can also be used to determine the risks of an incident to the reliability (i.e., confidentiality, integrity and availability) of information. The risk is defined here as the probability that a certain situation will happen in case of a particular incident. Determining the risk (i.e., probability) that certain situations may occur needs to be done by experts within the target organisation that have extensive knowledge of the information infrastructure, information need and state of information security of the organisation.

The answers to the questions of the method can be translated into the risks of an incident to reliability. The risk of each of the three attributes is the maximum value of risk that is determined for each particular attribute; e.g., if only one question yields a high *(H)* risk, the final attribute value will still also be high *(H)*. The final risk to reliability is described here by a triple in the form *(C,I,A)*, much like the reliability need discussed previously, where each component can take the value's High *(H)*, Moderate *(M)* and Low *(L)*. The values are presented below.

H: High risk, i.e., it is very likely that the incident will result in the presented situation.
M: Moderate risk, i.e., it is probable that the incident will result in the presented situation.
L: Low risk, i.e., it is very unlikely that the incident will result in the presented situation.

The SPRINT method for determining the risks to reliability of an incident is described in Appendix B.2.

4.2.2.3 Determining the severity

The combination of incident risk and the reliability need of the organisation can then assist in determining the overall severity of an information security incident (see Figure 4.2 on the following page), expressed as being severe (i.e., very high), high, moderate or low (**S, H, M** or **L**, see the

beginning of section 4.2.2*)*. Various methods for determining the actual severity will now be discussed, and one of these will be used.

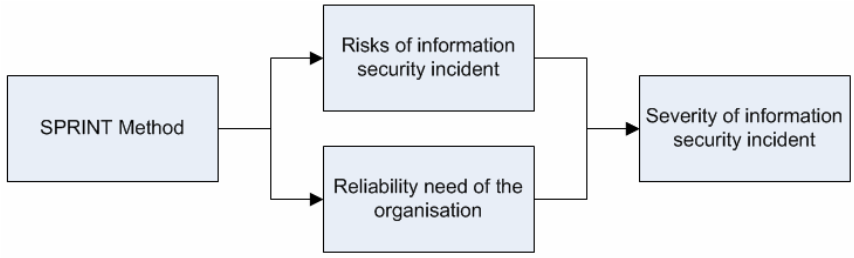

Figure 4.2 – Determining the severity of an information security incident

First, we could define the severity of an information security incident as the highest value of the triple representing an incident's risk to the three reliability aspects (i.e., (*C,I,A*)). This approach is described with the use of the following example. Say a fire has started in the main server room (e.g., due to short circuiting), which has already damaged one or more server-racks and is threatening to damage several more. This may have a significant impact on the availability of information and information processing, thus *A* = *High*. Confidentiality and integrity are most likely not at risk, so *C* = *Low* and *I* = *Low*. The maximum value of (*H,L,L*) is *High*, so the severity rating of the information security incident is marked *High* as well. This method does not take into consideration the reliability need of the organisation, which was said to determine the context in which an incident takes place. Thus, this method is deemed unsuitable.

It is also possible to determine the severity rating by calculating a sort of 'average' of the values of the triples. In the case of information security incident A, this would lead to something like a moderate severity rating: [AVG(*(M,L,L) (H,H,M)*)] = [AVG*(MH,M,LM)*] = moderate *(M)* severity. Following this approach, this may not seem to be such a severe incident. Nevertheless, remember that organisation A has a high need for confidentiality. This means that confidentiality is critical for business continuity. From that point of view, an information security incident that has medium impact on confidentiality

may be seen as severe due to the high need for confidentiality of organisation A. So, this second method is not suitable either.

This leads to a third approach, which can be generalised as follows: the average values of the attributes of the reliability need of an organisation and the possible risks of an incident to reliability, both specified by a triple *(C,I,A)*, result in a triple *(C,I,A)*.
In addition, the reliability need of an organisation takes precedence over the impact of an incident when there is no intermediate value between the two (i.e., *high* need and *low* risk means *moderate* severity, while *high* need and *moderate* risk means *high* severity, and also *moderate* need and *high* risk means *moderate* severity). This approach has the advantage that the severity (i.e., impact on the business) of an incident is not underestimated. Furthermore, it underpins the importance of an organisation's information reliability need as the context in which incidents take place.
Furthermore, an attribute that has a high value in both the reliability need and the incident risk, scores a very high (S) in the resulting severity triple. A few examples are described below.

RN = Reliability need; IR = Incident risk

RN = **H**, IR = **H** ➔ Severity value= **S**

RN = **H**, IR = **M** ➔ Severity value= **H**

RN = **M**, IR = **H** ➔ Severity value= **M**

RN = **L**, IR = **M** ➔ Severity value= **L**

Then the highest value of the resulting severity triple, which is thus based on the triples of the reliability need and the incident risks, determines the severity rating. This is because only one high (**H**) value is needed for the severity of the incident as a whole to be deemed high (**H**). Now a few examples of determining the severity using this approach are presented.

RN (C,I,A) = Reliability need; IR (C,I,A) = Incident risks to reliability

1) RN *(H,M,L)*, IR *(L,L,M)* ➔ Severity = Max *(M,M,L)* = **M**

2) RN *(H,M,L)*, IR *(M,H,L)* ➔ Severity = Max *(H,M,L)* = **H**

3) RN *(L,H,H)*, IR *(H,L,H)* ➔ Severity = Max *(M,M,S)* = **S**

In example number 3 we see that the organisation has a high need for availability and that the incident has a high impact on that availability. The high need for availability symbolises that it is crucial for the business, so when an incident then has a high impact on availability it could pose a threat to business continuity. This justifies the 'severe' (**S**) rating appointed to the incident.

4.2.2.4 Critical reflection

The benefit of the method for determining the severity of incident that was expounded here is that it takes both the context in which an incident takes place and its risk to the reliability of information into consideration. This leads to a severity rating that is tailored to the specific incident and the specific organisation and thus to a proper incident handling priority.

The disadvantage of the method is that it is rather theoretical and time-consuming. This leads to the question whether it is feasible to perform a severity determination with the expounded method. Expert feedback may offer further insight into the feasibility of the method.

Classifying information security incidents by severity can serve as a means for prioritisation of information security incident handling, where higher severity incidents are given priority over lower severity incidents. This proper prioritisation of information security incidents contributes to a more efficient response through quicker response to critical incidents.

Besides the obvious benefit for incident prioritisation, a severity rating gives an organisation an explicit notice on the possible consequences an incident could have on business continuity (when not handled appropriately and in a timely fashion).

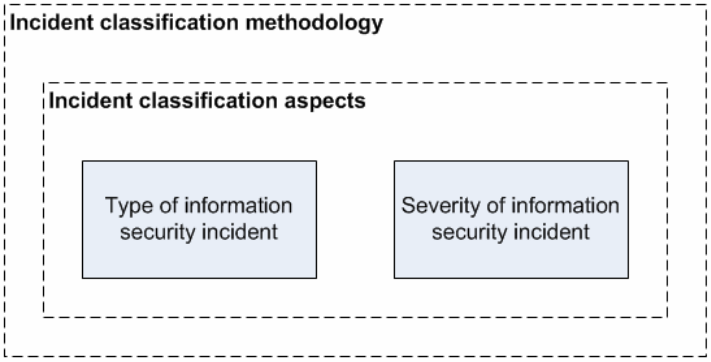

Figure 4.3 – The methodology including the two classification aspects type and severity.

4.3 Extending the classification

In the previous section a method for the classification of information security incidents was discussed. Information security incidents were classified based on their severity (possible negative impact on the business) and their incident category. The severity rating can be used to prioritise information security incidents and at the same time provides the organisation with an estimation of possible consequences for the business if an incident is not handled properly and in a timely fashion. The classification of incidents by incident category provides a means for logically structuring the list of possible information security incidents.

To have added value for the identification and response stages of the information security incident management process, the classification will be expanded by linking to possible indicators for incidents and appropriate responses to incidents. These indicators and responses are discussed in the following sections.

4.3.1 Adding indicators

Indicators for security incidents are events and/or activities that point towards a security incident that has taken place, is taking place or will take place. These indicators may help an organisation's help desk function (or other incident liaison) to identify security incidents more swiftly and more accurately, consequently leading to an improved response to an incident. This may thus provide added value for the identification stage of the information security incident management process.

Indicators can thus be used to perform a triage[3] process, to determine the type of incident and the necessary response. Triage is used on a daily basis by general practitioners to find out what kind of illness people have based on the symptoms they exhibit or experience, and based on the illness determined from the symptoms a treatment is devised. This process can also be applied to information security incidents. In this research the incident indicators are the symptoms, the information security incidents are the illnesses and the incident response is the treatment. Say, for instance, that users have reported symptoms (i.e., indicators) A and B. These symptoms point to information security incident C. In turn, information security incident C necessitates incident response D.

It may be that an information security incident itself is an indicator for another information security incident. In this case, linking incident indicators to the possible information security incidents, also adds value to the classification. The added value in this case is that the help desk (or other incident liaison) doesn't overlook the possibility that an information security incident is part of another, perhaps larger, incident.

It may not be abundantly clear what the indicators for a particular security incident are, or if an incident is in itself an indicator for another incident. To understand the linkage between indicators and incidents and what possible indicators for a specific incident are, requires a certain amount of expert knowledge.

The added value of linking indicators to the classified list of incidents will be illustrated by some examples. As the first example we choose an infection of computer systems with a computer worm[4]. If a worm infects a computer in an organisation's network, this may lead to an outbreak of that worm in

[3] Triage is a process for sorting injured people into groups based on their need for or likely benefit from immediate medical treatment. Triage is used in hospital emergency rooms, on battlefields, and at disaster sites when limited medical resources need to be allocated.

[4] A worm is a self-sufficient, self-replicating computer program. Worms are often designed to exploit the file transmission capabilities found on many computers. A worm uses a network to send copies of itself to other systems, thereby harming networks and consuming bandwidth.

the organisation's network, thereby harming the network and consuming bandwidth. There are multiple possible indicators for a worm infection. For example, a worm usually scans computers in the network on a certain port, e.g. to find out certain vulnerabilities. Usually, if the worm infects another computer, it installs a backdoor listener, through which it tries to find out if the infection was successful. If a worm outbreak occurs, an organisation can face total network service disruption, due to the worm clogging up the network bandwidth. However, if port scanning is linked as an indicator to the incident 'infection of systems by a computer worm', the worm may be stopped before it disrupts entire systems. This may significantly lower the possible negative consequences.
In another example, suppose that users find that their computers are running slower than usual; programs take longer to load, documents take forever to open or save, etc. Although this might just be a matter of e.g. hard-disk fragmentation, it could also be a symptom of a more serious cause: a virus infection. So, when users report this computer sluggishness to their help desk, and computer sluggishness is mentioned in the security incident classification as a possible indicator for a computer virus infection, the help desk will not overlook the possibility that it is indeed a virus infection. In that case, they are able to rapidly take measures to contain a possible outbreak.

The abovementioned examples, although brief, demonstrate that incident indicators may provide added value to the identification stage of the information security incident management process; incidents may be identified more effectively when indicators are linked to incidents. A help desk can check if certain activities, events and/or incidents that are reported, are possible indicators for an (other) information security incident. For instance, activity A and event B may point out that a security incident C has taken place, or security incident D may be an indicator for (major) security incident E.

Setting up a method for gathering and documenting incident indicators would justify a separate study and is an item for possible future research. For this research, adding indicators based on expert knowledge and details of previous information security incidents will have to suffice, and even though it is not an optimal solution it may still have added value for the identification part of incident management.

A brief example of how these indicators could be used will now be discussed. A possibility for using the indicators is in a decision tree. Say you have an indicator A. Now, is there an indicator B, or C? Say indicator B is also there. Now, we have another decision to make: is there an indicator D, E or F? Say, indicator E applies to the current situation. Taking the path A➔B➔E leads us to incident X, which is of ISF type Y and has severity Z. This process is visualised in Figure 4.4 below.

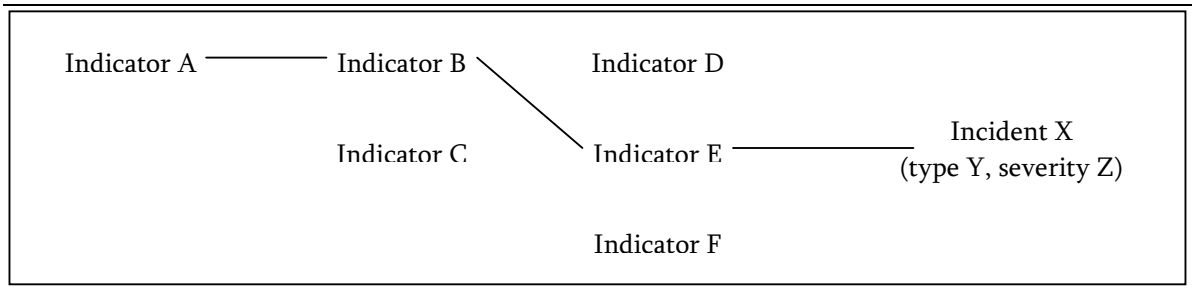

Figure 4.4 – Using the indicators in a decision tree.

The most useful application of this approach would be to automate the decision process as follows. An employee at the help desk inputs the first indicator, at which point he is presented by the system with a set of indicators and questioned which of them applies to the situation (i.e., possible incident). This ultimately leads to system displaying the incident (e.g., type and severity).

The aforementioned application is just an example of how indicators could be used to improve the identification of security incidents. The collecting, documenting and use of indicators is an interesting topic for future research.

4.3.2 Adding responses

It has already been discussed that not having documented the proper response to an incident in advance may lead to unnecessary damages to the organisation (see section 1.2). When organisations do define adequate responses to security incidents in advance, there is a greater likelihood that security

incidents are resolved properly and in a timely fashion, resulting in less damages from those security incidents. Proper incident responses help contain an incident and minimise an incident's possible negative consequences, such as loss of revenue and reputation damage.

