A Common Process Model for Incident Response and Computer Forensics

Felix C. Freiling
Laboratory for Dependable Distributed Systems
University of Mannheim, Germany
freiling@informatik.uni-mannheim.de

Bastian Schwittay
Symantec (Deutschland) GmbH, Germany
Bastian_Schwittay@symantec.com

Abstract: Incident Response and Computer Forensics are two areas with similar goals but distinct process models. While in both cases the goal is to investigate computer security incidents and contain their effects, Incident Response focusses more on restoration of normal service and Computer Forensics on the provision of evidence that can be used in a court of law. In this paper we present a common model for both Incident Response and Computer Forensics processes which combines their advantages in a flexible way: It allows for a management oriented approach in digital investigations while retaining the possibility of a rigorous forensics investigation.

1 Introduction

In the field of computer security, new threats such as bot networks and modular malicious code have emerged, while well-known attack tools such as rootkits and trojans remain dangerous by constantly using new and improved techniques. Since not all of these attacks can be prevented, *Incident Response* (IR) has become an important component of IT security management, because it provides procedures for detecting and containing computer security incidents to restore normal computing services as quickly as possible. Many attacks nowadays are already motivated by profit, attempting to achieve financial gain by identity theft, extortion, fraud, or theft of confidential information. With criminal acts becoming more common in computer related incidents, the need for valid evidence for these crimes also rises. The field of *Computer Forensics* (CF) aims at providing such evidence.

IR and CF are two highly related topics, and it is questionable whether a strict separation makes sense at all. Certainly, both investigate a large number of different computer security incidents or offenses, sometimes use identical tools and methods and also share some of the key phases of the investigation, although they may set different priorities. However, the specific view of the investigative process in IR and CF respectively may sometimes be too narrow to achieve optimal results in either case. On the one hand, a CF investigation may lack the proper management that is enforced in an IR investigation, where the investigative efforts are coordinated with all parts of an organization, e.g. legal counsel, Human Resources and business executives. Without these, the "bigger picture" of an incident might not be seen. On the other hand, the scientific standards of a CF investigation can give benefit to an IR procedure, as it promotes objectivity throughout the whole process and a precise and well-documented analysis.

In this paper, we take a unifying view of IR and CF. We take a step back from the two process models of IR and CF and develop a common process model for Incident Response and Computer Forensics which combines the advantages of both models in a flexible way: It allows for a management oriented approach in digital investigations while retaining the possibility of a rigorous CF investigation. At first sight one can argue that this approach seems to be "proposing the obvious" or that no experimental evidence is presented to show the superiority of the new approach. However, we feel that our Common Model (1) structures the phases of a digital investigation in a more logical way, (2) can help unify terminology and methods from different communities, and (3) offers further insight into the strenghts and weaknesses of different approaches. For example, because the large effort to conduct a full-scale forensic investigation may not be necessary in every case, the Common Model will identify certain parameters that determine the extent of an investigation. These parameters are *attacker threat level* and *potential damage* of an incident. We show how to evaluate them in order to develop a sensible response strategy.

The paper is structured as follows: We begin with a brief survey of the processes of Incident Response and Computer Forensics in Sect. 2. In Sect. 3, a Common Model for Incident Response and Computer Forensics is developed, which integrates forensic analysis practices into an Incident Response procedure. We argue in Sect. 4 that the Common Model is flexible enough to handle both IR and CF by providing the necessary criteria to parametrise the Common Model in a flexible way.

2 Background and Definitions

This section briefly recapitulates some important definitions and then gives an overview over IR and CF based on the standard literature by Mandia et al. [5] and Casey [3].

2.1 Definitions

A *computer security incident* can be defined as "a violation or imminent threat of violation of computer security policies, acceptable use policies, or standard security practices" [4]. Incident Response deals with computer security incidents in a well-defined manner to detect incidents, minimize the damage done to the organization, fix the weaknesses that were exploited and return to normal operations.

Computer Forensics is a scientific discipline which is concerned with the collection, analysis and interpretation of digital data connected to a computer security incident; it is sometimes also called *digital forensics*. Since any device, data or other resource subject to a forensic examination may be subsequently used as evidence in a court of law, Computer Forensics puts special emphasis on the correct treatment of potential evidence to prevent it from being altered or tampered with. Because of the same reason, any technique or tool used during an investigation has to meet strict standards, and any conclusions drawn must be scientifically reasonable. It is important to note that an Incident Response procedure can sometimes *include* a full forensic investigation.

2.2 Incident Response

As mentioned in the introduction, Incident Reponse is an organization's reaction to unlawful or unacceptable actions involving a computer or network component. Instead of being caught unprepared and starting a chaotic and possibly devastating response, a systematic and well-organized approach should be used to react. Therefore incidents are usually handled by a so-called *Computer Security Incident Response Team*, or *CSIRT*, which is comprised of personnel with different qualifications that are needed during the response process, in particular people with legal and technical expertise. The CSIRT will then coordinate the appropriate response to an incident, effectively providing a *Computer Security Incident Response Capability (CSIRC)*.

