PROJECT

Multi-tenant cybersecurity operations centre designed
as an open cloud service with machine learning elements

SOC centre concept

Linked to work package KPB2 – Design of the SOCaaS service

Linked to output
V6 – Description of dashboards, the monitoring and reporting system

Funded by the European Union NextGenerationEU

RECOVERY AND RESILIENCE PLAN

GOVERNMENT OFFICE OF THE SLOVAK REPUBLIC

Introduction to SOC 5
Concept of a cybersecurity operations centre 5
Improved collaboration 6
People working in the SOC centre 6
Tier 1 Security analyst – alert analyst 6
Tier 2 Security analyst – incident responder 7
Tier 3 Security expert – threat hunter 7
Security engineer 7
SOC manager – operations and management 7
CISO – chief information security officer 7
Certifications 8
Basic processes of the SOC centre 8
Event classification and triage 8
Ticketing process 9
Prioritization and analysis 9
Remediation and recovery 9
Assessment and audit 9
Tools 10
Prevention vs. detection 10
Network Security Monitoring (NSM) 11
Asset discovery 11
Vulnerability Assessment 11
Behavioural monitoring 11
IDS 12
Endpoint Detection and Response EDR 12
Network Detection & Response NDR 13
eXtended Detection & Response XDR 13
SIEM 13
SOAR 14
Other centres and units dealing with operations or security 14
SOC 14
NOC 15
CSIRT 15
Network security data 15
Types of security data 16
Alert data 16
Host logs 16
Session data 16
Transaction data 17
Full packet capture 17
Statistical data 17
Host logs 18

Syslog 18
Server logs 18
Logging from intermediary network devices 18
Proxy logs 18
SIEM 19
Next-Generation Firewalls (NGFW) 20
Information on network flows and applications 21
NetFlow 21
Application Visibility and Control 22
Logs as output of specific content filtering (Content filter logs) 22
Alert evaluation 23
Sources of alerts 23
Generation of security alerts 23
Overview of security alert evaluation 24
The need for alert evaluation 24
Evaluation of security alerts 24
Deterministic analysis and probabilistic analysis 26
Digital forensic analysis 26
Digital forensic process 27
Types of evidence 28
Recommendation for evidence collection 29
Chain of custody 29
Data integrity and protection 30
Attack attribution 31
The MITRE ATT&CK Framework 31
The Cyber Kill Chain 32
Diamond Model of Intrusion Analysis 33
Traversing the Diamond Model 34
The Diamond Model and the Cyber Kill Chain 34
Incident response 36
Incident response policy, plan and procedure 36
Elements of the incident response policy 36
Elements of the incident response plan 36
Elements of the incident response procedure 36
Stakeholders in the incident response process 36
Cybersecurity model certification 37
NIST incident response life cycle 38
Hardening based on lessons learned 38
Collection and retention of incident data 38
Reporting requirements and information sharing 39
References 39
Annex: List of tools 41

Introduction to SOC

Concept of a cybersecurity operations centre

A Security Operations Centre (SOC) is a centralized unit that deals with security and encompasses the people, processes and technologies whose main task is the continuous monitoring, prevention, detection, investigation of and response to cyber threats. SOC teams are responsible for monitoring and protecting the organization's assets and implement the organization's overall cybersecurity strategy.

Prevention and detection provide an advantage, because preventing attacks is always more effective than responding to attacks. Through monitoring, SOC teams are able to detect suspicious activities and eliminate them before they can cause damage. In this step, as much information as possible needed for further investigation is also obtained. During the investigation, the nature of the threat and the extent to which it has managed to penetrate the infrastructure are determined. Indicators and compromised areas are also sought before they are exploited. The final step is the response, or rather the reaction to the incident. As soon as the incident has been confirmed, the SOC team performs all the necessary actions, such as isolating endpoint devices, terminating malicious processes, deleting files and the like. After the incident has ended, the SOC team tries to restore the systems and data that were lost or compromised. If this step was successful, the network should be in the state it was in before the incident. When restoring systems, it is necessary to take into account the main components of disaster recovery and business continuity planning. One of them is the recovery time objective (RTO), which represents the maximum tolerated outage interval of a given device, system, network or application that may occur after a failure or critical event.

A SOC brings numerous benefits to organizations and their security. Here we mention a few of them.

Centralized knowledge

With a large number of attacks, there is a high chance that some of them will be overlooked. A SOC provides better equipment for cyber specialists, which ensures them a complete overview of the entire network and potential vulnerabilities. It also provides central

storage and sharing of information, thanks to which the whole team has the same information, which helps to identify and neutralize threats more effectively.

Cost control

In general, SOCs are more cost-effective for organizations. Whether it is locating the team in one place, in one building, or more effectively preventing attacks and damage, all of this can save the organization money.

Improved collaboration

Bringing the SOC team together in one place also benefits cooperation between team members, as they can better coordinate and divide work among themselves, thereby avoiding duplicate activities. Because the SOC is centralized in one place, the response time to threats is also shortened, which in turn shortens the time attackers have to carry out a successful attack.

Skilled expertise

A SOC brings together cybersecurity experts with various skills in this field, which they can combine and use in developing cybersecurity solutions. The last of the advantages we mention is sophisticated network monitoring, for which the SOC uses the latest monitoring and detection tools that help to better detect potential weak points in the network and address them before they become a problem.

People working in the SOC centre

Positions in the SOC centre have a hierarchical structure that corresponds to the level of expertise of the staff. This structure is shown in Figure 1.1. The people working in the SOC centre are divided into three tiers – Tier 1 to Tier 3 – which are complemented by the security engineer, the SOC centre manager and the CISO as the leader of the whole centre.

Tier 1 Security analyst – alert analyst

Manages and configures monitoring tools, runs vulnerability scans and reviews vulnerability reports, also reviews the latest alerts and determines their relevance and urgency,

and creates tickets for alerts that signal an incident and require review by Tier 2 analysts. More on tickets is given in the section Ticketing process in chapter 1.1.3 Basic processes of the SOC centre.

Tier 2 Security analyst – incident responder

Reviews tickets received from Tier 1, uses threat intelligence to identify the affected systems and the scope of the attack, and reviews and collects data on system assets for further investigation. Determines and manages the remediation and recovery of systems.

Tier 3 Security expert – threat hunter

Evaluates data on assets and vulnerabilities, investigates ways of identifying hidden threats that may have entered the network without being caught by any detection system, using the latest threat intelligence, performs penetration tests in order to verify resilience and identify weak points that need to be fixed, and recommends how to optimize monitoring tools.

Security engineer

Also referred to as a Security architect, this role completes the section of operational staff of the SOC centre. Manages and designs tools for monitoring and analysis. Builds the security architecture and cooperates with developers on securing this architecture. This may be a software or hardware specialist who pays due attention to security aspects in the production of information systems. Provides tools and solutions that enable the organization to respond effectively to attacks. [3]

SOC manager – operations and management

Oversees the activities of the SOC team, recruits, hires and trains employees, manages processes, reviews incident reports, measures SOC performance metrics and communicates the value of security operations to company executives.

CISO – chief information security officer

Defines the organization's security operations. Interacts with management on security issues and tasks related to regulatory compliance. The CISO provides the ultimate view of the policies, strategies and procedures concerning the organization's cybersecurity. The CISO also

has a primary role in regulatory compliance, risk management and the implementation of policies to meet specific security requirements. [3]

Certifications

Certifications available from various organizations can help us become a relevant SOC centre worker. Here we give a few examples:

- Cisco Certified CyberOps Associate
- CompTIA Cybersecurity Analyst Certification
- Information Security Certifications (ISC)²
- Global Information Assurance Certification (GIAC)

Other recommended skills include, for example, programming in Python or knowledge of Linux.

