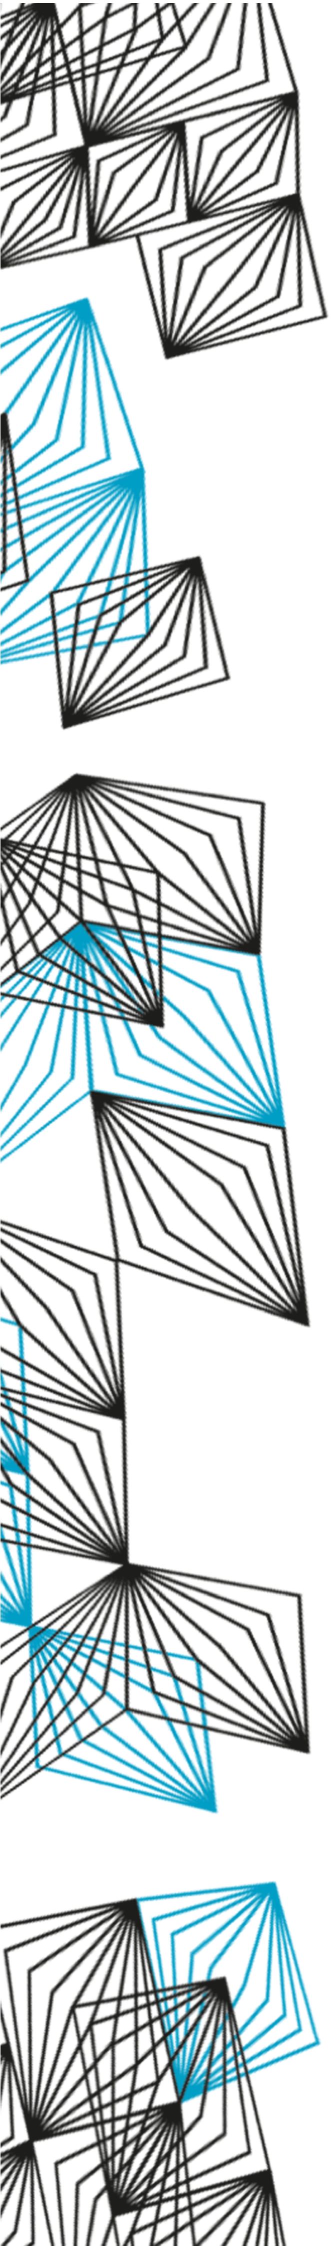

BEST PRACTICES IN CLOUD INCIDENT HANDLING

KIMBERLY HENGST

Submitted in partial fulfilment of the requirements for the degree of Master of Science

to the

Faculty of Electrical Engineering, Mathematics and Computer Science

February 7, 2020

UNIVERSITY OF TWENTE.

Kimberly Hengst: *Best Practices in Cloud Incident Handling,* © February, 2020. In cooperation with the Dutch NCSC.

SUPERVISORS:
Dr. A. Sperotto
Dr. J. van der Ham
Prof. Dr. M. Junger
M. van Piggelen

ABSTRACT

In the current trend of transitioning towards cloud environments, companies report issues with detecting and responding to cloud security incidents. Research has shown that organisations experience many challenges, among which are an insufficient overview of information, a lack of visibility, and an inadequate design and road map.

Therefore, this research aims to determine the current best practice in cloud incident handling. Furthermore, it aims to determine to what extent this practice is sufficient in the current Dutch incident handling landscape.

Based on a literature study of existing literature on cloud incident handling, 12 semi-structured interviews have been conducted with 14 participants from Computer Security Incident Response Teams (CSIRTs) of Dutch organisations. A thorough analysis of both literature and practice resulted in guidelines and recommendations. While all recommendations should be considered by organisations, the results indicate five important recommendations: (1) organisations should prepare for cloud incidents by informing themselves of the characteristics and features of the cloud environment, (2) organisations should obtain visibility into their cloud environment by implementing cloud management, (3) organisations should ensure proper cloud security, (4) all agreements, requirements, and responsibilities must be included in the Service Level Agreement (SLA), and (5) incident information should be shared as this is crucial in preventing incidents and holding Cloud Service Providers (CSPs) accountable. The presented recommendations can be used by companies to further improve their cloud incident handling strategy and contribute towards decreasing the gap between theory and practice.

ACKNOWLEDGEMENTS

First and foremost, I want to thank my supervisors, Jeroen, Marianne, Anna, and Marit. I wholeheartedly appreciate their guidance and valuable feedback.

Furthermore, I want to thank the participants of this research, for donating their time and expertise on the subject.

Finally, I wish to thanks my friends and family for their support and love. Special regards to Joost, for being my sparring partner and support. Eva, for her emotional support during our Skype calls. Kira, for her books and wisdom. Marike and Tara, for their support and distractions. Finally, Martijn for his motivation and inspiration. Without the help and support from those surrounding me, this thesis would not have been possible.

- Kimberly

CONTENTS

LIST OF FIGURES

LIST OF TABLES

LIST OF ABBREVIATIONS

AWS	Amazon Web Services
BEC	Business Email Compromise
CASB	Cloud Access Security Broker
CSA	Cloud Service Alliance
CSIRT	Computer Security Incident Response Team
CSP	Cloud Service Provider
CSU	Cloud Service User
IaaS	Infrastructure as a Service
IDPS	Intrusion Detection and Prevention Systems
IoC	Indicator of Compromise
MFA	Multi-Factor Authentication
NCSC	National Cyber Security Centre
NIST	National Institute of Standards and Technology
PaaS	Platform as a Service
RDP	Remote Desktop Protocol
SaaS	Software as a Service
SIEM	Security Information and Event Management
SLA	Service Level Agreement

Part I

BACKGROUND

1 INTRODUCTION

Investigating who leaked your company's Intellectual Property (IP) used to be a straightforward, albeit complex, task. In an on-premise environment, you could grab the server the IP was on, do a full forensic analysis of the system, find who accessed what file, and find the culprit. Moving to the cloud, especially to Software-as-a-Service (SaaS), has complicated this process. With all your IP spread across hundreds of cloud servers, where will you look? What logs are available? What server will you start to image? It might be too late if you think of these questions only after an incident has happened in your cloud environment.

In a continuously changing landscape of IT, incident response teams must always remain up to date on current trends. Not being up to date means a delay in getting back-to-business and a possibly large impact on revenue. On a larger scale, this can cause a significant economic impact if we do not sharpen our incident-handling strategies in the Computer Security Incident Response Team (CSIRT) community. Many companies move towards cloud computing [7, 17, 27, 34]. However, many experience issues regarding detecting and responding to cloud security incidents [42, 49]. In collaboration with the Dutch National Cyber Security Centre (NCSC), we intend to identify the current incident handling capabilities of companies and organisations in the Netherlands to provide a current best practice based on theoretical and empirical research. This way, organisations can be prepared with the proper tools before cloud incidents will happen, instead of filling this need when incidents have already occurred. This research, if adopted by organisations, will prevent or at least decrease the economic damage that criminals and state actors can have on our society.

This chapter introduces the research. Background on cloud computing and incident handling is provided in Section 1.1 and Section 1.2 respectively. Section 1.3 focuses on the challenges regarding cloud computing. These result in research questions which are presented in Section 1.4. The contribution of this research is described in Section 1.5. Finally, this chapter is concluded by providing an outline of this thesis in Section 1.6.

1.1 CLOUD COMPUTING

The National Institute of Standards and Technology's (NIST) SP800-145 [37] provides the most widely used definition and taxonomy of cloud computing. Their definition of cloud computing is as follows: *"Cloud computing is a model for enabling ubiquitous, convenient, on-demand network access to a shared pool of configurable computing resources (e.g., networks, servers, storage, applications, and services) that can be rapidly provisioned and released with minimal management effort or service provider interaction"* [37].

This definition is reflected by three service models in which Cloud Service Users (CSUs) have different levels of control: Infrastructure-as-a-Service (IaaS), Platform-as-a-Service (PaaS), and Software-as-a-Service (SaaS). The level of control is the main difference between cloud and on-premise computing, where the organisation has complete control in the latter over the infrastructure and deployed software. Although some recent publications argue these cloud service models should be more specific - such as Framework as a Service, Runtime environment as a Service, and Database as a Service [20, 30] -, the taxonomy presented by NIST is adopted in this paper. This taxonomy is used by the Dutch NCSC as well [41].

Due to the different models with different levels of control, cloud computing poses many challenges which are discussed in Section 1.3.

1.2 INCIDENT HANDLING

Multiple frameworks exist on incident handling such as ISO/IEC 27035 [43, 44], the ISACA Incident Management and Response Framework [46], and the NIST SP800-61 [13]. Most publications have based their work on NIST's SP800-61. It provides a Computer Security Incident Handling Guide [13] that describes guidelines for organisations to establish computer security incident response capabilities and handle incidents. These guidelines for handling an incident are structured according to their incident response lifecycle. This lifecycle describes four phases that occur when handling an incident: (1) preparation, (2) detection and analysis, (3) containment, eradication, and recovery, and (4) post-incident activity. NIST has illustrated this lifecycle, which is presented in Figure 1.1.

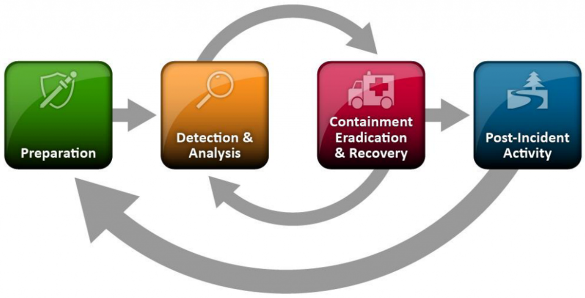

Figure 1.1: The NIST incident response lifecycle as presented in SP800-61 [13]

Although this lifecycle does not cover the cloud specifically, many publications have used the phases of the NIST incident response lifecycle to structure their work, and it has often been adapted for cloud incident handling strategies. A description of each phase is presented in Table 1.1.

Table 1.1: Descriptions of each phase of the NIST incident response lifecycle

Phase	Description
Preparation	This phase aims to establish an incident response capability and prevent incidents by ensuring existing applications, systems, and networks are secure.
Detection & Analysis	In this phase, it is determined whether an incident has occurred. This begins when anomalous behaviour is flagged when there are signs of an incident. Analysing the flagged incident determines if this behaviour is a valid threat and what priority it should receive
Containment, Eradication & Recovery	This phase ensures that the threat is contained to prevent it from infecting other systems. After containment, the threat needs to be eradicated from compromised assets. Finally, the normal operation of assets is restored.
Post-incident Activity	In this phase, incident response teams reflect on the incident to evaluate their incident handling process. This determines what occurred, what was done to mitigate the incident, and what should be done in the future.

Many publications are inconsistent regarding the terminology used for incident handling and incident response [2]. According to NIST, incident handling and incident response are synonyms [13]. However, the terminology used by Ab Rahman et al. [2] is adopted in this research, namely: incident response describes the containment, eradication, and recovery phase while incident handling is comprised of the aforementioned four phases.

1.3 CHALLENGES IN CLOUD INCIDENT HANDLING

Incident handling has traditionally been focused on on-premise environments. However, due to the different characteristics of cloud environments, cloud incident handling faces many challenges. Literature indicates many challenges exist with collecting, obtaining, and analysing data. Furthermore, it is difficult to understand the division of incident handling responsibilities and to obtain visibility into incidents and shadow IT (see Section 3.2).

When formulating the research questions security specialists and advisers at the NCSC and a security specialist at KPN were consulted on challenges with cloud incident handling in practice. The aim was to get a better understanding of issues that are currently experienced within companies and determine the scope of this research. Six issues were identified: (1) a lack of technology and process-related knowledge, (2) an insufficient overview of information, (3) a lack of visibility, (4) an inadequate design and road map, (5) a dependency on the vendor, and (6) conducting an investigation. These challenges from practice are further elaborated on in this research (see Section 4.2).

1.4 RESEARCH QUESTIONS

The challenges described in the previous section show that there are many issues regarding cloud incident handling, without clear solutions, and there is an interest in tangible support and recommendations. This is supported by other research (see Chapter 3). Therefore, this research aims to determine the current best practice in cloud incident handling which results in the first main research question (RQ 1): What is the current best practice in cloud incident handling?

The first step in answering this research question is identifying and evaluating how cybersecurity incidents are currently handled, and whether companies experience issues or shortcomings with their approach. Although the previous section described multiple challenges in cloud incident handling, they mainly showed a need for research into this area. Therefore, these issues are not considered final and are improved upon in this research. Furthermore, by identifying the current incident handling landscape of Dutch companies, it can be determined what knowledge exists within companies, how their current incident handling strategies are implemented, and what problems they encounter with their current implementations. This contributes to providing appropriate challenges and best practices. Therefore, identifying the current incident handling landscape of Dutch companies is one of the sub-questions for this research (RQ 1.1).

The second step is identifying the differences between cloud incident handling and on-premise incident handling (RQ 1.2). Using cloud environments changes how incidents are handled but can aid the incident handling process as well[24]. Determining these differences and accompanying challenges and opportunities results in more refined best practices and recommendations regarding cloud incident handling.

The third step of the research necessary for answering the first main research question is providing an overview of best practices identified by both literature and the practice. Therefore, this research considers defining best practices in cloud incident handling according to literature and the practice in the Netherlands as one of the research questions for this project (RQ 1.3).

The answers to these three sub-questions result in an overview of challenges and best practices in cloud incident handling. An essential part is combining cloud incident handling practices suggested in literature and practice. This provides organisations with practical guidelines based on the analysis of the current incident handling landscape of research question 1.2 and the challenges identified in research question 1.1. Companies can use these guidelines to evaluate their incident handling strategies. This can be used to answer the first main research question and determine the current best practice (RQ 1).

When a current best practice can be determined, it is important to determine to what extent that practice is sufficient in the current cloud incident handling landscape by comparing it with the identified challenges. Therefore, the second main research question (RQ 2) is: to what extent is the identified best practice sufficient in the current cloud incident handling landscape?

Summary

The aforementioned can be summarised into the following main research questions and sub-questions:

- **RQ 1:** What is the current best practice in cloud incident handling?
  - **RQ 1.1:** What is the current incident handling landscape of Dutch organisations?
  - **RQ 1.2:** What are the differences between cloud incident handling and on-premise incident handling according to literature and the practice in the Netherlands?
  - **RQ 1.3:** What are best practices in cloud incident handling according to literature and the practice in the Netherlands?
- **RQ 2:** To what extent is the best practice determined in RQ1 sufficient in the current cloud incident handling landscape?

1.5 CONTRIBUTION

The contributions of this research are threefold: it provides an insight into the current state of cloud incident handling in organisations in the Netherlands, it determines the differences between cloud incident handling and on-premise incident handling, and it identifies the current best practice in cloud incident handling. The first contribution shows solutions that need to be developed to satisfy the needs regarding cloud incident handling that exist in companies. Furthermore, it shows what areas research should focus on. The second contribution raises awareness that cloud incidents require an adapted incident handling strategy. Furthermore, it helps understand the differences between cloud incident handling and traditional incident handling. The third contribution provides companies with practical guidelines that can be used to implement or further improve their cloud incident handling strategy.

1.6 OUTLINE

Details of the method, previous research, and the results of this research can be found in subsequent chapters. This thesis is structured as follows:

Chapter 2 describes the method used in this research in three parts: the literature review, the interview process, and the comparison of the results from theory and practice.

Chapter 3 provides an overview of existing research into cloud incident handling. It introduces the results of existing interview surveys that have been conducted in the past and discusses cloud incident handling approaches that have been proposed and discussed in related work. An overview of the main findings is provided.

Chapter 4 describes the current Dutch incident handling landscape. It provides an overview of differences and challenges identified by organisations, best practices as identified by interviewees, and an overview of recommendations.

Chapter 5 combines the results from theory and practice and presents it in a concise overview.

Chapter 6 compares the results from literature and practice. In addition, it discusses the research and its limitations.

Chapter 7 concludes the report. It answers the research questions and presents directions for future research.

Part II

METHOD

2 METHOD

This chapter describes the method used to analyse existing literature, using a literature review, and practice, using an interview study.

Section 2.1 describes the method used to identify the differences between cloud incident handling and on-premise incident handling and the best practices in cloud incident handling in literature. Section 2.2 describes the method for identifying the current incident handling landscape of Dutch organisations. It describes the participants, the procedures taken and the details of the analysis. Finally, Section 2.3 describes the method used to derive best practices from literature and the practice in the Netherlands.

Methodological triangulation was used in this research, which means that more than one method was used to gather data, namely: conducting a literature review and an interview study. One of the criticisms on triangulation is that it inherently assumes that data can be compared while failing to take different circumstances into account. However, it adds another dimension, which fits the exploratory nature of this research [9].

2.1 CLOUD INCIDENT HANDLING IN LITERATURE

2.1.1 *Collecting data*

Papers relating to cloud incident handling were collected using Scopus and Google Scholar. Multiple search queries were used, which were a variation of the following basic search query: `( incident AND ( handling OR response OR management ) ) AND ( cloud OR "cloud computing" )`. Additional search terms such as "digital forensics" or "forensic readiness" were used to adapt the basic search query. Papers in the English language were a requirement. A selection of relevant studies was made based on title, abstract, and availability.

References were used to add additional papers that were not in the initial selection. The selection of relevant studies was refined by adding or removing papers based on a review of the full-text of each. During the full-text review, relevant information was highlighted.

2.1.2 *Analysing data*

Each paper was summarised using the highlighted information to compile the specific information relevant to the current research. These summaries were used to identify different topics within each phase of the NIST incident response lifecycle. For example, the "preparation" phase consists of "incident handling process", "SLA", "technical", and "reporting". Using the summaries, a second review of the content of each paper resulted in an overview of all exact quotes regarding each topic. Subtopics were created inductively by categorising all exact quotes per topic.

A third review of the content of the papers was used to identify information that might have been omitted from the summaries. Finally, the main findings for each subtopic were described based on the collected information.

2.2 CLOUD INCIDENT HANDLING IN PRACTICE

2.2.1 *Design*

A combination of interview survey and questionnaire survey design was deployed to analyse the current incident handling landscape of Dutch companies. This research has been approved by the ethics committee of the University of Twente. The number of the request is RP 2019-84.

2.2.2 *Participants*

CSIRTs were deemed the most relevant participants due to their expertise in incident handling. Listings of CSIRTs on the websites of FIRST, ENISA, and the Wikipedia page on Dutch CSIRTs were used to select companies and organisations [18, 15, 22]. Additionally, incident responders from the NCSC were asked whether they had suggestions for relevant companies. Some participants were approached as a result of the snowball method, where previous participants suggested relevant companies and organisations at the end of their interview.

Initially, a sample size of 20 participants was decided upon. However, it took more time than expected to find willing participants. Additionally, saturation was reached after ten interviews, as no new information was discussed in the final two interviews. The aforementioned, in addition to many companies declining due to the lack of experience with cloud incidents, led to the conclusion that the current amount of interviews was satisfactory and arranging and conducting more interviews to reach the intended sample size would not be beneficial.

A list of 33 potential participant organisations was compiled using the aforementioned sources. At least 27 of them were approached based on feedback received by the NCSC. The exact number is not known, as invitations were distributed by the NCSC to partners they did not disclose due to confidentiality. Six responded that the research was not relevant to them as - according to them - they did not have enough expertise i n cloud incident handling yet. In total, 14 participants from 12 companies volunteered their time for the study. Three organisations wished to remain anonymous. The organisations are, in alphabetical order, as follows:

- Fox-IT
- Informatiebeveiligingsdienst
- KPN
- Northwave
- Onderlinge
- Rabobank
- SURFcert
- Tesorion
- Universiteit Twente

The following occupations were represented: CERT manager, security specialist, incident handler, security manager, security officer, and security engineer.