Some of the key benefits that an organized CSIRC offers are for example a confirmation whether an incident actually occured or quick detection, containment and recovery. Achieving all this can seem like an overwhelming task, especially when one considers that often incidents put a lot of pressure on the team responsible for the resolution of the problem, which could for example threaten the very existence of a company. Another characteristic of computer security incidents is their wide variety, which includes, but is not limited to denial of service attacks or theft of confidential information. To be able to manage these different kinds of incidents and their unique complications regarding public law, business operations or even a company's reputation, the process of Incident Response has traditionally been broken down into a number of logical steps. The following model has been proposed in Mandia et al. [5], which is a major reference in the field of Incident Response.

The methodology describes the process of Incident Response using seven different phases, which are also illustrated in Figure 1. We now give a short summary of each step.

2.2.1 Pre-incident Preparation

Pre-incident preparation is an ongoing phase, and it is actually the only one that takes place even before an incident occurs. Its purpose is to prepare the organization and the CSIRT for a possible incident. Preparing the organizion may involve implementing host- and network-based security measures, e.g. an Intrusion Detection System. To make sure that a CSIRT is able to handle an incident well, all hardware and software needed have to be provided. Appropriate training for the members of the CSIRT is mandatory to maintain an effective Incident Response Capability.

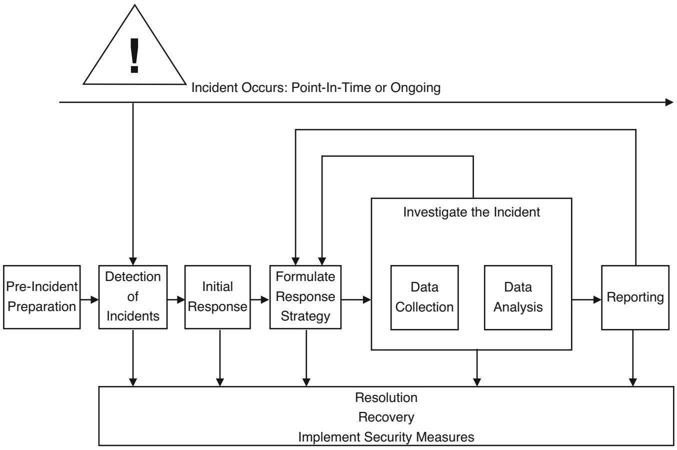

Figure 1: The Incident Response Process [5].

2.2.2 Detection of Incidents

Since obviously no incident can be responded to until it is detected first, detection of incidents is a very important part of IR. Generally speaking, anyone from a system administrator to a regular employee with no technical background could detect an incident. That is why it is important to have clear guidelines for all members of an organization explaining how to react to suspicions that something might be wrong. Most importantly it should be known who should be notified and what details have to be recorded, so that the CSIRT may take over the IR process as quickly as possible to further investigate the potential incident.

2.2.3 Initial Response

During Initial Response, the collection of information concerning the incident that has started in the previous phase continues. The goal is to gather enough information to allow formulation of an adequate response strategy in the next step. Typically, the data that is collected in this step consists of interviews of any persons involved in reporting the suspected incident, and available network surveillance logs or IDS reports, that may indicate that an incident occured. At the end of this phase, the CSIRT should have confirmed that an incident actually occured, should have identified the type of incident and the affected hosts and users, and should be able to assess the potential impact or damage. This will allow to make an informed decision about a viable response strategy.

2.2.4 Formulation of Response Strategy

The goal of this phase is to formulate a response strategy which best fits the situation, thus "considering the totality of the circumstances" [5] that surround the incident. These circumstances include the criticality of the affected systems or data, what kind of attacker is suspected and what the overall damage might amount to. Also, an organization's *response posture*, which defines its policy regarding the response to computer security incidents, may have a large impact on the choice of a response strategy. Because an incident might also induce legal or administrative action, the development of the strategy has to involve people responsible for the respective areas, which means e.g. upper management executives or the Human Resources department.

2.2.5 Investigation of the Incident

During the investigation of the incident, different types of evidence relevant to the incident, e.g. host- or network-based evidence, are collected in order to reconstruct the events that comprise the computer security incident. This reconstruction should provide explanations for what happened, when, how or why it happened and who is responsible. To achieve this, an investigation is typically divided into two steps: *Data Collection* and *Data Analysis*.

Examples of data collected during the Data Collection are host-based information retrieved from a live system, a duplication of a compromised host's harddrive or network surveillance logs. While collecting these information it is advisable to use forensically sound methods; this will be detailed later in this paper.

In the Data Analysis step, the previously collected information is reviewed to uncover the details behind the incident. It should be noted that often the collection and analysis of relevant data may take turns, effectively turning the investigation into a series of feedback loops until the final result is obtained.

2.2.6 Reporting

After the investigation of a computer security incident is finished, all the findings and results have to be documented in a written report. In this report, all investigative activities have to be documented, and all conclusions drawn have to be explained. The report should be written in a concise, yet understandable way, so that even non technical readers can follow. Because the results of the report might be used as evidence during a lawsuit, the report should be able to hold up against legal scrutiny.

2.2.7 Resolution

The purpose of the Resolution phase is to take the right measures to contain an incident, solve the underlying problems that caused the incident and to take care that a similar incident will not occur again. All the necessary steps performed should be taken and their progress supervised to verify that they are effective. It is important that changes to the affected systems are only done after collecting possible evidence, otherwise that evidence might be lost. After the resolution of the incident is complete, it may be necessary to update security policies or the IR procedures, if the response to the incident exposed a weakness in current practices.

