TRAFICOM
Finnish Transport and Communications Agency
National Cyber Security Centre

Instructions – Incident response for cloud environments

Traficom Publications
18/2023

Contents

Management summary

In many organisations, the use of cloud services is a part of normal activities. Cloud services are used for processes that are important to the business or even business critical. In many organisations, daily business and its smooth continuation requires that these cloud services operate reliably and with good information security.

However, maintaining information security requires constant work and development of the operations. These instructions describe how organisations can prepare for information security incidents involving cloud services and how they should act in situations in which an information security incident involving a cloud environment is suspected.

Key issues include especially the following:

1. Information security in the procurement of services

If you are in the process of purchasing new cloud services, make sure that enough attention is paid to the information security of the service already during the procurement phase. It is easier and often cheaper to implement solutions with a good level of information security when using cloud services than at a data centre. This requires sufficient expertise from the customer.

2. The shared responsibilities model in different types of services

In IaaS and PaaS cloud services, a large share of the information security responsibilities for the service falls on the organisation that uses them. For the maintenance of these services, it is very important to ensure that the organisation has enough maintenance resources available, and that their information security competence is at a sufficiently high level. The information security of cloud services often requires special expertise concerning these specific services. Many information security incidents originate from incorrectly configured services or a change that leads to the configuration becoming less secure. The potential for errors is reduced by ensuring that the people responsible for configurations have the competence necessary for the task.

3. Multi-factor authentication

It is worth investing into the protection of cloud services. One fundamental issue is ensuring that it is not possible to log in to the services with a simple combination of a user account and password. Make sure that multi-factor authentication is used in all of your cloud services for all user groups.

4. Training for incidents

You should prepare for incidents. Even if significant investments in information security are made, an incident can always occur. Drawing up plans on incidents and training for incidents provide security if a real situation should occur at some point.

5. Preparation measures

As a rule, incidents follow a certain life cycle that has different phases. The phases and the key tasks related to them are described in these instructions. In your organisation, ensure that you have sufficient capabilities for carrying out incident response measures. Often these measures require making preparations in advance, such as implementing sufficient logging and backup processes that have been tested to ensure that they function correctly.

6. Obtaining expert help for incidents

If an information security incident takes place in an organisation, expert helpinstru may be needed in order to resolve it. The agreements required for such expert help should be made already in advance, just in case.

1 Introduction

1.1 Purpose of the instructions

The goal of these instructions drawn up by the National Cyber Security Centre Finland (NCSC-FI) of the Finnish Transport and Communications Agency Traficom is to advise organisations on how to prepare for information security incidents involving cloud services and how to act in situations in which an information security incident involving a cloud environment is suspected. The instructions are intended for all organisations regardless of their size or sector.

These instructions cover the following phases of the life cycle of an information security incident:

- Detecting an information security incident – the instructions describe how information security incidents in cloud environments can be detected.
- How to act in case of an incident – the instructions cover the key measures an organisation should take in an incident.
- Recovery from an incident – the instructions describe how to recover from a cloud information security incident and return back to business.

The instructions are a part of the series of instructions by the National Cyber Security Centre Finland that describes recovery from information security incidents of different types. Recovery from incidents related to things other than cloud environments is described in the other parts of the series of instructions, which are found on the website of the NCSC-FI.

General information on the safety of cloud services can be found in the publication 'Ohjeita pilvipalvelujen turvallisuudesta yksityishenkilöille, pienyhteisöille ja -yrityksille'[1] ('Instructions on the security of cloud services for private individuals, small communities and businesses,' in Finnish). The National Emergency Supply Agency has published a guide for corporate decision-makers that are considering the adoption of cloud services or expanding their use to critical functions[2].

[1] https://www.kyberturvallisuuskeskus.fi/fi/julkaisut/ohjeita-pilvipalvelujen-turvallisuudesta-yksityishenkiloille-pienyhteisoille-ja
[2] https://www.huoltovarmuuskeskus.fi/files/dfd001d3135a6a37876a5afe88ba2a816156e8ae/huoltovarmuutta-pilvipalveluilla-230306.pdf

1.2 What does a cloud environment information security incident mean?

Information security incidents are unexpected or unwanted events or series of events, in which the security of protected data or services is endangered. In these instructions, we define a cloud environment information security incident as an information security incident in which the endangered data are stored in a cloud service or the service targeted in the security incident is a cloud service.

Cloud services refer to information technology applications or platform services used by the customer that are produced in the service provider's device environment and used over a telecommunications network, and whose technical details are the service provider's responsibility and hidden from the customer to some degree. From the customer's perspective, cloud resources are always available, they can be easily scaled to suit the need at a given moment, and their pricing is usually based on usage. Cloud services do not require the customer to invest in hardware or commit to long-term contracts.

Different kinds of information security incidents occur in cloud environments; their typical root causes include:

- Weak or leaked user account and password information or other secrets that an unauthorised party can use to access the data and services in the cloud. This risk is present especially if multi-factor authentication (MFA) is not used, or when secrets like access keys have been handled carelessly, such as a part of the source code of software.
- Use of the cloud environment's default configurations, which leads to the data, functions or applications being more widely visible in the public network than intended.
- Making previously used good information security settings less strict, which creates an unexpected opportunity for a cyber criminal into the services.
- A vulnerability in the customer's workload, such as a container or virtual machine.

The situations described above can be avoided by good expertise in cloud usage, and well-working configuration and change management, as well as monitoring information security vulnerabilities actively.

The European Union Agency for Cybersecurity ENISA[3] has estimated that attackers use the following methods, among others, to attack cloud environments:

- exploiting the vulnerabilities of a cloud environment
- using social engineering attacks to harvest credentials for cloud environments
- exploiting misconfigured containers
- targeting cloud infrastructure, interfaces or cloud-hosted backups to gain a foothold.

[3] ENISA's report: https://www.enisa.europa.eu/publications/enisa-threat-landscape-2022

1.3 Special characteristics of cloud environments for incident response

Cloud environment information security incidents have certain special characteristics. It is good for the customers of cloud services to take these special characteristics into account so that they can act as effectively as possible in case of an incident.

1.3.1 Visibility into the service provider's resources

Using a cloud service means outsourcing, where the customer assigns the development and maintenance of the purchased service to the service provider either completely or partially. It is good to be aware that the decision to implement a cloud service is also a decision to trust the cloud service provider and the level of information security of the services it provides.

There are terms of agreement concerning the division of responsibilities related to cloud services that set limits on the visibility into the service provider's processes or opportunities for active cooperation. From an information security perspective, this means that when a customer organisation is choosing a cloud service provider, the organisation should, before purchasing the service, carefully evaluate the service provider's documentation that describes the information security processes and management methods used, including operating models related to information security incidents and the level of support available in case of an incident.

It should be noted that some information security features of cloud services may be products subject to an extra charge, and their deployment requires both investments and special expertise.

1.3.2 Shared responsibility model

The cloud service provider and the cloud service customer always share the responsibility for the information security of the service. The type of cloud service (SaaS, PaaS, IaaS) and the agreement made with the service provider as well as the services used determine how the responsibilities are divided between the parties. The responsibilities are also reflected on incidents:

Infrastructure-as-a-Service (IaaS): In the IaaS model, the cloud service provider offers a scalable IT infrastructure where the customer deploys the desired services, such as virtual machines, network connections and storage capacity. In practice, the services are physically located in the service provider's data centres, and the service provider is responsible for the physical security of the hardware. The customer is responsible for ensuring the security of the services it uses and the data it stores, including e.g. specifying the users and access rights, configuring the information security of cloud resources, managing the vulnerabilities of containers and operating systems, detecting and combating attacks, duplicating services as well as sufficient backups. In case of an information security incident, the customer has a very extensive responsibility for resolving the incident and corrective measures, if the incident does not apply to the ordering and maintenance systems or production infrastructure that are under the service provider's responsibility.

Platform-as-a-Service (PaaS): In the PaaS model, the cloud service provider provides the platform intended for developing and releasing applications. The customer deploys the services it wants from the platform; they usually include storage and processing services for files and data in addition to development tools and the running platform for code. In the PaaS model, the customer is responsible for ensuring the security of its development environment, applications and the data it has stored in addition to the access management of the services. In case of an information security incident, the customer has the main responsibility for resolving the incident, but the cloud service provider may act as a partner in co-operation.

Software-as-a-Service (SaaS): In the SaaS model, the cloud service provider offers its customers a ready-made application that is usually used over the public network, i.e. with an internet connection. In the SaaS model, the service provider is responsible for the security of the infrastructure related to the application as well as the application itself. However, the customer is responsible to some degree for the security of the application during use, such as the management of users, access rights and identification means. The resolution of information security incidents is carried out in cooperation with the cloud service provider, because it is likely that the customer does not have access to the application's log or other data needed to investigate the incident. SaaS services may include additional information security features subject to a charge, related to e.g. identification means, the encryption of data and the availability of access logs.

A cloud service customer should always study the model of division of responsibilities of the service to be purchased in order to avoid misunderstandings and unpleasant surprises related to security and incident response.

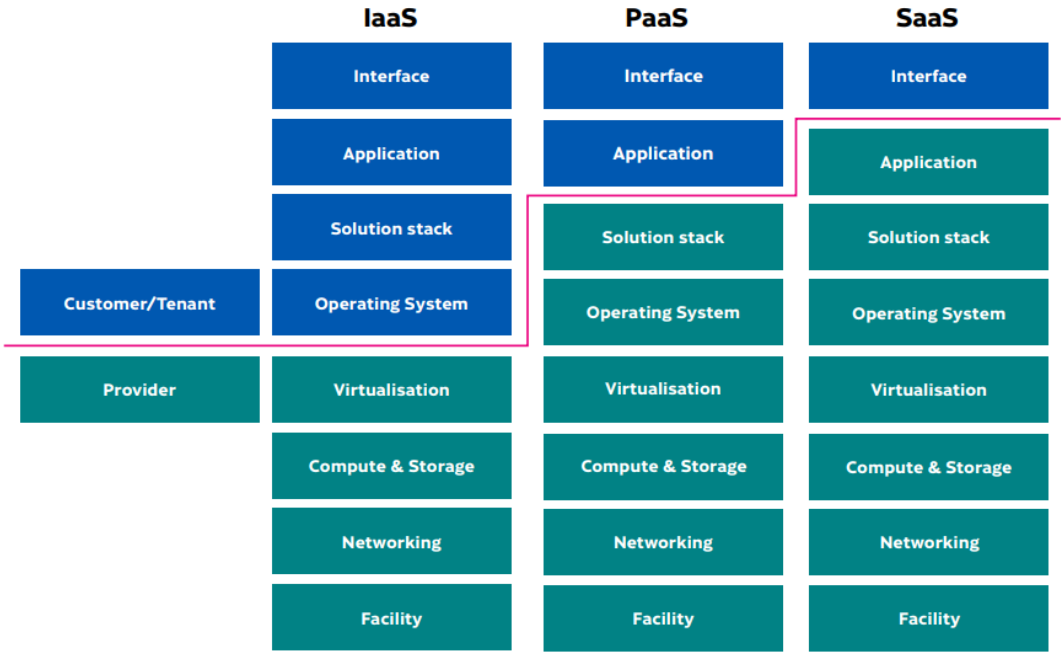

Image 1 A typical model for the division of responsibilities. Source: Criteria for Assessing the Information Security of Cloud Services (PiTuKri)

1.3.3 Special expertise needed for the resolution of incidents

Cloud environments can be complicated and contain a lot of data. In addition, they may be undergoing constant change, which makes incident resolution a task that requires special expertise.