Basic processes of the SOC centre

In this section, we briefly describe the basic categories of processes performed by analysts of the levels described above in the SOC centre.

Event classification and triage

To express the real value of log and data collection, the literature often states that effective collection gives analysts the opportunity to find the "signal in the noise". Key indicators of a suspicious event may be found, for example, in user activity or system events. Specific sequences and combinations of these events in specific patterns may signal an event that requires attention. We describe the individual monitored events in more detail in section 1.2.

The key to success in this phase is considered to be having a way to quickly classify each event so that we can prioritize and escalate critical events that require additional investigation.

Ticketing process

Tickets are used to track issues related to assets or (security) events. A ticket is a work request created in response to a situation that requires further investigation. When the analysed data on activities and events, mostly of a security nature, give rise to a suspicion of a possible problem, such a ticket is created and then assigned to another staff member, who investigates and resolves the problem. [4]

Prioritization and analysis

Setting priorities is the key to success in any field, especially in the field of cybersecurity. It is necessary to focus on those events that could have the greatest impact on business operations. Here it is necessary to know which assets are the most critical and to mark them as key to delivering business processes in the organization (business-critical assets). Maintaining the organization's business continuity is the most important responsibility of the SOC team.

Remediation and recovery

The faster the SOC team can detect and respond to an incident, the more likely it is to be able to limit the damage and prevent a similar attack in the future. At this point, it is necessary to be able to make the right decisions and to consider whether there is greater interest in recovery after the incident or in a thorough investigation of the cybercrime. That is why cooperation and communication with the management team and documentation of the necessary information are very important.

In some cases, the SOC team leaves the remediation and recovery process to other IT departments. In such a case, the SOC analyst creates a ticket and/or a change control request and delegates it to the responsible persons.

Assessment and audit

It is always optimal to find and fix weak points before an attacker exploits them to gain access to the given network.

The best way to do this is to perform regular vulnerability assessments and to examine their results in detail. These assessments identify technical rather than procedural shortcomings, so it is necessary to make sure that the team also addresses gaps in SOC processes that could also expose the team to risk.

In the process of information risk management, which is managed by the person in the CISO role. A necessary part of this person's remit is also the established risk management process, which covers both information and operational risks of the identified assets of the given organization. If some of the identified risks also concern the people, processes or technologies of the SOC centre, the person in the role of SOC centre manager may be designated as the owner of these risks, and it will then be their duty to address these risks in accordance with the agreed rules and procedures in the organization.

Tools

In this chapter, we present the basic division of the tools used by analysts in the SOC centre according to their purpose. In each category, there are several available solutions from various vendors. Some tools may be commercial, others are available as open-source solutions. For each category, a specific tool will be mentioned as an example, possibly also an open-source variant. Their alternatives and other options will be given in Annex A List of tools.

Prevention vs. detection

There are several different approaches to defence in cyberspace; broadly speaking, however, these approaches fall into two basic categories, namely preventive and detective. In a preventive strategy, the organization places emphasis on so-called hardening. In a detective strategy, the security team focuses on the proactive identification and remediation of threats that have breached the organization's security. The preventive approach is the most common, but nowadays it is no longer as effective as in the past, or rather it is not advisable to rely on it alone. Therefore, adding a detective approach to an organization's security strategy is becoming increasingly necessary for protection against modern cyber threats.

Network Security Monitoring (NSM)

This is an automated process that monitors devices (endpoint, network) and traffic for security vulnerabilities, threats or suspicious activities. Organizations can use it to quickly detect and respond to cybersecurity breaches. The primary objective is to provide a continuous service that checks the environment for the presence of suspicious activities and threats. Staff can then examine the reported activities and take appropriate corrective measures. [5]

Asset discovery

The first step in protecting a network is to know the network itself and what is in it. It is necessary to know which systems (laptops, servers) exist in it and what is installed on them (applications, services, active ports and others). A reliable asset inventory, together with an automated capability to discover new ones, is the foundation for building a SOC. At the KIS department, asset inventory tools such as Lansweeper and SpiceWorks have been tested in the past, as well as tools that also provide the rendering of topological maps of the entire network infrastructure, such as Auvik and LibreNMS [6]–[8] .

Vulnerability Assessment

Vulnerabilities represent small "cracks" that attackers exploit to infiltrate a network, devices, systems or applications. This is often referred to as the attack surface. It is therefore essential to continuously check and evaluate the vulnerabilities of the entire network. Moreover, this may be required by various regulations, laws and directives. These include, for example, American directives such as PCI-DSS or SOX. In Slovakia, these are mainly Act No. 69/2018 Coll. [9] and NBÚ Decree No. 362/2018 Coll. [10] In this category, tools such as Nessus, Nexpose, OpenVAS (Greenbone), Owas-zap, nmap and others can be used.

Behavioural monitoring

Effective cybersecurity monitoring consists of exception management and tracking the behaviour of the network itself. Exceptions mean, for example, violations of established rules, error messages, major changes in network activity, unexpected restarts and the like. For this, it is necessary to create a so-called baseline, which represents the behaviour of the network in normal operation. Creating such a baseline of system and network behaviour provides the

necessary basis for detecting anomalies, which often signal the presence of suspicious activity in the network. When creating it, it is important to combine several available behavioural monitoring technologies in order to provide a complete perspective. Moreover, applying correlation rules to these data will help to identify and classify the latest risks, as well as to capture data to support in-depth forensic investigation.

IDS

Intrusion detection at the point of entry into the network can have the greatest impact on reducing the risk of compromise of the systems located in the organization and of data leakage. This is why an Intrusion Detection System (IDS) is considered one of the essential SOC tools for identifying known attacks and known attacker activity. Unlike the behavioural monitoring tools we described above, an IDS works on the basis of signatures or rules that look for known patterns of suspicious network traffic and system activity that could signal an intrusion. When this activity or traffic is detected in the network, it is necessary to know about it immediately so that an investigation can begin with the aim of limiting any damage. It is a type of security software that automatically alerts administrators when someone or something attempts to compromise an information system through malicious activities, such as DDOS attacks or security policy violations. An IDS works by monitoring system activity through examining vulnerabilities in the system, file integrity and pattern analysis based on already known attacks. It also automatically monitors the Internet to look for the latest threats that could lead to a future attack. The tools are further divided into HIDS and NIDS. For HIDS, tools such as Tripwire and OSSEC are used, while Snort and Suricata are considered tools of the NIDS category.

The common denominator of the next three groups of tools is detection and response. In the text below, we present the division and definitions of these tools according to source [11] .

Endpoint Detection and Response EDR

Every device connected to the network represents a potential attack vector for threats from the Internet, and each of these connections is a potential gateway to the data stored on the given device. In general, EDR solutions collect data from endpoints, use it to identify potential

threats and provide useful ways to investigate and respond to these potential threats – modern solutions automate this process, followed by the creation of a report (reporting). Examples of EDR tools are SymantecEDR, SentinelOne or Enterprise Inspector from Eseta OpenEDR, which represents an open-source variant.

Network Detection & Response NDR

This group of tools has evolved from conventional network security and is a subfield of Network Traffic Analysis (NTA). It provides full visibility of known and unknown threats passing through the network. The solutions usually provide centralized, machine-based network traffic analysis and response solutions, including efficient workflows and automation. Positioning in the network and the assistance of machine learning provide a complete overview and analysis of the network with the aim of identifying and eliminating, in particular, unauthorized movement of attackers in the network (lateral movements). Tools in this category include Stellar Cyber, Vectra NDR, Reveal(x) and Flowmon.

eXtended Detection & Response XDR

This group of tools extends the potential of EDR systems by means of stronger artificial intelligence as well as an automation approach. XDR not only integrates endpoints but also provides an overview of the enterprise network by going beyond a single-vector point solution and encompassing network traffic between various devices, as well as applications for analysis and evaluation.

