Proposing a Digital Operational Forensic Investigation Process

M.A. Bihina Bella[1], M.S. Olivier[1] and J.H.P. Eloff[1, 2]

[1]ICSA Research Lab, Department of Computer Science, University of Pretoria, Pretoria, South Africa
[2]SAP Research Pretoria, Pretoria, South Africa
e-mail: mbihina@yahoo.fr

Abstract

The increasing complexity of IT systems can lead to operational failures with disastrous consequences. In order to correct and prevent the recurrence of such failures, a thorough post-mortem investigation is required to localise their root causes. However, the currently used troubleshooting approach fails to provide sound analysis of these causes. A promising alternative approach is the emerging field of operational forensics, which applies digital forensic techniques to failure analysis with a view to improve the faulty system. This paper proposes a process for an operational forensic investigation, and shows how the process could be applied to a real-life IT failure to provide the correct diagnosis of the problem quicker and with more accuracy than troubleshooting. It also revisits the current definition of operational forensics in order to make it more specific.

Keywords

Troubleshooting, operational forensics, digital forensics, forensic science, root cause analysis, failure analysis

1. Introduction

IT systems are getting more and more complex due to the advancement of technology and customer demands for innovative products to simplify or enhance their daily activities. This is evident with the increasing demand for convergent IP-based Next-Generation Network (NGN) services and mobile applications (Bihina Bella *et al*. 2009). The NGN services are developed from the integration of various applications from a range of vendors, adding complexity to the resulting products. This creates more security loopholes and potential malfunctions once the system is in operation, even though thorough testing has been performed during its development (Bihina Bella *et al*. 2009).

As incidents often originate from a sequence or combination of events and not a single action (Noon, 1992), identifying the root causes of a failure can be challenging and even more so in such composite systems. This may lead to longer system downtime and lost revenue (Trigg & Doulis, 2008), as well as various litigation issues over the liability of the different parties involved. Without an understanding of the underlying cause of the problem, preventing its reoccurrence will be hampered.

Currently, general system failures are usually handled through troubleshooting. A failure can also be addressed by an incident response program although this typically deals with security related incidents (Jordan, 2008), which is not the focus of this paper. Troubleshooting relies heavily on the investigator's experience with the target system and focuses primarily on restoring it to its operational state as quickly as possible (Turner, 2007). In so doing, valuable information that could pinpoint the root cause of the problem is often lost during rebooting and little time is given to a proper investigation (Trigg & Doulis, 2008).

However, various regulations, standards and best practices advocate for a root cause analysis of IT incidents (Stephenson, 2004). The following are examples of such recommendations.

- In Principle 14 of its section on Risk Management for Electronic Banking, the Basel Committee on Banking Supervision clearly specifies that banks should establish "a process for collecting and preserving forensic evidence to facilitate appropriate post-mortem reviews of any e-banking incidents as well as to assist in the prosecution of attackers" (Basel Committee on Banking Supervision, 2003).

- The ISO/IEC 27002 information security standard also recommends the post-incident collection and analysis of forensic evidence for future improvements in Part 13, control 13.2: "Responsibilities and procedures are required to manage incidents consistently and effectively, to implement continuous improvement (learning the lessons), and to collect forensic evidence" (ISO/IEC, 2007).

- The American National Institute of Standards and Technology (NIST) also stipulates that an incident post-mortem be conducted for improvement purposes. They mandate organisations to "emphasize the importance of incident detection and analysis throughout the organization" and to "use the lessons learned process to gain value from incidents" in their Incident Handling Guide (Scarfone *et al*. 2008).

Despite these strong regulatory requirements, root cause analysis is poorly addressed in the IT industry (Trigg & Doulis, 2008; Shedden *et al.* 2010). The absence of an in-depth investigation of the root cause of a significant system's failure can have disastrous consequences. This was the case with the Therac-25, a computer-controlled cancer radiation therapy machine. The poor initial investigation of the machine's malfunction through informal troubleshooting carries at least part of the blame for the series of massive overdoses of radiation which killed several patients (Leveson and Turner, 2002).

Operational forensics has the potential to prevent such disasters by providing a sound methodical approach to failure analysis. Corby (2000a), who coined the term *Operational Forensics*, defines it as "the application of digital forensic techniques to the identification of occurrences and underlying causes of observed computer based events." Note that operational forensics applies to events that occur once a system is

in production, thus after the design, development and testing phases. Unlike troubleshooting, operational forensics is based on objective scientific analysis methods, which increases the reliability of its results and makes the process repeatable.

After considering the issues above-mentioned in more detail, we provide our own definition of Operational Forensics in Section 4 as the following: *the application of scientific methods and legal principles to failure analysis of IT systems.*

Since the term Operational Forensics was coined in 2000 (Corby, 2000a), little research has been done in the field and it has remained at a conceptual level with no clearly defined end-to-end process.