Cloud service providers may offer service agreements of different levels that affect the kind of support they provide in case of an incident. In addition, there are specialised suppliers on the expert service market that can help with investigating incidents and recovering from them.

In many cases, it is useful for an organisation to draw up agreements in advance with service providers that can help in case of an incident. Drawing up agreements in advance is worth it to ensure that if an incident is in progress, there is no need to spend time on agreement negotiations.

The organisation must also ensure that the skills and knowledge of its own IT personnel are sufficient concerning cloud services and their safe use. Cloud expertise can be developed with training and by obtaining professional certifications related to cloud services.

Expertise related to incidents can be developed through training, which is described in more detail in Chapter 2.3 Training on information security incidents related to the cloud environment.

1.3.4 The information security features of a service may depend on the agreement or licence

Some cloud services may have additional features related to information security subject to a charge related to identification means, the encryption of data and the availability of access logs, among other things. In contrast, the information security features available to cloud services may also be linked to the type of licence that has been purchased for the service. In the IaaS and PaaS models, many information security technologies that offer more power for monitoring the environment and information security events, for example, are often subject to an extra charge.

An organisation purchasing or using cloud services should study the features, licences and service models of the service in order to ensure that the package as a whole meets the organisation’s information security requirements.

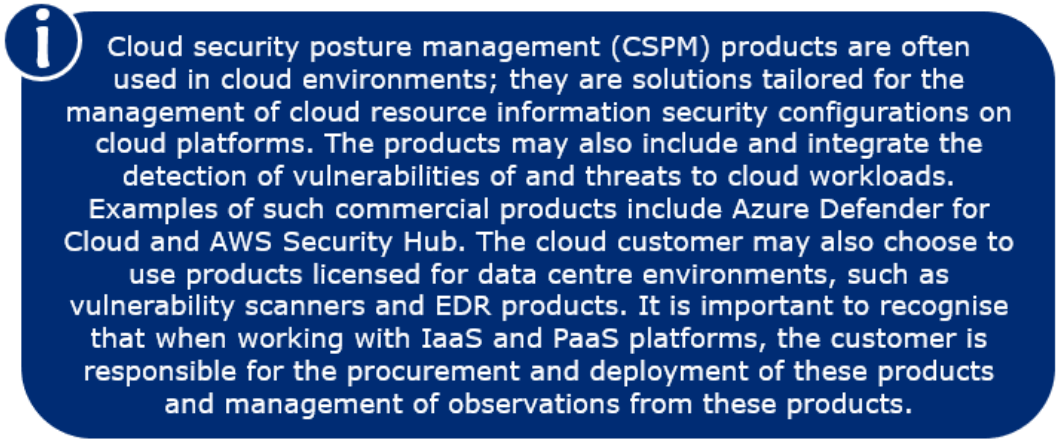

1.3.5 Gathering information during security incidents

The resolution of an incident in a cloud environment largely depends on the available log data. Whether a service generates and collects log data depends on the service and its configuration. If there are no logs collected by the customer, it may be difficult or impossible to prove that an information security incident has occurred.

Many cloud service platforms collect log data automatically from IaaS and PaaS services on the **control plane** (sometimes also known as the management plane). Control plane refers to the user interfaces, tools and interfaces that the customer's maintenance personnel use to order and configure cloud services. For example, the deployment of new services and authorisation of new maintenance personnel on the platform is typically recorded in a control plane log without the customer needing to switch on such an audit log.

The recording periods of logs vary, and it must be checked whether the period on the platform service is sufficient. For example, the storage period may be a few months, in which case investigating suspected misuse that has continued for a long time may be difficult. Typically, logs can be stored longer than the default for an additional charge.

Cloud services are used on the **data or application plane**. As for this level, access events include e.g. logging in to a virtual machine, downloading a file from cloud storage, and operations related to the data in a database service. The availability of access logs on this level often depends on the procedures adopted by the customer that the service platforms offer for an additional charge. The cloud service may not automatically generate any logs on such events at all; instead, the customer is responsible for their deployment and the log management life cycle.

In SaaS services, the service provider collects the log data. When deploying SaaS services, it must be ensured that the service provider's logging practices correspond to the requirements on logging set by the organisation. For example, the investigation of access events and suspected misuse should be possible. In some services, such logs may be available to the customer, while in others, they require search measures by the service provider, and in yet others, no logs are generated at all.

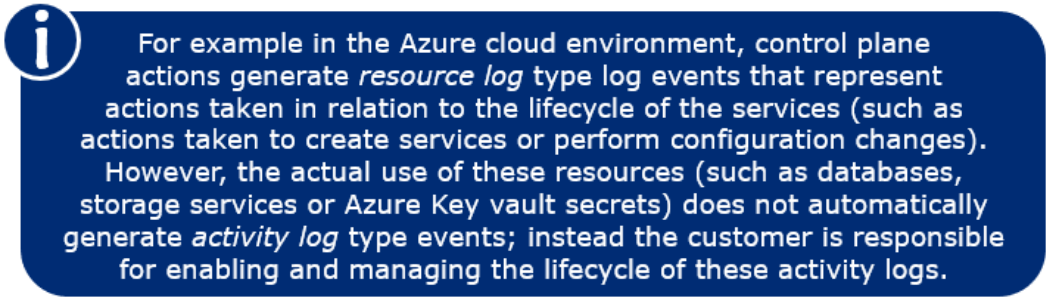

In addition to logs, cloud services may enable the monitoring of network traffic (network flow logs), which makes it possible to analyse the cloud workloads such as the telecommunications of virtual machines, containers and applications.

1.3.6 DevOps practices

Modern application development is based on DevOps practices, where the code, automation and repeatability are at centre stage. Such automation is especially common in cloud environments.

In practice, this means that the customer's share of the deployment, configuration, deployment in production and monitoring of the production of cloud services can be automated and created as a repeatable model. This makes recovery from incidents more flexible. The customer may have an option of restoring an application and its support services from source codes and data backups either in the same location, in another location by the same cloud service provider, or sometimes in the environment of another service provider with minor changes. According to good practices, the source code does not contain any authorisation information (such as accounts, keys or certificates); instead, the information is managed separately.

Important things to take into account concerning DevOps practices with regard to incidents include ensuring the availability of the source code and data backups. The best way to implement this is by saving backups in different services (e.g. database vs. object storage), protecting them with access rights of different levels and, if possible, also distributing them into their own data centre, a geographically separate storage location by the same cloud service, or in a completely different cloud service. More information on backups is available in Chapter 2.2.7 Ensure sufficient backup and recovery procedures.

2 Preparation

Preparing for security incidents is a key method of reducing their severity and making it possible to recover quickly and continue the business. Organisations can assess their own readiness by using the Kybermittari (Cybermeter) cyber security evaluation tool of the National Cyber Security Centre Finland, for instance[4]. An incident response plan that has been drawn up in advance is a good starting point for what to do in case of a security incident. The organisation must also ensure that measures such as locking user accounts, isolating servers and terminal devices from the network and restricting network traffic to harmful IP addresses or domain names are technically possible and that the personnel have the expertise and operating instructions required to carry them out.

Gathering, compiling and monitoring log data is important in order to detect incidents in time. Log data also make it possible to investigate incidents thoroughly, which speeds up the cleaning and restoration of the environment, if necessary. The National Cyber Security Centre Finland has drawn up a guide on how to collect and use log data[5]. Depending on the systems used by the organisation, comprehensive monitoring typically also requires network- and system-level solutions in addition to this.

2.1 Administrative measures

Incidents are nearly always hectic situations that are hard on the organisation. Measures can be taken in advance to make the situation easier by drawing up plans and practicing the operating models used during an incident. These measures ensure that when a situation is ongoing, the incident response team can focus on resolving the incident, because they are already familiar with the key operating methods, responsibilities, roles and procedures needed in the situation.

2.1.1 Know your cloud service

The first step of preparation is identifying which cloud services the organisation uses: a typical modern IT environment consists of e.g. several SaaS and one or more cloud platforms (such as Azure, AWS or Google Cloud). The organisation may also have implemented a multi-cloud strategy that takes advantage of the services of several cloud service platforms.

It is good for the organisation to build a picture of its own set of cloud services and their security by finding out, for example:

- Which cloud services and cloud service providers it uses.
- What the agreements and terms of service say on the division of security responsibilities between the organisation using the service and the cloud service provider (concerning the shared responsibility model, see Chapter 1.3.2).
- What information and types of information (e.g. customer information, confidential information) is processed in each service. It is especially important to know in which cloud services personal data within the scope of the General Data Protection Regulation (GDPR) are processed.

[4] https://www.kyberturvallisuuskeskus.fi/en/our-services/situation-awareness-and-network-management/kybermittari-cybermeter
[5] https://www.kyberturvallisuuskeskus.fi/en/ncsc-news/instructions-and-guides/collecting-and-using-log-data

- How have the access rights of cloud services been specified, and how compliance with the principle of least privilege is ensured in the access.
- Whether the environments of different types (e.g. testing and production) have been separated.
- How the cloud service is segmented on the network level.
- Whether any connections have been opened from the cloud to the data centre.
- Which identities the applications and automations use, and what kind of rights they have.
- Which different kinds of cloud services and resources are available, and what features they have from the perspective of information security.
- Where the services and data are located (geolocation). Have the services been duplicated at several locations in order to reduce the risk based on one single geographical location?
- How the services and data are being protected from cyber threats at the moment.
- What kind of logs the services generate, where they are saved, and how long they are available.
- How the incidents and alarms from log data and security products are monitored.
- What kind of backup and recovery procedures have been implemented in each service.