In documenting predefined proper incidents responses, organisations need to keep various matters in the back of their heads. Here, five types of responses are distinguished.
First, for certain types of security incidents, there are mandatory (legal) responses. For example, in most industries organisations need to report theft of privacy sensitive data to the right authorities and legal bodies. So, legal responses are a part of the predefined incident response plan.
Secondly, still considering the case of theft of privacy sensitive data (in this case, customer data), customers need to be informed as well. Therefore, organisations need to define in advance how, and if, they will inform their customers and the general public of a particular security incident. It is not intended that organisations inform the outside world of all information security incidents, but at least a certain amount of openness towards any stakeholders involved should be taken into account.
Third, the information security incident may need to be reported and communicated to employees, e.g., warning employees of an e-mail containing a virus or other malicious code.
Fourth, there may be a need to communicate the incident to the media to inform the public and stakeholders that can't be directly informed (see the second response type).
Finally, the most important response, as far as resolving an information security incident goes, is the containment response: what actions need to be undertaken to contain the information security incident? This response needs to be predefined as extensively as possible in order to make swift containment of the incident possible.
Summarising, the 5 response types are legal response, employee response, stakeholder response, media response and containment response.

In the end, not all responses can be predefined, but there are basic responses that have to, and can, be carried out independent of the incident specifics (e.g., legal responses). These responses can be predefined and can help swift containment and resolution of information security incidents.

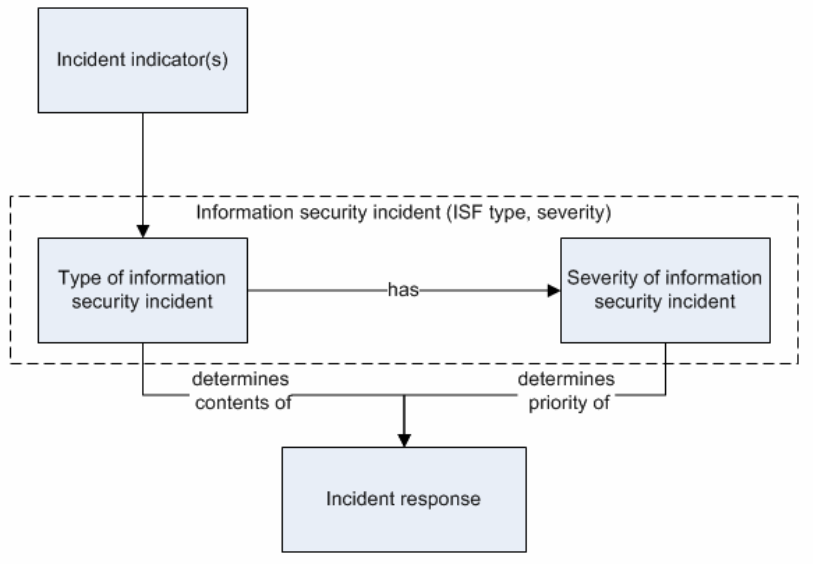

Figure 4.5 – The methodology with all four components.

Figure 4.5 shows a global overview of the various components of the methodology, namely the indicators that are used for identifying incidents, incidents classified by ISF type (determines the contents of the incident response) and by severity (determines the incident handling priority), and the proper response(s) linked to incidents.

Figure 4.6, depicted on the following page, shows a more detailed view of the interdependencies of the various components of the methodology, including supporting components such as SPRINT and the ISF typology and the placement of triage in the methodology. Figure 4.6 should not be seen as a detailed process overview, but as an overview of the various components.

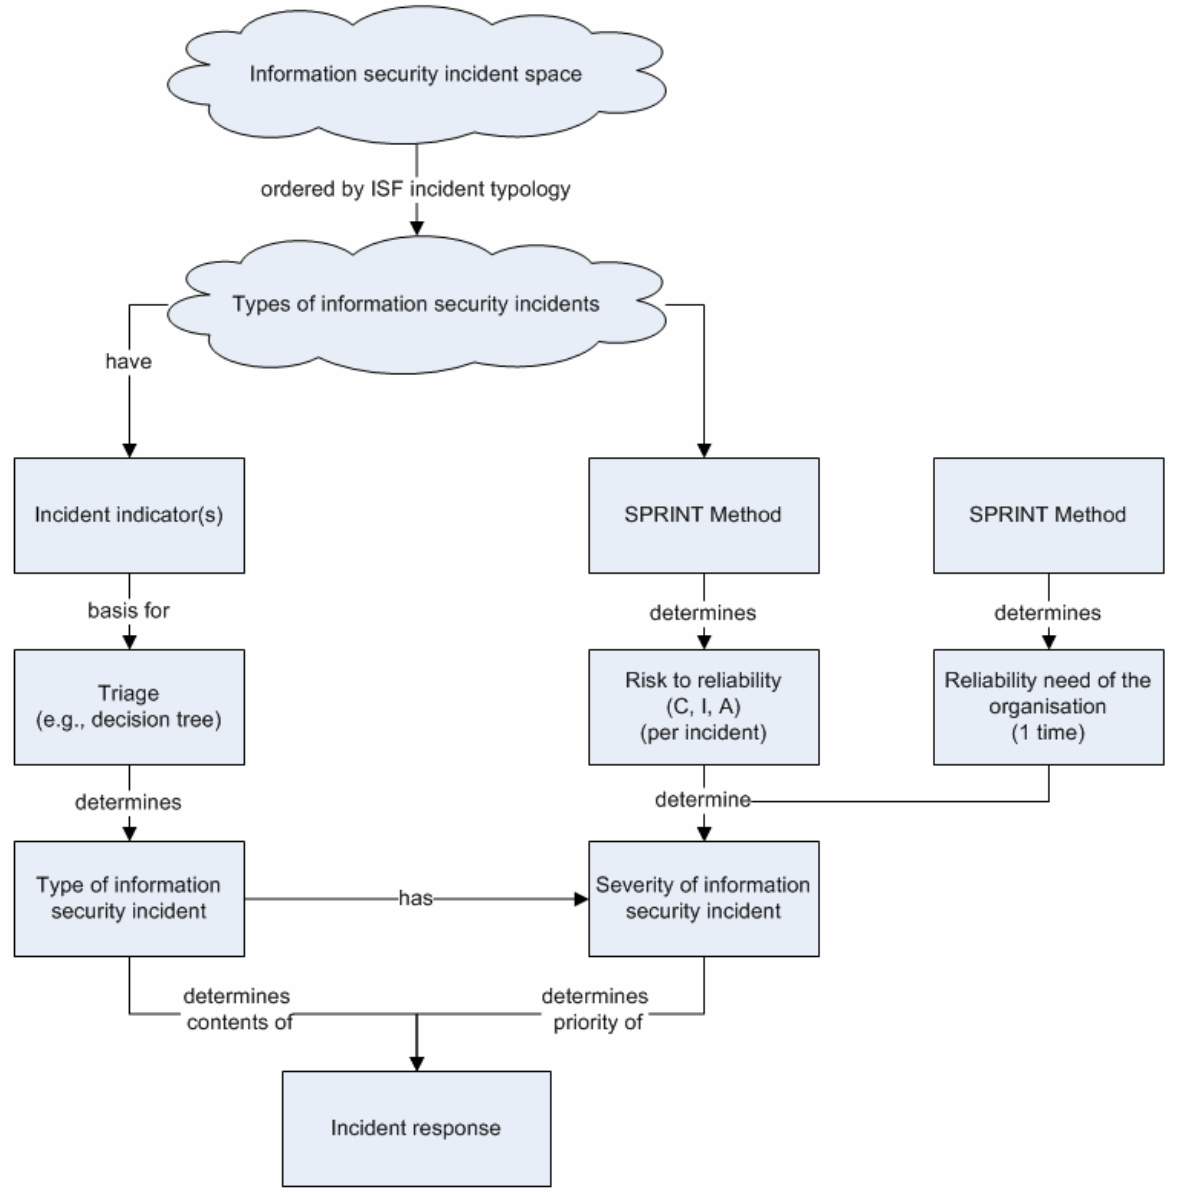

Figure 4.6 - Detailed overview of the classification methodology.

4.4 Placing the methodology within current best practices

The best practice methods that were discussed in chapter 3 mention the significance of classification of information security incidents. Therefore, a methodology for setting up an information security incident classification was set up in this chapter. It provides organisations with a clear, usable method of how to classify information security incidents, in this case by type and severity. The added value of the information security incident classification methodology that was laid out in the previous sections lies in the incident indicators and predefined incident responses linked to the list of possible information security incidents, which are classified by type and severity. This may aid in the identification and response stages of an organisation's information security incident management process (see section 3.2.3).

In essence, the methodology described in this chapter does not alter in any way the incident management processes of the best practices described in chapter 3. The methodology should be seen as an addition to those best practices, giving organisations a handle for the classification of information security incidents, combined with incident indicators and predefined incident responses. It is expected that the methodology supports information security incident management processes by providing improved identification of, and response to, information security incidents.

4.5 Chapter summary

In this chapter a method for classifying security incidents, linked to possible indicators and proper responses, was laid down. The output of the methodology is a classification which links security incidents, severity ratings of those incidents, potential indicators for security incidents and proper predefined responses for security incidents together. Additionally, it was explained how the methodology can be fitted in, or serve as an addition to, current best practices regarding information security incident management processes. In the following chapter we will describe a possible approach for the application of the methodology in a practical setting.

Chapter 5

APPLICATION OF THE METHODOLOGY

Research question 5. How can the information security incident classification methodology be applied in practice?

In this chapter we will describe a possible approach for the application of the methodology that was set up in the previous chapter; the approach will consist of possible steps for implementation and use. Not only a process for implementing and using the classification methodology is described, issues regarding the documenting (i.e., recording) and updating of the classification are also examined. In addition, an example process of using the (implemented) methodology is described.

5.1 Implementing the methodology

In this section a process is described which aids an organisation in setting up and implementing a classification methodology described in the previous chapter. The fact that such a classification needs to be 'implemented' before it can be used is dependent on two factors. The first is that the incident indicators need to be linked to information security incidents, so that the aforementioned process of triage can take place to determine the information security incident, it's type (for response purposes) and it's severity (for prioritisation purposes). Secondly, necessary responses for information security incidents need to be pre-defined to be more effective than ad-hoc determined responses. These responses are not meant to be exhaustive and covering all possible facets of an information security incident. Rather, these predefined responses are meant to enable an organisation to quickly contain the incident, provide proper obligatory responses (e.g., legal) and ultimately minimise damages resulting from the incident. The process needed to realise a classification of information security incidents and related indicators and responses is described below.

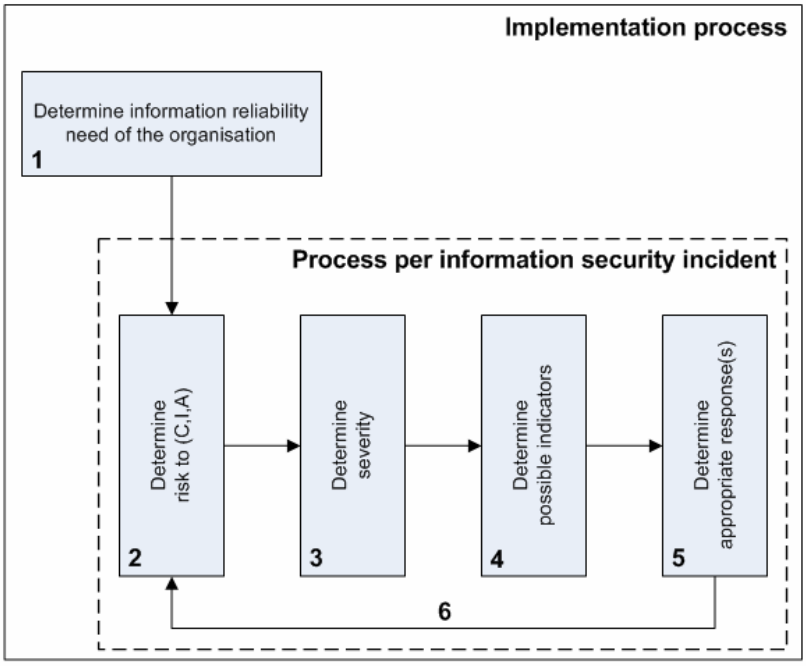

Figure 5.1 – Process for the implementation of the classification methodology

The steps of the implementation process are as follows:

1. Determine the information reliability need of the organisation *(C,I,A)* using the SPRINT method outlined in Appendix B.1

Per information security incident from the ISF typology:

2. Determine risks of the information security incident to *(C,I,A)* using the SPRINT method outlined in Appendix B.2
3. Determine severity of the information security incident *(S, H, M, L)* based on the reliability need and the risk to *(C,I,A)* of the incident.
4. Determine possible indicators (events/activities) for the information security incident
5. Determine appropriate response(s) for the information security incident
6. Repeat step 2 to 5 for other information security incidents

Step 1: Determine the information reliability need of the organisation (C,I,A)

Determining the information reliability need of the organisation will provide an organisation with the context in which information security incidents take place. The information reliability need can be attained through the adapted SPRINT method by [STOF04], a questionnaire on the impact of various reliability issues (regarding CIA) on the business. The SPRINT method can be found in Appendix B, at the end of this thesis.

Step 2: Determine risks of the information security incident to (C,I,A)

Determining the risks of an information security incident on the various information reliability aspects should be done by information security experts which have extensive knowledge of the organisation's (technical) information infrastructure. Risks may depend on such factors as e.g., the hard- and software that is used and the network infrastructure of the organisation. The risks of each aspect can be high *(H)*, moderate *(M)* or low *(L)*.