The purpose of this paper is two-fold. It proposes a more specific definition of operational forensics and presents a process for an operational forensic investigation as this does not yet exist.

The remainder of the paper is organised as follows: Section 2 provides background on previous research on operational forensics. Section 3 presents how the forensic approach is applied for failure analysis in other fields. Section 4 defines operational forensics based on its two main component fields: digital forensics and troubleshooting. It explains its commonalities and differences to these fields. Section 5 describes the proposed process to be followed during an operational forensic investigation. We apply this process to a case study of a real-life system failure in Section 6, the Therac-25 accidents mentioned earlier. The paper ends with a conclusion in Section 7.

2. Review of previous work on operational forensics

Operational forensics is an emerging field with little available literature. To the best of our knowledge, only the following two authors have addressed it in formal publications: Michael Corby (Corby, 2000b) and Barry Hodd (Hodd, 2010).

The first publication on operational forensics from Corby (2000a) defines the field and scope of this new discipline. His second publication (Corby, 2000b) presents a pre-incident operational forensic program to ensure that potential evidence is preserved and admissible in court. However, it does not indicate how to identify the root cause of the incident nor does it present a process for the investigation.

Hodd (2010), which references Corby's seminal work (Corby, 2000a), investigates the use of modeling tools such as Petri Nets for operational forensic analysis. His research is solely on crime investigation, mainly cases of insider threats and social engineering, while the present research focuses on general system failures. While his paper focuses on the formal modeling of a failure it does not indicate the investigation process.

In addition to the above research, the concept of applying digital forensic techniques to the investigation of computer failures has been explored before but mainly in the area of incident response (i.e. only for cases of security incidents). The possibility to extend it to more general cases of system failures is usually an afterthought (Kent *et al*. 2006; Turner, 2007).

Stephenson (2003) proposes a digital investigation process for security incidents based on the Digital Forensic Research Workshop (DFRWS) investigation framework (Palmer, 2001). The process, called EEDI (End-to-End Digital Investigation), is used to create narratives of the planned investigation steps. The narratives are then translated into the Digital Investigation Process Language (DIPL), which creates a structured model of the investigation process. The DIPL model can be used to simulate the possible outcome of the investigation. Although the EDDI could be applied to non-security related events, its scope is limited to the collection and analysis of evidence, rather than to the entire investigation which must include preservation of the evidence, presentation of the findings, and recommendations for improvements.

The NIST published a guide on how to use digital forensic techniques to assist in incident handling and troubleshooting (Kent *et al*. 2006). The guide explains how to establish a forensic capability but does not provide a process for the investigation nor does it mention the concept of operational forensics.

Turner (2007) unites the forensic approach with the incident response procedure through a Digital Evidence Bag to preserve digital evidence used in the investigation. It highlights the benefits of such an approach but, like Kent *et al* (2006), maintains the distinction between forensics and incident response.

Although it remains essentially conceptual and is not yet operational in the IT industry, the usage of forensics is wide spread in the domain of failure analysis and improvement in other industries from which valuable lessons can be learnt. This is the topic of the next section.

3. Lessons learned from other forensic disciplines

Although the field of operational forensics in the IT industry is still new with limited research available, the application of forensic science to investigations for improvement purposes in other industries is not a new concept. It is, for example, a professional field of practice in the engineering industry under the name of *forensic engineering* (Noon, 2001). Some specialised areas of forensic engineering even exist as disciplines of their own such as *forensic structural engineering* for failed structures (e.g. buildings, bridges) (Ratay, 2010), or *tire forensics* for tire failures (Gioppani, 2008). In each case, the goal is to improve product quality and limit litigation which is also the intention for operational forensics.

The emergence of formal failure investigations as currently conducted in forensic engineering can be traced to the Industrial Revolution during which many complex

machines were introduced. The added complexity led to many accidents that required expert analysis to understand their causes and prevent their reoccurrence. The types of accidents evolved as engineering products developed: from steamboats, railway trains and steel bridges in the 1800's to automobiles, home appliances and airplanes in the 1900's (Brown *et al*. 2003). Forensic engineering is now applied to significant failures of any engineering product; just as operational forensics could be used for any type of major IT system malfunction. As stated, the increasing complexity of IT systems also requires more formal investigations than are currently available with troubleshooting.

Forensic engineering has successfully demonstrated its effectiveness by applying various scientific examination tools and techniques, simulations and event reconstruction methods to identify the source of failure in numerous engineering disasters such as the Challenger Space Shuttle accident in 1986 (Rogers, 1986) and the Columbia Space Shuttle tragedy in 2003 (McDanels, 2006). Technical as well as organisational failures were found to be responsible for each accident which could not have been identified without a thorough forensic investigation.