It is good to identify the dependencies and impact that cloud services have had on other services and the business of the organisation, such as:

- How would an information security incident focusing on a specific cloud service affect the organisation's other IT systems?
- How would an information security incident targeting a specific cloud service affect the business processes of the organisation?

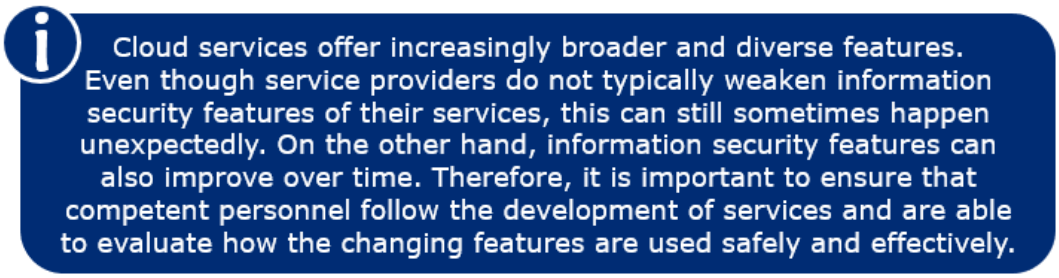

2.1.2 Plan and document incident response measures

Incident response plan: Drawing up an incident response plan is a good exercise in which the organisation can assess things such as the roles, communication channels and external operators (e.g. legal services, crisis communication services, determining the scope and severity of the incident) needed in an incident and specify what kind of an event the organisation defines as an incident that starts the incident response process. The incident response plan also describes the decision-making chain, according to which actions are taken in an incident. The finished incident response plan acts as an incident response support framework, and following it can be practiced in advance. The incident response plan makes it also possible to open up the operating model on how to act during an incident and provide training on it within the organisation.

In addition to the incident response plan, the following documentation can be useful:

Contact information: Contact information of the parties needed for incident response (e.g. the decision-makers of the organisation, technical support, contact information of the service providers).

Architecture and technical documentation: Descriptions of the architecture of the organisation's cloud environments and technical documentation, such as the configurations of cloud resources. When modern DevOps practices are in use, the configurations of cloud resources are typically available as code. It is usually also possible to generate automatic technical documentation by downloading the configuration of the SaaS and PaaS components from the cloud. This should always be done after changes, unless the changes are implemented with code and DevOps automation.

Data inventory: Data inventory or similar documentation can be used to determine which systems or storage locations contain personal data. This information is needed in order to determine the impact of the incident and identify an information security breach.

Instructions related to incident response: instructions on how technical measures can be implemented in incident response, such as isolating devices, gathering evidence or restoring data from a backup, for example. Technical investigation and restoration measures often include administrative elements related to roles, responsibilities and obligations, among other things, which is one important reason for planning and documenting these procedures in advance. For example, there are legal obligations related to the collection and processing of log data. So that such obligations can be taken into account at a sufficient level when operating during an incident, the measures should be planned and documented in advance. Also note that the instructions must be in line with the incident response plan, so that e.g. the decision-making chain matches the specifications of the incident response plan.

Incident response log template: a template for recording an event log of the incident. At the minimum, it must be agreed in advance which system or storage location is used during an incident to record and maintain an event log, and it must be ensured that sufficient procedures are in place for protecting the data.

Annual schedules for incident response and the security of cloud environments: Both the incident response measures and information security as a whole are processes that must be maintained and developed regularly. Make sure that updating the incident response plan, training for incidents and reviewing the best security practices of cloud platforms, for example, are included in the organisation's annual schedule.

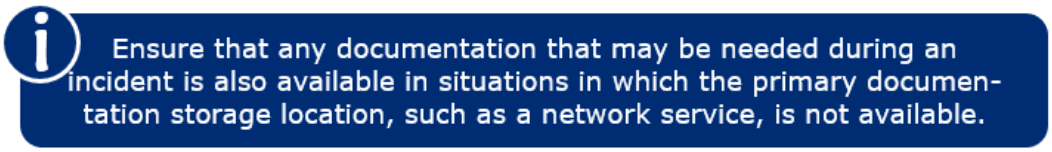

2.1.3 Define the roles and responsibilities of incident response

Make sure that your organisation has an up-to-date incident response plan in which the roles and responsibilities related to incidents have been agreed.

An incident requires cooperation between the different parties in the organisation. During an incident, an incident response team is assembled; it should include at least one representative of the management and one technical expert. If possible, the team should also include legal and communications experts as well as the technical expertise required for incident resolution. The roles related to information security incidents and the responsibilities related to the roles are described e.g. in the instructions of the Ministry of Finance 'Management of data security breach situations'[6]. If the organisation does not have the abovementioned resources available, the organisation can, at its discretion, obtain external experts to help with managing and resolving the situation. It is good to plan for the additional expertise that may be needed in an incident in advance.

When defining the roles and responsibilities, it is essential to take account of the role of the service providers, too, during incidents. The operating model for incidents should be agreed in writing with the service providers. With regard to cloud service providers, it is good to check the terms of the agreement or other documentation to see how the cloud service provider can help in case of an incident. Also make sure that the contact information of service providers can be found easily in case of emergency.

In situations in which a personal data breach is linked to the incident, the cloud service provider as a processor of personal data may be obliged to help with the investigation of the case.

2.2 Technical measures

This chapter describes the key technical measures used to improve the preparation for incidents.

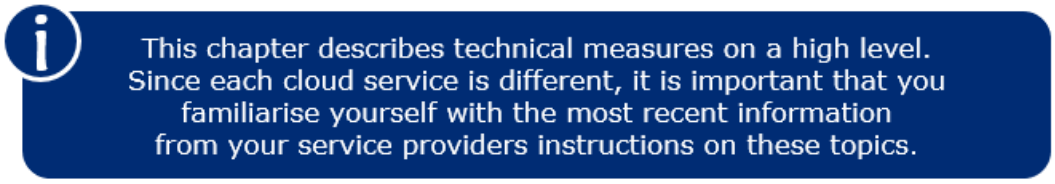

2.2.1 Design the cloud environment carefully

Attackers may gain access to the cloud environment by taking advantage of a vulnerability, for example, or by obtaining the user accounts and credentials of an authorised user. The service should be designed so that even if an attacker manages to access the service, accessing one target will not grant access to all of them. This is called a blast radius, i.e. how extensive an impact the measures caused by an attack can have.

In cloud services, it is typical that if an incorrectly configured storage service leaks data, all data are usually leaked at once. For this reason, it is especially important to take care of the access management of resources and restrict the

[6] https://julkaisut.valtioneuvosto.fi/handle/10024/79258

visibility of interfaces so that they cannot be accessed on the network level from the entire internet, which is usually the default for a cloud.

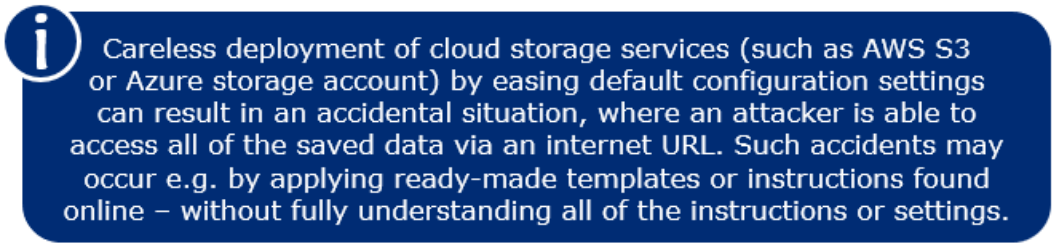

The impact of potential attacks can be minimised by using the Zero Trust design model, such as by minimising the access of users and services to different services:

- Grant users as few access rights as possible: only grant rights to the resources/services they need.
- Authenticate the user and device with several different pieces of authentication information at the time of use.
- Grant increased access rights only for the periods when they are needed (Just-In-Time authorisation).

Also aim to separate the services you use on the cloud platform into logical groups, such as:

- By dividing logical groups – such as each application, its data and support services – under separate accounts or orders.
- By dividing environments of different types based on usage (e.g. development, testing and production) into separate infrastructures.

Make sure that access between these parts is restricted or denied.

User accounts with access to several sections or environments should be restricted to administrative use only. The use of such rights should be implemented with e.g. the Just-In-Time model, and their use should be monitored.

It is usually possible to set limits on making unexpected changes in the cloud. For example, cloud services used for secret management, data services and container registers should be specified, as far as possible, so that their contents cannot be destroyed without a forced delay.

2.2.2 Follow the practices as instructed by the service providers

A part of the nature of cloud environments is that they reform and change frequently. Many cloud service providers aim to inform their users about changes and maintain up-to-date instructions on the best practices for the use and maintenance of the services, information security included. In order to ensure the information security of the service, take account of at least the following:

- Monitor the bulletins from the cloud services in use (e.g. email bulletins, web pages). The bulletins may contain information on e.g. new security features or known vulnerabilities.

- Follow the service providers' security recommendations during both the deployment of cloud services as well as their use. You can find instructions by international service providers by using the search terms "security best practices". Choosing not to comply with security recommendations should always be an intentional and reasoned decision by the organisation. Some cloud services have automations that support information security (e.g. security defaults or guardrails) that can be used to lock cloud configurations and changes or guide them to be safer. These possibilities vary by service provider.

- Harden services according to the service provider's hardening instructions (e.g. safe configuration, turning off default passwords and accounts, etc.). For international providers, you can find instructions with the search terms "hardening guide".

- Make sure that the new services and resources created for the cloud environment comply with the requirements by implementing standardised policies and templates. Standardised policies and templates establish standard settings for their targets in accordance with the specification. You can create the specifications yourself to match the requirements and architecture of your own organisation. Cloud platforms may also offer ready-made requirement templates based on e.g. good practices (such as AWS Well-Architected Framework, Azure Architecture Center).