Step 3: Determine severity of the information security incident (S, H, M, L)

The severity of an incident is determined by weighing the reliability need of the organisation and the risks associated with an information security incident. The method for this weighing of factors was explained in section 4.1.2. The severity of an incident may be deemed severe *(S)*, high *(H)*, moderate *(M)* or low *(L)*.

Step 4: Determine possible indicators (events/activities) for the information security incident

Indicators for security incidents are events and/or activities that point towards an information security incident, as was discussed in chapter 4. These indicators may help an organisation's help desk function (or other incident liaison) to identify security incidents more swiftly and more accurately through performing triage to determine the specific type of incident.

The determination of indicators should be done by experts, based on their knowledge of information security incidents in general and reports and documentation of previous incidents (as was argued in chapter 4). This step could be improved by a method for systematically collecting data on possible indicators for information security incidents (see also section 8.2 on future research).

Step 5: Determine appropriate response(s) for the information security incident

Determining appropriate responses to an information security incident beforehand ensures that the proper responses are carried out quickly after an incident has been identified. These responses are not meant to be exhaustive and covering all possible facets of an information security incident. Rather, these predefined responses are meant to enable an organisation to quickly contain the incident, provide proper obligatory responses (e.g., inform legal bodies) and ultimately minimise damages resulting from the incident.

The following types of responses that need to be determined for the specific information security incident in this step of the process, are: legal response, employee response, stakeholder response, media response and containment response (see section 4.3.2). It may be that a few of them are unnecessary for some incidents, e.g., not all incidents require a legal response. The containment response always needs to be defined as extensively as is possible when determining predefined responses, in order to contain the incident as quickly as possible.

The predefined response should preferably also contain a specific incident response team (IRT), that will carry out the response. Such an IRT should consist of people with the proper knowledge and skills to handle the incident (see Appendix A for more info on IRTs). Further specifics of assigning IRTs will not be provided here because it doesn't fall within the scope of this research.

Step 6: Repeat step 2 to 5 for other information security incidents

For all other information security incidents, e.g., those in the ISF list (see Appendix C), repeat the steps 2 to 5.

The result of this process is an overview in which incidents are classified (i.e., ordered) by type and severity, and linked to possible indicators (for triage) and proper predefined responses.

5.2 Documenting the classification

The classification that is obtained by an organisation that has followed the steps outlined in the previous section, needs to be documented in some way. There are various possibilities for the documentation of the classification. It can be documented electronically and on paper. Having at least one digital and one paper copy available, will ensure a greater certainty of the classification being available when it is needed, as opposed to having it documented only in one way.

It would also be advisable to communicate the (existence of) the security incident classification to all employees whom it concerns. The classification could be made a part of a security handbook that is already present in many organisations. It might also be necessary to put in place some policy and procedures for the use of the security incident classification. If there is no policy or procedure in place that mandates the use of the classification, it may happen that it not used when a security incident occurs, thus possibly leading to unnecessary negative consequences arising from such a security incident.

The classification (i.e., a classified list of information security incidents) needs to be present, physically and/or electronically, at an organisation's incident liaison, e.g. a help desk function, who will be the primary user of the classification. When a certain security related activity or event is reported to the incident liaison, it is checked out through the security incident classification and if needed, the appropriate procedures are put into motion to contain the incident and minimise any negative consequences.

It may also be possible to implement a classified incident list in such a way that an automated system can make use of it. This would make it possible for some types of security incidents to be automatically identified and prioritised, and that the proper incident response procedure is set in motion by the system (e.g., allocating proper resources and notifying the ones responsible for acting out the particular response). Such a system should play a supportive role, aiding the help desk function with security incidents that are hard to recognise for people, but not for systems (e.g., the scanning of ports by a worm).

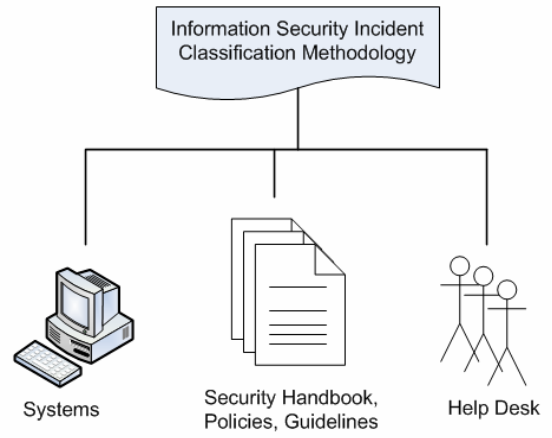

Figure 5.2 – Making the classified incident list available through various channels

5.3 Updating the classification

It is a known fact that the world of information technology and information security changes at a rapid pace. Technological innovations are commonplace and organisations that are implementing them are faced with the possible threats they pose to information security. Since these security threats can materialise into security incidents, it may be necessary to update the classification when a new technology is implemented by the organisation.

Keeping the classification updated is critical in maintaining its value to information security and the organisation; if it is not up-to-date, there may be security incidents for which there are no possible indicators known (swift and proper identification) and no predefined responses available (swift and appropriate incident response). This may lead to unnecessary damages resulting from those incidents.

5.4 Using the methodology

In this section an example process of how to use the classification methodology in case of an information security incident will be described. Practical usage of the classification in such a case

requires a structured approach. Such an approach is described here in the form of a process consisting of a number of steps.

A possible process from identification to response in case of a possible information security incident:

1. Suspicious event/activity reported
2. Perform triage to determine incident type and severity (i.e., identify and classify the incident)
3. Determine incident handling priority (based on severity)
4. Act out predefined response(s) when incident is being handled

The process is depicted in Figure 5.3 below.

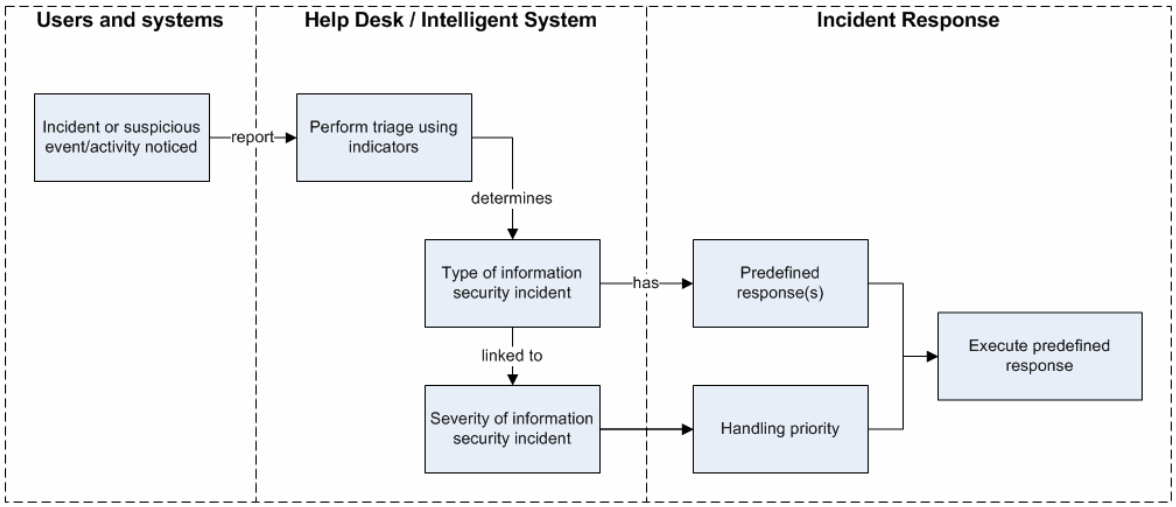

Figure 5.3 – From identification to response using the classification methodology

Step 1: Suspicious event/activity reported

When a suspicious event/activity is reported to, or discovered by, the organisation's help desk (or other incident liaison, like an automated system), there may be reason to determine if the situation

constitutes an information security incident. A reported suspicious activity/event is usually the starting signal for the information security incident management process.

Step 2: Perform triage to determine incident type and severity

Triage is performed to find out whether or not an information security incident has taken place, is taking place, or will take place. The outcome of the triage is whether or not it is an incident, and if it is an incident, it's type and severity are known (due to the classified list resulting from the implementation of the methodology, see section 5.1).

Step 3: Determine incident handling priority (based on severity)

The handling priority is based on the severity of the incident. It may be necessary to handle the incident immediately, due to the incident having a "Severe" (i.e., *(S)*) rating.

Step 4: Act out predefined response(s) when incident is being handled

When the response is to be carried out depends on the handling priority (see previous step). The actual predefined response is linked to the type of incident, which was determined by triage (see step 2). So, when the incident is being handled, the proper predefined response(s) can be carried out swiftly. Again, not all responses to an incident can be predefined, but there are basic responses that can and have to be carried out (e.g., legal responses) independent of the incident specifics. These responses can be predefined and can help swift resolution of incidents.

As a reminder, the abovementioned process is just an example of how the classification methodology could be used in case of an actual information security incident, it is not stated as the best or only way.

5.5 Chapter Summary

In this chapter we described a possible approach for the application and use of the methodology that was set up in the previous chapter. Not only a process for implementing the classification

methodology was described, issues regarding the documenting (i.e., recording) and updating of the classification were also examined. Furthermore, an example process of using the (implemented) methodology was discussed.

In the next chapter, information security experts will provide feedback on the methodology as it is, it's practical value, it's limits, etc. This feedback on the methodology may be used to adjust the methodology where possible and to uncover potential shortcomings of (the use of) the methodology. The feedback will also provide a basis for the specification of possible future research in the subject.

Chapter 6

FEEDBACK ON THE METHODOLOGY

Research question 6: Can the information security incident classification methodology improve security incident identification and response in practice?

Information security experts will provide feedback on the methodology's practical value, it's limits, etc. This feedback on the methodology will be used to uncover potential benefits and shortcomings of (the use of) the methodology. The test will also result in a value proposition for the classification methodology, which will concisely summarise the pro's and con's of the methodology. Potential shortcomings provide a basis for the specification of possible future research in the subject.

6.1 Test approach

The method cannot be applied so easily in practice, because it is more focused on design rather than implementation. Furthermore, to really test the method in practice, an actual information security incident would need to take place. So, another approach for testing is needed: experts will be asked to give feedback on the (various parts of) the methodology. Such an expert evaluation of possible benefits and drawbacks of the methodology needs to done by one or more information security experts of an actual organisation. These experts most likely have a practical view on the subject matter, i.e., the classification methodology, and can thus see the possible added value, possible benefits and the shortcomings of the methodology. The feedback achieved through such an evaluation can be used to uncover potential benefits and shortcomings of the methodology and provide valuable information for possible future research on the subject of information security incident management.

6.1.1 Qualitative feedback

With regard to attaining feedback there are two possible approaches: a quantitative and a qualitative approach. For the purpose of attaining feedback on the added value and feasibility of the classification methodology, and it's components, a qualitative approach is preferred. A qualitative approach places emphasis on understanding through looking closely at people's words, actions and records. A quantitative approach looks past these words, actions and records to their mathematical significance

and quantifies the results of these observations. In contrast, a qualitative approach is based on examining the patterns of meaning which emerge from the data and these are often presented in the participants' (i.e., the experts that will be evaluating the classification methodology) own words [SCDC].

Summarising, a qualitative approach for attaining the feedback needed to assess the added value and feasibility of the classification methodology in practice. Experts will be able to voice their opinion of the methodology in their own words, based on their own experience and knowledge regarding the subject matter.

6.1.2 Feedback format

Qualitative feedback can be obtained from various formats, among which are one-on-one interviews, focus groups and open-ended surveys. Each format has its own advantages and disadvantages, but all should yield useful end results if used properly.

One-on-one interviews are typically very in-depth and can be conducted anonymously to gather insights into a specific participant's mind. Due to the anonymity, participants will most likely voice their true opinions on and insights into the subject matter, resulting in very useful feedback.

Focus groups are an intense research method that can yield extremely interesting feedback. Focus groups (target groups that discuss a given topic together for some amount of time) provide an opportunity for like-minded individuals to interact with each other in order to draw out deeper issues regarding the subject matter. Focus groups can be hard to coordinate and moderate. Furthermore, it can be very hard to get experts together at the same place, at the same time due to e.g., their geographical dispersion and varying time schedules.

Surveys can be done to get a sense of people's perception of the subject matter. Specific questions can be asked about the various parts of the subject matter. The disadvantages of surveys are that there is no reciprocity between the surveyor and the participants and answers may be "led" by previous questions, leading to more superficial answers to the questions and thus more superficial feedback.

Based on the aforementioned, one-on-one interviews will be used to acquire the needed feedback on the classification methodology. Using one-on-one interviews leaves room for discussion on the subject matter and participants will remain anonymous.

6.2 Feedback on the methodology

It is to be expected that the classification methodology will receive mixed feedback, i.e., both positive and negative. This is due to the fact that these experts may have differing opinions regarding (the components of) the methodology. They have practical security knowledge which influences how they believe information security works. Different experts may see different benefits and drawbacks of the methodology.

For the purpose of evaluation, the methodology will be broken down into three components: 1) the classification of information security incident by type and severity, 2) linking possible indicators to information security incidents for the purpose of performing triage and 3) linking predefined responses to information security incidents. Each of these will be evaluated separately. Furthermore, the classification methodology as a whole needs to be evaluated, i.e., from identification to classification to response. Furthermore, experts will be asked if the method described in this research fits the research objective of improving information security incident identification and response.
The feedback received from experts will be described in the following sections. Feedback on the various components of the methodology will be presented.

6.2.1 Research objective

The main research objective was to improve information security incident identification and response. This was achieved by setting up a methodology for the classification of information security incidents and then linking these to predefined incident responses (improved response) and specific incident indicators (improved identification).

Experts agreed that predefined responses, to some extent, are necessary for a good information security incident management process. They concur that predefined responses facilitate a good foundation for swift and proper incident resolution, which minimises possible damages resulting from incidents.