Together with the use of SIEM, this correlation and visibility can be further strengthened. Alert data and events can be supplemented with further (meta)data that facilitate mitigation – the prevention of attacks – and remediation – the restoration of systems after attacks, e.g. by malware. This category includes Cynet Auto XDR, Trend Micro, InsightIDR and Cortex XDR.

SIEM

The collection and analysis of system events from the entire network provides a wealth of source material that can be used to obtain information about suspicious activities. SIEM tools were developed on the assumption that by looking for certain patterns of activities and

sequences of events, the SOC team can uncover a cyber attack. SIEM tools provide the foundation for building a SOC thanks to their ability to apply dynamic correlation rules against a large amount of heterogeneous log data in order to find the latest threats. We describe SIEM in more detail in subsection 1.2.2. Representatives of SIEM tools include, for example, Splunk, IBM Qradar, or the open-source SIEMonster and the ELK stack solution.

SOAR

SOAR platforms are similar to SIEM in that they aggregate, correlate and analyse alerts. In addition, SOAR technology integrates threat intelligence and automates the workflows of incident investigation and response, based on playbooks developed by the security team. SOAR places emphasis on integration tools and the automation of SOC workflows. It orchestrates many manual processes, such as the investigation of security alerts, which require human intervention only when necessary. This enables security staff to deal with more urgent matters and with the precise investigation and remediation of threats.

Later, in section 1.3.1, we will summarize all the tools mentioned above from the perspective of possible sources of alerts and present them in a clear diagram of tools.

Other centres and units dealing with operations or security

In the field to which this work belongs, the question often arises of the differences between a Security Operations Center (SOC), a Network Operations Center (NOC) and a Computer Security Incident Response Team (CSIRT). Although there are objectives that these centres have in common, their differences are evident. The differences are best captured by the following sources [12], [13] .

SOC

The security operations centre focuses on all matters related to security, i.e. threat detection, monitoring, analysis, incident response and forensic analysis. The SOC ensures availability and protects the network by creating and continuously improving the security architecture and infrastructure protecting the IT resources of the given organization. The SOC centre protects the network against human-made threats, such as malware, viruses and other

cyber attacks. The SOC is a centralized place where the team works 24/7/365 to protect the IT resources of security.

NOC

A Network Operations Center (NOC) focuses on the installation, maintenance, performance and availability of the network. Its task is to ensure that access to the network, servers, applications and data is always available and that they meet the needs of the organization and the Service Level Agreements (SLA). The NOC primarily focuses on the delivery of services and applications, operation, maintenance and prevention of/recovery from operational disasters and other unexpected events. Like the SOC, the NOC is a centralized place that provides the above-mentioned services.

CSIRT

A Computer Security Incident Response Team (CSIRT) is a centralized unit whose objective is to manage security incidents in the organization and to respond to them. It may fall under the SOC or may operate independently as the main security organization. A CSIRT differs from a SOC in that it is usually a conglomerate of roles within the enterprise that are involved in all types of incident response functions, including public relations (PR), marketing, customer support and management. The ultimate objective of a CSIRT is to minimize and control the damage resulting from an incident. A CSIRT must deal not only with the threat itself, but also communicate with customers, the board of directors and the public about the incident. And, for example, if the event was caused by an internal actor, it will be necessary to take disciplinary and possibly also legal action against the employees concerned.

In general, the inclusion of the above-mentioned SOC, NOC and CSIRT centres in the security portfolio of a company or organization depends greatly on its size, structure and needs. As the above-mentioned sources have stated, these centres may also be combined or merged.

Network security data

In this chapter, we deal with the data that analysts in the SOC centre work with and that can be collected from network devices, analysed, evaluated, handled and responded to. We describe the basic types of data and the sources from which they originate. SIEM, which

represents a grouping of tools working with these data, will also be mentioned here. Unless stated otherwise, the source of the following sections is the Cisco CyberOps Associate course [2] .

Types of security data

The basic types of data can be divided into six categories, namely alert data, data from host devices, session and transaction data, full packet capture data and statistical data. In this part of the work, we briefly describe the categories mentioned.

Alert data

Alert data consist of messages generated by intrusion prevention systems (IPS) or intrusion detection systems (IDS) in response to traffic that violates a rule or matches the description of a known exploit, a so-called signature. From now on, we will use only the term signature. Network IDS (NIDS) are supplied configured with rules for known exploits.

Host logs

Host-based intrusion detection systems (HIDS) – these intrusion detection systems run on individual user devices. From now on, we will use only the term host-based IDS, or HIDS. A HIDS not only detects intrusions but, in the form of so-called host-based firewalls (security gateways on user devices), can also prevent intrusion. This software creates logs and stores them on the host device. In enterprises, it is difficult to have an overview of what is happening on individual host devices, which is why many solutions include a function for sending logs to centralized log management servers. In this way, logs can be searched from a central location using NSM tools. HIDS systems can use agents to send logs to centralized servers.

Session data

Session data are a record of a conversation between two network endpoints (most often a client and a server). Session data are not the data themselves obtained and used by the client. Session data include identifying information, the five tuples: the source and destination IP address, the source and destination port numbers and the type (IP code) of the protocol used. Session data usually also include a session ID, the amount of data transferred by source and destination, and information related to the duration of the session.

Transaction data

Transaction data consist of messages that are exchanged during network sessions. These transactions can be viewed in packet capture transcripts. Device logs kept by servers also contain information about the transactions that occur between clients and servers. A session may, for example, include downloading content from a web server. The transactions representing requests and responses will be recorded in the access logs on the server or by the NIDS. A session is all the traffic involved in making a request; a transaction is the request itself.

Full packet capture

Full packet captures are the most detailed network data that are generally collected. Because of the amount of detail, they are also the most demanding types of data to store and search. Full packet captures also contain the actual content of conversations. Full packet captures contain the text of e-mail messages, the HTML of web pages and the files that enter or leave the network. Extracted content can be recovered from full packet captures and analysed for the presence of malware or of user behaviour that violates business and security policies.

Statistical data

Like session data, statistical data also concern network traffic. Statistical data are created through the analysis of other forms of network data. From these analyses, conclusions can be drawn that describe or predict network behaviour. Statistical characteristics of normal network behaviour can be compared with current network traffic in an effort to detect anomalies. Statistically significant differences should trigger an alarm and a prompt investigation.

Network Behavior Analysis (NBA) and Network Behavior Anomaly Detection (NBAD) are approaches to network security monitoring that use advanced analytical techniques to analyse network telemetry data from the NetFlow or Internet Protocol Flow Export (IPFIX) protocol. Techniques such as predictive analytics and artificial intelligence perform advanced analyses of detailed session data in order to detect potential security incidents.

Host logs

Important data that we have separated into a separate subchapter are host logs, which inform us about the state of host devices and provide us with another perspective on the network.

Syslog

Syslog contains specifications for message formats, the structure of the client-server application and the network protocol. This standard is used for logging events to centralized syslog servers. It is a client/server protocol supported by a wide range of devices.

The sender sends a small text message (<1kB) to a syslog receiver. The messages are sent via port 514 (UDP) and 5000 (TCP). With the exception of, for example, SSL, these data are usually sent in plain text.

Server logs

Server logs are a fundamental source of data for network security monitoring. Particularly important are the logs of DNS (proxy) servers, which document all DNS queries and responses occurring in the network. They are also useful for identifying host devices that may have visited dangerous websites and for identifying DNS data exfiltration and connections to malware command-and-control servers. The content of event log files depends on the type of server.

Logging from intermediary network devices

Security devices, as well as other intermediary network devices, can be configured to send events and alerts to security management platforms using the SNMP protocol or syslog.