- Take advantage of the automated code supported by the service provider (Infrastructure-as-Code, IaC) in the administration of cloud services, which makes the life cycle and change management of configurations clearer. The use of IaC is linked with good DevOps practices, and it speeds up the creation of new resources and the restoration of services in case of serious disruptions.

- Monitor the compliance of your cloud environment with good practices. You can use both automated assessments (e.g. Azure Security Score and Defender for Cloud, AWS Security Hub) or an external inspection for this purpose.

- Learn about the incident response documentation of the service providers. Cloud service platforms have service provider specific instructions related to incident management. Adopt the technical solutions in accordance with the instructions that are suitable for the solution of your organisation.

2.2.3 Ensure security of development and production environments

Attention should also be paid to the information security of application development in cloud environments. While there are plenty of instructions on the safe software development process model[7], attention should also be paid to the security of the development and production environments themselves, such as:

- Design development and production environments so that the environments for different purposes are in separate infrastructures.

[7] Example instructions: https://www.kyberturvallisuuskeskus.fi/en/publications/secure-development-towards-approval

- CI/CD automations are a common method of breaking into cloud production environments. Do not forget to take care of the information security of these systems, too.
- Schedule the information security work during development at as early a stage as possible (often referred to as a "shift left" on the project timeline). Take information security into account during development by using different kinds of tools and technologies for ensuring information security. When information security is integrated into the development pipeline correctly, it is possible to avoid using exposed components or configurations without information security.
- Design CI/CD automation, in which only a certain account is authorised to carry out installations and updates of the application. The use of this account must be within the scope of information security monitoring, and its use must be monitored in case of any incidents.

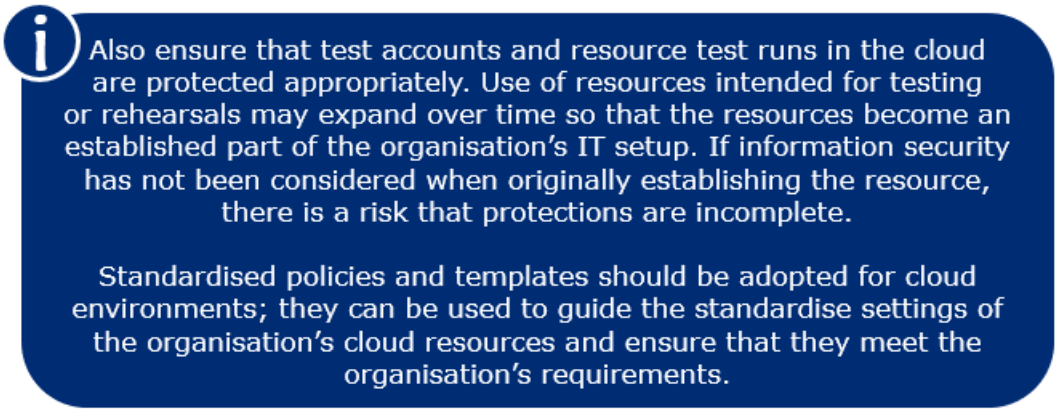

2.2.4 Reduce the attack surface

The attack surface of an organisation refers to its information systems, services and telecommunications ports that are open to the public network and against which the attacker could – at least in theory – attack. The concept of an attack surface is a way to concretise the idea that the more open information systems and services are towards the network, the better the organisation must also be able to detect potential attacks and prevent them. In order to reduce the attack surface:

- Ensure that only the cloud services and resources that the organisation needs are used.
- Harden the services used according to the hardening instructions (see Chapter 2.2.2 Follow the practices as instructed by the service providers).
- Note that access rights granted to users that are too extensive compared to the work tasks also increase the attack surface. Make sure that the access rights of users comply with the principle of least privilege, and that admin accounts and other increased access rights intended for specific work tasks are not used in the daily work.
- Create a process for removing unnecessary services and resources as well as users and access rights from use regularly.

2.2.5 Set a cost budget for the use of resources

Sometimes the measures taken by an attacker may include activating the cloud's resources so that the customer incurs significant financial damage due to the use of the cloud. In order to limit such an impact, rules should be set on the cloud's orders and accounts that limit the activation of expensive resources in addition to a cost budget that cannot be bypassed without the owner's approval. Cloud platforms have different features of this kind, and they should be studied in connection with deployment.

2.2.6 Ensure emergency access to services

For cloud service platforms, make sure that the organisation has configured emergency access for them. The purpose of these procedures is to ensure that the customer can administer its cloud services in all situations.

These access methods may be needed in exceptional situations, in which e.g. the multi-factor authentication method usually used by the organisation (such as an authentication application) is unavailable. In an emergency, main users have the option of authenticating themselves in the service by using a different procedure. These accounts are only used in exceptional situations, and procedures for their creation and storage must be established. The accounts should be intended solely for the organisation's emergency use, and therefore they should not be linked to the employees' accounts or identification means.

The strongest authentication method possible, such as the FIDO2 identification key, should be used for the accounts created in case of an emergency. The FIDO2 key is physical, which means that it can be stored behind lock and key, such as in the company's safe.

The use of emergency access methods should be within the scope of information security monitoring so that the use of the accounts trigger an alarm. This makes it possible to react if the accounts are misused.

Accounts created for emergency use should be tested at planned intervals to ensure that the necessary persons know the stages related to their use and that the accounts are still functional.

You also need to ensure that automation or risk-based access management practices (such as Conditional Access) related to the cleaning and removal of accounts and access rights do not prevent the operation of emergency accounts by mistake.

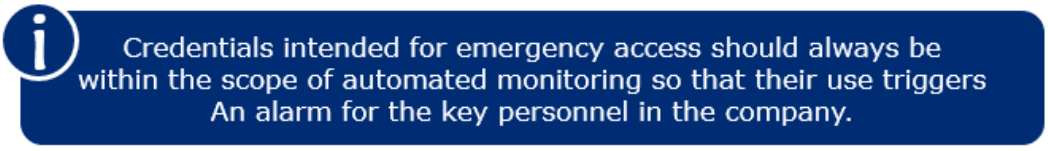

2.2.7 Ensure sufficient backup and recovery procedures

Make sure that sufficient backup and recovery procedures have been established for cloud services[8]. The cloud service platform may have some recovery procedures automatically in use, such as duplicating services in more than one single data centre. However, it is good to note that on cloud platforms, the customers get the services they pay for: in practice, the more advanced and robust backup and recovery procedures are usually services subject to an additional charge, and they are only used if the organisation has purchased them for its use and integrated them into its services.

If your organisation has not implemented backup services for the cloud services or ensured that the services are duplicated in several geographical locations[9], the cloud service provider probably has not automatically carried out these measures. In the worst case, the organisation may lose all of the resources it has built in the cloud and the data it has stored there as a result of an incident without sufficient backup and recovery measures.

When using container-based workloads (such as Docker), the services and applications can usually be restored faster than operating system based services. In recovery, it is critical that the register used for the distribution of containers is still available, or that one can be established and the codes recompiled from the source codes.

Ensure sufficient backup procedures by planning which cloud services require backups and how often the backups are taken. Cloud platforms have their own services and architecture recommendations for implementing the backups. When planning the implementation of backups, note that if you save the backup copies in the same cloud environment as their source data, the backups may be exposed to the same information security attacks. In the worst case scenario, an attack against the cloud environment could destroy both the original services and the data they contain as well as their backup copies. The backup copies can be distributed into accounts separated from the source systems (e.g. AWS account) or orders (e.g. an Azure subscription).

Testing restoration from backups is recommended regularly both as a technical measure and as a part of a crisis exercise. See also Chapter 2.3 Training on information security incidents related to the cloud environment.

2.2.8 Specify and implement logging

Sufficient logging plays a key role in detecting and investigating information security incidents. The logs collected by the customer may be the only way to prove an information security incident.

[8] The planning should be based on the Recovery Time Objectives (RTO) and Recovery Point Objectives (RPO) specified for the services. You can find further information on the management of the continuity of operations and its planning in the VAHTI 2/2016 instructions 'Control of continuity of operations' (in Finnish with a summary in English), for instance, at the address: https://julkaisut.valtioneuvosto.fi/handle/10024/75168.

[9] It is good to note that cloud service resources that have been placed in one single geographical location may be vulnerable to disruptions, if a disruption occurs in the service provider's infrastructure found in the location in question.

When planning logging, it is good to take into account that the logs may contain personal data, and the legislation sets requirements on the processing of logs[10].

Make sure that your organisation has introduced sufficient logging procedures in cloud services. Data collection rules that determine the collection of a log can be set for logging on cloud platforms. If the collection of log data from different services has not been specified, collecting log data in order to resolve an incident may prove very challenging.

Concerning logs, you should also consider the following issues:

- Can logs be used to see changes related to the configurations and main users of the cloud services?
- Based on the logs, is it possible to see who has used the service/resources, and when, where and why?
- For how long are the different kinds of logs available?
- Who can access the logs? Access must be restricted carefully and the processing of logs must be monitored.
- Are the logs and the services used to store them protected from changes? The integrity (i.e. constancy) of logs must be ensured to maintain the reliability of the data. Investigating information security incidents in a cloud environment clearly becomes more difficult if the attacker has an opportunity to destroy logs.
- How can network flow logs be implemented, and is any special expertise and/or tools needed to analyse them?

The matter is also discussed above in Chapter 1.3.5 Gathering information during security incidents.

2.2.9 Monitor environments and detect incidents

Detecting incidents during the normal operation and telecommunications of the cloud service requires monitoring the service. Different kinds of monitoring tools, usually subject to a separate charge, can be deployed on cloud platforms. In SaaS services that are located in IT environments maintained by the service provider, the service provider is usually responsible for detecting incidents.

The key issue in detecting incidents in the service is understanding the normal operation and telecommunications of the service (its baseline). When the normal operating modes, amounts and traffic are known, it is possible to detect activity and traffic that deviates from them.