6.2.2 Classification of information security incidents

Concerning the classification of information security incidents by type and severity, experts agree that these two are the minimal and crucial factors that are needed for properly responding to those incidents. These two aspects make it possible to prioritise the incident (severity) and determine the proper response (type).

A comment that was made by one expert, is that the aspect of chance could be incorporated as a third aspect (chance is the probability that an incident will occur). This chance plays a major role in determining the security measures that an organisation will implement to counter security threats. Since implementing such measures usually involves costs of some sort, these costs need to be evaluated with the estimated benefits of implementing such measures. Furthermore, these measures are usually necessary to implement because if they are not, some threats will exert themselves with a near certainty.

In the case of the methodology of this research, costs are limited to documenting how an incident can be identified and what to do in case such an incident occurs. Incorporating chance will additionally undermine the principle of the methodology, i.e., knowing what to do in case an incident does occur. For example, the risk of fire is very low in most modern office buildings, yet (nearly) all of them have extensive response plans in case such a fire occurs. The cost of such a predefined plan is relatively small with respect to the benefits it can provide in case of an actual fire, e.g., the lives of people that are saved by adhering to the plan.

In the classification methodology, the severity of an incident was based on the reliability need of an organisation and the specific risks to availability, integrity and confidentiality of an incident, both

determined by an adapted SPRINT method (see appendix B). While this did cover the context in which incidents take place and the specific risks of a particular incident, evaluation with experts revealed that this method was rather complicated and that it may be possible to merge the organisational context and risk determination. This approach will be further elaborated in section 6.3. Another comment regarding the use of the SPRINT method was that SPRINT is expert-driven: expert input is essential for it's use and application and it can't be done by just anyone. However, this can be interpreted as being both positive and negative depending on how one looks at it.

Moreover, regarding the severity of an incident, some doubts were raised as to how rigid the severity should be, i.e., severity may need to be adjusted when the incident occurs. For example, a computer virus outbreak may be labelled as having a High severity, but when a virus does break out it does not necessarily mean that it is very severe. It may be that only a few computers are infected, so that the severity is only moderate. This leads to two possible approaches to the matter: 1) adjust the severity, if needed, at the moment the incident takes place, and 2) define the various possible severity levels for an incident as separate incidents with their own set of predefined responses. These two possibilities will be further examined in section 6.3.

6.2.3 Indicators and triage

Regarding the aspect of linking specific indicators to information security incidents and performing triage with these indicators to determine the type of information security incident, experts agreed that this provided some new insights. As is the case with most information security incidents, it is most likely that not the incident itself is specifically reported to the incident liaison (e.g., the help desk) but rather a symptom, or indicator, for an information security incident. In various automated systems, such as IDSs[1], these indicators are already commonplace. In an IDS for example, events are monitored

[1] IDS means Intrusion Detection System. An Intrusion Detection System (or IDS) generally detects unwanted manipulations to systems. An IDS is composed of several components: Sensors which generate security events, a Console to monitor events and alerts and control the sensors, and a central Engine that records events logged by the sensors in a database and uses a system of rules to generate alerts from security events received.

and if these events possibly constitute a breach of security an alert of some sort is generated and logged.

The issue with indicators is that it is not enough to sit down and write down all the indicators you can come up with. A structured approach is needed to collect and organise these indicators. Sources for these indicators are e.g., documented incident reviews from which indicators can be extracted and information (indicator) sharing with other organisations. A structured and documented approach to collecting and organising information security incident indicators is a subject for recommended future research on the subject.

Furthermore, it may not be possible to determine indicators for each and every information security incident. On the other hand, this might not always be necessary, especially in the case of very basic incidents. Take a power outage for instance: the lights go out, computers are off, etc. In this case it is unnecessary to report these indicators, because it is apparent that there is a power outage.

Summarising, indicators may well provide a foundation for improving the information security incident management process by providing improved, i.e., swifter and more accurate, identification of information security incidents if a structured and documented approach to collecting these indicators has been defined. Additionally, indicators should be seen as an aid in the identification phase and not as obligatory for each and every information security incident.

6.2.4 Predefined responses

Considering the subject of having predefined incident responses available and linked to the classification of security incidents, experts agree that predefined responses are necessary and a crucial part of a good information security incident management process. From the available literature it became clear that predefined incident responses are essential for swift and proper incident resolution, which minimises possible damages resulting from incidents, a view shared among all experts that were interviewed. There were also several issues regarding predefined responses that came forward during the interviews with experts.

First, how to deal with practice situations that differ somewhat from the documented incidents. That said, in this research it has already been mentioned to predefine responses as much as possible and determine additional responses based on the incident specifics.

Secondly, there might be a need for different responses based on the severity level of an incident, which may be determined by how far it has spread, e.g., in the case of a computer virus. For such an approach, incident types could be further categorised into low, moderate and high severity, each with their own set of predefined responses. This will be discussed in more detail in section 6.3.
Furthermore, the idea was coined that it might be possible to generate fixed sets of predefined responses to the various incident types regardless of organisations specifics. This however may make the methodology too rigid due to the fact that most aspects are fixed.

6.2.4 Feasibility and added value

Experts found the methodology to be correct and having added value for information security incident management. They also mentioned that the concept of the methodology as it is could be applicable in practice, but it needs more structure to be actually applied as (part of) an information security incident management process.
A problem with the methodology might be how to estimate the actual severity of an incident in practice. One solution to this comes from a drawback in section 6.2.3. There an incremental approach to determining responses to incidents was proposed. This approach led to three variations of the same incident: low, moderate and high severity. This also meant having indicators linked to the different variations. Although this does amounts to extra work, it also solves the dilemma of how to estimate the actual severity of an information security incident.
Summarising, the added value of the methodology developed in this research is in its concept, not its direct applicability in practice. The concept of having a categorised list of information security incidents linked to indicators and predefined responses could improve information security incident management processes by providing improved identification and response. Discussions on the methodology provided valuable feedback and a multitude of options for future research.

6.3 Applying feedback to the methodology

In this section several proposed (possible) changes, based on the feedback described in the previous section, will be discussed and evaluated. First, a simplified severity determination method is examined. Then, we will take a look at the possibility of categorising each information security incident into three separate varieties, namely a (L)ow, (M)oderate or (H)igh severity version of an incident (e.g., a computer virus).

6.3.1 Simplified severity determination

In the classification methodology, the severity of an incident was based on the reliability need of an organisation and the specific risks to availability, integrity and confidentiality of an incident, both determined by an adapted SPRINT method.

It may be possible to merge the organisational context and risk determination into one risk questionnaire, thereby simplifying the severity determination. This questionnaire can be found in Appendix E. The questions work more or less the same as the SPRINT method, by determining for several factors if the risk is (L)ow, (M)oderate or (H)igh. The maximum value then determines overall severity, and when there are multiple (H)igh risks the severity should be (S)evere.

Although this method simplifies things, there is also a rather large drawback to this approach. In the original approach, severity of an incident was approached through separately determining the reliability need of an organisation and the specific risks to availability, integrity and confidentiality of an incident. This approach made the reliability need explicit. In the simplified approach however, this need has become implicit.

It can be concluded that the simplified approach is not necessarily wrong, as it is mostly similar to the method originally used, but it lacks the explicit reliability need and therefore will not replace the original approach used in the methodology.

6.3.2 Severity levels and predefined responses

Moreover, regarding the severity of an incident, two possible approaches to the matter were named: 1) adjust the severity, if needed, at the moment the incident takes place, and 2) define the various possible severity levels for an incident as separate incidents with their own set of predefined responses.

From these two possible approaches, the second fits the methodology best, due to linkage with different responses based on the severity level of an incident. For such an approach, incident types could be further categorised into low, moderate and high severity, each with their own set of predefined responses. Elaborating on this, it may only be necessary to define responses for a low severity version of the incident and then adding the necessary extra responses for the moderate and high severity types, i.e., using an incremental approach as depicted in Table 6.1.

A drawback to this approach is the sheer amount of work and effort that would go into it, because the amount of incidents that need to be classified would instantaneously triple (due to the possible low, moderate and high severity versions of all incident types). Moreover, this also means having distinct indicators for the low, moderate and high severity variations of the incident. Furthermore, how would the distinction be made between the various severity levels?

Considering the problems with the method that was just described, it will not replace the currently used SPRINT methods for determining severity. However, the severity that is determined through the SPRINT method should not be too rigid in practice; some amount of flexibility may prove useful as not to underestimate an incident on one hand, or overreact to an incident on the other. As mentioned before, this may be achieved by assessing an incident as it takes place and adjusting the severity if needed. Of course, incidents that were deemed to have (L)ow severity will not be adjusted to (H)igh severity and vice versa.

Responses	*Severity*		
	Low	Moderate	High
Incident A	Response X	Response X+Y	Response X+Y+Z

Table 6.1 – An example of an incremental approach to predefined responses for various severity levels of an incident.

6.4 Chapter Summary

This chapter discussed the feedback received from experts during various interviews concerning the classification methodology and incident management process. This feedback was both negative and positive and yielded most useful insights, points of attention and possibilities for improving the methodology. Furthermore, based on the feedback received, several possible options for adjusting the methodology were examined. However, these possibilities were not used to revise the methodology due to various problems and/or shortcomings, which deteriorate their possible added value.

Chapter 7

IT AUDITING AND THE METHODOLOGY

Research question 7. What are the implications of the methodology for IT auditing?

In this chapter we will take a look at possible implications of the developed classification methodology for IT auditing. First, a brief overview of IT auditing is presented, followed by an examination of the incident management part of the audit process.

7.1 Background on IT auditing

IT auditing (Information Technology Auditing) began as EDP Auditing (Electronic Data Process Auditing) and developed mainly as a result of the rise in technology in accounting systems, the need for IT control, and the impact of computers on the ability to carry out attestation (i.e., verification) services.

The introduction of computer technology into accounting systems changed the way data was stored, retrieved and controlled. During the mid-1950s to the mid-1960s, the auditing profession was still auditing around the computer, because at this time only mainframes were used and only a few people had the skills and abilities to program computers. This started to change in the mid-1960s with the introduction of new, smaller and less expensive computers which increased the use of computers in businesses. With this development came the need for auditors to become familiar with Electronic Data Processing (EDP) concepts in business. Along with the increase in computer use, came the rise of different types of accounting systems. In 1968, the American Institute of Certified Public Accountants (AICPA) had the Big Eight (now the Big Four[1]) accounting firms participate in the development of EDP auditing. The result of this was the release of Auditing & EDP, a book that included how to document EDP audits and how to process internal control reviews.

[1] Ernst & Young, Deloitte, KPMG and PricewaterhouseCoopers

Around this time EDP auditors formed the Electronic Data Processing Auditors Association (EDPAA). The goal of the EDPAA was to produce guidelines, procedures and standards for EDP audits. In 1977, the first edition of Control Objectives was published, which is now known as Control Objectives for Information and related Technology (CobiT; see chapters 2 and 3 for more information on CobiT). In 1994, the EDPAA changed its name to Information Systems Audit and Control Association (ISACA).

Several major events have had significant impact on the growth of IT auditing on a global scale, among which are the development and increasing popularity of the Internet and E-commerce, the 1998 IT failure at AT&T, and the Enron and Arthur Andersen LLP accounting scandals. These events brought a much needed focus to the importance of the accounting profession. Accountants certify the accuracy of public company financial statements and add confidence to financial markets. The heightened focus on the industry has brought improved control and higher standards for all working in accounting, especially those involved in IT auditing.

7.2 Auditing incident management

7.2.1 Introduction

In this section we will take a brief look at a part of the audit process, namely the audit of an organisation's internal control. COSO[2] defines internal control as "a process, influenced by an entity's board of directors, management, and other personnel, that is designed to provide reasonable assurance in the effectiveness and efficiency of operations, reliability of financial reporting, and the compliance of applicable laws and regulations".

As depicted in Figure 7.1, the system of internal control consists of 2 types of control, general and detail, which can be realised in 2 ways, manually or automated. The general internal control

[2] The Committee of Sponsoring Organizations of the Treadway Commission. http://www.coso.org.

measures, such as segregation of duties, determine the requirements for the detailed internal controls, i.e., the general control measures form the basis for detailed control measures.

Since this research is on information security incidents, we will focus on the General IT Controls. General IT Controls consists of three subjects: Change management, Logical access control and Other. Incident management is a part of the "Other" IT controls, which furthermore consists of backup and recovery, scheduling and monitoring. The control objective for (problem and) incident management is to determine that problems and incidents are identified and resolved in a timely manner.

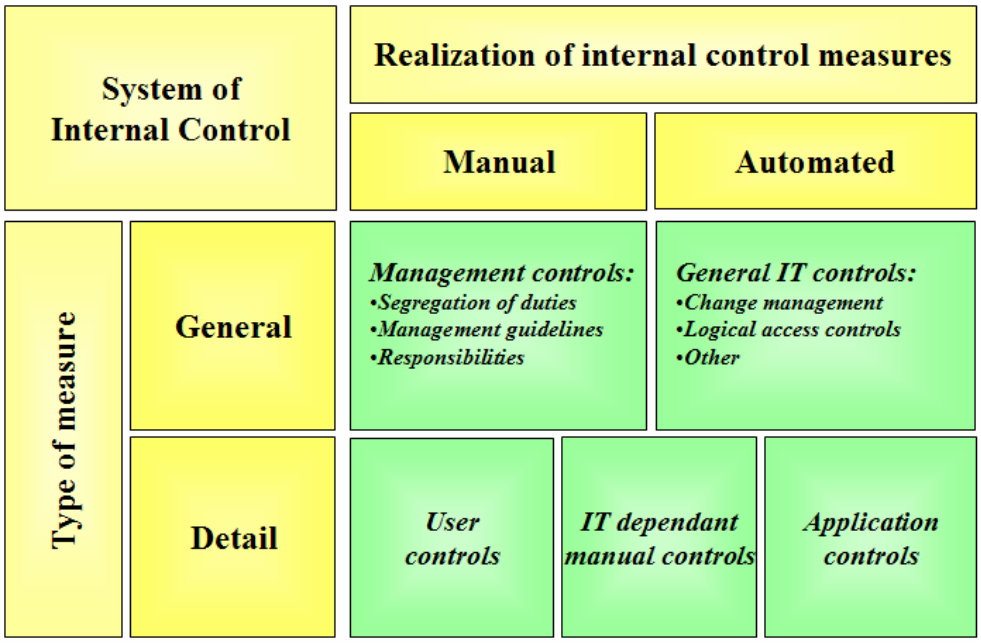

Figure 7.1 – Internal control framework overview (©Ernst & Young)

The control objective for incident management was to determine that incidents are identified and resolved in a timely manner. The feedback on the methodology developed in this research showed that it was a good foundation for proper identification of incidents and made swifter and more accurate responses to incidents possible, thus satisfying the aforementioned control objective for

incident management. As a result, an organisation that uses the methodology will most likely satisfy the control objective.