2.2.3 *Materials and Procedures*

Pilot

A questionnaire and interview guide were used to conduct the interviews. A pilot was conducted with two companies to evaluate the questionnaire (described in Section 2.2.3.1), the interview process (described in Section 2.2.3.2) and the time necessary to conduct the interview. Based on the feedback received from these companies, small changes were made to the questionnaire. However, this did not invalidate the already completed pilot-questionnaires, allowing those to be used during the research.

Recruiting participants

Requests for participation were sent via email to prospective companies - based on the relation between the NCSC and the respective companies - either by an NCSC representative or to the general email address on the company's website. This request contained information about the aim of the research, the target group characteristics, the time the person would spend when participating, and it stated all data would be treated anonymously.

When confirming the appointment via email, a questionnaire was sent together with the confirmation email. Participants were asked to send the completed questionnaire back before the interview appointment. They were specifically informed that the purpose of this questionnaire was for the researcher to prepare the interview.

2.2.3.1 *Questionnaire*

The questionnaire consisted of ten questions; four questions focused on the company's incident handling process, two focused on cloud incidents that occurred within the company, and four focused on the Cloud Service Provider (CSP). The questionnaire was in Dutch and can be found in Section A.1. Three different types of questions were used in the questionnaire: a ten-point scale, a five-point Likert scale, and a ratio scale. Translated examples of questions are given in the following paragraphs.

Question types

A ten-point scale was used with seven questions to rate the different phases of the company's incident handling process and their opinion on their CSP. An example of such a question is: "What grade do you give your organisation concerning the *preparation* phase when it comes to cloud security incidents?". Participants could score their answer from one to ten, where "1" signified a bad rating, and "10" a good rating. Additionally, participants could further elaborate on the strengths and weaknesses of their incident handling process.

A five-point Likert scale was used in two questions. An example of such a question is: "How clear is the division of roles and responsibilities with the CSP during an incident?". Participants could rate this question from "Very unclear" to "Very clear". In case a participant could not answer the question, an extra (sixth) option was added to allow for "I do not know".

A ratio scale question was used in one question to determine participants' impression of the ratio of cloud security incidents to total security incidents. The question was: "What do you estimate is the ratio of cloud incidents to total security incidents in the past year (in percentage)".

2.2.3.2 *Interviews*

Interviews were conducted to collect practices within Dutch organisations regarding cloud incident handling. For example, it identified which cloud services are used, which challenges companies encounter regarding cloud incident handling, and how companies have implemented each phase of the NIST incident response lifecycle.

Interview Guide

Semi-structured interviews were conducted in two types of interviews (in person or by phone) using an interview guide (see Section A.2). The interview guide consisted of six topics which were derived by studying existing literature (see Chapter 3): (1) general information, (2) differences between cloud and on-premise incident handling, (3) cloud usage within the organisation, (4) cloud incidents within the organisation, (5) incident handling process, and (6) the CSP. Although semi-structured interviews were conducted, example questions were formulated for each of these topics which were used as preparation for the interview by the researcher. An overview of the

Dutch topic list and the detailed preparation questions are provided in Section A.2.

The topic "general information" aimed to gather background information and to make the participant and researcher familiar with each other. For example, it focused on the daily activities of the participants and the security capacity of the organisation. The topic "differences between cloud and on-premise incident handling" served to introduce the participant to the topic and to define the scope. It provided the researcher with an indication of the participants' expertise which was then used to tailor questions. For example, when a participant focused on detailed technical differences during this topic, further questions would be focused more towards, but not solely on, technical details. The topic "cloud usage within the organisation" aimed to familiarise the researcher with the cloud services that the organisation uses. The resulting overview was used to ensure further questions would cover all services to prevent too much emphasis on a single service. The topic "cloud incidents within the organisation" identified which cloud incidents have occurred. This was used to tailor the conclusion of this research. Additionally, it served to obtain more information by focusing the participant on specific details surrounding an incident. The topic "incident handling process" discussed the organisation's incident handling process according to the four phases of the NIST incident response lifecycle. Allowing participants to discuss on-premise incident handling strategies prevented missing out on on-premise practices that companies use for the cloud as well but would not think to mention. Finally, the topic "CSP" would discuss participants' experiences with their CSPs. For example, this topic aimed to identify to what extent the CSP is involved in the participant's cloud incident handling process.

The aforementioned topics were not necessarily discussed in the presented order. In case participants mentioned other topics, the interview would continue on those other topics first to not disturb the flow of the conversation. In some cases, the topics and questions were adapted based on the type of company interviewed and the person interviewed. For example, when a company provides security services for their customers, questions would focus more on that service instead of the company's cloud services.

Interview procedure

The interviews were conducted either in person (n=6) or by phone (n=6). Due to time constraints, the choice was made to start conducting interviews by phone halfway through the interview process. The following describes the accompanying procedures of both:

INTERVIEW IN PERSON: When interviewed in person, the participants received an information brochure at the start of the interview containing the aim of the research, the details on withdrawing their participation, and whom to contact in case of further questions or complaints. Furthermore, the participants received a consent form to sign, which contained the option to allow audio recordings to be made. This was opt-in and it was explained that the purpose of the audio recording was to aid in processing the interview. In case the participants consented, the audio recording was started and the researcher made brief notes during the interview. In case the participants did not consent to the audio recording, the researcher made detailed notes on a laptop during the interview. Questions were then asked using the interview guide.

INTERVIEW BY PHONE: When interviewed by phone, the participants received the information brochure and the consent form, as described above, by email. In case the participants consented to audio recordings being made, after receiving the aforementioned explanation verbally, a laptop was used to record the call. In case the participants did not consent, the researcher made detailed notes on a laptop during the interview. Questions were then asked using the interview guide.

Debriefing

After concluding the interview, participants were asked whether their company name and job function could be mentioned in the research. When the questionnaire or consent form were unable to be completed before the interview, an email would be sent afterwards requesting the participant to provide these.

2.2.4 *Analysis*

A Qualitative Content Analysis (QCA) was used to analyse the interview data [32]. This is a widely used method to analyse qualitative data. It typically consists of the following six phases:

1. Preparing the data
2. Forming main categories corresponding to questions asked in the interview
3. Coding data with the main categories
4. Compiling text passages of the main categories and forming subcategories inductively on the material; assigning text passages to subcategories
5. Category-based analyses and presenting results
6. Reporting and documentation

Data preparation

The first phase consists of preparing the data. The audio of the interviews was transcribed to be able to better analyse the data. An initial read-through of the transcripts was conducted to get familiarised with the data. No extensive interpretation of the data was required. However, some interpretation was required when audio was distorted, only notes were available, or certain terminology needed to be further understood. Due to the confidentiality of the contents, the transcripts have not been included in this thesis.

Main categories formation

In the second phase, the main categories are formed based on the interview guide. The six phases presented above utilise concept-driven and data-driven development of codes in phase two and four respectively. This means that first, an initial coding frame with deductively formed codes was formed based on the interview guide. The following 11 categories were used in the first coding cycle:

1. Challenges
2. Cloud incidents
3. Differences
4. General information
5. General solutions
6. Incident Handling
7. Incident Handling: (1) preparation
8. Incident Handling: (2) detection and analysis
9. Incident Handling: (3) containment, eradication and recovery
10. Incident Handling: (4) post-incident activity
11. Incident Handling: (5) digital forensics

First coding cycle

In the third phase, the transcripts were coded according to these categories using the software ATLAS.ti [5]. Only text segments relevant to the research were coded [32].

Second coding cycle

In the fourth phase, subcodes were developed inductively on the data. First, all text segments corresponding to each category were compiled. Then, all text segments were coded once more with a method similar to open coding [16]. This means that the text segments were given conceptual labels, which grouped to form subcategories. The process of open coding is performed until the creation of subcategories appears saturated. Then, the data that has not been labelled with a subcategory will be coded.

Category-based analyses

In the fifth phase, concrete challenges and guidelines were derived from the coded text segments. By grouping and combining the text segments allocated to a category, guidelines could be derived.

Documentation

Finally, in the sixth phase, the results were documented. The participants of this research are described in Section 2.2.2. The results of the analysis are presented in Chapter 4.

2.3 COMPARING PRACTICES

The results from the literature review and the interviews were used to ultimately identify the current best practice in cloud incident handling. As mentioned in the introduction of this chapter, the two different methods of data collection were chosen to add more dimension to the results. This is more suited to the exploratory nature of this research instead of solely validating the results of either method. The comparison of the results from literature and practice was conducted in the following manner:

1. Identify the overlap in recommendations from literature and practice.
2. Identify which recommendations from literature are not applied in practice, and about which respondents expressed a negative opinion.
3. Identify which recommendations from literature are not applied in practice, but could solve issues mentioned by participants.
4. Identify which recommendations from practice are not mentioned, or not recommended, in literature.

First, to answer RQ 1.3, all four comparisons were compiled, resulting in a list of best practices. However, this includes practices which are not recommended by either literature or practice. Therefore, excluding the recommendations from literature which are viewed negatively in practice and vice versa leads to the current best practice in cloud incident handling (RQ 1). By analysing which challenges faced by organisations remain unsolved by this identified current best practice, its sufficiency can be determined (RQ 2).

Part III

RESULTS

3 CLOUD INCIDENT HANDLING IN LITERATURE

Over the last few years, some research has been done on incident handling in the cloud. Research initially focused on incident handling and emergency response in general. However, as cloud computing became more relevant, research into cloud incident handling strategies gained attention.

This chapter describes literature related to the current research. Section 3.1 describes the differences between cloud and on-premise incident handling. Section 3.2 describes the challenges of cloud incident handling. Section 3.3 and Section 3.4 describe existing interview surveys and literature on cloud incident handling. Section 3.5 covers literature on digital forensics strategies and Section 3.6 provides an overview of the recommendations from presented related literature.

3.1 DIFFERENCES BETWEEN CLOUD AND ON-PREMISE INCIDENT HANDLING

NIST describes a cloud model that is composed of five essential characteristics, three service models, and four deployment models.

The five cloud characteristics are: (1) on-demand self-service, (2) broad network access, (3) resource pooling, (4) rapid elasticity, and (5) measured service. A summary of the three service models and four deployment models is provided in the following paragraphs [33, 37].

INFRASTRUCTURE AS A SERVICE (IAAS) provides the CSU with fundamental computing resources such as servers, storage, and networking capability. The CSU cannot modify or control the underlying cloud infrastructure but can control operating systems and deployed applications. An example of IaaS is Amazon Web Services (AWS) EC2.

PLATFORM AS A SERVICE (PAAS) allows the CSU to deploy applications onto the cloud infrastructure. The CSU cannot modify or control the underlying cloud infrastructure but can control the deployed applications and possibly environment configuration settings. An example of PaaS is Google App Engine.

SOFTWARE AS A SERVICE (SAAS) allows the CSU to use applications running on a cloud infrastructure. The CSU cannot modify or control the underlying cloud infrastructure. Examples of SaaS applications are Google Drive, Dropbox, and Slack.

These service models differ from on-premise computing, in which the organisation has complete control over the infrastructure and deployed software. This results in a shift in control from the CSU to the CSP [24], depending on the service model chosen.

While cloud service models describe what CSUs are permitted within a cloud environment, cloud deployment models describe where cloud infrastructure is located and who controls it. Cloud environments can be located across multiple systems or jurisdictions, resulting in a distributed nature [35]. This leads to new data sources, but could also result in less information available as CSPs might withhold information [39, 24]. Cloud environments are often comprised of multiple tenants [28, 24] which could be a reason why CSPs might withhold data. Additionally, due to cloud characteristics such as autoscaling, information might be limited as well [39].

Deployment models

The four different deployment models are as follows [33, 37]:

PRIVATE CLOUD is exclusively used by a single organisation. It may be controlled by the organisation, a third party, or a combination. It may exist on or off premises.

COMMUNITY CLOUD is used by multiple organisations sharing computing resources within a community. Examples of such communities are universities, police departments, or hospitals. It may be controlled by multiple organisations, third parties, or a combination. It may exist on or off premises.

PUBLIC CLOUD is used by all types of users, including the general public, on a subscription basis. It may be controlled by businesses, governments, or academic organisations. It exists on the premises of the CSP.

HYBRID CLOUD is a combination of two or more of the above cloud deployment models.

3.2 CHALLENGES IN CLOUD INCIDENT HANDLING

Traditional incident handling does not apply to cloud incidents due to the aforementioned cloud characteristics such as its distributed nature, multiple parties involved, and the many sources of information. The main challenges that arise with cloud incident handling are briefly discussed in the following paragraphs.

Collecting data

One of the biggest challenges in cloud incident handling is the collection of data to conduct proper incident handling. To collect data, potential data sources need to be identified. Cloud environments introduce new data sources, such as cloud management planes, which pose a challenge as security teams need to be aware of which source provides what data [39]. This is further complicated when security teams have limited knowledge about CSPs' architecture, either due to the CSP not providing this information, or due to the security teams not knowing this information exists [24]. Additionally, logs are not universally available at each cloud service [39], which limits data collection capabilities.

Obtaining data

The second issue arises when trying to obtain the necessary data. Obtaining data is complicated by the distributed nature of cloud environments across multiple systems or jurisdictions [35]. Additionally, log availability is an issue when security teams do not have, full, access to CSPs' sources [24, 1]. Possible causes for not having the desired access are that security teams might not have direct points of contact with the CSP or are limited to standard support [39]. Obtaining data could also be a challenge as CSPs might not want to provide certain data. Cloud environments are often comprised of multiple tenants which means that the necessary data might include data of other tenants. Therefore, data might not be allowed to be disclosed due to privacy [28, 24]. Even if the necessary data is available to use, it could be lost due to cloud characteristics such as auto-scaling activity, or by terminating virtual machines [39]. Finally, a lack of knowledge results in issues with obtaining data. SANS reported that 20% of the interviewed organisations indicate a lack of skills, training, and certification which makes it the second biggest challenges organisations face in cloud environments [25].

Analysing data

Analysing data becomes difficult due to the distributed nature of cloud environments [35]. This is further complicated by the large volume of data [39, 42, 49]. Additionally, analysis requires context from correlating data, which requires organisations to use machine learning and advanced analytics solutions [42]. However, in a survey of 450 participants, SANS found that the biggest challenge in investigating cloud incident was a lack of standards, tools and training [25]. Additionally, in a survey of 1250 participants, Symantec found that 49% of the participants report insufficient cloud security manpower. This leads to the inability to address all incoming security alerts [49].

Understanding the division of incident handling responsibilities

As the different cloud service models shift some degree of control to CSPs, a challenge arises in understanding the division of incident handling responsibilities [24]. As the possibilities to work with the provider become limited, it is important to establish clear incident handling responsibilities early [28]. When a CSU uses a cloud platform, they do not inherit the CSP's compliance and should still ensure they are compliant with regulations themselves [42].

Obtaining visibility into incidents

Another challenge is obtaining visibility into incidents [25]. In a survey of 450 participants, the biggest challenge by far (38%) is detecting and reacting to cloud security incidents [42]. Additionally, 30% reported two areas which needed to improve the most: identifying software vulnerabilities, and identifying noncompliant workload configurations [42]. Another challenge is limited network visibility, where for example network logs might not contain full packets [39].

Obtaining visibility into shadow IT

This issue is further complicated by the challenge to obtain visibility into shadow IT. Shadow IT "refers to the adoption and use of SaaS apps without the IT department's oversight or sanction"[49]. One of the most astonishing findings of the survey conducted (N=450) by Oracle and KPMG is that while 92% of the respondents require all cloud usage to be preapproved, 82% are concerned that these policies are violated [42]. Lack of visibility into cloud application usage is mentioned by 25% of the respondents [42]. Additionally, Symantec's survey of 1250 participants found that, on average, 28% of employees indulge in high-risk behaviour such as using personal accounts for cloud services and sharing credentials [49].

3.3 existing interview surveys on cloud security

In the current research, interviews are used to determine the current incident handling landscape of Dutch organisations (see Chapter 4). To our knowledge, such research does not exist yet. However, some comparable studies exist that surveyed organisations either worldwide or from different countries. This section provides a concise summary of existing research and their results.

Worldwide survey on digital forensics

The SANS Institute is a research and education organisation. They provide information security training and security certifications. In 2013, they surveyed 450 organisations on digital forensics practices [25]. These organisations were active in many different primary industries. The largest group was governmental organisations, covering 24% of the respondents. The main finding of this survey is that lack of tools, training, standards, and visibility are considered fundamental challenges in investigating cloud incidents. Furthermore, 31% of the respondents relied on their CSP to collect evidence for them, and only 16% of the respondents have a Service Level Agreement (SLA) that allows them to gather evidence themselves [25]. Based on the results of their survey, the SANS Institute provides many recommendations to organisations such as ensuring appropriate training, reviewing existing SLAs, and closing capability gaps in tools. Furthermore, organisations should consult their legal team to retain the validity of potential evidence.

Nordic survey on CSUs' incident handling needs

In 2015, Jaatun and Tøndel published a study in which they surveyed four critical infrastructure organisations based in Norway and Sweden [28]. This study for example highlights incident handling needs that apply to a CSU, mechanisms that would aid the collaboration between the CSU and the CSP during incident handling in the cloud, and the type of incidents companies experience. Although none of the surveyed companies make use of cloud services, Jaatun and Tøndel mention the relationship with the CSP, automation, and secure exchange of incident information as topics pertaining to cloud scenarios. Furthermore, they describe a tool that could improve the quality of incident information that a CSU receives.

Worldwide survey on security challenges of hybrid and multi-cloud environments

The Cloud Security Alliance (CSA) is an organisation dedicated to defining best practices in cloud security. Furthermore, they operate a cloud security provider certification program, a cloud security user certification program, and a framework that maps security controls to standards, regulations, and best practices. In 2019, they surveyed approximately 500 organisations to identify challenges in managing security in hybrid and multi-cloud environments [7]. The respondents were mainly larger organisations. The primary industry was "IT and technology", covering 38% of the respondents. The survey led to four key findings: (1) lack of visibility in cloud resources, (2) cloud computing complexity, (3) lack of security expertise, and (4) regulatory compliance and legal concerns. As more organisations are migrating more workloads to cloud-based resources, visibility into these resources and adequate security expertise of a CSU and CSP's staff becomes critical. Furthermore, more than three-quarters of the respondents considered compliance and audit preparation the main challenges of managing the security of their public cloud resources. The CSA stresses the need for organisations understanding how to leverage cloud platforms and how to use tools provided by the CSP. Furthermore, organisations should retain awareness of the latest developments in cloud service features. Automation is highlighted as aiding in the lack of expertise by, for example, detecting security gaps, compliance violations, and service misconfigurations.

3.4 Cloud Incident Handling

This section will categorise existing research on cloud incident handling according to the four phases of the NIST incident handling lifecycle (see Section 1.2). We provide a concise summary of each paper. Table 3.1 provides an overview of which papers describe what areas regarding cloud incident handling.