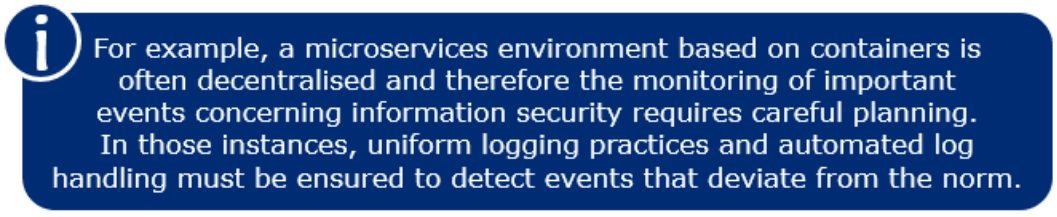

[10] Further information on the handling of log data can be found e.g. in the instructions by the National Cyber Security Centre Finland: https://www.kyberturvallisuuskeskus.fi/en/ncsc-news/instructions-and-guides/collecting-and-using-log-data

More information on detecting incidents can be found in Chapter 4 Detecting an information security incident.

2.2.10 Plan the technical implementation for investigating an incident

As a part of the technical preparations, you need to plan how the technical investigation into an incident could be implemented in practice:

- Are there sufficient logs available, and who can access them in case of an incident? How are the log data stored in a way that enables further investigation and the use of the data as evidence?
- Is there an environment available, in which the incident could be studied? In an ideal situation, the investigation into an incident does not take place in the same environment that the attacker can access. For example, cloud platforms may be able to create a separate forensic investigation environment.
- Is there enough expertise to use the right kind of searches on the log data to investigate the events? Plenty of log data are accumulated, and finding the necessary data may be challenging. It is good to study the logging format and making log searches when there are no incidents in progress.
- Make sure that there are written procedural instructions for the technical investigation of an incident, and that the persons carrying out the technical investigation have the sufficient competence for the purpose.

2.3 Training on information security incidents related to the cloud environment

In addition to administrative and technical preparation, an effective way of preparing for incidents is training on how to operate during an incident. In an incident, an organisation must investigate the incident in addition to running its normal business operations, which requires both quick decision-making as well as taking measures. Training for incidents familiarises the organisation with the practices followed during an incident, the decision-making model, the roles and responsibilities as well as the key measures. Training allows the participants to already have a good idea of the matters mentioned above during an actual incident, which helps with acting as agreed.

There are many different kinds of incident exercises, and their range covers both tabletop exercises as well as simulated attacks (red teaming, purple teaming) and different types of exercises in between. Organisations, in which incident response is not yet established, may benefit from a tabletop exercise that illustrates e.g. the responsibilities and requirements related to different roles during an incident. A tabletop exercise also makes it possible to review the measures recorded in the incident plan, discuss their effectiveness and find development targets in the plans in this way.

Training also makes it possible to detect potential deficiencies in the incident response capabilities. For example, training makes it possible to test if the organisation has taken the special characteristics related to the incident response in cloud environments described in Chapter 1.3 into account, and whether the incident response team is able to carry out the measures listed in Chapter 5 Instructions.

Regardless of the type of exercise, a training scenario relevant to the organisation must be selected for the information security incident training. For cloud environments these may include e.g. the following types of initial scenarios, in which the incident response measures are carried out:

- A data dump containing confidential data of the organisation is found in the dark web. How did it end up there?
- The attacker has encrypted some of the cloud service's resources and is now demanding a ransom. How can you recover from the situation?
- The cloud environment monitoring tool gives a notification of suspicious behaviour by the admin account. How do you act in such a situation?

Incident training can also be arranged together with the organisation's IT service providers, if they play an active role in case of an incident.

You can find more instructions for cyber training and several training scenarios on the website of the National Cyber Security Centre Finland[11].

[11] https://www.kyberturvallisuuskeskus.fi/en/ncsc-news/instructions-and-guides/instructions-and-guides-organisations-and-companies

3 The life cycle of an information security incident

The life cycle of information security incidents typically consists of the following stages:

- Detection
- Investigating the incident
- Recovery
- Post-incident review.

All information security incidents roughly follow this life cycle regardless of which services or information are affected by the incident. However, the role of an organisation targeted in the incident during the different phases of the life cycle depends on the type of cloud service model used and what the agreement says concerning operation during exceptional situations. The table below covers the key differences.

Table 1 Typical responsibilities of the customer related to the life cycle of an information security incident for different types of cloud services

Phase	IaaS	PaaS	SaaS
Detection	The customer is responsible for observing information security incidents in the cloud resources used by the customer. Typical methods for this include monitoring log sources, using information security monitoring tools, monitoring incidents in telecommunications as well as monitoring changes in the configurations.	The customer is responsible for the information security detection in its applications. Typical methods include monitoring the event logs of applications and application programming interfaces (API), monitoring configuration changes, as well as application-specific implementations of monitoring logics that describe potential misuse.	The customer's responsibility for detecting information security incidents is limited. Usually, a customer may detect one if the contents of the service change unexpectedly, there are strange users in the service, or if the use of the service is prevented.
Investigating the incident	The customer uses the available tools, such as information security, configuration and log analysis tools, to make observations and investigate the situation. The service provider usually helps with investigating the incident, at least when the incident may also affect its operations or the data and services of other customers.		During the investigation phase, the cloud service provider plays a central role. The customer must report its observations to the service provider, if necessary.
Recovery	The customer activates the measures in accordance with the customer's recovery plan and restores the data and cloud resources. Usually, recovery in the cloud is carried out by establishing new service instances. The service provider may help with the recovery measures. The scope and contents of service agreements may affect the help available.		The ability to recover largely depends on the solutions of the service provider.
Post-incident review	The customer finds out the root causes of the incident and analyses how successful the incident response was. The aim is to use the observations to improve the operations in order to prevent a similar situation.		The post-incident review is carried out in cooperation between the organisation and the service provider.

A key issue is that in SaaS services, the ability of the organisation using the services to investigate the incident and recover from it is often limited: in the SaaS model, often only the service provider has these capabilities. For this reason, only some of the measures listed in Chapters 4–6 are suitable for the users of SaaS

services. In the SaaS model, the organisations focus on investigating the situation together with the service provider in incident response, in addition to communicating about the situation to parties both within the organisation as well as outside it, if necessary, such as to the authorities.

4 Detecting an information security incident

Information security incidents, including cloud environment information security incidents, can be detected in many different ways. The organisation may detect the incident itself with e.g. monitoring or by detecting disturbances in the service. The organisation may also find out about the incident from external parties, such as customers, the media, white hat hackers or the service provider.

This chapter describes how cloud environments can promote the detection of information security events and incidents.

Alarms and security views

Large cloud platforms often have views or pages visible to the named users of the organisation that describe the security status of services, where administrators can see both security recommendations related to their services as well as alarms concerning information security events and incidents.

Various services used on cloud platforms may have their own security dashboards, and they can generate alarms concerning suspicious activity. For example, cloud-based centralised identity management technology (e.g. Azure AD, Google Cloud Identity, Okta) can offer the user organisation alarms in case of suspicious activity. However, such features may require user licences of a certain level and the deployment of automations.

In a SaaS service, alarms concerning suspicious activity may come from the service provider or the users of the service.

Cloud-based information security solutions

Large cloud platforms often offer separate security solutions that can protect the technologies located in the cloud (e.g. Microsoft Defender for Cloud, AWS GuardDuty). These solutions can be used to protect and monitor the security of cloud environments and resources. Such solutions can also be used to implement automated workflows that react to suspicious activity. For example, the solutions can automatically isolate a device on which suspicious activity has been detected. Creating, developing and maintaining the rules related to these solutions requires sufficient resourcing from the organisation.

Information security monitoring system

The organisation can increase its detection capabilities by investing into Security Information and Event Management (SIEM) as a separate system in addition to the cloud solutions described above. These kinds of centralised log systems can analyse log data and use them to detect information security events. Creating, developing and maintaining rules related to these solutions also requires sufficient resourcing from the organisation.

Notifications by service providers

Sometimes, the first observation of an information security incident comes from the cloud service provider. In these cases, it is important to follow the cloud service provider's instructions, such as by changing the passwords and other access information.

However, be careful and make sure that the observation actually does come from the cloud service provider – cyber criminals often use the names of familiar and well-known actors such as Microsoft and Amazon in scam messages that can be easily confused with real messages.

Notifications by customers and other parties

Sometimes the first observations about an information security incident come from the customers, white hat hackers, independent information security researchers or other unexpected parties. The organisation should have ways of processing such notifications and appropriate instructions for the customer service, for example. Publishing suitable contact information is useful for organising the process[12].

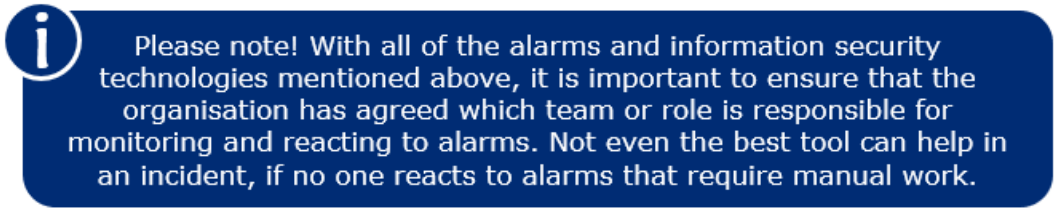

[12] For example, the organisation can specify a security.txt file on its website, in which it states that it would like to receive observations related to information security and lists the desired notification channel. Further information available e.g. at the address: https://www.kyberturvallisuuskeskus.fi/en/news/practice-facilitating-reporting-vulnerabilities-not-yet-widely-adopted-finland.

5 Instructions

You can use the checklists with measures in this chapter if you suspect that a cloud environment information security incident has occurred.

The checklist helps organisations prioritise and use a phased approach when investigating information security incidents.

5.1 Workflow of investigating a cloud environment information security incident

The diagram on the following page shows the key measures related to incident response. The diagram supports the use of the checklist, which is in table format. It is useful for the organisation to draw up a flowchart or swimlane diagram suitable for its own incident response process that corresponds to the procedures described in the incident response plan.

During the incident response process, it is also important to keep an accurate event log of the measures taken. The log should show the measure taken, the timestamp and the party that implemented the measure.

The gathering of potential evidence should also be documented carefully. You should record who gathered the data, what it was, and when and how it was gathered. A carefully drawn up event log makes the investigation as well as the cooperation with the police and information security investigators significantly easier.