7.2.2 Testing framework

E&Y developed a testing framework , based on e.g., ITIL and COBIT, which makes it possible to score the control objectives of an organisation in a qualitative manner [EYRAPP]. Control objectives are divided into multiple normative control measures, which are score on a 0 (non-existent) to 5 (excellent) scale. In this section the incident management part of the framework is examined, and adjusted with new normative control measures to reflect the results of the methodology developed in this research. The incident management control objective with its current set of normative control measures is described in Appendix F (in Dutch).

There are two components in the incident management control objective that are affected by this research, namely 'Classification and allocation' and 'Diagnosis and solution'. Each of these components will be dealt with separately.

In the Classification and allocation component, incidents need to be classified by priority and solution group. Furthermore, incidents need to be prioritised the moment they come in. In the methodology developed in this research, incidents are classified in advance and linked to an appropriate response (set). Additionally, the handling priority of an incident in the methodology is determined beforehand based on its severity, which is determined by a risk analysis. It may be that the actual severity is slightly different but it will not vary much.
It is hereby proposed to replace two normative measures, namely InM 10 (incidents need to be classified by priority and solution group) and InM 11(incidents need to be prioritised the moment they come in) by the following two normative control measures:

- Classify incidents by type and priority through triage.
- Assess if actual severity (priority is in accordance with pre-determined severity. Adjust severity, and thus priority, if needed.

In the 'Diagnosis and solution' component, it is stated that in the case of recognition and handling of known errors, the help desk function needs to make use of documented solutions. If only responses to a small set of known errors are documented, there will most likely be incidents that cause more damage than necessary. In the methodology, incidents are classified beforehand and linked to appropriate responses to minimise damages for all types of incidents.

It is hereby proposed to replace normative measure InM 12 (in the case of recognition and handling of known errors, the help desk function needs to make use of documented solutions) by the following normative control measure:

- The help desk function needs to make use of the documented responses for all incidents.

7.3 Summary

This chapter examined the possible influences of the methodology that was developed in this research on IT auditing. It became clear that an organisation that uses the methodology will most likely satisfy the control objective for incident management, which is to determine that incidents are identified and resolved in a timely manner. Furthermore, it was shown that the methodology could impact the way that incident management is evaluated, by updating several normative control measures to reflect the findings of this research.

Chapter 8

CONCLUSION

8.1 Summary

First, we looked at information security in general and what its role in organisations is. Information plays a huge role in organisations these days and that protecting that information has become vital for the business continuity of most organisations. Information security tries to protect that information and has to be approached as a continuous process, not as a separate set of activities.
Aspects that are important for information security are the reliability aspects of information (confidentiality, availability and integrity), different security controls (preventive, detective, and corrective versus logical, technical and organisational) and requirements for a successful security programme (e.g.,). Furthermore, there are various factors that impose limits on information security, such as security spending that is lagging behind total IT expenditures. Finally, several best practices and standards were briefly addressed, such as COBIT and ITIL.

Then, information security incident management and current security incident management best practices were examined. Information security incidents are those events that can cause damage to confidentiality, integrity or availability of information or information processing, and in the process may threaten business continuity as a result. Incidents can be defined by various aspects, such as their nature (acts of men and God) and intentionality (intentional or unintentional).
Furthermore, the incident management processes of ITIL, COBIT were examined, as well as the ISF information security incident management process. Several shortcomings were discovered. One recurring shortcoming was that no method for the classification of security incidents was given by these best practices, although they recognised the importance of classification for the incident management process. This is where the goal of this research comes into play: to design a methodology that can improve information security incident management. Such a methodology can possibly provide a practical contribution to the incident management processes of the aforementioned best practices.

The fourth chapter was about how to improve information security incident identification and response. A methodology for classifying security incidents (by ISF type), linked to possible indicators (for triage and faster identification) and proper responses (for better incident response), was laid down. The output of the methodology is a classification which links security incidents, severity ratings of those incidents, potential indicators for security incidents and proper predefined responses for security incidents together.
It was explained that the methodology serves as an addition to current best practices regarding information security incident management processes, not as a replacement.

Chapter 5 revolved around applying the methodology in practice. A possible approach for the application and use of the methodology was set up. Not only an example process for implementing the classification methodology was described, issues regarding the documenting (i.e., recording) and updating of the classification were also examined. Furthermore, an example process of using the methodology, once implemented, was discussed.

Chapter 6 dealt with the question whether or not the methodology could improve security incident identification and response in practice. This question was answered through the feedback received from experts during various interviews concerning the classification methodology and (information security) incident management.
The theoretical foundation of the methodology was found to be solid and experts deemed that the methodology could improve information security identification and response.
Furthermore, based on the feedback received, several possible options for adjusting the methodology were examined. However, these possibilities were not used to revise the methodology due to various problems and/or shortcomings, which deteriorate their possible added value and in turn the value and use of the methodology.

8.2 Conclusions

The objective of this research was:

> *… to formulate a methodology that can be used to improve information security incident identification and response in organisations.*

To reach the research goal mentioned above, various research questions were answered over the course of several chapters (see the summary in the previous section).

From the research it can be concluded that a methodology for improving incident identification and response has been successfully set up. The methodology that was set up during the course of this research is summarised in Figure 8.2 on the following page (see also Figure 4.6, chapter 4).
Possible practical usage of the methodology is depicted in Figure 8.1. This shows that the methodology can be applied in practice. (note that Figure 8.1 is just an example of usage).

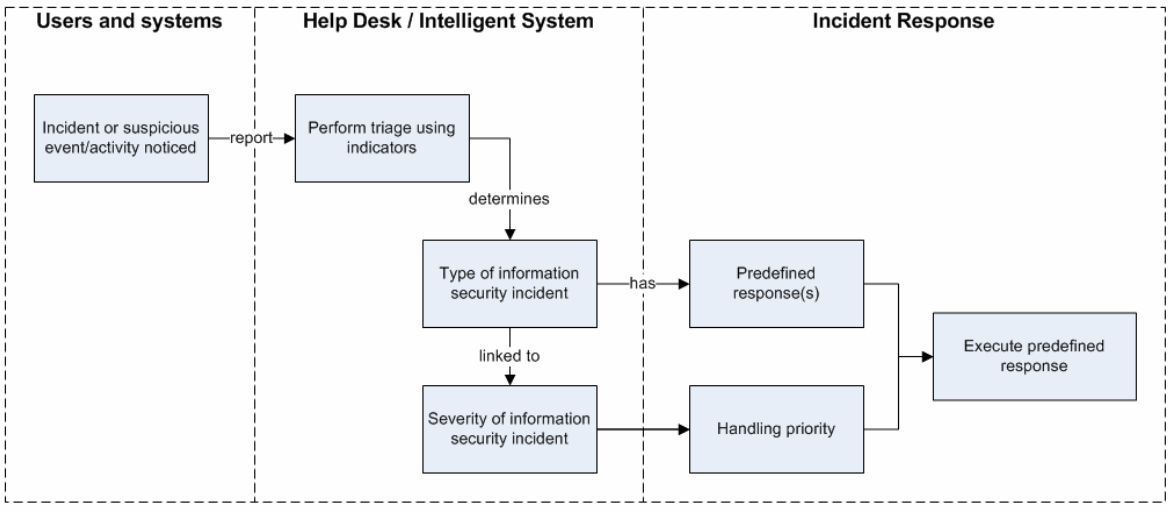

Figure 8.1 – From identification to response using the classification methodology

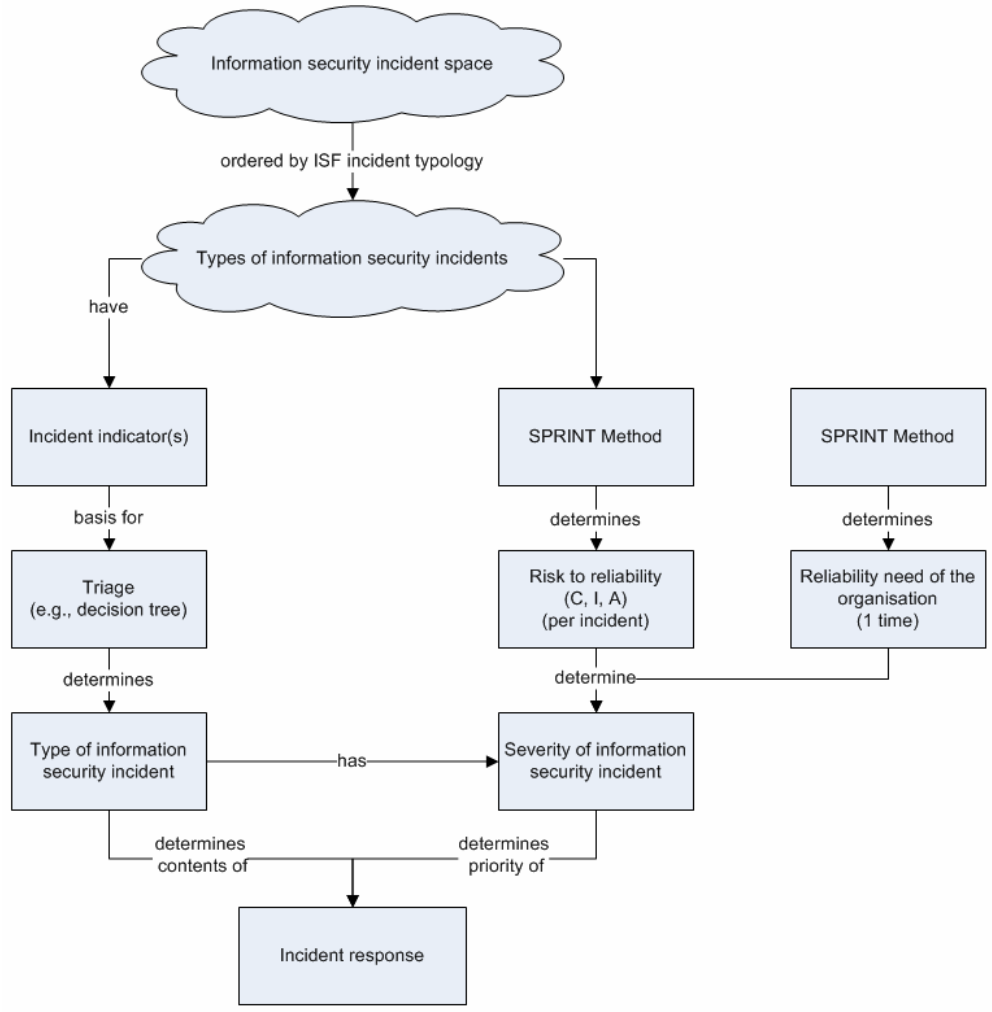

Figure 8.2 - Detailed overview of the classification methodology.

Figure 8.2 above shows a detailed view of the interdependencies of the various components of the methodology, including supporting components such as SPRINT (for severity determination) and the ISF typology (for the incident types) and the placement of triage (for identifying the incident) in the methodology.

Concluding, it can be said that the target of this research was achieved. Although the methodology was not applied in practice, experts found the methodology to add value to the information security incident management process and improve information security incident identification and response.
The methodology provides a new look at information security incident management compared with conventional methods, by linking all the various components together and providing a method for classifying information security incidents. The methodology's added value lies in the method for classification and the linkage between indicators ➔ incidents ➔ response(s).

Of course, there are also several things that need to be researched in order to apply and implement the methodology in practice as a part of an organisation's information security process. Some of these issues are discussed in the next section on topics for future research.

8.3 Future research

This research yielded, besides the developed methodology, several interesting subjects for possible future research. Some of them will be briefly discussed here.

Probably the topic most in need of research is the topic of the so-called incident indicators. The main questions here are how to attain, organise and apply these indicators in the incident management process.
As far as the attaining (data collection) of indicators is concerned, two possibilities are: 1) examination of documented post-incident reviews, for which indicators for particular incidents can be derived, and 2) cooperation between organisations, sharing data on incidents; data sharing means that organisations can learn from each other, which may lead to better information security incident management processes for all the organisations involved.
Then there is the matter of how these indicators need to be organised and linked together. A good approach to this could be the use of incident indicators in a decision tree, since that is basically how triage works (see Figure 8.3).