Table 3.1: Overview of discussed topics in related literature

	Preparation				Detection & Analysis		Containment, Eradication & Recovery		Post-Incident Activity		Digital Forensics			
Papers	Incident handling process	SLA	Technical	Reporting	Organisational	Technical	Cloud capabilities	Technical	Evaluation	Report sharing	Models	Legal	Technical	Organisational structure
Ab Rahman, Cahyani, and Choo [1]											x			
Ab Rahman and Choo [2]	x		x	x	x	x			x		x			
Ab Rahman et al. [3]											x			
Chung et al. [12]												x		
Dekker, Liveri, and Lakka [19]	x	x		x					x	x				
Dykstra and Sherman [21]												x		
Frøystad et al. [23]										x				
Grobauer and Schreck [24]	x	x	x			x	x	x					x	
Jaatun and Tøndel [28]	x									x				
Kent et al. [31]											x			
Martini and Choo [35]	x								x		x	x	x	
Mckemmish [36]											x			
Mogull et al. [39]	x	x	x		x	x	x	x	x					
Monfared and Jaatun [40]			x											
Ruan et al. [47]		x										x	x	x
TaheriMonfared and Jaatun [50]	x	x	x	x				x						
Zhang et al. [51]	x	x							x					
Zhang, Patwa, and Sandhu [52]										x				

3.4.1 *Preparation*

In this section, papers that cover points of attention for cloud incident handling in the preparation phase are discussed. The preparation phase aims to establish an incident response capability and prevent incidents by ensuring existing applications, systems, and networks are secure [13].

3.4.1.1 *Incident Handling Process*

Cloud computing has a significant impact on the incident handling process, technologies and methods. Therefore, it is imperative that CSUs establish a clear incident handling process [24, 2]. The first step is establishing a CSIRT [24]. These teams are responsible for determining what happened, which response needs to be taken and executing these responses [2].

Establishing a CSIRT

The establishment of a CSIRT is supported by the publication of Ab Rahman et al. which states that establishing a CSIRT is a key backbone in the preparation phase [2]. In their survey, 139 publications on incident handling and digital forensics are analysed aiming to contribute to the knowledge gap regarding incident handling in the cloud. They propose a conceptual cloud incident handling model by integrating digital forensics principles, the Capability Maturity Model Integration for Services (CMMI-SVC) [14], and the cost involved in the incident handling process. For example, investment costs are associated with the preparation phase such as setting up a dedicated department, establishing security objectives, and purchasing technology. To determine the effectiveness of the security investment risk can be incorporated through, for example, the Return On Security Investment (ROSI) formula.

Risk Assessment

A well-structured risk management methodology can help CSIRTs and management identify appropriate controls [51]. In the paper published by X. Zhang et al. [51], an information risk management framework covering several cloud services and deployment models is proposed. They provide information about the implementation and operation of risk management, which consists of three processes that are to be followed consecutively: (1) risk analysis, (2) risk assessment, and (3) risk mitigation [51].

The first process decides whether it is advisable to proceed. It consists of threat and vulnerability identification. The second process assesses the outcome from the risk analysis. This process considers the following four actions: (1) likelihood determinations, (2) impact analysis, (3) risk determination, and (4) control recommendations. In the final process, the CSP rates the risks by labelling them with possible actions (avoidance, reduction, retention, transfer, and acceptance). X. Zhang et al. [51] stress the importance of documenting risk treatment plans into SLAs because different cloud computing models handle vulnerabilities and threats differently.

ENISA's framework for reporting about major cloud security incidents is focused on providing government authorities with guidance on handling security incidents [19]. However, some of their recommendations to government authorities can be applied to organisations as well. They state that risk assessments should be conducted periodically to identify critical parts of the IT infrastructure. Additionally, this assessment identifies which core services are cloud-dependent [19].

ENISA further recommends to prioritise critical sectors and interdependencies, and to identify cyclic dependencies [19]. Some incidents might affect other infrastructure - power outages might affect ICT infrastructure - which in turn is necessary to resolve the incident [19]. However, it is difficult to assess what impact failing cloud computing services could have. Critical services should be targeted by translating ICT infrastructure risk assessments into thresholds for specific services [19].

Documentation of information

Additionally, CSUs should understand and document what information will be available and is needed during an incident [28, 39]. As noted above, different cloud computing models handle vulnerabilities and threats differently which makes it all the more imperative that CSUs know which data and logs will be available from which CSP. Additionally, CSUs should know which tools should be used [24, 39].

It is difficult to determine what data sources are necessary to successfully handle an incident. Therefore, it is imperative that information gathering does not start when first encountering a security incident [24]. The CSU needs to identify relevant data sources to add security-specific event sources which can be accomplished by analysing the CSP's support capabilities [24]. Furthermore, CSUs should understand the types of data that can be obtained, and where to obtain them. It is important to understand the format of the data obtained from the CSP for the data to work with the incident handling tools used by the CSU [24, 35].

Incident response selection

Tabletop exercises and threat modelling are highlighted by Mogull et al [39] to determine the most effective response to different types of attack. This should cover the different responses needed for IaaS, PaaS, and SaaS as well.

Response selection is supported by the publication of Ab Rahman et al. [2] which stresses the importance of response selection techniques. The use of these response selection techniques should ideally ensure that the incident response process can rapidly be deployed. They identify three response selection techniques: static mapping, dynamic mapping, and cost-sensitive mapping. Static mapping maps a predefined incident alert to predefined responses. This could be done by using probabilistic cognitive maps, or by applying ontology. While the static mapping ensures that a response can rapidly be deployed, it enables attackers to anticipate response actions as well. Dynamic mapping prevents this. In this mapping process, a variety of more advanced approaches is deployed. Risk assessment methods, machine learning, and game theory are mentioned as possible mapping strategies. However, the drawback when using dynamic response selection is that it does not consider the cost of damage and response. Cost-sensitive mapping balances damage and response cost. Four key factors that must be minimised in this mapping technique are as follows: (1) cost of implementation, (2) the level of resources that are needed, (3) time effectiveness, and (4) the cost of induced modification. Examples of cost-sensitive mapping are modelling dependency graphs of services, an adaption of the Return of Investment index, and a weighted linear combination [2].

CSP requirements

The aforementioned mainly focuses on CSU requirements, although many CSP requirements exist. Two of such CSP requirements are identified by Monfared and Jaatun [50] who describe an approach to handle compromised components in an IaaS cloud installation. Their research analyses two case studies based on an adapted form of the NIST incident handling guide [13] and the Cloud Computing model [37]. This analysis leads to two CSP requirements regarding preparing the incident handling: (1) provide a security service, and (2) provide information about their architecture.

CSP provided security service

Monfared and Jaatun posit that CSUs may not be interested in developing security mechanisms [50]. This issue can be alleviated by the CSP by providing a security service [24]. Security services developed by the CSP can be more reliable and less challenging to deploy [50]. Additionally, CSPs should establish incident reporting services including acceptance and forwarding of external reports [24, 50].

CSP provided details

However, some CSUs prefer to develop their incident detection and analysis mechanisms in which case it is necessary to know the details of the CSP's cloud service [50]. Additionally, handling an incident can be facilitated by having the necessary cloud service details. Therefore, CSPs should provide information about their architecture and infrastructure [24, 50]. Grobauer and Schreck further argue that CSPs should provide access to relevant data sources and open interfaces for incident data exchange, although this is complicated by the challenge to build and maintain the required infrastructure [24]. An additional challenge is the standardisation of event information, which is further addressed in Section 3.4.1.4.

3.4.1.2 *Service Level Agreement*

According to the CSA "the most important security consideration is knowing exactly who is responsible for what in any given cloud project" [39]. These roles and responsibilities are defined in a formal written agreement, which may be required by law, often called the Service Level Agreement (SLA). The importance of the SLA is supported by X. Zhang et al. [51], who state that until CSUs can easily switch between CSPs, allowing customers to overcome any issue, SLAs can be used to alleviate current concerns. In addition to describing roles and responsibilities, SLAs should cover measures to be taken in case of compromised or stolen data, the permitted and prohibited use of data, and expectations of both parties [39].

Different CSPs provide different services resulting in different relationships; using a CSP that offers a custom private cloud will have a different relationship than a CSP that offers a generic SaaS application [39]. Therefore, it is important for security teams to understand the terms defined in the SLA. By doing so, security teams are better prepared regarding what services they can utilise [39].

Forensic investigations

SLAs furthermore facilitate communication and collaboration in regards to forensic activities [47]. Therefore, they must clarify procedures that are to be followed during forensic investigations. Ruan et al. [47] define three terms that should be included in SLAs: (1) the services provided, techniques supported, and access granted by the CSP during investigations, (2) trust boundaries, roles, and responsibilities between the CSP and CSU, and (3) the process for conducting investigations that span multiple jurisdictions without breaching laws, regulations, and policies. It is important that the chain of separation is standardised and that cryptographic keys are formalised in (service level) agreements between CSPs, CSUs and law enforcement [47]. During the process of drafting the SLA, legal teams should be involved to ensure all jurisdictions in which a CSP might operate are covered to prevent jurisdiction issues in a later stage (see Section 3.5.2) [47].

Roles and responsibilities

Defining roles and responsibilities becomes more important in a cloud computing setting and should be addressed in the SLA [24, 39, 47, 50, 51]. This is due to multiple organisations being responsible for implementing and managing different parts which causes security responsibilities to be divided across these multiple organisations as well [39]. Further complications regarding these roles arise when cloud brokers or other intermediaries are used [39]. To prevent complications, CSUs should build a responsibilities matrix to document in which they document who implements what part in what manner [39].

Another model has been proposed by X. Zhang et al [51]. They present an information risk management framework covering several cloud services and deployment models. This framework describes three phases, each consisting of several processes: (1) architect and establish, (2) implement and operate, (3) monitor and review. Part of the architect and establish phase are *selecting relevant critical areas* and *strategy and planning*. These processes are crucial to designing and planning effective security risk management and ensure that roles and responsibilities are clearly defined, critical areas to focus on are identified, and management can make clear choices for resource allocation. X. Zhang et. al further recommend incorporating the result of risk assessments into SLAs [51].

Reporting requirements

Grobauer and Schreck [24] cover multiple phases in their overview of challenges and approaches in the cloud. While they mainly cover challenges and solutions in the detection, analysis, and incident response phase, they briefly focus on defining reporting requirements in SLAs. They state that an SLA must provide a well-defined incident classification scheme. Furthermore, it should cover reporting obligations such as what is reported and the response time that can be expected [24]. This can be aided by the CSP providing an incident detection service (see the CSP requirements in Section 3.4.1.1).

The European Union Agency for Cybersecurity (ENISA) describes a framework for reporting about major cloud security incidents [19]. Their goal is to provide government authorities with guidance on how to implement reporting of network and information security (NIS) incidents. They surveyed stakeholders to identify views and best practices in incident reporting. To illustrate the cloud incident reporting process, four use cases are depicted in which they present the challenges and solutions for implementing incident reporting per scenario. While their framework focuses on the details of reports and the reporting process - such as root cause, systems affected, and mitigation actions - they emphasise that reporting only occurs when the SLA obliges the CSP to do so. Therefore, they recommend CSUs to address the specific requirements of incident reporting in their SLAs. Further details of the requirements pertaining to reports are described in Section 3.4.1.4.

3.4.1.3 *Technical*

As noted above, CSUs might prefer to develop their security mechanisms [50]. Two different development options are identified by Monfared and Jaatun; either base security mechanisms on reports from various sources - incident reports, end-users' reports, or third parties - or base them on CSPs' APIs [50]. This further shows the importance of CSUs determining what data sources are relevant for incident handling [24, 39].

CSP provided precursor and indicator sources

For CSUs to be able to base their security mechanisms on CSPs' APIs, the CSP should allow for this functionality. Monfared and Jaatun recommend that CSPs should develop such APIs that provide event monitoring capabilities and forensic services [50]. Furthermore, CSPs should provide precursor and indicator sources. These sources are a result of mechanisms that can be implemented by CSPs such as intrusion monitoring sensors, log files, and firewall statistics [50].

Cross-layer security approach

Monfared and Jaatun's earlier work intro distributed cloud environments determines a cross-layer security approach to be effective in a distributed cloud environment [40]. Allowing CSUs to implement CSPs' security agents into their resources facilitates such a cross-layer security approach [50]. It allows the CSU to know what information has been disclosed while neither the CSP nor the CSU needs to know details about each other's infrastructure or architecture design [50].

Proper configuration and architecture

Implementing proper configuration and architecture is necessary to support incident response. CSA's security guide provides five configuration implementations that can support cloud incident response: (1) enable instrumentation (such as logging) of which a backup is stored in a secure location, (2) utilise isolation, (3) use immutable servers, (4) implement application stack maps, and (5) perform threat modelling and tabletop exercises to determine the most effective containment strategies [39]. Implementations should be tested with the CSP, to ensure it functions as intended [39].

Logical security control

Finally, Ab Rahman et al. state that logical security control is crucial [2]. Examples of security controls are malware protection, vulnerability assessments, firewall implementation, and network monitoring [2]. Their conceptual cloud incident handling model addresses hard- and software examples corresponding to these security controls, although not all pertain to all cloud models. Recommended hardware are routers and backup servers. Recommended software technologies are, for example, firewalls, System Information and Event Management (SIEM), Intrusion Detection and Prevention Systems (IDPS) and anti-virus software [2].

3.4.1.4 *Reporting*

Cloud systems are comprised of several actors. This may lead to poorly coordinated activity correlation or it can cause misdirection in incident reporting. Therefore, CSUs should define a clear incident reporting strategy [2]. A widely known cloud reporting framework is ENISA's Cloud Security Incident Reporting Framework [19]. Their framework focuses on government authorities, but its recommendations can be utilised by organisations as well. They make two notable recommendations; (1) national reporting schemes for NIS should be set up by authorities, and (2) attention should be paid to the harmonisation of incident reporting.

National reporting schemes

ENISA is strongly in favour of national reporting schemes. They argue that authorities should provide this possibility as national reporting can be used for better understanding of security and resilience [19]. Because CSPs often work across borders, customers and regulations from many different countries may be involved. This might lead to unnecessary costs [19]. Therefore, ENISA stresses the necessity of harmonised incident reporting.

Harmonisation

While incident reporting legislation must be harmonised, several more areas should be addressed. The need for a common reporting template is evident and could be a starting point towards harmonising incident reporting [19]. In addition to alleviating the CSPs reporting workload due to standardisation, report sharing becomes more effective [19]. For example, vocabulary, format, and terminology could be standardised [19]. Details and recommendations about report sharing are further discussed in Section 3.4.4.2.

Harmonisation of incident reporting is supported by Monfared and Jaatun, who argue that a standard communication protocol is required to achieve systematic incident detection and analysis mechanisms [50].

Content of reports

As noted before, the requirements of CSPs' reports to customers should be described in the SLA. Furthermore, ENISA recommends having at least the following included in reports [19]:

- Technical information
- Duration of the incident
- Area impacted
- Remediation time
- Systems affected
- Root cause
- Mitigation action
- Confidential information

Additional suggestions to include in the contents of reports are: key findings from forensic analysis, documentation compiled during the incident, and the analysis methods and techniques used [2].

3.4.2 *Detection and Analysis*

This section discusses the detection and analysis phase. In this phase, it is determined whether an incident has occurred. This begins when anomalous behaviour is flagged, either automatically by a tool or manually by people. Detection occurs when there are signs of an incident. These signs can be categorised as precursors and indicators. Precursors are signs that an incident might occur in the future while indicators are signs that an incident might be occurring now. Analysing the flagged incident is important to determine if this behaviour is a valid threat and what priority it should receive [13].

Challenges

In previous sections, many issues regarding cloud incident handling were described. Issues pertaining to the detection of incidents are for example: no, or insufficient, access to CSP sources, the inability to add specific security measures, and the misdirection of reports [24]. Cloud analyses can be complicated due to a lack of knowledge on the architecture, unclarity about the division of responsibilities, and missing access to relevant data sources [24]. These could be solved by the CSP by providing the CSU with access to relevant sources, implementing an IDPS, and improving communication [24]. It is important that CSUs identify what logging is needed and if the CSP is willing to provide these [39].

3.4.2.1 *Organisational*

Risk management is key to estimating the damage following an incident and incident prioritisation [2]. When an incident is detected, the evidence collection process will be started by forensic examiners. This process will further determine the incident's severity level such that an appropriate escalation strategy can be assigned [2]. Examples of costs associated with this phase are wages, the acquisition of evidence, and digital forensics software and hardware [2].

Legal team

Although not every incident will lead to legal action needing to be taken, CSUs should start consulting their legal team to understand possible issues in the post-incident activity phase [39]. More details on obtaining legally acceptable evidence can be found in Section 3.5.2.

3.4.2.2 *Technical*

Live compromised systems and log files analysis techniques

Improving techniques for analysing live compromised systems and log files are deemed priorities [24]. With cloud computing, forensics must often be performed on running systems, in which case valuable information can be obtained through live analysis. However, according to Grobauer and Schreck, there are no suitable approaches yet [24]. Furthermore, they state that, especially for PaaS and SaaS, important evidence sources are CSP's log files. Therefore, it is essential to improve the generation and analysis of logging [24]. Other potential data sources are CSU's devices and off-site CSP data centres [2].

Knowledge gaps

Furthermore, monitoring the cloud management plane is important to identify changes in the environment and configuration. Knowledge gaps can occur when the CSU is missing information, either because the CSU is not aware that the CSP can provide this or because the CSP is not able to provide it [39]. For example, network logs might only be flow records instead of a full packet capture. Furthermore, the information that a CSP provides might not meet legal standards [39].

Cloud capabilities to leverage

Several cloud capabilities could be leveraged in this phase. Automation offers functionalities such as creating a snapshot of the storage of the virtual machine, capturing any metadata at the time of the alert, and pausing the virtual machine to retain the memory state [39]. To identify the extent to which the cloud platform was affected, cloud platform capabilities can be used such as analysing network flows to determine if network isolation was successful, examining configuration data to identify similarly affected instances, and reviewing data access logs to determine whether the attack affected the cloud platform itself [39].

3.4.3 *Containment, Eradication, and Recovery*

The impact of an incident can be mitigated by containing it. Containment ensures that the threat does not infect other systems. After the threat is contained, it needs to be eradicated from compromised assets. Finally, the normal operation of assets is restored. This might involve actions such as: installing patches, changing passwords, and replacing compromised files with clean versions. In this phase, data that helps resolve the incident, and aids the possible legal process after is collected. This phase often cycles back to the detection and analysis phase, to for example identify whether the threat has spread to other systems. [13]

3.4.3.1 *Cloud Capabilities*

Grobauer and Schreck [24] find it difficult to provide general advice regarding this phase because every incident happens under different circumstances with different attack vectors. Instead, they examine frequent scenarios divided into issues and solutions for IaaS and SaaS/-PaaS. They describe challenges introduced by using a cloud environment - such as configuration capabilities offered by the CSP - and opportunities that a cloud environment offers. For example, the elasticity of a cloud environment is described as an opportunity, where the resources of an asset can be expanded or limited based on the active threat [24].

Technical capabilities

Further cloud capabilities aiding this phase are: enabling infrastructure to be quickly rebuilt in a clean environment, snapshots for rollbacks of virtual machines, and API calls for changing virtual networks or machine configurations. However, CSUs using SaaS and PaaS are cautioned, as those tend to be limited in functionality. This causes the CSU to be more dependent on the CSP [39].

3.4.3.2 *Technical*

The first step after an incident has been identified should be ensuring that the exploit path is closed. There is no need to immediately eradicate because the cloud offers more flexibility in this phase than during on-premise incident response [39]. The CSU needs to make sure in their eradication and recovery step that their data is purged from the attacker's activity. The CSP can aid the incident response process by for example providing the ability to configure networking (IaaS), access to snapshot features (IaaS), and direct read and write access to customer data (PaaS and SaaS) [24].

Monfared and Jaatun [50] adapted several actions described in the NIST guidelines [13] using the cloud model presented by Mell and Grance [37]. The case studies they analyse, and by extension their recommended responses, consist of two types of incidents: a compromised compute worker (compromised via unauthorised access and malicious code), and a bogus component.