2.3 Computer Forensics

Computer Forensics, or Digital Forensics, is a forensic science that deals with obtaining, analyzing and presenting *digital evidence*, which can be defined as “any data stored or transmitted using a computer that support or refute a theory of how an offense occured or that address critical elements of the offense such as intent or alibi” [3]. By employing accepted and proven techniques and principles, which are also applied in other forensic sciences, admissibility in a court of law and credibility of the evidence is achieved. This includes for example a high level of objectivity in every step of an investigation, and the use of reliable, repeatable and well-documented methods throughout the examination.

The Investigative Process Model [3] (see Fig. 2) which will be described now is in fact a more general approach to investigation of a computer-related offense, as it also includes steps which normally tend to fall into the responsibility of law enforcement personnel rather than the forensic analyst. We now briefly describe each step in the process model and already point out some similarities and overlaps with the steps of the IR process model.

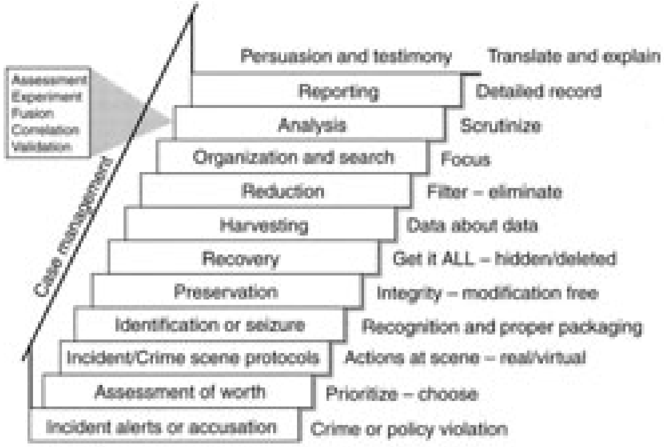

Figure 2: The Investigative Process Model [3].

2.3.1 Incident Alerts or Accusation

The CF process starts with an incident alert or an accusation which could be an automated alert from an IDS or a concerned citizen reporting possible criminal activity. At this point, an initial assessment of the reliability of the source of the alert is done, and some facts surrounding the suspected incident are gathered to get an idea of what the investigator is dealing with. Notice that there is no Pre-Incident Preparation Phase because an investigator is not necessarily working for the organization that reported the incident, but rather contacted when an incident is suspected.

2.3.2 Assessment of Worth

During this phase, it is decided whether a detailed investigation should take place or if the suspicion of an incident that raised the alarm can be discarded. Generally this depends on how severe the problem appears to be, i.e. what the potential damage might be or whether the incident can be contained easily.

2.3.3 Incident/Crime Scene Protocols

If the decision in the previous step is to conduct a full investigation of the incident, all potential evidence at the crime scene has to be secured and documented using accepted protocols and procedures. The goal is to "freeze" the evidence in place and provide a "ground truth for all activities to follow" [3]. Actual analysis is not done at this point.

2.3.4 Identification or Seizure

After securing the crime scene, all items that may contain potential evidence for a suspected incident have to be seized and a *chain of custody* for all evidence must be started.

There are numerous legal implications and restrictions to securing a crime scene and seizing material from a crime scene. Excellent resources regarding these two steps during an investigation are U.S. Department of Justice [8] and UK Association of Chief Police Officers [7], which describe the recovery of digital evidence from law enforcement's point of view. This knowledge can also be valuable to a digital investigator and helping him to work together with the other parties involved in an investigation.

2.3.5 Preservation

After all potential evidence has been seized and is available to the digital investigator, duplicate copies of all data sources are made to ensure integrity of the original evidence. Actual analysis is only done on the copies, while the original evidence is securely stored and cataloged. This step is in fact the first step that is part of the actual digital forensic investigation, and it makes use of trusted and reliable tools to perform the duplication of evidence.

2.3.6 Recovery

The first step of a computer forensic analysis consists of recovering any data that may have been deleted, hidden, encrypted or is otherwise unavailable to the forensic investigator in its current state. The goal is to recover anything possible — "throwing out a large net" — so that the maximum of data is available for further analysis, hoping that it will contain valuable evidence.

2.3.7 Harvesting

During the Harvesting phase, metadata, i.e. data about data, are collected and used to organize the large amount of data that the last step produced. Often, files are grouped according to their filetype, or by their temporal relationship with respect to file timestamps. By slowly starting to make sense of all the available data, the reconstruction of events and the development of different hypotheses about what might have happened starts. The structured data that is the output of this phase represents the first step towards extracting the evidence out of all the data that was collected in the beginning.

2.3.8 Reduction

Since the volume of data under investigation may be very large, it is crucial for a forensic examiner to eliminate any unnecessary data in order to focus on the more important pieces of information. Still, this does not mean that evidence is reviewed based on content, but rather that data is eliminated based on filetype or by using hash databases of known good files. The desired result is "the smallest set of digital information that has the highest potential for containing data of probative value" [3].

2.3.9 Organization and Search

Organization of the data extracted in the Reduction step helps to make the actual analysis easier and can also simplify referencing to specific data during the following phases. A searchable index is often used to allow efficient retrieval of data and thus speed up the analysis.

2.3.10 Analysis

In the Analysis step, all the data that has been prepared in the steps before will be analyzed in detail, this time also incorporating the actual content, e.g. the content of text files. Finding out how and why an offense occured is investigated by establishing links between different pieces of evidence, ultimately trying to identify the offender and offer comprehensive proof that supports the conclusion. With standard scientific methods such as showing correlation between certain events, experimenting with the data and rigorously validating the results, an investigator can present digital evidence that can be relied upon even under high standards such as that of a court of law.