Forensic engineering has evolved from its initial focus on legal investigations in product liability cases to its current focus on failure analysis for product and system quality improvement purposes. Presently, most forensic engineering investigations do not reach the courtroom and are done mainly with a view to prevent similar future accidents (Carper, 2000). One example is the forensic investigation of the September 11, 2001 World Trade Center collapse which was undertaken to understand the impact of the fire on the collapse of the twin towers although the responsible parties were already known (Usmani *et al*. 2003). Various elements such as the construction design, fire properties of materials used, and their thermal expansion were examined through simulations and computer-based structural analysis. Based on this analysis it was determined that using reinforced concrete instead of lightweight steel as well as providing an energy absorbing structure could prevent the collapse of such tall buildings in the future (Zhou, 2004).

In this regard, forensic engineering is comparable to forensic pathology, another application of forensic science to the improvement of a field's procedures and products. Forensic pathology is a branch of medicine that investigates the cause of death upon a legal request. However, when applied to public health and safety, it is used for the prevention and control of diseases. For instance, a forensic autopsy may uncover a previously undetected contagious disease to prevent an outbreak or pandemic. It may also help identify an hereditary condition that will enable family members proactively to seek treatment (Dolinak *et al*. 2005).

In summary, as demonstrated for many years in the engineering and medical fields, through the integration of scientific methods and legal principles, the forensic approach has ensured objective, comprehensive investigations producing reliable results and has created opportunities for improving product quality.

In light of the above benefits in other fields of practice, we can expect similar benefits for the IT industry based on the application of operational forensics. The next section defines the scope of operational forensics based on its relation to its two components: digital forensics and troubleshooting.

4. Link between digital forensics, operational forensics and troubleshooting

In order to define operational forensics as a new field, it helps to relate it to existing fields with which it shares certain elements. As its definition suggests, operational forensics uses digital forensic techniques to analyse the cause of an event. It thus contains elements of both digital forensics and troubleshooting. Section 4.1 presents the link between operational forensics and digital forensics and Section 4.2 explains the link between operational forensics and troubleshooting.

4.1. Link between digital forensics and operational forensics

The Digital Forensics Research Workshop defines digital forensics as "the use of scientifically derived and proven methods toward the preservation, validation, identification, analysis, interpretation, documentation and presentation of digital evidence derived from digital sources for the purpose of facilitating or furthering the reconstruction of events found to be criminal, or helping to anticipate unauthorized actions shown to be disruptive to planned operations" (Palmer, 2001).

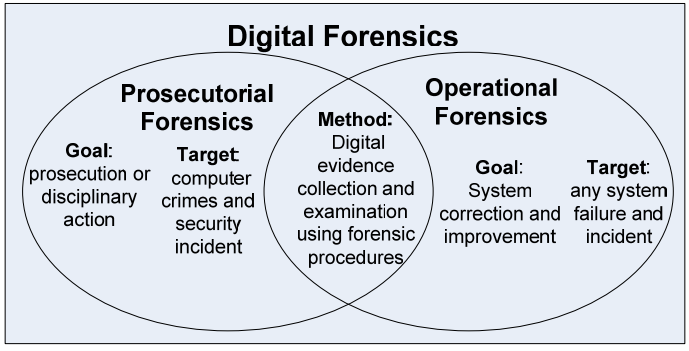

Figure 1: Relationship between prosecutorial forensics and operational forensics

According to Corby (2000b), digital forensics can be categorised into two branches based on the goal of the investigation: operational forensics and prosecutorial forensics. Traditionally, digital forensics has been prosecutorial by nature with the objective of collecting evidence for prosecution or disciplinary action. By contrast, the main goal of operational forensics is to gather evidence for system correction and improvement (Hodd, 2010). Unlike prosecutorial forensics, which only deals with computer crimes and security incidents, operational forensics handles any kind of

computer event. It can be argued that in prosecutorial forensics, the stress is placed on the legal aspect of the investigation, while in operational forensics the emphasis is on the scientific approach of the analysis. We have summarised the distinctions between the two branches of digital forensics in Figure 1.

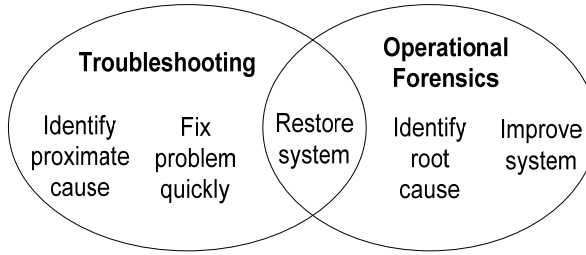

Figure 2: Relationship between operational forensics and troubleshooting

Figure 1 shows the differences between the two branches and highlights their commonality, which is the collection and examination of digital evidence using forensic procedures. A more detailed comparison of operational forensics and prosecutorial forensics is provided in Table 1. The table shows our analysis of the two fields and presents the differences that have a direct impact on the outcome of the investigation.