Specific actions

Monfared and Jaatun provide four adapted containment actions: (1) identifying and isolating other infected hosts, (2) blocking particular hosts, (3) soliciting user participation, and (4) disabling services. Furthermore, they describe two eradication actions: (1) disinfect, quarantine, delete, and replace infected files, and (2) mitigate the exploited vulnerabilities for other hosts within the organisation. Finally, they describe two recovery actions: confirm that the compromised systems are functioning normally, and implement additional monitoring to look for future related activity if necessary. Multiple approaches - such as filtering, disinfecting components, and component authentication - are described in detail, accompanied by an overview of advantages and disadvantages, and how these approaches can be implemented [50].

3.4.4 *Post-Incident Activity*

The final phase in the NIST incident response lifecycle is post-incident activity. In this phase, incident response teams reflect on the incident to evaluate their incident handling process. This determines what occurred, what was done to mitigate the incident, and what should be done in the future. Furthermore, it identifies developments in either technology or threats [13].

3.4.4.1 *Evaluation*

The post-incident activity phase requires a high degree of proactiveness from relevant personnel [2]. They should take the initiative in this phase to recognise and defend against new threats, and to improve existing protection measures.

Learning in organisations

Ab Rahman et al. recognise a lack of research into incident learning [2]. They briefly mention organisation learning theory and ontology as concepts that could aid organisations in the post-incident phase.

CSP involvement

Additionally, CSPs can aid in the evaluation process. Martini and Choo argue that preservation of digital evidence in cooperation with the CSP is one of the most critical steps in digital forensics investigations [35]. The security guide of the CSA [39] recommends that CSUs should work with the CSP and the incident response team to evaluate the handling of the incident. A key point is the limitations that were encountered, and how these can be addressed in the future. Although they emphasise the difficulty of adapting SLAs, the CSU should try to negotiate with their CSP when agreements have not been met such as response time, data provision, and other support [39].

Re-evaluate risk assessment

In addition to reviewing the SLA, CSUs should reevaluate their risk assessment whenever significant changes have been made to improve the existing security strategy [39]. Therefore, internal audits are necessary to determine if the risk assessments need to be modified. [51, 19]. This applies to both CSP and CSU. Costs that can be associated with this phase are direct or indirect losses due to the security incident [2].

3.4.4.2 *Report Sharing*

Sharing incidents or summaries of security incidents will lead to a discussion of best practices and improve security incident handling [19]. Furthermore, sharing incident reports improves accountability of CSPs [28, 23]. Frøystad et al. state that the introduction of new regulations such as the GDPR increases the need for effective incident information sharing [23]. Therefore, they propose a simplified method of incident sharing. While their approach does not ensure that all CSPs and CSUs involved understand all the available information, it ensures that every party involved understands the information that pertains to them [23].

Governmental responsibility

ENISA's framework for reporting about major cloud security incidents [19] recommends - based on their surveys with experts - that authorities should be responsible for sharing incident reports across borders and should act as a filter (which is supported by 80% of their respondents, n=40) [19]. This recommendation builds on their recommendation to harmonise incident reporting and national reporting schemes as described in Section 3.4.1.4.

Community-based secure information and resource sharing

According to Y. Zhang et al. it is likely that cyber attacks will happen to other organisations that share the same cloud platform [52]. By sharing incident reports with other organisations, this risk can be mitigated. Their paper focuses on a detailed secure access control model for AWS that can be used to securely share incident information. In contrast to other papers, where it is not defined whom to share reports with, this paper focuses on community sharing. This means that incident reports are shared among organisations who use AWS [52].

3.5 DIGITAL FORENSICS

Digital forensics is the process of identifying, preserving, analysing, and presenting digital evidence in a manner that is legally acceptable according to the definition presented by McKemmish [36]. It does not necessarily aid in the incident handling process, but it is important to obtain trustworthy evidence that will be admissible in court. Due to the distributed nature of cloud environments, acquiring and analysing digital evidence is more difficult in a cloud computing environment than for traditional server-based systems [35].

3.5.1 *Models*

The Digital Forensic Research Workshop (DFRWS) Investigative Model [8], the Abstract Digital Forensics Model (ADFM) [45], and the Integrated Digital Investigation Process (IDIP) [10] are well known digital forensics models. However, they do not cover cloud environments.

Integrated digital forensics framework

The conceptual cloud incident handling model proposed by Ab Rahman et al. integrates forensic activities into each phase of the incident handling model [2]. In the preparation phase, forensic readiness is the main digital forensics activity. This state in digital forensics prepares for further digital forensics activities by determining in advance what evidence is required and how to obtain it. They describe several forensic readiness activities, recommend having dedicated digital forensic workstations and software, and recommend preparing an incident handling strategy. In the detection and analysis phase, the forensic analysis takes place. They mention examples of reports that support the forensic analysis process such as incident report forms, digital evidence analysis reports, and incident management action reports. In the containment, eradication, and recovery phase, the evidence is handled. Finally, in the post-incident activity phase, an indepth analysis, evidence retention, and the presentation of evidence are the main digital forensics activities.

In subsequent studies, Ab Rahman et al revised their model [3] and performed a case study to validate the model [1]. The Conceptual Forensic-by-Design Framework they devised describes how to integrate digital forensic requirements into cyber-physical cloud systems to ensure that forensic investigations can be facilitated. The framework consists of six factors: (1) risk management principles and practices, (2) forensic readiness principles, (3) incident handling principles and practices, (4) laws and regulations, (5) cyber-physical cloud systems hardware and software requirements, and (6) industry-specific requirements. They demonstrated how this framework can be used in a case study of controlled experiments in a cloud setting [1]. They conclude that the model is useful for both CSUs and CSPs and that the model enables CSUs to undertake incident investigations.

Integrated iterative digital forensics framework

Another framework is the digital forensic framework proposed by Martini and Choo [35]. It is based on two of the most widely used forensic frameworks by McKemmish [36], and NIST [31]. McKemmish's digital forensics framework consists of four key elements: the (1) identification, (2) preservation, (3) analysis, and (4) presentation of digital evidence. Martini and Choo [35] argue that the extraction and processing described by McKemmish's framework in the analysis phase are critical and potentially time-consuming in a cloud computing environment. Therefore, they represented the extraction of evidence as a separate step [35].

NIST's framework consists of the following four phases: the (1) collection, (2) examination, (3) analysis, and (4) reporting of digital evidence. They discuss identification and preservation as part of the collection phase. Martini and Choo suggest that the identification of cloud computing as evidence source and preservation in cooperation with the CSP are the most critical steps in digital forensics investigations and should be conducted simultaneously [35].

This resulted in the following framework presented by Martini and Choo: (1) evidence source identification and preservation, (2) collection, (3) examination and analysis, and (4) reporting and presentation [35]. Its key difference with the frameworks of McKemmish and NIST is the iteration phase. This phase is possible when evidence of cloud computing use is discovered in the examination and analysis phase. A second iteration will then start at the first phase via the CSP. If further evidence sources are then identified in the examination and analysis phase, another iteration would start.

3.5.2 *Legal*

Another important factor in digital forensics is the legality of the obtained evidence. This is more complicated in a cloud computing environment where legal issues become exacerbated [47].

Distributed nature challenges

For example, the collection and preservation of evidence could prove to be difficult as the evidence could be located in another jurisdiction or spread over multiple systems [35]. To alleviate this problem, Ruan et al. recommend that regulations and agreements should be developed to ensure laws and regulations in the jurisdictions where the data is stored are not broken [47]. Law enforcement agencies might have to rely on CSPs to provide evidence but this could break requirements set by courts [35].

Acquiring trustworthy and legal forensic evidence

To best adhere to such requirements set by courts, Dykstra and Sherman analysed technical and trust issues that occur when acquiring forensic evidence from an IaaS cloud computing environment [21]. Their paper focuses on acquiring forensic evidence that can be proven trustworthy in court, and analyses forensic acquisition tools in an Amazon EC2 setting. First, they identify six different layers in a cloud computing environment, accompanied by an acquisition method, and the trust required. The six cloud layers are as follows: (1) network, (2) physical hardware, (3) host OS, (4) virtualisation, (5) guest OS, and (6) guest application/data. They argue that only using technology to obtain forensic evidence is insufficient to obtain trustworthy data. Therefore, they provide four alternatives that combine technology and CSP support: (1) trusted platform modules, (2) the management plane, (3) forensics-as-a-service, and (4) legal solutions. Their recommendation is to use the management plane to acquire forensic evidence, as this balances speed and control with trust.

Chung et al. [12] provide a detailed procedure for investigating a cloud storage service. This procedure shows a workflow to obtain forensic evidence and takes legal issues into account, such as considering search and seizure warrants, and jurisdictions. Furthermore, this paper provides an extensive analysis of artifacts and presents methods for collecting and analysing forensic evidence on multiple devices using multiple cloud storage services.

3.5.3 *Technical*

As noted before, obtaining forensic evidence in a cloud computing setting is more difficult. Therefore, Martini and Choo [35] stress the need for metadata retention and recommend potential evidence sources such as centralised auditing, extensive logging, and file integrity checking. Teams must understand what data can be extracted, requested, or converted to a format that can be used with traditional digital forensics tools [24, 35].

Cloud characteristics to consider

Three cloud characteristics are highlighted by Ruan et al. that require consideration: rapid elasticity, resource pooling, and virtualisation [47]. Tools and procedures - adapted to these characteristics - should be developed that are elastic, segregate forensic data between multiple tenants in various cloud structures and locate forensic data with timestamps [47].

Proactive measures

Several proactive measures exist that can be taken to facilitate cloud investigations. These are, for example, regular snapshot retention, continually authentication and access tracking, and object-level auditing of all accesses [47].

3.5.4 *Organisational Structure*

Multiple entities may be involved in cloud forensic investigations. Ruan et al. [47] provide an overview of such entities and their dependencies on external parties such as academia, third parties (e.g. for auditing), and law enforcement. They describe five roles that should be fulfilled to establish a cloud forensic capability: (1) investigators, (2) IT professionals, (3) incident handlers, (4) legal advisers, and (5) external assistance [47].

3.6 OVERVIEW OF RECOMMENDATIONS

Section 3.3 to Section 3.5 provided an overview of work related to this research. This section provides an overview of recommendations related to cloud incident handling specifically.

3.6.1 *Cloud Incident Handling*

Considering the focus of this research on the best practice in cloud incident handling, it is important to know what challenges, opportunities, and solutions have been covered on this topic in the existing literature. Furthermore, this literature can be compared to the results from RQ1.2, to identify the overlap and gaps between practice and literature. Most studies have based their work on the NIST incident response lifecycle (see Figure 1.1) and provide cloud-specific recommendations that complement traditional incident handling methods. An overview of the main recommendations for each stage in the NIST incident response lifecycle is presented in Table 3.2. These are grouped per incident handling phase and whether they apply to CSUs, CSPs, or in general.

3.6.2 *Digital Forensics*

Digital forensics is important in incident handling, as this ensures evidence is legally admissible in court. Cloud environments are more complex than traditional on-premise systems due to their distributed nature. Therefore, digital forensics must adapt as well to be able to ensure the trustworthiness of evidence obtained in cloud environments. Many publications focus on specific technical solutions. However, this section focused on models, challenges, and opportunities regarding digital forensics in the cloud to identify how digital forensics must adapt. The main findings are as follows:

1. Acquiring trustworthy evidence that is legally admissible in court is more difficult in a cloud computing setting.
2. The collection and preservation of evidence potentially become more difficult when confronted with multiple jurisdictions and multiple systems.
3. Teams should understand what data can be extracted and requested, and how to convert this data to a format that can be used in traditional digital forensics tools.
4. Forensic investigators might have to rely on CSPs to provide evidence. However, this could lead to the obtained evidence rendered inadmissible.
5. CSUs should pay attention to the following in the SLA with the CSP:
    - The services and access to data that are provided by the CSP.
    - The roles and responsibilities of both CSU and CSP.
    - The process for conducting investigations that span multiple jurisdictions.

Table 3.2: An overview of the recommendations from literature regarding cloud incident handling

	CSU	CSP	GENERAL
		PREPARATION	
Incident handling	Establish an incident handling process according to a risk assessment. The following should be clear: •The relevant data sources •The information that can be provided by the CSP (additionally the CSU should understand the content and format of the provided data) •The services and functionality CSPs and cloud platforms provide •The necessary tools •The communication channels that are to be used •The responses that correspond to incidents	Provide a security service Provide information about their architecture / infrastructure Establish an incident reporting service (which should allow external reports)	

This table continues on the next page

	CSU	CSP	GENERAL
SLA	Pay attention to the following in the SLA with the CSP: •The roles and responsibilities of both CSU and CSP •Measures to be taken in case of compromised or stolen data •Permitted and prohibited use of data •The forensic investigation process •The incident reporting process •The response time that can be expected		
Technical	Ensure proper configuration and architecture	Develop APIs that provide event monitoring capabilities and forensic services	
Reporting	Identify all necessary elements required in incident reports		Governments should establish national reporting schemes An independent standard communication protocol should be developed to harmonise information exchange
		DETECTION AND ANALYSIS	
Organisational	Confer with the legal team to identify post-incident issues (such as evidence admission in court)		

This table continues on the next page

	CSU	CSP	GENERAL
Technical	Incident response actions that utilise cloud platform capabilities are: •Analysing network flows to determine if network isolation was successful •Examining configuration data to identify similarly affected instances •Reviewing data access logs to determine whether the cloud platform itself is affected	Provide CSUs with information necessary in the incident handling process (information about architecture, access to data sources)	Automation can be beneficial in this phase (snapshots, capturing metadata, pausing VMs) Improve techniques for analysing live compromised systems and log files
		CONTAINMENT, ERADICATION, AND RECOVERY	
Cloud capabilities			The cloud offers more flexibility in this phase than with on-premise incident response. Technical capabilities that aid this phase are: •Infrastructure can be quickly rebuilt in a clean environment •Snapshots facilitate rollback functionality •API calls can be used for changing virtual networks of machine configurations
Technical		Provide the ability to configure networking (IaaS), access to snapshot features (IaaS), and direct read and write access to customer data (PaaS and SaaS)	Response selection techniques contribute to a rapid incident response

This table continues on the next page

	CSU	CSP	GENERAL
		POST-INCIDENT ACTIVITY	
Evaluation	Work with the CSP and the incident response team to evaluate the incident handling process Try to renegotiate the SLA with the CSP when agreements have not been met. Reevaluate the risk assessment whenever significant changes have been made bases on internal audits	Reevaluate the risk assessment whenever significant changes have been made bases on internal audits	
Report sharing			Share incident reports (across borders) to improve security incident handling and hold CSPs accountable

4 CLOUD INCIDENT HANDLING IN PRACTICE

This chapter presents the results from analysing cloud incident handling in practice. Section 4.1 discusses the differences described by participants. These differences are accompanied by challenges, which are discussed in Section 4.2. Section 4.3 describes cloud incidents encountered by organisations. Section 4.4 and Section 4.5 describe the results from the interviews, discussing the cloud incident handling strategies that are used by Dutch organisations. This chapter is concluded by Section 4.6, which provides an overview of the findings.

4.1 DIFFERENCES BETWEEN CLOUD AND ON-PREMISE INCIDENT HANDLING

Level of control

The difference between cloud incident handling and on-premise incident handling most often mentioned by interviewees is the different levels of control. Where organisation retain full control over on-premise environments, more control is shifted to the CSP with each cloud service level. More information on the different cloud service levels is provided in Section 3.1. Multiple interviewees express the diminished level of control in the following sentiment: *"you cannot walk over and pull the plug"*. Although CSUs relinquish some degree of control with each service model, one interviewee mentions considering their level of control in IaaS and PaaS environments to be almost equal to their level of control in on-premise environments. However, they mention that CSPs prohibit CSUs from looking behind the scenes which they consider a limitation. The different level of control pertains to data storage as well. One interviewee notes it is more difficult for CSUs to control where their data is stored, making cloud environments less transparent than on-premise environments.

Ease-of-use

The reduced level of control contributes to an experienced ease-of-use as well. On-premise environments need to be installed and maintained by organisations themselves. This often involved cost in storage space, extinguishers, and cooling installations. Additionally, updates need to be installed by organisations themselves as well. If these are not installed, the system could become vulnerable and thus compromised. However, this is less of a concern in cloud environments. Interviewees mention cloud services are easy to purchase, expand, and upgrade. In addition, the responsibility for maintaining - and in some cases updating - the cloud service is shifted from the CSU to the CSP. One interviewee remarks that CSUs will not notice issues occurring at the CSP's datacenters.

Log retention

In addition to the experienced ease-of-use, interviewees highlight cloud capabilities that aid the incident handling process. For example, multiple interviewees note that in on-premise environments, log retention sometimes is either not configured, or has a short retention period. However, default log retention in cloud is often ninety days and often enabled by default. This considerably assists a cloud investigation, as investigation teams have more data to analyse. One interviewee mentions that a short log file retention in on-premise environments is often due to high storage costs, whereas cloud services often store these at no added cost.

Cost

Multiple interviewees doubt whether cloud services are cheaper than on-premise environments. Many costs associated with on-premise environments are eliminated such as equipment, storage, and certain staff. In addition, interviewees consider the costs associated with cloud services to be more clear than in on-premise environments, as cloud services are often purchased based on the number of users. However, upgrading the subscription-based service can be expensive and CSPs charge for some features, such as downloading log files.

4.2 CHALLENGES IN CLOUD INCIDENT HANDLING

The aforementioned differences require an adapted incident handling process and result in many challenges with cloud incident handling as well. The main challenges mentioned by interviewees are discussed in the following paragraphs.

Lack of technology and process-related knowledge

The interviews indicated that companies experience *a lack of technology and process-related knowledge* regarding cloud incident handling. First, companies indicate that they found it difficult to adapt their incident handling process to cloud environments. Second, some interviewees indicate their incident response teams are lacking technical knowledge - or do not possess knowledge at all - of the different cloud platforms that are used. Workshops exist that try to bridge this gap, but these are considered expensive, time-consuming, and do not match the need of each company individually. Third, companies often do not know which specific tooling is available at the CSP as well as tools they could purchase themselves. Finally, a lack of knowledge can occur when deciding to migrate processes to the cloud. Companies might have unrealistic expectations of the services a CSP provides, which results in an inadequate, or non-existing, response processes. This manifests itself in the misconception that cloud services are more secure by default. For example, if CSUs migrate a vulnerable version of an operating system to the cloud, they might think it becomes secure. However, it remains vulnerable.

Insufficient overview of information

The second issue that interviewees reported is *an insufficient overview of information*. This information includes, for example, the architecture of the cloud environment, what services are used in the cloud environment, and whom to approach (both at the CSP as well as the user of the cloud environment). Information, and retaining its overview, is necessary to successfully and timely respond to an incident. Another factor contributing to an insufficient overview of information is not knowing where to find the information. Furthermore, interviewees reported issues with an abundance of information. For example, one interviewee mentions that if a CSU uses multiple cloud services, this results in multiple monitoring solutions that the CSU should monitor as well. An abundance of information pertains to the availability of documents as well. Because teams lack knowledge regarding the cloud environments, documentation is necessary to successfully respond to incidents. However, documentation is often extensive, making it difficult to obtain the necessary information.

Lack of visibility

The third issue reported regarding cloud incident handling is *a lack of visibility*. Multiple interviewees mention having difficulty with obtaining and maintaining visibility into their cloud environments. Interviewees indicate that the lack of visibility is more prevalent in cloud environments than in on-premise systems because they are more restricted in their cloud environments. This leads to incident response teams being unaware of incidents occurring, impacting the security of assets.