Proxy logs

Proxy servers, such as those used for web and DNS requests, contain valuable logs that are a primary source of data for network security monitoring. A proxy server generates logs of all requests and responses. These logs can then be analysed to determine which host devices are sending requests and whether the destinations are safe or potentially malicious, and also to obtain an overview of the kind of resources that have been downloaded.

Web proxy servers provide data that help determine whether responses from a website were generated in response to legitimate requests or were manipulated to look like responses but are in fact exploits. It is also possible to use web proxies to inspect outgoing traffic as a means of data loss prevention (DLP). DLP involves scanning outgoing traffic to determine whether the data leaving the web contain sensitive, confidential or secret information.

SIEM

It is used in many organizations to provide real-time reporting and long-term analysis of security events, as shown in Figure 1.3. The purple arrows represent the inputs, which include information collected from individual devices. The orange arrows represent the outputs that SIEM offers after processing the information (e.g. dashboards and reports)

The objective of SIEM tools is to provide a comprehensive view of the enterprise network by means of the following functions:

- **Log collection** – Event records from sources (assets) across the organization provide important forensic information and help in addressing compliance reporting requirements.
- **Normalization** – Maps log messages from different systems into a common data model, which enables the organization to connect and analyse related events even if they were originally recorded in different source formats.
- **Correlation** – Links logs and events from different systems or applications, thereby speeding up the detection of and response to security threats.
- **Aggregation** – Reduces the volume of event data by consolidating duplicate event records.
- **Reporting** – Presents correlated and aggregated event data in real-time monitoring and long-term summaries, including graphical interactive dashboards.
- **Compliance** – This is reporting to meet the requirements of various regulatory compliance mandates.

Next-Generation Firewalls (NGFW)

Next-Generation Firewalls extend network security beyond IP addresses and transport layer port numbers to the application layer, or even deeper beyond this level of analysis. NextGen firewalls are advanced devices that provide many more functions than previous generations of network security devices. One of these functions is dashboards with interactive functions that enable quick reports on specific information without the need for SIEM or other event correlation tools.

Common NGFW events include:

- **Connection Event** – Connection logs contain session data that are detected directly by the Next-Generation IPS (NGIPS). Connection events include basic connection properties, such as timestamps, source and destination IP addresses, and metadata about why the connection was logged, for example which access control rule logged the event.
- **Intrusion Event** – the system examines packets passing through the network for malicious activity that could affect the availability, integrity and confidentiality of a host device and its data. When the system identifies a possible intrusion, it generates an intrusion event, which is a record of the date, time, type of exploit and contextual information about the source of the attack and its target.
- **Host or Endpoint Event** - When a host device appears in the network, the system can detect it and record details about the device hardware, IP addresses and last known presence in the network.
- **Network Discovery Event** - Network discovery events represent changes that have been detected in the monitored network. These changes are recorded in response to network discovery policies, which specify the kinds of data to be collected, the network segments to be monitored and the hardware interfaces of the device that should be used to collect events.

- **Netflow Event** - network discovery can use a number of mechanisms, one of which is the use of exported NetFlow flow records to generate new events for host devices and servers.

Information on network flows and applications

Information on network flows provides us with more detailed data and a view of the network, or rather of network traffic. Although the NetFlow protocol was not originally conceived as a network security monitoring tool, it is considered a useful tool in the analysis of network security incidents. It can be used to understand the behaviour of individual users, or to track the movement of an attacker or host-to-host exploitation within the network.

NetFlow

NetFlow is a protocol from Cisco, conceived as a tool for network troubleshooting and session-based accounting of transferred data. NetFlow effectively provides an important set of services for IP applications, including network traffic accounting, usage-based network billing, network planning, security, denial-of-service monitoring capabilities and network monitoring. NetFlow provides information about network users and applications, peak times in the network and traffic routing. NetFlow does not perform full packet capture, nor does it capture the actual content in a packet. NetFlow records information about the packet flow, including metadata. Cisco allowed its use (in version 9) as the basis for the IETF standard called IPFIX.

NetFlow is based on a set of 5 to 7 IP packet attributes:

- Source IP address
- Destination IP address
- Source port
- Destination port
- L3 layer protocol type
- Class of service
- Router or switch interface

All packets with the same source/destination IP address, source/destination ports, protocol type and class of service are grouped into a flow, and the packets and bytes are then counted. All NetFlow flow records will contain the first five items in the list above and the

timestamps of the start and end of the flow. Additional information is highly variable and can be configured. A large number of attributes are available for a flow. The IANA IPFIX entity registry lists several hundred, of which the first 128 are the most frequently used.

Application Visibility and Control

Identifying network applications by port provides very little granularity and insight into user behaviour. Application visibility through the identification of application signatures, however, identifies what users are doing. Vendors usually state for their products how many applications they can recognize, with the number ranging from 1000 upwards.

The Application Visibility and Control (AVC) system from Cisco combines several technologies to recognize, analyse and control more than 1000 applications. AVC uses the new generation of network-based application recognition version 2 (NBAR2), also known as Next-Generation NBAR, to discover and classify the applications used in the network. Similar tools are also provided by other vendors, but we will not deal with them further in this work.

Logs as output of specific content filtering (Content filter logs)

Devices that provide content filtering, such as the Cisco Email Security Appliance (ESA) and the Cisco Web Security Appliance (WSA), provide a wide range of security monitoring functions. Logging is available for many of these functions.

The ESA, for example, has more than 30 logs that can be used to monitor most aspects of e-mail delivery, system operation, antivirus and antispam operations, and blocklist and allowlist decisions. Most logs are stored in text files and can be collected on syslog servers or sent to FTP or SCP servers. In addition, alerts concerning the operation of the device itself and its subsystems can be monitored by e-mail by the administrators who are responsible for monitoring and operating the device.

WSA devices offer a similar level of functionality. The WSA effectively functions as a web proxy server, which means that it logs all incoming and outgoing transaction information for HTTP traffic. The WSA can be configured to send logs to a server in various ways, including syslog, FTP and SCP. Other logs available to the WSA include ACL decision logs, malware scanning logs and web reputation filtering logs.

Alert evaluation

In this section, we describe working with the data mentioned in the section Network security data, using the tools described in the section Tools

Sources of alerts

In this section, we deal with the sources of alerts and alert data. Our objective in this work is to provide SOC analysts with all three basic functions, i.e. full packet capture and data types, intrusion detection systems and alert analysis tools. We have also succeeded in covering these functions, with intrusion detection systems addressed only at the network level and not at the level of individual host devices and critical servers.

Detection tools for collecting security alerts

Our objective is to create an integrated environment that simplifies the deployment of a comprehensive NSM solution as much as possible and makes use of existing systems already deployed at the department in the past, so that they do not have to be significantly changed or replaced, provided that they fulfil the required purpose.

Generation of security alerts

Security alerts are notification messages generated by NSM tools, systems and security devices (alert generation). They can be used to alert cybersecurity analysts to events that require attention. A cybersecurity analyst can work with a queue of such alerts, which they can further examine, classify, escalate or dismiss.

Alerts will generally contain the five tuples, if available, as well as timestamps and information identifying which device or system generated the alert.

Further information may include whether a permit or deny decision was applied to the traffic, some captured data from the packet payload, a hash value for a downloaded file or another type of data.

Depending on the security technology, alerts may be generated on the basis of rules, signatures, anomalies or behaviour. Regardless of how they are generated, the conditions that trigger an alert must be predefined in some way.

The information found in the alerts displayed in ELK will differ in message format, because it comes from different sources. It is important for a cybersecurity analyst to be able to interpret what triggered an alert so that the alert can be investigated. For this reason, a cybersecurity analyst should understand the components of Suricata IDS rules, which are the main source of alerts in our deployed ELK system.