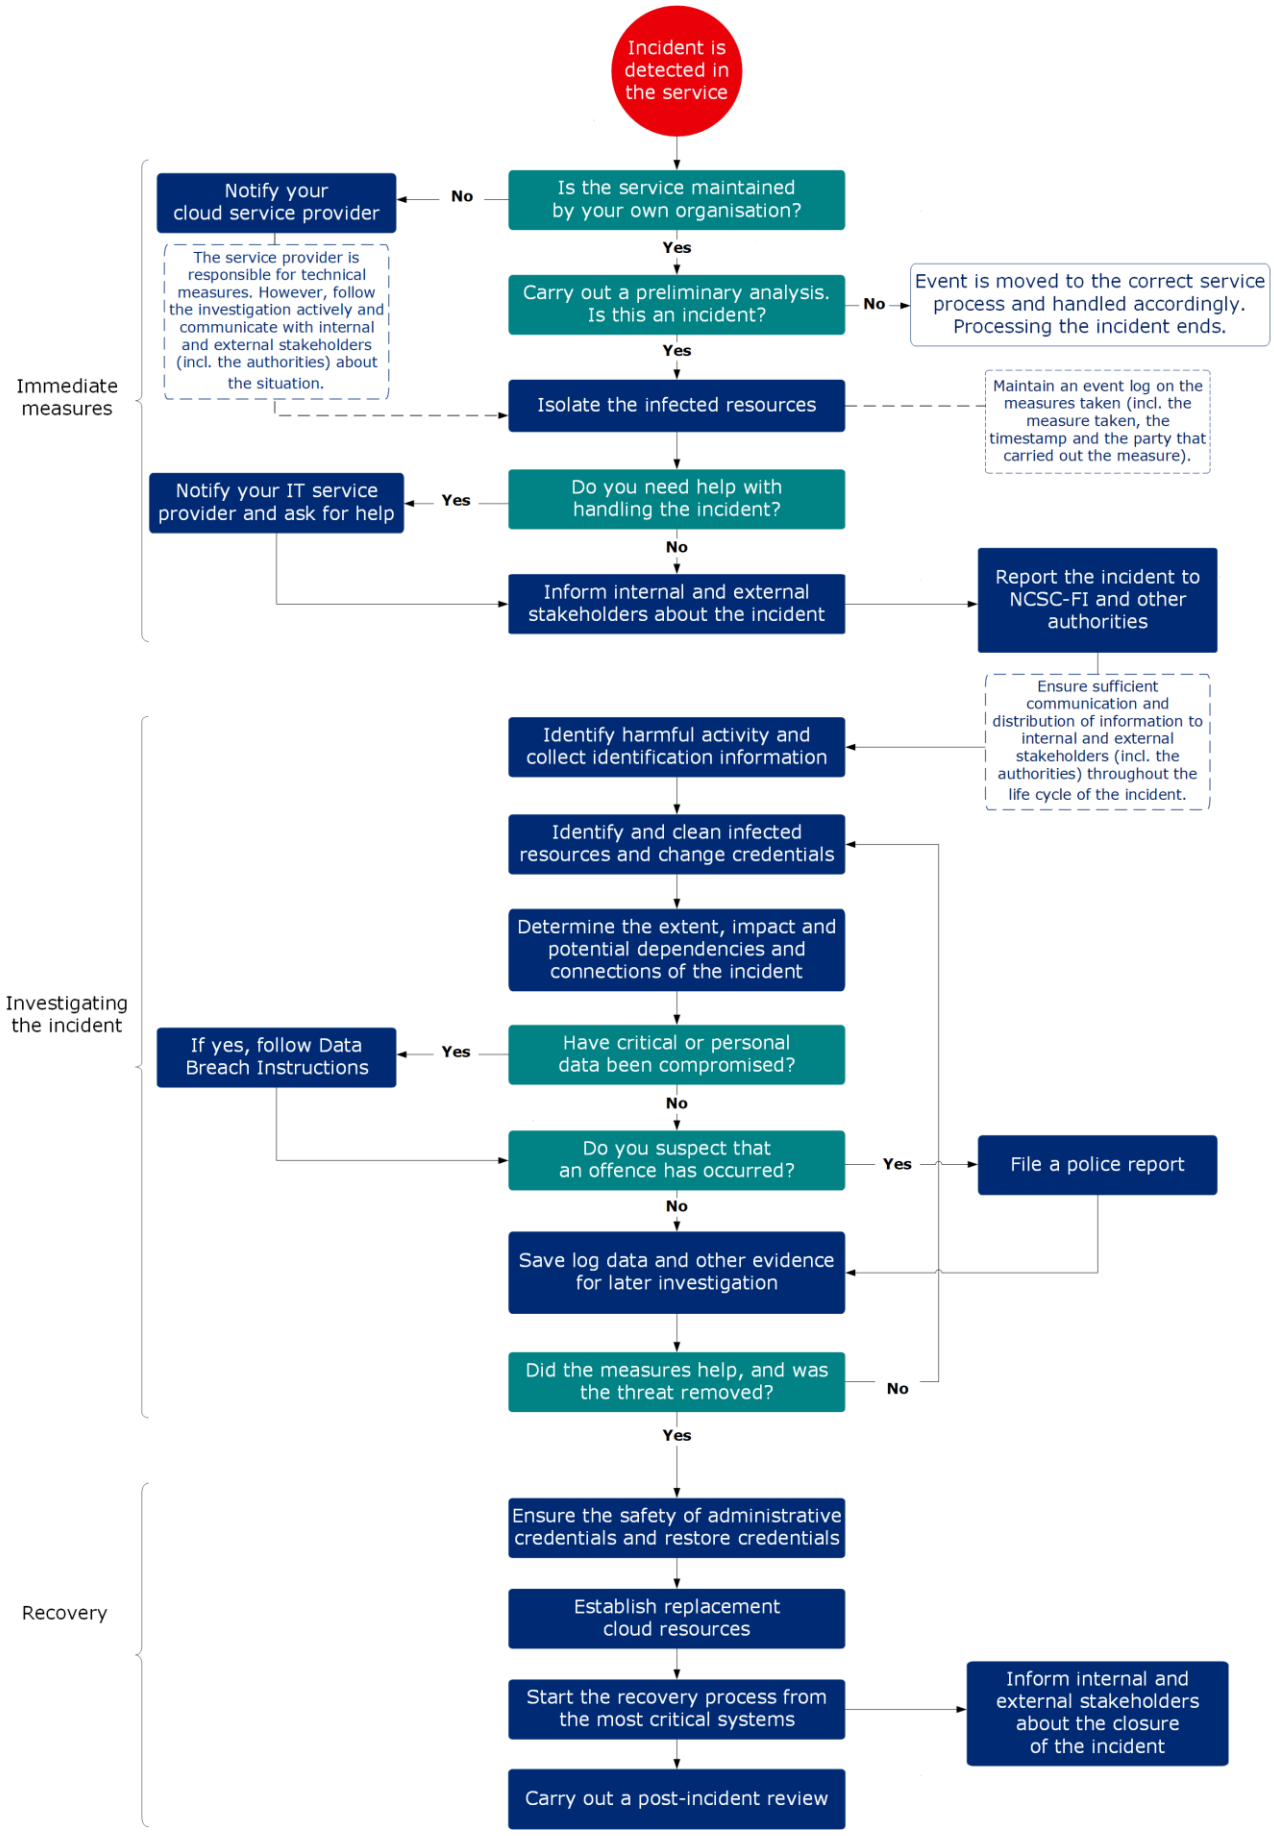

Image 2 Cloud environment incident investigation workflow, simplified diagram

5.2 Immediate measures

Goals of the phase	The purpose of the immediate measures is to protect the critical data in the environment, prevent the attacker from gaining a foothold in the environment and prepare for initiating the recovery process.	
Phase	**Purpose**	**Measures**
Incident is detected in the service		
Is the service maintained by your own organisation? **Contact your cloud service provider**	The help of the cloud service provider or IT partner may be needed to resolve the situation. The available support services vary depending on the technology area affected by the incident, and whether the customer or the service provider is responsible for it. IaaS, PaaS: Contact the IT partners or cloud service provider, if any, if the incident is in their area of responsibility. SaaS: The cloud service provider organisation is necessary to resolve the situation. Note that the service agreement may specify that the organisation is obliged to notify the service provider about the incident.	Locate the occurrence of the incident: identify the service or services in which the potential incident is occurring, and find out if they are maintained by your own organisation or if your cloud service provider or IT partner is responsible for them. If the service is maintained by the cloud service provider or IT partner, **contact the service provider's support channel**. The service provider is responsible for the technical measures, but you should still monitor the investigation into the incident and take care of the communications with the internal and external stakeholders (incl. the authorities), among other things. Go to section "Isolate the infected resources". If the service is maintained by your organisation, continue the incident response process (go to section "Carry out a preliminary analysis").
Carry out a preliminary analysis. Is this an incident?	Not all information security alarms and observations lead to incident response measures. Some alarms are so-called false alarms. Certain alarms and disruptions are taken care of as a part of the organisation's normal operations. The organisation should define what kind of an information security event or chain of events is severe enough to initiate the incident response process. <u>If necessary, initiate continuity management measures</u> Depending on the severity of the incident and in order to ensure the continuity of the business, it may be necessary to initiate exceptional procedures, start using an alternate website or carry out other similar continuity management measures. It may also be necessary to start partial recovery measures on a case-by-case basis in order to ensure the continuity of the business. In an ideal situation, restoring the systems only starts after the incident root cause analysis has been completed to ensure that the attacker no longer has access to the environments. However, for the continuity of the business, it may not be possible to wait for the investigation to be completed; instead, it may be necessary to start partial recovery earlier to ensure continuity. This creates the risk, that the system remains vulnerable after recovery and a new incident occurs.	Do a preliminary analysis: investigate the information on the alarm or other indication you have received, and use the necessary additional sources to evaluate whether this seems like an information security incident. If you think that this is or may be an information security incident, continue the incident response process. In other cases, the workflow specified for the purpose should be used. Processing the incident ends. Make an assessment on the severity of the situation and, if necessary, initiate continuity management measures (exceptional procedures, alternate website, etc.). In order to ensure the continuity of the business, it may be necessary to start partial recovery measures on a case-by-case basis even before the investigation is complete. In order to reduce the risk of an attack, the temporary restoration should be implemented, as far as possible, in a separate infrastructure by using isolated maintenance measures (accounts, devices, etc.). Monitor and assess the situation and update measures throughout the incident response process.