Inadequate design and road map

The fourth issue identified is *an inadequate design and road map*. While companies often transition to the cloud because of cost-saving measures, additional security features are often expensive. This can lead to companies not purchasing these features or making changes during the transition when these are deemed too expensive, impacting the incident response capabilities. The latter often occurs without re-evaluating their risk analysis, which coincides with the issues regarding the lack of overview of information and a lack of visibility.

Dependency on the vendor

The fifth issue that interviewees report is a *dependency on the vendor*. An example is vendor lock-in, which occurs when a company cannot, easily, switch from CSP. Companies often do not consider the situation where their CSP is acquired by another party that does not fit the needs of the company. In this situation, vendor lock-in occurs because migrating environments to another CSP is a time-consuming and expensive process. Another issue arises when the company's CSP goes out of business. Companies might not be aware to what extent their business continuity is impacted, and whether contingency measures exist or not. Furthermore, coinciding with the issue regarding a lack of knowledge, CSPs might not be able to support the CSU's incident handling process in case of an incident. This could be due to a lack of knowledge or resources at the CSP, or because the CSP cannot be contacted at all. Additionally, when a CSP assigns the incident a low priority, companies might not receive the necessary support in time to successfully resolve the incident. This can lead to disruptions in the business continuity, a loss in revenue, and high damages. Finally, CSUs are dependent on their CSP for feature development in their cloud environments and access to security features. For example, one interviewee mentions having developed security tools that interact with a CSP's API, and this CSP suddenly blocked this access.

Conducting an investigation

Finally, the aforementioned differences and challenges result in difficulties with *conducting an investigation*. Due to the shift in control from the CSU to the CSP, CSUs cannot physically access environments. Therefore, it is more difficult to acquire evidence. In addition, security teams are dependent on the CSP, and therefore on the resources a CSP provides. For example, if a CSPs log file service becomes, temporarily, inaccessible, incidents cannot be detected or analysed. Additionally, one interviewee mentions having had difficulty processing CSP's log files to their desired formatting in order to enable alerting.

4.3 CLOUD INCIDENTS ENCOUNTERED

Almost all interviewees report using the cloud service Office 365. Additionally, some mention using AWS, Microsoft Azure, or hosting their own cloud environment. Multiple interviewees mention utilising multiple cloud services, therefore adopting a multi-cloud philosophy. Although some cloud services have not been mentioned by (all) interviewees, it should not be assumed these are not used by them as well.

Interviewees do not often encounter cloud incidents. Moreover, one interviewee mentions not knowing how they would respond to encountering a cloud incident. Cloud incidents are caused by various reasons. However, interviewees mention that the few cloud incidents they do encounter are often not high-profile or advanced, but caused by human errors such as configuration issues, or by sharing links that should not be shared. The interviews resulted in four types of incidents that are often encountered: (1) unauthorised access, (2) Business Email Compromise (BEC), (3) phishing, and (4) ransomware. A trend described by one interviewee is that cloud attacks often are large-scale attacks focusing on specific cloud services instead of targeting organisations. Large-scale attacks are easier as cloud services often have many users, such as Office 365. However, the advantage of this is that large scale attacks can be detected by CSPs more easily. The four types of cloud incidents observed by interviewees are briefly discussed in the following paragraphs.

Phishing

Phishing is an incident where attackers are, for example, able to gain access to the environment by sending emails that appear legitimate. Attackers can, for example, make these phishing emails appear more legitimate by purchasing domain names via anonymous registrars. According to an interviewee, attackers specifically focus on obtaining Office 365 credentials.

Business Email Compromise (BEC)

In BEC attacks, attackers gain access to email accounts - for example as a result of phishing or password spraying - and use these legitimate accounts to obtain information or money. For example, attackers can compromise a supplier and subsequently monitor payments. They can use the legitimate email account to deceive customers of the supplier into paying them money by sending them, legitimate, invoices where the bank account number is changed to one belonging to the attacker.

Unauthorised access due to wrong configurations

Unauthorised access can occur as a result of configuration errors as well. Interviewees mention organisations that accidentally allow their employees to access features or storage spaces of other departments in their cloud service. Additionally, wrongly configured cloud environments could lead to cloud services that are unintentional remotely accessible. For example, one interviewee describes Amazon S3 data buckets that are often accidentally connected to the internet, allowing unauthorised access. They caution that unauthorised access can also happen due to cloud services being wrongly configured by the CSP.

Ransomware

Unauthorised access can lead to organisations being infected with ransomware which has been mentioned by multiple interviewees. One interviewee remarks that, although it does not happen often, cloud services may become infected. They mention this often happens due to attackers being able to log in via remote desktop protocol (RDP). An other interviewee describes an incident where an entire environment was compromised with ransomware while not having back-ups.

4.4 CLOUD INCIDENT HANDLING

4.4.1 *Preparation*

4.4.1.1 *Risk Management*

Organisations must shape their cloud incident handling process based on the organisation's needs. Part of this is determining which risk is acceptable. Therefore, organisations should conduct a risk assessment before they start using the cloud.

Design

There are multiple risks involved with cloud computing, depending on the design of the cloud environment. It is important to carefully consider the type of subscription an organisation purchases. As an example, one interviewee mentions Microsoft's Business Premium accounts where users receive access to all applications. Because data is stored on multiple locations, due to automatic synchronisations between Microsoft Teams and SharePoint, they deem it easier for users to make mistakes. Additionally, organisations should evaluate available subscriptions to determine which is the best fit. For example, some subscriptions come with security features that might not be necessary for an organisation. While it might seem a good idea to purchase a subscription that includes incident response, one interviewee mentions the situation where the CSP assigns a low priority which leaves the organisation vulnerable longer. Therefore, it is important to design the organisation's cloud environment well, including the type of subscription purchased.

Exit strategy

While organisations should carefully evaluate their cloud environment subscriptions, it is important to think about an exit strategy as well. CSPs could collapse, change their pricing model, or become compromised. Organisations should be prepared to move their assets, although one interviewee cautions this is unrealistic as this is nearly impossible due to it being a time and money consuming process.

Availability

Another risk to consider is the availability of the cloud. Although cloud environments provides many benefits, such as flexibility, business continuity is impacted when the cloud environment becomes unavailable. An organisation should determine what the acceptable downtime of the cloud environment is, as often no guarantees can be given by the CSP. Therefore multiple interviewees recommend determining what data to store in the cloud. For example, when non-critical applications are run in the cloud, business continuity is less impacted than when high-critical applications are run in the cloud. Furthermore, multiple interviewees indicated that availability is a serious consideration to not use cloud environments for high-critical applications. However, this is not limited to applications, but applies to data as well, such as customer and financial data. If these are obtained by an attacker, it has severe consequences for the organisation. In addition, an organisation's internet connection could be down as well. Therefore, organisations mention having multiple physical internet connections at different providers to ensure redundancy and prevent connection problems.

Cloud backup

Another important consideration is the backup of cloud environments. While a CSP could have some downtime, it could be that a CSP drops their service, or they could have a long-term issue with their service. This is a small chance, but should still be considered. Therefore, organisations have, or are working on implementing, an on-premise backup system of the data in their cloud environment. Furthermore, a backup of data that might be important for forensic analysis should be stored on an on-premise system. One of the interviewees mentions that their backup of logs can only be stored for three months. However, not all log sources are covered, which might cause their on-premise log retention to be reduced to one month in the future. Therefore, they recommend identifying which logs are the most important.

Compliance

Compliance is a concern for many of the organisations. They have to be GDPR compliant (or "AVG" in Dutch). Before storing personal data at third parties, a processing agreement ("verwerkingsovereenkomst" in Dutch) has to be agreed upon per Dutch law [6]. It is important to note that the CSP does not become responsible for the data, as this remains the responsibility stays of the organisation. When such a processing agreement cannot be agreed upon, the service cannot be used for storing personal data. Therefore, processing agreements have to be carefully considered when storing data in the cloud, as this can leave the organisation liable. In addition, the extent of sharing data with foreign CSPs' support teams should be considered as well, as this could violate regulations.

Users

Risk can be mitigated by many technical solutions. However, the human element remains a complicating factor. Therefore, many interviewees stress the importance of users knowing how to properly use a cloud environment. For example, one company provided Office 365 training to their employees which covered how to securely share documents. Additionally, users might try to circumvent set policies. This leads to many complications in the incident handling process such as visibility issues. Therefore, it is important that users adhere to the set cloud policies and understand the dangers of not doing so.

4.4.1.2 *Cloud Management*

A lack of visibility is one of the biggest challenges in cloud incident handling. This could be solved by proper cloud management; knowing exactly who uses what cloud service for which purpose. However, this is a difficult task as cloud environments change easily. Instances can be deployed quickly, owners change, and unused instances are kept running. Therefore, many interviewees deem cloud management challenging.

Asset management

Some cloud services provide modules that assist in asset management (SaaS), but this is challenging to centralise when using multiple services. However, one interviewee highlights automation as a solution, as many APIs exist that can be used to retrieve information. Different types of asset management are deployed by interviewees. One organisation only uses asset management with IaaS infrastructure where they have deployed an on-premise management solution. Another keeps an internal wiki page up-to-date. A third organisation has developed a tool that links cloud instances to network traffic. This shows that many types of asset management can be used. However, most importantly is the information is kept up-to-date. Information that should be tracked is the owner of the instance, the organisational unit it belongs to, the department which finances it, and its purpose. This way, the owner can verify flagged behaviour or the organisational unit can be approached to claim the instance when the owner has left the organisation. Requiring users to provide this information before being allowed to use a cloud instance facilitates asset management. Additionally, taking down cloud instances when this information is not provided within a certain time frame tightens an organisation's control of cloud assets (IaaS, PaaS). The aforementioned asset management strategies that can be implemented by CSUs mitigate the visibility challenge. However, one interviewee states the solution to visibility issues should come from CSPs.

Identity and access management

Identity and access management ensures the right users have the appropriate access. Interviewees describe two different approaches for allowing access to cloud environments. The first approach uses an on-premise environment where users have to request permission and provide details to be allowed to use cloud instances. The second approach requires all new assets to be deployed in the cloud and users have to request permission and provide details to be allowed to use on-premise systems. Although their approaches differ, they both have in common that users need permission to use cloud environments. One interviewee recommends restricting permissions to start new instances to the minimum amount of people to prevent the proliferation of cloud instances. Additionally, one interviewee recommends restricting access rights of cloud instances. For example, their CSIRT is allowed to access the cloud monitoring plane but not cloud instances themselves.

4.4.1.3 *Cloud Security*

A pitfall often seen by multiple interviewees is the mentality that cloud environments are more secure than on-premise solutions. However, this is a misconception, as running a vulnerable version of an OS in the cloud instead of on-premise will not make the system more secure. While many CSPs provide security services, organisations should ensure their cloud is properly secure themselves. Therefore, while cloud security is not part of the cloud incident handling process, cloud security recommendations made by interviewees are addressed in the following paragraphs.

Design

A secure cloud environment starts with a good design. Organisations should identify all features of a service to prevent blind spots. As an example, one interviewee describes the situation where organisations might not know that SharePoint is enabled for all users within a subscription, leaving the organisation more vulnerable. They caution that cloud services are more opt-out than opt-in, enabling too many features per default which could compromise security.

Security assessment

CSPs offer many services that help protect cloud environments such as app security or threat protection services. However, besides offering paid services, CSPs often provide public best practices resources such as Microsoft 365 security for business decision-makers [11] or AWS best practices in architecting for the cloud [4]. One interviewee mentions the Microsoft Secure Score, an interface that checks if legacy protocols are still enabled, old administrator accounts still exist, and multi-factor authentication (MFA) is enabled. When improving upon security issues, the score will improve. The best practices are updated often, and organisations should evaluate them periodically to keep their cloud security up-to-date. Another way to measure the security is by assessing the security maturity level. For example, one interviewee states that their required security maturity level for their crown jewels is level four.

Cloud Service Provider

Organisations should furthermore check CSPs on security and be critical of the services they purchase. Interviewees argue that organisations should have the opportunity to verify and assess the CSP's security by for example conducting penetration tests on their services. One interviewee recommends checking if systems and firewalls are only used by your organisation, or if these are shared with other customers which might cause issues due to interfering settings of other users. Another organisation mentions requiring their CSP to implement certain security measures and protection against certain attacks. However, not all providers allow security requirements to be made or security assessments to be conducted such as penetration testing. Whether this influences an organisation's decision to use their services depends on the CSP's reputation at some organisations.

Additionally, CSPs are responsible for securing their cloud service. One interviewee mentions that CSPs can actively assist CSUs by notifying them when, for example, MFA is not enabled or when best practices are not applied. Furthermore, they suggest a security quality mark that is awarded to CSPs after being checked by independent organisations.

Policy

Policies are important as these are the organisation's guidelines to ensure the security and integrity of information. It is important that users adhere to these policies, although one interviewee notices this does not happen in practice. One policy consideration organisations often start with is determining what data is allowed to be stored in the cloud and what is not. As mentioned in Section 4.4.1.1, customer and financial data should be considered critical and their storage should be discussed.

Internet connectivity

Many security incidents occur due to instances being connected to the internet when they have no reason to be. Therefore, organisations should carefully consider which instances should be internet-facing and which not. This prevents incidents where, for example, attackers can access an instance using RDP, which might not be a necessary protocol for this instance. Interviewees recommend implementing sandbox environments which are separated from production environments when developers require access to the internet to, for example, test their applications.

Multi-factor authentication

Multiple interviewees stress the need for MFA, considering it a must-have feature. Implementing strong passwords significantly contributes to secure systems, but does not prevent them from being cracked. One interviewee mentions that cloud service usernames are often the same as the user's email address. Therefore, attackers only need to crack passwords after obtaining a list of corporate email addresses. Adding another layer of protection using MFA prevents password cracking. However, one interviewee describes the situation where the MFA is implemented incorrectly on-premise. This way, someone can connect to the cloud using other protocols and thus circumvent MFA. Therefore, organisations should ensure MFA is implemented correctly.

Updates

However, not all protocols support MFA. One interviewee cautions that older devices, such as iPhones before iOS 11, many native Android clients, and older versions of Outlook for Mac, do not support the newest exchange online protocols. These devices use legacy protocols that do not support MFA. Organisations should ensure that as many devices as possible are up-to-date, to patch older vulnerabilities, but to allow newer security techniques to be used as well.

4.4.1.4 *Service Level Agreement*

Organisations should write the right contract, not the cheapest, stresses one of the interviewees. If it is not in the contract, the CSP has no incentive to share information with the organisation and without information, there can be no cloud incident response.

Although interviewees indicate that it is difficult - albeit not impossible - for organisations to include requirements in SLAs with larger CSPs, such as Amazon and Microsoft, organisations can often influence smaller CSPs to include their desired security requirements.

Security

Multiple interviewees stress the need for verifying the security of CSPs by, for example, conducting penetration testing. Such assessments should be agreed upon in the SLA. Additionally, one interviewee mentions their organisation requiring their CSPs to implement mitigation methods to certain types of attack.

Requirements aiding the incident handling process

During incident response, time is crucial. Therefore, multiple interviewees have included certain requirements in the SLA to be able to respond as quickly as possible. Knowing whom to contact, the expected notification time depending on the incident severity, and details of the CSP's cooperation, such as what information they will provide, are considered essential parts of the SLA. One caveat mentioned by an interviewee regarding response time is that although, for example, the CSP has agreed to a response time of four hours, this can still be unsatisfactory depending on the investigation.

Legal

Response times, such as breach notifications, are important when it comes to complying with the GDPR (the Dutch AVG as mentioned in Section 4.4.1.1) as well. CSUs should ensure a processing agreement is agreed upon by the CSP when storing personal data at third parties. An interviewee describes not being able to use a certain cloud service due to not being able to agree upon a processing agreement. Additionally, an other interviewee remarks having implemented a responsible disclosure policy. This means that third parties are allowed to find and report vulnerabilities to the organisation. This organisation assures that people who report vulnerabilities will not be sued, therefore they require service providers to support that policy as well. Finally, non-disclosure agreements can be included in the SLA.

4.4.1.5 *Incident Handling Process*

In the preparation phase of the incident handling process, CSUs determine and implement the necessary policies and processes. It is crucial that CSUs determine which information is needed and available to implement the appropriate security controls. For example, CSUs should identify the logs that are available and how to incorporate them into existing security solutions.

Cloud security familiarisation

CSUs should identify which cloud services are used within their organisation and become acquainted with the security features each CSP offers. Interviewees caution that many organisation are often not aware of certain security services that are readily available or provided with their subscription. While CSUs should become acquainted with the security features offered with the used cloud services, they should also identify how to best secure the service. This can be achieved by training the organisation's security teams, preferably before moving to the cloud. The interviewed organisations either provide certification courses, such as SANS, or bring in specialised companies to provide training.

Incident response exercises

Multiple organisations have created playbooks, which contain responses to certain security incidents. These playbooks are used to perform exercises to evaluate the organisation's incident response process. During these exercises it becomes clear whether all necessary information is available, security teams know whom to contact, and what could be improved in the incident response process.

4.4.1.6 *Technical*

In addition to determining how to respond to incidents, organisations should identify which tools are necessary for the incident handling process. Many tools described by interviewees pertain to the detection and analysis phase and are therefore discussed in Section 4.4.2.2. However, which tools should be used need to be identified in the preparation phase as well.

Cloud Access Security Broker

A tool often mentioned by interviewees is a Cloud Access Security Broker (CASB). According to Gartner, CASBs "are on-premises, or cloud-based security policy enforcement points, placed between cloud service consumers and cloud service providers to combine and interject enterprise security policies as the cloud-based resources are accessed". They are used by interviewees to, for example, identify which cloud services are used within their organisation, thus improving visibility, and to enforce policies. An example of a cloud-based

CASB highlighted by interviewees is Microsoft Cloud App Security, which is described as an intelligent log viewer that alerts on events that should be examined. However, one interviewee cautions that CSUs should evaluate whether such tools are necessary or that, for example, logs can be used with existing tools such as a SIEM instead to achieve the same desired result.

Another interviewee describes creating their own tools to manage assets. Their tool associates cloud instances to network traffic to link alerts to specific instances.

Identify important log sources

Log files are important in order to obtain actionable information. Interviewees recommend assessing which log files are important and should be evaluated more in-depth. Log files can, for example, be analysed by a SIEM, which is a tool that identifies and categorises events by correlating data. Therefore, multiple organisations require logging to be connected to their SIEM, unless it is deemed to not add value. Additionally, one interviewee recommends enabling the backup of log files on-premise as long term investment. This ensures that important events can still be detected offline. Multiple interviewees would like to see log sources harmonised, as this facilitates the investigation. However, one interviewee warns that attackers can use the harmonised logs to determine how they can evade detection better.

4.4.2 *Detection and Analysis*

4.4.2.1 *Process*

Logging

Log files are necessary to be able to detect incidents. In cloud environments, multiple log sources from multiple services exist, such as Amazon API gateway, AWS Lambda, and Amazon DynamoDB. Therefore, as stated before, CSUs should know which log files are important and should be analysed. Multiple interviewees recommend to log and inspect audit logs. For example, when an attacker gains control over an administrator account or can log in to mailboxes, these actions are logged in the audit logs. Another interviewee argues that all log files are important as long as they can be used to provide context to events such as timestamps and IP addresses.

Detection

Part of the detection process is identifying anomalous behaviour. For example, interviewees mention forwarding rules in mailboxes. They argue that it is unusual for employees to set these up in their work mailbox and could thus be regarded as anomalous behaviour. Another example is analysing e-mail behaviour. For example, if an employee's usual behaviour is to send ten e-mails a day, and they start sending a hundred e-mails per hour, alerts should trigger on this anomalous behaviour. Some cloud services assist by sending these alerts automatically. However, CSUs should inspect all features of their service to identify the options and ensure these are enabled.