Overview of security alert evaluation

In this section, we describe the motivation for data evaluation, the evaluation itself and the possible results of this process. We also mention two types of approach to analysis. In preparing this section, we drew information mainly from source [2] .

The need for alert evaluation

Threats are constantly changing, because new vulnerabilities emerge and new threats are developed. As the needs of users and the organization change, so does the attack. Threat actors have learned how to quickly change the attributes of attacks in order to avoid detection. It is not possible to design measures to prevent all possible threats. Threat actors will always make the greatest possible effort to evade protective measures, regardless of how sophisticated these measures may be. Therefore, situations arise in which an analyst detects an attack while it is in progress or only after an incident has occurred. It is recommended that detection rules be more conservative. In other words, it is better to have alerts that are sometimes generated by legitimate traffic than to have rules that miss malicious traffic. For this reason, experienced cybersecurity analysts need to review alerts to determine whether an exploit has actually occurred.

Evaluation of security alerts

Security incidents are classified using a scheme borrowed from medical diagnostics. The point is that a diagnosis can be either accurate, true, inaccurate or false. False diagnoses are either expensive or dangerous. In network security analysis, the cybersecurity analyst is presented with an alert. The cybersecurity analyst asks: “The system says that an exploit has occurred. Is this true?"

Alerts can be classified as follows:

- True positive
    - The alert has been verified as a genuine security incident.
    - True positives are the desired type of alert. This means that the rules that generate the alerts have worked correctly.
- False positive
    - The alert does not indicate a genuine security incident. It is legitimate activity that leads to a false positive result.
    - They are not desirable. Although they do not indicate that an undetected exploit has occurred, they are costly, because cybersecurity analysts must investigate false alarms, which takes time away from investigating alerts that point to genuine exploits.

The alternative situation is that no alert was generated. The absence of an alert can be classified as:

- True negative
    - No security incident has occurred. The activity is legitimate.
    - They are desirable. They indicate that legitimate, normal traffic is correctly ignored and no erroneous alerts are issued.
- False negative
    - An undetected incident has occurred.
    - They are dangerous. They indicate that the security systems in place have failed to detect an exploit. These incidents may remain undetected for a long time, and data loss and damage may occur.

Legitimate events are those that should not trigger alerts. Excessive legitimate events indicate that some rules or other detectors need to be improved or removed. False negatives may be discovered only long after an exploit. This can happen through retrospective security analysis (RSA). RSA may occur when newly obtained rules or other threat intelligence are applied to archived network security data. For this reason, it is important to monitor threat intelligence in order to learn about new vulnerabilities and exploits and to evaluate the

likelihood that the network has ever been vulnerable to them in the past. In some cases, adding new mitigation techniques may be sufficient; in others, a more detailed analysis needs to be performed.

Deterministic analysis and probabilistic analysis

To evaluate the risk of whether intrusions or exploits will be successful in a given network, statistical techniques can be used, employing two general approaches, namely deterministic and probabilistic analysis.

Deterministic analysis is based on the assumption that for an intrusion to be successful, all the preceding steps must also be successful. The cybersecurity analyst knows the steps for a successful exploit.

Probabilistic analysis uses statistical techniques to determine the probability that a successful exploit will occur, based on the probability with which each step of the exploit can be successful.

Digital forensic analysis

In order to protect the organization and prevent computer crime, it is necessary to identify threat actors, report them to the competent authorities and provide evidence to support prosecution. Tier 1 cybersecurity analysts are often the first to uncover wrongdoing and must therefore know how to handle evidence properly and attribute it to threat actors.

Digital forensics is the recovery and investigation of information found on digital devices that is related to criminal activity. This information may be data on storage devices, in volatile computer memory, or traces of computer crime that are preserved in network data, such as full data captures (pcaps) or the logs we mentioned in section 1.2.1 Types of security data. It is essential that all indicators of compromise are preserved for future analysis and attack attribution.

Cybercrime activity can generally be characterized as activity originating from within specific organizations or from outside them. The investigation of offences by persons not covered by any organization concerns individuals inside the organization. These individuals may behave in a way that violates user agreements, or other behaviour that cannot yet be

classified as a criminal offence. However, if individuals are suspected of involvement in criminal activity involving the theft or destruction of intellectual property, the organization may decide to involve law enforcement authorities, in which case the investigation becomes public. Internal users may also have used the organization's network to carry out other criminal activities that are unrelated to the organization's mission but are in breach of various legal regulations. In this case, public officials will carry out the investigation.

When an external attacker exploits the network and steals or alters data, it is necessary to collect evidence to document the extent of the exploitation. Various regulatory bodies specify the scope of the measures that an organization must take if various types of data have been compromised. The results of the forensic investigation can help identify the measures that need to be taken.

In some cases, the organization itself may be the subject of an investigation. Cybersecurity analysts may find themselves in direct contact with digital forensic evidence that describes in detail the behaviour of members of the organization. Analysts must know the requirements concerning the preservation and handling of such evidence. Failure to do so may result in criminal penalties for the organization and even for the cybersecurity analyst if an intent to destroy evidence is proven.

Digital forensic process

It is important for an organization to develop well-documented processes and procedures for digital forensic analysis. Regulatory compliance may require this documentation, and this documentation may be inspected by the authorities in the event of a public investigation.

The special publication NIST 800-86 Guide to Integrating Forensic Techniques into Incident Response is a valuable resource for organizations that need guidance in developing digital forensic plans. It recommends, for example, that forensic analysis be carried out using a four-phase process.

- Collection – this is the identification of potential sources of forensic data and the acquisition, handling and storage of these data. This phase is critical, because particular care must be taken to ensure that important data are not damaged, lost or omitted.
- Examination – i.e. assessing and extracting relevant information from the collected data (e.g. decompressing or decrypting the data). It may be necessary to remove information that is irrelevant to the investigation. Identifying actual evidence in large collections of data can be very demanding, including in terms of time.
- Analysis – this means drawing conclusions from the data. Significant elements, such as people, places, times, events, etc., should be documented. This step may also include the correlation of data from multiple sources.
- Reporting – This involves preparing and presenting the information that resulted from the analysis. Reporting should be impartial and, where necessary, alternative explanations should be provided. The limitations of the analysis and the problems encountered should be included. Suggestions for further investigation and next steps should also be made.

Types of evidence

In legal proceedings, evidence is generally classified as direct or indirect. Direct evidence is evidence that was undoubtedly in the possession of the accused, or is eyewitness evidence from someone who directly observed the criminal behaviour.

Evidence is further classified as:

- Best evidence – This is evidence that is in its original state. This evidence may be storage devices used by the accused or file archives that can be proven to be unaltered.
- Corroborating evidence – This is evidence that supports a claim that is based on the best evidence.

- Indirect evidence – This is evidence that, in combination with other facts, establishes a hypothesis. (For example, evidence that an individual has committed similar crimes may support the claim that the person committed the crime of which they are accused)

Recommendation for evidence collection

IETF RFC 3227 provides guidelines for the collection of digital evidence. It describes a recommendation for collecting digital evidence based on the volatility of the data. Data stored in RAM are the most volatile and are lost when the device is switched off. In addition, important data in volatile memory may be overwritten by routine machine processes. The collection of digital evidence should therefore start with the most volatile evidence and proceed to the more stable. An example of the order of evidence collection, from the least stable to the more stable, is as follows:

1. Memory registers, caches
2. Routing table, ARP table, process table, kernel statistics, RAM
3. Temporary file systems
4. Non-volatile media, fixed and removable
5. Remote logging and monitoring data
6. Physical interconnections and topologies
7. Archival media, tapes or other backups

Details of the systems from which the evidence was collected, including who has access to these systems and at what level of permissions, should be recorded. Such details should include the hardware and software configurations of the systems from which the data were obtained.