Operational forensics	Prosecutorial forensics
Similarities	
Digital evidence collection and examination	
Use sound forensic process and techniques	
Evidence admissible in court	
Differences	
Applicable in any computer-based event	Applicable only in computer crimes and security incidents
Used for system improvement and correction	Used for prosecution and disciplinary action
Proactive collection of evidence (mostly "live" forensics)	Reactive collection of evidence (mostly "dead" forensics)
System usually remains operational during the investigation	System is frozen during the investigation
Investigation sequence affected by internal factors (the system or company)	Investigation sequence affected by external factors (e.g. trial)
Investigation can be extended to other domains relevant for improvement (e.g. organisation culture and management)	Investigation focused only on the crime scene
Finding (root cause identification) and decision (recommendation for improvement) are an integral part of the investigation.	Identifying and sanctioning the perpetrator are not part of the process, but are determined in trial outside the digital investigation
Conclusion drawn by investigator	Conclusion drawn by judge

Table1: Operational forensics versus prosecutorial forensics

4.2. Link between operational forensics and troubleshooting

Troubleshooting is a logical search for the source of a problem in order to fix it so the system can immediately resume working. Isolating the cause of the problem typically involves a process of elimination which starts with the most visible or easiest problem to fix depending on the investigator's experience (TechTerm.com, 2011). Operational forensics, meanwhile, focuses on improving the system so the failure does not reoccur, and uses scientific forensic techniques to identify the origin of the problem. In order to prevent the reoccurrence of the issue, its root cause must be identified. Troubleshooting can be satisfied with a proximate cause as long as it helps solve the problem at hand.

These distinctions are summarised in Figure 2. The diagram shows from left to right the main stages of an investigation in both fields. Troubleshooting starts by identifying a proximate cause of the problem based on the investigator's suspicion and ends with a system restoration, while operational forensics first restores the system, identifies its root cause based on an analysis of the acquired digital evidence, and ends with recommendations to improve the system.

As shown in Figure 2, the common denominator between the two fields is to restore the system to its working state. A more detailed comparison of operational forensics and troubleshooting is provided in Table 2. The table shows our analysis of the two fields and presents the differences that have a direct impact on the outcome of the investigation.

Operational forensics	Troubleshooting
Similarities	
Used to solve the problem at hand	
Includes system restoration	
Differences	
Find root cause (s)	Find proximate cause (s)
Relies on scientific forensic analysis	Relies on investigator's experience with the target system
Focus on improving the system	Focus on restoring the system
Includes a formal post-event investigation	No formal post-event investigation
Process is repeatable	Process is not repeatable, case by case
Evidence collected and documented before, during and after the event in a planned manner	No evidence collected

Table 2: Operational forensics versus troubleshooting

Based on the above analysis, we can establish that operational forensics is an interdisciplinary discipline which lies between the fields of troubleshooting and prosecutorial forensics. It is a new area of digital forensics aimed at addressing the limitations of conventional troubleshooting with regard to complex IT systems. This

is illustrated in Figure 3, which only shows the main connections between operational forensics and the other two disciplines.

As we recall from Section 1, our own definition of operational forensics, which stems from the above mentioned interrelation, is as follows: *the application of scientific methods and legal principles to failure analysis of IT systems.*

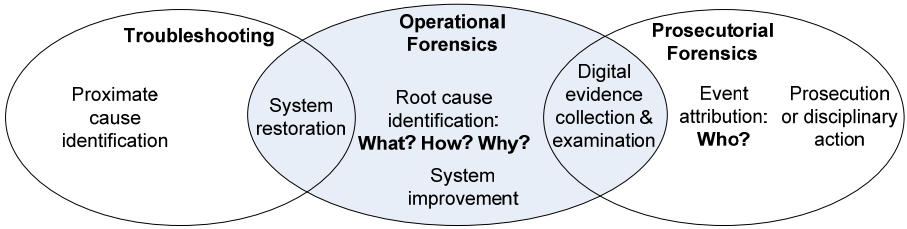

Figure 3: Link between operational forensics, troubleshooting and prosecutorial forensics

Like in other fields, an operational forensic investigation follows a specific process, which is described in the next section.

5. The operational forensic investigation process

As demonstrated in previous sections, operational forensics is composed of digital forensics and troubleshooting and can be compared to forensic engineering. It is therefore expected that an operational forensic investigation combines elements of those three fields. Operational forensics has two facets: the forensic preparation and the investigation. Section 5.1 presents the forensic preparation and Section 5.2 describes our suggested high-level investigative process.