For example, one interviewee recommends inspecting the detected risky sign-ins in Microsoft Azure when suspecting an Office 365 account is compromised. One interviewee mentions that the alert feed their cloud service produces are sent to their organisation's alert platforms as well, such that these can be analysed in the same manner as on-premise alerts.

Another example of alerts that interviewees have implemented is regarding certain incoming e-mails. Suspicious e-mails are, for example, identified by analysing whether certain senders and subject headers occur frequently. This is used to prevent phishing attempts from succeeding. Alerts are not limited to e-mails but can apply to systems as well. An interviewee mentions Amazon S3 buckets as an example. These buckets contain data and outsiders can read this data when the buckets are, accidentally, connected to the internet. When an open data bucket is created within their company, they receive an alert.

Analysis

During the alert analysis, interviewees recommend involving the affected user as well. They could help explain suspicious behaviour. If this behaviour cannot be explained by the user, they should be involved in the further analysis as well. They could, for example, help identify the application, the risk involved, the impact when compromised.

4.4.2.2 *Technical*

Automation

When faced with a possible incident, time is crucial. Automation can ensure incidents are detected faster by, for example, conducting big data analysis. Conducting big data analysis on log files can detect incidents extremely fast, while it could take many hours to conduct analysis manually. Therefore, interviewees indicate they want to automate as much as possible.

Automation can pertain to workflow as well. Shared Indicators of Compromise (IoC) can automatically be added to the CSU's detection strategy to quickly adapt their security to detect potential new threats. Regarding the aforementioned example of forwarding rules, alerts or tickets could, for example, automatically be created when these are added to mailboxes. One interviewee mentions having created PowerShell commands that automatically detect which users have forwarding rules active.

Tools

Many tools exist that assist in automating the detection and analysis phase. One of the most widely known tools is a SIEM. This tool correlates data to identify potential incidents. For example, SIEMs that are used by interviewees are Splunk [48] and IBM QRadar [26]. Both are used in combination with JIRA [29], which is an issue tracking tool and can help track security incidents as well. Other interviewees use analysis tools provided by CSPs, such as the analysis tool in the security centre of Microsoft Azure, or third-party analysis tools that were not specified further.

One interviewee mentions that detection and monitoring depend on properly registering cloud usage. As previously mentioned, cloud users could assist in analysing potential incidents. However, this becomes nearly impossible if the user cannot be identified or is not registered. One interviewee recommends to properly register users during the onboarding process to the cloud to save considerable time and effort. Cloud management approaches are discussed in Section 4.4.1.2.

4.4.3 *Containment, Eradication, and Recovery*

4.4.3.1 *Organisational*

Organisations often make a distinction in which incidents will be handled by the CSIRT. Interviewees describe different levels of authorisations regarding CSIRTs. Multiple interviewees describe their CSIRTs have been given the mandate to perform certain actions. This allows the CSIRT to, for example, take servers offline. However, this mandate is provided more often to CSIRTs of smaller organisations. Another interviewee describes an upscaling capability, where a specially trained crisis team - which handles other types of security incidents as well - is convened to handle incidents that are beyond the CSIRTs abilities. A cloud-specific consideration mentioned by an other interviewee is deciding, based on the type of incident, whether other parties, such as the user or CSP, should be involved. Multiple interviewees mention not doing so due to the rapid nature of incident response.

Legal

Organisations consult legal teams as well in this phase, although these are not always involved. According to the interviews, most legal teams are, if deemed necessary, consulted on legal issues such as relevant regulations, data breaches, and criminal law. One interviewee mentions that their legal team is automatically involved when it concerns a major incident. When asked about jurisdiction, no organisation indicated that this is considered an issue. One interviewee indicated that jurisdictions are covered in their processing agreements, and therefore not a concern for their CSIRT. One interviewee remarks that sharing data with foreign support teams should carefully be considered by organisations before doing so.

4.4.3.2 *Technical*

Containment

According to one interviewee, organisations should not deviate too much from existing incident response processes because it is difficult to consider each exception. However, some specific containment actions are considered relevant to cloud environments. Interviewees mention containment actions such as removing forwarding rules from mailboxes, disabling specific user accounts, and removing active sessions. These actions could be automated as well. Furthermore, CSUs should carefully consider when to turn servers off. Malicious actors could be alerted due to certain containment actions which could exacerbate the situation. Instead, one interviewee mentions short term solutions, where alerts are enabled to monitor the activity of the malicious actor. An other interviewee mentions rarely turning off servers because this could eradicate important evidence such as random access memory. They might turn off servers - or disconnect the virtual network adapter - when dealing with, for example, ransomware or compromised servers which are used to pivot into the network.

Eradication and recovery

Automation can aid in the eradication and recovery process as well. One interviewee mentions that automation does not necessarily needs to be utilised to identify the worst incidents, but can recover environments that were affected by smaller incidents as well. For example, it can quarantine and eradicate malware but can perform rollbacks in virtual machines as well.

Identifying the extent of the attack

To identify the extent of the attack, CSUs should know how many systems are vulnerable. Having implemented a cloud management system (see Section 4.4.1.2) aids this process. However, attacks might not be contained to one single environment. One interviewee recommends to include on-premise environments in the investigation when faced with advanced attacks. They remark that the focus of a traditional attacker is financial gain with as little effort as possible. Therefore, advanced attacks could indicate other motives, such as espionage, which require the investigation to broaden its scope. Attacks can pertain to multiple parts of a cloud environment as well. Therefore, one interviewee recommends, for example, examining all registered Microsoft Azure apps to identify potentially suspicious activity.

CSP involvement

Additionally, CSUs should consider whether to involve the CSP or not. One interviewee mentions that CSUs should still ensure that their CSP has proper security configurations, such as a properly tailored cloud firewall. The CSPs security implementations could prevent large attacks, such as WannaCry, from spreading.

4.4.4 *Post-Incident Activity*

4.4.4.1 *Evaluation*

In the post-incident activity phase, teams evaluate the incident. Multiple interviewees mention only evaluating high or critical incidents. Smaller incidents are either evaluated based on its impact or evaluated in a smaller capacity. In the full evaluations, they evaluate what happened, how the incident was resolved, and how it could be prevented from happening in the future. Examples of considerations that have been mentioned in the interviews are patching and firewall rules. Furthermore, interviewees mention assessing the risk of recurrence, whether more security controls are needed, and which improvements can be made, such as which process could be automated. An other interviewee mentions not having a formal evaluation process but wishes to implement this in the future. Multiple interviewees mention taking these evaluations into account in future procurement projects.

Short-term matters to take into account

Interviewees mention several short-term matters CSUs should take into account after finishing the incident and the evaluation. The first is implementing short-term alerting. If the identified mitigation strategy for the incident cannot be implemented yet, short-term alerting - even using simple e-mails - can be implemented to alert on identified characteristics, such as the attacker's modus operandi. Secondly, when risks have been identified, these should be noted down. For example, when vulnerabilities have been identified in the cloud service, CSUs should try to monitor for it as well as involving the CSP. Finally, CSUs can share certain incident information with other parties, either to inquire whether those third parties have encountered the same incident or to warn them of certain IoCs. More information on incident sharing is discussed in Section 4.4.4.2.

CSP involvement

CSP involvement is often difficult to achieve according to the interviewees. Large CSPs are often more difficult to involve than smaller CSPs. One interviewee mentions sharing specific points from the evaluation with the CSP, to try to remediate the vulnerability. They surmise that CSPs probably do not want to involve the CSU in their investigation, as this might be considered too intrusive by the CSP. An other interviewee remarks that if the CSU can identify a vulnerability in the CSP's environment, either the CSP did not know about it, or they did not want to share the vulnerability with the CSU. Therefore, they argue it is important to engage the CSP.

4.4.4.2 *Report sharing*

The results from the evaluation should be recorded in a report. Interviewees mention it differs how comprehensive the report is based on the severity of the incident. Furthermore, interviewees try to share reports as much as possible. One interviewee mentions two considerations in sharing information, the CSU might be able to obtain faster results by sharing incident information and it can ensure other organisations are not affected. However, sharing depends, for example, on the trust established with other parties and the confidentiality of the data. Interviewees mention several parties that information is often shared with: the Dutch police, the NCSC, the public prosecution service, Microsoft, and trusted communities within sectors. Examples of information that can be shared are:

- Modus operandi
- Attack patterns
- IP addresses
- Threat explanations
- Threat responses
- Entire reports
- Checksums

The above can be shared with different methods and with different parties. For example, security organisations publish threat explanations and possible mitigation responses on their website to help other organisations. Two different methods of sharing are mentioned by interviewees: (1) sharing electronically, and (2) sharing in person.

Electronic sharing of Indicators of Compromise (IoC)

Electronic sharing is often done using the threat intel sharing platform MISP [38]. This platform allows organisations to add IoCs, which can be used by other organisation to better detect incidents. For example, if a malicious IP address is shared, other organisation can detect whether this IP address has connected to their environment as well to identify possible security incidents. Several communities exist that different levels of information can be shared with.

In-person sharing of information

Communities can conduct in-person meetings as well. These meetings are confidential and might require a non-disclosure agreement to be signed before being allowed to attend. During these meetings, progress on certain investigations can be discussed, threat information can be exchanged, and security policies can be addressed. Interviewees remark that meeting the other organisations helps establish trust and shortens the communication lines. One interviewee suggests that, for smaller CSUs, branch associations could organise such meetings as well or share information using newsletters.

4.5 DIGITAL FORENSICS

Digital forensics is different in cloud environments because of its dynamic nature. One interviewee considers obtaining forensic evidence from cloud environments - especially when it is not known which and how services are used - nearly impossible. They question whether the obtained evidence would be legally admissible in court, due to the environment being in constant flux and several tenants possibly being involved in the data.

Automatic evidence collection

Multiple interviewees want to automatically collect evidence. One interviewee mentions that some automation possibilities exist for cloud incidents occurring in popular services, such as Office 365. However, this is in its infancy and needs to be expanded to other services. Their observation is that when a lot of incidents occur in a service, the community tends to start developing automation possibilities for those services. Another interviewee considers automating forensic investigation to be part of a higher level of security maturity, and will, therefore, implement it in the future.

Forensic analysis

Regarding the forensic analysis of cloud environments, one interviewee warns against shutting down systems. This eradicates forensic evidence such the random access memory. They mention virtual machine disks (VMDK) can be downloaded and analysed, either by using a tool or manually. Furthermore, an other interviewee mentions that CSUs should understand that multiple log files are involved to improve the analysis method. For example, they consider the unified audit log a crucial source to analyse.

Forensic investigation tools

One interviewee describes tools used specifically for conducting forensic investigations in cloud environments. The first is a set of scripts developed by an other organisation which have been publicised. The other tool is commercially available, which connects to many cloud services and retrieves data. It can be used to analyse this data as well. However, the interviewee notes this tool often does not function as expected.

4.6 overview of recommendations

4.6.1 *Cloud Incident Handling*

Many interviewees have provided insights in practices in cloud incident handling. These are summarised in Table 4.1. As opposed to literature, most recommendations are only focused on CSUs instead of towards CSPs as well. However, two recommendations for the CSP can be distilled from the interviews: (1) CSPs should provide publicly available best practices for their service and (2) CSPs should take a pro-active role in notifying CSUs of security vulnerabilities. Furthermore, a general recommendation from one interviewee was to implement a security quality mark to assess CSPs' cloud security.

4.6.2 *Digital Forensics*

Obtaining legally admissible forensic evidence cloud environments is considered difficult. Multiple interviewees want to automatically collect evidence. However, automation possibilities are limited and considered relevant to a higher level of security maturity. CSUs should consider when to shut-down systems, as this could eradicate important evidence. Furthermore, CSUs should understand that multiple log sources are involved and should be analysed.

Table 4.1: An overview of the recommendations from practice regarding cloud incident handling

PREPARATION	
Risk management	Conduct a risk assessment and decide what risk is acceptable Evaluate available cloud service subscriptions and decide which is best suited for the organisation Determine a cloud exit strategy Consider the data that will be stored with regards to compliance. This includes processing agreements. Possible risk mitigation strategies: •Identify which log files are important •Determine the contents of (on-premise) back-ups •Consider physical internet connections at multiple providers •Provide training on safely using cloud environments to users
Cloud management	Consider implementing cloud management with regards to the following: •Utilise APIs for data retrieval •Consider the information necessary to identify instance owners •Keep the information up-to-date •Restrict who is able to create and modify cloud environments
Cloud security	Identify all characteristics of a cloud service to prevent blind spots Regularly evaluate the cloud service's security. This can include verifying the CSP's cloud security. Consider implementing cloud policies with regards to the following: •Data storage •Internet access of cloud instances •Sandbox environments for developers •Multi-factor authentication •Updates

This table continues on the next page

SLA	Pay attention to the following in the SLA with the CSP: •Security verification of the CSP (such as penetration testing) •Information that can be expected from the CSP •Whom to contact •Maximum response time •Non-disclosure agreements •Processing agreements
Incident handling process	Determine what information is necessary and available Determine how to incorporate information into existing solutions Become acquainted with the security features each CSP offers Train security teams on securing the cloud service Create playbooks and conduct exercises to assess the incident response process
Technical	Identify which tools are necessary for the incident handling process. Some cloud specific tools exist, such as a CASB. However, CSUs should assess whether logs can be used with existing tools instead Identify the log sources that need to be used
DETECTION AND ANALYSIS	
Process	Analyse log files that provide context to events such as: •Amazon API gateway •AWS lambda •Amazon DynamoDB •Unified Audit Log Utilise cloud service features that generate alerts Consider involving the affected user
Technical	Utilise automation: •Big data analysis on log files •Incorporate shared IoCs Consider using tools such as a SIEM or analysis tool provided by the CSP
CONTAINMENT, ERADICATION, AND RECOVERY	

This table continues on the next page

Organisational	Consider providing the CSIRT a mandate to perform certain actions Determine if the legal team needs to be consulted on the following: •Relevant regulations •Data breaches •Criminal law
Technical	Organisations should use existing processes. However, the following are relevant to cloud environments: •Removing forwarding rules from mailboxes •Disabling specific user accounts •Remove active sessions •Use rollbacks to recover compromised systems Consider when to shutdown systems. This could alert the attacker or eradicate important evidence such as random access memory. Consider investigating the on-premise environment when faced with an advanced attack
	POST-INCIDENT ACTIVITY
Evaluation	Evaluate the following: •Patching and firewall rules •Risk of recurrence •Security controls •Improvements Take relevant incident evaluations in account in future procurement projects Consider a cloud service vulnerability a risk and implement short-term alerting Although it might be difficult, engage the CSP
Report sharing	Share incident information (electronically or in person) as much as possible Information that could be shared are as follows: •Attack patterns •Threat explanations •Threat responses

5 THE CURRENT BEST PRACTICE IN CLOUD INCIDENT HANDLING

Chapter 3 and Chapter 4 presented best practices according to literature and the practice in the Netherlands. Combining the results from both leads to the current best practice in cloud incident handling. Table 5.1 presents the current best practice in a concise manner.

Many of the sub-topics overlapped between literature and practice, such as *evaluation* and *report sharing* in the post-incident activity phase. However, the results from practice added sub-topics, such as *risk management*, *cloud management*, and *cloud security*. The more detailed, practical recommendations from practice refined the more abstract recommendations from literature. Therefore, the overlap, addition, and refinement facilitated the integration of the recommendations from literature and practice. No conflicting recommendations could be determined.

Table 5.1 provides recommendations and remarks on a large variety of topics, ranging from technical details to contracting. These are grouped per incident handling phase and whether they apply to CSUs, CSPs, or in general. While all recommendations should be considered by CSUs, five recommendations are emphasised most by either literature or practice.

1. CSUs should prepare for cloud incidents by informing themselves of the characteristics and features of the cloud environment
2. CSUs should obtain visibility by implementing cloud management, which in addition supports in contacting users and the CSP
3. CSUs should ensure proper cloud security. Not only securing the environment at a technical level but ensuring proper security policies are in place as well, such as MFA
4. All agreements, requirements, and responsibilities must be included in the SLA
5. Incident information should be shared as this is crucial in preventing incidents and holding CSPs accountable

Table 5.1: An overview of the recommendations from literature and practice regarding cloud incident handling

	CSU	CSP	GENERAL
	PREPARATION		
Risk management *Section 4.4.1.1*	Conduct a risk assessment and decide what risk is acceptable		
	Evaluate available cloud service subscriptions and decide which is best suited for the organisation		
	Determine a cloud exit strategy		
	Consider the data that will be stored with regards to compliance. This includes processing agreements.		
	Possible risk mitigation strategies: •Identify which log files are important •Determine the contents of (on-premise) back-ups •Consider physical internet connections at multiple providers •Provide training on safely using cloud environments to users		

This table continues on the next page

	CSU	CSP	GENERAL
Cloud management *Section 4.4.1.2*	Consider implementing cloud management with regards to the following: •Utilise APIs for data retrieval •Consider the information necessary to identify instance owners •Keep the information up-to-date •Restrict who is able to create and modify cloud environments		
Cloud security *Section 4.4.1.3*	Identify all characteristics of a cloud service to prevent blind spots		
	Regularly evaluate the cloud service's security. This can include verifying the CSP's cloud security.		
	Consider implementing cloud policies with regards to the following: •Data storage •Internet access of cloud instances •Sandbox environments for developers •Multi-factor authentication •Updates		

This table continues on the next page

	CSU	CSP	GENERAL
Incident handling *Section 3.4.1.1, 4.4.1.5*	Establish an incident handling process according to a risk assessment. The following should be clear: •The relevant data sources (and how these could be incorporated in existing solutions) •The information that can be provided by the CSP (additionally the CSU should understand the content and format of the provided data) •The services and functionality CSPs and cloud platforms provide •The necessary tools •The communication channels that are to be used •The responses that correspond to incidents. Exercises should be conducted to assess the responses. Train security teams on securing the cloud service	Provide a security service Provide information about their architecture / infrastructure Establish an incident reporting service (which should allow external reports)	

This table continues on the next page

	CSU	CSP	GENERAL
SLA *Section 3.4.1.2, 4.4.1.4*	Pay attention to the following in the SLA with the CSP: •The roles and responsibilities of both CSU and CSP •Measures to be taken in case of compromised or stolen data •Permitted and prohibited use of data •The forensic investigation process •The incident reporting process •The response time that can be expected •The security verification of the CSP •Information that can be expected from the CSP •Whom to contact •Non-disclosure agreements		
Technical *Section 3.4.1.3, 4.4.1.6*	Ensure proper configuration and architecture	Develop APIs that provide event monitoring capabilities and forensic services	
	Identify which tools are necessary for the incident handling process		
	Identify the log sources that need to be used		

This table continues on the next page

	CSU	CSP	GENERAL
Reporting *Section 3.4.1.4*	Identify all necessary elements required in incident reports		Governments should establish national reporting schemes
			An independent standard communication protocol should be developed to harmonise information exchange
		DETECTION AND ANALYSIS	
Organisational Process *Section 3.4.2.1, 4.4.2.1*	Confer with the legal team to identify post-incident issues (such as evidence admission in court)		
	Analyse log files that provide context to events such as: •Amazon API gateway •AWS lambda •Amazon DynamoDB •Unified Audit Log		
	Utilise cloud service features that generate alerts		
	Consider involving the affected user		

This table continues on the next page

	CSU	CSP	GENERAL
Technical *Section 3.4.2.2, 4.4.2.2*	Consider using tools such as a SIEM or analysis tool provided by the CSP Incident response actions that utilise cloud platform capabilities are: •Analysing network flows to determine if network isolation was successful •Examining configuration data to identify similarly affected instances •Reviewing data access logs to determine whether the cloud platform itself is affected Utilise automation: •Big data analysis on log files •Incorporate shared IoCs	Provide CSUs with information necessary in the incident handling process (information about architecture, access to data sources)	Automation can be beneficial in this phase (snapshots, capturing metadata, pausing VMs) Improve techniques for analysing live compromised systems and log files
	CONTAINMENT, ERADICATION, AND RECOVERY		
Organisational *Section 4.4.3.1*	Consider providing the CSIRT the mandate to perform certain actions Determine if the legal team needs to be consulted on the following: •Relevant regulations •Data breaches •Criminal law		

This table continues on the next page

	CSU	CSP	GENERAL
Technical *Section 3.4.3.2*, *4.4.3.2*	Organisations should use existing processes. However, the following are relevant to cloud environments: •Removing forwarding rules from mailboxes •Disabling specific user accounts •Remove active sessions •Use rollbacks to recover compromised systems Consider when to shutdown systems. This could alert the attacker or eradicate important evidence such as random access memory. Consider investigating the on-premise environment when faced with an advanced attack	Provide the ability to configure networking (IaaS), access to snapshot features (IaaS), and direct read and write access to customer data (PaaS/SaaS)	Response selection techniques contribute to a rapid incident response

This table continues on the next page

	CSU	CSP	GENERAL
Cloud capabilities Section 3.4.3.1			The cloud offers more flexibility in this phase than with on-premise incident response. Technical capabilities that aid this phase are: •Infrastructure can be quickly rebuilt in a clean environment •Snapshots facilitate rollback functionality •API calls can be used for changing virtual networks of machine configurations
		POST-INCIDENT ACTIVITY	
Report sharing Section 3.4.4.2, 4.4.4.2	Share incident information (electronically or in person) as much as possible Information that could be shared are as follows: •Attack patterns •Threat explanations •Threat responses		Share incident reports (across borders) to improve security incident handling and hold CSPs accountable

This table continues on the next page

	CSU	CSP	GENERAL
Evaluation *Section 3.4.4.1, 4.4.4.1*	Work with the CSP and the incident response team to evaluate the incident handling process Try to renegotiate the SLA with the CSP when agreements have not been met. Evaluate the following: •Patching and firewall rules •Risk of recurrence •Security controls •Improvements Take relevant incident evaluations in account in future procurement projects Consider a cloud service vulnerability a risk and implement short-term alerting		Reevaluate the risk assessment whenever significant changes have been made bases on internal audits

Part IV

DISCUSSION

6 DISCUSSION

In this chapter, the results of this research presented in Chapter 3, Chapter 4, and Chapter 5 are discussed. The results from literature and practice are discussed, followed by the implications for literature and practice. This chapter is concluded by a discussion of the research's limitations.