Chain of custody

Although evidence may have been collected from sources that support attribution to the accused individual, it may be argued that the evidence could have been altered or fabricated after it was collected. In order to counter this argument, a strict chain of custody must be

defined and followed. The chain of custody covers the collection, handling and secure storage of evidence. Detailed records should be kept, in terms of answers to the following questions:

- Who discovered and collected the evidence?
- All details about the handling of the evidence, including times, places and the personnel involved.
- Who has primary responsibility for the evidence, when was the responsibility assigned and when did custody change?
- Who has physical access to the evidence while it was stored? Access should be restricted to the most essential personnel only.

Data integrity and protection

When collecting data, it is important that they are preserved in their original state. The timestamp of files should be preserved. For this reason, the original evidence should be copied and the analysis should be performed only on copies of the original. This is to prevent accidental loss or alteration of the evidence. Since timestamps may be part of the evidence, opening files from the original media should be avoided.

The process used to create copies of the evidence used in the investigation should be recorded. Whenever possible, the copies should be direct bit-level copies of the original storage units. It should be possible to compare the archived disk image and the image of the investigated disk in order to determine whether the content of the examined disk has been tampered with. For this reason, it is important to archive and protect the original disk so that it is preserved in its original, undamaged state.

Volatile memory may contain forensic evidence, so special tools should be used to preserve this evidence before the device is switched off and the evidence is lost. Users should not disconnect or switch off infected computers unless explicitly permitted to do so by security personnel.

Following these processes will ensure that all evidence of wrongdoing is preserved and that any indications of compromise can be identified.

Attack attribution

After the extent of the cyber attack has been evaluated and the evidence has been collected and preserved, the incident response can move on to identifying the source of the attack. There is a wide range of threat actors, from disgruntled individuals, hackers, cybercriminals and criminal gangs to individual states. Some criminals act from inside the network, while others may be on the other side of the world. The sophistication of computer crime also varies.

Threat attribution refers to the act of determining the individual, organization or nation responsible for a security breach. The identification of the responsible threat actors should be carried out through a systematic investigation of the evidence.

In an evidence-based investigation, the incident response team correlates the tactics, techniques and procedures (TTP) used in the incident with other known exploits. Cybercriminals, like other criminals, have specific traits that are common to most of their crimes. Evidence of computer crime is rarely direct evidence. Identifying common traits between the TTPs of known and unknown threat actors is indirect evidence.

Some aspects of a threat that can help with attribution are the location of the originating host devices or domains, the characteristics of the code used in the malware, the tools used and other techniques. Sometimes, at the level of national security, threats cannot be openly attributed, because this would reveal methods and capabilities that need to be protected.

In the case of internal threats, asset management plays a major role. Uncovering the devices from which the attack was launched can lead directly to the threat actor. IP addresses, MAC addresses and DHCP protocols can help trace the addresses used in the attack back to a specific device. AAA logs from devices with a user database used for authentication, authorization and accounting (AAA) are very useful in this regard, because they track who accessed which network resources at what time.

The MITRE ATT&CK Framework

One way of attributing an attack to specific threat actors, i.e. identifying those actors, is to model the behaviour of threat actors. The MITRE Adversarial Tactics, Techniques &

Common Knowledge (ATT&CK) framework provides the ability to detect attackers' TTPs as part of threat defence and attack attribution. The MITRE ATT&CK Framework is a global knowledge base of threat actor behaviour. It is based on the observation and analysis of real network intrusions with the aim of describing the attacker's behaviour, not the attack itself. It is designed to enable automated information sharing by defining data structures for the exchange of information between the user community and the MITRE organization. At the time of writing this work, the framework contains 14 tactics, 185 techniques and 367 sub-techniques [14] .

The Cyber Kill Chain

The Cyber Kill Chain was developed by Lockheed Martin to identify and prevent cyber intrusions. Seven steps lead to the Cyber Kill Chain. Focusing on these steps helps analysts understand the techniques, tools and procedures of threat actors. When responding to a security incident, the objective is to detect and stop the attack as early as possible in the course of its progression (kill chain progression). The earlier the attack is stopped, the less damage is caused and the less the attacker learns about the target network.

The Cyber Kill Chain specifies what an attacker must accomplish in order to achieve their goal. If the attacker is stopped at any phase, the chain of attack is broken. Breaking the chain means that the defender has successfully thwarted the threat actor's intrusion. Threat actors are successful if they complete step 7:

1. Reconnaissance
2. Weaponization
3. Delivery
4. Exploitation
5. Installation
6. Command and control
7. Actions on objectives

We have described more about this framework, or rather the steps of the Cyber Kill Chain, in the document entitled KrokyCyberKillChain, which is stored in Annex B.

Diamond Model of Intrusion Analysis

The Diamond Model is intended for the analysis of a security intrusion, a security incident or an event (diamond model of intrusion analysis) and consists of four parts. In the Diamond Model, an event is a time-bound activity that is restricted to a specific step in which the adversary uses a capability over an infrastructure to attack a victim in order to achieve a specific result.

The four core features of an intrusion event are adversary, capability, infrastructure and victim:

- Adversary – These are the parties responsible for the intrusion.
- Capability – This is the tool or technique that the adversary uses to attack the victim.
- Infrastructure – This is the network path or paths that adversaries use to establish and maintain a system for the command and control of target devices, which they exploit for this purpose to generate the attack.
- Victim – This is the target of the attack. However, the victim may initially be the target and may then be used as part of the infrastructure to launch other attacks.

Meta-features slightly expand the model to include the following important elements:

- Timestamp – indicates the start and end time of the event and is an integral part of grouping malicious activities.
- Phase – this is analogous to the steps in the Cyber Kill Chain; malicious activity involves two or more steps performed in succession to achieve the desired result.
- Result – determines what the adversary gained from the event. Results can be documented as one or more of the following: compromise of confidentiality, compromise of integrity and compromise of availability.
- Direction – indicates the direction of the event across the Diamond Model. These include the directions: adversary to infrastructure, infrastructure to victim, victim to infrastructure and infrastructure to adversary.

- Methodology – used to classify the general type of event, such as port scanning, phishing, a content delivery attack, a flood of TCP synchronization request segments and others.
- Resources – these are one or more external resources that the adversary uses to carry out the intrusion, such as software, the adversary's knowledge, information (e.g. user names and passwords) and assets for carrying out the attack (hardware, funds, facilities, network access).

Traversing the Diamond Model

A cybersecurity analyst may be asked to use the Diamond Model of intrusion analysis to create a diagram of a series of intrusion events. The Diamond Model is ideal for illustrating how an adversary moves from one event to another (pivoting).

For example, in Figure 1.8 on the following page, an employee reports that their computer is behaving abnormally. A host device check performed by a security technician indicates that the computer is infected with malware. Analysis of the malware shows that it contains a list of Command and control (CnC) domain names. These domain names resolve to a list of IP addresses. These IP addresses are then used to identify the adversary, as well as to examine logs to determine whether other victims in the organization are using the given CnC channel.

The Diamond Model and the Cyber Kill Chain

Adversaries do not operate in just a single event. Instead, events are linked into a chain in which each event must be successfully completed before the next event. This thread of events can be mapped to the Cyber Kill Chain.

The following example, shown in Figure 1.9 on the following page, illustrates the entire process of an adversary as they move vertically through the Cyber Kill Chain, use a compromised user device to move horizontally to another victim, and then begin another thread of activity.

1. The adversary performs a web search for the victim company Gadgets, Inc., obtaining the domain name gadgets.com as part of the results.

2. The adversary uses the newly discovered domain gadets.com for a new search, “network administrator gadget.com”, and discovers forum posts by users claiming to be gadget.com network administrators. The user profiles reveal their e-mail addresses.