5.1. The operational forensic preparation

In order to maximise the effectiveness and speed of the investigation, a forensic capability needs to be established in the organisation prior to the investigation. The organisation must be "forensic ready" by taking the following actions: (a) equip personnel with necessary forensic skills; (b) identify, acquire and maintain potential evidence such as log files and (c) develop supportive policies and procedures (Corby, 2000a; Kent *et al.* 2006). The organisation must also ensure that all system documentation is available and up-to-date. This includes system specifications, user manuals, licensing information, test plans, and a history of changes and reported incidents (Trigg & Doulis, 2008). This operational forensic program ensures that when a problem occurs, information that can be used as evidence during the investigation is readily available and the responsible parties know how to collect it in a forensically sound manner.

5.2. The operational forensic investigation process

The proposed investigative process consists of three basic stages. The first two occur during the event or immediately after it has been detected: firstly, collect evidence and secondly, restore the system. The third phase is the root cause analysis which is conducted once the system has been restored.

Phase 1: **Information collection**

This phase corresponds to the first step of a digital forensic investigation. Shortly after a failure has been detected, all information that can assist in the investigation needs to be collected in a forensically sound manner by maintaining the chain of custody and preserving its integrity. For the purpose of this paper, we classify the information to be collected as either primary or secondary. The primary information is the electronic data that can serve as potential evidence (e.g. audit logs, network and system configuration settings) while the secondary information is background information regarding the system and the issue at hand. This includes the documentation indicated previously in the forensic readiness program as well as a recording of the state of the scene (e.g. screenshot of the error message) and interviews with the system administrator and the users who reported the failure (Kent *et al*. 2006).

Phase 2: System restoration

Once all relevant information has been acquired, the problem is fixed and the system is restored to its operational state as quickly as possible. This reduces its downtime, which limits any associated negative consequence such as financial loss. A restoration might be as simple as rebooting the system or it might necessitate some preliminary diagnostic of the failure to fix it. This will follow a typical troubleshooting process, which requires a recreation of the problem to isolate its cause (Juniper Networks, 2011).

Phase 3: Root cause analysis

An operational forensic investigation is a failure analysis of an IT system. It follows closely the process of a forensic engineering investigation, which includes laboratory examination, simulations and reconstruction of the incident to determine its root cause (Brown, 1995). These steps are also applicable to the analysis of an IT event and are thus part of our process. The electronic data collected in the 1st phase of the investigation is examined in a computer lab in conjunction with the secondary information to understand the failure. The examination follows the scientific method, which consists of formulating hypotheses for all the probable causes of the failure and predicting and testing evidence for each hypothesis. The root cause is the hypothesis that accounts for most of the evidence (Noon, 2001).

In case the responsibility for a system failure is attributed to a criminal or malicious intent, the investigation becomes a prosecutorial forensic case to identify and

prosecute the perpetrator. A representation of the complete investigation process is provided in the flowchart in Figure 4.

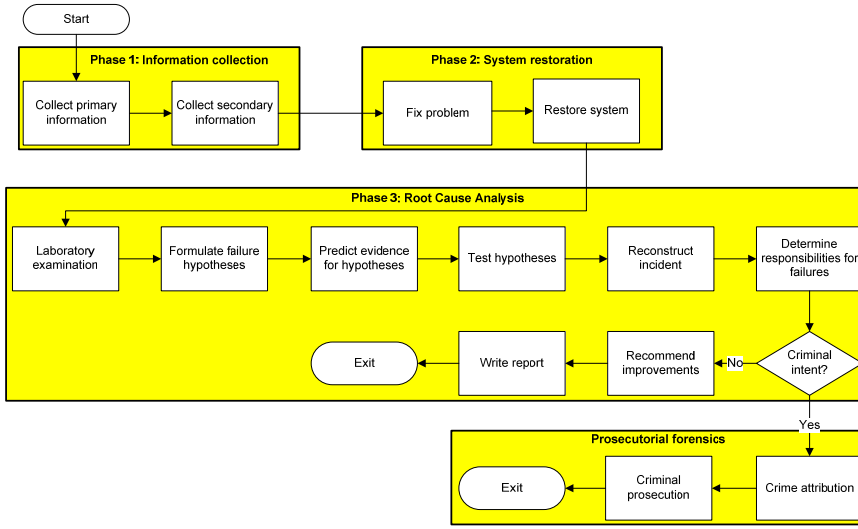

Figure 4: The operational forensic investigation process

The process does not indicate a technique to examine the collected evidence but rather steps that can lead to a thorough investigation of all possible sources of the problem and how to choose the best one. As it is a lengthy process it is best suited to significant failures with high impact.