2.3.11 Reporting

The final report should accurately document each step of the investigation, back up any conclusions drawn during the examination with evidence and report on the methods used to obtain the evidence. Whenever possible, accepted and known protocols and methods applied during the analysis should be referenced to increase the credibility of the investigation and its results.

2.3.12 Persuasion and Testimony

Sometimes it may be necessary for an examiner to testify in court or to answer to decision makers regarding the results of his report. Trying to explain the details of an investigation to a mostly non-technical audience can be difficult and therefore special techniques exist that allow an expert to do so in an understandable fashion.

2.4 Comparison between IR and CF

The Incident Response process clearly focusses on management issues and integration of the investigative process into a business environment. For this reason, it has to take into account things such as a possible hit to an organization's reputation and weighing off the monetary cost of an investigation against the potential damage the incident may cause. When proposing a common model for Incident Response and Computer Forensics in Sect. 3, the management issues will be addressed in a similar way as in the Incident Response model described here. The different emphasis is also reflected in the way a response strategy is developed in the IR model, compared to the Assessment of Worth step in the CF model. While in the case of IR, there is a wide variety of possible response strategies, in the case of the Investigative Process the question at this point will simply be if it is necessary to assign additional resources to the case to investigate further, or if the investigation can be stopped altogether.

Another striking discrepancy between the IR and CF model is, that all the steps starting from Preservation up to the Analysis step in the CF process model correspond to just one step in the IR process model, the Investigation phase. This is because during Incident Response, the actual investigation can take on various forms or might also be left out altogether in some situations. A CF investigator on the other hand will always stress the importance of scientific and forensic methods during an investigation in order to obtain evidence that is admissible, i.e. that it can be presented in a court of law. For that reason the analysis steps are clearly seperated and structured in a subtle way to allow a systematic approach to an investigation. It is mandatory to use the scientific principle when drawing conclusions from the evidence, which means that not only has to be reconstructed what has happened, but also it has to be shown why other explanations can be ruled out; this practice is also known as *falsification*. The common model for IR and CF will incorporate these forensic principles and procedures into the investigation step that was described rather loosely in the IR model.

3 Common Process Model for Incident Response and Computer Forensics

We now present a Common Process Model for IR and CF and explain the different steps of this model. More details on the function of these steps, the methods and specific techniques that are used today can be found in Schwittay [6].

3.1 Overview

The Common Process Model for Incident and Computer Forensics is a proposal for a new process model to investigate computer security incidents, and its aim is to combine the two concepts of Incident Response and Computer Forensics to improve the overall process of investigation. In fact the Common Model somewhat resembles a Computer Forensic investigation which is embedded into an Incident Response procedure.

The Common Model consists of three main phases: Pre-Analysis, Analysis and Post-Analysis (see also Fig. 3).

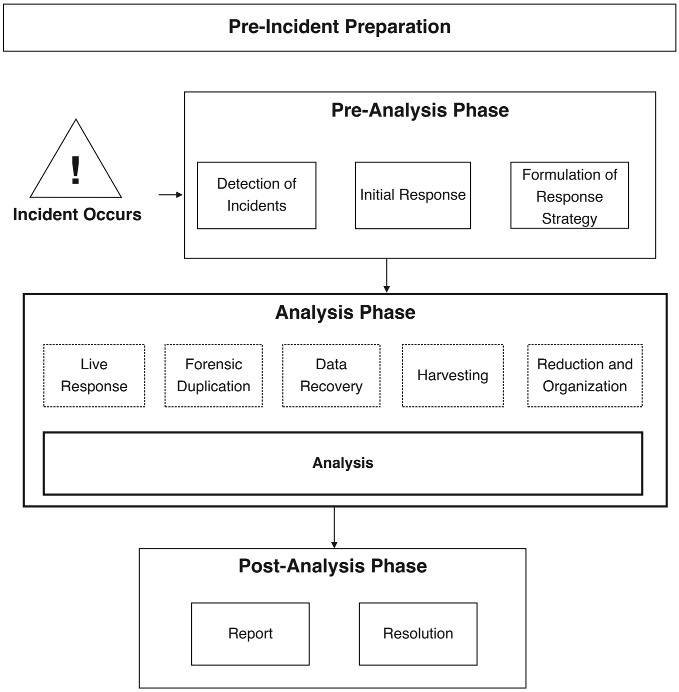

Figure 3: Common Process Model for Incident Reponse and Computer Forensics

Pre-Analysis Phase The *Pre-Analysis Phase* contains all steps and activities that are performed before the actual analysis starts; in this case, analysis means that compromised hosts or data are reviewed in detail with the intention to reconstruct the reason for the computer security incident in question. Using this notion, a quick survey of the affected host during the *Initial Response* step does not qualify as actual analysis. The three steps of the Pre-Analysis Phase correspond to the respective steps in the Incident Response process model introduced earlier, although the last step, *Formulation of Response Strategy*, differs slightly.

Analysis Phase Actual analysis takes place in the *Analysis Phase*, which uses part of the Investigative Process Model's steps presented earlier. The sequence of steps is based on the Computer Forensics process model, except for the step of *Live Response*, which did not appear explicitly in either of the two previous models. In short, Live Response means collecting information about an incident on hosts that are still running, i.e. "live" — as opposed to a "dead" analysis of devices after the host containing it has been powered down.