3. The adversary sends phishing e-mails with an attached Trojan horse to the gadget.com network administrators.

4. One gadget.com network administrator (NA1) opens the malicious attachment. This executes the enclosed exploit, which enables further code execution.

5. The compromised user NA1 sends an HTTP post message to an IP address and registers it with the CnC controller (control server). The compromised user NA1 receives an HTTP response.

6. Reverse engineering revealed that the malware has additional IP addresses configured, which serve as a backup if the first CnC server does not respond.

7. Through a CnC HTTP response message sent to host device NA1, the malware begins to act as a web proxy for new TCP connections.

8. Using information from the proxy server running on user NA1, the adversary searches the web for “most important research ever” and finds victim 2, “Interesting Research Inc”.

9. The adversary checks NA1's list of e-mail contacts for any contacts from “Interesting Research Inc.” and discovers a contact for the company's lead researcher.

10. The company's head of research receives a spear-phishing e-mail from NA1's Gadget Inc. e-mail address, sent by the company user NA1, with the same unwanted content as observed in event 3 (Trojan horse).

The adversary now has two victims under their control from which further attacks can be launched. For example, the adversary could mine the head of research's e-mail contacts for further potential victims. The adversary may also set up another proxy to exfiltrate all of the head of research's files.

Incident response

Incident response comprises the methods, policies and procedures that an organization uses when responding to a cyberattack. The goal of incident response is to limit the impact of the attack, assess the damage caused and implement recovery procedures. It is essential that organizations create and maintain detailed incident response plans and designate the employees who will be responsible for carrying out all aspects of this plan.

The recommendations of the US National Institute of Standards and Technology (NIST) for incident response are set out in detail in its Special Publication 800-61, Revision 2, entitled "Computer Security Incident Handling Guide" [15] .

Incident response policy, plan and procedure

In this section we briefly describe the elements of a cybersecurity incident response policy, plan and procedure.

Elements of an incident response policy

An incident response policy sets out in detail how incidents should be handled based on the organization's mission, size and function. The policy should be reviewed regularly so that it is adapted to meet the goals of the established plan.

Elements of an incident response plan

A good incident response plan helps to minimize the damage caused by an incident. It also helps to improve the overall incident response programme by adjusting it according to lessons learned. It ensures that every party involved in the incident response clearly understands not only what it will do, but also what the others will do.

Elements of an incident response procedure

The procedures followed during incident response should be governed by the incident response plan.

Stakeholders in the incident response process

Other groups and individuals within the organization may also be involved in incident handling (Incident Response Stakeholders). We mentioned some centres and units in section

1.1.5 Other centres and units dealing with operations or security. It is important to ensure that they cooperate before an incident occurs. Their expertise and capabilities can help the Computer Security Incident Response Team (CSIRT) handle an incident quickly and correctly.

Cybersecurity Maturity Model Certification

The Cybersecurity Maturity Model Certification (CMMC) framework was created to assess the ability of organizations performing functions for the US Department of Defense (DoD) to protect the military supply chain from breaches or losses resulting from cybersecurity incidents. Security breaches involving DoD information indicated that NIST standards were not sufficient to mitigate the growing and evolving threat landscape, especially from nation-state threat actors. In order to receive contracts from the Department of Defense, companies must be certified. Certification consists of five levels, with different levels required depending on the degree of security required by the project.

CMMC certifies organizations by level. For most domains there are five levels, but for incident response there are only four. The higher the certification level, the more mature the organization's cybersecurity capability.

- Level 2 – Create an incident response plan that follows the NIST process. Detect, report and prioritize events. Respond to events according to predefined procedures. Analyse the cause of incidents in order to mitigate future problems.
- Level 3 – Document and report incidents to the stakeholders identified in the incident response plan. Test the organization's ability to respond to incidents.
- Level 4 – Use knowledge of attacker tactics, techniques and procedures (TTPs) to refine incident response planning and execution. Establish a security operations centre (SOC) that enables 24/7 response capability.
- Level 5 - Use accepted and systematic computer forensic data collection techniques, including the secure handling and storage of forensic data. Develop and use manual and automated real-time responses to potential incidents that follow known patterns.

NIST incident response life cycle

NIST defines four steps in the life cycle of the incident response process, as shown in the figure below.

- Preparation – CSIRT members are trained in how to respond to an incident. CSIRT members should continually develop their knowledge of emerging threats.
- Detection and analysis – through continuous monitoring, the CSIRT quickly identifies, analyses and validates an incident.
    - Precursor – This is a sign that an incident may occur in the future.
    - Indicator – This is a sign that an incident may already have occurred or is currently occurring.
- Containment, eradication and recovery – The CSIRT implements procedures to contain the threat, eradicate the impact on organizational assets and use backups to restore data and software. This phase may loop back to detection and analysis in order to gain more information or expand the scope of the investigation.
- Post-incident activities – The CSIRT then documents how the incident was handled, recommends changes for future response and specifies how to prevent a similar incident from recurring in the future.

Hardening based on lessons learned

After a major incident has been resolved, the organization should hold a stakeholder meeting to review the effectiveness of the incident handling process and identify the necessary tightening of existing security controls and procedures.

Collection and retention of incident data

Thanks to stakeholder meetings, the collected data can be used to determine the cost of an incident for budgetary reasons, as well as to determine the effectiveness of the CSIRT and to identify possible security weaknesses throughout the system. The data that should be collected are those that can be used to define and refine the incident handling process.

Every organization should have a policy in place that states how long evidence of an incident is retained. Evidence is often retained for many months or years after an incident. In

some cases, regulations and directives may mandate the retention period. These are some of the decisive factors for evidence retention:

- Prosecution – if an attacker is to be prosecuted because of a security incident, the evidence should be retained until all legal actions have been completed. In legal proceedings, no evidence should be overlooked or considered insignificant. The organization's policy may state that no evidence relating to an incident that was associated with legal action may ever be deleted or destroyed.
- Data type – an organization may state that certain types of data should be retained for a certain period. Items such as e-mail or text may only need to be retained for 90 days. It is possible that more important data, such as data used in incident response (that has not been prosecuted), will need to be retained for three years or longer.
- Cost – If a lot of hardware and storage media need to be stored for a long time, their storage can become expensive. It must also be borne in mind that as technologies change, functional devices capable of using the obsolete hardware and storage media must also be stored.

Reporting and information-sharing requirements

The legal team should consult government regulations to determine precisely the organization's responsibility for reporting the incident. In addition to legal requirements and stakeholder considerations, NIST recommends that the organization coordinate the details of the incident with organizations. For example, the organization could record the incident in the VERIS community database.

List of references

[1] Damoglu Burak, 'Introduction to Security Operations Center | by Burak Damoglu | Medium', Sep. 25, 2020. https://medium.com/@BurakDamoglu/introduction-to-security-operations-center-73e2e50f3ba7 (accessed Mar. 13, 2023).

[2] 'Course: 2022_CyberOps', 2020. https://lms.netacad.com/course/view.php?id=1418088 (accessed Mar. 29, 2023).