6. Case study: Therac-25 accidents

This section illustrates the application of the proposed investigative process in a real-life scenario. The example used is the infamous disaster of the Therac-25, a computer-controlled radiation therapy machine used to treat cancer patients. The machine, designed by AECL (Atomic Energy of Canada Limited), was installed in 11 hospitals throughout the USA and Canada in the mid 1980's. Due to several software bugs in the machine, a series of six (6) severe overdoses of radiation occurred between 1985 and 1987 in different hospitals killing several patients (Leveson & Turner, 2002). The author chose this event as a case study as it is well documented with publicly available reports on the various accidents and resulting investigations. This is not the case for more recent incidents as companies keep this type of information confidential for legal reasons and to uphold their public image.

An in-depth investigation of the Therac-25 disaster is beyond the scope of this paper. Reports of such investigations are publicly available (Leveson, 1995; Leveson & Turner, 2002) and the root causes of the accidents have been identified. In this section, we demonstrate how our process could have been used for the investigation. The information presented is from Leveson (1995) and Leveson and Turner (2002).

The Therac-25 was used to administer a radiation beam to the patient in either one of two modes depending on the depth of the tumour:

- Low energy or electron therapy: electron beam of 200 rads aimed at patient directly. The computer controls the beam energy (5 to 25 MeV1) and current.
- High energy or X-ray: 25 MeV through a metal plate between beam and patient. Metal plate transforms beam into an x-ray. Electron beam 100 times greater than for low energy mode. The positioning of the metal plate is determined by a turntable.

Prior to the accidents, the Therac-25 had already been in use for two (2) years and had successfully treated hundreds of patients in various hospitals. The operators of the machine, however, had become accustomed to its frequent malfunctions which had never affected any patient before the deadly accidents. In such cases, the operator would call a hospital technician to reset the machine and restore it to service. This was the troubleshooting approach commonly used and which they initially followed after each accident.

First event: Kennestone Oncology Center, Marietta, Georgia, USA, 3 June 1985

The machine did not show any sign of unusual activity and did not generate an error message. However, the patient felt a high heat sensation after receiving treatment and accused the machine's operator of having burnt her. Shortly after returning home, the patient's skin reddened and swelled and she was in great pain. This was initially attributed to her disease. Weeks later, the patient's breast was removed, and her shoulder and arm were paralysed due to obvious radiation burn but the doctors could not explain its cause. It was later estimated that 15 000 to 20 000 rads had been administered instead of the set 200 rads.

No investigation was conducted for this accident as there was no information to indicate the machine was responsible for the patient's condition. The operational forensic process would not have yielded any result either as there was no primary or secondary information available. Indeed, the system was not forensic ready as the logs were not activated due to memory constraints. There was no system documentation available and no previous case had been reported. What could have been done (but was not done), however, was to interview the patient and the machine operator and file a report of the incident for future reference.

Second event: Ontario Cancer Foundation, Hamilton, Ontario, Canada, 26 July 1985

The machine paused after 5s of activation and displayed HTILT error message, NO DOSE and TREATMENT PAUSE. As the machine indicated that no radiation had been administered, the operator retried four (4) times until the machine stopped. The patient complained of burning electric sensation after the treatment. On 30 July, she was hospitalised as her skin was swollen and burnt and the machine was put out of

service. She died on 3 November 1985 from cancer but the autopsy revealed that the radiation burn would have necessitated a complete hip replacement had she survived. It was later estimated that she had received 13 000 to 17 000 rads. Table 3 shows how this accident could have been investigated with our proposed operational forensic process.

Investigation	
What was done with troubleshooting	**What could have been done with the operational forensic process**
Phase 1: Information collection	
No information was collected.	- **Collect primary information**: No log files, but record error messages. - **Collect secondary information**: No system specification and test plans, but obtain user manual and case history. Also interview the machine's operator and the patient.
Phase 2: System restoration	
The machine was reset by the hospital's technician who did not find anything wrong. Operation of the machine was discontinued five (5) days later.	- First reset the machine so that it can resume working. - Discontinue usage of the machine once patient started developing skin reddening and swelling after the treatment. - Only put the machine back into service once the investigation has been completed and the implemented improvements have been tested.
Phase 3: Root cause analysis	
- AECL first tried to recreate the problem with no success. - AECL suspected a mechanical failure and hardwired its error conditions. They found some hardware design flaws and fixed them. - They also modified the software to better control the positioning of the turntable. - Based on these changes, AECL claimed a significant improvement of the machine, although they concluded that they could not ascertain the exact cause of the accident. The machine was put back into operation despite this uncertainty.	**Laboratory examination of collected data**: - *User manual*: the user manual's description of many error messages was cryptic. The meaning of HTLILT was unclear. NO DOSE indicates that no dose of radiation has been delivered. - *Report of 1st accident:* based on the patients' testimony and symptoms, a correlation could have been established between the two events. **Formulation of hypotheses**: 3 possible scenarios - *Electrical problem* since patients experienced electrical shock. - *Hardware failure.* E.g. Incorrect positioning of the metal plate - *Software error* since the software controlled the machine. **Prediction of evidence to support hypotheses** - The electrical shock theory was ruled out by a thorough inspection by an independent engineering company which did not find any electrical problem in the machine. - AECL's test identified some hardware design flaws, which supported the hardware failure theory. - AECL identified some weaknesses in the software, supportive of the software error theory. **Test the hypotheses** - Thorough testing of the improved machine with the corrected mechanical flaws would result in another overdose as other accidents followed the 2nd one despite this improvement. This would have ruled out the mechanical failure theory. - The only theory remaining was the software error. Further examination of the software would be necessary to identify the bugs responsible for the failure. The last four (4) steps of the investigation (**incident reconstruction, responsibilities for failures, recommendations for improvement and report writing**) depend on the results of the thorough examination of the software to identify the bugs. As this was done following the 6th and last accident, they are not covered in this paper.