Post-Analysis Phase Following the Analysis Phase, the *Post-Analysis Phase* is first of all concerned with documentation of the whole activities during the investigation in a written report; both classic models include this step.

Having introduced the Common Model for Incident Response and Computer Forensics, the specific steps of the Common Model will now be described in more detail, pointing out their significance within the whole process. For more details, see Schwittay [6].

3.2 Pre-Analysis Phase

Not surprisingly, the Pre-Analysis Phase (see Fig. 4) is comprised of all steps that are performed before the actual analysis of an incident starts. This includes Pre-incident Preparation, which is different from the rest of the steps because it is an ongoing phase, containing preparation procedures for a possible incident. Incident Detection is concerned with incident detection mechanism and procedures, and Initial Response deals with gathering initial information that allow to choose an appropriate response strategy in the Formulation of Response Strategy step.

3.2.1 Pre-incident Preparation

The goal of Pre-incident Preparation is to enable an organization to handle a computer security incident in a well-defined manner that allows quick and effective resolution.

Usual measures taken during Pre-incident Preparation can be divided into two main categories: those that prepare the personnel that is responsible for responding to incidents, in particular the CSIRT, and those that concern preparation of the organization or environment that an incident may occur in.

Another aspect concerning the preparation of an organization for possible computer security incidents is the definition of proper policies. This includes both a general statement of an organization's response stance, called Response Posture, as well as policies and rules regarding the monitoring of end users or network traffic. It has to be assured, that a proper investigation can take place without breaching privacy rights or violating public law. So-called Acceptable Use Policies define what is considered an acceptable or unacceptable

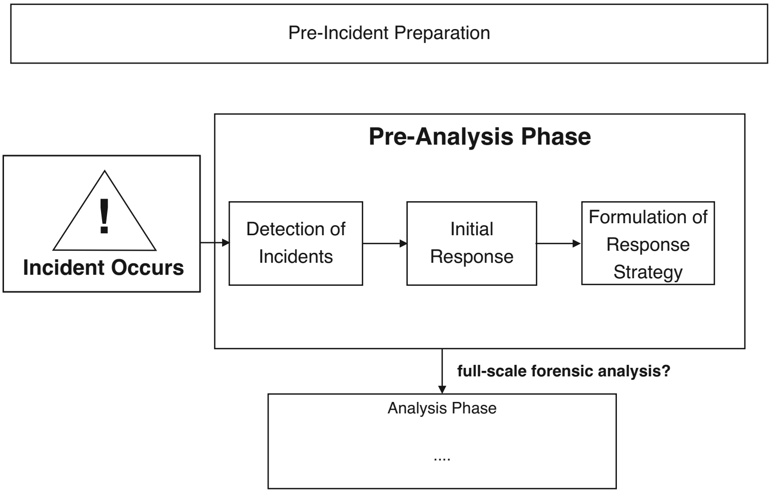

Figure 4: Pre-Analysis Phase of the Common Model

use of computer resources, and thus which behaviour can be considered a computer security incident.

3.2.2 Incident Detection

Incident Detection is about establishing proper incident detection guidelines, allowing for quick detection of computer security incidents. Proper notification and reporting procedures have to be in place, making it possible to transfer control over the remaining investigation to the CSIRT as soon as possible.

Incident Detection normally occurs whenever a person or security mechanism suspects an unauthorized or unlawful action involving a computer system or network. Actually suspicion of an incident can come from a lot of different sources, e.g. end users, security personnel, Intrusion Detection Systems, or system administrators. This is why there have to be clear guidelines for everyone that may detect a possible incident, whom to contact and how to react in such a situation, so that no potential evidence is destroyed. If an incident is reported, some initial information concerning the detection should be recorded, including date and time of detection, who is involved in the incident (by detection or otherwise), what kind of incident is suspected, and/or which hosts/networks are involved. Ideally, appropriate forms for this checklist should be provided, to allow a standardized way of documenting the incident detection. Once completed, the CSIRT should be alerted and informed about the suspected incident, so that they can take control of any further investigation.

3.2.3 Initial Response

During the Initial Response step, the goal is to have the CSIRT confirm a suspected incident or discard the suspicion by collecting and reviewing information related to the incident. If the incident is confirmed, its type and scope should be determined, in order to be able to develop a suitable response strategy in the following step of the Pre-Analsis Phase. Initial Response also includes initializing well-documented containment measures to limit the potential damage of an ongoing incident.

Often network monitoring will be initiated to confirm an ongoing incident and possibly collect additional data. In rare cases, the Initial Response may also include procedures normally taken during *Live Response*; in this case all precautions that apply for Live Response also have to be respected in this earlier part of the process. If necessary, possible harm for an organization can be avoided by containing the incident, e.g. by removing compromised hosts from the network, initializing packet filtering at routers or firewalls or by keeping suspected persons away from the "crime scene". As always all activities have to be documented accurately to maintain well-organized incident handling.

Once an incident has been verified, the information collected is used to estimate the incident's impact on users, systems and business operations. This assessment will then be used to formulate an adequate response strategy.

3.2.4 Formulation of Response Strategy

The goal of the Formulation of Response Strategy step is to determine the most appropriate strategy to handle the incident in question. In particular, the response team will have to decide whether a full-scale forensic analysis of the incident is warranted.