[3] Tyagi Aakanksha, 'Role And Responsibilities of a SOC Analyst - InfosecTrain', May 17, 2021. https://www.infosectrain.com/blog/role-and-responsibilities-of-a-soc-analyst/ (accessed Mar. 15, 2023).
[4] 'Ticketing process - IBM Documentation', Mar. 03, 2021. https://www.ibm.com/docs/en/sss/3.1.1?topic=siteprotector-ticketing-process (accessed Mar. 15, 2023).
[5] Carklin Nicolette, 'Network Security Monitoring: A Complete Guide', Feb. 07, 2022. https://www.parallels.com/blogs/ras/network-security-monitoring/ (accessed Mar. 15, 2023).
[6] Ploštica Marek, 'Simulátory počítačových sietí – analýza II', Bachelor's thesis, Žilinská univerzita v Žiline , 2021.
[7] Čulík Štefan, 'Prehľadová analýza dostupných nástrojov pre mapovanie prepojení a vizualizáciu zariadení v sieti', Bachelor's thesis, Žilinská univerzita v Žiline , 2022.
[8] Husová Simona, 'Implementácia bezpečnostných opatrení pre zmiernenie rizík pri poskytovaní IaaS služieb v akademickom prostredí', Diploma thesis, Žilinská univerzita v Žiline , 2022.
[9] '69/2018 Z.z. - Zákon o kybernetickej bezpečnosti a... - SLOV-LEX'. https://www.slov-lex.sk/pravne-predpisy/SK/ZZ/2018/69/ (accessed Jan. 19, 2021).
[10] '362/2018 Z.z. - Vyhláška Národného bezpečnostného úradu'. https://www.slov-lex.sk/pravne-predpisy/SK/ZZ/2018/362/ (accessed May 10, 2020).
[11] 'EDR, NDR, XDR, MDR - Different concepts of Detection &… | Nomios Group', Sep. 26, 2021. https://www.nomios.com/news-blog/edr-ndr-xdr-mdr/ (accessed Mar. 15, 2023).
[12] 'NOC vs SOC: Comparing Network & Security Operations Centers | Splunk', Jan. 30, 2023. https://www.splunk.com/en_us/blog/learn/noc-vs-soc.html (accessed Mar. 18, 2023).
[13] 'What is the Difference Between a SOC and a CSIRT? | Rapid7 Blog', Apr. 19, 2017. https://www.rapid7.com/blog/post/2017/04/19/difference-between-a-soc-and-a-csirt/ (accessed Mar. 18, 2023).
[14] 'What is MITRE ATT&CK: An Explainer'. https://www.exabeam.com/explainers/mitre-attck/what-is-mitre-attck-an-explainer/ (accessed Apr. 13, 2023).
[15] P. Cichonski, T. Millar, T. Grance, and K. Scarfone, 'Computer Security Incident Handling Guide', Aug. 2012, doi: 10.6028/NIST.SP.800-61R2.

Annex: List of tools

Vulnerability Assessment			
Name	**Company**	**Licence**	**Note**
Intruder	Intruder Ltd.	Proprietary	free trial 14 days
Acunetix	Invicti Security	Proprietary	Contact for Quote
Nessus	Tenable Inc	Proprietary	free trial 7 days
Invicti (Netsparker)	Invicti Security	Proprietary	Contact for Quote, free trial cloud only
Nexpose	Rapid7	Proprietary	free trial 30 days
Burp Suite	PortSwigger Ltd.	Proprietary	free trial
Tripwire IP360	Tripwire	Proprietary	Contact for Quote
OpenVAS	Greenbone Networks	Open-source	
Nikto2	-	Open-source	
W3AF	-	Open-source	

HIDS			
Name	**Company**	**Licence**	**Note**
Security Event Manager	Solarwinds	Proprietary	
Samhain	Samhain	Open-source	
OSSEC	-	Open-source	
Wazuh	Wazuh	Open-source	Built on top of OSSEC
Tripwire	Tripwire	Open-source	Currently not maintained
Tripwire Enterprise	Tripwire	Proprietary	Replaced Tripwire and is paid
NIDS			

Name	Company	Licence	Note
Snort	Cisco	Open-source	
Zeek/Bro	-	Open-source	
Suricata	-	Open-source	
Open WIPS-NG	-	Open-source	
Sagan	-	Open-source	

EDR			
Name	**Company**	**Licence**	**Note**
EDR Security Service	Cynet	Proprietary	also an XDR solution, free trial 14 days
Endpoint Central (Desktop Central)	ManageEngine	Proprietary	Contact for Quote, free trial 30 days
CrowdStrike - Falcon Insight	CrowdStrike	Proprietary	free trial, EDR and XDR in a single platform
Carbon Black	VMware	Proprietary	
SentinelOne	SentinelOne	Proprietary	EDR and XDR
Symantec EDR	Broadcom	Proprietary	
Cybereason	Cybereason	Proprietary	
Cisco Secure Endpoint (AMP for Endpoints)	Cisco	Proprietary	free trial 30 days
OpenEDR	-	Open-source	

NDR			
Name	**Company**	**Licence**	**Note**
Stellar Cyber	Stellar Cyber	Proprietary	NDR within XDR

ExtraHop Reveal(x)	ExtraHop networks	Proprietary	
Vectra NDR	Vectra.ai	Proprietary	
Hillstone Networks	Hillstone Networks	Proprietary	NDR, also offers XDR
IronNet	IronNet	Proprietary	
Darktrace	Darktrace	Proprietary	
Lastline	Lastline	Proprietary	
Flowmon	Flowmon networks	Proprietary	Get a Quote, free trial 30 days

XDR			
Name	**Company**	**Licence**	**Note**
Cynet Auto XDR	Cynet	Proprietary	also an EDR solution, free trial 14 days
Cortex XDR	Palo Alto Networks	Proprietary	
Sophos	Sophos	Proprietary	Contact for quote
Symantec	Broadcom	Proprietary	
InsightIDR	Rapid7	Proprietary	free trial, “It was XDR before XDR was even a thing” (also referred to as SIEM)
Fidelis Elevate	Fidelis Cybersecurity	Proprietary	free trial 15 days
Trend Micro	Trend Micro	Proprietary	free trial 60 days
SecureX	Cisco	Proprietary	
SIEM			

Name	Company	Licence	Note
Splunk	Splunk	Proprietary	Free trial 60 days, then free version (<500MB/day)
Exabeam Fusion	Exabeam	Proprietary	Licence required
Log360	ManageEngine	Proprietary	Free trial 30 days, then free version (25 workstation, 5 log sources...)
Logpoint	LogPoint	Proprietary	Price calculated using a calculator
USM	AlienVault / AT&T Cybersecurity	Proprietary	
LogRhythm SIEM	LogRythm	Proprietary	Subscription, Contact for Quote
Qradar SIEM	IBM	Proprietary	Price based on the number of workstations / servers, licences based on a memorandum with FRI
Elastic Stack	Elastic	Open-source	Standard functionality free, commercial add-ons available
SIEMonster	SIEMonster	Open-source	Single server up to 100 endpoints, plus other paid editions

OSSIM	AlienVault / AT&T Cybersecurity	Open-source	

PACKET ANALYZERS			
Name	**Company**	**Licence**	**Note**
Deep Packet Inspection and Analysis Tool	SolarWinds	Proprietary	Free trial 30 days
PRTG	Passler	Proprietary	Free up to 100 sensors
Omnipeek	LiveAction	Proprietary	Free trial
tcpdump	-	Open-source	CLI
WinDump	-	Open-source	tcpdump clone for Windows
Wireshark	-	Open-source	GUI
tshark	-	Open-source	CLI version of Wireshark
NetworkMiner	Netresec	Open-source	
Fiddler	Telerik	Proprietary	Free trial 10 days
Capsa	Colasoft	Proprietary	Free version with limitations available

FLOW ANALYZERS			
Name	**Company**	**Licence**	**Note**
SolarWinds NetFlow Traffic Analyzer	SolarWinds	Proprietary	Free trial 30 days

ManageEngine NetFlow Analyzer	ManageEngine	Proprietary	Free version available
nProbe	ntop	Proprietary	Requires a per-system licence
Ntopng	Ntop	Proprietary	Licence required
Scrutinizer	Plixer	Not determined	
Noction Flow Analyzer	Noction	Proprietary	Licence required
Kentik	Kentik technologies	Proprietary	Free trial 30 days
PRTG	Passler	Proprietary	Free up to 100 sensors
Nagios	Nagios	Proprietary	Free trial 30 days