6.1 COMPARISON BETWEEN LITERATURE AND PRACTICE

Many of the recommendations presented in Chapter 5 are not limited to incident handling in cloud environments. Several recommendations are made in existing on-premise incident handling guides as well, such as conducting risk assessments. These continue to be important in cloud environments. Moreover, some of the recommendations presented in existing incident handling guides are considered to be more important in cloud incident handling, such as SLAs.

Interviewees mainly describe using two services: Office 365 and AWS. Due to the large user base of Office 365, many of the recommendations are tailored towards SaaS applications. Moreover, interviewees indicate they mainly need guidelines on handling security incidents in SaaS environments as they consider IaaS and PaaS to be the same as on-premise environments which therefore do not require adjustments in the incident handling process. In addition, practice places more emphasis on cloud security as opposed to the literature.

Interviewees furthermore describe on-premise environments as a thing of the past, "vroeger" in Dutch. This highlights a shift towards cloud computing. However, interviewees report not having encountered many cloud incidents. One of the drawbacks of cloud services is that attackers can more easily launch large scale attacks, such as BEC, due to being able to reuse attack strategies. Therefore, the number of organisations affected by cloud incidents can suddenly change drastically.

Practice furthermore emphasises CSP accountability, while literature only mentions it briefly. Literature makes multiple recommendations for CSPs, while practice indicates organisations are mostly focused on improving their incident handling capability and barely address improvements CSPs could make. In addition, multiple interviewees indicated that it is nearly impossible to involve CSPs in the cloud incident handling process, although other interviewees succeed. However, the interviewees that indicate it is nearly impossible to involve CSPs mainly refer to global CSPs such as Microsoft and Amazon. Although CSUs could have difficulty contacting CSPs, involving them and succeeding in vulnerability mitigation on the CSP's end could help other organisation that are not able to involve the CSP. Further differences between interviewees mainly pertained to technical preferences. Some interviewees prefer access to a CSP's API, and then develop tools themselves. Other interviewees prefer CSPs to combine information and present it in a clear dashboard.

Both literature and practice emphasise the need for information sharing. While literature focuses on a more abstract level such as national reporting schemes, interviewees stress the need for building and maintaining trust relationships with other organisations. Furthermore, both literature and practice contain extensive recommendations in the preparation phase. Literature mainly restricts itself to the incident handling process, while practice considers risk management, cloud management, and cloud security to be additional important elements. Although risk management is briefly addressed in literature, practice provides more concrete considerations for CSUs. Furthermore, literature stresses the need for including roles and responsibilities in SLAs, but this was not mentioned by interviewees. This could be due to the operational background of many interviewees.

Nevertheless, there was a large overlap between literature and practice, with practice often providing a little more detail. It is clear that recommendations from practice are focused more on the needs of the organisation itself, such as organisational difficulties, while literature assumes a more abstract level.

6.2 implications of this research

This study provided an overview of differences and accompanying challenges related to cloud incident handling. The main difference between the challenges from literature and practice is that challenges from practice are practical problems, while challenges from literature are complex challenges. Therefore, this research mainly solves, or alleviates, challenges from practice and provides directions for solving the complex challenges from literature.

Regarding challenges from practice (see Section 4.2), this research provides recommendations for conducting an investigation. For example, each phase of the incident handling lifecycle is addressed, and process and tool-related recommendations are provided. Additionally, recommendations are provided to alleviate the challenge of an insufficient overview of information, such as cloud management strategies. Although not many specific recommendations could be made, considerations are provided regarding the design and road map of cloud strategies, and a lack of technology and process-related knowledge that CSUs should take into account. A lack of visibility has been addressed by, for example, discussing cloud management strategies, although the challenge of not detecting incidents remains. Finally, a dependency on the vendor is a difficult challenge, which cannot easily be solved by either this research or CSUs themselves. However, the current research provides direction by emphasising the need for CSUs to consider which data to be stored in the cloud and determine an exit strategy. Many recommendations are focused on establishing or improving the cloud incident handling strategies of CSUs themselves. These recommendations apply to a variety of organisations, regardless of familiarity with the cloud or security maturity level. However, some recommendations require certain philosophy changes from some CSUs, such as establishing trust to facilitate information sharing. Many recommendations, such as information sharing, are not limited to cloud environments, which suggest organisations could use these recommendations to improve their on-premise incident handling as well.

Regarding challenges from literature (see Section 3.2), this research addressed solutions for the challenges of collecting data, obtaining visibility into incidents, and obtaining visibility into shadow IT. It briefly discussed how to obtain data, improve the understanding of the division of incident handling responsibilities, and analyse data. However, this research provides a comprehensive overview of all topics related to cloud incident handling. In addition, it provides an overview of the challenges faced by CSUs. Although mitigation strategies for these challenges have been discussed in this research, they could be detailed further. Further directions in future research are provided in Section 7.1. These future research directions and the results of this thesis contribute to the advancement of knowledge on cloud incident handling and provide ample opportunity for further exploration.

6.3 RESEARCH LIMITATIONS

Although this research yields many interesting results, certain limitations should be considered. These do provide a direction for future research. Cloud incident handling is a very broad topic in retrospective, as it included many elements such as contracting, regulations, and technical details. In addition, research into cloud incidents at organisations yielded that not many have experienced cloud incidents yet. Therefore, it was difficult to provide specific guidelines on each associated topic in the allocated time. Additionally, the characteristics of organisations vary widely such as size, security maturity, and how likely they are to be attacked. This made it difficult to tailor recommendations for each and leads to many recommendations which should each be evaluated by organisations to determine whether these are relevant for them. Furthermore, even when specific recommendations can be made, they can be rendered obsolete when CSPs decide to alter services or cloud architecture.

The aforementioned leads to some questions on the completeness and the validity of the results. Furthermore, multiple organisations were hesitant to take part in this research due to the sensitivity and confidentiality surrounding the topic. This increases the risk that participants provide general or politically correct answers in order not to highlight drawbacks in their security implementation. Many organisations indicated they would not participate in the research if anonymity was not guaranteed. Therefore, the participants and their contribution to this research have been made as anonymous as possible. Names of respondents have not been noted on the transcripts and the results have been generalised as much as possible to ensure responses cannot be traced to participants.

Regarding the process itself, the apparent broad nature of the topic made it difficult to discuss each aspect of it in depth with participants. Furthermore, participants were often specialised in a few aspects and could sometimes only provide speculative insights into other aspects. For example, some participants use log files more intensively than others and could, therefore, discuss them more in detail. The broad nature of the topic occasionally prevented in-depth discussions as well. Due to multiple aspects needing to be discussed during the semi-structured interviews, sometimes no time was left to discuss them all. Therefore, not all aspects could be discussed with each participant as well. However, the results from this research can still be considered valid, as each topic was discussed by multiple interviewees. Additionally, it can be said that saturation was reached as no new information was discussed in the final interviews.

Expert sampling was used in this research to select participants by approaching CSIRTs of organisations. The drawback of this sampling method is that it may not be possible to generalise results to the entire incident handling community. However, identifying best practices in cloud incident handling was the goal of this research and not mapping the entire security landscape of organisations. Therefore, the chosen sampling method was best suited to achieve this goal.

7 CONCLUSION

This research aimed to identify best practices in cloud incident handling. Through analysing literature and practice, a current best practice in cloud incident handling could be identified. The results of this research can be used by organisations to establish or improve their cloud incident handling strategies. In addition, the results provide directions for future research and contribute to the advancement of knowledge on the topic. To be able to identify a current best practice, three sub-questions and two main research questions were answered in this research.

> **RQ 1.1:** What is the current incident handling landscape of Dutch organisations?

Semi-structured interviews were conducted with 14 participants from CSIRTs at different organisations to gain insight into the current incident handling landscape. Chapter 4 describes the results of these interviews. Although there is a shift towards cloud environments, organisations have not encountered many cloud incidents yet. The incidents that occur mainly consist of BEC and phishing, and occasionally organisations encounter ransomware. Although these types of attack are prevalent in on-premise environments as well, it is easier for attackers to launch large scale attacks, as attack strategies can be reused for other targets, due to cloud services often having many users.

Organisations adhere to different cloud philosophies (e.g. cloud-first, on-premise first, multi-cloud). IaaS and PaaS applications are mainly considered to be similar to on-premise systems, which therefore do not require many adaptions to the incident handling process. SaaS applications, however, are difficult as CSUs depend on CSPs to a great extent, with the majority describing Office 365. Two main CSPs are mainly discussed: Microsoft and Amazon.

Finally, six challenges in cloud incident handling could be distilled from the interviews: (1) lack of technology and process-related knowledge, (2) insufficient overview of information, (3) lack of visibility, (4) inadequate design and road map, (5) dependency on the vendor, and (6) conducting an investigation.

> **RQ 1.2:** What are the differences between cloud incident handling and on-premise incident handling according to literature and the practice in the Netherlands?

To answer this sub-question, the conducted interviews and existing literature were analysed (see Section 3.1, Section 4.1). The main difference between cloud and on-premise environments is the diminished level of control. Although IaaS and PaaS environments are considered to be similar to on-premise systems, SaaS provides (almost) no control to CSUs, which therefore leads to CSIRTs needing to adapt their incident handling process. This is complicated by the fallacy that cloud environments are inherently more secure. Cloud is less transparent than on-premise environments, which means that CSUs have to prepare well by ensuring they are aware of all the risks and possibilities of the cloud service.

Additionally, cloud incident handling requires more cooperation from third parties as opposed to on-premise environments. Many CSUs involve the CSP in their incident handling process as not doing so could leave the CSU vulnerable.

> **RQ 1.3:** What are best practices in cloud incident handling according to literature and the practice in the Netherlands?

Best practices in cloud incident handling according to literature and practice are described in Chapter 3 and Chapter 4 respectively. Cloud incident handling involves many topics such as contracting, legal, and technical details. Literature is focused towards an abstract level compared to practice, with literature providing recommendations for CSPs and governments as well. Practice focuses more on the practical issues, and their recommendations are focused towards what CSUs could do themselves regarding cloud incident handling. Additionally, practice emphasises many topics that are not necessarily part of incident handling, such as cloud security. Nevertheless, there was a large overlap between literature and practice, with practice often providing a little more detail. The recommendations made by literature and practice have been condensed and are presented in Section 3.6 and Section 4.6 respectively.

> **Main research question 1:** What is the current best practice in cloud incident handling?

Combining the results of the aforementioned sub-questions led to the current best practice in cloud incident handling. This extensive overview covers many topics involved and is presented in Chapter 5. There are five important recommendations. First, CSUs should prepare for cloud incidents by informing themselves of the characteristics and features of the cloud environment. Second, CSUs should obtain visibility by implementing cloud management, which in addition supports in contacting users and the CSP. Third, CSUs should ensure proper cloud security. Not only securing the environment at a technical level but ensuring proper security policies are in place as well such as MFA. Fourth, all agreements, requirements, and responsibilities must be included in the SLA. Fifth and final, incident information should be shared as this is crucial in preventing incidents and holding CSPs accountable.

> **Main research question 2:** To what extent is the best practice determined in RQ1 sufficient in the current cloud incident handling landscape?

This research provides solutions for many challenges presented by literature and practice as discussed in Section 6.2. Many of the presented recommendations are practical recommendations, which therefore address many of the practical problems experienced by CSUs. When organisations (plan to) use cloud environments, the presented recommendations could be used to identify which are relevant to the organisation. A checklist could be created based on this selection, to which their cloud incident handling process can be evaluated.

The difficulty in identifying a current best practice is that CSUs have different characteristics such as size, security maturity, and how likely they are to be attacked. This makes it difficult to identify a best practice that solves each problem for all CSUs. However, the recommendations presented in this research provide a direction, if not solution, for all challenges. Therefore, while it can be concluded that this identified best practice does not solve all problems, it will prevent or at least decrease the impact that malicious actors have on organisations.

7.1 FUTURE WORK

Due to the broad nature of the results, future research should focus on specific phases in the incident handling process instead to be able to provide detailed guidelines. Currently, organisations do not encounter cloud incidents often which led to more general recommendations. When organisations have encountered more cloud incidents in the future, combined with research focused on specific phases, this could yield detailed, relevant recommendations. Additionally, the findings from this research provide sufficient ground for further quantitative research which could, for example, assess whether suggested guidelines, including those proposed in this research, are effective. For example, the reinfection rates could be analysed to identify the effectiveness of CSUs' incident handling strategies.

In addition to focusing on specific phases of the incident handling process in detail, future studies could focus on best practices for CSPs as well. This could, for example, focus on assessing the implementation and effectiveness of a cloud quality mark.

Finally, this research identified the need for information exchange between CSUs and the difficulties associated with it. Therefore, future research should focus on secure information exchange between organisations. This could focus on implementing a standardised secure exchange protocol, as well as identifying how to establish trust relations between CSUs to accomplish incident information exchange.

A MATERIALS

A.1 QUESTIONNAIRE

Voorbereidende Vragenlijst Cloud Incident Response

Deze vragenlijst dient ter voorbereiding op het interview. Indien een antwoord op een vraag niet gegeven kan worden, of niet van toepassing is, dan kan deze leeg gelaten worden.

Bij de onderstaande vragen wordt gebruik gemaakt van de NIST Incident Response Lifecycle, deze is geïllustreerd in Figuur 1.

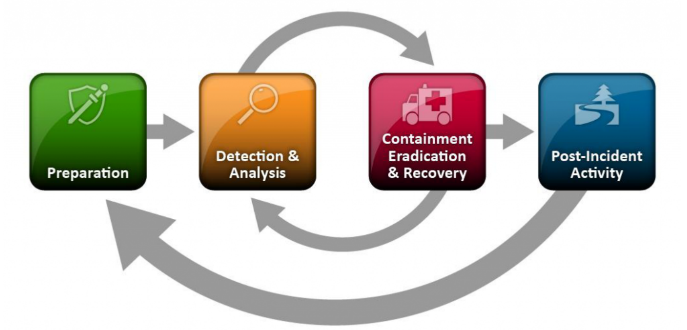

Figure 1: De fases binnen het incident afhandelingsproces zoals beschreven door NIST

Incidenten afhandelingsproces

1. Welk cijfer geeft u uw organsiatie met betrekking tot de *preparation* fase als het gaat om on-premise incidenten?

   ontevreden **tevreden**

   □ □ □ □ □ □ □ □ □ □

   1 2 3 4 5 6 7 8 9 10

   Welk cijfer geeft u uw organsiatie met betrekking tot de *preparation* fase als het gaat om cloud incidenten?

   ontevreden **tevreden**

   □ □ □ □ □ □ □ □ □ □

   1 2 3 4 5 6 7 8 9 10

   Algemeen gezien, wat zijn sterke punten?

   Algemeen gezien, waar is ruimte voor verbetering?

2. Welk cijfer geeft u uw organisatie met betrekking tot de *detection & analysis* fase als het gaat om on-premise incidenten?

ontevreden **tevreden**
□ □ □ □ □ □ □ □ □ □
1 2 3 4 5 6 7 8 9 10

Welk cijfer geeft u uw organsiatie met betrekking tot de *detection & analysis* fase als het gaat om cloud incidenten?

ontevreden **tevreden**
□ □ □ □ □ □ □ □ □ □
1 2 3 4 5 6 7 8 9 10

Algemeen gezien, wat zijn sterke punten?

Algemeen gezien, waar is ruimte voor verbetering?

3. Welk cijfer geeft u uw organisatie met betrekking tot de *containment, eradication & recovery* fase als het gaat om on-premise incidenten?

ontevreden **tevreden**
□ □ □ □ □ □ □ □ □ □
1 2 3 4 5 6 7 8 9 10

Welk cijfer geeft u uw organsiatie met betrekking tot de *containment, eradication & recovery* fase als het gaat om cloud incidenten?

ontevreden **tevreden**
□ □ □ □ □ □ □ □ □ □
1 2 3 4 5 6 7 8 9 10

Algemeen gezien, wat zijn sterke punten?

Algemeen gezien, waar is ruimte voor verbetering?

4. Welk cijfer geeft u uw organisatie met betrekking tot de *post-incident activity* fase als het gaat om on-premise incidenten?

ontevreden										**tevreden**
	□	□	□	□	□	□	□	□	□	□
	1	2	3	4	5	6	7	8	9	10

Welk cijfer geeft u uw organsiatie met betrekking tot de *post-incident activity* fase als het gaat om cloud incidenten?

ontevreden										**tevreden**
	□	□	□	□	□	□	□	□	□	□
	1	2	3	4	5	6	7	8	9	10

Algemeen gezien, wat zijn sterke punten?

____

Algemeen gezien, waar is ruimte voor verbetering?