Table 3: Operational forensic investigation of the 2nd Therac-25 accident

As this example demonstrates, the operational forensic process offers many advantages over the troubleshooting method used in the case of the Therac-25. It could have located the source of the problem as a software error and not a hardware failure as suspected by AECL. In addition, further software examination would have identified the software bugs responsible for the overdoses before other accidents occurred. Unfortunately, due to overconfidence in their software, AECL refused to consider this option until several other accidents occurred after they had fixed the hardware.

In essence, a comprehensive forensic investigation would have provided the following benefits. Firstly, ensure that the results of the investigation were reliable as they were based on objective scientific analysis. Secondly, ensure that the root cause and not a proximate cause for the failure was identified before restoring the machine to operation. This would have prevented further accidents. Thirdly, improve the quality of the machine and AECL's procedures for failure analysis. AECL had no forensic capability and no mechanism to follow-up on reported incidents. Besides, there was no audit log activated on the machine and documentation of the software design and test plan were lacking. The user manual was also poorly designed. In addition, late during the actual investigation, it was established that AECL did not perform a thorough testing of the software before installing the machine.

7. Conclusion

Operational forensics has the potential to solve some of the main limitations of troubleshooting for cases of complex IT system failures. The proposed process shows how the forensic method can lead to a proper diagnosis of the problem. It is not a silver bullet which guarantees the identification of the root cause but it ensures that all aspects of the problem are taken into account before reaching a conclusion. This is illustrated by the case study. The lack of an objective and sound failure analysis process supported by appropriate evidence was among the main errors AECL made that led to the multiple accidents. Based on their experience with the system, AECL engineers "diagnosed" the accidents without supporting evidence, which is typical of troubleshooting. Future research involves specifying appropriate methods to examine the evidence, reconstruct the failure and localise its root cause.

Acknowledgement

The support of SAP Research Pretoria/Meraka Unit of Technology Development towards this research is hereby acknowledged. Opinions expressed and conclusions arrived at are solely those of the authors and should not necessarily be attributed to SAP Research Pretoria/Meraka UTD.

8. References

Basel Committee on Banking Supervision (2003), "Risk Management Principles for Electronic Banking", Bank for International Settlements website, July 2003, http://www.bis.org/publ/bcbs98.pdf, (Accessed 21 February 2011).

Bihina Bella, M.A., Olivier, M.S. and Eloff, J.H.P. (2009). "A Fraud Management System Architecture for Next-Generation Networks", *Forensic Science International*, Vol. 185, pp.51-58.

Brown, J.F., Obenski, K.S. and Osborn, T.R (2003), *Forensic Engineering Reconstruction of Accidents*, 2nd edition, p4, Charles C. Thomas Publisher, Springfield, Illinois, USA.

Brown, S. (1995), *Forensic Engineering Part 1 – An Introduction to the investigation, analysis, reconstruction, causality, prevention, risk, consequence and legal aspects of the failure of engineered products*, ISI Publications, Texas.

Carper, K.L. (2000), *Forensic Engineering*, Second Edition, pp.2-4, CRC Press, Boca Raton.

Corby, M. J. (2000a), "Operational Forensics – The New Frontier", *Proceedings of the 23rd National Information Systems Security Conference*, Baltimore, USA, 16-19 October 2000, http://csrc.nist.gov/nissc/2000/proceedings/papers/317slide.pdf, (Accessed: 10 May 2010)

Corby, M. J. (2000b), "Operational Forensics", *Information Security Management Handbook*, 4th Edition, vol. 2, chapter 28, Auerbach Publications, Boca Raton.

Dolinak, D., Matshes, E.W and Lew, E.O. (2005), *Forensic pathology: principles and practice*, p68, Elsevier, Oxford.