Isolate the infected resources	The aim of isolating the threat is to stop the attack from progressing and protect the data in the system.	Isolate the infected resource by using the cloud environment management tools. In a cloud environment, isolation measures may mean e.g. the following: • Shutting down the account or access token. • Stopping the operation of the application. • Limiting the activity of the application's interfaces and network traffic. • Closing the integration interface. • Shutting down the operation of the virtual machine or container-based service. Maintain an event log of the measures taken. The log should show the measure taken, the timestamp and the party that implemented the measure.
Do you need help with handling the incident? **Notify your IT service provider and ask for help**	External help may be needed in investigating the incident with e.g. technical measures, incident response or organising the measures. If the necessary expertise is not available within the organisation itself or directly from the IT service providers, you should consider getting external help. External expertise may be required e.g. for collecting identification information and investigating the threat based on it. External assistance may also be useful with checking whether the attacker was able to obtain data important to the business, and if so, what kind of data.	If you think that you will need external help with handling the incident, **contact a service provider** specialising in incident management. You can find Finnish service providers in the resources listed in the footnote[13].
Inform internal and external stakeholders about the incident	During an incident, it is important that the key decision-makers know what is going on. An information security incident may cause the partners, customers and service providers risks or problems with the availability of services.	Follow the incident response plan of your organisation and make sure that the incident response team is established and that information flows within the team. The necessary amount of communication is largely dependent on the situation: especially if a serious situation is suspected, the management should be notified of the situation immediately. Notify the contact persons in case of crisis situations of different stakeholders about the incident if you believe that it may affect the availability of their services or endanger their security. If the server has also had connections to other organisations, notify them about the issue, too. This will allow them to invalidate the accounts, keys or certificates used on the infected server. It is also important that they check the integrity of their own data.

[13] You can find Finnish service providers on the following websites: https://dfir.fi/; https://www.fisc.fi/fi; https://www.hansel.fi/en/framework-agreements/tiedonhallinnan-ja-digiturvallisuuden-asiantuntija/

Submit a report to the NCSC-FI and other authorities	The infrastructure operators and service providers critical to the security of supply that are subject to the NIS directive of the EU on the security of network and information systems must notify the authorities about information security incidents in network and information systems[14]. In addition, information security incidents should always be reported to the NCSC-FI, which offers support and advice to organisations in case of incidents. Please note that the report you have sent to the NCSC-FI is confidential in principle, and parties such as the police will not be notified about it. If you suspect an offence, filing a report of an offence with the police is a separate process. <u>Ensure sufficient communication and distribution of information throughout the life cycle of the incident</u> Incident response requires coordinated communication both within the organisation as well as with external stakeholders and the authorities. These communication responsibilities should already be agreed in advance and listed in the incident response plan.	If your organisation is an operator critical to security of supply, notify the supervisory authority of your own sector about the information security incident. In addition, the NCSC-FI should also be notified of information security incidents as early as possible.[15] The National Cyber Security Centre Finland can help organisations especially during the first response to the incident as well as by offering additional information on similar cases in Finland and abroad. Name the role(s) or person(s) tasked with sharing status information with internal and external stakeholders (incl. the authorities) and keeping them up to date throughout the life cycle of the incident.

[14] https://www.kyberturvallisuuskeskus.fi/en/services/report-security-incident-nis-notification-obligation
[15] Submit a report to the NCSC-FI with an online form https://www.kyberturvallisuuskeskus.fi/en/report, via email (cert@traficom.fi) or by telephone at the number +358 295 345 630 (local network rate/mobile charge); the number is open 9:00–15:00 on weekdays

5.3 Investigating the incident		
Goals of the phase	The goal of investigating the information security incident is to determine the extent of the attack and its impact on the organisation. A thorough investigation ensures that malware and potential backdoors have been removed from the environment.	
Phase	**Purpose**	**Measures**
Identify harmful activity and collect identification information	Identification information is collected to make it possible to map how widely the infection has spread to devices and how the stolen access rights have been used. Once attackers have gained a foothold, they may use different kinds of attack methods. In fact, identification information should be collected extensively and signs of their use should be studied carefully to ensure that the cleaning of the environment can be done reliably. Recovery can only start safely after the attacker has been removed from the environments.	Identify harmful activity and collect identification information as comprehensively as possible. The identification information gathered includes, among other things, the time when the incident occurred, when a login to the server occurred, or when a certain command was run on the server. The malware often communicates with the attacker's command and control server. By studying the network traffic of infected devices or domain name resolution (DNS logs), the source IP addresses or domain names used by the attacker can be identified. When harmful files are identified, their hashes (MD5/SHA256) can be extracted and used to identify harmful files on other devices, too. Authentication events related to infected devices and measures taken by the user accounts linked to them can be used to determine the accounts used to spread the malware. Endpoint detection and response often has features for collecting and using the identification information mentioned above. Otherwise, the measures should be taken manually by using a centralised log server. If no such server is available, either, the logs of individual servers and terminal devices should be examined.
Identify and clean infected resources and change credentials	The identification information can be used to find out how far into the organisation the attacker was able to penetrate. By collecting identification information and searching for it in the target systems, it is possible to ensure that all infected devices, identifiers, keys and certificates are found and cleaned. Identification information can be used to find infected devices, such as by using the endpoint detection and response features that often offer the option of searching for events on devices based on different identifiers. If the organisation has also implemented centralised log management, it can be used to search for events efficiently based on identifiers from several different devices at the same time. There is a risk that the attackers have attempted to cover their tracks by disabling logging after gaining access to a device. In that case, it may not be possible to find all of the collected identification information in the device logs. For this reason, it is important to aim to use a wide variety of identification information and event sources.	Use the identification information to help with identifying all infected systems and the identifiers, keys and certificates known to the attacker. You can use the data on the monitoring of terminal devices or log management, for example. If neither of the solutions mentioned above is available, identifiers should be searched separately from each device. Different kinds of remote control solutions can be used for the purpose, however; they often enable running PowerShell commands simultaneously on several servers, for instance. Clean the infected systems. Change the identifiers, keys and certificates known to the attacker.

Determine the extent, impact and potential dependencies and connections of the incident	Cloud resources often have active connections to other systems. Such connections may include a database connection or different kinds of API calls, keys and certificates. The integrity of the connected systems and the impact of the incident on business must be verified as soon as possible to determine how serious the situation is.	If the service has active connections to other systems, ensure the integrity of the data by reviewing the logs of the connected systems. The investigation should focus on the accounts, certificates or access tokens with which the integrations have been created.
Have critical or personal data been compromised? **Follow the data breach instructions**	As a part of the investigation, it should be determined whether the attackers were able to access important data of the organisation or potentially the personal data of customers or employees. Note that even if attackers have not destroyed or stolen the data, they may have edited them. The attacker may also have stolen data with a small size but great importance, such as accounts. In situations in which the cloud service provider acts as a processor of personal data and the incident may also involve a data breach, the provider is obliged to help with investigating the situation and filling in the report.	Find out if the accounts, certificates or keys used for the connections have been used to log in from a place other than the server on which they should be used. Find out if the attacker was able to access the data and steal them by reviewing the information service or interface logs. Based on the overload or the searches made, you can determine if the attacker intended to retrieve or edit information. If critical or personal data have been endangered, **follow the 'Tietomurto' data breach instructions of the NCSC-FI[16] (in Finnish).**
Do you suspect that an offence has occurred? **File a police report**	Report the incident to the authorities. The organisation may have an obligation to report the incident based on regulations or the terms of the cyber insurance. Please note that the report you have sent to the NCSC-FI is confidential in principle, and parties such as the police will not be notified about it. A report of an offence must always be filed separately.	If you suspect an offence, **file a report of an offence about the incident with the police[17].**
Save log data and other evidence for later investigation	The aim of collecting and storing evidence is to guarantee a high-quality investigation after the incident so that the root causes of the incident can be determined. Evidence may be needed for filing a report of an offence and the court proceedings. If the organisation has a cyber insurance policy, the insurance company may also require more detailed information on the security incident as well as evidence for the investigation.	Save log files that contain information relevant to the investigation of the incident on a hard drive isolated from the network. Also collect harmful email and other messages, if any. Aim to keep the evidence, such as complete disk images and memory samples, as intact as possible. Extract integrity hashes from them to ensure this. Aim to save samples of the malware detected. They should be handled with extreme care. Professional expertise is often required to carry it out safely. Send the samples to the National Cyber Security Centre Finland[18].
Did the measures help, and was the threat removed?	Before starting recovery, make sure that the measures taken have helped.	If the measures did help, start the recovery process from the most critical systems and services.

[16] https://www.kyberturvallisuuskeskus.fi/sites/default/files/media/publication/TietomurtoToimintaohje.pdf
[17] https://poliisi.fi/en/report-a-crime
[18] https://www.kyberturvallisuuskeskus.fi/en/news/transmitting-e-mail-and-sending-samples-national-cyber-security-centre-finland

5.4 Recovery		
Goals of the phase	During the recovery phase, the organisation aims to restore the business back to normal as quickly as possible once the recovery can be carried out safely. The recovery process should start from the systems and services that are most critical for business.	
Phase	**Purpose**	**Measures**
Ensure the safety of administrative credentials and restore credentials	Ensure that the login information of all of the potentially infected accounts is changed so that the attacker can no longer use the accounts to access the organisation's systems. Make the user login requirements stricter, if possible.	Change the passwords of infected accounts and start using the accounts again. Do the same to keys and certificates. To make sure, change the password of administrator accounts and service accounts in case some of them have fallen into the hands of the attackers. Deliver the new passwords to users either verbally in person, in a text message or by telephone. Do not use the organisation's email or instant messenger, because the attacker may still have access to them. Make sure that multi-factor authentication is in use for all accounts. Nevertheless, monitor the accounts used in the attack more carefully after the attack in case the attacker gains control of them again. If it is still unclear to the organisation how the attacker was able to gain control of certain accounts, consider creating completely new accounts.
Establish replacement cloud resources	Create new cloud resources by reinstalling the services and applications. Do not try to clean an infected system by using anti-malware or automated tools, because there is no guarantee that they will be able to clean the system completely.	One of the benefits of operating in a cloud environment is that reinstalling the resources can usually be done either completely or nearly completely automatically. In addition, when following the DevOps practices, applications and containers can also be distributed automatically. This makes restoring workloads considerably faster, because there is no need to clean and verify servers for use again. Of course, recovery must be implemented so that the risk of reactivating the attack is managed. This means e.g. ensuring that the attacker has not changed the contents of the container register or source codes so that the new resources contain a backdoor, for instance.
Start the recovery process from the most critical systems	The aim is to restore the data and systems and return to normal operation. The restoration should be done as safely as possible to ensure that the attacker cannot get back into the system. If it is suspected that the attacker has edited the contents of the database, the database must be restored from a backup copy to invalidate the attacker's changes, if cleaning the data is not possible.	When operating in a cloud environment, it is often enough to restore the data, and the infrastructure services, running platforms, containers and application codes that are run in the cloud can be reinstalled automatically. This is the recommended restoration method. If necessary, also restore systems, such as virtual servers, from backups. Also take account of the risk that previous daily (incremental) backups may already have been infected. When restoring old backups, keep in mind that the backup may include the vulnerabilities that the attacker used in the attack.