When trying to find an optimal response strategy to deal with a computer security incident, a lot of factors will influence the process of reaching a decision; this is also sometimes called "considering the totality of the circumstances" [5]. Some of the more obvious properties of a specific incident, which have an influence on the adopted response strategy are the criticality of the compromised hosts/data, potential perpetrators, apparent skill level of the attacker, downtime caused by the incident, or monetary loss. These are factors that are directly related to the incident and will influence how many resources are deployed to investigate the case and what kind of approach will be used to tackle the problem. Apart from these factors, there are also a number of more subtle ones that nevertheless can have a large impact on the choice of a response strategy, as for example political considerations (e.g. what happens if the incident becomes public), legal constraints (e.g. liability for not reporting certain incident to law enforcement), business objectives, acceptable use policy, monitoring policy, previous enforcement of policies, etc.

In the Common Model, the Formulation of Response Strategy step is also the time when it has to be decided, whether a full-scale forensic analysis should be conducted. We will describe later in this paper, what criteria is relevant to such a decision.

3.3 Analysis Phase

During the Analysis Phase (Fig. 5), the actual analysis of the compromised computer systems takes place, as outlined in a response strategy that was developed in the previous phase. Starting with Live Response, data concerning the incident are first gathered while the computer systems in question are still running. The rest of the Analysis Phase deals with "dead" analysis of systems, i.e. when they have been powered down.

If the previous phase resulted in opting for a full-scale forensic analysis, then all of the steps of the Analysis Phase have to be performed without taking any shortcuts. If such a thorough investigation is not wanted, some of them may be skipped or shortened; this will be mentioned specifically for each step later on. Notice that some response strategies may not even demand any analysis at all and a response team will go straight to the resolution phases.

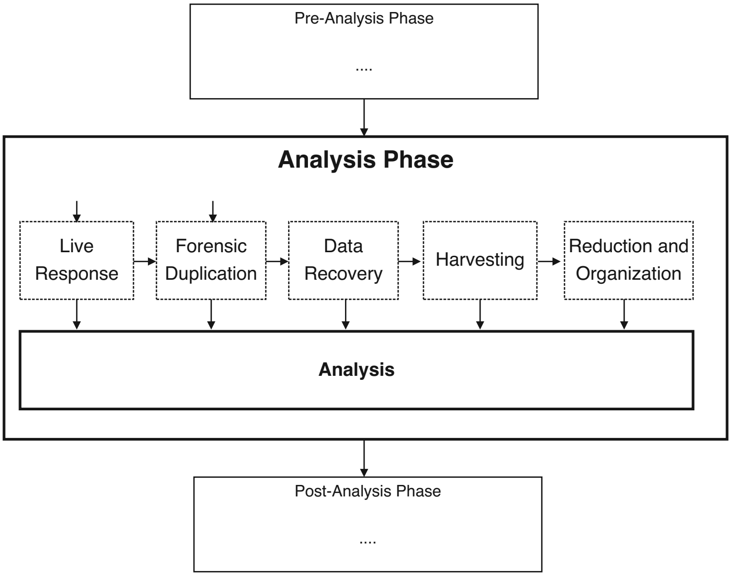

Figure 5: Analysis Phase of the Common Model

3.3.1 Live Response

Live Response is concerned with collecting data from "live" computer systems, that means the systems under analysis are still powered on and running, in contrast to "dead" analysis on systems that have been powered down. The goal is to collect *volatile* data, i.e. data that can not be recovered from a forensic duplication alone after having cut off the power supply. In addition to this it has been common to collect a number of non-volatile information too for convenience. All these data collection activities should modify the running system as little as possible so that the original evidence is not altered more than absolutely necessary.

Certain pieces of data on a running computer are only present in main memory, that means that they will be erased once the computer is switched off. There is no way to reconstruct these information later, including for example the system date and time, list of open network connections etc. In addition to collecting these volatile data, it has become common to collect non-volatile information from running systems, usually out of convenience. Non-volatile data could also be recovered during a standard dead analysis, but collecting it from a live host may have certain advantages. For example, if a rootkit is installed on an evidence system, the output of Live Response tools may be manipulated to hide certain information. So checking for rootkits should can quickly validate or invalidate the corresponding hypothesis.

The combination of volatile and non-volatile information collected in Live Response can often contain valuable clues for an investigator to estimate the extent of an incident and take measures to contain it.

3.3.2 Forensic Duplication

In the Forensic Duplication step, exact copies of all storage media that are involved in the incident are generated, while the original evidence has to stay unaltered. A chain of custody for all media is started and the original media are stored in a safe place along with the duplicates.

3.3.3 Data Recovery

Working on the output of the Forensic Duplication step, mostly disk image files, Data Recovery is concerned with the recovery of data that in the current state of the image is not available for analysis. This includes recovery of deleted, damaged, hidden, or otherwise inaccessible data on a file system image, as well as uncovering data that is hidden otherwise, e.g. in unallocated space. For more details, see Carrier [2], which is also a major reference for forensic file system analysis.

3.3.4 Harvesting

During Harvesting, the investigator begins to gather *metadata* (data about data, like file names, file types, file sizes etc.) about the preserved material which is the output of the Data Recovery step. This allows to structure the largely unorganized material based on certain criteria, typically file timestamps, permissions and other file attributes, and the file type. This structure will help during the rest of the analysis, because often sets of data with certain properties "seem or are known to be related to the major facts of the case or incident known to this point in the investigation" [3].