Giapponi, T.R. (2008), *Tire forensic investigation: analyzing tire failure*, pp.xv-xvii, SAE International, Warrendale, USA.

Hodd, B. (2010), "Modelling for operational forensics", *Digital Forensics Magazine,* Issue 3, 1st February 2010.

ISO/IEC (2007), "Information technology - Security techniques - Code of practice for information security management", International Standard ISO/IEC 27002, ISO Copyright office, Geneva.

Jordan, S. (2008), "Mining gold... A primer on incident handling and response". *SANS Institute InfoSec Reading Room*, http://www.sans.org/reading_room/whitepapers/incident/mining-gold-primer-incident-handling-response_32818, (Accessed: 09 February 2011).

Juniper Networks (2011), "Basic Approaches to Troubleshooting", Juniper website http://www.juniper.net/techpubs/software/junos/junos42/swcmdref42/html/strategies2.html, (Accessed 26/02/2011)

Kent, K., Grance, T., Chevalier, S. and Dang, H. (2006), "Guide to Integrating Forensic Techniques into Incident Response", *NIST Special Publication 800-86*, National Institute of Standards and Technology, Gaithersburg, USA, http://csrc.nist.gov/publications/nistpubs/800-86/SP800-86.pdf, (Accessed 09 February 2011)

Leveson, N. (1995), "Medical Devices: The Therac-25", *Safeware: System Safety and Computers*, Appendix A, pp515-548, Addison-Wesley.

Leveson, N. and Turner, C. (2002), "An Investigation of the Therac-25 Accidents", *IEEE Computer*, Vol. 26, Issue 7, August 2002, pp.18-41.

McDanels, S. J. (2006), "Space Shuttle Columbia Post-Accident Analysis and Investigation", *Journal of Performance of Constructed Facilities*, Vol. 42, Issue 3, pp.159-163.

Noon, R.K. (1992), *Introduction to Forensic Engineering*, 1st edition, p1, CRC Press, Boca Raton.

Noon, R.K. (20001), *Forensic Engineering Investigation*, 1st edition, p1, CRC Press, Boca Raton.

Palmer, G. (2001), "A Road Map for Digital Forensics Research", *Report from the First Digital Forensics Research Workshop*, 7-8 August 2001, New York, http://www.dfrws.org/2001/dfrws-rm-final.pdf, (Accessed 25/02/2011).

Ratay, R.T. (2010), *Forensic structural engineering handbook*, 2nd edition, McGraw-Hill, New York. p.xi.

Roger, W.P. (1986), "Report of the Presidential Commission on the Space Shuttle Challenger Accident", US Government Accounting Office, http://history.nasa.gov/rogersrep/genindex.htm, (Accessed 27/02/2011).

Scarfone, K., Grance, T. and Masone, K. (2008), "Computer Security Incident Handling Guide", *NIST Special Publication 800-61*, Revision 1, March 2008, National Institute of Standards and Technology, Gaithersburg, USA, http://csrc.nist.gov/publications/nistpubs/800-61-rev1/SP800-61rev1.pdf, (Accessed 22 February 2011).

Shedden, P., Ahmad, A. and Ruighaver, A.B. (2010). "Organisational Learning and Incident Response: Promoting Effective Learning Through The Incident Response Process", *Proceedings of the 8th Australian Information Security Management Conference*, Edith Cowan University, Perth Western Australia, 30th November 2010.

Stephenson, P. (2003), "Modeling of Post-Incident Root Cause Analysis", *International Journal of Digital Evidence*, Vol. 2, Issue 2.

Stephenson, P. (2004), "The Application of Formal Methods to Root Cause Analysis 0f Digital Incidents", *International Journal of Digital Evidence*, Vol. 3, Issue 1.

TechTerm.com (2011), "Troubleshooting", computer and technology terms online dictionary, http://www.techterms.com/definition/troubleshooting, (Accessed 10 February 2011).

Trigg, J. and Doulis, J. (2008), "Troubleshooting: What Can Go Wrong and How to Fix It", *Practical Guide to Clinical Computing- Systems: Design, Operations, and Infrastructure,* chapter 7, pp 105-128, Elsevier Inc, London.

Turner, P. (2007), "Applying a forensic approach to incident response, network investigation and system administration using Digital Evidence Bags", *Digital Investigation*, Vol. 4, Issue 1, March 2007, pp.30-35.

Usmani, A. S., Chung, Y. C. and Torero, J. L. (2003), "How did the WTC towers collapse: A new theory", *Fire Safety Journal*, Vol. 38, Issue 6, pp.501-533.

Zhou, Q. And Yu, T.X. (2004), "Use of high –efficiency energy absorbing device to arrest progressive collapse of tall building", *Journal of Engineering Mechanics*, Vol.130, issue 10, pp.1177-1187.