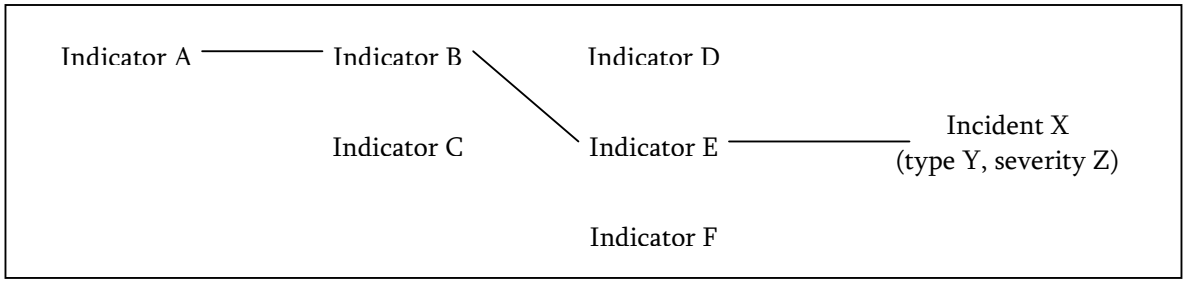

Figure 8.3 – Using the indicators in a decision tree.

How to apply this decision tree is the third matter. Such a decision tree could be implemented in an automated intelligent system, in which someone can enter an indicator and is then asked which other indicators apply to the current incident situation, or the intelligent system performs triage itself. This then eventually yields the specific type of incident, and since the priority and responses are already linked to that incident, containment and resolution of the incident can commence immediately.

Furthermore, the methodology described in this research should be evaluated further. Although several experts were interviewed on the methodology, it has not yet been practically applied. A study on how to exactly apply the methodology in practice will probably lead to a better understanding of it's working in practice and what the advantages and disadvantages are when using the methodology. For the methodology to be used in practice it needs to be made operational in some way. The theoretical foundation of the methodology was found to be sound by experts that gave their feedback on it. An example of what needs to be made operational is the implementation of the categorised list of information security incidents, mapped to the ISF information security incident categories.

Another interesting topic for future research is the possibility to classify beyond only incidents. Now, we only classified the incidents in the incident space to the ISF categories. But maybe there is a possibility for classifying indicators and responses in a way similar to incidents, into similar groups. Take the responses for example: all the responses possible form the response space, similar to the incident space containing all the incidents. The question is whether there is a way to organise these

responses in to groups of (seemingly) related ones, like the external attack group in the ISF information security incident list which consists of e.g., hacking, cracking, spoofing and social engineering.

Setting up a method for determining the information security incident management maturity, in which the classification methodology described in this thesis may play a role may also be interesting. [MEUL05] set up a method for the determination of information security maturity; such a method could also be valuable for determining the information security incident management maturity of an organisation. This is mostly due to the relatively small influence of incident management on the maturity score of an organisation. And in [ISIM06] it is stated that incidents cost organisations e.g., revenue, customers and public trust.

Another possibility may be to adjust the information security maturity model in such a way that incident management becomes a more important part in determining the overall information security maturity level of an organisation, given the possible impact of information security incidents on the business.

The list of topics for future research that has been described in this section is not exhaustive, nor is it meant to be. It does however provide insight in the possible research that still can, and needs, to be done.

APPENDICES

Appendix A
INCIDENT RESPONSE TEAMS

A.1 Guidelines for setting up an IRT

According to [NIST04] an IRT should be available for anyone who discovers or suspects that a security incident involving the organisation has occurred. One or more team members, depending on the magnitude of the incident and availability of personnel, should then handle the incident. The incident handlers analyze the incident data, determine the impact of the incident, and act appropriately to limit the damage to the organisation and restore normal services.

The NIST guide states that the structure of an IRT falls into one of three categories. For small organisations and for large organisations with minimal geographic diversity in terms of computing resources, it is advisable to have a central Incident Response Team. Distributed Incident Response Teams are effective for large organisations and for organisations with major computing resources at distant locations. Nonetheless, distributed incident response teams should be part of a single centralized body in order for the incident response process to be consistent organisation-wide. The final structure is a coordinating IRT, which provides guidance to other teams without having authority over them.

Furthermore, the NIST guide distinguishes three ways to staff IRTs: use only employees, partly outsource IRTs or fully outsource IRTs. The choice for each alternative depends on the need for security, the expertise within the organisation and the jobs that specific IRTs have to perform. For example, 24/7 monitoring of firewalls is usually outsourced. In the selection of a staffing method for the IRT(s), an organisation should consider such factors as the need for 24/7 availability, full-time versus part-time members, employee morale and expertise, organisational structure and the costs of the various alternatives.

Regardless of the incident response model an organisation chooses (central/distributed/coordinating), a single employee should be in charge of incident response (with one other employee acting as backup for that employee). Furthermore, members of the IRT should have excellent skills, both technical and otherwise, because these skills are critical to the team's success. Critical skills include for example system administration, network administration, technical support, good problem solving skills, teamwork skills, communication skills and writing skills. It is important to continuously build and maintain skills, for example by conducting simulated incident handling exercises for the team. Exercises are particularly important because they not only improve the performance of the incident handlers, but also identify issues with policies, procedures and communication.

It is essential to identify other groups within the organisation that may be needed to participate in incident handling so that their collaboration can be requested before it is needed. Every incident response team relies on the expertise, judgment, and abilities of others, including management, the information security team, IT support, business continuity, HR and the legal department, among others. As was mentioned earlier in this thesis, management's role in information security, and thus incident response, is pivotal. In the most elementary sense, management establishes incident response policy, budget, and staffing and is ultimately held responsible for coordinating incident response among various stakeholders and minimising damage.

Although the main focus of an IRT is incident response, there are usually various other services offered by an IRT. Such services include advisory distribution (i.e. distributing information regarding (the mitigation of) new vulnerabilities), vulnerability assessment (examining systems for possible exploits, risk assessment), education and awareness (workshops, seminars, awareness programmes) and patch management (acquiring, testing, and distributing patches throughout the organisation).

A.2 Incident response handling roles

'A framework for incident response' (by [ISTD02]) describes, among other things, incident response handling roles. It proposes that one full-time member of an organisation's security team should act as

Incident Investigator and Coordinator (IIC). Additionally, any member of the Incident Response Team (IRT) may serve as the Incident Liaison (IL). The IL has to be assigned by the IIC based on the area the incident affects (e.g. hardware, software, documentation).

The IIC has to assign the severity levels to incidents and perform all investigative duties. IIC duties further permit unlimited access to directly affected resources, where that access has to be granted and monitored by the IL for the extent of the investigation. The IIC and IL have to determine the requirements and necessities of disrupting further services as part of the recovery from incident. The IIC coordinates evidence documentation, the incident lifecycle and finally evidence storage.

The IL has to act as coordinator and liaison to resources required by the IIC. The IL must also act as secondary witness to all system modifications when forensic analysis will be performed and verify that evidence has been collected without disruption or corruption. The IL should oversee documentation and reporting of all factual information and verify that said documentation is delivered to the appropriate people. Finally, the framework proposes the response handling and investigation stages should proceed in a methodical manner.

A.3 IRT operational issues

Four important aspects of Incident Response Team operations, as given by [CSIR03], will be discussed here. This will give some more insight in how IRTs should operate in practice. Although the handbook focuses on Computer Security Incident Response Teams, its content is nevertheless very much applicable to general IRTs. The Handbook for CSIRTs distinguishes four main operational issues, namely fundamental policies, continuity assurance, security management and staff issues, which will be dealt with in the following sections.

Fundamental policies

Some policies are fundamental for any organisation and need to be in place regardless of organisation specifics. However, these policies' contents should reflect the service and quality specifics of an organisation. We will discuss three of them here. (Of course, the security policy (see chapter 2) should not be forgotten).

A code of conduct provides basic direction about how to react in certain situations and sets the basis for interactions both within and external to the team. The code of conduct should become incorporated into the behaviour of incident handlers.

It is imperative for an IRT to define an information disclosure policy for incident management and response. Such a policy describes what IRT members can say to whom and when. There are basically three factors that determine if, to what extent and how information will be disclosed. These factors are purpose (someone has a "need to know" the information), target (whom it concerns) and category (decided by the information categorisation policy) of the information.

Finally, a human error policy can help in minimising and containing damages caused by human errors. It should clearly state what possibilities a staff member should consider in case of an error he/she made; it should clearly state the proper reactions from management; and it should outline the consequences.

Continuity assurance

The continuity of dependable services is essential to the successful operation of an IRT. This directly reflects on the perceived competence and level of trust of a team by its constituency. The timeframe for which one seeks to assure continuity may make quite a difference for the kind of problems encountered, the services that can be provided, and the measures to be taken. Regarding the timeframe, continuity threats can be divided in short-term issues (days/weeks, e.g. lack of time), medium term issues (months, e.g. staff burnout) and long-term issues (years, e.g. limited ability to adapt to change due to stabilisation).

Workflow management is important to consider for helping to assure a CSIRT's continuity of work. It is defined as managing the flow of events that are part of work. Incident handling continuity problems arise as CSIRTs have to deal with a lot of problems over longer periods of time, because of the continually changing events and also changing team members working on these problems.

Other factors that are important for assuring IRT continuity are out-of-hours coverage (as opposed to business hours coverage) and off-site coverage (e.g. when the IRT is at a backup facility in case of a crisis situation).

Security management

An IRT obviously has to place great emphasis on guarding its own security. When considering the goals for IRT security management, factors such as confidentiality, availability, integrity, privacy, exclusivity, etc. need to be taken into account.

Other elements that are important for an IRT's security management are e.g. isolated networks for tests, providing off-site access to local facilities, physical security and handling internal security incidents.

Internal security incidents may irreparably damage the IRT's reputation, so it is essential an IRT is prepared for such incidents and addresses them when they occur.

Staff issues

IRT work is basically service based and as a result, there is an inherent dependence on competent and trustworthy staff. Hiring appropriate staff thus plays a key role in ensuring the mission and service of the operation of the IRT.

A critical criterion for IRT staff is an individual's willingness and ability to follow procedures and to provide a professional interface to all parties interacting with the IRT. Good interpersonal and communication skills are far more important than technical expertise.

When hiring IRT staff (whether internal or external to the organisation), it is important to decide in advance the hiring process that will be used. The hiring process should be designed to ensure that candidates have suitable interpersonal skills and has, or can be trained in, the needed technical skills.

Training of IRT staff is necessary from three perspectives: bringing new staff members up to the necessary skill level, broadening the abilities of IRT staff, and keeping the overall IRT skill set up-to-date with emerging technologies and security trends. Furthermore, it will be important for every member of the team to be trained in areas specific to the incident handling functions and the local team environment.

Appendix B
ADAPTED SPRINT METHOD

The methods described in this appendix are based on the adapted SPRINT (Simplified Process for Risk Identification) method by [STOF04]; First, the SPRINT method for determining an organisation's reliability need will be discussed. Secondly, the SPRINT method for determining the risks of an information security incident will be discussed.

B.1 Reliability need

The SPRINT method in this section can be used to determine the organisational (specific) context in which information security incidents may take place. The context used here is the information reliability need of an organisation, i.e., the importance of confidentiality, integrity and availability of information for an organisation's business (continuity). The answers to the questions of the method can be translated, by an expert, into the reliability need of an organisation. The final reliability need is described here by a triple in the form *(C,I,A)*, where each component can take the value's High *(H)*, Moderate *(M)* and Low *(L)*. For example, the reliability need of organisation X is *(H,L,M)*.

H: High impact, ranging from severe damages to threatening to business continuity.
M: Moderate impact, ranging from slight damages to moderate damages.
L: Low impact, ranging from no damages to slight damages.

The questions are presented on the following pages.

B.1.1 Confidentiality

	Impact
Competitive disadvantages To what extent can the organisation suffer damages when a competitor has access to information (contained in the information facilities)?	
Direct loss of revenue To what extent can revenue be lost as a result of leaking of information?	
Public trust To what extent can the trust of customers, the public, the stockholders or the suppliers be damaged as a result of leaking of information?	
Additional costs To what extent can additional costs arise as a result of leaking of information?	
Legal liability To what extent can the leaking of information result in breaches of legal, contractual or other required guidelines?	
Employee morale To what extent can leaking of information have damaging effects on employee morale and/or motivation?	
Fraud To what extent can goods or money be embezzled as a result of leaking of information?	
Need for confidentiality of information (H, M or L)	

B.1.2 Integrity

	Impact
Decisions of management To what extent can decisions of management be negatively influenced as a result of information inaccuracies?	

Direct loss of business To what extent can revenue be lost as a result of information inaccuracies?	
Public trust To what extent can the trust of customers, the public, the stockholders or the suppliers be damaged as a result of information inaccuracies?	
Additional costs To what extent can additional costs arise as a result of information inaccuracies?	
Legal liability To what extent can information inaccuracies result in breaches of legal, contractual or other required guidelines?	
Employee morale To what extent can employee morale and/or motivation be damaged if they can't rely on the integrity of information?	
Fraud To what extent can the embezzlement of goods or money be the result of, or remain unnoticed due to, unauthorised changes of information?	
Disruption of business processes To what extent can the primary business processes be disrupted because of information inaccuracies?	
Need for integrity of information (H, M or L)	

B.1.3 Availability

	Impact		
Duration of information unavailability	*hour*	*day*	*week*
Decisions of management To what extent can decisions of management be negatively influenced as a result of information not being available?			

Direct loss of revenue To what extent can revenue be lost as a result of information not being available?			
Public trust To what extent can the trust of customers, the public, the stockholders or the suppliers be damaged as a result of information not being available?			
Additional costs To what extent can additional costs arise as a result of information being unavailable?			
Legal liability To what extent can unavailability of information result in breaches of legal, contractual or other required guidelines?			
Employee morale To what extent can employee morale and/or motivation be damaged if information is unavailable?			
Fraud To what extent can the embezzlement of goods or money be the result of, or remain unnoticed due to, unavailability of information?			
Disruption of business processes To what extent can the primary business processes be disrupted because of information being unavailable?			
Repairs To what extent can additional costs arise related to the catching up of work backlogs as a result of information unavailability?			
Need for availability of information (H, M or L)			

B.2 Information security incident risk to reliability

The adapted SPRINT method in this section can be used to determine the risks of an incident to the reliability (i.e., confidentiality, integrity and availability) of information. The risk is defined here as the probability that a certain situation will happen in case of a particular incident. Determining the risk (i.e., probability) that certain situations may occur needs to be done by experts within the target organisation that have extensive knowledge of the information infrastructure, information need and state of information security of the organisation.