When investigating files, the file type information can then be used to group files of a certain type or class, e.g. when investigating a case of suspected distribution of child pornography one would focus on image and video files, whereas in case a computer intrusion with a suspected rootkit installation the investigator would rather focus on executables or driver files.

3.3.5 Reduction and Organization

During Reduction and Organization all data that can be identified as irrelevant to the case is eliminated, and the remaining data is organized to alleviate access to the data later. The goal is to reduce the large set of structured data that was the output of the previous step to "the smallest set of digital information that has the highest potential for containing data or probative value" [3], and then to organize the data to allow efficient search, identification and reference to relevant data in the Analysis step and the whole Post-Analysis Phase.

3.3.6 Analysis

After having recovered, harvested, reduced and organized the data related to an incident in the previous steps of the Analysis Phase, the actual reasoned analysis start in the Analysis step. An investigator develops a detailed reconstruction of the events that comprise the incident and tries to answer the questions of what happened, when and how did it happen, and who is responsible. During an investigation of suspected criminal activity, the perpetrator has to be identified, along with determining the means, motivation and opportunity. By reviewing the actual contents of the data, different pieces of evidence are correlated to establish links between them. By carefully documenting the activities and validating the results, this results in a thorough reconstruction of the incident based on objectivity and scientific principles.

To guarantee an objective analysis and interpretation of the results, the scientific method has to be applied, that means an investigator should test multiple theories regarding an incident and try to disprove them, instead of trying to verify them. By eliminating theories that conflict with intermediate results, a remaining theory has a high possibility to represent a correct reconstruction of the events.

Another very important property of analysis results is that they are *repeatable*, meaning that any observer could make the same observations as the investigator using the same methods. To allow a smooth and accurate Post-Analysis Phase, every action taken, technique applied or tools used should be documented precisely and immediately, along with its results and impact on the considered theories.

3.4 Post-Analysis Phase

The Post-Analysis Phase (see Fig. 6) starts when all activities regarding the collection and analysis of digital evidence have ended, and the objectives set by the response strategy for the Analysis Phase have been fulfilled.

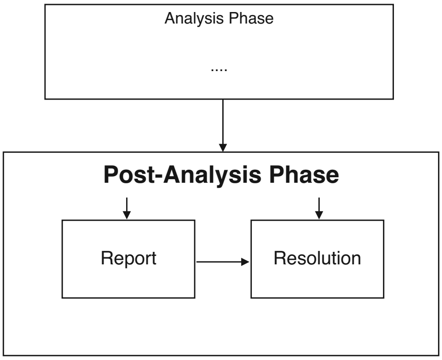

Figure 6: Post-Analysis Phase of the Common Model

3.4.1 Report

The Report step deals with writing a precise report that describes the details of an incident, is understandable to non-technical readers or executives and meets the legal standard for admissibility in court. It contains all of the documentation material that (hopefully) was produced during the Pre-Analysis and Analysis Phase respectively, and assembles the different pieces into a comprehensive overview of the whole case, while pointing out the most important analysis results and implications for resolution of the incident.

3.4.2 Resolution

The goal of Resolution is to contain the problem that caused the incident, solve it and take appropriate measures to keep it from occurring again in the future. In case of an ongoing incident, these measures may also be taken prior to conclusion of the analysis, but if possible they should only be put in place after all potential evidence has been secured for further analysis, because the evidence will generally be altered by the implemented security measures. To verify that the chosen security mechanism have been integrated successfully and that they are effective, the implementation should be supervised and the effectiveness of the measures should be validated in the future.

3.5 Discussion

The Common Model for Incident Response and Computer Forensics offers a way to conduct proper Incident Response while applying principles known from Computer Forensics during the actual analysis phase. In this way it addresses the shortcomings of both seperate models, effectively integrating a forensic analysis into an Incident Response framework. This model makes use of the management procedures developed in Incident Response process models [5], which mainly concern the Pre-Analysis and Post-Analysis phases. Typically these measure like Pre-Incident Preparation and Incident Detection are not included in traditional investigative models, since those tend to approach the problem from a law enforcement's point of view.

4 Unifying IR and CF

After a computer security incident has been detected and some initial information concerning the incident has been collected, a suitable response strategy has to be chosen; this step has been called *Formulation of Response Strategy* above. In the case of the Incident Response process model, this simply means that after assessing the potential damage, criticality of the affected hosts and data and taking into account an organization's response posture, a strategy is formulated which seems to be the optimal way to resolve the incident. It has to be kept in mind, that a viable response strategy might in some cases be to just reinstall a compromised hosts's operating system, apply all patches and then resume normal operations again. Also, such a response strategy does not have to respect forensic principles, it may be perfectly adequate to disregard them altogether, depending on the particular incident.

In the Common Model, Formulation of Response Strategy includes an additional decision regarding the further investigation. In this step, an investigator or team of investigator has to decide, whether a *full-scale forensic analysis* is necessary.

It is clear that this type of investigation in general requires a lot of resources, because due diligence has to be invested for the analysis, and it may even cause downtime of its own, because hard drives have to be imaged, systems might have to be isolated or taken to the lab, and so on. This is why the decision to go for the full-scale forensic analysis should not be made lightly, but it should be made if the details of the incident in question call for it. The Common Model provides a comprehensive way to make the decision whether a full-scale forensic analysis is warranted or whether other response strategies may be more appropriate.