____

Cloud incidenten binnen de organisatie

5. Wat is uw inschatting over het aandeel cloud incidenten van het totale aantal incidenten in het afgelopen jaar? (in procenten van het totaal)

□□□ %

6. Wat is uw inschatting over de hoeveelheid cloud incidenten die hebben plaatsgevonden in de afgelopen drie jaar?

Veel minder incidenten	Minder incidenten	Gelijke hoeveelheid incidenten	Meer incidenten	Veel meer incidenten	Geen idee
□	□	□	□	□	□

Cloud Service Provider

7. In hoeverre is het duidelijk wat de rolverdeling is met de CSP in het geval van een incident?

Zeer onduidelijk	Onduidelijk	Neutraal	Duidelijk	Zeer duidelijk	Geen idee
□	□	□	□	□	□

8. Welk cijfer geeft u de communicatie met de CSP?

ontevreden **tevreden**

□ □ □ □ □ □ □ □ □ □

1 2 3 4 5 6 7 8 9 10

9. Welk cijfer geeft u de informatievoorziening (bijv. over de architectuur en de mogelijkheden van het platform) vanuit de CSP?

ontevreden **tevreden**

□ □ □ □ □ □ □ □ □ □

1 2 3 4 5 6 7 8 9 10

10. Welk cijfer geeft u de reactietijd van de CSP?

ontevreden **tevreden**

□ □ □ □ □ □ □ □ □ □

1 2 3 4 5 6 7 8 9 10

Eventuele toelichting over de CSP(s):

A.2 INTERVIEW GUIDE

Interview Cloud Incident Response

Algemene informatie

1. Wat is uw naam?
2. Wat is de naam van uw organisatie?
3. Wat is uw functie bij deze organisatie?
4. In welke sector opereert de organisatie voornamelijk?
5. Wat is uw mening over de bezetting binnen de organistie voor het afhandelen van security incidenten?

Verschillen tussen de afhandeling van cloud en on-premise incidenten

6. Wat zijn volgens u de grootste verschillen tussen traditionele incident handling en cloud incident handling?
7. Wat zijn voordelen van het gebruik van cloud met betrekking tot incident handling?
8. Wat zijn nadelen van het gebruik van cloud met betrekking tot incident handling?

Cloudgebruik binnen de organisatie

9. Welke services maken gebruik van de cloud binnen de gehele organisatie?
   - 9.1. In hoeverre is er zicht op wat voor gebruik er wordt gemaakt van de cloud?
   - 9.2. In hoeverre is er zicht op welke services toegevoegd worden door gebruikers/afdelingen?
10. Mocht de cloud onbereikbaar zijn, in hoeverre wordt de organisatie beperkt in haar bedrijfscontinuïteit?

Cloud incidenten binnen de organisatie

11. Wat is uw inschatting over het aandeel cloud incidenten van het totale aantal incidenten in het afgelopen jaar? (hoeveel procent)
12. Wat is uw inschatting over de hoeveelheid cloud incidenten die hebben plaatsgevonden in de afgelopen drie jaar? (minder / meer)
13. Wat is het meest noemenswaardige / relevante incident dat heeft plaatsgevonden?

Incidenten afhandelingsproces

14. Worden incidenten door de organisatie zelf afgehandeld, of wordt dit uitbesteed?

Deze worden zelf afgehandeld . . □→ Vraag 15

Dit wordt uitbesteed . . . . □→ Hoe tevreden bent u hierover? ⟶ Vraag 34

Preparation

15. Hoe wordt ervoor gezorgd dat de organisatie optimaal is voorbereid op een incident?
16. In hoeverre verschilt de aanpak van cloud incidenten ten opzichte van traditionele incidenten?

17. Hoe vaak wordt deze aanpak geëvalueerd?
18. Wordt er gebruik gemaakt van een framework?
19. In hoeverre is er bij het in gebruik nemen van de cloud nagedacht over welke processen er in de cloud gaan draaien?
20. Welke security measures die in het eigen netwerk zijn geïmplementeerd worden ook toegepast in de cloud omgeving?
21. Welke support tools worden er gebruikt?
22. Welke support tools zou u willen zien?
23. Hoe is het afhandelingsproces buiten kantooruren ingericht?

Detection and Analysis

24. Hoe worden momenteel incidenten gedetecteerd?
25. In hoeverre zijn er aanpassingen gemaakt voor cloud incidenten?
26. Zijn er plannen voor het uitbreiden/upgraden van deze detectie? Zoja, welke?
27. In hoeverre wordt er gebruik gemaakt van automatisering?

Containment, eradication, recovery

28. Hoe worden gedetecteerde incidenten afgehandeld?
29. In hoeverre verschilt dit afhandelingsproces voor cloud incidenten ten opzichte van traditionele incidenten?
30. In hoeverre wordt een juridisch team betrokken in het incident afhandeling process? Waarom?

Post-incident activity

31. In hoeverre worden incidenten geëvalueerd?
32. Zijn er aanpassingen gemaakt naar aanleiding van eerdere cloud incidenten? Zoja, welke?
33. Wordt er informatie over incidenten met andere partijen uitgewisseld?

Cloud Service Provider

34. Hoe kunt u de CSP bereiken in het geval van een incident?
35. Hoe vaak is er al contact nodig geweest met de CSP?
36. Zijn er specifieke punten opgenomen in de SLA met de CSP?
    36.1. In hoeverre hebt u het idee dat de CSP voldoet aan deze afspraken?

Overig

37. Hebt u verder nog opmerkingen of toevoegingen die relevant zijn voor dit onderzoek?
38. Mag ik u benaderen voor verdere vragen mocht er iets niet helder zijn?

REFERENCES

[1] Nurul Hidayah Ab Rahman, Niken Dwi Wahyu Cahyani, and Kim Kwang Raymond Choo. “Cloud incident handling and forensic-by-design: cloud storage as a case study.” In: *Concurrency Computation* 29.14 (July 2017). ISSN: 15320634. DOI: 10.1002/cpe.3868.

[2] Nurul Hidayah Ab Rahman and Kim Kwang Raymond Choo. *A survey of information security incident handling in the cloud*. 2014. DOI: 10.1016/j.cose.2014.11.006.

[3] Nurul Hidayah Ab Rahman, William Bradley Glisson, Yanjiang Yang, and Kim Kwang Raymond Choo. “Forensic-by-Design Framework for Cyber-Physical Cloud Systems.” In: *IEEE Cloud Computing* 3.1 (2016), pp. 50–59. ISSN: 23256095. DOI: 10.1109/MCC.2016.5.

[4] “Architecting for the Cloud : Best Practices.” In: *Amazon Web Service* October (2018), pp. 1–45. URL: https://d1.awsstatic.com/whitepapers/AWS%7B%5C_%7DCloud%7B%5C_%7DBest%7B%5C_%7DPractices.pdf.

[5] *ATLAS.ti*. URL: https://atlasti.com (visited on February 5, 2020).

[6] Autoriteit Persoonsgegevens. *Wanneer moet ik een verwerkersovereenkomst afsluiten?* URL: https://autoriteitpersoonsgegevens.nl/nl/onderwerpen/algemene-informatie-avg/verwerkers%7B%5C#%7Dwanneer-moet-ik-een-verwerkersovereenkomst-afsluiten-7101 (visited on January 6, 2020).

[7] Hillary Baron, Sean Heide, Shamun Mahmud, and John Yeoh. *Cloud Security Complexity: Challenges in Managing Security in Hybrid and Multi-Cloud Environments*. Tech. rep. Cloud Security Alliance, 2019. URL: https://cloudsecurityalliance.org/artifacts/cloud-security-complexity/.

[8] Vanansius Baryamureeba and Tushabem Florence. “The Enhanced Digital Investigation Process Model.” In: *Proceedings of the Fourth Digital Forensic Research Workshop* (2004).

[9] Alan Bryman. “Encyclopedia of Social Science Research Methods "Triangulation".” In: *Encyclopedia of Social Science Research Methods* (2004), pp. 1143–1144. DOI: http://dx.doi.org/10.4135/9781412950589.n1031.

[10] Brian Carrier and Eugene Spafford. "An event-based digital forensic investigation framework." In: *Digital forensic research workshop* (2004), pp. 1–12. DOI: 10.1145/1667053.1667059. URL: http://www.digital-evidence.org/papers/dfrws%7B%5C_%7Devent.pdf.

[11] Brenda Carter and Denise Vangel. *Microsoft 365 Security*. 2019. URL: https://docs.microsoft.com/en-us/microsoft-365/security/microsoft-365-security-for-bdm (visited on January 9, 2020).

[12] Hyunji Chung, Jungheum Park, Sangjin Lee, and Cheulhoon Kang. "Digital forensic investigation of cloud storage services." In: *Digital Investigation* 9.2 (2012), pp. 81–95. ISSN: 1742-2876. DOI: 10.1016/j.diin.2012.05.015. URL: http://dx.doi.org/10.1016/j.diin.2012.05.015.

[13] Paul Cichonski, Tom Millar, Tim Grance, and Karen Scarfone. "Computer Security Incident Handling Guide : Recommendations of the National Institute of Standards and Technology." In: (2012). DOI: 10.6028/NIST.SP.800-61r2. URL: https://nvlpubs.nist.gov/nistpubs/SpecialPublications/NIST.SP.800-61r2.pdf.

[14] CMMI Product Team. *CMMI for Services, version 1.3 (CMU/SEI-2010-TR-034)*. Tech. rep. November. CMMI Institute, 2010. URL: http://cmmi.center/library/SDocs/CMMI%7B%5C_%7DSVC%7B%5C_%7D1%7B%5C_%7D3.pdf.

[15] *Computer emergency response team*. URL: https://nl.wikipedia.org/wiki/Computer%7B%5C_%7Demergency%7B%5C_%7Dresponse%7B%5C_%7Dteam (visited on November 26, 2019).

[16] Juliet M. Corbin and Anselm Strauss. "Grounded theory research: Procedures, canons, and evaluative criteria." In: *Qualitative Sociology* 13.1 (1990), pp. 3–21. ISSN: 0162-0436. DOI: 10.1007/BF00988593. URL: http://link.springer.com/10.1007/BF00988593.

[17] Katie Costello. *Gartner Forecasts Worldwide Public Cloud Revenue to Grow 17.5 Percent in 2019*. April 2019. URL: https://www.gartner.com/en/newsroom/press-releases/2019-04-02-gartner-forecasts-worldwide-public-cloud-revenue-to-g.

[18] *CSIRTs by Country - Interactive Map*. URL: https://www.enisa.europa.eu/topics/csirts-in-europe/csirt-inventory/certs-by-country-interactive-map%7B%5C#%7Dcountry=Netherlands%20(The) (visited on November 26, 2019).

[19] Marnix Dekker, Dimitra Liveri, and Matina Lakka. *Cloud Security Incident Reporting - Framework for reporting about major cloud security incidents*. December. 2013, p. 38. ISBN: 9789279000775. DOI: 10.2788/14231.

[20] Yucong Duan, Guohua Fu, Nianjun Zhou, Xiaobing Sun, Nanjangud C. Narendra, and Bo Hu. "Everything as a Service (XaaS) on the Cloud: Origins, Current and Future Trends." In: *Proceedings - 2015 IEEE 8th International Conference on Cloud Computing, CLOUD 2015* June (2015), pp. 621–628. DOI: 10.1109/CLOUD.2015.88.

[21] Josiah Dykstra and Alan T Sherman. "Acquiring forensic evidence from infrastructure-as-a-service cloud computing : Exploring and evaluating tools , trust , and techniques." In: *Digital Investigation* 9 (2012), S90–S98. ISSN: 1742-2876. DOI: 10.1016/j.diin.2012.05.001. URL: http://dx.doi.org/10.1016/j.diin.2012.05.001.

[22] *FIRST Teams*. URL: https://www.first.org/members/teams/?%7B%5C#%7Dnetherlands (visited on November 26, 2019).

[23] Christian Frøystad, Erlend Andreas Gjære, Inger Anne Tøndel, and Martin Gilje Jaatun. "Security Incident Information Exchange for Cloud Services." In: Scitepress, May 2016, pp. 391–398. DOI: 10.5220/0005953803910398.

[24] Bernd Grobauer and Thomas Schreck. "Towards incident handling in the cloud." In: Association for Computing Machinery (ACM), October 2010, p. 77. DOI: 10.1145/1866835.1866850.

[25] Paul Henry, Jacob Williams, and Benjamin Wright. *The SANS Survey of Digital Forensics and Incident Response*. Tech. rep. SANS Institute, 2013. URL: https://blogs.sans.org/computer-forensics/files/2013/07/sans%7B%5C_%7Ddfir%7B%5C_%7Dsurvey%7B%5C_%7D2013.pdf.

[26] *IBM QRadar*. URL: https://www.ibm.com/security/security-intelligence/qradar (visited on February 5, 2020).

[27] Cisco Global Cloud Index and CV Cisco Visual Networking Index. *Forecast and Methodology, 2016-2021; White Paper; Cisco Systems*. Tech. rep. 2017.

[28] Martin Gilje Jaatun and Inger Anne Tøndel. "How much cloud can you handle?" In: *Proceedings - 10th International Conference on Availability, Reliability and Security, ARES 2015*. Institute of Electrical and Electronics Engineers Inc., October 2015, pp. 467–473. ISBN: 9781467365901. DOI: 10.1109/ARES.2015.38.

[29] *Jira | Issue & Project Tracking Software | Atlassian*. URL: https://www.atlassian.com/software/jira (visited on February 5, 2020).

[30] Steffen Kächele, Christian Spann, Franz J. Hauck, and Jörg Domaschka. "Beyond IaaS and PaaS: An extended cloud taxonomy for computation, storage and networking." In: *Proceedings - 2013 IEEE/ACM 6th International Conference on Utility and Cloud*

Computing, UCC 2013 (2013), pp. 75–82. DOI: 10.1109/UCC.2013.28.

[31] K Kent, S Chevalier, T Grance, and H Dang. *Guide to integrating forensic techniques into incident response*. Tech. rep. 2006. DOI: 10.6028/NIST.SP.800-86. URL: https://nvlpubs.nist.gov/nistpubs/Legacy/SP/nistspecialpublication800-86.pdf.

[32] Udo Kuckartz. "Qualitative Text Analysis: A Systematic Approach." In: *Compendium for Early Career Researchers in Mathematics Education*. Ed. by Gabriele Kaiser and Norma Presmeg. Cham: Springer International Publishing, 2019, pp. 181–197. ISBN: 978-3-030-15636-7. DOI: 10.1007/978-3-030-15636-7_8. URL: https://doi.org/10.1007/978-3-030-15636-7%7B%5C_%7D8.

[33] Tom Laszewski and Prakash Nauduri. "Migrating to the Cloud." In: *Migrating to the Cloud* (2012), pp. 1–19. DOI: 10.1016/b978-1-59749-647-6.00001-6.

[34] LogicMonitor. *Cloud Vision 2020 : The Future of the Cloud*. URL: https://www.logicmonitor.com/resource/the-future-of-the-cloud-a-cloud-influencers-survey/.

[35] Ben Martini and Kim Kwang Raymond Choo. "An integrated conceptual digital forensic framework for cloud computing." In: *Digital Investigation* 9.2 (2012), pp. 71–80. ISSN: 17422876. DOI: 10.1016/j.diin.2012.07.001. URL: http://dx.doi.org/10.1016/j.diin.2012.07.001.

[36] Rodney Mckemmish. *What is Forensic Computing?* Australian Institute of Criminology Canberra, 1999.

[37] Peter Mell and Tim Grance. "The NIST definition of cloud computing." In: *Cloud Computing and Government: Background, Benefits, Risks*. Elsevier, 2011, pp. 171–173. ISBN: 9781617617843. DOI: 10.1016/b978-0-12-804018-8.15003-x. URL: https://linkinghub.elsevier.com/retrieve/pii/B978012804018815003X.

[38] *MISP - Open Source Threat Intelligence Platform & Open Standards For Threat Information Sharing*. URL: misp-project.org (visited on February 5, 2020).

[39] Rich Mogull, James Arlen, Adrian Lane, Gunnar Peterson, Mike Rothman, and David Mortman. *Security Guidance for Critical Areas of Focus in Cloud Computing v4.0*. Tech. rep. Cloud Security Alliance, 2017. URL: https://cloudsecurityalliance.org/download/security-guidance-v4/.

[40] Aryan Taheri Monfared and Martin Gilje Jaatun. "Monitoring intrusions and security breaches in highly distributed cloud environments." In: *Proceedings - 2011 3rd IEEE International Conference on Cloud Computing Technology and Science, CloudCom 2011* (2011), pp. 772–777. DOI: 10.1109/CloudCom.2011.119.

[41] National Cyber Security Center. *Cloudcomputing & security*. Tech. rep. 2012, p. 78. URL: https://www.ncsc.nl/binaries/nl/...advies/.../whitepapers/whitepaper.../....

[42] ORACLE and KPMG. *Oracle and Kpmg Cloud Threat Report, 2018*. Tech. rep. 2018, p. 41. URL: http://www.oracle.com/us/dm/oraclekpmgcloudthreatreport2018-4437566.pdf.

[43] *Part 1: Principles of incident management ISO/IEC 27035-1*. Tech. rep. ISO, 2016. URL: https://www.iso27001security.com/html/27035.html.

[44] *Part 2: Guidelines to plan and prepare for incident response ISO/IEC 27035-2*. Tech. rep. ISO, 2016. URL: https://www.iso27001security.com/html/27035.html.

[45] Mark Reith, Clint Carr, and Gregg Gunsch. "An examination of digital forensic models." In: *International Journal of Digital Evidence*. Vol. 1. 3. May 2002, pp. 1–12. DOI: 10.1.1.13.9683.

[46] Salomon Rico, Francis Kaitano, Munyaradzi Mufambisi, and Miguel Villegas. *Incident Management and Response*. Tech. rep. ISACA, March 2012. URL: http://www.isaca.org/Knowledge-Center/Research/ResearchDeliverables/Pages/Incident-Management-and-Response.aspx.

[47] Keyun Ruan, Joe Carthy, Tahar Kechadi, and Mark Crosbie. "IFIP AICT 361 - Cloud Forensics." In: *International Federation of Information Processing* 361 (2011), pp. 35–46. URL: https://link.springer.com/content/pdf/10.1007%7B%5C%%7D2F978-3-642-24212-0%7B%5C_%7D3.pdf.

[48] *SIEM, AIOps, Application Management, Log Management, Machine Learning, and Compliance | Splunk*. URL: https://www.splunk.com/ (visited on February 5, 2020).

[49] Symantec. "Adapting to the New Reality of Evolving Cloud Threats." In: *Cloud Security Threat Report* 1.June (2019).

[50] Aryan TaheriMonfared and Martin Gilje Jaatun. "Handling compromised components in an IaaS cloud installation." In: *Journal of Cloud Computing: Advances, Systems and Applications* 1.1 (2012), p. 16. ISSN: 2192-113X. DOI: 10.1186/2192-113x-1-16.

[51] Xuan Zhang, Nattapong Wuwong, Hao Li, and Xuejie Zhang. "Information security risk management framework for the cloud computing environments." In: *Proceedings - 10th IEEE International Conference on Computer and Information Technology, CIT-2010, 7th IEEE International Conference on Embedded Software and Systems, ICESS-2010, ScalCom-2010*. 2010, pp. 1328–1334. ISBN: 9780769541082. DOI: 10.1109/CIT.2010.501.

[52] Yun Zhang, Farhan Patwa, and Ravi Sandhu. “Community-Based Secure Information and Resource Sharing in AWS Public Cloud.” In: *Proceedings - 2015 IEEE Conference on Collaboration and Internet Computing, CIC 2015*. Institute of Electrical and Electronics Engineers Inc., March 2016, pp. 46–53. ISBN: 9781509000890. DOI: 10.1109/CIC.2015.42.