	Inform internal and external stakeholders about the closure of the incident Notify internal and external stakeholders (incl. the authorities) about the end of the incident. Notify stakeholders (e.g. customers) if it has been necessary to restore data to an older version starting from a specific date so that the interested parties can update their data.	You can try to avoid the risk of a new breach during recovery by restoring cloud services and workloads so that their interfaces and network connections are opened in a more limited way. For example, an interface that was previously open to the internet can be opened only to the necessary addresses or networks. Take advantage of logs to determine if the attacker has edited records. If the accuracy of the logs is not sufficient for cleaning up the changes, restore the data in the database from the latest safe backup. If the attacker has stolen data, all of the passwords contained by the stolen data must be changed. This should also be done even if passwords were only stored in the form of hashes.
Carry out a post-incident review	When the crisis is over and business operations have returned to normal, it is important to start the post-incident review of the incident and learn as much as possible about what happened for the future.	Carry out a post-incident review, identify the necessary measures and update the incident response plan.

6 Post-incident review

When the crisis is over and business operations have returned to normal, it is important to start the post-incident review of the attack and learn as much as possible about what happened for the future. At the same time, crisis management systems should be updated based on the observations made. The organisation may become a victim of a similar attack again, if the root causes of the incident cannot be determined and no lessons are learned from it.

During the post-incident review, the activities during the crisis are studied: what measures were done well, what could have been done better, and how the plans and the security level could be improved. A report should be drawn up on the post-incident review that examines at least the following questions in addition to the course of the events:

- Root causes of the incident:
  - What technical or functional weaknesses led to the situation?
- Effectiveness of the organisation's own protection:
  - Were the controls used to detect attacks sufficient?
  - Did the attacker's actions raise any alarms?
  - What was the reaction to the alarms like? Was the information about alarms transmitted to the right responsible persons?
- Actions during the crisis:
  - Was the crisis plan followed? How usable was it?
  - Were the responsibilities of the crisis management team assigned to the right people?
  - How successful was limiting the scope of the attack and removing the attacker?
  - How successful were the communications of the crisis management team? How were the stakeholders taken into account?
- Recovery:
  - How did the recovery of critical information and services go?
- Post-incident review:
  - Have the course of events and the investigation work been documented?
  - Was the technical investigation of the incident sufficient? Has it been possible to submit sufficient data on the attack for the use of the authorities, for example?
  - Evaluate the actions of the service providers. Were the response time and the services that were agreed upon sufficient for the investigation of the incident?

The organisation should update its own incident response plan and more detailed playbooks designed for combating different types of security incidents after the fact. Practicing different scenarios at regular intervals is also recommended to ensure that you can benefit from them in crisis situations.

The National Cyber Security Centre Finland hopes that the companies and organisations share the most important lessons they have learned from the incident with the Centre, too. With incident reports, the National Cyber Security Centre Finland can help other organisations in Finland as well as internationally to investigate similar cases. The lessons learned from recovery help with developing the preparedness of all organisations.

Appendix

Appendix 1

Informational content of the diagram in Image 1 (a typical model for the division of responsibilities, source: Criteria for Assessing the Information Security of Cloud Services (PiTuKri)) in table format (accessible version).

Table 2 A typical model for the division of responsibilities in table format

Cloud service type	Division of responsibilities
IaaS (Infrastructure-as-a-Service)	Customer / Tenant is typically responsible for: Interface, Application, Stack and Operating System. Provider is typically responsible for: Virtualisation, Compute & Storage, Networking and Facility.
PaaS (Platform-as-a-Service)	Customer / Tenant is typically responsible for: Interface and Application Provider is typically responsible for: Stack, Operating System, Virtualisation, Compute & Storage, Networking and Facility
SaaS (Software-as-a-Service)	Customer / Tenant is typically responsible for: Interface Provider is typically responsible for: Application, Stack, Operating System, Virtualisation, Compute & Storage, Networking and Facility

Appendix 2

Informational content of the diagram in Image 2 in table format (accessible version).

Table 23 Cloud environment incident investigation workflow presented in table format

Phase	Additional information or an alternative progression of the process
Immediate measures	
Incident is detected in the service	
Is the service maintained by your own organisation?	If no, notify your cloud service provider. The service provider is responsible for technical measures. However, follow the investigation actively and communicate with internal and external stakeholders (incl. the authorities) about the situation. Go to section "Isolate the infected resources". If yes, go to section "Carry out a preliminary analysis".
Carry out a preliminary analysis. Is this an incident?	If no, the event is moved to the correct service process and handled accordingly. Processing the incident ends. If yes, go to section "Isolate the infected resources".
Isolate the infected resources	Maintain an event log on the measures taken (incl. the measure taken, the timestamp and the party that carried out the measure).
Do you need help with handling the incident?	If yes, notify your IT service provider and ask for help. If no, go to section "Inform…".
Inform internal and external stakeholders about the incident	
Report the incident to NCSC-FI and other authorities.	Ensure sufficient communication and distribution of information to internal and external stakeholders (incl. the authorities) throughout the life cycle of the incident.
Investigating the incident	
Identify harmful activity and collect identification information	
Identify and clean infected resources and change credentials	
Determine the extent, impact and potential dependencies and connections of the incident	
9. Have critical or personal data been compromised?	If yes, follow Data Breach Instructions. If no, go to section "Do you suspect that an offence has occurred?"
10. Do you suspect that an offence has occurred?	If yes, file a police report. If no, go to section "Save log data…".
Save log data and other evidence for later investigation	
Did the measures help, and was the threat removed?	If no, go back to section "Identify and clean infected resources and change credentials". If yes, continue to section "Ensure the safety of administrative credentials and restore credentials".
Recovery	
Ensure the safety of administrative credentials and restore credentials	
Establish replacement cloud resources	
Start the recovery process from the most critical systems	Inform internal and external stakeholders about the closure of the incident.
Carry out a post-incident review	
(The process ends)	

Appendix 3

Information content of the information boxes in this instruction in table format (accessible version).

Table 4 Information content of the information boxes in this instruction in table format

Page	Information box
Page 3	Please note! These instructions have been created in 2023. The features of cloud services develop rapidly, and some of the technologies or practices mentioned in these instructions may no longer be up to date at the time of reading the instructions. The most up-to-date instructions on e.g. the safety features of a specific cloud service platform can be found in the instructions of the service in question. You can also check if there is a later version of these instructions available in the collection of instructions by the NCSC-FI at https://www.kyberturvallisuuskeskus.fi.
Page 7	Cloud security posture management (CSPM) products are often used in cloud environments; they are solutions tailored for the management of cloud resource information security configurations on cloud platforms. The products may also include and integrate the detection of vulnerabilities of and threats to cloud workloads. Examples of such commercial products include Azure Defender for Cloud and AWS Security Hub. The cloud customer may also choose to use products licensed for data centre environments, such as vulnerability scanners and EDR products. It is important to recognise that when working with IaaS and PaaS platforms, the customer is responsible for the procurement and deployment of these products and management of observations from these products.
Page 8	For example in the Azure cloud environment, control plane actions generate resource log type log events that represent actions taken in relation to the lifecycle of the services (such as actions taken to create services or perform configuration changes). However, the actual use of these resources (such as databases, storage services or Azure Key vault secrets) does not automatically generate activity log type events; instead the customer is responsible for enabling and managing the lifecycle of these activity logs.
Page 11	Cloud services offer increasingly broader and diverse features. Even though service providers do not typically weaken information security features of their services, this can still sometimes happen unexpectedly. On the other hand, information security features can also improve over time. Therefore, it is important to ensure that competent personnel follow the development of services and are able to evaluate how the changing features are used safely and effectively.
Page 12	Ensure that any documentation that may be needed during an incident is also available in situations in which the primary documentation storage location, such as a network service, is not available.
Page 13	This chapter describes technical measures on a high level. Since each cloud service is different, it is important that you familiarise yourself with the most recent information from your service providers instructions on these topics.
Page 15	Careless deployment of cloud storage services (such as AWS S3 or Azure storage account) by easing default configuration settings can result in an accidental situation, where an attacker is able to access all of the saved data via an internet URL. Such accidents may occur e.g. by applying ready-made templates or instructions found online – without fully understanding all of the instructions or settings.
Page 16	Also ensure that test accounts and resource test runs in the cloud are protected appropriately. Use of resources intended for testing or rehearsals may expand over time so that the resources become an established part of the organisation's IT setup. If information security has not been considered when originally establishing the resource, there is a risk that protections are incomplete. Standardised policies and templates should be adopted for cloud environments; they can be used to guide the standardise settings of the organisation's cloud resources and ensure that they meet the organisation's requirements.
Page 17	Credentials intended for emergency access should always be within the scope of automated monitoring so that their use triggers An alarm for the key personnel in the company.
Page 19	For example, a microservices environment based on containers is often decentralised and therefore the monitoring of important events concerning information security requires careful planning. In those instances, uniform logging practices and automated log handling must be ensured to detect events that deviate from the norm.

Page 24	Please note! With all of the alarms and information security technologies mentioned above, it is important to ensure that the organisation has agreed which team or role is responsible for monitoring and reacting to alarms. Not even the best tool can help in an incident, if no one reacts to alarms that require manual work.

Finnish Transport and Communications Agency
Traficom
National Cyber Security Centre Finland
PO Box 320, FI-00059 TRAFICOM
tel. +358 29 534 5000

kyberturvallisuuskeskus.fi

ISBN 978-952-311-879-9
ISSN 2669-8757 (online publication)

NATIONAL EMERGENCY SUPPLY AGENCY

TRAFICOM
Finnish Transport and Communications Agency
National Cyber Security Centre