The answers to the questions of the method can be translated into the risks of an incident to reliability. The risk of each of the three attributes is the maximum value of risk that is determined for each particular attribute; e.g., if only one question yields a high *(H)* risk, the final attribute value will still also be high *(H)*. The final risk to reliability is described here by a triple in the form *(C,I,A)*, where each component can take the value's High *(H)*, Moderate *(M)* and Low *(L)*. The values are presented below.

H: High risk, i.e., it is very likely that the incident will result in the presented situation.

M: Moderate risk, i.e., it is possible that the incident will result in the presented situation.

L: Low risk, i.e., it is very unlikely that the incident will result in the presented situation.

B.2.1 Confidentiality

Could the information security incident affect confidentiality of information? If yes, determine the overall risk to confidentiality by answering the questions below. If no, the risk to confidentiality is non-existent *(N)*.	
	Risk
Competitive disadvantages What is the risk of a competitor gaining access to information (contained in the information facilities) due to the incident breaching confidentiality of information?	

Direct loss of revenue What is the risk that revenue will be lost as a result of the incident breaching confidentiality of information?	
Public trust What is the risk that the trust of customers, the public, the stockholders or the suppliers is damaged as a result of the incident breaching confidentiality of information?	
Additional costs What is the risk that additional costs arise as a result of the incident breaching confidentiality of information?	
Legal liability What is the risk of breaches of legal, contractual or other required guidelines as a result of the incident breaching confidentiality of information?	
Employee morale What is the risk that the breach of confidentiality of information by the incident negatively affects employee morale and/or motivation?	
Fraud What is the risk of goods or money being embezzled as a result of the incident breaching confidentiality of information?	
Risk to confidentiality of information (= Max(Risk))	

B.2.2 Integrity

Could the information security incident affect integrity of information? If yes, determine the overall risk to integrity by answering the questions below. If no, the risk to integrity is non-existent *(N)*.	
	Risk
Decisions of management What is the risk that decisions of management are negatively influenced as a result of the incident affecting information integrity?	
Direct loss of business What is the risk that revenue will be lost as a result of the incident affecting information integrity?	
Public trust What is the risk that the trust of customers, the public, the stockholders or the suppliers is damaged due to the incident affecting information integrity?	
Additional costs What is the risk that additional costs will arise as a result of the incident affecting information integrity?	
Legal liability What is the risk that legal, contractual or other required guidelines are breached due to the incident affecting information integrity?	
Employee morale What is the risk that employee morale and/or motivation are damaged due to the incident affecting information integrity?	
Fraud What is the risk that the incident affecting information integrity results in the embezzlement of goods or money?	
Disruption of business processes What is the risk that primary business processes are disrupted because of the incident	

affecting information integrity?	
Risk to integrity of information (= Max(Risk))	

B.2.3 Availability

Could the information security incident affect availability of information? If yes, determine the overall risk to availability by answering the questions below. If no, the risk to availability is non-existent *(N)*.			
	Risk		
Duration of information unavailability due to the incident	*hour*	*day*	*week*
Decisions of management What is the risk that decisions of management are negatively influenced as a result of unavailability of information due to the incident?			
Direct loss of revenue What is the risk of revenue being lost as a result of unavailability of information due to the incident?			
Public trust What is the risk that the trust of customers, the public, the stockholders or the suppliers are damaged as a result of unavailability of information due to the incident?			
Additional costs What is the risk that additional costs will arise as a result of information being unavailable due to the incident?			
Legal liability What is the risk that legal, contractual or other required guidelines are breached due to the unavailability of information due to the incident?			
Employee morale			

What is the risk that employee morale and/or motivation are damaged as a result of information unavailability due to the incident?			
Fraud What is the risk that the unavailability of information due to the incident results in the embezzlement of goods or money?			
Disruption of business processes What is the risk that information being unavailable due to the incident disrupts primary business processes?			
Repairs What is the risk that the incident leads to additional costs related to the catching up of work backlogs as a result of information unavailability?			
Risk to availability of information (= Max(Risk))			

Appendix C

ISF INCIDENT TYPOLOGY

Incident category	Incident type	Incident description
External attack	Carrying out denial-of-service attacks	Deliberately overloading systems and network devices or re-directing network traffic.
	Hacking	Penetrating systems and networks.
	Undertaking malicious probes or scans	Probes or scans of network devices and systems to gather information that could be used to undertake an attack.
	Cracking passwords	Determining the plaintext version of an encrypted password.
	Cracking keys	Determining the plaintext version of an encrypted key (e.g. WEP keys in wireless networks).
	Defacing web sites	Unauthorised modification of web site content.
	Spoofing web sites	The creation of a bogus web site - that masquerades as a genuine web site - to which users are directed.
	Spoofing user identities	The unauthorised use of valid user identity information to gain access to a system (typically as a result of 'identity theft').
	Modifying network traffic	Modifying the content of network traffic in transit or falsifying the source or destination address of network traffic.
	Eavesdropping	The unauthorised interception of information in transit.
	Distributing computer viruses (including worms)	Self-replicating programs that propagate between systems and carry out an unauthorised action or set of actions (typically referred to as the payload).
	Introducing Trojan Horses	Computer code that masquerades as an authorised program but which carries out an unauthorised action (or set of actions).
	Introducing malicious code	The introduction of malicious code (e.g. rootkits), malicious mobile code (e.g. unauthorised active content), 'spyware' or 'adware'.
	Carrying out social engineering	The deliberate manipulation of staff to elicit information that can be used to undertake an attack (e.g. by providing UserID and password details).
	Distributing SPAM	Excessive distribution of unsolicited (commercial) messages (including e-mail, instant messaging and telephony).
Internal misuse and	Gaining unauthorised access	Deliberately gaining access to computer systems

abuse	to systems or networks	or networks to which a user is not authorised (e.g. by means of password theft or other covert action).
	Changing system privileges without authorisation	Changing system privileges to either enable or deny access to information or functionality.
	Changing or adding software without authorisation	Changing or adding software to produce unauthorised system behaviour or action (e.g. to bypass mechanisms introduced to prevent fraud).
	Modifying or inserting transactions, files or databases without authorisation	Changing or adding transactions, files or databases to produce unauthorised system behaviour or actions.
	Misusing systems to cause disruption	Using authorised system facilities in a malicious or excessive manner to adversely affect system performance and availability (e.g. by downloading or uploading high volume .mp3 or .mpeg files).
	Misusing systems to commit fraud	Using authorised systems to defraud the organisation (e.g. diverting goods to false delivery addresses).
	Downloading or sending of inappropriate content	Using authorised communication services, such as e-mail and instant messaging, to carry out unauthorised activities (e.g. downloading or sending of obscene, discriminatory or harassing content).
	Installing unauthorised software	Installing software from an untrusted source (e.g. from an e-mail attachment).
	Disclosing authentication information	Unauthorised disclosure of authentication details (e.g. sharing of UserID and password).
	Disclosing business information	Unauthorised disclosure of business information (e.g. confidential financial information).
Theft	Software piracy	Software piracy includes the illegal copying and use of unlicensed software.
	Theft of business information	Theft of business information (e.g. customer lists, product designs, intellectual property).
	Theft of identity information (e.g. as a result of Phishing)	Theft of personally identifiable information (e.g. credit card numbers, employment IDs, personal health details).
	Theft of computer equipment	Theft of computer equipment (e.g. laptops, PDAs, mobile phones).
	Theft of authentication information	Theft of authentication information (e.g. UserID, passwords or PIN numbers).
	Theft of software	Theft of software (e.g. programs or methodologies).
System malfunction	Malfunction of business application software developed in-house	Incorrect execution or failure of software developed or integrated in-house (e.g. a software bug or system 'abend').
	Malfunction of business application software acquired	Incorrect execution or failure of software acquired from a third party (e.g. SAP R/3, Baan

	from a third party	IV, Oracle Financials).
	Malfunction of system software	Incorrect execution or failure of system software (e.g. operating system software or utilities).
	Malfunction of computer/network equipment	Malfunction of computer/network equipment (e.g. routers, hubs, switches).
Service interruption	Damage to, or loss of, computer facilities	Damage to or loss of computer facilities (e.g. data centres, computer/network rooms, trading floors or process control systems).
	Damage to, or loss of, communications links/services	Damage to or loss of communications links/services (e.g. Internet connections, Gigabit Ethernet networks, wireless networks).
	Loss of power	Failure of mains electricity or back-up power supply.
	Damage to, or loss of, ancillary equipment	Damage to or loss of ancillary equipment (e.g. computer cooling equipment).
	Natural disasters	Natural disasters include earthquakes, fires and extreme weather (e.g. flooding or freezing).
	System overload	Excessive system activity causing performance degradation or failure.
Human error	User errors	Mistakes made by staff who use systems (e.g. mistakes in inputting data, incorrect operation of workstations, sending material to the wrong address).
	IT/network staff errors	Mistakes made by staff responsible for operating and maintaining computers or networks.
Unforeseen effects of change	Unforeseen effects of introducing new/upgraded business processes	Adverse or damaging impact upon information and/or systems from the introduction of new/upgraded business processes.
	Unforeseen effect of changes to software	Adverse or unwanted system results from newly implemented or recently changed software.
	Unforeseen effect of changes to business information	Adverse or damaging impact upon systems and/or information from changes made to business information (e.g. customer lists, product designs and other intellectual property).
	Unforeseen effects of changes to computer/communication equipment	Adverse or unwanted system results from newly implemented or recently changed computer or network hardware.
	Unforeseen effect of organisational changes	Adverse or damaging impact upon systems and/or information from organisational changes (e.g. as a result of mergers, acquisitions, outsourcing or internal reorganisation).
	Unforeseen effect of changes to user processes or facilities	Adverse or damaging impact upon systems and/or information from the changes to user processes or facilities (e.g. user/operating procedures, staffing or accommodation).

Appendix D
VRAGENLIJST

Introductie

Om de methode voor het classificeren van informatiebebeveiligingsincidenten, zoals beschreven in deze thesis, te kunnen beoordelen op zaken als toepasbaarheid, haalbaarheid, juistheid, volledigheid en toegevoegde waarde, zijn enkele experts geïnterviewd over deze zaken. Zij werden gevraagd naar hun mening over o.a. de methode in relatie tot de doelstelling van het onderzoek en de methodiek, met al zijn componenten, zelf. De feedback verkregen van deze experts in de interviews vormt de basis voor het vormen van een oordeel over de waarde van de methodiek. De vragen die zijn gebruikt tijdens de interviews zijn in deze appendix opgenomen.

Doelstelling

De doelstelling van de incident classificatie methode uit het onderzoek is om informatiebeveiligings-incidenten sneller en beter te identificeren en er sneller en adequater op te reageren in vergelijking met een situatie waar er geen vooraf gedefinieerde respons voor handen is. In een onderzoeksrapport van het ISF kwam naar voren dat nog veel organisaties een onvolwassen incident management proces hebben. [ISIM06]

Vraag: *In hoeverre sluit de methode voor het classificeren van informatiebeveiligingsincidenten aan op de doelstelling? D.w.z., maakt de methode het mogelijk om informatiebeveiligingsincidenten sneller en beter te identificeren en er sneller en adequater op te reageren?*

Classificatie methode

Classificatie aspecten

De classificatie methode bestaat uit het classificeren van informatiebeveiligingsincidenten naar type en naar ernst. Voor het type wordt gebruik gemaakt van de ISF typologie voor informatiebeveiligings-incidenten. De ernst wordt bepaald op basis van de informatie betrouwbaarheidsbehoefte van een

organisatie en de risico's voor de betrouwbaarheid van informatie van een incident. De betrouwbaarheidsbehoefte en de risico's worden beiden bepaald door middel van (een aangepaste versie van) de SPRINT methode.

Vraag: *In hoeverre is het classificeren naar type en ernst van een incident juist en afdoende? (Zou er bijv. naar andere en/of meer aspecten geclassificeerd moeten worden?)*

Vraag: *In hoeverre is de ISF typologie geschikt voor het classificeren naar type? (Dekt deze typologie het gehele spectrum van incidenten?)*

Vraag: *In hoeverre vormen A) de betrouwbaarheidsbehoefte m.b.t. informatie van een organisatie en B) de risico's van een informatiebeveiligingsincident een goede basis voor het bepalen van de ernst van informatiebeveiligingsincidenten?*

Vraag: *In hoeverre is de methode SPRINT geschikt om de betrouwbaarheidsbehoefte m.b.t. informatie van een organisatie te bepalen?*

Vraag: *In hoeverre is de methode SPRINT geschikt voor het bepalen van de risico's van een incident op de betrouwbaarheid van informatie?*

Incident indicatoren

Daarnaast worden incidenten gekoppeld aan indicatoren, welke kunnen helpen bij het vaststellen van het type incident door middel van een zogenoemd "triage"-proces (simpel voorbeeld: een brandalarm dat afgaat is een indicator voor een brand). Het is namelijk vaak zo dat niet het incident wordt gemeld, maar eerder een situatie of activiteit die een voorbode dan wel gevolg is van een incident.