4.1 Determining factors

When deliberating the particular question of whether to conduct a full-scale forensic investigation, the Common Model focusses on two main factors, which will also be called "soft" factors, because they can only be estimated and depend largely on the properties of each particular incident:

1. attacker threat level
2. potential damage

The *attacker threat level* denotes the estimated threat of an attacker, which for this matter is any person who is directly responsible for the incident in question, whether it be by actively attacking a computer system or for example by violation of a security policy.

Such a threat measure can be derived from a more detailed *attacker model*, which is a model of the attacker with his skills and intentions, thus allowing to assess the threat he poses for the organization under attack. In the Common Model, a high attacker threat level favors a full-scale forensic analysis, because a highly skilled attacker is likely to cover his tracks or even plant false clues. It may also be desirable to completely reconstruct the incident if a skilled attacker is involved, in order to find out what exactly the attacker did and how he did it; only a thorough forensic analysis can fulfill this requirement.

The *potential damage* of the incident under investigation is the second key factor to consider when considering the type of analysis to go for. Damage in this case means financial loss as well as loss of reputation or credibility, and since the actual damage of in incident is generally not known exactly in such an early state of the response process, it can only be approximated. Generally, a high potential damage will favor the decision for a full-scale forensic investigation, because in these cases, criminal prosecution or filing for civil complaint is a likely outcome. Hence, reliable and admissible evidence may be needed to fully resolve the incident.

With these two factors in mind, an abstract equation can be used to illustrate the decision:

$$AttackerThreatLevel \times PotentialDamage > Threshold \tag{1}$$

If this equation evaluates to true, a full-scale forensic analysis should be conducted. That means that the combination of a skilled attacker and high potential damage will influence the response strategy in a way that favors a forensic analysis, while in the case of an unskilled attacker or low potential damage, such an analysis is not justifiable.

4.2 Constraints

Apart from the above soft factors, there are also two “hard” factors which are mostly independent from each other and from the soft factors, and which can call for a full-scale forensic investigation: an organization’s Response Posture and legal constraints.

A *Response Posture* describes an organization’s general stance towards responding to incidents. It contains general guidelines as well as prefered response procedures for a specific type of incident, e.g. what has to be done if confidential customer data has been compromised, or the organization’s web site has been defaced. It may also include recommendations regarding administrative actions should an employee violate a corporate policy, e.g. how to deal with theft of trade secrets or illicit internet use.

In some cases, a full-scale forensic analysis may be demanded explicitly by the Response Posture, for example because the organization pursues a “zero-tolerance” policy towards computer security incidents and will always try to get to the bottom of each case. In other cases a forensic investigation might be necessary because of the likely outcome of the investigation, in particular whenever the resolution could include administrative or legal action. In these cases admissible evidence is required, regardless of the attacker threat level or the potential damage of then incident, therefore it is a hard factor that can indicate the need for a full-scale forensic analysis.

Legal constraints can be another hard factor that will demand a full-scale forensic analysis of the incident. In some cases, an organization will choose to involve law enforcement and make a report against a suspected attacker. It may even be mandatory to report a suspected crime, e.g. the failure to notify the authorities about suspected possesion of child pornography may make an organization liable [5]. There have also been cases where the failure to objectively analyze an incident has led to negligence charges against the investigators [3]. Because of these possible complications, a well documented and objective approach to analysis of such incidents is necessary; that means a full-scale forensic analysis has to be conducted.

4.3 Discussion

Equation 1 is related to the well-known *Risk Equation* which is often used to define risk [1]:

$$Risk = Threat \times Vulnerability \times Cost \tag{2}$$

The Threat value corresponds approximately to the Attacker Threat Level and the Cost value corresponds to Potential Damage. Vulnerability normally denotes the likelyhood that a particular attacks succeeds, which in the case of Incident Response has in fact already happened; therefore this factor could be set equal to 1, and the two equations would draw even nearer to each other. In this sense, the value that determines whether a full-scale forensic analysis is necessary equals Risk under the condition that a vulnerability has already been exploited.

Acknowledgments

We wish to thank Maximillian Dornseif for helpful discussions.

References

[1] The Risk Equation. `http://www.icharter.org/articles/risk_equation.html`.

[2] Brian Carrier. *File System Forensic Analysis.* Addison-Wesley, 2005.

[3] Eoghan Casey. *Digital Evidence and Computer Crime - 2nd Edition.* Academic Press, 2004.

[4] Tim Grance, Karen Kent, and Brian Kim. *Computer Security Incident Handling Guide.* National Institute of Standards and Technology, 2004.

[5] Kevin Mandia, Chris Prosise, and Matt Pepe. *Incident Response & Computer Forensics - 2nd Edition.* McGraw-Hill, 2003.

[6] Bastian Schwittay. Towards automating analysis in computer forensics. Master's thesis, RWTH Aachen University, Department of Computer Science, 2006.

[7] UK Association of Chief Police Officers. *Good Practice Guide for Computer based Eletronic Evidence*. National Hi-Tech Crime Unit, 2003.

[8] U.S. Department of Justice. *Electronic Crime Scene Investigation: A Guide for First Responders*. National Institute of Justice, 2001.