Vraag: *In hoeverre is het nuttig dan wel mogelijk om indicatoren aan informatiebeveiligings-incidenten te koppelen?*

Vraag: *In hoeverre kunnen indicatoren (triage) helpen om informatiebeveiligingsincidenten sneller en beter te identificeren in vergelijking met een situatie waar geen indicatoren worden gebruikt?*

Vooraf gedefinieerde incident responses

Verder worden informatiebeveiligingsincidenten gekoppeld aan vooraf gedefinieerde responses. Door de responses voor informatiebeveiligingsincidenten zoveel mogelijk vooraf te bepalen, kunnen incidenten sneller en adequater worden beheerst en opgelost. Dit resulteert in beperking van de eventuele schadelijke gevolgen van een dergelijk incident (bijv. verlies van klanten, reputatieverlies en omzet). De responses op incidenten zijn in de thesis ingedeeld in de volgende 5 categorieën: juridisch, intern (binnen de organisatie), stakeholder (klanten, leveranciers, etc), media, en beheersing van het incident.

Vraag: *In hoeverre kunnen vooraf gedefinieerde responses helpen om informatiebeveiligingsincidenten sneller en beter te beheersen en op te lossen in vergelijking met een situatie waar ad hoc bepaalde responses worden gebruikt?*

Vraag: *In hoeverre is het mogelijk om responses vooraf te definiëren?*

Haalbaarheid

Dat de methodiek correct is en zelfs toegevoegde waarde heeft voor incident management, wil niet zeggen dat hij ook geschikt is om daadwerkelijk geïmplementeerd te worden. Bij daadwerkelijke toepassing spelen zaken als tijd en geld een belangrijke rol, maar ook de werkelijke efficiëntie en bruikbaarheid.

Vraag: *Is het mogelijk deze methode in de praktijk in een organisatie te implementeren? Wat zijn de vereisten, eventuele bezwaren, wat zou er moeten worden aangepast, etc.?*

Vraag: *In hoeverre is het mogelijk de methode daadwerkelijk te gebruiken in het geval van informatiebeveiligingsincidenten? Waar zitten de beperkingen?*

Toegevoegde waarde

Na de verschillende componenten van de methode te hebben behandeld, is het interessant om te weten of experts echt een toegevoegde waarde zien in de methodologie, en waar deze toegevoegde waarde dan in zit.

Vraag: *In hoeverre biedt de classificatie methode toegevoegde waarde t.o.v. bestaande best practices voor information security incident management?*

IT auditing

Het is belangrijk om te weten of de methode gevolgen heeft voor de audit van het incident management processen van organisaties.

Vraag: *In hoeverre heeft de classificatie methode impact op de manier van auditen van incident management processen, met name de beoordeling van de onderdelen identificatie en respons van die processen?*

Stel: een organisatie gebruikt de classificatie methode of één of meer componenten daarvan (classificatie, indicatoren, vooraf gedefinieerde responses) als onderdeel van zijn incident management proces.

Vraag: *In hoeverre heeft dit gevolgen voor de audit van dit proces? Is er een andere aanpak vereist dan normaal?*

Appendix E
SIMPLIFIED RISK QUESTIONNAIRE

This Appendix described the questions used to attain a severity rating for an incident using a simplified approach compared to the SPRINT methods of Appendix B. For each question in the risk questionnaire below, determine whether there is a (L)ow, (M)oderate or (H)igh risk (i.e., probability) of the incident causing the described situation or event.

Risk questionnaire:

	Risk
Competitive disadvantages Can the incident cause competitive disadvantages?	
Direct loss of revenue Can the incident cause direct loss of revenue?	
Direct loss of business Can the incident cause direct loss of business?	
Disruption of business processes Can the incident cause disruption of business processes?	
Public trust Can the incident cause a breach in public trust?	
Legal liability Can the incident cause a breach of legal, contractual or other required guidelines?	
Fraud Can the incident cause embezzlement of goods or money?	
Employee morale Can the incident have a negative impact on employee morale and motivation?	

Additional costs Can the incident cause additional costs to arise?	
Repairs Can the incident cause repairs (e.g., to affected systems) and work backlogs?	
Risk to business	

Determine severity by taking the maximum value of the risks determined from the questions above. For example, if there are various (L)ow and (M)oderate risks, and one (H)igh risk, overall severity is still (H)igh, since this method implicitly contains the reliability need of the organisation. In other words, this approach directly measures possible impact on business continuity. Furthermore, if there are multiple (H)igh risks, the severity may need to be deemed (S)evere, as such an incident may constitute a serious risk to business continuity when it occurs and thus needs immediate action to contain and resolve it.

Appendix F

AUDITING INCIDENT MANAGEMENT

Beheersdoelstelling	**Totaalscore**
Dit proces dient ervoor te zorgen dat er goede communicatie tussen de gebruikers en de IT-organisatie plaatsvindt. Een helpdesk-faciliteit, die eerste-lijns ondersteuning en advies aan gebruikers levert, dient te zijn ingevoerd. Incident Management dient te zorgen voor continuïteit in de dienstverlening door een zo spoedig mogelijk herstel van het afgesproken dienstenniveau, wanneer een afwijking hierop wordt geconstateerd.	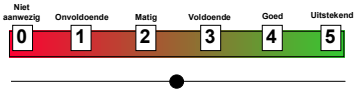

Normen, c.q. eisen	**Referentie**	**Detailscores**					
		0	**1**	**2**	**3**	**4**	**5**
Organisatie							
Bij storingen, vragen, klachten of andere opmerkingen over de IT-dienstverlening dient de melder contact op te kunnen nemen met een helpdesk.	InM 1						
Een heldere, eenduidige procedure dient te voorzien in een vaste werkwijze bij de behandeling van de diverse soorten meldingen.	InM 2						
De openingstijden van de helpdesk dienen aan te sluiten op de behoefte van afnemers van diensten en de afspraken in SLA's.	InM 3						
De helpdesk dient over voldoende personele bezetting met een toereikende opleiding en vaardigheden te beschikken om haar taken te kunnen uitvoeren.	InM 4						
Het incidentmanagement dient in voldoende mate te zijn afgestemd op de processen problem, change en configuration management.	InM 5						
Communicatie en registratie							
Voor een optimale communicatie dient de helpdesk te voldoen aan:	InM 6						
• goede bekendheid bij de melders;							
• goede bereikbaarheid;							
• klantvriendelijkheid;							
• duidelijkheid en begrijpelijkheid bij het verstrekken van informatie.							
De helpdesk dient de beschikking te hebben over voldoende hulpmiddelen voor registratie en communicatie.	InM 7						
Alle meldingen dienen te worden vastgelegd in een registratietool, ook als de helpdesk de melding direct kan afhandelen.	InM 8						
Alle gegevens over de afhandeling van een incident dienen te worden vastgelegd in een registratietool.	InM 9						
Classificatie en toewijzing							
Bij ontvangst van een incident dient dit te worden geclassificeerd naar prioriteit en oplossingsgroep.	InM 10						

De vaststelling van de prioriteit dient plaats te vinden op basis van :	InM 11						
• impact;							
• urgentie;							
• verwachte inspanning.							
Vragen die niet direct kunnen worden beantwoord, dienen te worden gerouteerd naar deskundige oplossingsgroepen (tweede- of derdelijns-support).	InM 11						
Diagnose en oplossing							
Voor herkenning en afhandeling van bekende fouten ('known errors') dient de helpdesk gebruik te maken van de registraties van oplossingen.	InM 12						
Afspraken - over het oplossen van incidenten - met externe specialisten dienen te worden vastgelegd in een (onderhouds)contract.	InM 13						
De helpdesk dient de voortgang van de afhandeling van incidenten op basis van de prioriteiten te bewaken.	InM 14						
De acties die dienen te worden ondernomen ten aanzien van langdurig openstaande incidenten, dienen te zijn vastgelegd in procedures (escalatie, her-routering).	InM 15						
De helpdesk dient de melder volgens een vastgestelde periodiciteit (bijvoorbeeld dagelijks) te informeren over de voortgang van de afhandeling.	InM 16						
Afsluiting							
Indien het incident is opgelost dient dit door de oplosser/specialist te worden gemeld aan de helpdesk.	InM 17						
Bij afsluiting van een incident dient de oorzaak en de oplossing van het incident te worden vastgelegd in een registratietool.	InM 18						
De helpdesk dient de oplossing van het incident met de melder op een begrijpelijke manier te communiceren.	InM 19						
De helpdesk verifieert of het incident daadwerkelijk is opgelost, waarna afsluiting pas plaatsvindt nadat de melder daarmee akkoord is gegaan.	InM 20						
Indien de melder niet akkoord gaat met de oplossing dient een escalatie-/ klachtenprocedure te voorzien in de te ondernemen vervolgstappen.	InM 21						
Rapportage							
Standaard rapportages dienen periodiek inzicht te verschaffen in het functioneren en de kwaliteit van de helpdesk. Deze rapportage dient onder andere de volgende gegevens te bevatten:	InM 22						
• aantal incidenten uitgesplitst naar dienst, impact, et cetera;							
• gemiddelde en maximale storingstijd;							
• verdeling storingstijd;							
• bestede tijd/kosten als gevolg van incidenten door helpdeskmedewerkers, tweedelijns-ondersteuning en toeleveranciers.							
De wijze waarop periodiek wordt gerapporteerd over de status van de gemelde incidenten is vastgelegd in procedures en richtlijnen.	InM 23						

BIBLIOGRAPHY

APPL03 Applegate, L., Austin, R. and McFarlan, F. 2003. Corporate Information Strategy and Management – The Challenges of Managing in a Network Economy. McGraw-Hill.

BABI05 Babiak, J., Butters, J. and Doll, M. 2005. Defending the Digital Frontier – Practical Security for Management, 2nd Edition. John Wiley & Sons / E&Y.

BERG04 Berg, J. van den, and Pijl, G. van der. 2004. Security and ICT Audit 2004/2005, college slides, EUR.

BISH05 Bishop, M. 2005. Introduction to Computer Security. Addison-Wesley.

BLOE06 Bloem, J., Doorn, M. van, and Mittal, P. 2006. Making IT Governance Work in a Sarbanes-Oxley World. John Wiley & Sons.

CERT CERT Coordination Center, CarnegieMellon Software Engineering Institute. www.cert.org.

CSIR03 West-Brown, M., Stikvoort, D., Kossakowski, K., Killcrece, G., Ruefle, R. and Zajicek, M. 2003. Handbook for Computer Security Incident Response Teams (CSIRTs), 2nd edition. Carnegie Mellon SEI.

CSO05 2005 E-Crime Watch Survey – Survey Results. 2005. CSO magazine / U.S. Secret Service / CERT® Coordination Center.

COBI05 COBIT4.0. 2005. The IT Governance Institute, www.itgi.org.

DTIS04 Information security breaches survey 2004. 2004. DTI & PricewaterhouseCoopers.

DTT05 2005 Global Security Survey. 2005. Deloitte Touche Tohmatsu.

EY05 Global Information Security Survey 2005 – Report on the Widening Gap. 2005. Ernst & Young.

EYRAPP Rapportage toetsingskader ITIL-processen. 1999. Ernst & Young.

FBI05 Gordon, A., Loeb, P., Lucyshyn, W. and Richardson, R. 2005. 2005 CSI/FBI Computer Crime and Security Survey. CSI/FBI.

GIBS05 Gibson, C. 2005. Practical IT Management, college slides. MIT Sloan School, Center for Information Systems Research.

GREE03 Greenwell, W., Knight, J. and Strunk, E. 2003. Risk-Based Classification of Incidents. Department of Computer Science, University of Virginia (USA).

ISFS05 The Standard of Good Practice for Information Security. 2005. Information Security Forum (ISF).

ISIM06 Information Security Incident Management – Establishing an Information Security Incident Management capability. 2006. Information Security Forum (ISF), April 2006.

ISTD02 A framework for incident response. 2002. Information Security Team, DePaul University. https://infosec.depaul.edu.

ITIL99 ITIL Security Management. 1999. OGC (formerly CCTA).

MEUL05 Meulstee, A. 2005. Information Security Maturity. Master Thesis Informatica & Economie, Erasmus Universiteit Rotterdam

MOLL05 Mollema. 2005. Postdoctorale opleiding EDP Audit

NIST04 Grance, T., Kent, K. and Kim, B. 2004. NIST Computer Security Incident Handling Guide - Recommendations of the National Institute of Standards and Technology. National Institute of Standards and Technology, USA

OVER00 Overbeek, P. 2000. Informatiebeveiliging onder controle. Pearson Education.

PALM05 Palmgren, K. 2005. Controlling Internal Abuse Through The Process Of Security. SecurityDocs.com, http://www.securitydocs.com/library/2998

PATH05 Pathak, J. 2005. Information Technology Auditing – An Evolving Agenda. Springer Verlag, Berlin, Germany.

PELT05 Peltier, T., Peltier, J. and Blackley, J. 2005. Information Security Fundamentals. CRC Press.

PWC05 The Global State of Information Security 2005. 2005. PriceWaterhouseCoopers / CIO Magazine.

RADC05 Radcliff, D. 2005. After a Security Breach. NetworkWorld, 24-10-2005, www.networkworld.com.

SCDC The Qualitative Paradigm. School of Computing, Dublin City University. http://www.computing.dcu.ie/~hruskin/RM2.htm.

SCHN00 Schneier, B. 2000. Secrets and Lies. John Wiley and Sons

SCMA06 SC Magazine. 2006. Personal info of 26.5 million veterans lost, May 22nd, & VA secretary 'mad as hell' over breach, May 25th. http://www.scmagazine.com.

STEP05 Stephenson, P. 2005. Incident analysis and recovery. Computer Fraud & Security, March 2005.

STOF04 Stofbergen, E. 2004. Informatiebeveiliging op een hoger niveau - Onderzoek naar een methode voor het systematisch opstellen van informatiebeveiligingsbeleid. master thesis Bestuurlijke Informatica, Erasmus Universiteit Rotterdam

VIDA04 Vidalis, S. 2004. A Critical Discussion of Risk and Threat Analysis Methods and Methodologies. School of Computing, University of Glamorgan, Wales, UK.

WIAN05 Wiant, T. 2005. Information security policy's impact on reporting security incidents. Computers & Security Vol. 24, p. 448-459.